



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

**THE DEVELOPMENT
OF LARGE SCALE
IN SILICO MODELS
TO EXPLORE
BIOLOGICAL
MECHANISMS
IN RESEARCH
AND TEACHING**

Richard Fitzpatrick

Doctor of Philosophy

The University of Edinburgh

2021

I declare:

- (a) That the thesis has been composed by myself, and
- (b) Either that the work is my own, or, if I have been a member of a research group, that I have made a substantial contribution to the work, such contribution being clearly indicated, and
- (c) That the work has not been submitted for any other degree or professional qualification except as specified, and
- (d) That any included publications are my own work, except where indicated throughout the thesis and summarized and clearly identified on the declarations page of the thesis.

Richard Fitzpatrick

25/02/2021



Supervisors: Dr Melanie Stefan; Dr Peter Kind

LAY SUMMARY

Computer modelling is a way of testing ideas about biology without needing to use living tissue. This is potentially a very powerful way of exploring biological ideas because it is quick, can be done at any time, and can make predictions about what experimenters might find if they repeat the same experiments in an organism or test tube.

One type of computer modelling is biochemical modelling, which explores how molecules in a cell interact with one another over time. A cell can have many millions of different molecules, and simulating all of these has proven impossible to achieve in the past due to *how* models are simulated, as well as how powerful computers were. This has led to most models only having a handful of molecules and the chemical reactions they undergo, which has limited their use by researchers. This in turn reduces the perceived value of models by researchers. As biochemical models are built on experimental data, this perception has arguably reduced the ability to request this data be collected to allow models to become better.

The first half of the thesis explores what a biochemical model is, and how we can make them better using new tools and approaches. I show two attempts to construct very large models which try to help explain how learning and memory works in the nervous system, focussing on structures called dendritic spines. These receive information from other neurons and will 'fire' if the right biochemical conditions are met. One molecule crucially important for many dendritic spines to work properly is an AMPA receptor (AMPA), which changes the voltage within a dendritic spine when open. The first model explores building a model of a dendritic spine which contains this receptor and all the other molecules known to be affected by the activity of AMPARs that have been previously modelled. From this attempt, I found several issues that makes combining models like this very difficult to do reliably, as well as many gaps in available values for concentrations and reactions. As a result, I developed a set of guidelines as to how so-called 'maximal' models should be built.

I used these to refine and redevelop a new AMPAR model. AMPARs can be in many different 'states' and are found in different locations on the surface of the dendritic spine. Traditional model approaches cannot model all of these at once, because each possible state has to be written out for each reaction, even if the changes to states has no effect on the reaction rate. This leads to a massive increase in the number of molecules and reactions, quickly making writing out the model

impossible. I use a new approach called rule-based modelling, which changes reactions into rules, and has allowed me to construct the first comprehensive model of AMPARs at a dendritic spine. I show the development of this model over time, exploring the addition of parts of the first model and how I have tested it to be very close to what happens inside neurons. I then use the model to explore how the states and numbers of AMPARs changes over time under different conditions which mimic how a dendritic spine is stimulated in a real neuron. I confirm that some states are necessary for successful changes to occur, and that the way in which AMPAR travelling is controlled in the dendritic spine has a big impact on the response.

The second half of the thesis looks at how building models can be used in different teaching contexts, including in public engagement. The first method is similar to that of the first half of this thesis, where students were asked to build models of biochemical molecules based on their own research. These are very successful in showing that model development can provide a novel and accessible way to understand complex biological concepts even in secondary school students. They also show that these models once constructed are useful academically, providing a template for more rigorous models and highlighting gaps in experimental literature on the biochemistry that underpins a biological concept.

The second method is less typical, investigating how using popular sandbox games such as Minecraft may provide a creative approach to both modelling and biological education across multiple age groups. I demonstrate several different investigations into this approach, which includes student-led model building and public engagement aimed at young children. I also introduce pilot studies into how to use Minecraft in the development of more complex activities such as virtual and augmented reality, and how the approach may even be used in human-based experiments on learning and memory.

This thesis aims to show how the development, and not just the use of, models is a powerful approach to both improving understanding of biological concepts scientifically as well as personal understanding of these concepts in a creative manner.

ABSTRACT

Computational models are increasingly valued as a method to help with both describing and explaining biological phenomena. Models work by collating information (both qualitative and quantitative) which has been discovered experimentally to relate to a biological output of interest. Models can then be used to test the validity of hypotheses or findings about this output, as well as making predictions of future experimental findings. Models come in many flavours, usually focussing on a particular level of biological inquiry (molecular, cellular, systems etc.). Biochemical modelling aims to offer explanations of biological phenomena through simulating the actions of biochemical networks, chiefly by defining molecular concentrations, interaction rates, and localisation. To date however, this type of model has been restricted to only truncated models representing only partial elements of biochemical networks. This has limited their utility, as well as the perception amongst non-modellers as being a useful parallel method of analysis.

In the first half of this thesis I present a reappraisal of biochemical modelling, highlighting its strengths as well as its underutilisation. In particular, I note how the development of these models requires collaboration with experimenters, both in the acquisition of data to improve models, as well as in better explanations of how models work currently when kinetic interaction data is scarce. I show some of these principles in the development of a merged model that collates pre-existing synaptic plasticity models into a single “maximal” model, the AMPAR-centric Model of Postsynaptic Signalling (AMPS). I use this experience to develop a sketch of how “maximalist modelling” should be, defining its core principles and how it may be used in both improving collaborative practice and biological understanding.

One large restriction of models like AMPS is that many of the molecules associated with synaptic plasticity may adopt multiple functional states. In traditional modelling approaches, these states must all be explicitly defined, leading to an exponential rise in model size as the numbers of components increases. However, thanks to recent advances in computational power and algorithms, one can generate rule-based models, which are better suited to explain complex phenomena such as synaptic plasticity, and that align more closely to maximalist modelling principles. I adapt AMPS to capitalise on this and show how I have iteratively constructed and developed the model to explore how AMPARs are trafficked at the synapse during plasticity events. I show that the stability of elevated levels of AMPARs in the early phase of long-term potentiation (LTP)

requires depression of phosphatase activity, which can be achieved through high levels of CaMKII autophosphorylation. Moreover, I show that this stabilisation does not require any new protein synthesis of AMPARs, or any change to other molecular concentrations. I confirm that trafficking is a complex mixture of receptor-modification, synaptic stabilisation and localisation that is best suited to being explored *in silico*. I also confirm that mGluR-mediated pathways are crucial for a complete understanding of synaptic plasticity, with their absence leading to potentiation when paired-pulse low-frequency stimulation is applied.

The second half of the thesis arises from one defined aspect of maximal models, which is that they should be educationally valuable tools as well as being valuable in research. I have developed both a typical and atypical approach to biological modelling that can be used in secondary through tertiary pedagogy. The typical approach is to make use of common, well characterised model software such as COPASI to allow learners to construct their own model (and subsequently their own understanding) of complex biochemical systems. Adopting a constructivist stance, I show how learners with no prior knowledge of these systems can generate accurate, academically useful models that are shaped by the learners' interests developed over the course of model construction. I also provide evidence for the utility of this approach in improving theoretical understanding and exposing experimental gaps in the literature for the field as a whole, one which benefits from the models being constructed by non-specialists.

The atypical approach focusses more on how complex biological concepts can be abstracted into more visual and tactile environments in a way that does not result in any meaningful loss of their complexity. I present an extended case study in using the sandbox game Minecraft, where I have developed applications for its use in undergraduate learning, multi-audience outreach and academic experimentation. I show how the use of such an environment allows for creativity in a learner's approach to biology, which is often difficult to include in more traditional pedagogical approaches. I also present a pipeline for how to incorporate actual biological data into the game world, and how this may help to reduce oversimplification of concepts in gamified settings.

Together, this work demonstrates that the development process of models provides a hereto underexploited method of theoretical collation, critique and understanding which has notable benefits for both academics and educators.

ACKNOWLEDGEMENTS

Firstly, a huge thank you to my supervisor, Melanie Stefan. Thank you for letting me take ownership of these four years of work, and indulging my slightly wild whims! Your support and guidance have been invaluable throughout all of this.

This second half of this thesis would have been impossible without the hard work and insight of my student collaborators. Lewis and Kyle, thanks for your enthusiasm and for showing me how models could be used for so much more than just research. James and Siobhan, thanks for all your hard work on your Honours projects, showcasing just how much undergraduates can achieve when given the freedom to work independently.

Simon, Okuk, Chikku, and Connor, thanks for helping me start on this rather surreal journey into using games to teach, and for your insightfulness, conscientiousness, and enthusiasm. Simon, a special thanks for letting me use some of your work as inspiration, and for all your constructive comments along the way! Also Justin – thank you for all your help in making the AR applications a reality.

All PhDs need a good support network, and I want to especially thank all members of Room 176 over the years for letting me vent, chatting, and just generally being awesome people – Drs Cat Hume, Amy Warnock, Katherine Bonnycastle, Nathalie Vladis, and soon-to-be Drs Danny Schnitzler, Atanaska Velichkova and Kirsten Wilson.

Another huge thank you goes to my teaching mentors Ruth Deighton and Jane Taylor, who never let me get complacent about reflecting on my teaching practice, and helping me see just how much I enjoyed teaching! Tom Little – thanks for taking an interest in my Minecraft work and letting me add “comprehensive turtle audit” to the list of bizarre phrases that make up this PhD!

Matthew Tse, thanks for all your terrible jokes over the years, and I hope to keep our monthly pub chats for a long time to come. Laura Underwood, I am honoured to be able to call you my friend, and thank you for always stopping me from becoming boring, and for all the chats and experiences we had and will have.

Big thanks to all my family too. Mum and Dad, you’ve always let me be me, and supported me every step of the way; thank you for always being there for me. Jamie – I’ll avoid anything unnecessarily soppy here, so I’ll just say thank you for letting me spend my life with you.

As for my more leftfield acknowledgements, I wish to thank the hosts of the Off Menu podcast (James Acaster, Ed Gamble, and the Great Benito), and Adam Buxton for his podcast. Your immensely enjoyable chats have been a welcome accompaniment to my writing over the past few months.

My other great love is music, and none of this work likely could have happened without it. I also can't resist the chance to make a playlist, so for those interested, head to the data repository for this thesis

(https://github.com/yishusing/Thesis_Repository) to find some of my favourite "maximalist" music!

PUBLICATIONS AND AUTHOR CONTRIBUTION

Chapter 6 is submitted and undergoing review, as

Fitzpatrick R., & Stefan, M.I., (2021) Validation through Collaboration: Encouraging team efforts to ensure internal and external validity of computational models of biochemical pathways. *Neuroinformatics*.

Richard Fitzpatrick co-designed the paper and co-wrote the manuscript. Melanie Stefan developed the overall theory, co-designed the paper, and co-wrote the manuscript.

TABLE OF CONTENTS

<i>LAY SUMMARY</i>	III
<i>ABSTRACT</i>	V
<i>ACKNOWLEDGEMENTS</i>	VII
<i>PUBLICATIONS AND AUTHOR CONTRIBUTION</i>	IX
<i>FIGURE LIST</i>	XIII
<i>TABLES</i>	XVII
 <i>CHAPTER 0</i>	 XVIII
<i>LIST OF ACRONYMS</i>	XX
 <i>CHAPTER 1: AN INTRODUCTION TO AMPARs</i>	 1
Function	2
Structure	8
Subunits	9
Lifecycle	24
Theoretical Summary	30
Functional/Mechanistic Gaps	31
 <i>CHAPTER 2: AN INTRODUCTION TO COMPUTATIONAL BIOCHEMICAL MODELLING</i>	 33
Why Model?	33
Biochemical Modelling	42
To Build A Model	48
Types of Biochemical Model	53
 <i>CHAPTER 3: MAXIMALIST MODELLING AND AMPS</i>	 56
Models, now	57
What I talk about when I talk about Maximalist Modelling	60
A Maximalist Model of the Synapse	72
An AMPAR-centric Model of Receptor Signalling (AMPS)	77

<i>CHAPTER 4: KAPPA MODEL OF AMPARS</i>	<i>106</i>
A Recap of AMPAR Function	106
Kappa Modelling	107
Experiment #1: Model Construction and Component Verification	113
Experiment #2 Piloting the Model	121
Experiment #3 Generating Naturalistic Glutamate and Calcium Fluctuations	135
Experiment #4 Exploring Different AMPAR Population Responses to LTP	140
Experiment #5 CaMKII Parameter Set Testing	147
Experiment #6 Modelling AMPAR Trafficking during Plasticity Events	162
Experiment #7 A Merged Model of AMPAR Trafficking	174
Experiment #8 Modelling AMPAR Trafficking during Plasticity Events, Part II	181
Experiment #9 A Monomeric Model of Synaptic AMPAR Trafficking	194
General Discussion	206
 <i>CHAPTER 5: BIOLOGICAL MODELS IN TEACHING</i>	 <i>210</i>
The Nuffield Scholar Programme	211
A Biochemical Model of FMRP	211
Sonic Hedgehog Signalling	218
Constructing Knowledge Through Modelling	231
 <i>CHAPTER 6: VALIDATION THROUGH COLLABORATION</i>	 <i>234</i>
 <i>CHAPTER 7: MINECRAFTED</i>	 <i>246</i>
Learning Through Play	247
Minecraft	249
But Why Minecraft?	251
Pilot Study	252
Epidemiological Modelling in Minecraft	282
Neuron Safari	287

World Structures	292
Activity Run #1: Leith Labs	296
Activity Run #2: Edinburgh Science Festival	299
Activity Run #3: Science Festivals and Organic Visits	304
Reproduction of Animal Behavioural Experiments in Minecraft	309
MinecraftEdAR: Integrating Minecraft into the Classroom	322
Project Millport: Virtual Ecological Fieldwork	326
General Reflections	344
 <i>CHAPTER 8: EPILOGUE</i>	 346
 <i>REFERENCES</i>	 349

FIGURE LIST

Chapter 1

Figure 1 Synaptic Plasticity and AMPARs	4
Figure 2 AMPAR Subunit C-terminal Sequences.	9
Figure 3 AMPAR subunit distribution and dimerisation K_D values.....	10

Chapter 3

Figure 4 HB Model: Key Components.	80
Figure 5 HB Model: AMPAR Phosphorylation States.....	80
Figure 6 HB Replication: cAMP-AMPA relationship resembles original model	84
Figure 7 HB Replication: Receptor phosphorylation profiles and Protein Phosphatase 1 (PP1) levels in a bistability analysis are replicated exactly.	85
Figure 8 HB Replication: CaMKII switches in the PSD Compartment successfully observed.	86
Figure 9 HB Replication: Calcium and cAMP effects on CaMKII and AMPARs in PSD closely resemble the published model.	87
Figure 10 HB Replication: Stability of CaMKII and AMPAR high states in weakly coupled version closely resembles published model.....	88
Figure 11 IB Model: Key Components.	90
Figure 12 JB Model: Key Components.....	90
Figure 13 RKTVD Model: Key Components.	91
Figure 14 An Overview of AMPS Modules.	92
Figure 15 An accurate replication of Jain and Bhalla (2009) mTOR model, LTP variant ...	94
Figure 16 An accurate replication of Jain and Bhalla (2009) mTOR model, LTD variant...	95
Figure 17 The Padala et al. (2017) model can be faithfully replicated.....	96
Figure 18 Replication of Iyengar and Bhalla (1999) mGluR model is comparable in the merged model	98
Figure 19 Replication of Iyengar and Bhalla (1999) mGluR model is comparable in the merged model cont.....	99
Figure 20 Replication of Jain and Bhalla (2009) model LTP variant is only partially successful in the merged model.....	100
Figure 21 Replication of Jain and Bhalla (2009) model, LTD variant is only partially successful in the merged model.....	101
Figure 22 Replication of Padala et al. (2017) model fails in the merged model.....	102

Chapter 4

Figure 23 Model Components, Properties, and States.....	112
Figure 24 Simplified Schematic of Key Model Components.	115
Figure 25 Replication of the Dutta-Roy (2015) concerted opening model of AMPARs.	118
Figure 26 Initial AMPS rule calibration and assessment.....	120
Figure 27 Model S831 phosphorylation levels require high calcium concentrations.	123
Figure 28 Model S845 phosphorylation levels are sensitive to even low calcium concentrations	124
Figure 29 Synaptic AMPAR composition is overall similar to initial conditions when calcium is present.	125

Figure 30 Model kinase-phosphatase levels show sensitivity to small calcium level changes.	126
Figure 31 PP2B is more sensitive to calcium levels than CaMKII in the model.	127
Figure 32 Background calcium levels are not affected by injection of calcium into the model.	128
Figure 33 Active Calmodulin levels are only modestly affected by a transient increase in calcium levels.	129
Figure 34 Ubiquitination of AMPAR K868 Residues is reduced indirectly through calcium mediated PKA activation.	130
Figure 35 Very high calcium causes strong receptor phosphorylation but no change to synaptic AMPAR composition.	131
Figure 36 Very high calcium saturates CaMKII activation and increases cAMP without permanently raising basal calcium levels.	132
Figure 37 Activity-Induced AMPAR open state changes are possible in a rule-based model.	138
Figure 38 Model responses to different AMPAR compositions in basal and LTP conditions.	142
Figure 39 Slot addition during LTP induction stabilises higher synaptic AMPAR populations only when added concurrently with the induction protocol.	143
Figure 40 Original (AMPS) values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTP)	152
Figure 41 Original (AMPS) values with and without steric hinderance: AMPAR phosphorylation and PP1. (LTP)	153
Figure 42 Original (AMPS) values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTD)	154
Figure 43 Original (AMPS) values with and without steric hinderance: AMPAR phosphorylation and PP1. (LTD)	155
Figure 44 Stefan and Coomber values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTP)	156
Figure 45 Stefan and Coomber values with and without steric hinderance: AMPAR Phosphorylation and PP1. (LTP)	157
Figure 46 Stefan and Coomber values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTD)	158
Figure 47 Stefan and Coomber values with and without steric hinderance: AMPAR phosphorylation and PP1. (LTD)	159
Figure 48 The Slot-centric model variant.	162
Figure 49 The Receptor Phosphorylation model variant.	163
Figure 50 The Exocytosis model variant.	164
Figure 51 Slot model variant shows activity-driven insertion of AMPARs in the synapse.	168
Figure 52 Receptor Phosphorylation model variant shows modest stable potentiation of AMPARs in the synapse.	169
Figure 53 Exocytosis model variant fails to maintain initial AMPAR numbers.	170
Figure 54 LFS responses differ from expected LTD in all variants.	171
Figure 55 Merged Model (LTP) shows no AMPAR trafficking to the membrane.	178
Figure 56 Merged Model (LTD) suggests all protocols are too weak to maintain synaptic AMPARs.	179
Figure 57 Exocytosis variant with updated PP2A values and PP1-CaMKII interactions: activity-driven AMPAR insertion now possible.	185

Figure 58 Receptor variant with updated PP2A values and PP1-CaMKII interactions: more CaMKII activation and protocol differentiation in AMPAR numbers.	186
Figure 59 Slot variant with updated PP2A values and PP1-CaMKII interactions: strong CaMKII activation and synaptic AMPAR numbers.	187
Figure 60 Exocytosis variant with reduced GRIP numbers: greater AMPAR insertion that lasts for longer.	188
Figure 61 Receptor with reduced GRIP numbers: greater CaMKII activation and stabler GluA1/2 synaptic numbers.	189
Figure 62 Slot variant with reduced GRIP numbers: synaptic AMPAR numbers more stable.	190
Figure 63 Reduced GRIP Slot variant with no decay: even greater stability of synaptic AMPARs.	191
Figure 64 AMPA-tetramer complexes form and show stable populations in the monomeric model.	198
Figure 65 Monomeric Model: Receptor Variant (LTP): HFS can trigger formation of GluA1/2 synaptic AMPARs.	199
Figure 66 Monomeric Model: Receptor Variant (LTD): Enhanced stable populations of GluA1/2 receptors.	200
Figure 67 Monomeric Model: Slot Variant (LTP): AMPAR trafficking insensitive to induction protocols.	201
Figure 68 Monomeric Model: Slot Variant (LTD): AMPAR trafficking insensitive to LFS but triggers high levels of CaMKII activation.	202

Chapter 5

Figure 69 Lewis' model of FMRP molecular interactions.	217
Figure 70 Kyle's model of Shh signalling.	221
Figure 71 James' Shh Model.	225

Chapter 7

Figure 72 Example of student work in Minecraft.	254
Figure 73 Code quantities for each student report on the MinecraftED pilot.	256
Figure 74 Code Co-occurrence across all reports on the MinecraftED pilot.	257
Figure 75 Student B's redstone-based CaMKII model.	263
Figure 76 Video demonstration of the DNA Factory.	265
Figure 77 Siobhan's Minecraft SIR Models.	284
Figure 78 Data Conversion Pipeline for importing neuroscientific models into Minecraft.	289
Figure 79 Neuron Safari Zones.	291
Figure 80 Leith Labs Feedback: strong engagement and content comprehension.	298
Figure 81 Neuron Safari @ Glasgow Science Festival.	305
Figure 82 Neuron Safari Website Statistics.	307
Figure 83 Virtual Event Arena.	315
Figure 84 Virtual Water Maze.	315
Figure 85 Development Version of Water Maze and Event Arena.	320
Figure 86 Development Version of AR DNA Factory App.	323
Figure 87 Final Island Design for Project Millport.	330
Figure 88 Beach Maps.	332

Figure 89 Fintray Beach.....	332
Figure 90 EcoLab.	335
Figure 91 Tutorial Island Zone.	335
Figure 92 Millport Station.....	339
Figure 93 Island Train Network.	339
Figure 94 Millport Aesthetical Additions.	340

TABLES

Chapter 4

Table 1 AMPAR compositions in the model	141
Table 2 Variable testing sets for CaMKII interactions	148

Chapter 5

Table 3 Species and concentrations of Lewis' model.	213
Table 4 Reactions and parameters of Lewis' model.....	214
Table 5 A list of species in Kyle's model.	222
Table 6 Reaction rate parameters for Kyle's model.	223
Table 7 Species list for James' model.....	226
Table 8 Reaction list for James' model.....	227

CHAPTER 0

General Introduction

This thesis has the overarching aim of emphasising the utility of biological modelling to address theoretical questions within neuroscience. Within this remit, I intend to introduce several different modelling approaches to understand the activity of AMPA receptors (AMPARs) at the level of a single synapse, as well as how models more generally can be utilised to aid in the understanding and clarification of complex biological concepts.

My approach to writing this document has been that it should encompass the whole range of my PhD training, which included a considerable amount of teacher training at a tertiary level. This is my justification for the latter half of my remit, and includes work carried out with students at the University of Edinburgh. A model that cannot be understood is no better (perhaps worse) than no model at all, and I believe the insights gleaned from this aspect of my work feeds into how computational modelling within biology could progress.

These threads of neuroscience and education have synthesised into my current and potentially far-reaching work of utilising sandbox game environments as a device for teaching biomedical concepts within university and outreach settings. I explore this approach through detailing several ongoing activities, and thus look to where I intend to take the skills and tools from this PhD in my career.

In **Chapter 1** I will outline the theoretical understanding to date about AMPAR structure and function, with a specific focus on synaptic plasticity. The continual thread across all the studies here has been how much of synaptic plasticity can be attributed to the modifications made just to the AMPA receptor, and begin to determine the importance of its many different compositions in regulating learning.

Chapter 2 introduces computational modelling within biology. This serves as an introduction to the concept of modelling in biology, which exists largely to explore specific questions within a diverse range of subfields. I attempt here to draw together the commonalities of these approaches and argue that broader, more question-driven (as opposed to hypothesis-driven) models are needed to facilitate

the general acceptance of modelling as equal to, and not just supportive of, experimental research.

The methodology behind the two major models I built are the respective focus of **Chapters 3 & 4**. Both attempt to create a larger model of synaptic function from the synthesis of pre-existing models. I critically assess this methodology and the difficulties that arise from assumptions and missing information from model source material that in some cases make this additive method more difficult than constructing a new model. These chapters also cover the specific way in which each model was written. One was constructed in a well-used software (COPASI). The second is in a newer and less intuitive language but is nonetheless more powerful and geared towards the size of models generated in this PhD.

My teaching of biological ideas using computer models constitutes **Chapter 5**. This touches upon the vocational aspect of my PhD, which required at least 25 % of time be spent towards teaching-related activities. I have highlighted here the teaching I carried out that made use of the modelling concepts and software outlined in Chapters 2 and 3. These involved the supervision of undergraduate Honours students and high school Nuffield scholars in generating their own models of signalling pathways.

Chapter 6 is a commentary article that argues for a way of validating models both internally (can the model be reproduced?) and externally (how close is the model output to experimental evidence?). I link this with the work in Chapter 5 to show how students are more than capable of carrying out both methods of model verification, and how this would impact model development.

Chapter 7 outlines my investigations into the use of Minecraft for the modelling of biological concepts. This was initially run with a small group of undergraduate biomedical science students from the University of Edinburgh, and has since expanded into a larger, wide-reaching project that fuses complex data and theory with a user-friendly and widely used software platform.

Finally, in **Chapter 8** I provide a summary of the whole thesis.

LIST OF ACRONYMS

Acronym	Name
ADAR2	Adenosine Deaminase RNA Specific 2
AKAP	A-kinase anchoring protein
Akt	Protein kinase B
AMPA	alpha-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid
AMPA	AMPA receptor
APC ^{Cdh1}	Anaphase-promoting complex
ApoER2	Apolipoprotein E receptor 2
ASD	Autism Spectrum Disorder
BDNF	Brain-derived neurotrophic factor
CA1	cornu Ammonis 1
CaM	Calmodulin
CaMKII	Ca ²⁺ /calmodulin-dependent protein kinase II
CaMKIII	Ca ²⁺ /calmodulin-dependent protein kinase III
cAMP	Cyclic adenosine monophosphate
CASK	Calcium/calmodulin-dependent serine protein kinase 3
CNIH	Protein cornichon homolog
CYFIP	Cytoplasmic FMR1-interacting protein
DAG	Diacylglycerol
eIF4E	Eukaryotic translation initiation factor 4E
EPSPs	Excitatory post synaptic potentials
ERK	Extracellular signal-regulated kinase
F-BAR	Fes/CIP4 homology-Bin-Amphiphysin-Rvs
fEPSP	Field excitatory postsynaptic potential
FMRP	Fragile X Mental Retardation Protein
FRAP	Fluorescence recovery after photobleaching
FRAS1	Fraser syndrome 1
FXRP	Fragile X mental retardation syndrome-related protein
FXS	Fragile X Syndrome
GRIP	Glutamate receptor-interacting protein

HEK	Human Embryonic Kidney
HMG-CoA	β -Hydroxy β -methylglutaryl-CoA
K_D	Dissociation constant
KIF5	Kinesin-related protein 5
LTD	Long term depression
LTP	Long term potentiation
MAPK	Mitogen-activated protein kinase
MFB	Medial forebrain bundle
mTOR	Mechanistic target of rapamycin
Nedd	Neural precursor cell expressed, developmentally downregulated
NEEP21	neuron-enriched endosomal protein 21
NMDAR	N-methyl-D-aspartate receptor
NPRAP	Neural plakophilin-related armadillo protein
NSF	N-ethylmaleimide-sensitive factor
ODE	Ordinary differential equations
PDZ	Post synaptic density protein; Drosophila disc large tumour suppressor; zonula occludens-1 protein
PICK1	Protein interacting with C-kinase1
PKA	Protein kinase A
PKC	Protein kinase C
PP1	Protein Phosphatase 1
PP2A	Protein Phosphatase 2A
PP2B	Protein Phosphatase 2B
PSD	Postsynaptic Density
RNF167	Ring finger protein 167
S6K	Ribosomal protein S6 kinase
SBML	Systems biology Markup Language
SED-ML	Simulation Experiment Description Markup Language
Shh	Sonic Hedgehog
TARP	Transmembrane AMPAR regulatory protein
TSC2	Tuberous Sclerosis Complex 2

CHAPTER I

AMPARe: An Introduction

The α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) receptor is a ligand gated ion channel that allows the flow of sodium, potassium, and sometimes calcium ions into a cell. The action of this receptor has long been linked to the success of information transmission between neurons, and still has a dominant role in current theories of synaptic plasticity. This chapter briefly sketches the current theory of AMPA receptors' (AMPARe) structure and function, focussing finally on their function within plasticity and memory¹.

In the past 20 years alone, over 320 review articles and almost 4000 research papers have been written where AMPA receptors are the major focus². Such intense research has led to a rapid acceleration in our understanding not just of the receptor, but also excitatory synapses in general. The rest of this chapter will try to capture a snapshot of our current understanding, which will help make sense of the models introduced in Chapters 3 and 4.

¹ Readers should look to (Lodge, 2009) for an excellent history of the discovery of ionotropic glutamatergic receptors more generally.

² Based on a filtered search of the NIH PubMed database with "AMPA Receptors" as the primary MeSH search criteria.

Function

The primary function of AMPARs is to control fast synaptic transmission within the central nervous system. Binding of glutamate allows for influx of sodium ions and the generation of excitatory post synaptic potentials (EPSPs). Heightened responses by a neuron to an input are achieved through modulation of the properties and composition of synaptically localised AMPARs (Diering & Huganir, 2018).

All the theory discussed here, unless explicitly mentioned, refers to hippocampal neuronal AMPAR function. The broad mechanisms may hold across different brain regions and cell types, but there is likely to be multiple minor changes between them.³

Synaptic Unsilencing

The major 'job' often ascribed to AMPARs in development is to provide a voice to the synapse. Signalling events from the presynaptic neuron need to be communicated and transferred down the dendrite for integration and potential transmission elsewhere in a neural network (**Figure 1**). Prior to AMPAR insertion into the synapse, the spine is 'silent'. NMDARs can interact with glutamate but are normally prevented from allowing ion influx due to a Mg^{2+} blockade in the pore region. Only a very large depolarisation event can relieve this blockade and allow sodium, potassium and calcium ions into the spine (Isaac, Nicoll, & Malenka, 1995).

The calcium influx triggers a signalling cascade which leads to rapid insertion of nearby AMPARs into the synaptic membrane. The calcium ions bind to calmodulin, which in turn activates CaMKII. The activity of this kinase results in the activation of exocytotic machinery and vesicle transfer to the cell membrane (Huganir & Nicoll, 2013). The presence of AMPARs reduces the required threshold for postsynaptic firing considerably, and thus 'unsilences' the synapse. This was thought to largely be a one-off event- once a synapse possesses AMPARs, the change is irreversible. Yet recent evidence suggests that this may not be a universal law (B. R. Lee et al., 2013; Xiao, Wasling, Hanse, & Gustafsson, 2004). What is clear however is that the presence of AMPARs is a requirement for the type of synaptic activity critical for learning and memory in the brain.

³ There are to date few comparative reviews exploring the differences between AMPAR function between brain regions (see Hey-Kyoung Lee & Kirkwood (2011) for an example). Cerebellar functionality and subunit composition are considerably different however to cortical regions, so a review is overdue. Gallimore et al., (2016, 2018) provide a good overview in the meantime.

Synaptic Plasticity

Synaptic plasticity is the modulation of synaptic response thresholds or magnitudes. These create fine-grained, dynamic synaptic profiles which can be taken as components of memory when situated within a network of neurons. There are two reliable and observable readouts of synaptic plasticity – spine morphology (Desmond & Levy, 1990; Engert & Bonhoeffer, 1999; Maletic-Savatic, Malinow, & Svoboda, 1999) and changes to surface proteins (Bosch et al., 2014; H.-K. Lee & Huganir, 2008). Both are linked (Hruska, Henderson, Le Marchand, Jafri, & Dalva, 2018), but here we will focus on how AMPARs are modulated in plasticity events (**Figure 1**).

The postsynaptic mechanism behind the role of AMPARs in plasticity is interpreted currently through two positions (Tim V.P. Bliss & Collingridge, 2013). The first to be expressed was that AMPARs are modified to alter their function during plasticity, with both channel conductance (T. A. Benke, Lüthi, Isaac, & Collingridge, 1998) and open probability implicated (Banke et al., 2000; Greengard, Jen, Nairn, & Stevens, 1991). Second, the idea that the number of AMPARs might change and be the primary mechanism by which plasticity is realised was also proposed (T. A. Benke et al., 1998), but confirmation experimentally came later (W. Y. Lu et al., 2001) and this idea has arguably been more widely accepted and studied in the years since (see Bliss & Collingridge, (2013) for a historical review).

Plasticity is split into two forms: long-term potentiation (LTP) and long-term depression (LTD). To best explain the different responses of AMPARs to these phenomena, I provide a brief description below.

LTP

Activity-driven movement and modification of AMPARs has been shown repeatedly to be a necessary component of hippocampal neuronal LTP (taken here to be a sustained increase in postsynaptic responses to inputs) (Diering & Huganir, 2018; Huganir & Nicoll, 2013). Basally the synaptic zone will contain a mixture of GluA1/2 and GluA2/3 AMPARs (Wenthold, Petralia, Blahos, & Niedzielski, 1996). Upon patterned high synaptic signalling activity, both NMDA-dependent and




Figure 1 Synaptic Plasticity and AMPARs.

One possible measure of synaptic plasticity is the number of AMPARs present in the postsynaptic membrane. (A) Active synapses will show the clustering of AMPARs and NMDARs directly below the presynaptic vesicle release sites. (B) Under strong patterns of incoming activation, the number of AMPARs increases, largely GluA1/2 AMPARs. (C) Low frequency activation gradually promotes the removal of AMPARs from the synaptic zone, returning to extrasynaptic sites or being degraded by the ubiquitin-proteasome pathway.

independent calcium level changes can trigger the process of AMPAR modifications and transport necessary for early LTP induction and maintenance.

These changes act to begin the process of increasing the spine volume, and therefore the area of the synaptic zone (Holtmaat & Svoboda, 2009). AMPARs are required to both induce and maintain this change. The induction appears to result in an increase in GluA1-containing receptors, and a possible decrease in GluA2/3 receptors (Tanaka & Hirano, 2012). This GluA1 population, given the rapid (within minutes) onset of this increase is thought to come from perisynaptic pools of GluA1-homomeric receptors (Diering & Huganir, 2018). These allow for great changes to calcium levels in the spine and may boost the likelihood of activating the necessary downstream pathways for spine growth, protein synthesis and actin remodelling (Jędrzejewska-Szmek & Blackwell, 2019; Matsuzaki, Honkura, Ellis-Davies, & Kasai, 2004).

These homomeric receptors appear to be only transiently expressed and are soon replaced by GluA1/2-containing AMPARs. These are calcium impermeable and less conductive than the homomers (T. Benke & Traynelis, 2019). They are driven into the synaptic zone by phosphorylation profiles of the GluA1 subunits either directly or indirectly leading to stabilisation (Diering & Huganir, 2018). The movement of GluA1/2 receptors is commonly referred to as being activity-dependent, which is that changes to signalling patterns are thought to induce movement of this receptor type. This I think is more accurately described as being sensitive to kinase and phosphatase activity – trafficking is altered significantly when the kinase-phosphatase balance is adjusted, independent of if this is driven by changes to synaptic activity (Sanderson, Scott, & Dell'Acqua, 2018).

The maintenance phase of LTP is so named because the changes to biochemical networks downstream of AMPAR signalling are unstable until sufficient changes in copy number of certain proteins has been achieved. In this phase, any perturbation to the trafficking or stability of AMPARs in the synaptic zone usually results in the failure of LTP (Boehm et al., 2006; D.-T. Lin et al., 2009).

LTD

The main outcome of LTD is the controlled reduction of synaptic efficacy. There are multiple reasons why such an outcome is desirable. One is that the neuronal connection is no longer very active – maintaining a large spine for a weakly active connection is energetically costly for a neuron (Harris, Jolivet, & Attwell, 2012).

The second is that LTP events can reach a ceiling where efficacy is at a maximum⁴ (Collingridge, Peineau, Howland, & Wang, 2010; Sheng & Ertürk, 2014).

The response of AMPARs to an LTD induction event is largely observed as a removal of AMPARs from the synapse. This is a two-step process – the removal from the synaptic zone, and the removal of receptors from the extrasynaptic zone via endocytosis (Moretto & Passafaro, 2018; Shepherd & Huganir, 2007). This reduces the ability for the synapse to respond to presynaptic activity, and therefore limits the signal transduction occurring.

AMPA removal is thought to be mediated by PICK1, and more specifically its F-BAR domain. S880 phosphorylated, GluA2-containing AMPARs show much reduced GRIP binding, whilst maintaining PICK1 binding. The F-BAR domain interacts with the membrane and associated proteins to allow for clathrin-type endocytosis mechanisms to promote AMPAR removal (Fiuza et al., 2017; Kastning et al., 2007).

GluA1-mediated removal of AMPARs is less well characterised, but S845 phosphorylation appears to be a necessary step for LTD-induced removal. Additional evidence comes from CaMKII phosphorylation of S567, which is in the first intracellular loop of the GluA1 subunit. This phosphorylation site is only phosphorylated by the non-Ca²⁺/CaM-bound, autophosphorylated (at T286) form of CaMKII, which is sustained by low but persistent activity associated with LTD (Coultrap et al., 2014).

Homeostatic Plasticity

The above two forms of synaptic plasticity are usually classified as “Hebbian” (Hebb, 1949). This is an associative form of plasticity, usually pithily defined as “cells that fire together, wire together”⁵. This idea however predicted a problem: positive feedback loops which would lead to ceiling effects (for LTP) or complete loss of the synapse (in LTD). This suggested that additional mechanisms were necessary to explain both why and how synapses did not reach these ceilings, which would drastically limit the informational capacity of the brain over time (Martin, Grimwood, & Morris, 2000; Moser & Moser, 1999).

Homeostatic plasticity was theorised as a solution to this problem. This is a mechanism by which synaptic ‘weights’ (the relative strength of each synapse) are

⁴ This falls, I admit, into mere depotentiation/homeostatic plasticity territory. LTD is less well understood than LTP, and also appears at a glance to involve many more mechanisms than have been supposed for LTP.

⁵ Derived from a phrase coined in Löwel & Singer, (1992): “neurons wire together if they fire together”

normalised to avoid catastrophic loss or lack of gain. This acts on a principle of negative feedback – activity of a given synapse is returned to a set point over time to maintain network stability. Two models explaining this are sliding threshold and synaptic scaling (Turrigiano, Leslie, Desai, Rutherford, & Nelson, 1998) but was not conclusively shown to occur *in vivo* until relatively recently (see H. K. Lee & Kirkwood, 2019 for a list of studies). These studies all show a clear *in vivo* phenomenon that is explainable through one or both homeostatic plasticity models.

Many papers exploring the relationship between Hebbian and homeostatic plasticity highlight the high level of convergence of the two mechanisms onto known biochemical pathways (Fox & Stryker, 2017). There is also a consensus that the two act at different timescales. The idea that one may be a form of the other has until recently been dismissed/not considered (although see Dammasch, 1989). This is being challenged by evidence that associative properties can emerge from homeostatic plasticity (Gallinaro & Rotter, 2018).

These might be not two competing mechanisms, but the same mechanism being explored at different timescales and in different locations (Galanis & Vlachos, 2020; Vianna Gallinaro, 2020). Homeostatic plasticity is better equipped to explain memory as most people use the term, with Hebb or Hebb-like plasticity better equipped for the special cases of working memory and information encoding. This latter term describes the updating of memories depending on the predictions made by the homeostatic instantiation of it, one which over time becomes less important as synapses stabilise, and the system becomes less plastic.⁶

⁶ The additional advantage of homeostatic memory is that unlike Hebbian plasticity forgetting is the default mode, and any stabilisation requires the repeated presence of a stimulus, or a way of activating it in part. These “dispositional” memories (term taken from De Brigard (2014)) are attractive because they do not require any special biomolecular storage mechanism, the synapse is important only for encoding/updating, and they are much better able to explain semantic memory and its necessity in allowing stabilisation through partial activation.

Structure

AMPARs are tetramers made of four subunits (GluA1-4)⁷ The relative composition varies prominently across brain regions (Schwenk et al., 2014), but in the hippocampus is predominantly a heteromer of GluA2 and either GluA1 or GluA3 subunits. The differences in structure between each subtype is only minor, but sufficient to result in marked changes in functionality and in the localisation of the receptor.

Each subunit is comprised of four transmembrane domains, with a large (> 8(50% of the total protein) extracellular N-terminal domain and smaller intracellular C-terminus (García-Nafria, Herguedas, Watson, & Greger, 2016). Part of the extracellular loop between domains III and IV forms a portion of the ligand binding domain (LBD) along with the N-terminus. All major binding domains for intracellular interaction lie on the C-terminus (**Figure 2**). These regions are prone to phosphorylation, ubiquitination and palmitoylation (Anggono & Huganir, 2012).

Most AMPARs are impermeable to calcium. A single amino acid change in the TMII domain of the GluA2 subunits in the post-translation mRNA is the root cause of this (Henley & Wilkinson, 2016). Impermeability guards against accidental excitotoxicity from large influxes of calcium in potentiation events. The absence of GluA2 subunits results in receptors with a higher channel conductance and probability of being in the open state, as well as calcium permeability (Anggono & Huganir, 2012).

These are all subtly different. Channel conductance is a measure of ion flow through the receptor channel (Bezannila, 2008). The probability of being open determines if the channel pore, through which ions flow across the membrane, prevents this flow. This is stabilised into open configurations through the binding of endogenous ligands (glutamate for AMPAR). Calcium permeability is the ability of calcium to pass completely through the channel pore – the abovementioned amino acid change (glutamine to arginine), makes the pore more positively charged, repelling calcium from passing through (Pachernegg, Munster, Muth-Kohne, Fuhrmann, & Hollmann, 2015).

⁷ The nomenclature of subunits shifts between GluRx and GluAx depending on the researcher. There appears to have been some shift towards A over R in the past decade, but it is not enforced. I will stick to Ax in this thesis.




Figure 2 AMPAR Subunit C-terminal Sequences.

The major identified amino acid residues modified post-translationally for all for AMPAR subunits.

Thus, the presence of certain AMPA subunit configurations at the synapse helps determine the current activity profile of the neuron itself, and is something we shall see later in relation to receptor trafficking during above-baseline activity.

Subunits

GluA1

GluA1 is found in practically all brain regions. It is particularly abundant in the hippocampus and olfactory bulb, where it has almost equal abundance with GluA2 (**Figure 3**) (Schwenk et al., 2014). In the hippocampus this correlates well with the 80:20 ratio of GluA1/2:GluA2/3 AMPARs considered to be the composition of AMPARs in this region (W. Lu et al., 2009). Whilst the kinetics of dimerisation favour heterodimers, where large numbers of GluA1 subunits are produced GluA1 homomeric AMPARs have been detected and even implicated in functional roles within synaptic plasticity (Plant et al., 2006; Sanderson, Gorski, & Dell'Acqua, 2016). The GluA1 subunit is perhaps the most well-studied of the four subunits. Three phosphorylation sites have been implicated in plasticity. These sites (S831, S845, and S818) are acted upon by three separate kinases-CaMKII, PKA, and PKC, respectively.




Figure 3 AMPAR subunit distribution and dimerisation K_D values.

Top: The distribution of AMPAR subunit types across the brain, adapted from Schwenk et al., 2014. The size of the chart relates to the density of functional AMPARs. Adapted from Schwenk et al., (2014), with permission from Elsevier. Bottom: Derived K_D values for dimerisation of AMPAR subunits at their amino-terminal domains. GluA1/2, GluA2/3 and GluA1/3 dimers are energetically the most favoured dimer forms. Data from H. Zhao et al., 2017.

1997; Derkach, Barria, & Soderling, 1999; Kristensen et al., 2011). Phosphorylation of the S831 residue has been implicated in the increasing of the conductance of a given AMPAR (Barria, Muller, Derkach, Griffith, & Soderling, 1997; Derkach, Barria, & Soderling, 1999; Kristensen et al., 2011). There is some debate as to whether this residue also confers stabilisation of AMPARs within the synaptic membrane (H.-K. Lee & Huganir, 2008). S845 is implicated in trafficking AMPARs to the membrane in the phosphorylated state (Oh, Derkach, Guire, & Soderling, 2006) and away from the membrane via endocytosis in the unphosphorylated state (Man, Sekine-Aizawa, & Huganir, 2007). The phosphorylation of S845 also increases single channel open probabilities (Banke et al., 2000). S818

phosphorylation appears to play a role in this recruitment step to the membrane and may also be contributing to the maintenance of AMPARs within the synaptic zone (D.-T. Lin et al., 2009).

It is the modification of these three sites which appears to be critical for the induction and perhaps maintenance of long-term potentiation (LTP) within a dendritic spine. Convincing evidence comes from Zhou et al., (2018) where the C-tail of a GluA1 subunit was spliced onto GluA2 subunits and was sufficient to induce LTP phenomena at the synapse. Conversely, the splicing of a GluA2 tail onto a GluA1 subunit did not result in LTP induction. This suggests that the channel element of the AMPAR does not play a major role in modulating its function and that the GluA1 C-tail is both necessary and sufficient for hippocampal LTP. There is some conflicting evidence however that the trafficking of extrasynaptic AMPARs, regardless of subunit type, is a necessary step for long-term potentiation (Granger, Shi, Lu, Cerpas, & Nicoll, 2013). Some reconciliation of these findings may lie in the new discovery of control elements on the N-tail of AMPARs (see next section).

In the case of S831 and S845 phosphorylation, these mechanisms may be individually sufficient to result in LTP. There appears to be a basal level of phosphorylation at these two sites, although estimates for this range anywhere between almost 0 through to 15 % of membrane bound AMPARs, and are averages across a large number of synapses (Diering, Heo, Hussain, Liu, & Hugarir, 2016; Hosokawa, Mitsushima, Kaneko, & Hayashi, 2015; Oh et al., 2006). The range of different results may be reflective of methodological problems, or reflective of the highly transient nature of these modifications. No studies to date show any real-time analysis of phosphorylation, but all show an increase in phosphorylation in an activity-dependent manner. It may be the case that basal phosphorylation does not play a specific role, but may be reflective of recent synaptic activation patterns⁸, or be reflective of subunits phosphorylated by freely moving PKA complexes (Michel & Scott, 2002), and play no role in basal transmission.

Developmentally, PKA-mediated and CaMKII-mediated plasticity work in parallel. The removal of one pathway in young animals has only minor impacts of observed hippocampal learning, with at least two of the three main kinases (PKC being the third) needing to be removed to block LTP induction (Wikström, Matthews, Roberts, Collingridge, & Bortolotto, 2003). Both PKA and CaMKII are

⁸ We shall see later in Chapter 4 that in a simplified synaptic environment basal phosphorylation does indeed appear to be close to zero if input patterns are not plasticity-inducing.

able to enter bistable states independent of the other biochemically (Kotaleski & Blackwell, 2010; Song, Smolen, Av-Ron, Baxter, & Byrne, 2006). In later adult animals, CaMKII appears to be predominant, being both necessary and sufficient for LTP induction (Bortolotto & Collingridge, 1998; Incontro et al., 2018). This does not appear to be completely true however, as particular LTP inductions require PKA in older animals (Blitzer, Wong, Nouranifar, Iyengar, & Landau, 1995; Y. Lu et al., 2007).

Interestingly these mechanisms appear to be differentially preferred depending on cell type and brain region. Neocortical regions show a much greater sensitivity to changes in PKA expression, with inhibition leading to failures of LTP and LTD phenomena (Liu, Rao, & Daw, 2003; Seol et al., 2007). PKA mediated plasticity elsewhere is important up until a few weeks post-natally, when the effect of PKA blockade is limited to late maintenance phases of plasticity (H. K. Lee & Kirkwood, 2011).

The role of S818 phosphorylation appears to be restricted to activity-dependent exocytosis. PKC, which is activated in an LTP induction protocol, triggers an increase in the phosphorylation levels of S818. This allows for 4.1N protein binding, which is correlated with an increase in membrane localised GluA1 subunits (D.-T. Lin et al., 2009). Preventing S818 phosphorylation or 4.1N binding show quicker hippocampal LTP decay back to baseline suggesting a critical importance of this exocytotic control in maintaining LTP (Boehm et al., 2006; D.-T. Lin et al., 2009). It should be noted that beyond these two studies there has been no further published experiments on S818 phosphorylation, so unlike S831 and S845 the functional sketch is much less complete.

Whilst dephosphorylating S845 is a necessary step for hippocampal LTD of GluA1 subunits, there is a requirement for this site to be *phosphorylated* for LTD to occur (Huganir & Nicoll, 2013). This stems from a study where an S845A mutation (preventing phosphorylation) abolishes the induction of LTD (H. K. Lee, Takamiya, He, Song, & Huganir, 2010). The mechanistic reasons for this are still not definitively proven, but it appears that the initial kinase activity either recruits or promotes the mechanisms that trigger removal of GluA1 subunits from the membrane upon dephosphorylation.

The best evidence to date comes from Sanderson, Gorski, & Dell'Acqua (2016), who investigated the function of AKAP-bound PKA on AMPARs during LTD. They show that PKA activity on calcium permeable (non GluA2-containing) AMPARs situated at perisynaptic sites allows them to occupy the synaptic zone. This receptor incorporation enhances activation of phosphatases such as calcineurin, which is also bound to AKAP. This acts as a feedback loop and

removes these receptors as well as inducing a destabilizing wave of calcium into the synaptic network. This brief disruption of the calcium dynamics increases the probability of the endocytosis of stabilised AMPARs within the membrane without requiring further presynaptic inputs, which would be more likely to promote potentiating dynamics.

GluA2

GluA2 appears to be one of the major component subunits of AMPARs across a wide number of neuronal types, combining commonly with GluA1 or GluA3. Most brain regions possess large copy numbers of GluA2 compared to the other subunit types apart from the brainstem (where GluA4 is more abundant) and the cerebellum (where there is almost no GluA2) (**Figure 3**) (Schwenk et al., 2014).

Unlike GluA1, the C-terminal of these subunits is relatively short and does not possess as many modification sites, with three phosphorylation sites and two ubiquitination sites (Diering & Huganir, 2018). This subunit forms interactions via its PDZ domains with molecules such as GRIP1 and PICK1 (Dong et al., 1997; Xia, Zhang, Staudinger, & Huganir, 1999), which are an additional layer of control of the movement of AMPARs in membranes. These two molecules work in a complex dynamic coupling, being able to bind to one another (W. Lu & Ziff, 2005) and regulate their function as well as bind independently to GluA2 (and GluA3) subunits (Bissen, Foss, & Acker-Palmer, 2019)⁹.

Much like the major role GluA1 appears to play in hippocampal LTP, GluA2 seems to be the major player in LTD (at least in NMDA-dependent forms of LTD). Splicing of GluA2 C-tails onto GluA1 subunits both reduced LTP and enhanced LTD (Zhou et al., 2018). For the GluA1/2 AMPARs that predominate hippocampal synapses, it may be that the two subunits act in dynamic opposition to one another, with the kinase/phosphatase balance of the surrounding environment determining which subunits determine the short-term fate of the receptor.

The GluA2/3 AMPAR subpopulation appears to be electrically quiet in basal conditions, requiring cAMP-mediated pathways to become active (Renner et al., 2017). A theory still reported in articles (first proposed in Shi, Hayashi, Esteban, & Malinow, 2001) is that that these AMPARs are not only constitutively cycled in the membrane, but slowly replace the more sensitive GluA1-containing receptors over time. However, the population of non-GluA1 containing AMPARs is significantly lower (W. Lu et al., 2009), and the copy numbers of GluA3 make it difficult to see how this replacement would operate (Helm, 2019). Renner et al.

⁹ This review is one of the best collection of papers on the topic of GRIP and PICK interactions with AMPARs.

(2017) however, suggest that they may contribute to plasticity, with the advantage to be electrically quiet in basal conditions. This would allow for stable large spines that can respond to differences in presynaptic firing profiles whilst overall being resistant to spontaneous firing fluctuations leading to destabilisation of the spine.

Regarding long-term depression, there appears to be a phosphorylation site which plays a role in the expression of LTD. The S880 residue when made phospho-null results in a reduction in long-term depression (Seidenman, Steinberg, Huganir, & Malinow, 2003). This is likely due to this site regulating the receptors' interaction with GRIP1 and PICK1 proteins. There is evidence that the phosphorylation of the site prevents the binding of GRIP1 and promotes binding to PICK1 instead (Seidenman et al., 2003). This destabilizes the receptor and promotes endocytosis, although the exact mechanism for this is yet to be fully understood.

GluA3

The evidence for how GluA3 subunits contribute to AMPAR function both basally and in dynamic events is still not well characterised. They have a very low open probability under basal conditions (Renner et al., 2017) and yet as previously stated they may play heightened roles in synapses after plasticity events, particularly LTP events. There is growing evidence however that cAMP may mediate a pathway which increases the open probability of AMPARs containing this subunit, therefore adding an extra layer of control changes to the biochemical network through the regulation of these relatively inert channels (Renner et al., 2017). This mediation would be quite rapid and may be a useful way of increasing the stability of synapses without sacrificing the ability to respond to major changes as a result of increased presynaptic activity.

In LTD the S885 site acts similarly to the S880 GluA2 site (Diering & Huganir, 2018). Knock-out studies indicate GluA3 has no non-compensated role in LTD and basal transmission, but does lead to enhancement of LTP (Meng, Zhang, & Jia, 2003)¹⁰. In disease, GluA3 knock-outs *in vitro* and *in vivo* slow the Amyloid- β -induced cognitive decline in mice, implicating this subunit in human Alzheimer's disease outcomes (Reinders et al., 2016).

GluA3 subunits preferentially dimerise with GluA1 and GluA2, with a higher K_D for GluA4 (**Figure 3**). Homomers are very unlikely given the extremely high K_D in comparison (H. Zhao et al., 2012), and the normal co-localisation with other subunit types. Intriguingly, whilst GluA1/3 heteromers are not kinetically

¹⁰ Behaviour in GluA3 KOs is impaired however, suggesting the receptor is not redundant and may be necessary to mediate dopamine and serotonin levels through reducing synaptic excitability (Adamczyk et al., 2012)

unfavourable (H. Zhao et al., 2017), and have been detected in small amounts in hippocampal neurons (Wentholt et al., 1996), at the time of writing no follow-ups have been published on the possible function of these receptor types. In the absence of GluA2 subunits this would be the preferred dimer and thus tetrameric population, so it is quite likely there are specific roles to play for this receptor subtype in neuronal or glial cells.¹¹

GluA4

Homomeric GluA4 dimers and tetramers have a highly unfavourable K_D if other subunits are present but do form under controlled experimental conditions (Yu, 2018; H. Zhao et al., 2017). The existence of GluA4 homomers has been implicated as being a possible explanation for multiple experimental outcomes, although to date there is only indirect evidence that these homomers exist naturally *in vivo* at all (Lesperance, Yang, & Wang, 2020; J. J. Zhu, Esteban, Hayashi, & Malinow, 2000).

GluA4 subunits are highly expressed in the first few post-natal days in rodent hippocampal neurons (J. J. Zhu et al., 2000), and appear to mediate the unsilencing of immature synapses in the absence of the GluA1-containing AMPARs. These subunits are much more sensitive to spontaneous activity (J. J. Zhu et al., 2000). GluA4 incorporation into the membrane is mediated by PKA, which phosphorylates the S842 residue in an activity dependent manner (Carvalho, Kameyama, & Huganir, 1999; Esteban et al., 2003).

GluA4 preferentially dimerises with GluA1 and GluA3 subunits. GluA2 binding is possible, but the K_D is higher than for all other subunits (including GluA2 homomer formation; **Figure 3**) (H. Zhao et al., 2017). GluA1/4 heteromers have been detected at interneuron glutamatergic synapses, which unlike the hippocampal principal cells see a shift from GluA1 homomers to the GluA1/4 heteromers two weeks post-natally. (Akgül & McBain, 2016; Pelkey et al., 2017). These receptors possess extremely fast kinetics and high Ca^{2+} permeability, which may be essential for the tight regulation of principal cell firing and the generation of oscillatory activity (Pelkey et al., 2017).

Dimerisation and Tetramerisation

The formation of dimers in all subunits known appears to be mediated by the N-terminal domain. More specifically, this has been proposed to be due to the binding of the two ligand binding domains (Gan, Salussolia, & Wollmuth, 2015). The affinity of subunit binding at these sites is quite high, ranging from 2.9 nM to

¹¹ I should note for any distant future readers that some evidence for this is in a thesis here: <https://jscholarship.library.jhu.edu/handle/1774.2/62424>, but is embargoed until 2024.

5.2 μM (**Figure 3**). Whilst there is no explicit physical reason preventing the formation of homodimers, as we have seen their existence is thought to be rare, with only GluA1 homomers to date being reliably detected *in vivo*. This is driven by higher affinities for inter-subunit type interactions than intra-subunit type (**Figure 3**). Very recent data (Y. Zhao, Chen, Swensen, Qian, & Gouaux, 2019) however shows that both homo -dimer and -tetrameric GluA2 configurations are detectable *in vivo*.

Tetramers form not from N-terminal domain interactions but from contact points across the whole subunit. Removal of the N-terminal domain does not prevent the formation of tetramers (Pasternack et al., 2002), and isolated N-terminal and ligand binding domain regions do not form tetramers (Gan et al., 2015). This suggests a key role for the transmembrane domain in tetramer formation (Salussolia et al., 2013).¹² The prevailing dimer to dimer theory dictates that tetramers form only from dimers, and is well supported in a number of structural studies (Mayer, 2011; Salussolia et al., 2013). Recent analyses of AMPAR subunit stoichiometry however has shown that trimers and monomers are viable forms outwith the endoplasmic reticulum (Yu, 2018). These findings suggest that whilst dimer to dimer receptor formation is the most common, it is not the only manner of receptor construction¹³

Returning to Y. Zhao et al., (2019), a large proportion of detected receptors were triheteromeric, with confirmed evidence for an A1A2;A3A2 receptor, and indirect support for A1A2;A4A2, A4A2;A3A2, A2A4;A1A2 and A2A4;A3A2. These findings strengthen a recent radical new hypothesis which we will discuss throughout this chapter. This experimentally backed hypothesis (Morise et al., 2019) (henceforth transient-AMPAR) proposes AMPARs to not exist as tetramers for any longer than a few hundred milliseconds, and instead to be continually forming transient complexes that allow for rapid responses to system perturbations without any need for intricate control mechanisms.

Ligand Binding

Each subunit possesses a ligand binding domain. These are clamshell-type structures which in a dimer are back-to-back. The binding of glutamate to a ligand binding domain induces the lower lobe of the domain to detach from the neighbouring one, leading to channel opening (Popescu, 2014).

¹² We saw previously that the C-terminal region appears to have no role in AMPAR assembly.

¹³ This finding is particularly important given very recent evidence that AMPARs are unstable complexes, which is discussed later.

The mechanical nature of how this opening operates is debated. One idea is that each subunit may open or close independently of the others. This gives five different states of opening. Glutamate's presence in a domain stabilises the open state for a quadrant of the channel. Each glutamate that binds increases the probability of opening (Rosenmund, Stern-Bach, & Stevens, 1998).

The second hypothesis is that each subunit participates in the concerted conformational change of all the subunits in the receptor. This gives four different states of opening. The binding of glutamate shifts the equilibrium towards larger open states, but these are not subunit specific. The pore diameter emerges from a concerted conformational change and does not require glutamate binding to occur (although it is kinetically highly unlikely to open fully) (Dutta-Roy, Rosenmund, Edelstein, & Le Novère, 2015).

NTD interactions

The combined ligand binding domain/N-terminal domain elements of AMPAR subunits are so extensive they reach halfway into the synaptic cleft (~130Å) (García-Nafria et al., 2016). These extracellular components may play a role in the anchoring of AMPARs within the synaptic zone¹⁴ (Buonarati, Hammes, Watson, Greger, & Hell, 2019). Deletion of the N-terminal domain affects the clustering of AMPARs in synaptic zones, being more critical for GluA1 than GluA2 (J. F. Watson, Ho, & Greger, 2017). The binding partners responsible for this stabilisation are not yet known. It is possible this is related to previously associated auxiliary proteins such as stargazin (Cais et al., 2014), NPRAP (Silverman et al., 2007), and CNIH-2 and -3 (Shanks et al., 2014), but given the size of the domain, may be presynaptic proteins. This would be fitting with the Brownian motion dynamics attributed to AMPARs/subunits (Tolle & Le Novère, 2010), which would not have any way of detecting that a membrane region is a synapse without both pre- and post-synaptic binding partners. Presynaptic partners may help not only to confirm the existence of a synapse, but also initially to signal that the synapse is now functional.

This research is still new, so the number of proposed partners is low. Early promising candidates are cadherins and pentraxins (García-Nafria et al., 2016). N-cadherin is closely associated with NPRAP and appears to colocalise with AMPARs (Tai, Kim, & Schuman, 2008). The binding domains on N-cadherin when they become 'floppy'¹⁵ move closer to where N-terminal domains would be (García-Nafria et al., 2016; Tai et al., 2008), and a number of N-cadherin binding sites have been identified on the GluA2 N-terminal domain (Saglietti et al., 2007; Y. Takeuchi, Morise, Morita, Takematsu, & Oka, 2015). This would point to conditions that generate 'floppy' N-cadherins lead to protective effects on GluA2-containing AMPARs by securing them to the membrane.

Pentraxins are secreted proteins that have been shown to promote AMPAR clustering through binding to N-terminal domain sites, and if linked to the neuronal pentraxin receptor could provide a transsynaptic linkage network (García-Nafria et al., 2016). The strongest evidence for this interaction comes from

¹⁴ With the C-terminal domains being more associated with synaptic *targeting*, that is, the movement of receptors *to* the synaptic zone.

¹⁵ Caused by the dissociation of calcium, which alters the structure such that it no longer spans the synaptic cleft and instead leans towards the postsynaptic membrane.

GluA4-containing AMPARs (Sia et al., 2007), but there is evidence this is not restricted to this subunit (O'Brien et al., 1999; Xu et al., 2003).

Auxiliary Proteins

Investigations into the regulation of AMPARs via immediate binding partners stemmed from earlier research into the interactions of NMDA receptors with membrane-associated guanylate kinases (E. Kim, Niethammer, Rothschild, Nung Jan, & Sheng, 1995; Kornau, Schenker, Kennedy, & Seeburg, 1995; Niethammer, Kim, & Sheng, 1996). This led to investigations in the late 90s that began to build a picture of how AMPARs are regulated by several binding partners (Lu Chen et al., 2000; Dong et al., 1997; Xia et al., 1999). Three groups – PICK1, GRIP1/2 and the TARP family were both the first to be discovered and the ones studied the most. Recent advancements in proteasomal analysis have expanded the number of associated proteins to a large degree and are in turn reshaping our understanding of what an AMPAR is.¹⁶

GRIPs

The Glutamate Receptor Interacting Proteins (GRIPs) were first characterised as scaffold proteins that may maintain glutamatergic receptors in the membrane and link them to other proteins (Dong et al., 1997). This linking function is strongly supported by the presence of seven PDZ domains present on both GRIP1 and GRIP2. Outwith the synapse, GRIP1 acts as a transporter through its binding to KIF5 (Bissen et al., 2019), and as a binding partner with Fras1 (Long et al., 2008).

The function of GRIPs at the synapse is very diverse, implicated in several AMPAR-associated phenomena. All relate to the scaffolding properties mentioned and have to date only been associated with GluA2 and A3-containing AMPARs. The AMPAR-GRIP binding interfaces are at the end of the C-tail of GluA2 and 3, and PDZ4¹⁷ and PDZ5 on GRIP1 (Dong et al., 1997). This leaves five PDZ domains free to bind to proteins which may influence AMPAR function or location.

GRIP binding has repeatedly been shown to both facilitate receptor insertion into the membrane and limit AMPAR movement both within the membrane and intracellularly (see Bissen et al., (2019) for references). Brain-specific knock out of GRIPs results in decreased membrane AMPAR detection and a failure of cerebellar LTD (Takamiya, Mao, Huganir, & Linden, 2008). There is evidence that

¹⁶ PICK1 and GRIP1/2 will both feature in Chapter 4. Cornichons and TARPs do not, but will be important in any future enhancement of the AMPAR model.

¹⁷ Often unreported in follow-up papers is that 30 amino acids on the N-terminal side of PDZ4 are necessary for binding (Dong et al., 1997).

GRIP binding does not guarantee the retention of AMPARs in a given location (States, Khatri, & Ziff, 2008)¹⁸. However, given the molecule's complex dynamics with PICK1 it is hard to distinguish molecular properties from those that are emergent from this small interaction network. It is also further complicated by the collectively seven known isoforms of GRIP1 (five) and GRIP2 (two), all of which show differences in localisation, and some evidence for divergent functions (Bissen et al., 2019).

In addition to PICK1 and AMPAR subunits, GRIP binds to many other plasticity-related proteins. The most immediate of these is ABP, which is GRIP2 without the seventh PDZ domain, and is often associated with GRIP1 (S. Srivastava et al., 1998)¹⁹. Similarly, GRASP1 is found often as a complex with GRIP1 and AMPARs and appears to amplify the functions of GRIP1 (Chiu et al., 2017; W. H. Lu, Yeh, & Huang, 2017). Other binding partners include CASK, NEEP21, liprin- α , Thorase, Ephrins and ApoER2, all of which have been associated with AMPAR trafficking (see Bissen et al., 2019).

Loss of GRIP1 in mice is embryonically lethal in some (Bladt, Tafuri, Gelkop, Langille, & Pawson, 2002) but not all strains (Takamiya et al., 2004). This lethality is linked to cell-matrix adhesion and is not because of any catastrophic neural phenotype. As mentioned earlier, brain-specific knock outs of GRIP1 have been developed. These show links to murine autistic behaviours (Mejias et al., 2019), suggesting that dysfunction of GRIP may be either causal for or a symptom of autistic-associated mechanisms.

PICK1

PICK1, like the GRIPs, binds to the PDZ domain on both GluA2 and 3 subunits (Xia et al., 1999), and plays an important role in the localisation and trafficking of AMPARs both intracellularly and within the membrane (Citri et al., 2010; Iwakura et al., 2001). It acts as a complement to GRIPs, acting more to reinforce internalisation or instability of AMPARs over retention and membrane

¹⁸ This appears to have been largely ignored in recent findings, with most papers referencing earlier work by the Hirai lab (Matsuda, Launey, Mikawa, & Hirai, 2000; Matsuda, Mikawa, & Hirai, 2002) which show that binding is necessary and that S880 phosphorylation reduces GRIP1 affinity. The States paper however shows that loss of GRIP/ABP binding does not lead to a definite internalisation of phosphorylated receptors, and that it is possible that once complexes have been formed with other proteins, the need for GRIP is diminished. This may be supported by the extremely low levels of GRIP detected in recent copy number analysis (Helm, 2019).

¹⁹ Why or indeed if this splice variant is necessary is yet to be determined.

localisation. However, like GRIPs, the intensive amount of research renders this interpretation a little naïve.

PICK1 has been found to bind to multiple synaptically-associated proteins (see Bissen et al., 2019 for references). This makes defining the role of PICK1 in relation to AMPARs impractical in a short space here. To sketch out a brief summary, one can conclude that PICK1 is likely, due to its F-BAR and PDZ domains (Salzer, Kostan, & Djinoić-Carugo, 2017), to localise to the synaptic membrane. Here, it binds to GluA2 or A3 subunits, and acts as a major docking hub. Its binding partners are all associated with plasticity and morphological changes at the synapse (Bissen et al., 2019). PICK1 therefore may act to aggregate multiple proteins into a defined location, which increases the probability of binding and therefore the feasibility of the biological phenomena associated with these proteins.²⁰

Loss of PICK1 is not lethal, but does result in a loss of Hebbian plasticity (Steinberg et al., 2006; Terashima et al., 2008; Volk, Kim, Takamiya, Yu, & Huganir, 2010), and synaptic scaling (Anggono, Clem, & Huganir, 2011). These losses appear to be restricted to adult neurons, which fits nicely to the proposal that early glutamatergic synapses are predominantly GluA1/4, and shift to GluA2-containing after two weeks of post-natal development (Luchkina et al., 2017). This would position PICK1 as being integral for a type of plasticity, with regions with low GluA2 and 3 concentrations being relatively unaffected²¹.

TARPs

Transmembrane AMPAR Regulatory Proteins (TARPs) were the first transmembrane proteins shown to directly associate with and modulate AMPAR function, with stargazin (TARP $\gamma 2$) being the very first (Lu Chen et al., 2000). To date, $\gamma 2 - 4$, 7 and 8 have all been shown to both bind to AMPAR subunits (Schwenk et al., 2012), and that their knockdown or removal results in a loss of typical functionality (although there does appear to be some compensatory effects) (Bissen et al., 2019). As well as AMPAR subunits, all known TARPs have been shown to also bind with PSD proteins (Dakoji, Tomita, Karimzadegan, Nicoll, & Bredt, 2003; Kato et al., 2007; Pedersen et al., 2017).

²⁰ PICK1 is found throughout the whole body, and PICK1 KO mice have been found to be infertile. Rescue of this infertility by adding back PICK1 results in a rescue of fertility (X. Li, Mao, Wu, & Xia, 2013), lending some weight to this physical localisation hypothesis.

²¹ Some admittedly ambiguous evidence for this is that PICK1 KO studies to date have only shown deficits in cerebellar LTD, not LTP (Steinberg et al., 2006). The cerebellum has comparatively lower populations of GluA2 and 3 than the rest of the brain (Schwenk et al., 2014).

Stargazin is particularly well-studied, implicated in almost all aspects of AMPAR trafficking (Bissen et al., 2019). Like AMPAR subunits, the C-tail possess a large number of residue modification sites, including nine identified serine phosphorylation sites (Tomita, Stein, Stocker, Nicoll, & Brecht, 2005), and at least one threonine phosphorylation site (Chetkovich, Chen, Stocker, Nicoll, & Brecht, 2002; Stein & Chetkovich, 2010).

TARPs appear to have some overlapping functionality, with the loss of certain types resulting in only minor losses of typical system behaviour. $\gamma 2$ and $\gamma 8$ show the greatest effects when knocked out/down, with other forms at typical levels being unable to rescue function sufficiently (Bissen et al., 2019). $\gamma 8$ knockout mice possess only 15 % of normal GluA subunit levels in the hippocampus, yet still show 65 % of normal AMPAR synaptic transmission (Rouach et al., 2005).

Cornichons

The research on cornichons is to date quite muddled. What is known is that they appear to play a role similar to TARPs but are not specialised for any influence of synaptic stabilisation or incorporation. Their critical role is in the transporting of AMPARs (or their subunits) from endoplasmic reticulum (ER) and Golgi structures into the membrane (Brockie et al., 2013; Y. Shi et al., 2010). They act as both a transport and auxiliary protein, with the latter function assumed to be related to the local TARP isoforms present. Cornichons are only co-localised in the membrane with AMPARs when $\gamma 8$ is present (Herring et al., 2013). Cornichon binding appears to regulate the stoichiometry of TARPs with AMPARs, allowing only two to bind (out of a possible four), which alters the resensitisation of the receptor (Gill et al., 2011; Kato et al., 2010).

There is only evidence for interactions with GluA1 subunits at the time of writing (Herring et al., 2013; Lussier et al., 2012), but this has been challenged as being as a result of a methodological error (Boudkkazi, Brechet, Schwenk, & Fakler, 2014), with proteomic analysis showing cornichon expression even in the absence of GluA1 (Schwenk et al., 2012).

In the original report on cornichons, CNIH-2 and 3 were found to be widespread in the brain (Schwenk et al., 2009), and this has been supported in subsequent studies. The Schwenk paper and follow-ups (e.g., Boudkkazi et al., 2014) also report that cornichons appear to slow the EPSP decay time course. This would enhance considerably the probability of short-term plasticity through temporal summation. Knockdown of CNIH-2 is sufficient to remove the slow EPSP decay of MFB-mossy cell synapses, suggesting this is a causal molecule in determining the difference between slow and fast EPSP decay at synapses (Boudkkazi et al., 2014).

As the above section shows, AMPARs are much more than the GluA subunits. There are multiple binding partners forming complexes which appear to be necessary for typical AMPAR function. These partners shape every aspect of known AMPAR function, and in many cases appear to be vital for the typical functioning of AMPARs in biochemical networks.

The transient-AMPA hypothesis makes the distinction between AMPARs and their binding partners even more blurry. AMPARs in this form exist for only transient periods and are likely almost always coupled to one or more auxiliary proteins mediating their localisation, trafficking and stabilisation. If confirmed, it will require a reimagining of AMPARs not as monolithic permanent structures, but emergent entities that appear only for the time that two or more protein complexes associate and interact.

Lifecycle

Generation & Localisation

AMPARs appear to be tightly regulated post-translation before being allowed to exit the ER. Evidence suggests that tetramerization of subunits is a necessary, but insufficient step for ER exit (Greger & Esteban, 2007). Once constructed, there is some form of testing which detects the successful formation of the ligand binding domain, that glutamate can bind to the ligand binding domain, and that the receptor is able to desensitise (Greger & Esteban, 2007).

Once these conditions are met, the AMPAR is thought to then be transported through the Golgi network and localised predominantly in dendrites (Buonarati et al., 2019). The long-standing position is that once an AMPAR has been built, it exists as a complete complex that only dissociates upon proteolytic breakdown. As we have already seen however, recent work from Morise et al. (2019) has provocatively suggested that these tetramers may be quite unstable, and that once situated in the plasma membrane, they readily dissociate into dimer and monomer configurations. These may still tetramerise, but exist for only a few hundred milliseconds, rather than the hours and days that the orthodox view holds.²²

Regardless of the stability of the AMPAR, populations are found to be present at synapses, extrasynaptically, and within endosomes either close to, or within dendritic spines (Cooney, Hurlburt, Selig, Harris, & Fiala, 2002; Granger & Nicoll, 2014).

Trafficking

Trafficking studies make up a large proportion of AMPAR investigations, and we have covered much of this in parts earlier in this chapter. To briefly summarise; there are two ways in which AMPARs may be transported: vesicle-mediated transport and subsequent exocytosis, or lateral diffusion within the cell membrane. Whether there is a distinct mechanistic reason for both has yet to be determined.

Trafficking as we have seen appears to be modulated quite heavily through interactions with many different proteins, which modulate the probability of AMPARs interacting with transporter molecules. To date, none of these proteins are essential however, with most *-/-* forms showing phenotypes that range from mild to severe, but rarely lethal due to their absence for AMPAR interactions

²² We shall look at the impact of this alternative at the relevant points in the remainder of this chapter.

alone, and only lethal post-natally²³. Because of this, it is reasonable to conclude that the trafficking system exhibits a strong degree of robustness and redundancy. As such, the interactions and causal relationships of the many different proteins currently implicated in this process will be hard to untangle.

Synaptic Stabilisation

The maintenance of AMPAR numbers in the synapse is essential for the short-term stability of potentiated and depressed states, and arguably for synaptic function overall. Whilst trafficking and stabilisation are often conflated in descriptions, I feel it best here to separate the two to highlight the different proteins involved in the stabilisation process.

AMPA receptors are highly mobile in a membrane and appear to obey Brownian diffusion dynamics (Borgdorff & Choquet, 2002; Tolle & Le Novère, 2010). Several hypotheses currently exist as to how these receptors are then stabilised within the synaptic zone. One of the more popular hypotheses is the idea of slots present in the postsynaptic density (Huganir & Nicoll, 2013; Malinow, Mainen, & Hayashi, 2000; Opazo, Sainlos, & Choquet, 2012). These slots briefly stabilize receptors but do not irreversibly maintain them there, with receptors eventually diffusing and moving back into the extrasynaptic space where they may return or be endocytosed.

What exactly a “slot” is is still not fully determined and may involve the interaction of the receptors with multiple proteins. There is still speculation on how this stabilization process works²⁴. One idea is it is limited by the number of slots available which is adjusted depending on long-term or potentiation or depression events. Another stipulates that it is the relative affinity of the given receptor for the slots which is the important factor in stabilisation. It is thought that in basal/biochemically stable conditions the number of slots is also stable and that any changes to the number requires a change to the underlying biochemical dynamics (Diering & Huganir, 2018). Current evidence points to roughly four nanodomains within a synaptic zone, each of 20 – 25 receptor capacity (Broadhead et al., 2016; Compans, Choquet, & Hosy, 2016; Hruska et al., 2018; Nair et al., 2013). It may be that the capacity of a nanodomain is what is changed during plasticity,

²³ The only lethal loss of function I could find for a single protein that had a strong correlation to being related to AMPAR function was ADAR2, which edits GluA2 mRNA to remove the calcium permeability of this subunit (Higuchi et al., 2000). Mice were still viable however and only died after several weeks. Another protein, ATAD1, if mutated also was lethal, but not in all animals, and also has a role in mitochondrial function (Piard et al., 2018).

²⁴ Both put forward in (Diering & Huganir, 2018)

and that these “trap” any glutamatergic receptors stably for up to 15 minutes at a time (S. H. Lee et al., 2017).

Slots do not appear to be particularly specific to AMPARs. Experiments where kainate receptors are used in place of AMPARs show comparable long-term potentiation and receptor movement within the synaptic zone (Granger et al., 2013). The slots themselves are probably promiscuous (Huganir & Nicoll, 2013) and the extent to which any receptor type dwells within these slots is likely dependent upon the binding partners of that receptor or its post-translational modifications. This also lends support to the idea that the presence of slots does not change much in plasticity, as during cerebellar LTD events blocking endocytosis is sufficient to block LTD (Y. T. Wang & Linden, 2000), suggesting that receptors diffuse in and out and that the number of slots is not reduced at the onset of any long-term depression event.

A recent modification to this idea has promoted the notion of a hierarchy of AMPARs, where the composition and modification profile determines the chance of probability of being incorporated into these slots (Diering & Huganir, 2018). This relates to insertion-centric hypotheses which stipulate that the movement of receptors from intracellular pools to perisynaptic regions is important and may be the major source of new receptors during plasticity, with the initial extrasynaptic receptors playing only a minor role. Evidence for this comes from Penn et al., (2017) where the immobilization of AMPARs already at the surface resulted in a blockade in short term potentiation but still produced (an albeit smaller) long-term potentiation. Allowing for extrasynaptic movement but blocking exocytosis results in no such long-term potentiation but still allows for short term potentiation, suggesting that exocytosis is a critical component in the expression of long-term potentiation.

The transient-AMPA hypothesis (Morise et al., 2019) requires a much greater role for this stabilisation process. Either subunits would be captured by scaffolders such as GRIP1 or PICK1 already present in the synaptic zone or recruited by capture prior to entry into the synaptic zone. Evidence would point towards complex formation prior to synaptic zone insertion but may be complicated by whether dimerisation and tetramerisation are possible whilst within these complexes. The capture in the synaptic zone may be partially true in that the AMPAR complexes might be partially formed and tethered to membrane-associated guanylate kinases. These would capture incoming complexes which contain trafficking-associated proteins such as GRIP. This improves the probability of capture, which can be fine-tuned by the relative copy numbers of complex parts.

If slots are seen to change dynamically under plasticity conditions, rather than being a simple case of “opening” and “closing” slots, these dynamical changes can be considered as the composition of both incoming and tethered complex proteins. Perhaps even more intriguing is the fact that as these complexes are forming independently in different spatial locations, there is the possibility of a mismatch between tethered complex composition, and the incoming complexes. The response of a synapse to a plasticity event therefore is not guaranteed if the AMPAR complex composition does not match well with the current synaptic capacity.

Recycling and Breakdown

This section overlaps with trafficking of receptors, as much of the process is the same. However, it is important to understand how such a process differs in how AMPARs may be maintained and stored as contingencies. At the edge of many synaptic regions are areas where the diffusion of AMPARs results in their immediate removal from the plasma membrane. These zones are maintained through the interaction of the proteins Homer and Shank with the postsynaptic density (PSD) (J. Lu et al., 2007). They allow for the clustering of a pool of AMPAR-filled endosomes close to the membrane surface. Upon the presence of an LTP-inducing stimulus, this repository of vesicles may be rapidly brought to the surface and used to continue the inductive response, as well as the subsequent LTP-maintenance period.

The slot-centric hypothesis of LTP sustaining ultimately sees the total number of available receptors as being more important than the compositional makeup of the receptors themselves. Evidence for this is twofold. The first is that the replacement of all AMPARs with kainate receptors does not result in the loss of normal spike generation (D. Yan, Yamasaki, Straub, Watanabe, & Tomita, 2013). Second, the removal of the recycling zones results in a marked aberration in LTP/LTD processes (J. Lu et al., 2007).

AMPARs are ultimately degraded in converted endosomes, which become lysosomes and cause breakdown of the receptor complex (Widagdo et al., 2015). The ability to detect which receptors are to be degraded is dependent upon the ubiquitination of the component subunits. Ubiquitin is a small protein which can bind with itself to form polymers on any of seven lysine sites present in the protein (W. Li & Ye, 2008).

After being used at the synapse (i.e., after glutamate binding), each subunit is rapidly tagged with ubiquitin- this does not occur for any surrounding receptors not used or ones remaining outside of the plasma membrane (Widagdo et al.,

2015). Mutating the sites of ubiquitination affects only the mechanisms leading to degradation, leaving trafficking and function unchanged (Widagdo et al., 2015).

Ubiquitination is not a confirmation of immediate lysosomal breakdown. Some studies have indicated that this tagging is necessary for receptor internalisation in general (e.g. Widagdo et al., 2015). This is juxtaposed with work that suggests ubiquitination occurs only after internalisation has occurred (Lussier, Nasu-Nishimura, & Roche, 2011). This leaves the tagging by ubiquitin as a filter for further processing leading to degradation, or to recycling via transfer to endosomes. Given that the latest study indicates methodological issues with the preceding work, the veracity of any strong claims in this area is open to debate.

The trafficking of AMPARs appears to be further controlled by the presence of ubiquitination tags on several residues. It appears that ubiquitination may be playing a major role in the trafficking intracellularly of AMPARs with this regulation of this pathway being implicated in disorders such as Alzheimer's disease (Zhang, Guo, Huo, Wang, & Man, 2018), chronic stress (Wei et al., 2016), and epilepsy (J. Zhu et al., 2017).

Ubiquitination of a protein depends on several different enzymes. These include ubiquitin activating and ubiquitin conjugating enzymes, as well as ubiquitin ligases. In the case of AMPARs the ligases Nedd-1 & 2 (A. Lin et al., 2011; J. Zhu et al., 2017), APC^{Cdh1} (Fu et al., 2011) and RNF167 (Lussier et al., 2012) have to date been shown to mediate ubiquitination in neurons. This ubiquitination pathway is calcium-dependent and involves the activation of CaMKII, which may be phosphorylated in these cases (the actual mechanism is to date not fully understood)(Widagdo, Guntupalli, Jang, & Anggono, 2017). There is also evidence for other activation methods suggesting that further work needs to go into investigating how many possible routes there are in activating these ligases for use on AMPARs (Widagdo et al., 2017).

Nedd4 has been shown to reduce surface AMPARs in overexpression studies leading to accumulation of internalized AMPARs (Schwarz, Hall, & Patrick, 2010). The two forms of Nedd4 both act on the GluA1 subunit at the K868 site (Widagdo et al., 2017). Despite their similarity these two forms appear to have different functions within the cell (B. Yang & Kumar, 2010). Our current understanding appears to suggest a convergence on the GluA1 subunit, but there may be a difference in the recruitment profile of these forms depending on neuronal type, stimulation, dynamic profile etc. RNF167 has a much more passive role in the contribution of ubiquitination of AMPARs. It acts as a scaffold to allow for E2 enzymes to directly transfer ubiquitin to (in our case) the GluA2 subunit

(Widagdo et al., 2017). Loss of this ligase results in overexpression of GluA1-3 subunits (Lussier et al., 2012).

Normally ubiquitination tags are used as a marker of protein degradation and eventually reach a threshold which results in transfer of proteins to proteasomes or lysosomes. This threshold is taken to be a chain of four or more, although this does not appear to be an actual biological mechanism (Clague & Urbé, 2010). This notion has led to speculation that fewer tags, or different chain structures, may be encoding important information related to AMPAR localisation via endocytosis and intracellular sorting. Indeed, mutant K868 sites that only allow monoubiquitination result in a failure of GluA1-containing AMPAR endocytosis (Patrick, Bingol, Weld, & Schuman, 2003). Evidence against this comes from a few studies which suggest the endocytosis precedes the transfer of ubiquitin onto AMPARs, with mutants deficient in the ubiquitin sites unaffected in terms of surface expression or internalization of receptors (Lussier et al., 2011; Widagdo et al., 2015).

Complicating matters further is evidence that phosphorylation of nearby residues such as S845 can prevent or slow the rate of ubiquitination (Guntupalli et al., 2017). This crosstalk, coupled with the fact that each subunit type can be ubiquitinated and may each be sufficient to maintain normal functionality, makes it difficult to assess the veracity of the multiple roles these tags are purported to play beyond mere degradation signalling.

What the studies do suggest however is that ubiquitin tags appear to play a considerable role in the determining of how many (and what type of) AMPARs are present within the membrane of a dendritic spine. This can be further shown through analysis of pathological disorders in which ubiquitination mechanisms are perturbed. Amyloid beta appears to trigger an increase in the level of ubiquitination particularly at the K868 site, with mutant receptors showing no increase in the removal of AMPARs from the membrane (Guntupalli et al., 2017; Rodrigues, Scudder, Goo, & Patrick, 2016). In addition, there is evidence that Nedd4 is upregulated (or redistributed) in Alzheimer's disease (Kwak et al., 2012; Rodrigues et al., 2016), providing a useful target pharmacologically for preventing the loss of AMPARs from spines and subsequently maintaining synaptic function for longer.

The rate of ubiquitination and breakdown appear to be important in maintaining the balance of AMPARs within the synaptic membrane, and subsequently LTP/LTD maintenance. The extent to which this system may play a causal role is still undetermined (Louros & Osterweil, 2016).

Theoretical Summary

As we have seen, AMPARs are the primary driver of change to biochemical dynamics in glutamatergic cells. They possess a prodigious number of different modifications, subtypes and states, which is reflected in the wide heterogeneous population of cells they are found in. This adaptation has led to a broad range of similar but distinct mechanisms of phenomena such as synaptic plasticity and is perhaps responsible for the complexity of neural networks as a result.

These mechanisms appear to centre on the regulation of the conductance, open probability and number of AMPARs in functionally viable areas of the cell membrane. These alterations drive calcium dynamics locally. These changes are amplified and modulated through a small number of highly sensitive molecules which act as filters to determine what biochemical changes result from AMPAR activity.

AMPARs are therefore critically important for change, but also subservient to it. AMPARs are not long-lived molecules, and much of the evidence now suggests them to be part of much larger molecular complexes. They may be better thought of as transient collection of properties that define a system state over time as well as being a historical record of previous system states.

The role of AMPARs is complexly embedded into information processing and transfer. Depending on your point of view, their role is aided by or encoded upon by many different molecules such that the pattern or code records the local network state. The incredibly large number of possible configurations may act as dispositional properties to drive the biochemical network to previous (effectively unique) states (De Brigard, 2014; Fitzpatrick, 2016), and therefore be a possible instantiation of synaptic memory.

This is not something within the scope of this Chapter (or indeed the whole thesis). It does suggest an interesting avenue of investigation for proponents of both synaptic memory and synaptic tagging hypotheses, however (Redondo & Morris, 2011; T. Takeuchi, Duzskiewicz, & Morris, 2014). Could the set of AMPAR subunits and their associated proteins be the biomarker of a memory? Is it correct to associate memory with LTP and LTD when the composition of these macromolecular complexes respond very differently depending on the experimenter's induction protocol?

Functional/Mechanistic Gaps

The picture presented in this chapter is of a molecule with a vast array of binding partners, modifications and localisations, which makes making sense of it much more difficult, and very difficult in a short piece of text. Our investigations have revealed that AMPARs are a major player in the dynamics of dendritic spine biochemistry, synaptic plasticity, and information transfer. With this complexity and depth of study come a range of questions yet to be answered.

Our understanding of N-terminal residue modifications and possible non-glutamate binding sites is still in its infancy, but points towards this aspect of AMPARs having important roles in understanding trafficking and signal transduction. We have seen that AMPAR trafficking is still a tangle of incomplete hypotheses, which whilst not mutually exclusive are difficult to reconcile with current evidence and understanding.²⁵ Ubiquitination may be a critical element of this trafficking away from the membrane, but this too is a debated area. So too is the pattern of phosphorylation, although much of this appears to be subunit and brain region-specific.

Although rarely mentioned in reviews and commentary on AMPAR studies (see Moretto & Passafaro, 2018 for a rare instance), many findings attribute properties to single subunit types whilst neglecting to isolate these subunits. This leads to conflation of dimer and tetramer properties to single subunits, as well as the misattribution of functions of one subunit (or the concerted action of multiple subunits) to another.

Associated binding proteins are possible extensions of this same principle of modifying AMPAR trafficking and function. These, like AMPAR subunits, are very differentially expressed across different brain regions, and come in a much greater number of variants. We are beginning to get closer to understanding *how* these proteins influence AMPAR function, but the question of *why* they are necessary has not yet received serious consideration.

We may even have to rethink our whole idea of how AMPARs are formed and maintained, or even what “an AMPAR” is. The evidence we have about AMPAR function is only as good as our methodologies used to investigate them. The next steps in our understanding require a more holistic picture of AMPARs that embed them more explicitly within the mechanisms they are part of. One way we can begin to approach this is by bringing computational methods of simulating

²⁵ We will tackle approaches to resolving this in Chapter 4.

AMPARs into our arsenal of theory-building tools. The next few chapters will deal with this idea head-on.

CHAPTER 2

An Introduction to Computational Biochemical Modelling

Biological modelling (as defined here) is the simulating of biological processes *in silico*. The underlying rationale is that actions of biological processes are inherently mechanistic – that a behaviour may be explained through the description of a mechanism which itself is made up of many different parts (Craver, 2013). Biochemical modelling focusses on the dynamic and complex reaction networks which are part of these mechanisms and therefore underpin biological phenomena.

The extent to which computational modelling is necessary for theory building and conceptual understanding has been, and still is, fiercely debated. In some fields modelling is a necessary part, in others, a misleading distraction. This has led in some cases to a dichotomy between modellers and physical experimenters²⁶ which has led to a ‘mutual suspicion’ driven by a lack of understanding by one side of the other (Mogilner, Wollman, & Marshall, 2006).

This chapter’s aim is to offer a primer of how biochemical modelling works, what it shows, the challenges and problems with it, and how it integrates with *in vivo* and *in vitro* work²⁷. It also acts as an introduction to what I perceive to be the underlying philosophy of modelling, and to some of the flavours of modelling currently in use.

Why Model?

There are many surface advantages to modelling a biological system. You can run experiments with identical conditions multiple times, alter parts on the fly, and avoid the need for animal experimentation or tissue collection (Berro, 2018; Sneddon, Faeder, & Emonet, 2011). Multiple hypotheses or questions can be addressed rapidly and accurately within the same model.

These models can be built from a reductionist (from genes/molecules up) or holistic (from phenomena/mechanism down) angle. They can look at anything

²⁶ I may explain this in more detail later, but I strongly oppose the exclusive use of “experimenter” or “experimentalist” to those who deal with living tissue. It’s a modern-day “featherless biped” definition that needlessly exacerbates the perceived gulf between the two methodologies. Unfortunately for me it’s commonly used, and so for the sake of readability I will maintain this usage. Begrudgingly.

²⁷ Therefore, unless explicitly stated, the use of “model”, “modelling” and “modeller” all refer to biochemical variants only.

from a few molecules (e.g. Li, Stefan, & Le Novère, 2012) through to a whole cell (e.g. Karr et al., 2012). A computational model can make testable predictions which can be used to narrow down the number of ‘experimental’ investigations required to reach an understanding of a given phenomenon. Alternatively, it can be used to accommodate existing work and assess its coherence and comprehensiveness.

The aim of constructing a model is to offer possible explanations for phenomena from reaction dynamics, to work out *how* things work together over time to give rise to a given phenomenon. By understanding the features which emerge from network complexity we can begin to understand why and how the phenomenon arises. Similarly, we can see if the same principles can be applied to other reaction systems and discover templates or frameworks common to complex networks of interacting molecules.

In its purest sense then, models drive an understanding of phenomena by both defining and explaining the composition of mechanisms. I would define mechanism as follows (influenced by (Craver, 2002, 2013; Nicholson, 2019; Varela, Maturana, & Uribe, 1974)):

***Mechanisms** – a bounded organisation of a collection of agents and activities which together give rise to a behaviour that defines the mechanism. These are explained in space, time, and the sequence of activity.*

Mechanisms are difficult to experimentally assess. Yet without a way of testing and investigating mechanisms, the functional definitions of its components are in danger of becoming too free – you can fit any observable function into a mechanism which fits a subjective view.

This is not a new problem in biology – we have long struggled with the mapping of intricate biochemical work to higher level phenomena:

The modern student finds it difficult to see the wood for the trees ... He does not always have a synoptic concept of the nervous system in his mind ... If we subject a clock to minute analysis by the methods of physics and chemistry, we shall learn a great deal about its constituents, but we shall not discover its operational principles, that is, what makes these constituents function as a clock. Physics and chemistry are not competent to answer questions of this order, which are an engineer’s task ... Both modes have their place and limitations; and they complement one another. (Walshe, 1961, p. 131) (from Chirimuuta, 2018)

What is being highlighted here is the difference between reduction and reverse engineering – the former is difficult to use if there is no understanding of what it is supposedly doing. This issue at first glance is not much of a problem – we have working theories of neural function, phenomena such as memory etc, which we then look for the component parts of. And there is a recognition that the two are necessary. What we lack sometimes is the correct translation between the two (Wimsatt, 2006).

Modellers are as guilty of this as the rest of biology, and especially in neuroscience. However, models are I believe in the best position to (and should strive to) represent the context and scale that these ‘mechanisms’ operate in.

Extended Example: Fragile X Syndrome

There are some key examples of this emerging within certain fields which have been to date driven largely by physical experiments alone. I believe looking at one of these in detail may help frame the importance of modelling more clearly.

Fragile X Syndrome (FXS) is a mono-genetic disorder arising from an abnormal amount of CGG repeats in the gene coding for Fragile X Mental Retardation Protein (FMRP). A genetic definition of FXS could be described as an organism with greater than 200 of these repeats present in the untranslated region of the gene. There is a critical threshold at which symptoms occurs: prior to this an individual acts as a carrier of FXS (between 55 - 200 repeats) (Dahlhaus, 2018).

In those with FXS, cognitive deficits arise largely during and after adolescence. The IQ range 20 - 70 is well below the general average of 100 for the population (Cornish, Munir, & Cross, 2001; Van der Molen et al., 2010). Male FXS IQ appears to be correlated with the level of methylation of the gene- the lower the methylation, the higher the IQ (Merenstein et al., 1996). In females, the chance of intellectual disability on a clinical scale is much lower, although there are still deficiencies in learning. The major manifestations arise in attentional and executive function, which has a major knock-down effect on educational learning (Van der Molen et al., 2010).²⁸

Outwith these symptoms, there are several other neurological differences which are somewhat specific to FXS. Epilepsy is comorbid in roughly 20% of young FXS patients and is often generalised (tonic-clonic) (Musumeci et al., 1999). The uniqueness of this phenotype has been utilised in animal models as a reliable indicator of FXS. Audiogenic seizure experiments make use of a common audio

²⁸ The intellectual disabilities of FXS are not that of autism. The two are comorbid, with around 30% of FXS individuals exhibiting autistic traits (Hagerman et al., 2010).

hypersensitivity exhibited in FXS to induce seizures (L. Chen & Toth, 2001; Musumeci et al., 2000; Rotschafer & Razak, 2014; Q. J. Yan, Asafo-Adjei, Arnold, Brown, & Bauchwitz, 2004). Indeed, there is evidence of a global hypersensitivity to sensory input (Merenstein et al., 1996).

The effect of the CGG repeats on the FMR1 gene is that there is a silencing at the translational level, leading to a global decrease or loss of FMRP (Hagerman, Hoem, & Hagerman, 2010). This is particularly detrimental to the structure and functional dynamics of neurons. FMRP is especially densely localised to postsynaptic sites, where it is thought to repress translation of proteins through complementary binding to mRNAs (O'Donnell & Warren, 2002). What makes this protein so intriguing to investigate is that despite this seemingly important role in synaptic function, its loss is confined to detrimental, but not lethal, changes to brain activity.

Indeed, the loss of FMRP only appears to result in FXS and has to date not been associated with any other major disorder. Within the brain, the loss of this protein manifests itself most visibly in the abnormal morphology of dendritic spines (e.g. Irwin et al., 2001). In addition to these modifications through globally increasing protein levels, the absence of FMRP also affects lipid pathways (Tabet, Moutin, et al., 2016; H. Wang, 2014). Somatostatin is present both centrally (Urban-Ciecko & Barth, 2016) and peripherally (Yasaka, Tiong, Hughes, Riddell, & Todd, 2010), and has been linked to both cognitive and the physical abnormalities seen in FXS.

The long-held, and well supported, view is that the dysregulation of FMRP directly leads to an overexpression of certain proteins which are usually controlled (Bear, Dölen, Osterweil, & Nagarajan, 2008; Bear, Huber, & Warren, 2004). The absence of FMRP leads to unchecked translation, and the overabundance of proteins triggers the syndrome's characteristics. Much of the past decade of work in FXS research has been devoted to understanding if the dysregulation of mGluR signalling is the sole underlying issue for the FXS phenotype (Bear et al., 2004; Telias, 2019).

The loss of FMRP leads to a marked increase in concentration levels within dendrites. Efforts to reset mGluR levels to normal physiological values has produced many interesting results. A reduction of 50 % in FMR1 KO mice is enough to rescue the phenotype, achievable both pharmacologically and genetically (Dölen et al., 2007). The double mutant rescue of both the FXS and TSC2^{-/-} phenotype is perhaps the strongest evidence for the direct interaction between mGluR signalling and protein expression levels (Auerbach, Osterweil, & Bear, 2011).

Therapeutically, the effects of statins on mGluR have proved the most effective way of treating the disorder. Statins are inhibitors of HMG-CoA reductase, an enzyme critical for the generation of cholesterol (Osterweil et al., 2013). It has been hypothesised that this activity disrupts the development of lipid rafts critical for the expression of receptors at the synaptic surface (van der Most, Dolga, Nijholt, Luiten, & Eisel, 2009). Additionally, further work found that statins may be playing a role more directly in the suppression of abnormal protein synthesis by acting on the Ras/ERK pathway (Osterweil et al., 2013).

Work outwith the mGluR theory is quite widely distributed across several other receptor types and specific signalling pathway targets. In the interests of brevity and to avoid unnecessary repetition, much of this other work is collected in an excellent review by Wang et al. (2015).

Understanding the Mechanism

These findings provide a detailed insight into some of the dynamics and functions of very different molecules seemingly related to FXS. They all show an ability to rescue the phenotype through interventions which suggest the studied molecule to be of causal importance. What they do not show is how we reconcile these ideas together. The mono-genetic facet of FXS was attractive to research because it suggested an easy fix – find the molecule most affected, rescue it, and translate this fix into human clinical work.

So why are so many molecules important? There are a few possible explanations to propose here, of varying degrees of plausibility. One is that the rescue effects we see in these studies are short term, and that prolonged studies would find a gradual return to the FXS phenotype due to shifts in signalling dynamics elsewhere. Another is that mGluR's relation to all these pathways is not constitutive. Admittedly, taking the position that the causal role of mGluRs has been overstated is controversial, but a reappraisal of the evidence may be beneficial to unpick exactly how it interplays with the complexity of other pathways.

As you may have now surmised, it may be that no single molecule is important at all, given the complexity of the synaptic environment. Placing any part as more influential over the others when one or more are absent is to in turn simplify the biochemical pathways to assume any single aspect somehow carries more information about the global state than the rest.

Indeed there is a slowly building body of evidence indicating different neuronal and synaptic types respond differently to the lack of FMRP (Dahlhaus, 2018; Sawicka et al., 2019). As such, our attempts to presume one molecule's rescue is

enough is to mask the variability in compensatory mechanisms throughout the organism. This does not exclude some weighted importance of certain signalling molecules and/or receptors, but these would be in a situation-specific manner.

So where does this leave the mGluR theory? The dramatic success of mGluR studies in mice and initial clinical studies has led to great progress in understanding the molecular pathways that underpin glutamate-modulated changes to the synaptic environment. However, mGluR5 treatments have not shown any benefits to phenotypes outwith the nervous system (such as macroorchidism). Additionally, the clinical trials have all failed at stage II or III thus far (Nicoletti, Bruno, Ngomba, Gradini, & Battaglia, 2015), despite the confirmation time and again that in rodents this rescue is complete and readily replicable (Jeste & Geschwind, 2016; Mullard, 2015, but see Thomas et al., 2011). The narrative of mGluR is persuasive, but as part of a greater whole.

We can appeal to the complexities of drug tolerance and dosage potentially influencing the results. Yet we should begin to question whether this narrow functional focus is still optimal. The difficulties in transferring work from rodents to humans notwithstanding, there is also the sometimes-overlooked issue that in rodents the cognitive impairments exhibited are very mild. Much of the behavioural work in the literature makes use of carefully crafted testing protocols to tease out small changes in cognitive flexibility within FXS models (Dahlhaus, 2018). Yet the defects found in the human condition would result in significant differences in much more robust and well-established protocols.

The functional investigations in animals are being defined by mechanisms which are in turn defined by behaviour observed in humans. Assuming that mechanisms in model organisms map onto human variants is common practice, as we are usually unable to repeat the same kind of experiments that would show these mechanisms in human tissue. But neuroscience's definition of function is embedded within the mechanism that the function is discovered in. If the mechanism is not the same, then the functional definition is potentially misleading or incorrect.

The behavioural output of FXS synaptic dysfunction is very different in rodents compared to humans. If this behaviour is because of the exact same mechanism as in rodents, then the behaviour should be comparable. Even more crucial is that the function of FMRP is the same in humans and rodent models. Yet a look at the literature shows there are considerable differences, all the way down to the genetic level (excellently reviewed in Dahlhaus, 2018).

The work is a sketch of possible ways in which FMRP loss causes FXS – a range of hypotheses as to the mechanism(s) FMRP is situated in. It is not possible to say more than this because the range of mechanisms has not been constrained to those realisable in humans. What is needed is a way of exploring these myriad hypotheses by situating them within proposed mechanisms, and seeing how they work across time, and in response to the varied perturbations most biological systems undergo.

Here is where modelling comes in. Without modelling, the field cannot be sure of the mechanisms that exist in the rodent, and subsequently even less sure of those in humans. The function of FMRP in rodents needs to be situated within a rodent-specific model. The mechanism that emerges from this will not be the same as the one in humans but may be a template that grows with the additional complexity of structure and biochemical composition into the mechanism seen in humans.

Without modelling to help constrain the possible mechanisms at work, research will be confused by this mapping of functions from humans to animals, then to mechanisms that may not exist in humans. The success of mGluR interventions in rodents does not map to humans in part because of too much faith that the considerable difference between rodents and humans does not have much of an impact on the mechanism.²⁹

This faith that a mono-genetic disorder should be readily explainable through ‘experimental’ work alone is perhaps why to date no biochemical model exists for FMRP, and only a handful of broader computational models (e.g. O’Donnell, Tiago Gonçalves, Portera-Cailliau, & Sejnowski, 2017). This has provided a wealth of *in vivo* functional data but is no longer tractable to comprehend through piecemeal experimentation alone.

Part of the tractability problem is that focussed and standardised work acts to isolate molecular function from context to avoid ascribing incorrectly functions to one molecule that are derived from another. Ultimately however, these molecules work as part of complexes and networks in which emergent properties arise from the dynamics that play out over time (Nicholson, 2019). Isolated molecular analysis is also at risk of attributing causal relevance of a function (or loss of it) to explain ‘downstream’ outcomes (Wimsatt, 2006).

The attribution of causality however is muddled because biochemical networks do not have “start” or “end” components (Villalobos & Ward, 2015). Nor do they

²⁹ It also is because the definitions of “success” in rodent vs clinical trials show almost no overlap (Erickson et al., 2017)

have any internal sense of directionality, of pathways or intentions. We write about them in this way to make sense of them but fall too easily into unconsciously attributing these properties onto the network itself. These molecular networks are the outcomes of random collisions between molecules, which is constrained by location, concentration, temperature, pH and the electrical charges of component amino acids or surrounding ions. The organisation of these molecules is not ordered, not sequential, and certainly not passive (Nicholson, 2019).

Our common definition of a mechanism is one of parts that compose a whole – the gears and springs that make up a clock for example. Each component is unlikely to change when separated from the mechanism, and only a few pieces are redundant. Furthermore, most parts are maintained in a stable environment, with a wide safety range and carefully constructed to work in a range. Cellular environments are bounded and less chaotic than a simple soup of molecules, but are still very much not a clock-like machine. A growing number of molecules are being identified as making use of the “Brownian storm” – the random fluctuations in shape and location determined by the environment’s thermal profile and available space (Ait-Haddou & Herzog, 2003; Chowdhury, 2013; Nicholson, 2019).

Whilst we are largely aware of these chemical properties in principle and upon reflection, they are rarely incorporated into our calculations as to how phenomena are realised. We collapse a macro-level mechanistic explanation (and the subsequent casual mechanism relations between molecules) onto this micro-scale and try to make it fit.³⁰ This has to date worked quite well, but faced with the sheer volume of data now available, we need to make a concerted effort to reemphasise that different scales may require different ways of explaining and defining³¹.

Modelling is a way of exploring these now rather unwieldy functional definitions by situating them within the whole mechanism (in this case, synaptic function). This can be done initially rather simply through the construction of a model that describes the known binding partners, reactions, localisations etc. of molecules best characterised as relating to FMRP-mediated synaptic function. These models do not require any real parameters. They act as visual reviews of the experimental evidence and help to see how things currently fit together.

³⁰ This may help highlight that almost no neuroscientist is a reductionist wholesale. A purely reductionist position would work solely from micro up and work with the biochemical properties alone and observe what emerges. There would be no fitting to a higher level mechanism, or appeal to higher levels.

³¹ An obvious analogy here is that between quantum and classical mechanics.

These types of diagrams act like jigsaws, helping with defining both the current boundaries of a mechanism (as understood through a set of functional descriptions) and the gaps (the missing pieces). The latter is particularly useful for identifying misconceptions, broad assumptions, and just simple lack of experimental data.³² The former helps with how modellers sketch out initial mechanism descriptions, and ensures the model produced is relatable to the physical experimental field.

In some cases, these models are enough to narrow down research focusses and glean more nuanced aspects of the mechanism the model highlights as missing. This is an example of models as ‘experimental’ aids, and perhaps to some readers the main purpose of modelling more generally. The next section (and the following chapters) however argue how models are much more than aids *to* ‘experimental’ work and can be themselves ‘experimental’. I will show how models can both complement and challenge hypotheses and theory, and why this is both beneficial and crucial for future research in neuroscience.

³² See Chapter 5 for some examples.

Biochemical Modelling

Biochemical modelling explores biological function through the simulation of the chemical reactions thought to occur within a defined 'system'. There are thus three main parameters required for any biochemical model: the molecules (or 'species') interacting, the concentrations of these molecules, and how quickly they interact (the reaction rate).

A model derives information from the simulation of these interactions over time. A modeller can assess the rate of change in species, visualise stable system-state configurations, and look for indirect causal relationships, amongst others. Furthermore, a modeller can perturbate the system by altering any of these parameters and see how they alter the outputs of the model.

There are two overarching methods of simulation: deterministic and stochastic. Deterministic models are completely defined by the parameters inputted by the modeller. Running a model with the same set of parameters and initial conditions n times will give you the same results each time. Stochastic modelling differs in that uncertainty is injected somewhere into the model system. Running the same set of parameters n times can produce different results.

Biochemical software packages such as COPASI (Hoops et al., 2006) now provide a compromise hybrid simulator. Fast reactions, which are less likely to be affected by spatial/environmental fluctuations, are modelled deterministically. Slower reactions are modelled as stochastically.

Language Standardisation

Many of the published models of recent years are written in the systems biology markup language (SBML), which is designed to standardise the format of models so that they are readily exchangeable, do not require editing, and survive longer than the lifespan of the software used to create them (Hucka et al., 2003; Zi, Zheng, Rundell, & Klipp, 2008). SBML focusses on representing the biological aspects of a model accurately and clearly – it does not define the model mathematically. This agnosticism allows models to be easily transferred between software packages, which only need to be able to read the entries in the SBML file and match to their own internal file structures.

An SBML file defines model agents as ‘species’, and interactions as ‘reactions’. There are also definitions for compartment, (global) parameters, units, and rules³³. In so doing, the SBML format preserves the minimally necessary model components to allow for model reuse. More recent upgrades to the format (known as ‘levels’) also provide the space to annotate model components (Keating et al., 2020). This enrichment of model metadata provides an opportunity to link model species to databases such as UniProt (2019), as well as provide sources for concentrations and reaction rates where applicable. A model can contain within it all the necessary information to verify, translate and reuse in both part and whole, without the need for a complementary database.

Other xml-based representation formats, such as the Simulation Experiment Description Markup Language (SED-ML) (Waltemath et al., 2011), builds on this formalisation by also requiring explicitly defined information about simulation software and algorithms used, all data processing steps, and a definition of the original simulation outputs in the format first published. This more rigorous representation is designed to minimise the loss of crucial but usually unpublished or unrecorded information necessary to allow for model replication and reproduction. In so doing, models are preserved and may be used ‘off the shelf’ much more readily, with no need to redetermine initial parameter settings, and verification steps becoming much more rapidly processed.

Multiscale Modelling

Simply put, a multiscale model is one where sections operate at scales of time or space that differ by some order of magnitude. The model may span these in terms of its variables, or in terms of resolution, where the operations at one level are inserted as inputs into another. In biomedical settings, this usually means a model

³³ These rules are different from the rule-based modelling approach in the following section and Chapter 4.

which couples fast, micro to millisecond movements of ions with biochemical signalling pathways that operate from seconds through to minutes, hours, and days.


Such models are desirable because they combine the richness and complexity of the smaller-scale components(s) with the less complex but more readily observable outcomes of the more macroscopic component(s). Furthermore, the reductionist take on biology assumes that the outputs of a lower ‘level’ of analysis can explain the observed behaviour of higher levels, thus affording a check on how lower levels effect the behaviour of higher ones.³⁴

Biochemical models, when sufficiently detailed, can be argued to operate in a multiscale environment. Synaptic reactions may operate on millisecond through to second timescales, but the observed outcomes of a model may run across minute and hours. These slow-onset observed features (for example, receptor trafficking within the membrane), then influence the inputs which drive those molecular reactions which gave rise to the phenomena in the first place.

Multistate Modelling



Biologically important proteins may exist in a multitude of different configurations depending on binding partner availability, environmental conditions (e.g., temperature, pH, ion concentrations), space and post-translational modification. These different configurations are often functionally relevant, allowing for networks of a relatively small number of proteins to have a range of responses to changes in context (or “input”). Furthermore, these configurations can exist simultaneously, and influence the likelihood of other properties. This leads to a dynamic interplay of signals and functionality within a single protein, ranging from the straightforward and simple to the incredibly complex.

Modelling molecular species has often had to abstract away much of this intra-complexity to make investigations into their functionality in networks tractable. A problem quickly arises for any complex definition of a species wherein each possible state must be accounted for in any reactions defined within the model. Take for example a simple reaction like the one below:



³⁴ That said, most models of this type are dynamically coupled, and agree that higher level behaviour can in some manner influence the lower one. Unpicking what this means mechanistically is much beyond the scope here, I direct readers here for a primer (Craver, 2007) and here for a rebuttal (Romero, 2015).

If we assume A, B, and C have only one configuration each, then we just write the one reaction. If A has a phosphorylation site for example however, which is phosphorylated in a separate reaction, then we must write two reactions:



We would have to write these out regardless of if the phosphorylation state had any impact on the reaction rate. A is now two different molecular species in the model, with their own separate concentrations. If B also had a phosphorylation site, we would have to write out a reaction for every possible configuration, and have two species for B. For every new reaction that involves A or B, the more reactions need to be written. For every new state or property we attribute to a molecule, if these are not mutually exclusive, adds to the number of explicitly defined reactions and species required.

This problem is exponential, and leads to a phenomena dubbed ‘combinatorial explosion’ (Danos, Feret, Fontana, Harmer, & Krivine, 2007). A small number of properties, when independent of one another, can result in a set of possible combinations that far outstrips any ability to define them in the nature required for a standard biochemical modelling approach. The problems modellers face has been previously defined by Stefan, Bartol, Sejnowski, & Kennedy, (2014) as being two types of problem: one of specification and of computation.

The specification problem is that of defining all states *a priori*. Of these states, some may never exist, but a modeller cannot assume this without strong experimental evidence. They can also choose, if the numbers are not too high, to define each state and construct a model. This however runs the risk of human input error, which can be hard to detect and correct. One can also only choose those properties relevant to a specific research question. This then removes the model’s ability to accurately model both the molecule and the system however, as it necessarily assumes that the experimental data used as a basis is correct, that all other properties have no influence on the system, which acts to bias the model from the onset towards a particular output which may not represent reality.

The computation problem is that of simulation – if a single molecule can have many millions of possible configurations, how can we track those over time in a way which is computationally tractable? If a simulation takes significantly longer to run than the real-time equivalent, it may not be a good use of that time. This is especially true for molecular networks which do not have a depth of experimental parameters to source from – it may take longer to estimate and test parameters than it would to liaise with and organise a lab to conduct the necessary interaction

experiments. In an era of fast science also, researchers wanting to use models to test their experimental findings are less likely to be able to afford to wait for a model whose simulations take days rather than minutes.

Another issue is simply computational power – a protein with nine phosphorylation sites (i.e. 512 possible states) would require over a million reactions and species, and proved unable to be simulated using high end computers a decade ago (Sneddon et al., 2011). Even with the rapid advancements in processing power and memory, such as new GPU approaches (e.g., Tangherloni, Nobile, Besozzi, Mauri, & Cazzaniga, 2017) complex multistate modelling is still a challenge to achieve with traditional approaches on appropriate timescales for use.

Tackling the Dilemma

Over the past two decades there has emerged a growing number of modelling methodologies which tackle one or both horns of this dilemma. All operate differently, with different strengths and weaknesses, but they share the same underlying approach, which is to define a biochemical system by a set of rules rather than reactions. A rule stipulates those elements which are important for (or altered by) an interaction, and ignores the rest. This “don’t care, don’t write” (Danos et al., 2007) approach sees molecules as bundles of contexts, to which rules may be logically applied. Parts of a molecule which is not part of the rule (the invariant context) does not need to be specified, because it undergoes no change.

Rules can therefore recondense a reaction set like our previous example back to a single line of code, removing our need to explicitly define except when required. This does not necessarily address the issue of computation however. Being able to tractably write out a complex multistate reaction network does not solve the problem of simulating, as the set of possible states still exists within the model. There are a number of methods implemented to simulate these rule-based models which work to reduce the computational load and allow for simulations that run at acceptable rates. One such method is to define each molecule in the simulation as its own entity, which can then be acted on by the rules when conditions match up. This has the advantage of reducing at simulation time the state space the simulator must consider, as only those currently present need to be considered when applying rules. For larger models however the tracking of each entity can override this advantage. As computational power increases, explicitly spatial variants such as MCell (Kerr et al., 2008; Stiles, Van Helden, Bartol, Salpeter, & Salpeter, 1996) have become more tractable for researchers also.

Another method is to take rules and defined agents and simulate them using the continuous time Monte-Carlo algorithm, or other similar algorithms for solving

rules stochastically. In Kappa, for example, each rule is given an instantaneous probability rate constant for any instance of the left-hand side of the rule in the system. Additionally, the activity of a rule is determined by the number of instances of the left hand side currently within the system (therefore adhering to the law of mass action) (Boutillier et al., 2018).

All approaches do not completely resolve the dilemma but do considerably reduce both to comparative levels to those seen in deterministic modelling.

To Build a Model

How a model is built, much like how a traditional experiment is designed, can play an outsized role in its outcomes. If built too much towards a given theory, the model is skewed towards a mechanism too early and remove alternative explanations prematurely.³⁵ This is important to note because modellers rarely have a complete and biologically derived set of parameters for a given behaviour.

Of course, having a *complete* set of parameters would likely mean we knew a lot already about a mechanism and render the model moot. For most biological systems, this knowledge is still incomplete. Thus, a model must wrestle with being outwith a biological environment and make assumptions on what we think happens, without assuming too much that the model is built to fit to one interpretation of the data.

We can take robust measures to check and ensure a model does not ‘overfit’ the data it uses. Models are often made however in support a hypothesis, or some experimental result. Showing the proposed mechanism works *in silico* is used to validate the supposed biological instantiation of it. Yet this does not *prove* this mechanism is what is in play within an organism. There are two reasons why: the multiple realisability of molecular function, and the implicit assumptions modellers make during model construction. The two tend to co-occur.

Multiple Realisation of Function

As the FMRP example shows us, the function of a molecule in one setting does not mean the mechanism it is part of is the same in another. FMRP function may be the same (or close to the same) in both rodents and humans, but the mechanism they are embedded in may be different.

Just because a mechanism sketch produces the seen output, it does not mean that this *is* the mechanism used by the system, or that it is the *only* one. A neural phenomenon might be possible through the action of different mechanisms. Mechanisms are derived from observation, defined by the detection of casual patterns within a group of interacting components. The patterns we see may not be representative of the true mechanism at play.

Furthermore, unlike human-made mechanisms, these systems are not created *for* a specific function. Rather, the functions arise from the affinity and efficacy of molecular interactions inducing further binding, breakdown or creation of other

³⁵ The reverse extreme to this (building free of any theory) is wholly implausible, as mentioned earlier.

molecules. Interactions which preserve or improve the ability of a cell to survive are retained, and those that do not are lost or repressed.

Success of a set of functions is strongly linked to the intracellular and extracellular environments. The sensitivity of a mechanism to the environment is largely because of this link – they become tuned to a range of environmental parameters. Cells may however over time³⁶ experience several very different environments. As such, it is possible for additional mechanisms to exist that perform the same function but do so optimally under different conditions. Where there is overlap, these mechanisms may work in tandem. Alternatively, one or more may be repressed by the activity of another mechanism.

The cell therefore has a set of redundancy mechanisms which may exist in parallel, working optimally in environments less commonly experienced, or emerge only upon the loss or damage to another, more dominant mechanism that achieves the same result. In both cases, the function is preserved within acceptable limits. This is however not a deliberate choice of the cell, but rather an emergent property of the reorganisation of biochemical dynamics brought about by an extreme change in a parameter.³⁷

A biochemical model allows us to discover these hidden or undefined mechanisms, at least in part. By taking the experimental evidence for a mechanism, and more importantly the disputed aspects of it, we can begin to sketch alternative or parallel explanatory mechanisms, ones which we can predict via model outcomes that can then go on to be experimentally verified.

Implicit Modelling Assumptions

Until recently, many of these assumptions were borne out of necessity. This however has led to a normalisation of these assumptions, so to the new modeller they appear to be neutral or even positive aspect of model creation. This section is intended mainly to highlight these, I will leave it to others to provide a more in-depth discussion.

Models are simple because of computational tractability. When computer modelling began, you could only simulate a small number of variables, which

³⁶ This here refers to a single cell over time, and an organism historically – the number of different environmental exposures and ability to survive within them over time becomes hard-wired into the genetic/epigenetic system of a whole species.

³⁷ There is an argument here for going further and debating how the higher levels of organisation act to rescue the sensitivity of the whole organism to local cell environmental changes. This feeds nicely into the robustness arguments outlined in this chapter but is too far beyond the scope attempted here.

even then could take a considerable amount of time to complete (Marrink et al., 2019). Models were thus reduced to the perceived smallest number of entities or parameters necessary to generate the phenomena being investigated.

This does bring certain advantages. It is much easier to determine the minimal possible mechanism(s) for a given system output. It requires less experimental data sourcing, and it allows for lots of replication studies in a short timescale.

What it assumes however is that all system phenomena can be explained by looking at a smaller subset of the system. Not only that, but that the mechanism *necessarily* must work independently of the system it is part of. A model can go even further and suggest a mechanism must also be at a lower “level” than the system output is. These all may be true, but they have not been arrived at from any reasoning on the potential mechanism at play.

A biochemical model, in an idealised construction, is defined by its experimentally derived kinetic and concentration parameters alone. The existence of a pattern of function that is expected to give rise to a mechanism, if not predicted by such a model, is not a refutation of the model, but rather a refutation of the mechanism sketch used as a template for that model.

Rarely does such an idealised model exist. Too many molecules have no direct kinetic or concentration data, with values often derived from large averaged or normalised datasets, or indirectly from a confluence of data sources. Thus, the ability for a model to challenge a model directly is limited, and can be challenged itself for using too many assumptions. Modellers therefore will tend to generate models which initially fit to the dominant mechanism sketch in the literature.

Whilst not ideal, the exponential complexity of biochemical networks allows these models to still have powerful influences on shaping mechanism definitions. A modelling paper can highlight distinct outputs which do not yet have an experimental corollary, raise points of contradictory evidence, or even times where seemingly contradictory evidence sets are in fact both valid. Even if a model is generated with a complete bias towards a given mechanism, it may still produce behaviours not expected, and produce outputs that conflict with prevailing thought.

In this it shares much with traditional experimental work and design. We should however be keener on removing this assumption where possible. A modeller’s largest boon is that their experiments cost a fraction of the cost (and sometimes also time) of an experimental approach. There is therefore no burden on a modeller to assume that the current mechanisms thought to be at work are true or complete.

Models are rarely, if ever, produced to refute or falsify a mechanism. Even weaker aims, such as exploring the extent to which two mechanisms may work in parallel (both supportive and in dynamic opposition), is not a common amongst published models. This in part comes from the worries outlined earlier, that the lack of solid experimental parameter evidence makes any claims in this domain open to strong challenges.

I suggest that these problems however are no different to those seen in “wet” (*in vivo* and *in vitro*) approaches, but they are more openly visible to a wider and less specialised audience. Living organisms are automatically perceived as being more reliable sources of evidence, given they are the biological reality a model is attempting to represent *in silico*. This neglects the rather large fact however that “wet” experimentation is also an indirect representation of organismal reality.

Take for example electrophysiological recordings of hippocampal neurons. These take the movement of ions across cell membranes as a correlate for neuronal activity, and are commonly referred to as postsynaptic potentials (both excitatory and inhibitory depending on brain region and neuronal type). Reliably recording the output of a single neuron is extremely difficult, and so most electrophysiological outputs are of field potentials. In hippocampal plasticity experiments, this field excitatory postsynaptic potential (fEPSP) is usually taken from the stratum radiatum of the CA1.

Interpretation of this signal is reliant on a wide range of assumptions that are not readily accessible without prior experience with this methodology and are not to be found in any recent methods section³⁸. The measure recorded is the grouped signalling of a population of neurons and synapses. The largest assumption is that a glutamatergic signal is the only one being detected in the fEPSP readout, or arguably that the level of glutamatergic signalling is so overwhelmingly large to remove any complications from other co-current neurotransmitter signalling events. An fEPSP is a coarse readout of ion flow – it detects the movement of ions either away from or towards the recording electrode. It can be modulated only by movement of the electrode or the strength of the neuronal activity stimulation.

fEPSPs assume that the grouped activity measured is representative of a ‘standard’ response by a synapse in this region, and that this is common to all neurons present. As we have alluded to previously however, there is a growing range of evidence showing a broad range of cellular and synaptic heterogeneous functionality (Altschuler & Wu, 2010; Lengler, Jug, & Steger, 2013). That is to say

³⁸ Indeed, my own lack of prior hand-on experience with this methodology makes me defer here to (Herreras, 2016)

that the response of a given synapse may not be the same as another sampled, but that the type of information carried may be the same. The average readout may be falsely attributing a synaptic response profile which is much more diverse than can be detected through this manner.

Within the field itself, there are examples of diverse *in vivo* fEPSP profiles depending on if an animal is anaesthetised or not (Devonshire, Grandy, Dommett, & Greenfield, 2010; Raz et al., 2014; Sellers, Bennett, Hutt, Williams, & Fröhlich, 2015; Tian, Xu, Liu, Wang, & Hu, 2015), and even evidence that plasticity profiles may be drastically altered by surgery-induced changes to the cellular environment (Guéniot et al., 2020; Saxena & Maze, 2018). These confounds are hidden and not improved through reproduction in other labs. Laboratories use similar but not necessarily identical setups, with considerable variation in equipment, experimenter skill and environmental conditions. This is useful for ensuring the robustness of a finding, as this noise element ensures that the observed outcome was not a result of highly specific experimental setup. It does however mask the former confounds, as these are common to the general method.

This issue, much like that in modelling, does not require the rejection of all prior data collected, and an immediate change in practice. It merely reminds us that our measures are indirect and require multiple approaches in different tissues. Modelling should not be held to any higher level of accountability – the burden should not *solely* be on the model.

This issue, in part, stems from the relatively small number of biochemical modellers and models which exist. A computational model is vulnerable to the vagaries of how a model should be built, the hidden assumptions and lost information, as well as the relatively poor comprehension of model theory and implementation in the wider biomedical community. Norms are developed locally within lab groups, and only slowly integrate into frameworks for ‘standard’ practice (Hucka et al., 2003; Novère et al., 2005). This, coupled with the terminology such as parameter ‘estimation’ can make models comparatively look less robust and rigorous than *ex silico* approaches, relegating them to a supporting role³⁹. Models however have the potential to offer a way of bridging the gap between experiment and theory. On one hand, they offer a distillation of knowledge about a particular system that is hard to visualise or write about concisely. On the other, they provide and space to analyse competing experimental findings and uncover implicit methodological confounds. We shall explore these now.

³⁹ We will explore a potential solution to this in Chapter 6.

Types of Biochemical Model

Whilst all sharing the same characteristics, biochemical models are built for subtly different reasons. Each one requires relation of a model to experiment, but the interpretation of the outputs differs between each.

Models as Constrainers

Particularly true of more descriptive biochemical network models, these models highlight the necessary parts (or at least the parts most sensitive) within the set of molecules associated with a biological phenomenon. These broad, often large models help to form the boundary points that constrain future purported mechanisms. Unlike the other two model types, this one is sometimes 'hidden' in a longform, descriptive setting (e.g. a review paper) and may not necessarily be considered by some as a biochemical model. These models however are crucial to ground any more complicated biochemical models and are especially necessary for the reduction in molecular species and reactions used in early mathematical models of biological phenomena.

Models as Mechanism Sketches

This type of modelling is arguably the most dominant form for models that use experimentally derived parameters. The goal is to both validate the model to experiment, and then use that validation as grounds to infer potential mechanisms based on the internal model components and experimental findings unrelated to the model in question. This usually presents itself in the form of a set of predictions which require experimental validation.

This type plays important roles in determining both the limits of current experimental work and gaps which may not be readily apparent. Descriptive constraining models may be found, upon the attempt to generate a biochemical model, to be not based on any particular experimental finding, or to have presumed information about a mechanism which has since been refuted or challenged. Modellers, in having to acquire a broad understanding of a particular body of experimental work, will often find these in the process of model construction.

This can be as simple as a lack of evidence and kinetic data for the association of two molecules, up to a mechanism layout which is rendered implausible by the concentrations and localisations of molecules. This is an underappreciated use of models, and one I will bring up again shortly.

Models as Theory Builders

A model is only useful if there is a reason for its existence. In practice, a model is designed to explore a theory- molecule x performs function F in the mechanism M which realises phenomenon y . Alternatively, a question: does molecule x have the set of functions commonly attributed to it? These must be explicitly mentioned as the rationale behind the model's construction. Biochemical models for the sake of exploring just the properties of a network, without any formal theory or hypothesis, are both not built nor (usually) publishable. Yet the testing of hypotheses in 'experimental' work has for a long time accepted that having no formal theory or rationale behind an investigation beyond looking for interesting patterns is valid. Large -omic data sets are published with the hope of being useful, but are in themselves not directed towards a given theory or hypothesis beyond a broad sense that the data will inform about biological function (Coveney, Dougherty, & Highfeld, 2016).

Models are therefore much more strictly structured – they must be explicit in their assumptions and in most cases the theory they are built to explore (Guest & Martin, 2020). Building a model and *a posteriori* generating a hypothesis is all but impossible, given the model invariably requires coding and specificity in terms of its parameters and outputs.

Whether these positions are well-founded or not, they do often lead to models which are clear in their assumptions, questions, and rationale. This in turn generates a clear theory to which the model is designed to show support for⁴⁰. The outputs of the model are therefore clearly and readily tied to a given viewpoint on the nature of the molecular interactions under investigation.

Models are therefore well-placed to provide both a snapshot of current theory⁴¹ in a field and lend support to/refute it. This has led some to consider computational models as a 'third way' ⁴² methodology, reciprocally linking experimental work to

⁴⁰ To date, little biochemical modelling has set out to actively refute a theoretical position. Most are supportive of the prevailing view at the time of publication. Note that here theory may be substituted by mechanism.

⁴¹ Theory is rarely explicit in biology and is often a collection of working hypotheses that best explain the data, and the opinions of the senior researchers whose labs undertake the research (which may only be a handful of people). This can be seen as both a good and bad thing. I personally think solid theories may be necessary in the age of large datasets – see (Coveney et al., 2016) for a complementary view on this.

⁴² Others have noticed that wet labs increasingly incorporate computational approaches and have collapsed wet and dry into a so-called "moist" lab descriptor (Penders et al., 2008). I shall not be using that here. The idea, incidentally, is nothing new; early

nascent theoretical biology and other related disciplines (Chowdhury, 2013; Penders, Horstman, & Vos, 2008).

The advantage here is that biochemical models operate neither as wholly reductionistic or holistic. Kinetic parameters, concentrations and localisation constrain and shape the way molecules interact. But the pattern or pathway of interaction is in turn not deterministic, and dependent on surrounding context and fluctuations brought about by Brownian dynamics (Nicholson, 2019). Macroscopic properties may emerge from these microscopic stochastic interactions. For example, a molecule may have multiple binding partners. Unbalancing the concentrations of one of these partners may overwhelm the molecule, preventing binding of other molecules indirectly. This will in turn impact the system both up and down-stream of these other binding partners. Setting off a potentially complex chain reaction that explains a phenomenon.

Models, therefore, are vital in generating predictions when the molecular interactions become non-linear⁴³. They can dissect emergent properties in ways difficult/intractable in experimental contexts, and link purely theoretical or high-level data to an interaction network, thus providing a solid correlate.

While I mentioned earlier that models are built to one of these categories, it may be better to clarify that they will contain aspects of all three. Each category iteratively builds on the preceding form of model, and thus internalises its assumptions and properties as a result. I feel however it is useful for the purposes of this thesis to highlight how these elements are distinct, and hopefully help make better sense of the models to follow. In the following two chapters, I show two modelling attempts. Each tried to capture AMPAR dynamics in a dendritic spine, using data and models from extant literature sources to attempt to map current biochemical data on AMPARs with some of the experimentally derived observations detailed in Chapter 1.

neuroscientific practitioners such as Hodgkin and Huxley were blends of dry and wet science.

⁴³ Arguably, this is all interactions of molecules in biological systems.

CHAPTER 3

Maximalist Modelling and AMPS

The main experimental aim of this PhD has been to define and construct a model of AMPAR function at the synapse. Such a model I determined required a much more comprehensive signalling network to be modelled than had been attempted previously. In the first half of this chapter, I will explain the rationale behind “maximalist” modelling, arguing why I believe it to be a necessary next step in biochemical modelling of neural phenomena.

The model discussed in the second half of this chapter is the AMPAR-centric Model of Receptor Signalling (AMPS). It is an attempt to construct a comprehensive model of the postsynapse using standard biochemical modelling software. I will discuss the challenges that arise and arose from this kind of model, and how these should (and are) being addressed.

Part I

Models, now

As Chapter 2 showed, biochemical models are invaluable tools that can help understand key mechanisms underpinning biological phenomena. Commonly, their usage tends to maintain a close relationship between the model and complementary experimental work, ideally forming a feedback loop. This can be a powerful way of checking hypotheses and finding promising areas of further exploration.

A model like this is built to help validate an experimental paradigm (i.e., confirm that an experimental approach and finding(s) can be explained and therefore validated through a biochemical network model). This automatically restricts its usage. Models are never perfect and will always include some element of user bias. This applies to the construction of a model – it will be built in this situation to help explore the experimental data. But in doing so the model will assume certain elements of the experimental paradigm to be true facts about the phenomena being tested regardless of the paradigm. What is actually occurring however might not hold to these truths, which are derived from the experimental paradigm (or even from a wider theoretical cluster).

Alternative mechanisms and interpretations are potentially lost as a result, and a researcher may spend many years working with a powerful method of validating a paradigm which turns out to be incorrect. Internally, it will work, and data will validate the model which will direct towards further experimental work. But much like overfitting, the risk is that the paradigm is too narrow. It neglects the *context* of the phenomena.

We tend to adopt a narrow focus on our research day-to-day. This is necessary – you cannot contemplate every possible other interpretation of data without then changing your methodology part way, consciously or not. The longer we spend on a specific paradigm, the more likely we accept our hypotheses to be true. It becomes more a question of how data fits within *my* paradigm, rather than how much does it support or refute it.

This leaves us open to committing some lapses in judgement about our work. The main one to be concerned with here is that we forget that in many cases the components of the phenomena we are looking at have multiple roles outside the one we are investigating. There are two traps we may fall into here. One is to presume that the method of action in our paradigm holds in all cases. That is, the way in which we observe a component working *is* how it works *in vivo* normally,

under physiological conditions. Second, we presume that the surrounding interactions (and other roles that the component may have) have no effect on the component in the situation being explored.

Contextual Importance

This second is what I wish to focus on here. To reiterate, this omission of the bigger picture is cognitively and experimentally a sound one to begin with. Once a possible mechanism (or set of mechanisms) is established however, this approach stymies understanding. When we neglect to place things in context, we slowly neglect that the context is important (Heil et al., 2018).

In a broader sense we are quite aware of it. A paradigm wholly done *in vitro* does not make strong claims about the potential *in vivo* substantiation. Studies are critiqued for testing outwith physiological parameters such as temperature, pH and oxygen levels. These provide useful approximations to the true mechanisms of a phenomena, and to sketch out the framework quickly and in isolation.

The reasoning hidden behind all this is that the whole, complete mechanism that generates a phenomenon is wholly reducible to its parts. If you take each component in turn, understanding them in isolation, you can then piece them back together and understand the whole mechanism.

This is much like a jigsaw – each piece contains a unique part of the whole, which when fitted together correctly shows you everything about the jigsaw. You start by sketching out the boundaries, which are well defined and easy to spot. You then start to fill in the centre using different strategies. Sometimes you stick to the edges, using the pieces in place as clues to what fits and working inwards. Sometimes you take a cluster of similar pieces and fit together a section without knowing quite where it goes in the final picture.

This is fairly close to how paradigms work, with people using different strategies and working together to fill in the picture piece by piece. Completion means a full comprehension of the phenomena being investigated. In so doing however we are making some rather large assumptions.

The first is that there is only one correct way to complete the jigsaw. Each piece fits into a rigid structure that is the only possible combination of them that works. The second is on the surface slightly contradictory. The piece is only for the jigsaw being looked at, but we can use it to help work on other jigsaws. Our understanding of the piece in jigsaw A tells us something about jigsaw B. This might be the actual piece (the component) is the same in both, or that the shape of the piece is important (the content of the piece itself might not be the same).

The final assumption is related to the first. As we have all the pieces and put them together in the ‘right’ way, where the jigsaw is does not matter. The whole mechanism is self-contained and remains the same in different contexts. There is only then the need to work out how it fits into other mechanisms at its edges.

Going back to the case of FXS, these three assumptions can lead to important experimental evidence taking longer to parse and process into theory (e.g. Maurin, Zongaro, & Bardoni, 2014; Tabet, Vitale, & Moine, 2016). We all assume we are looking at the same jigsaw, that the findings in model animals are analogous to humans and that findings of a component in other contexts map to FXS mechanisms currently identified (see Dahlhaus (2018) for more on why this is problematic). And these are all up to a point, *necessary* assumptions to make (or at least reasonable).

I argue that the current state of things shows that we don’t have all the pieces, *and* our jigsaw mechanism is not merely reducible to the components within the synapse. Our working model of FXS is beginning to stabilise around assumptions that may no longer hold true. We need to take the experimental evidence we have and reassess it. Ideally in a transparent manner, that does not presume that the existing mechanism sketch is correct. One possible tool to help (but by no means solve) is a model that contains all the available evidence as well as data from other paradigms. This “maximalist” model can aid in the redrawing of mechanistic boundaries and highlight indirect emergent properties not obtainable from simpler models.

Part II:

What I talk about when I talk about Maximalist Modelling

A maximal model definition needs to be broad enough to work at multiple levels of neuroscience, but in some way distinguish them from striving towards a 1:1 model⁴⁴. In this way, I wish to make it clear from the onset that this approach is not the same as the whole-cell modelling theory of Karr (Karr et al., 2012; Karr, Takahashi, & Funahashi, 2015), although they do share some similar features.

I propose the following seven facets to maximalist models:

1. A maximal model intends to capture the mechanism(s) underpinning phenomena. This may extend beyond physical boundaries but is constrained *functionally*⁴⁵.
2. Maximalist models are transparent – they provide a way of exploring hypotheses without themselves imposing a bias on the results.
3. A maximalist model is maximal for a single network (e.g. signalling, metabolic, morphological). It does not presume to simulate every aspect of a system or organism⁴⁶.
4. Maximalist models are always deemed incomplete- they offer likely explanations based on a non-exhaustive set of initial conditions.
 - a. New data can be inserted into the system at any time.
5. Maximalist models should where possible be constantly running. The system runs and updates itself, and the user changes things to see what occurs. Outcome x at time t_1 may not hold at time t_2 when alteration A is made at both.

⁴⁴ Which is only particularly useful when you already know quite a lot about the system being modelled – enough that the actual model probably does not really need to exist.

⁴⁵ The level of this constraint is up for debate.

⁴⁶ i.e., maximalist models are horizontally maximal (within a level), not vertically maximal (between levels).

6. Maximalist models enable more interesting multi-scale and multi-state modelling.
 - a. Data integration is from a fluid self-maintaining system that will differ from point to point naturally (Viswan et al., 2018).
 - a. Emergent properties of one scale may be explainable by the fluctuations of another.
7. Maximalist models are educationally valuable tools

Let's discuss each in turn.

Capturing a Mechanism

A suitable explanation of a phenomenon is a description of a mechanism ('how does this mechanism explain the phenomenon?') (Eronen, 2008).⁴⁷ Mechanistic explanation aims to understand a function of the brain by description of a mechanism that makes it occur. These may need to be in turn described by mixtures of smaller mechanisms, but reduction is not always necessary. What is critical is that the mechanism is focussed on the phenomenon.

This might seem a rather banal statement. But what it is capturing is that the mechanism should not be constrained by physical boundaries, such as a cell membrane, nor restrict the level of explanation the mechanism ends up encompassing⁴⁸. The latter is particularly important here. A mechanism of synaptic plasticity for example may include the network of neighbouring neurons, glial interactions, as well as the more studied molecular components.

A maximalist model is thus comprised of the parts needed to explain the phenomenon. A synapse alone is not sufficient to explain plasticity at the synapse, and so the other components (both downstream and upstream causally) must be taken into account.

⁴⁷ There is an interesting debate that rages on how biologists (particularly neuroscientists) explain phenomena. In this thesis I take Craver's viewpoint to hew the closest to what we do (Craver, 2002, 2013). Or at the very least, what we *ought* to do. My definitions of a mechanism stem from his work, but other definitions are available (Nicholson, 2012).

⁴⁸ "Level of explanation" refers to the molecular cellular, systems etc. hierarchy commonly used to broadly define research approaches in biology.

Transparency

Bias is hard to avoid in modelling, as outlined earlier. A maximalist model, by incorporating contextual elements into its construction, mitigates some of the common biases. Selection of molecular components for example is no longer constrained to just those thought to be involved, and molecules that might have been actively disregarded are still included. The aim is to capture the environment the interactions that make up the phenomenon are situated within, without assuming too much as to what those interactions may be.

The real benefit of this aspect is the inclusion of alternative roles for molecules that may be playing key roles in the instantiation of a phenomenon. A molecule with many interactions will be restricted by concentrations of it and its interactors (Edelstein, Stefan, & Le Novère, 2010; Romano, Pharris, Patel, & Kinzer-Ursem, 2017; R. D. Roy, Rosenmund, & Stefan, 2017). The reaction speed also plays a role. These parameters in isolation do not capture the true activity of the molecule but is the common way of approaching biochemical modelling.

Transparency is thus an avoidance of isolating molecules from too many of the roles they play *in vivo*, thus capturing a more realistic possible system profile.

Horizontally Maximal

A maximalist model is *not* an attempt to capture everything in a biological organism. It is maximal in the things it includes that are a part of (or very close neighbours to) a mechanism that underpins a phenomenon. Returning to our synaptic plasticity example, this would not include the whole metabolic system or genetic translation. They are important inputs to the system in places, but simulation of them is not necessary within the maximalist model. They are situated at boundaries where the system takes in inputs.

These boundaries are, as previously mentioned, not restricted to physical boundaries, but rather explicitly defined and defended by the modeller. Maximalist models do need constraining however, and so a maximalist model should keep to being specific to a particular type of model. In this manner, maximalist models are *not* multi-scale models⁴⁹, and do not intend to be complex across levels of explanation. Rather, the model strives for comprehensiveness within an explanatory level (hence the term 'horizontal'). Returning to our synaptic model, this would mean keeping *only* to biochemical signalling pathways that are not involved in metabolic roles. Within this constraint, a maximalist model can maintain a sharp focus on the biochemical properties of a phenomenon, rather than necessarily striving to *explain* a phenomenon. These two may be the

⁴⁹ Although, as we will see, can be made part of one.

same, but to assume so in the construction or interpretation is misleading and reduces the usefulness of the model to other investigators.

Incompleteness

A maximalist model is likely to be always incomplete in its coverage, and this should be recognised during construction. There may be components not yet discovered which play a role in the mechanism, there may be a complex interaction with another mechanism, or we might be looking at the wrong mechanism entirely.

Recognising that the model is never complete or veridical changes how we update the model. The model changes shape and size depending on new information and the outcomes are assessed. There is no pre-filtering of the new information based on how well it fits to a hypothesis or presumed method of action, and no need to rebuild the model from scratch. Data can be inputted and updated, and the outcomes assessed by researchers independently (Viswan et al., 2018).⁵⁰

Scientists arguably use abductive reasoning when interpreting data. Otherwise known as inference to the best explanation, we use incomplete information to presume the most likely conclusion (Baral, 2000; Magnani, 2001). However, the interpretation of models tends toward more inductive reasoning (van Riel, 2006). This is taking the model observations as a specific complete unit and forming a general conclusion. The difference is that abductive reasoning is pragmatic- it is aware its 'inputs' are incomplete, and that its interpretations are predictive⁵¹.

Because of this, it is much more open to moving away from what might be the standard form of explanation for a phenomenon. What matters is what best explains the observations. As such, this form of thinking can be subject to change – given new observations, the best explanation may shift. The same occurs in inductive reasoning – however it is easy during researching for the induction to favour the explanation assumed prior to the induction. If the observations lead to

⁵⁰ How decisions are made on if this data should be incorporated more permanently into the base model is something I have not found a satisfactory answer to. One possible method is that all is included and updated every time, but these reactions can be turned off.

⁵¹ To use a wonderful example from Preyer & Mans, (1999) (cited in Walton, (2001): Deductive reasoning is taking a marble at random from a bag that only contains red marbles and inferring that the marble you have selected is red. Inductive would be where the contents of the bag were unknown to you, but you select a few at random and they are all red. You could infer that all the marbles are therefore red. Abductive reasoning would be finding a red marble near to a bag of red marbles and inferring this marble is from the bag you can see.

a general conclusion consistent with current understanding, there is no need to explore an alternative explanation.

Maximalist models are positioned much more solidly in abductive reasoning. They are transparent in their adherence to any preferred hypothesis; thus, they may provide multiple explanations. Despite their complexity, they should not be striving towards deductive explanation either. They are tools to make sense of, and guide future gathering of, data derived from *ex silico* experiments. Maximalist models are a way of opening biological explanations to not having to assume they need eventually reduce to only molecular components, whilst not dismissing that the reductionist position may be true.

Always Active

Biological systems are almost always constantly running features of an organism which carry out a set of functions that maintain organismal integrity and therefore themselves (Varela et al., 1974). This is something the Maximalist model idea is trying to emphasise; there is no neatly defined system that we can model in isolation, nor should we. A model needs to consider the fact that a system does not only run once, that it is an active and dynamic process.

Take for example a signalling pathway. A model might define the inputs as the binding of a ligand to a receptor and a change in ion concentrations that leads to a biochemical cascade resulting in protein synthesis. The system starts at the ligand and has clearly defined outputs from the system's "end".

The actual system has none of these properties, of a start and end and a proscribed sequence of events. This is all imposed on the system to make sense of it (Villalobos & Ward, 2015; Weber & Varela, 2002). Any single part of a system can be seen as the start or end of it – there is always something upstream or downstream affecting or being affected. Similarly, just because the components are there, does not guarantee they interact successfully⁵². Modellers need to start to realise this more explicitly when constructing simulations.

Practically, this entails a model that is always "on" – users modify the system through changes to parameters, but not from a standing start. The way a system acts at one time point to an intervention may not be the same at another, due to the indirect effects context provides; maximalist models consider the history of the system, as well as the currently observed state. Current models (and indeed the models in this thesis) will approximate this by allowing models to reach steady state values before manipulating or perturbing the system, or by at least giving

⁵² Here meaning the reaction the modeller expects to happen, happens.

the model an arbitrary time to settle post-initialisation. This does limit the starting conditions however to then a modeller-defined stable or settled state, assuming that doing so captures the effect of the perturbation more clearly. Which is indeed true, but the perturbation is based on physiological events which do not operate under such idealised conditions.

It may be the case that the Brownian storm of molecular interactions is a key regulatory part in determining how physiological events play out. For example, there are many aspects of cognition which are now thought to ride a so-called ‘wave’ of uncertainty, and that systems within the brain are close to being chaotic so as to maximise the ability to adapt and change responses based on information flow between an organism and its environment (see Clark, 2016). Why is this not the case biochemically? Stochasticity can cause spontaneous state changes that can dramatically alter the responses or localisation of other molecules, and may be necessary to prevent a system becoming ‘stuck’ in a stable state (Nicholson, 2019).

Rather than seeing a model as being initialised, run, analysed, and then rerun, we should begin to think about perturbing a constantly running system. It may be that in most cases this has no effect on the outcome, but that assumes too much that is not known. Especially as a maximalist model becomes more complex in terms of molecular species types, previous system inputs may have very long-term changes which prime the system to respond optimally (or at least as observed experimentally) to future inputs, both similar or different.

Enhancing Multi-x Modelling

The rationale behind maximalist models is very similar to some of those argued for multi-state models (Pharris et al., 2019; Stefan et al., 2014). In the latter, the argument given is that a model must avoid presuming the likely ways in which the model components interact. This has been often out of necessity, as modelling all possible states of a molecule with multiple different binding or modification sites becomes exponentially difficult the more molecules and binding partners are present in the model (see Chapter 4 for more on this).

Maximalist models on the other hand emphasise that a model should avoid presuming what molecules are the most important for understanding a biochemical mechanism. This therefore does not exclude a multi-state model from being a maximalist model, but means all maximalist models will likely be multi-state models. A multi-state model is already trying to avoid constraining model complexity, and a maximalist model is the ultimate version of this.

To address the other “multi-”modelling type, multi-scale models are, as we have seen, not maximalist models. This is not to say maximalist models reject the multi-

scale model process however. Two or more maximalist models can interact by supplying each other with inputs, as in a multi-scale model. These models are both independent maximalist models however, meaning they are reusable, have many components not perceived as important for a specific mechanism, and are constantly active.

Educationally Valuable Tools

A good explanation of a biological phenomenon generated by a model is one interpretable from the model without any need for highly specialised knowledge of the model or the system it is simulating. A user should be able to manipulate the model and make sense of how their interventions influence how the model acts. What component A does to B that causes C is not a necessary fact to possess to understand that A interacting with B causes C at rate x over time t .

An established maximalist model offers an alternative understanding of a biological system than traditional models. A traditional model is simple and emphasises the importance of the components; A does this, B does this etc. You understand the system by learning a causal chain of events that lead to a known outcome. In teaching situations, models are generally used to teach how the properties of certain things relate to how these things interact (Greca & Moreira, 2000).

An example might be a model of action potential firing. The model shows how the flow of sodium and potassium ions alters the relative voltage across the cell membrane, triggering signal transduction. By manipulating the concentration levels of these ions, or the number of ion channels, you can start to relate a taught system into mechanical terms, treating the components as parts and processes of a mechanism.

Maximalist models do not allow for this linear chain notion of cause and effect. They much more emphasise the complexity of the situation and allow the *user* to determine what has cause and effect. The user has much more control on constructing what is important to understand how the system works, and how it fits into the context. Rather than telling the user what the 'right' system is, and allowing just manipulation of those parts, the user of a maximalist model has to discover that system.

If this is coupled with the always active ideal of a maximalist model, you begin to have a model which is interacted with differently and will respond differently to user interventions over time. Within even relatively simple models, the number of possible combinations of parameters, and states of model components, become intractable immensely quickly. Even well-chosen constraints based on data can

leave a large range of possible system states. A maximalist model, by being transparent in construction, cogent of system history, and contextually rich, offers a way of exploring this complexity.

Rebuttals and Challenges

I wish to now briefly address several rebuttals to this approach some readers may have.

Why bother?

On the surface, there appears to be no need to approach models in any way different to now. However, if we continue to produce kinetic models in piecemeal stages of only a handful of reaction species, we are hampering ourselves unnecessarily. There is the ability to readily acquire kinetic detail for biochemical reactions, and whole databases of proteins and molecules localised to specific cells or areas (Glont et al., 2020; Jeske, Placzek, Schomburg, Chang, & Schomburg, 2019; Sivakumaran, Hariharaputran, Mishra, & Bhalla, 2003; Wittig et al., 2012). The small model approach was valid prior to computational and other technological advances. Now, it is in danger of providing static pictures of functional processes which only partially reflect what really occurs *in vivo*.

Maximalist models are a way of seeing the whole picture- they can still be compartmentalised or reduced to a smaller model. What it does do though is allow any such modifications to be compared to the whole system, to check if there are differences that occur indirectly as a result or need to be present for the reduced system to compare. Even in its weakest form, a maximalist model is a closer approximation to an *in vivo* experiment than another model type.

How do you know when to stop adding to a maximalist model?

The extent to which a maximalist model needs to be exhaustive in what it includes is admittedly not well defined. I lean towards a hard boundary in that if the model is on signalling pathways, it should not begin to include metabolism or DNA modifications, for example. These are separate maximalist models. The maximalist notion only applies horizontally.

This does raise boundary questions. For example, at what point does DNA modification or metabolic reaction networks differ from signalling networks? And how is that reliably and consistently defined across modellers? This is not something I have an answer to, and is an area where the whole-cell modelling approach has no such problem (Karr et al., 2015).

Is continuous modelling necessary?

The systems we are investigating do not stop and start, they are part of a self-maintaining process. Continuous modelling abstracts away from seeing processes as hard-wired pathways and emphasises more interactions and emergence. It is not necessary for a maximalist model to do this, but its omission does remove some of the advantages of this approach.

Much of this advantage is in the perception of the data by the user. By having no control over the exact state prior to the changes made, there is a much stronger emphasis on understanding how the pre-change state influenced the outcome, if the outcome is different to previous results or unexpected in other ways. The relationship between the system state and the outcome now must be considered, rather than just focussing on the relationship between the changes made and the outcome.

What new insights might maximalist models bring right now?

Plenty of proteomic data is available for specific cell types and functional units. If this can be converted into maximalist models, they are a way of exploring these large datasets for indirect and emergent properties that are not possible to see in data snapshots (Büchel et al., 2013). It also allows for this data to be used without any specific hypothesis in mind, which increases the possibility of observing biologically relevant patterns. This does risk finding patterns where there are none however, and as such should be seen as a preliminary method of reducing the hypothesis space and pinpoint areas that warrant robust multi-method investigation.

Any examples of where maximalist models might be useful?

There are a few examples of work in neurobiology which lend themselves to a maximalist model approach. As of 2021, individuals exhibiting persistent deficits in social behaviour, communication, and a restriction in their habits triggered by altered sensitivity to sensory information are diagnosed with a single neurodevelopmental disorder, under the term Autism Spectrum Disorder, or ASD (American Psychiatric Association, 2013). This disorder is extremely heterogeneous⁵³, with no single causal trigger. Genetics plays some role, but mutations that can be linked to ASD currently only make up 20 % of reported cases, yet still implicate hundreds of genes (Huguet, Ey, & Bourgeron, 2013). Other links have been found related to prenatal exposure to air pollution and infection, amongst many others (Lyall et al., 2017), but this again offers no solid common theme or underlying mechanism to be targeted.

Maximalist models based on the proteome of affected individuals may help in determining where differences arise in biochemical network, and how they manifest. In instances where ASD is prevalent in a family, comparisons of neurotypical and neurodivergent members' maximalist models can be used to determine if any changes that present themselves might be a result of stochastic

⁵³ This is reflected in earlier definitions, which treated it as a diverse group of distinct disorders.

adjustments to biochemical networks at critical points that then over long timescales result in divergent phenotypes.

This is in effect taking the snapshots of genetic and epigenetic profiles of individuals and simulating them to assess *what* these differences in expression result in biochemically. More specifically, it is exploring whether the differences are observable only as emergent properties that result from a mixture of stochasticity and signalling pathway interactions.

Recap

A Maximalist Model is a horizontally complex model of a given system looking at one dimension (e.g., biochemical). It should be designed to run continually and based upon neutral parameter values. Interpretation of a maximalist model is dependent on user alterations at given times and parts of the model. These may not match with the same alterations at different time points done by others. A maximalist model critically does not presume any understanding of how its networks should operate.

It is then up to the user to determine how to use it. There are no constraints from the model itself. This requires no major alterations to how such models are constructed. For example, kinetic models run may already run indefinitely in this manner, as the equations they use do not have any defined “end” to them. But if you alter the perception of what ‘starting’ and ‘ending’ a model is, you begin to consider the history of the network into any interpretations of it. This is key to how best to incorporate holistic methodology into mechanistic modelling. These changes are ones of perspective. It’s seeing the house, not just the bricks.

Thus the “maximal” part of “maximalist model” is not about modelling *everything*. It is about modelling all the information on a given scale within a defined functional system. Information flow is constant, and is determined by the history, not just arbitrary user-defined triggers. From this emerges a richer interpretation of a system as an *activity*, not a structure.

Part III:

A Maximalist Model of the Synapse

The maximalist modelling idea was borne out of work I conducted on the development of a large-scale model of postsynaptic function. The aim was to use an aggregation of published models to generate a complex signalling snapshot of current knowledge of postsynaptic signalling, and use this to explore the effect of FMRP dysfunction across the whole network. In this section, I cover the component models, their validation, and the challenges that arose in the aggregation process.

Before exploring the model in detail, I have provided a brief background into some of the key principles behind this model type. Readers can skip to *An AMPAR-centric Model of Receptor Signalling (AMPS)* if this is already well known to them.

The Law of Mass Action

For biochemical models, two types of parameters are necessary, the rate constant and molecular concentrations. The rate constant of a reaction is a temperature sensitive parameter that in conjunction with the molecular concentrations, determines the reaction's rate and direction. The concentration of a molecule in a modelled system is the measure of the amount of a molecule within a volume.⁵⁴

Kinetic rates and concentrations are linked together in the law of mass action. Let's look at a simple reaction:



The forward and reverse reaction rates (or velocities) here can be written as:

$$\begin{aligned}v_1 &= k_1[A][B] \\v_2 &= k_2[C]\end{aligned}$$

If we take the reaction to be in a dynamic equilibrium, the forward and reverse velocities are equal:

$$k_1[A][B] = k_2[C]$$

⁵⁴ Given the diversity in sizes of biological volumes of the same type, and the time-consuming element of copy number calculation, concentrations have long been the default measure of molecular abundance. Biochemical models are usually solved in terms of particle numbers however and are preferable if available.

Rearranging this equation allows us to collapse the two rate constants into a singular dissociation constant: K_D ⁵⁵:


$$K_D = \frac{k_2}{k_1}$$

$$K_D = \frac{[A][B]}{[C]}$$

This law thusly defines a reaction rate in a well-mixed medium to be proportional to the concentrations of the interacting molecules.

Michaelis-Menten Enzyme Kinetics

Enzymatic reactions can be simplistically defined as:



where S is the substrate, P the product, and E the enzyme⁵⁶. In this chemical relationship, complex formation and conversion of S to P are defined by two kinetic parameters: K_m and K_{cat} respectively. The rate of product formation (aka flux) is thus defined as:

$$V_{max} \frac{[S]}{K_m + [S]}$$

Where $V_{max} = K_{cat}[E_T]$; ($[E_T]$ is the total concentration of enzyme present). This law is still based upon the mass action law, and is derived from an approximation of three differential equations:

$$\frac{d[S]}{dt} = -k_1[S][E] + k_2[ES]$$

$$\frac{d[ES]}{dt} = k_1[S][E] - (k_2 + k_3)[ES]$$

$$\frac{d[P]}{dt} = k_3[ES]$$

Where k_1 and k_2 are the reversible reaction constants, k_3 the irreversible reaction constant, and $[E]$ the concentration of unbound enzyme. Using this knowledge, one can convert any Michaelis-Menten reaction constants into their k_1 , k_2 and k_3 constituents (Upi Bhalla, personal communication):

⁵⁵ Biochemical models of the type illustrated in this thesis are concerned primarily with the k_n parameters, rather than K_D .

⁵⁶ For more details please see Sterratt, Graham, Gillies, & Willshaw, (2011).

$$\begin{aligned}k_1 &= \frac{(k_2 + k_3)}{K_m} \\k_2 &= K_{cat} \cdot (k_2/k_3) \\k_3 &= K_{cat}\end{aligned}$$

Conversion of enzymatic reactions into mass-action ones is done in all my models featured in this thesis. Michaelis-Menten equations hold when the system shows steady-state behaviour, but this occurs within a range of concentrations and rates which are unlikely to exist commonly *in vivo* (W. W. Chen, Niepel, & Sorger, 2010). Michaelis-Menten assumes (amongst other things) that the enzyme-complex concentrations exist in much lower numbers than the substrate, which holds for metabolic networks but less so for synaptic signalling pathways (Ciliberto, Capuani, & Tyson, 2007).

Ordinary Differential Equations (ODEs)

Biochemical models operate over time, and so require the conversion of the aforementioned laws into a set of differential equations that can be solved. The solution of these equations determines the rate of change in the molecules within the model over a user-defined timescale⁵⁷. This allows us to quantitatively analyse the model and relate findings to experimental data.

The important aspects of a reaction for the purposes of ODE formulation are the same as mentioned for model building: the concentrations of reacting molecules (or species), and the rates of the model reactions. Solving these equations is done using an ODE solver, a number of which are available, and are normally packaged up within a biochemical modelling software package such as COPASI (Hoops et al., 2006) or MOOSE (Dudani, Bhalla, & Ray, 2014; Ray, Deshpande, Dudani, & Bhalla, 2008). The differences between each solver are not considered by modellers during model construction, and as such will not be discussed here.

Compartments

A further manner of naturalising a model is to define semi-closed compartments that loosely represent different locations the modelled system may work in. For example, synaptic models may have a presynaptic and postsynaptic compartment, and may further subdivide these. By doing so, and setting the appropriate volume for these compartments, a model can mitigate the well-mixed medium assumption mentioned earlier. This can have large impacts where the concentration of molecules is measured as low relative to the whole investigated

⁵⁷ In so doing we make the further assumption that the reactions behave deterministically – one which has its drawbacks when looking at sub-cellular biochemical network scales. We shall return to this with some solutions later.

volume but known to be highly localised. A small reaction volume may greatly increase the chance of molecular interactions and make signalling networks viable.

This principle is seen within the gross morphology of almost all cells, which possess a singular thin protrusion known as a primary cilium. Signalling pathway proteins associated with Hedgehog, Wnt, mTOR and other G-protein coupled receptor-associated pathways have all been shown to either localise or migrate to this area to interact with other molecules, and to instigate crosstalk between signalling pathways (Wheway, Nazlamova, & Hancock, 2018).

Diffusion is also a key part of model-generated compartments, which require a way of acquiring and transmitting reaction results through the diffusion of molecular components. This diffusion parameter is often estimated as the compartment is man-made but may be calculable from experimental methods such as FRAP (Kang, Day, Kenworthy, & DiBenedetto, 2012).

Stochastic Modelling

As we have mentioned, the law of mass action approach requires several assumptions about molecular interactions and their environment: that the compartment volume is well mixed (i.e., no localisation), thermally stable, and that interactions between molecules is always equally successful. This is a deterministic modelling approach – the outcomes of the model will always be the same when the parameters and initial conditions are unchanged.

The deterministic approach works very well for large systems, where behaviour operates on a macroscopic scale. Any noise or fluctuations that are present in chemical interactions tend to cancel out, leaving a consistent output. This approach is arguably less appropriate for smaller, microscopic systems. No cell is well-mixed, nor thermally stable, and molecular species range from the highly diffuse through to extremely localised. Furthermore, molecular copy numbers for some species can be very low (Aldridge, Burke, Lauffenburger, & Sorger, 2006).

Low numbers (< 1000) of any species within a system makes the reactions they are part of much ‘noisier’ – just because they are present does not mean they are accessible to or sufficiently close to a reactant, on top of the fact that an interaction does not always elicit a ‘successful’ outcome. The environment of a reaction system at microscopic levels has a much greater impact on outcomes too – Brownian motion and specialised semi-closed compartments can vastly alter the probability of a successful interaction (Ait-Haddou & Herzog, 2003; Singh & Bhalla, 2018). The environment itself is also dynamic – the compartment may change in volume or area, may be created or disappear entirely, new molecules

may enter the vicinity which might alter the likelihood of another reaction's success.

Stochastic modelling is the attempt to capture some of this noise within a model, to better represent these fluctuations and determine if they are part of, or impact, the outcome of a proposed mechanistic pathway. Simply put, the model shifts from being about rates of change to probability of change.

Again, much like deterministic models, the solvers used for simulations are varied and part of standard software packages and are rarely considered during model construction. They may also be combined with deterministic solvers in so-called hybrid solvers. This works by the fact that if a reaction involves species present in large numbers⁵⁸, then the deterministic solution is quicker and equivalent to a stochastic one, by virtue that the noise is cancelled out. By separating deterministic from stochastic in this manner an especially complex model can be simulated much more quickly without any loss of complexity.

I have made use of both deterministic and stochastic models in Chapter 3, and a slightly different variant to stochastic modelling in Chapter 4.

⁵⁸ For reference, COPASI defines high to be 1000 particles, although this value can be adjusted if required.

An AMPAR-centric Model of Receptor Signalling (AMPS)

Synaptic plasticity is thought to be the basis of learning and memory (Martin et al., 2000). The synaptic environment changes as a result of the activity levels at the synapse. These changes are felt to act as a “store” of information that can be accessed through the reactivation of the synapse (Heynen, Abraham, & Bear, 1996; Milstein et al., 2020). A configuration of synapses firing in a certain pattern is thus the neural correlate of a memory.

Much research has gone into the study of the mechanisms that play a role in plasticity, particularly focussed on the postsynaptic region. A single dendritic spine contains at least 1000 different proteins in the postsynaptic density alone (Roy et al., 2018), and dendritic spines undergo dramatic changes to their morphology depending on their activity level (Bosch & Hayashi, 2012). The complexity of such overlapping mechanisms limits the ability of experimental work to investigate closely how experimentally derived mechanisms interact, interplay and dynamically alter under different conditions.

At the time of conceiving of this model, very few attempts had been published that looked at multiple postsynaptic signalling pathways simultaneously (Heil et al., 2018). There were however several different models created by Upinder Bhalla and his colleagues that over time have created models related to different key signalling pathways in glutamatergic spines (e.g., Bhalla & Iyengar, 1999; Hayer & Bhalla, 2005; Jain & Bhalla, 2009). Others have modelled broad but moderately simplified reaction networks (Gallimore, Aricescu, Yuzaki, & Calinescu, 2016; Gallimore, Kim, Tanaka-Yamamoto, & De Schutter, 2018; Jędrzejewska-Szmek & Blackwell, 2019), others more complex but highly constrained to a handful of key molecules (L. Li et al., 2012; Pharris et al., 2019; Stefan, Edelstein, & Novere, 2008).

The aim of this research programme was to develop a comprehensive model of AMPARs and their role in synaptic plasticity. This would include as many signalling pathways as possible and extend to local translation of mRNA at the spine base. This was done with the idea that this model could explore potential AMPAR-associated changes that might map to FXS aetiology, and help in understanding the current experimental evidence within a more explicitly mechanistic framework.

To construct this model, I decided to combine the existing models developed by the Bhalla lab, which would provide most of the key molecules associated with both AMPARs and synaptic plasticity. Initially, this required understanding and

reproducing a model which focussed on the interaction of AMPARs with their major kinase and phosphatase partners, which I shall now discuss.

AMPARs, Enzymes, Switches

The model in question comes from Hayer & Bhalla (2005) (henceforth HB). The original outcome of this model was the discovery of the feasibility of molecular switches for both AMPARs and CaMKII that were both stable and defined distinct “high” and “low” states for a synapse. The existence of switches was first postulated as a way of preserving synaptic information in the face of high molecular turnover (J. E. Lisman, 1985). Whilst proteins can have half-lives approaching six days or longer (L. D. Cohen & Ziv, 2019), in highly active conditions the turnover rate can increase dramatically, likely in part to the concurrent activation of the ubiquitin tagging complexes (Djakovic, Schwarz, Barylko, DeMartino, & Patrick, 2009; Widagdo et al., 2015).

This high turnover leaves a problem for the assumption that memory is, in part, synaptically “stored”, and that the activation of a network of synapses is required for the reconstruction of the original perceived input. The timescales generate a ‘Ship of Theseus’-esque dilemma – how can a memory trace be preserved if everything is replaced?

A switch works with the phenomenon of turnover, rather than against it, to preserve internal states of the synapse long after the original conditions that generated have ceased (J. E. Lisman, 1985). Molecular turnover appears to be remarkably robust in terms of maintaining absolute numbers of molecule for a given area. So long as the number of available molecules is constant, newly synthesised molecules can be ‘entrained’ to the current system state, for example through post translational modifications, localisation, or complex formation (Bhalla, 2014).

This alone, HB note, is insufficient for memory stabilisation. Molecules may move about the synaptic bouton, diffuse in and out, and be localised to different areas or parts (e.g., intracellular vs membrane, PSD vs cytosol). Stabilisation of molecules in certain locations would also appear to be a necessary component – AMPARs can only be affected by glutamate when within the PSD-defined synaptic zone (Y. Yang, Wang, Frerking, & Zhou, 2008). The modifications to a molecule may help define its localisation by adjusting reaction rates or interaction feasibility. These adjustments too may help with shaping some of the stochasticity (‘noise’) that becomes more impactful when dealing with small numbers of

molecules. A switch may change states rapidly when present in small numbers, or in a large volume.⁵⁹

HB's model shows that an interpretation of experimental work down to a biochemical level creates a switch phenomenon in both AMPARs and CaMKII, and that these are stable, allow for protein turnover, and require/accommodate movement of molecules within the membrane. Some of the key biochemical molecules can be seen in **Figure 4** and **5**.

Model Insertion

HB demonstrate their evidence using six variants of the model, which is constructed modularly. These were constructed and implemented using the General Neural Simulation System (GENESIS) (Vayttaden & Bhalla, 2004), which was developed by Upi Bhalla. Whilst a successor to GENESIS, MOOSE (Dudani et al., 2014), was in development at the time of instigating this new model, it was not sufficiently complete to accommodate all of the original model's features. The model was instead implemented into COPASI, which similarly allows for the simulation of biochemical networks via both GUI and command-based methods (Hoops et al., 2006).

A copy of the models from the HB paper were downloaded from the Database of Quantitative Cellular Signalling (DOQCS) (Sivakumaran et al., 2003) in the SBML format (Hucka et al., 2003), which is compatible with COPASI. The first step was to ensure the model outputs matched with the findings in the paper. This had two purposes. The first was to ensure the model downloaded was the one used in the production of the paper figure. The second was to ensure that the differential solvers used by both GENESIS and COPASI, which are different, did not impact on the outputs.

⁵⁹ I do not want to make any claim this is an unwanted outcome here, as it is likely this phenomenon is exploited on occasion in biochemical systems.




Figure 4 HB Model: Key Components.

Much of the HB model is centred around AMPAR trafficking and how this relates to the kinase and phosphatase activity in both the PSD and cytosol. The model shows the complex interrelationship and feedback mechanisms between the main kinases (PKA, CaMKII) and phosphatases (PP1, PP2B), with high PKA and CaMKII activation saturating PP1 activity and enabling CaMKII to reach a stable highly phosphorylated state.




Figure 5 HB Model: AMPAR Phosphorylation States.

The HB model focusses on two serine residue sites on the GluA1 subunit, phosphorylated by CaMKII (S831) and PKA (S845). Each of the two GluA1 subunits in a GluA1/2 AMPAR are modelled, giving sixteen possible AMPAR state profiles in two locations (PSD and cytosol).

Model Verification

Although GENESIS and COPASI are related in how they solve model files, they do operate differently, and this may have an impact on the outputs in AMPS compared to the published models used as a base. The HB model is the most important model to AMPS, as it focusses on AMPARs and contains many major enzymatic and signalling pathways associated with synaptic function. Whilst the HB model was provided in the paper (and in DOQCS) as a SBML-compatible *xml* file, this was a fixed export of the GENESIS model, and as such there was no record of it being tested and validated by an external group.⁶⁰

The downloadable model file initially required a partial renaming of almost all species and reactions. This was to remove file paths which made up the names within the GENESIS file structure. For example, the species name *AC_slash_CaM_dot_PDE1_slash_CaM_dot_PDE1_slash_CaM_dot_PDE1_cplx* was reduced to *CaM_dot_PDE1_cplx*. Other species and reactions included strings such as “minus” (for a hyphen) and “star” (for a *, representing phosphorylation). As these names were not representative of the names in the model paper at a glance, and with many sharing the same initial string (being part of a group within the GENESIS model construction), the first steps with the model were to reduce and translate these names into one which was less ambiguous and closer to the ones listed in the supplementary data of the original paper.

This process also allowed for a manual check of all the species, reactions, and parameters based on the supplementary files, to ensure that the sourced model version was the same as the one used in the creation of the paper figures.

Figure Replications: HB Model

Graphs were generated in RStudio (<http://www.R-project.org/>; <http://www.rstudio.com/>), using the *ggplot2* (Wickham, 2016), *readxl* (Wickham & Bryan, 2019), *writexl* (Ooms, 2020), and *viridis* (Garnier, 2018) libraries, or in COPASI.

Component graphs of Figures 5 – 9 in the HB paper were selected to be recreated as part of the verification process. These were identified as being the most important ones for the purposes of AMPS, as they primarily show observed changes to AMPARs in the model as well as changes over time for their major kinases and phosphatase binding partners. Each figure was replicated based upon

⁶⁰ Two variants of the model were available in the EMBL BioModels Database (Malik-Sheriff et al., 2020), but were at the time not curated.

the paper's description of a) the model variant used, and b) any inputs or perturbations made during the simulation.

The new model shows a very strong qualitative fit to the published HB model, managing to replicate almost every figure (**Figures 6 -10**). The key difference is within the magnitude of the CaMKII activation response in **Figure 9** and **10**, which fail to reach the same particle number as in the HB paper. The key findings of the paper however: that CaMKII and AMPARs can be stabilised into "high" active states in the synapse, and that these may be complexly related to one another, are maintained.

Original HB Paper Figure Details

This section briefly explains the figures from the original HB paper, to help understanding the phenomena seen in the replications that follow. Readers can go to the original paper for more detailed information.

Figure 5

The bistability of AMPAR levels in the membrane and the relationship to calcium and cAMP levels is the main concern of this figure. AMPS verification was focussed on parts F and G. These look at the effects of cAMP pulses in deterministic and stochastic modes of the model. In both instances, raising cAMP levels can initiate the translocation of AMPARs to the surface of the dendritic spine. This “high” state can be reset by either temporarily removing all cAMP in the system (deterministic) or spontaneously after many hours (stochastic).

Figure 6

Derived from a simplified version of the model, with fewer phosphorylation sites and lower concentrations of PP1, no CaMKII, and altered rates for PKA and PP1 activity. If PKA levels are raised to 40 active molecules for 2400 s, PP1 activity becomes saturated, allowing for an increase in phosphorylated receptors localised both internally and to the membrane. This is maintained even if PKA levels return to basal levels. If PKA is then subsequently removed entirely, PP1 can return most of the receptors to a dephosphorylated state.

Figure 7

This figure explores the bistability of CaMKII, and how CaMKII phosphorylation can retain itself in a self-sustaining “high” state when PP1 activity is saturated, and PKA levels are kept low. Adding in a more detailed PKA model in stochastic simulations results in flips of the CaMKII population between “high” and “low” (unphosphorylated) states.

Figures 8 & 9

Both figures deal with the AMPAR and CaMKII bistability together. In figure 8, both AMPAR and CaMKII share the same pool of PP1. Transient high calcium levels are only sufficient to transiently increase AMPAR and CaMKII to a high state. cAMP increases can stabilise both populations in the high state. In figure 9, AMPARs and CaMKII have separate pools of PP1. Calcium triggers a slow increase in AMPAR levels indirectly through CaMKII, which quickly reaches a stable high state. cAMP triggers a faster increase of AMPARs to the high state, but no concurrent CaMKII activation. Subsequent increased calcium levels however can flip CaMKII into the high state. AMPARs can therefore exist in three populations – low, high with low S831 phosphorylation, or high with high S831 phosphorylation.




Figure 6 HB Replication: cAMP-AMPAR relationship resembles original model

The COPASI model is able to replicate the cAMP-mediated switches in AMPAR numbers from “low” to “high” activity states in both the deterministic (A) and stochastic (B) simulations. Note however that in AMPS the high state values exceed those of the HB model, and that a loss of cAMP does not lead to a switching of the AMPAR population back to the low state. In the deterministic simulation, 0.2 μM cAMP is added for 1000s. In the stochastic simulations, 4000 s of 0.2 μM cAMP is used to trigger the high state. Left hand figures adapted from Hayer and Bhalla (2005).




Figure 7 HB Replication: Receptor phosphorylation profiles and Protein Phosphatase 1 (PP1) levels in a bistability analysis are replicated exactly.

The COPASI instantiations on the right show an identical replication of the HB model (left hand figures). This is a bistability analysis of a simplified form of the AMPAR submodule, where species of internal (IR) and membrane (MR) receptors are phosphorylated by PKA and dephosphorylated by PP1 (only in the membrane species). In all four replications, PKA is raised to 40 molecules for 2400 s, and transiently set to 0 for 3600 s, indicated by the black horizontal bars on the original graphs. Left hand figures adapted from Hayer and Bhalla (2005).




Figure 8 HB Replication: CaMKII switches in the PSD Compartment successfully observed.

A key feature of the HB model is that CaMKII can switch between high and low states, and that the high state is largely stable over long timescales. AMPS is able to show that CaMKII in the PSD compartment is capable of maintaining both low and high states, as well as spontaneously switching between the two. (A) shows two runs where PKA concentrations are fixed, and (B) when PKA activation is driven by the system. The latter is much more likely to show CaMKII flipping between high and low states. Left hand figures adapted from Hayer and Bhalla (2005).




Figure 9 HB Replication: Calcium and cAMP effects on CaMKII and AMPARs in PSD closely resemble the published model.

AMPS exhibits a broadly faithful replication of the response of CaMKII and AMPAR species in response to calcium and cAMP inputs. Calcium injections into the model do not exhibit the same amplitude in activated CaMKII in AMPS, but this does not impact the AMPAR profile in either the calcium then cAMP or cAMP then calcium tests. AMPS also shows the same tight relationship between active CaMKII levels and the number of AMPARs in the membrane, with high AMPAR numbers unsustainable when CaMKII activation returns to a low state. Deterministic runs (A – B) use $2.7 \mu\text{M}$ calcium (500 s) and $0.108 \mu\text{M}$ cAMP (2000 s). Stochastic runs (C – D) show the model with no perturbation or with added cAMP approximately 1 hour in to trigger high states respectively. Left hand figures adapted from Hayer and Bhalla (2005).




Figure 10 HB Replication: Stability of CaMKII and AMPAR high states in weakly coupled version closely resembles published model.

This is a variant of the model, where the PP1 population within the PSD is split into two separate pools for AMPARs and CaMKII respectively to represent immobilised PP1 populations. The model shows a high level of stability in this format for both AMPARs in the membrane and active CaMKII levels. This finding holds for both deterministic and stochastic simulations. Both runs add $2.7 \mu\text{M}$ calcium (500 s) and $0.108 \mu\text{M}$ cAMP (2000 s) at the times indicated. Top two figures adapted from Hayer and Bhalla (2005).

Additional Models

Once confirmed that the model in COPASI matched closely to the published outcomes, the next phase of model construction was to begin to add in extra signalling pathways. Initially this began with models taken from DOQCS. This had the advantage of being in the same format as the HB model (as most were from the Bhalla lab), to have several overlapping components, and to consistently derive their kinetic sources in the same manner. The models were chosen for the relevance of the pathways they modelled to AMPAR function, not due to their findings or original aims.

The models chosen were the mGluR component of Iyengar & Bhalla (1999) (**Figure 11**), the Jain & Bhalla (2009) model (**Figure 12**), and the Wnt signalling model from Padala *et al.* (2017) (**Figure 13**) (Henceforth, IB, JB and RKTVD respectively). The full model therefore has a large number of AMPAR-related signalling networks, as well as those implicated in gene expression and protein translation (**Figure 14**).

The JB model is made up of nine sub-components:

- 1) BDNF receptor signalling
- 2) AKT signalling
- 3) 4E-BP model
- 4) S6 Kinase model
- 5) CaMKIII model
- 6) Protein synthesis model
- 7) CaM
- 8) PKC
- 9) MAPK model.

The RKVTD model is comprised of the ERK, PI3K/Akt, and Wnt/ β -catenin signalling pathways. A model with all the above components incorporated resulted in a model of 587 species and 1077 reactions spread across five compartments. Links to species, reactions, parameters and compartments can be found at https://github.com/yishusing/Thesis_Repository, along with the full model in SBML and COPASI formats (which are also available from https://github.com/yishusing/AMPA_model).




Figure 13 RKTVD Model: Key Components.

Unlike the previous models, the RKTVD model was not generated to explore synaptic plasticity, but to show how cancerous mutations in one or more of the pathway components disrupted the signalling network. Both EGF and Wnt have been implicated in synaptic function (Cerpa, Gambrell, Inestrosa, & Barria, 2011; Farías et al., 2009; Tang et al., 2015; Vargas, Fuenzalida, & Inestrosa, 2014; L. Wang et al., 2012), but are not commonly added to models of synaptic function, which was why it was identified as a candidate for merging into the larger AMPAR-centric model. Rounded arrowheads represent inhibition.




Figure 14 An Overview of AMPS Modules.

AMPS is built from four pre-existing models. Three of these are related, and from the same laboratory, focussing on the major synaptic signalling cascades associated with synaptic plasticity in the early 2000s. The other module is based on more recent findings on the involvement of Wnt signalling pathway components in synaptic plasticity across several brain regions and cell types.

Individual Model Verification

The first step prior to adding these models with HB was to verify them as best as possible within COPASI. As most of these models were run and tested using GENESIS, there was a possibility of some differences in response profiles that would need to be considered during the merging process. Replication also ensures that the model I had downloaded was the one used in the generation of the published paper's figures.

As with the HB verification, each figure was replicated based upon the paper's description of a) the model variant used, and b) any inputs or perturbations made during the simulation. AMPS outputs for all comparisons are on the right-hand side of the figures.

The JB model was tested by attempting to replicate Figure 8 from the original paper, which covers the model's response to LTP and LTD induction protocols which modify the available free calcium. The RKVTD model was tested by replicating Figure 2 of the original paper, which looks at the model function in normal conditions.

The JB model was downloaded from the paper's supplementary materials, and imported into COPASI. Some adjustments were required to allow for the replication of the paper figures. Twelve events were created, which allowed for the control of adding BDNF and Calcium into the model at specified time points. The pool of calcium was also changed from a reaction-based to a fixed parameter. This means that the concentration of calcium does not change throughout the simulation, unless manipulated by the manually added events.

We show that the model figures for both LTP and LTD can be accurately replicated within COPASI (**Figure 15****Figure 16**). The outputs in particular for the LTP contexts however do show some differences in the magnitude of the changes resulting from the context. The model was unchanged from the downloaded version, but the introduction of transient concentrations of Ca and BDNF had to be implemented from scratch in COPASI. I assume therefore that the differences are due to differences in how MOOSE/GENESIS and COPASI manage interventions.

We find a similarly strong fit with the RKVTD model when run in COPASI (**Figure 17**). This was expected due to the original published model being run in COPASI, but does indicate that there was no loss of model information or specific parameter setups that were not included in the published paper.




Figure 15 An accurate replication of Jain and Bhalla (2009) mTOR model, LTP variant

Replication of LTP protocol to generate responses in several different molecular pathways in the model. A 10 μM stimulus of Calcium is added for 1 second at 300 second intervals, paired with a 5 s 0.0037 μM BDNF pulse, and the response of different components of the model recorded. All graphs show very similar response profiles to the original paper (left hand figures), although Akt, CaMKIII and Protein traces do not reach the same maximal values as in the original model. Adapted from Jain and Bhalla (2009).




Figure 16 An accurate replication of Jain and Bhalla (2009) mTOR model, LTD variant.

Replication of LTD protocols to generate responses in several different molecular pathways in the model. The LTD protocol consists of a 900 s pulse of 1 μM or 0.2 μM Calcium combined with a 900 s pulse of 0.0037 μM BDNF. All graphs show very similar response profiles to the published model (left hand figures), although the Akt traces do not reach the same maximal values as in the original model. Adapted from Jain and Bhalla (2009).




Figure 17 The Padala et al. (2017) model can be faithfully replicated

Replication of the model in COPASI show a faithful output (right) compared to the published findings (left), which look at levels of phosphorylated ERK (pERK) (A) and Akt (pAkt) (B), and the formation of β -catenin/TCF (C). The W parameter stands in for the presence of Wnt in the model, which affects two reactions – the activation of the Dishevelled protein, and the phosphorylation of Raf. Paper figures show data normalised to show percentage change, which they define as “[setting] the largest and smallest value in a data set as 100 and 0 percent respectively.” My outputs are kept to absolute values. Adapted from Padala et al. (2017) with permission from The Royal Society of Chemistry.

Merged Model Verification

After verifying that the component models could be reliably replicated in COPASI, I began the process of merging the models into the HB model. In this process, there were several instances where models shared reaction species. In most of these cases, because the models came from the same lab, these components could be merged quickly, as the parameters matched. However, there were some instances where the components were more complex, adding in extra reactions and modulators to existing pathways in the HB model.

In these instances, the newer model components were chosen to supersede the older ones where they had been changed due to new experimental evidence. These changes had no impact on the findings of the HB model, as the AMPAR system was unique to this model and unchanged by biochemical changes elsewhere.

To assess the quality of the merging process, the verification steps for JB and RKVTD models were repeated to assess how closely the outputs matched to the previous replications. Some change is to be expected, as all the chosen models have a degree of overlap in molecules and reactions used. There should however be at least some similarity in the outcomes seen.

IB, which had not been verified in COPASI previously, was explored in the merged model by comparing to Figure 4 of the original paper. Whilst the IB model components were broadly similar in how they responded to calcium inputs (**Figure 18 Figure 19**), the remaining model components showed drastically altered response profiles and concentrations (**Figure 20 Figure 22**). This made assessing the quality of the merging process very difficult from these figures alone. Whilst some qualitative aspects remained in the directionality of changes, the difference in complexity and increased use of certain species in reactions appears to have an observable impact on the model outputs. What this highlighted was that the progression of AMPS would require much more work to determine the extent to which these models could be successfully merged.




Figure 18 Replication of Iyengar and Bhalla (1999) mGluR model is comparable in the merged model

Above are examples of the effect of simulating a 100 Hz stimulus generating a 1 s 10 μM Ca transient in the postsynaptic neuron three times, with a 600 s interval. All show the same qualitative shape, but not the same as published. Graphs are compared with the filled triangle traces from the IB model graphs, representing the complete model. The non-compared square traces are for a model with feedback blocked by fixing concentrations of AA. These molecules have a considerable overlap in other incorporated models, increasing the number of reactions they are part of and reducing the capacity to replicate the published findings. From (Bhalla & Iyengar, 1999). Reprinted with permission from AAAS.




Figure 19 Replication of Iyengar and Bhalla (1999) mGluR model is comparable in the merged model cont.

MAPK, CaMKII, and Calcineurin dynamics are again altered compared to the isolated model here, perhaps in part due to the much greater range of reactions they are involved in globally within AMPS. Graphs are compared with the filled triangle traces from the IB model graphs, representing the complete model. The non-compared square traces are for a model with feedback blocked by fixing concentrations of AA. From (Bhalla & Iyengar, 1999). Reprinted with permission from AAAS.




Figure 20 Replication of Jain and Bhalla (2009) model LTP variant is only partially successful in the merged model

Replication of LTP protocol to generate responses in several different molecular pathways in the model. A 10 μM stimulus of Calcium is added for 1 second at 300 second intervals, paired with a 5 s 0.0037 μM BDNF pulse, and the response of different components of the model recorded. Graphs are compared with the clear square traces from the JB model graphs (left). Note that in none of the instances here were the reaction profiles the same as in the published paper or the previous replication. Adapted from Jain and Bhalla (2009).




Figure 21 Replication of Jain and Bhalla (2009) model, LTD variant is only partially successful in the merged model

LTD simulations carried out on the AMPS model showed some similarity between the more complex model and that originally published. Graphs are compared with the clear square traces from the JB model graphs (left). As in the LTP protocol, whilst some qualitative similarities are present, the merged version does not closely fit with the original. The LTD protocol consists of a 900 s pulse of 1 μM Calcium combined with 900 s pulse of 0.0037 μM BDNF. Adapted from Jain and Bhalla (2009).




Figure 22 Replication of Padala et al. (2017) model fails in the merged model

The model (right) shows no comparison with the original model (left) or the previous replication when in AMPS, either in transient concentrations, time scale or response profile (A – C). (C) shows both with (W=1) and without (W=0) Wnt – the results are close to identical, which is why only one trace is showing. Adapted from Padala et al. (2017) with permission from The Royal Society of Chemistry

Problems with Merging

The premise of this comprehensive model was to simply add the models together, their merging occurring along the shared species and reactions. From this base a model could then be generated focussing specifically on the biochemical relationship between AMPARs and newly incorporated FMRP signalling networks. Issues arose however due to several model assumptions in the source models, both explicit and hidden, which rendered this approach intractable.

As mentioned earlier, despite the model being compatible with SBML formatting, the model names for species and reactions were derived from the GENESIS format, whose structure led to many of these names having additional unwanted grouping information. Initially, these names were maintained, with only edits to replace unnecessary repetitions and words (i.e., the word 'slash' used in the species' name path), to clarify between species of similar names, and to make some names clearer (e.g. replacing 'star' with 'phos' or 'phosphorylated').

Much of this made the model clearer but there were still complications with the conversion from GENESIS to COPASI regarding kinetic parameter calculation. In the original model files, each reaction is assigned a unique function, which took the general form of:

$$\text{kinetic rate} * \text{substrates} / \text{compartment volume}$$

Whilst this converted into COPASI as a custom function for each reaction, these were not guaranteed to translate to other software, nor did the values now match the values listed in the publications, making checking if these values matched more time consuming.

The JB model consists of 17, equally sized compartments, six of which are empty. The model species are mostly in a single compartment, with only a few in the remaining ones. The volumes used (1×10^{-12} L) are consistent with the volume of just over one complete dendrite (van Elburg & van Ooyen, 2010). This rather generalised volume size does not lend itself well to merging into a model which has specific volumes for spine and immediately local dendritic space. Despite only the volume changing in the incorporation of these models, the outputs no longer matched to published findings. This is further compounded by a lack of information about localisation of molecular components of the pathways. For example, MAPK is found in PSD, spine and dendritic fractions, and where the reactions take place in this model is not identified. The relative volumes of each of these compartments differs greatly, ranging from 1×10^{-17} L – 1×10^{-12} L, and all have a different contextual background which may impact on MAPK interactions with other molecules.

The RKVTD model was generated for a generic cell, and so used a basic volume of 1. Whilst incorporating the model reactions was easier than the others, given a more transparent nomenclature, the model was unstable with the synaptic volume. Attempts to run the model when Wnt was present (indicated by changing the W parameter from 0 to 1) resulted in the model software crashing. The crashes occurred with no predictable pattern, making using this component difficult.

All these issues, and the attempts to ensure a true fit of these different models began to take on a much larger role than anticipated and detracted away from the main aim, that of modelling AMPARs in a rich contextual environment to explore the impact of FMRP-mediated changes to the biochemical network. To attempt to fix this, I communicated with Nisha Viswan, who was also working on validating the Bhalla lab models as part of the FindSIM pipeline development (Viswan et al., 2018). I worked on a MOOSE-compatible version of the model, removing all custom functions and restoring simple mass action kinetics for all reactions, with the parameter values found in the publications of each model.

This version of the model exceeds the computational capabilities of COPASI, with both deterministic and stochastic simulations failing to run. MOOSE however should be capable of simulating this model, despite the size. Attempts to run in MOOSE and verify the merging errors was passed on to Nisha. The merged model is to date still being tested and assessed as to why the merging process failed to show predictable and comprehensible outputs.

The failure of my endeavours here however has proven useful in exemplifying the difficulties of model merging, and the importance of consistent documentation and verification of models. They also emphasised that for models of this scale, a new approach to construction may help in preventing some of the challenges that I came upon here.

Taking Stock, and Becoming more AMPAR-centric

As alluded to earlier, the sheer complexity of the model also generated its own problems with being able to simulate reliably, and within a viable timeframe. Whilst AMPS was being tested and validated externally, it was necessary to rethink my modelling approach, and refocus on a still complex, but more constrained model of AMPAR function. The design of this model would still consider existing models. The key thing was to think more carefully about the design and engineering of the model throughout, such that the problems encountered in AMPS were less likely to occur.

In focussing more specifically on AMPARs the issue of combinatorial explosion needed to be addressed. In AMPS, just the S831 and S845 phosphorylation

AMPAR profiles required 32 different AMPAR species to be explicitly modelled (16 for both cytosolic and PSD compartments). Adding more residue modifications, glutamate binding, and binding partners such as GRIP and PICK, would make any COPASI-based model as complex and difficult as AMPS, with millions of possible receptor state profiles.

The regulation of AMPARs, and the interaction network, was at the heart of the original aims of AMPS. The new model therefore was conceived to capture the core features of how AMPARs act and are acted on in a functioning dendritic spine. In the next chapter, I outline this engineering process in detail, showing how the model evolved, and some examples of how it was tested through exploration of how AMPAR trafficking may occur *in vivo*.

CHAPTER 4

Kappa modelling of AMPARs

As mentioned in the previous two chapters, combinatorial explosion is a major confound in generating complex naturalistic biochemical models. My initial attempt at creating a maximalist model highlighted that an SBML approach would require an intractably large number of reactions written to fully capture AMPARs in all possible states. To overcome this, the latter half of this PhD has involved developing a rule-based model of AMPARs at the synapse.

This chapter explores the translation of elements of AMPS into Kappa, and the subsequent development of the model in detail. The model is designed with a much more specific focus on AMPAR activity, in an attempt to capture the whole lifecycle of the receptor in a single model. I provide a use case here based on investigating three major hypotheses of AMPAR trafficking, and how they match with current paradigms of synaptic plasticity. This research is deepened through the use of multiple stimulation protocols from *in vitro* and *in vivo* work, and the development of a naturalistic mechanism of varying free calcium concentrations to better assess feasible synaptic states. We demonstrate a model which is suited to investigate a wide range of AMPAR-related phenomena, is user-friendly, and quickly modularizable or adaptable to a user's needs.

A Recap of AMPAR Function

To refresh readers (and avoid needing to flip back to Chapter 1), I will begin with sketching the key aspects of AMPARs relevant to understanding the model development to follow.

AMPARs are tetrameric structures which form a gated ion channel whose open state can be stabilised by the binding of glutamate. They allow for ion flow across the membrane (largely sodium ions) when open, depolarising the dendritic spine and triggering changes in calcium levels that influence local signalling pathways (Shepherd & Huganir, 2007).

AMPARs are tightly linked with synaptic plasticity, with the numbers of AMPARs present (and their properties) in a synaptic zone usually representative of the activity level of that synapse in the recent past (Diering & Huganir, 2018; Huganir & Nicoll, 2013; Shepherd & Huganir, 2007).

LTP is a process that increases the sensitivity of a dendritic spine, causing it to respond more strongly (or 'loudly') to patterns of presynaptic activity. This is

associated with an increase in synaptic GluA1-containing AMPARs, which possess several modifiable amino acid residues on the intracellular C-terminal (Diering & Huganir, 2018; S. H. Shi et al., 2001; Tanaka & Hirano, 2012). These modulate the channel conductance, localisation, and synaptic stability. This process can be triggered by a high frequency presynaptic signalling pattern, or similar patterns applied to postsynaptic free calcium (Tim V.P. Bliss, Collingridge, Morris, & Reymann, 2018; Grover, Kim, Cooke, & Holmes, 2009; Nicoll & Roche, 2013).

LTD is an opposing process which is most often triggered by long-lasting low frequency stimulation. This reduces the level of synaptic AMPARs, making the dendritic spine less responsive to incoming information (less 'loud') (Cooper & Bear, 2012; Gallimore et al., 2018; Kemp, McQueen, Faulkes, & Bashir, 2000; Moretto & Passafaro, 2018).

In both cases, the trafficking of models both intra and inter-spine rely on the complex interplay of enzymatic networks, calcium level control, and specific protein machinery such as PICK1 or GRIP, which effect the ability of AMPARs to move through the membrane (Bissen et al., 2019). Whilst experimental output to date has told us much about where AMPARs are located, their responses, and their associated proteins, the actual biophysical mechanism behind this trafficking is still unclear, with several components being ascribed as being important (Huganir & Nicoll, 2013).

This model is developed to provide a way for multiple properties of AMPARs can be assessed simultaneously, to begin to better understand how AMPAR trafficking works at the synapse. To enhance the model, AMPAR properties stretch across the entire lifecycle of the receptor, from tetramer formation through to degradation mediated by ubiquitin tagging.

Kappa Modelling

Kappa is a rule-based language for complex biochemical systems (Danos & Laneve, 2004). As seen in Chapter 2, these languages help to reduce the number of reactions requiring explicit specification in a computer model by utilising rules which capture a set of reactions instead. Kappa treats the biochemical network as one in which agents have multiple properties and states. These agents interact with one another based on the rules written into the model (Danos et al., 2007). A rule specifies the necessary context for a reaction to occur. For example, an agent may have multiple properties, of which only two are required for a reaction. The rule will therefore explicitly reference these two properties and the state they need

to be in. The remaining properties, having no bearing on the reaction, will not be mentioned. If the context is met, the reaction may occur stochastically⁶¹. The reaction may change one or both properties, or be responsible for a change on another agent.

A rule can therefore apply to any molecule that fits those properties, and properties may be part of many rules. This greatly reduces the explicit number of reactions needing to be specified in a model, allowing for a much more naturalistic look at a biochemical network.

This approach allows modellers to explore multi-state complex molecules without needing to bias the model with assumptions on what is occurring. We take into account the whole context of the molecule, without needing to specify the whole context in the rules. This combination of avoiding the problem of combinatorial explosion and minimising modeller bias makes this language useful in exploring the idea of maximalist models in a more controlled setting.

Kappa has been used in previously published models to understand MAPK signalling (Suderman & Deeds, 2013), RNA Polymerase II phosphorylation (Aitken, Alexander, & Beggs, 2013), DNA repair (Köhler, Krivine, & Vidmar, 2014), and clathrin polymerisation (Sorokin, Heil, Armstrong, & Sorokina, 2018).

Understanding Kappa

A Kappa model, on top of the writing of rules, requires the declaration of model agents, their initialisation, variables, observations, and any modifications. Agents (*%agent:*) are the language's analogue for species or molecule – any model component which is part of a rule. When declaring an agent it is necessary to also declare all possible binding sites, properties and states that agent may possess in the model. For states, these must be declared between {} with spaces between each state.

Each agent requires an initialisation (*%init:*) that determines the number of the agent available on simulation start. This must also specify the state profile of any property that possess multiple states, as well as any binding partners. Variables (*%var:*) define elements such as rule rates, initial quantities of agents, compartment size etc. These may be numerical or a combination of other variables, and are not ascribed any units.

⁶¹ More formally, Kappa's syntax is context-free, meaning a rule's computational load is determined by the complexity of the left hand side rule pattern, rather than the total reactant numbers.

Rules (which have no declaring descriptor, but can be named) are written as a pattern of agents to be detected on the left-hand side. A match of this pattern applies the rule, which is the outcome listed on the right-hand side of the rule. Only those states explicitly mentioned in both sides of the rule determine both the match, and any changes (simplified to 'don't care, don't write') (Danos et al., 2007). The pattern of declared properties and bindings need to be the same on both sides for a valid rule.

A rule may be uni- or bi-directional. This is written as either a single headed or doubled head arrow respectively, which sits between the left-hand side and right-hand side patterns. The rate of a rule is specified after the right-hand side pattern, preceded by an @. The rule rate may be a numerical value, a variable name, an observable or a combination of these. For bidirectional rules, two values are needed, separated by a comma.

The simulation will run if all the above components are characterised, but will not provide any recorded output. The user must specify the observables (*%obs:*), which are normally a name (in single quotation marks) followed by the pattern to be matched (within two vertical lines). This pattern is counted, and the number of examples saved as a numerical value under the name provided by the user. As such, an observable may be the concatenation of patterns, or any other mathematical operation desired.

A final optional element is modifications (*%mod:*). These are a powerful advanced aspect of Kappa that allows a user to perturb the system at points determined by the state of the model, as well as by simple points in time. For example, an agent may be added at t seconds into the simulation, but it may also be added when the number agent C is higher than n . These perturbations can be single time events or primed to trigger every time the conditions specified hold true. This provides a user with many fine-grained ways of manipulating a model simulation. Expressions are Boolean in nature, and range from adding and deleting agents, through to updating of variables or the application of a rule under specific conditions. This element may also be used to provide snapshots of observables during simulation run-time, or even to interrupt the model and allow for changes to be made mid-run (which meets the always-on maximalist modelling criteria seen in Chapter 3).

Reaction rates are not explicitly provided units in Kappa. Caution must be exercised when entering variables to ensure they are correct. As Kappa works with particle numbers, and not concentrations, a stochastic rate constant is required for non-unimolecular reactions. This constant can be derived for a bimolecular reaction as follows:

$$\gamma = \frac{k}{N_A V}$$

where k is the concentration rate constant (in $M^{-1}s^{-1}$), N_A is Avogadro's number and V the volume of the compartment.

These components are listed here in a logical order, but it is not a requirement for simulation of the model. Kappa collates all specified files before running a simulation, allowing a user to separate out each model declarable element into separate files. Alternatively, a modeller may prefer to separate a large model by module. The latter makes it easier to add and remove model elements, whereas the former makes editing and later comparisons of model changes much easier to follow.

The model collation is always saved to the same location as where the eventual output of the model will be saved, as an *inputs.ka* file. Additional input files are named *inputs~n.ka*, with n being how many inputs since the original one. These files are automatically generated upon a simulation being initialised. Within each of these files is the full model, ordered into agent, variables (including observables), plot (variables to be observed), rule, modification, initialisation, and complete number of events before simulation termination. In addition, the head of the file provides the run's unique identifier (UUID), the random seed for the stochastic elements of the simulation, the plot period (i.e., the timescale for observation recording), the output file name, and two definitions of how the simulator should respond to event clashes and locked simulations. Running this input file within the simulator will produce an identical result to the original output, so long as no changes are made to this file or later on to the Kappa simulation package.

This preservation of model components in these snapshots helps record all model instantiations over time, leaving a traceable model history that can be used to compare model versions and check for unintended changes, or those not documented. It also provides an efficient way to assess model replicability and reproducibility, avoiding again any hidden model changes that can be common and prevent reuse of models by others.

Kappa Simulation

Kappa files (all of which end in the extension *.ka*) are simulated using KaSim (Boutillier et al., 2018). This is a command-line based stochastic simulator, for which a GUI is available. Experiments 1 – 5 were run on a HP z240 SFF workstation running Ubuntu 16.04 LTS, and later 18.04 LTS. Experiments 6 – 9

were run using the ECDF Linux Compute Cluster (Eddie) provided by the Edinburgh Compute and Data Facility (<https://www.ecdf.ed.ac.uk/>).

When running in command line, the user must specify the location of KaSim, the input files to be used, the timescale of observable recording, the duration of the simulation, and the output file name (saved as a *.out* file). Simulations may be run in parallel, although those run on a computer may experience slowdown of simulation run-time when there are multiple ones active. Currently, complex models are slow – many of the chapter’s simulations were set to 5000 s, which would take up to 48 hours to simulate.

This slowness in run time makes traditional methods of parametrisation impractical to deploy. As such, and in keeping with the spirit of model recycling, where at all possible parameters were derived from existing models or the DOQCS database (Sivakumaran et al., 2003). In instances where parameters required changes or had no such source, these were iteratively carried out manually until the model output met with expected behaviour derived from experimental outcomes.

Figure Generation

All graphs were generated in RStudio (<http://www.R-project.org/>; <http://www.rstudio.com/>), using the *ggplot2* (Wickham, 2016), *readxl* (Wickham & Bryan, 2019), *writexl* (Ooms, 2020), and *viridis* (Garnier, 2018) libraries. The latter library was used in all model-output figures apart from **Figures 25 – 36**.

Data Archiving

All graphs, output files, model components and associated files are found at https://github.com/yishusing/Thesis_Repository. All the experiments below can be recreated exactly from the material at this address.

Kappa can be downloaded from <https://kappalanguage.org/>

Experiment #I: Model Construction and Component Verification

Construction of an AMPAR-centric Model of a Generic Excitatory Spine

As no models of synaptic plasticity in the most recent Kappa language existed at the time of model construction, the entire model was constructed from scratch (including translation of old Kappa models). The model was built around the reactions necessary to explain the full lifecycle of an AMPAR, from synthesis through to ubiquitin-mediated breakdown. Synthesis rules allowed for the construction of GluA1/1, GluA1/2 or GluA2/3 dimers, which further combine into either GluA1 homomeric AMPARs, GluA1/2 AMPARs or GluA2/3 AMPARs. Each AMPAR subtype has a range of different properties which can possess one or more states, which act as a useful way to introduce the model's composition:

```
%agent: AMPAR_12(S831{u p pp},S845{u p pp},S880{u p pp},Y876{u p pp},S818{u p pp},Open{B S M L},Glu{g0 g1 g2 g3 g4},Anchor{endo, X, EZ, PSD},K868{u0 u1 u2 u3 u4 u5},K870{u0 u1 u2 u3 u4 u5}, K882 {u0 u1 u2 u3 u4 u5},DUB,Nedd4,RNF,pick1,grip,nsf)62
```

S831, S845, S880 + S818, and Y876 are residues targeted by CaMKII, PKA, PKC and SFK respectively. They are all important in controlling the localisation of AMPARs at the synapse. The number of possible phosphate groups depends on the number of subunits present with an appropriate residue. **Open** and **Glu** are Kappa translations of the different possible open configuration states of AMPARs determined in Dutta-Roy *et al.* (2015). **Anchor** is the location of the AMPAR in the spine. **K868**, **K870** and **K882** are the residues targeted for ubiquitination by Nedd4 and RNF167, which along with the de-ubiquitinating enzyme DUB make up the next three properties. Ubiquitination is a modifier which plays an unclear role in AMPAR trafficking beyond its role in regulating AMPAR degradation (Widagdo *et al.*, 2015, 2017). **pick1**, **grip**, **nsf** are binding site properties for three key molecules that regulate AMPAR movement within the membrane and cytosol.

Properties and states in **red** are Gallimore (2016) derived values, and were gradually added over the first few iterations of the model.

CaMKII autophosphorylation rules were taken from an in-house model generated by Tara O'Driscoll (O'Driscoll, 2017). Ubiquitination rules were similarly

⁶² Properties highlighted in red were not part of the original model. Anchor's original states were "tethered" and "untethered".

developed in-house by Ayisha Mahmood. PKC activation originally was modelled on the simple version in (Bhalla & Iyengar, 1999). I modified this after Experiment 4 to feature two distinct populations of PKC in cytosol and membrane. PKC could only interact with AMPARs (and thus be perceived as “active”) if bound to three calcium ions (Lipp & Reither, 2011; Torrecillas, Laynez, Menéndez, Corbalán-García, & Gómez-Fernández, 2004). The parameters for this model are derived from Antunes, Roque, & Simoes-de-Souza (2016). There is no data for which phosphatase acts on S818 nor the rate of PKC activity on S818, so I assume the same reaction rates as for S880⁶³. All other kinase and phosphatase reactions were taken from DOQCS. All enzymatic reactions and their rates were converted into mass-action kinetics using the following formula (Upi Bhalla, personal communication):

$$\begin{aligned}k_1 &= \frac{k_2 + k_3}{K_m} \\k_2 &= K_{cat} * 4 \\k_3 &= K_{cat}\end{aligned}$$

In all the findings represented here, the model is in a single-compartment configuration that represents a single dendritic spine. The volume of the spine (1e-16 L) was taken from AMPS and remained unchanged throughout all experiments. Kappa v4 is to-date not capable of multi-compartment setups⁶⁴, and so assumes (at least upon initialisation) a well-mixed system. In accordance with common practice, this issue is mitigated through the different properties of some agents in the model which have states that represent different locations within the spine. The **Anchor** tag above is one such example.

The model components are represented in **Figure 23**, and a complete list of agents and variables (and their sources) is available in the Thesis Repository. The model can be broken down into different component modules, although these are not any inherent property of the model’s construction. All the elements directly relate to one or more properties of the AMPARs present in the model, and include post-translational modification of primarily serine residues, trafficking of receptors, and control of receptor open states both in the absence and presence of glutamate (**Figure 24**). More information on all of these is found in Chapter 1.

⁶³ As a consequence this part of the model does not play a major role in many of the experiments seen in this chapter, but is an area necessary to develop in the short-term.

⁶⁴ An older version of Kappa does have a spatial variant which would allow for localisation (Sorokina, Sorokin, Armstrong, & Danos, 2013). This model was designed to be in the most recent (version 4) Kappa language however, which makes use of a slightly different syntax.




Figure 24 Simplified Schematic of Key Model Components.

- AMPA channel opening and glutamate binding. This is a translation of the Dutta-Roy et al. (2015) allosteric model of AMPAR open states (see Figure 25). Glutamate binding biases the channel towards more open states.
- Calcium regulated enzyme activation. Calcium directly activates CaM, Nedd4, and RNF167 indirectly via CaMKII. cAMP activity in the model is tied to calcium levels as an approximation for adenylate cyclase activity, represented by the dotted arrow.
- AMPA trafficking-related proteins. A translation of the Gallimore (2016) model, these molecules are well characterised regulators of AMPAR movement within the synapse between cytosol and membrane.
- Key AMPAR residues. The model has eight different residues that have been heavily implicated in AMPAR function at the synapse.

Verification of Model Translation

Three pre-existing models were utilised in the construction of the first κ AMPAR model⁶⁵. As discussed in Chapter 3, the combination of models can prove difficult in cases where models overlap, or critical information related to model construction is missing. To avoid this situation here, I made use of models which had almost no shared species, but all were connected via either calcium or AMPARs. Incorporation of these models could then be applied wholesale once reproduced within Kappa with a minimised risk of unforeseen incompatibility. It also allowed for more effective ‘silencing’ of a component if required.

Dutta-Roy (2015) Model

This model was developed to explore how AMPARs moved between different conductance states in the presence of an agonist. It assumes a Monod Wyman Changeux model of allosteric regulation (Monod, Wyman, & Changeux, 1965), which states the receptor must maintain symmetry. A change in the confirmation of one subunit necessitates the change of all subunits to preserve symmetry.

This model assumes four possible AMPAR configurations: Closed (Basal) or Open (Small, Medium, or Large). Within these configurations, a receptor may have up to four agonists bound, one for each subunit’s binding site. When no agonist is present, a receptor is overwhelmingly likely to be in the Basal state.

The model was translated into Kappa and was validated through the replication of the kinetic profile of AMPARs in the presence of 1 μ M of glutamate. **Figure 25** shows the average of 5 runs and visually shows a strong fit to the original model. This data was sufficient to satisfy that the model was reproducible in Kappa and therefore usable in this model.

O’Driscoll (2017) Model

Developed by Tara O’Driscoll, this model characterises CaMKII as individual subunits which are part of hexameric rings, that themselves can form into a stacked dimer. Each subunit is its own defined agent in the model, and may influence its neighbour’s phosphorylation state. A defining feature of CaMKII is its ability for Thr286-phosphorylated subunits to phosphorylate neighbouring subunits at Thr286, conferring autonomous activity on the molecule (Schworer, Colbran, Keefer, & Soderling, 1988). The mechanism by which this can occur

⁶⁵ Many of the AMPAR reactions are taken from the HB model, but due to the nature of rule based models, and the fact only a subset of the model was used, it is more correct to say that the new model shares the same parameter sets, taken from DOQCS.

however is still not fully characterised, and the O'Driscoll model was designed in part to investigate different possible ways in which this may occur.

The four mechanisms are unidirectional intrahexamer autophosphorylation, bidirectional intrahexamer autophosphorylation, unidirectional intrahexamer with interhexamer autophosphorylation, and bidirectional intrahexamer with interhexamer autophosphorylation. In the interests of choice, all four mechanisms were retained in the model, but only one (unidirectional intrahexamer autophosphorylation) is used by default.

This is the most conservative of the four mechanisms, as it only allows for autophosphorylation to travel in one direction around a hexamer. This makes it more sensitive to phosphatase concentrations, and requires heightened levels of calcium for long periods to achieve full phosphorylation. *Ex silico* evidence for both bidirectional and interhexamer autophosphorylation is currently constrained by a lack of fine-grained methodological tools available. Structurally, the interhexamer hypothesis also assumes physical configurations which are unlikely given our current understanding of the holoenzyme (Bradshaw, Hudmon, & Schulman, 2002). As they are not *de facto* ruled out however by these issues, there was no reason to omit them entirely from future consideration within this model.

Mahmood (2018) Model

Activity-driven ubiquitination of AMPARs appears to be just as crucial to synaptic plasticity as the diffusion and trapping of AMPARs within the synaptic zone. As the underpinning aim of this model is to capture AMPARs *in toto* at the synapse, having no capacity for their breakdown would have been an oversight, even if it as we shall see is not a defining feature of the findings of this chapter.

I have used a model developed in-house by Ayisha Mahmood, which provides a slightly simplified mechanism of enzyme mediated ubiquitination and deubiquitination of AMPARs. Three characterised ubiquitination sites are modelled: K868 on GluA1 subunits, and K870 and K882 on GluA2 subunits. As a simplification, only one relevant subunit site per receptor agent was defined.

Each site may have up to 5 ubiquitin 'tags' added. Ubiquitination of K868 is mediated by Nedd4 (A. Lin et al., 2011; J. Zhu et al., 2017), and K870 and K882 sites by RNF167 (Lussier et al., 2012). Deubiquitination of all three sites is mediated by a deubiquitinating enzyme, DUB (Hu et al., 2002). These interactions may occur anywhere in the reaction volume.




Figure 25 Replication of the Dutta-Roy (2015) concerted opening model of AMPARs.

The Kappa instantiation (bottom right) shows a very strong qualitative fit to the published outputs of the Dutta-Roy (2015) model (bottom left), which simulates a possible relationship between glutamate binding and the size of AMPAR opening (top). This is a replication of a simulation where 50 AMPARs and 1 μ M glutamate are present (Some open states removed for clarity as in the original figure). Figure adapted from (Dutta-Roy et al., 2015)

Prior to tagging, both Nedd4 and RNF167 require the transfer of a ubiquitin protein from an E2 ligase. The E2 ligase is part of a reduced version of the early ubiquitination system implemented in the model. Nedd4 and RNF167 are activated through the binding of phosphorylated CaMKII, based on the view of Widagdo et al., (2017). These tags are transferred as state properties of the respective agents, as opposed to being simulated as individual agents. The model assumes that an unlimited supply of ubiquitin is available to the system. Ubiquitin is an incredibly stable protein, and free ubiquitin levels make up a very large proportion of neuronal tissue fractions (C. W. Park & Ryu, 2014), which we feel justifies this assumption.

One change to model made at this stage was to include the prevention of ubiquitination at K868 sites when S845 is phosphorylated (based on Guntupalli et al., 2017) (**Figure 26A**). This does not remove any previously added tags, but they may still be removed by DUB. K870 and K882 are always available to RNF167 and DUB.

AMPA Synaptic Incorporation

Recent evidence has increasingly converged upon there being no directed movement of AMPARs through the membrane. Instead, halting their movement appears to be dictated by trapping through binding to proximal complexes, which is modified by auxiliary proteins such as TARPs and cornichons. Outwith these trapping events, receptors (or their subunits) traverse through the membrane in a manner consistent with Brownian motion (Tolle & Le Novère, 2010).

The AMPS model is based on work which does not explicitly investigate AMPAR trafficking dynamics in this way, and was instead focussed on the requirement of phosphorylation states for localisation⁶⁶. As such, there was no parameters in the model for a general rate of AMPAR movement.


To capture an estimate, test runs were generated in an attempt to reproduce model data from Tolle & Le Novère, (2010). Taking their diffusion constants as a guide, I estimated a parameter which was able to closely replicate their rate of AMPAR incorporation into the synaptic zone (which was dependent on the ratio of receptors:scaffold proteins) seen in the paper (**Figure 26B**).⁶⁷

⁶⁶ This as we saw in Chapter 1, and will again later in the chapter, is not mutually exclusive with the diffusional trapping (or ‘slot’) hypothesis.

⁶⁷ Movement was measured by the switching of an AMPAR tag from “untethered” to “tethered”. In later stages of the model, these parameters were replaced with those found in Gallimore (2016).

Figure 26 Initial AMPS rule calibration and assessment

(A) AMPAR ubiquitination is slowed if phosphorylated S845 is ruled to prevent E3 ligase binding. The model shows a decrease in the number and rate of polyubiquitination (U1 – U5) in 500 s (left) when compared to a model where S845 states have no effect on ligase binding (right). (B) AMPAR movement within the membrane (in the model the change from an “untethered” to a “tethered” state on the AMPAR) was calibrated to match with rates determined by Tolle & Le Novère, (2010), which assume movement of AMPARs within the membrane by Brownian motion alone. Each trace is representative of a different scaffold/AMPA ratio (0.25 – 2). Bottom left figure from Tolle & Le Novère, (2010) Red plus, 0.25; Green times, 0.5; Blue stars, 0.75; Purple squares, 1.0; Cyan filled squares, 1.25; Red circles, 1.5; Yellow bullets, 1.75; Blue triangles, 2.0.



Experiment #2 Piloting the Model

Introduction

Given how closely the model's structure adhered to the AMPAR-centric parts of AMPS, the first task was to investigate how the model responded to injections of calcium of different amplitudes and durations. This would allow me to assess how the model responded to its principal molecular driver, whilst also introducing how Kappa worked.

One open question was if the emergent properties observed in AMPS (e.g., saturation of PP1 required for CaMKII switch behaviour in physiological calcium concentrations) would map to a rule-based model. The extent to how localisation of molecules (albeit coarsely defined in COPASI) mattered was also under consideration, as was how easy it was to maintain a comparable level of basal calcium through a calibration of an artificial rule.

All these considerations went into the pilot stages of the model, outlined in this Experiment, and through to Experiment 6. The aim was to show a preservation of the phenomena shown by AMPS, and to use this as a springboard for further investigation into AMPAR function. As the simulations in Experiment 2 were an early pilot, only an n of 1 is used. As such, the following findings are treated as trends, and any conclusions as tentative.

Results

The model shows sensitivity across different calcium puncta events

As calcium is the start point for almost all the reactions in the model, it was important to ascertain how much calcium was required to see changes to the system state. I began by assessing what calcium level was sufficient to trigger a state change in the system beyond basal. A range of concentrations of calcium were injected at 100 s in separate simulations, from 0 through to 10000 molecules. There was distinct change in the profile of the system as calcium injection amounts increased (**Figures 27 – 36**). 100 molecules of calcium were sufficient to induce an observable change in cAMP, CaMKII and PKA levels, key molecules in the regulation of AMPAR phosphorylation (**Figures 30 – 33**). PKA activity was much more sensitive to low calcium concentrations than CaMKII, which only resulted in a maximum of 10% subunit activation in < 3000 calcium simulations.

Calcium injections importantly had no long-term impact on the observed amount of calcium within the model, with the basal average of less than 5 free molecules being maintained in all cases (**Figure 32**).

Synaptic AMPARs may be complexly related to calcium levels

In basal conditions, the default state of the model appears to replace GluA1/2 AMPARs with GluA2/3s, matching to early models of AMPAR population shifts over time (S. H. Shi et al., 2001). This maintains a stable number of synaptic receptors, but the newly incorporated GluA2/3 AMPARs are less responsive to incoming inputs. For low calcium injections (≤ 1000 ions), the replacement is halted and both populations remain stable for the duration of the simulation. Between 2000 – 3000 ions, the system destabilises, and GluA1/2 AMPARs are removed and at least partially replaced by GluA2/3 AMPARs. Past this level, the system again shifts, with GluA1/2 levels stabilising or increasing (**Figure 35**).

High calcium levels may be required for CaMKII switch

The model shows a distinct shift in response profile when calcium concentrations are briefly raised above 5000 molecules (**Figure 36**). In these conditions, CaMKII enters a high state, represented by being or approaching complete autophosphorylation. In the 10000 calcium ion simulation, this is stable for the duration. These findings are consistent with the CaMKII switch, which was shown in the HB model to be biochemically possible, and based on the work of Lisman and others as a possible site of information storage (J. E. Lisman, 1985; J. Lisman & Raghavachari, 2015; Singh & Bhalla, 2018).




Figure 27 Model S831 phosphorylation levels require high calcium concentrations.

CaMKII-mediated phosphorylation (single (S831_P) and double (S831_PP)) of GluA1/2 AMPARs showed a dose-dependent response, with the onset of phosphorylation also shortened with >1000 calcium ions. Receptor loss exceeds receptor generation as the simulation proceeds, so the receptor number changes throughout. All graphs are the result of single simulations (n=1), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 28 Model S845 phosphorylation levels are sensitive to even low calcium concentrations

PKA-mediated phosphorylation (single (S845_P) and double (S845_PP)) shows a broadly dose-dependent shift from unphosphorylated to doubly phosphorylated AMPARs. This response suggests that the single phosphorylated state is uncommonly observed. The 100 calcium ion response is more similar to the 3000 calcium ion response than the intermediate outputs. Receptor loss exceeds receptor generation as the simulation proceeds, so the receptor number changes throughout. All graphs are the result of single simulations ($n=1$), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 29 Synaptic AMPAR composition is overall similar to initial conditions when calcium is present.

Basal responses show a preference for GluA2/3 AMPARs (Synaptic_23) over GluA1/2 (Synaptic_12) in the synaptic zone, with the latter declining over time. Calcium appears to stabilise the populations, although high (>1000) calcium levels suggest a more unstable state for GluA1/2 AMPARs to be followed up in further work. However, in all cases GluA1/2 AMPARs predominate for the duration of the simulation. All graphs are the result of single simulations (n=1), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 30 Model kinase-phosphatase levels show sensitivity to small calcium level changes.

In all levels of calcium, the saturation of active PP1 (Free_PP) activity is necessary for PKA (PKA_a) activity to increase. PKC (actPKC_level) activity appears more frequently in higher calcium scenarios, but with no increase in the magnitude of activated molecules seen. All graphs are the result of single simulations ($n=1$), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 31 PP2B is more sensitive to calcium levels than CaMKII in the model.

Both CaMKII (phos_CK_hexamers and phos_CK_subunits) and PP2B (PP2B) show a dose-dependent increase in the level of active molecules. Activated PP2B occurs much quicker than CaMKII, and saturates at 2000 calcium ions, whilst CaMKII subunits are still only fractionally activated at high calcium scenarios. Fully active hexamers are only seen in high calcium scenarios. All graphs are the result of single simulations ($n=1$), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 32 Background calcium levels are not affected by injection of calcium into the model

Calcium injections do not result in elevated levels of background calcium, with all states averaging below 5 ions, consistent with experimental data. cAMP levels are dose dependent, showing constant fluctuation but staying elevated in all non-basal conditions for the duration of the simulation. All graphs are the result of single simulations ($n=1$), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 33 Active Calmodulin levels are only modestly affected by a transient increase in calcium levels.

Active calmodulin is defined as being bound to four calcium ions (Active_CaM). Only modest effects of high calcium scenarios are detected, with greater calcium levels largely increasing the rate at which active CaM is first observed. CaM does not remain bound to CaMKII in any large numbers in all scenarios (CaM_CaMKII). All graphs are the result of single simulations ($n=1$), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 34 Ubiquitination of AMPAR K868 Residues is reduced indirectly through calcium mediated PKA activation.

AMPA K868 ubiquitination (K868_ub – K868_ub5) is shown to be reduced in non-basal conditions due to the protective effect of S845 phosphorylation on the K868 site (Guntupalli et al., 2019). Very little ubiquitination is present in high calcium scenarios. All graphs are the result of single simulations (n=1), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 35 Very high calcium causes strong receptor phosphorylation but no change to synaptic AMPAR composition.

Extremely high levels of calcium drive residue phosphorylation to saturate and prevent any loss of GluA1/2 receptors (Synaptic_12) from the synaptic zone. PKA activation (PKA_a) shows a brief but pronounced spike. All graphs are the result of single simulations (n=1), with the extra calcium added at 100s, and show molecular copy numbers.




Figure 36 Very high calcium saturates CaMKII activation and increases cAMP without permanently raising basal calcium levels.

CaMKII subunits (phos_CK_subunits) exhibit saturation in activation only in the very highest calcium scenario. CaM (Active_CaM) and cAMP levels show transient spikes in response to high calcium, with background calcium levels raised but still low for the duration of the simulation. Ubiquitination of AMPARs is further stymied by high calcium levels. All graphs are the result of single simulations ($n=1$), with the extra calcium added at 100s, and show molecular copy numbers.

Discussion

The findings of this study help suggest that the phenomena seen in AMPS were not due to the model type and methodology, but rather aspects of the biochemical mechanisms surrounding AMPAR function within the synapse. The model is sensitive to calcium, registering changes in enzyme activity even in the presence of 100 calcium ions. This sensitivity however is graded, with saturation effects only reached with > 3000 calcium ions injected into the model.

This pilot also shows tentative evidence of an interplay between synaptic receptor composition and stability and calcium levels. Low calcium levels preserved the status quo, whilst moderate levels induced instability into the GluA1/2 population. High levels reversed this, with GluA1/2 AMPARs supplanting GluA2/3. This may be due to the dynamics between PKA and PP2B action on the S845 residue. In low calcium levels, PKA and PP2B activity appears to be evenly matched. At moderate levels, PP2B activation plateaus, but this is significantly higher than available PKA, leading to greater S845 dephosphorylation and therefore increased chance of GluA1/2 loss from the synapse. High levels see a spike in PKA activation, which overcomes PP2B activity temporarily for the simulation duration as the PP2B concentration is at its maximum. These trends are tentative – with only a single replication, these are informative for the following experiments but in no way definitive.

The model shows the ability to preserve average calcium levels to a physiologically accurate level, whilst still allowing for the random fluctuations that exist *in vivo*. This does come with the price of slowing down the model runtime, but this is not a large an issue given this early stage of model development.

Whilst these findings were promising, they also highlighted a few required changes to the model in order to ensure it more faithfully represented *in vivo* conditions. The particle numbers used here for calmodulin, CaMKII and AMPARs are too low for a spine of this size if we compare to similar synaptic models (Hayer & Bhalla, 2005), although evidence for the quantities of AMPAR in particular are not well-characterised.

The evidence from this pilot overall shows that the key aspects of the component models have been successfully incorporated and match broadly to expected outputs in a short-term study, but require more rigorous testing.

Summary

- The model shows a broadly physiologically relevant response to different concentrations of calcium.

- Model components work together, and no errors were reported
- Molecular copy numbers require adjusting to better represent *in vivo* environment

Experiment #3 Generating Naturalistic Glutamate and Calcium Fluctuations

Introduction

A key attribute of maximalist models (outlined in Chapter 3) is that they attempt to capture naturalistic conditions as much as possible. Rule based models capture the fluctuations in molecular copy numbers, but do not alone capture the context surrounding the model which may be defined as its input.

In the case of this model, the key input comes from pre-synaptic release of glutamate into the synaptic cleft. Glutamate binding to both ionotropic and metabotropic receptors post synaptically are a major driver of synaptic change, and the physical manner in which information is transferred (Diering & Huganir, 2018; Mukherjee & Manahan-Vaughan, 2013). Glutamate release events may be both activity-driven (i.e., by the arrival of an action potential to the presynaptic bouton) or spontaneous (i.e., triggered by local ion fluctuations) (Traynelis et al., 2010).

For the purposes of this model, the main outcome of these release events is that they alter the calcium levels of the postsynapse away from a baseline of 5 ions (approximately 0.08 μM , based on the dendritic spine volume used in Hayer & Bhalla (2005)). This may come from NMDA receptor channel opening, release from intracellular stores, or a combination of the two (Bender, Bender, Brasier, & Feldman, 2006; Mahajan & Nadkarni, 2019; Y. Wang, Rowan, & Anwyl, 1997). Presentation of the initial model instantiation led to questions asked about this aspect of the model, and whether it would be possible to include a way both glutamate release probability and calcium levels would fluctuate in manner consistent with spontaneous release events, and scale with AMPAR numbers.

Results

As both glutamate and calcium dynamics are very fast, a complex ruleset for either would be time-prohibitive, and in pilot tests often led to simulation failure. To reduce the load on the system, I decided that the level of calcium would fluctuate based upon the number of AMPARs present in the synaptic zone. This acts as a surrogate for the NMDAR-mediated changes to calcium, which is thought to be modulated *in vivo* by AMPAR numbers. More AMPARs lead to more (or more likely to lead to) strong depolarisation of the dendritic spine, which leads to greater fluxes in calcium levels mediated by NMDARs.

The model handles this by checking every 10 ms for the open states of all AMPARs currently identified as being within the synaptic zone. It also checks for the number of glutamate bound to the receptor. The combination of the two

properties generates a value for the number of calcium ions to be inserted into the system⁶⁸:

```
%obs: 'CaGlu'(0.15 * 'S1') + (0.6 * 'S2') + (1.05 * 'S3') + (1.5 * 'S4') + (0.3 * 'M1')
      + (1.2 * 'M2') + (2.1 * 'M3') + (3 * 'M4') + (0.45 * 'L1') + (1.8
      * 'L2') + (3.15 * 'L3') + (4.5 * 'L4')
```

```
%mod: alarm 0.01 [true] do $ADD CaGlu Ca(); repeat [true]
```

Glutamate release is determined by the detection of a particle dubbed “Random”. The creation and destruction of this particle is calibrated so as to exist very briefly.

```
.<-> Random() @ Random_In, Random_Out
```

When observed by the simulator, a modification command adjusts the rate variable ‘Glu_Release’:

```
%mod: 'Glu_Trigger' > 0 do UPDATE 'Glu_Release'
      3000000000000000/Theta; repeat [true]
%mod: 'Glu_Trigger' = 0 do $UPDATE 'Glu_Release' 0; repeat [true]
```

‘Glu_Release’ is tied to a very fast reaction which begins to add glutamate to the system for as long as the random particle is present.

```
.<-> Glutamate() @ Glu_Release, Glu_Remove
```

The random generation of a single trigger molecule causes the release of glutamate, which lasts for 1 ms – 0.5 ms to accumulate and 0.5 ms for removal.

⁶⁸ This was inspired in part by Benke & Traynelis (2019), but is not strictly tied to observed values. For a population of 40 AMPARs, the formula has a maximal output of 180 Ca²⁺, representative of 3 µM input.

This decouples the glutamate release (which is analogous to presynaptic firing) from postsynaptic responses, requiring the two to coincide for a calcium induction to be triggered. This is to reflect the fact that information transfer at synapses is not 100 %. Presynaptic terminals can fail to release neurotransmitter, or calcium fluxes may fail to change biochemical networks sufficiently for a change to be detected.

Note that in this instance there is an approximately 10 % chance of detecting a glutamate signal and eliciting a postsynaptic response. This rather low value is on the lower bound of some reports of *in vivo* release probability (Borst, 2010), but quite consistent with many studies on rat hippocampal neurons (Branco & Staras, 2009). Some examples of this release mechanism in action can be seen in **Figure 37** which shows the different open states of the AMPARs (see **Figure 25**) in eight exemplar simulation runs. Spontaneous release of glutamate allows AMPARs to open transiently, which can be seen partially at the seconds timescale. Outwith these events almost all AMPARs are in the basal state, with a small proportion open in the small state at any given time.




Figure 37 Activity-Induced AMPAR open state changes are possible in a rule-based model.

AMPArs are found most in the basal (closed) or small open configuration, with more open configurations much more rarely seen. Shown here are eight example runs of the model over 5 s to illustrate. AMPAR open state figures adapted from (Dutta-Roy et al., 2015).

Discussion

This addition to the model led to no substantial increase in the run-time of the simulation. The generation of the trigger molecule, and thus the frequency of glutamate release events, can be controlled by altering the reaction rate. The success of these events transferring information can be controlled by adjusting the time of sampling the synaptic receptor profile.

The Random rule at the rate given results in a ~1 Hz pulse of calcium entry. This can be altered to suit the requirements of spontaneous and basal signalling events, which have been estimated to occur between 0.1 to 1 Hz for CA1 pyramidal cells (Mizuseki & Buzsáki, 2013; Waters & Smith, 2002). Similarly, the detection of these events and subsequent calcium flux can also be adjusted by altering the *\$ADD CaGlu Ca()* alarm time. In the above example, only 10 % of glutamate events will lead to calcium changes, which can be seen as a 10 % success rate of presynaptic transmission of an action potential into glutamate release, and a 0.1 Hz rate of basal postsynaptic firing.

The frequency of glutamate release events, which as mentioned earlier is independent in the model from the detection mechanism, was set for all future simulations to 0.5 Hz, which lies at about the mean value for firing rates of neurons in the hippocampus and entorhinal cortex (Mizuseki & Buzsáki, 2013). This system, whilst not a true representation of calcium dynamics, is still much more 'noisy' than comparable published models, where such fluctuations are tightly controlled. In so doing however the model adheres more closely to the maximalist model requirement of naturalistic contexts and allows us to capture the possibility that spontaneous firing events may impact on the likelihood of successful changes to (or maintenance of) postsynaptic plasticity.

Summary

- Calcium levels can be successfully altered by tying synaptic AMPAR numbers and open state profiles to a standardised influx of calcium.
- Further filtering is possible by adjusting the rate of spontaneous release events, which can be altering during a simulation as well as during initialisation.
- These two mechanisms can be coupled together more or less through adjustments to the timescales at which the simulation checks for AMPAR open states.

Experiment #4 Exploring Different AMPAR Population Responses to LTP

Introduction

The model has shown itself to be similar to AMPS, to incorporate imported models correctly, and to be able to both generate and respond to naturalistic inputs on very small timescales that have observable impacts on the system state. This is all an ideal platform to therefore proceed to investigate how the system responds to calcium dynamics that are correlated with synaptic plasticity.

Changes to AMPAR numbers, subunit composition, and conductance state have all been shown to tightly correlate with long-term changes to synaptic responsiveness and downstream effects which are triggered by activity patterns at the synapse (see Chapter 1). In learning and memory, these are broadly sorted into LTP and LTD phenomena, which correspond to heightened and dampened synaptic responses to input respectively. If the model is to capture AMPAR function within the synapse, it is vital that changes in ions such as calcium should trigger changes to synaptic AMPAR composition and states which at least qualitatively resemble those seen experimentally.

In this preliminary investigation, a transient high amount of calcium injected into the system. Additionally, the population makeup of the initial synaptic AMPARs was also altered to assess whether this had any impact on subsequent LTP expression. The aim was to determine if the model agents were sufficient for LTP expression, and if not, how best to upgrade the model to accommodate likely candidates that are necessary.

Plasticity

LTP was induced through the delivery of calcium in a defined “burst” of 100Hz activity, triggered after 2500 s of simulation run time.

180 Calcium ions were representative of 3 μ M input. The duration of this LTP-inducing stimulus was altered to last 1, 2, 5 and 10 s in some simulation runs. To investigate how the capacity of the synaptic zone may influence receptor profiles under HFS conditions, two runs added 10, 25, 50 or 100 slots either at the time of, or 30 minutes post-induction.

AMPAR population composition

Five different population makeups were tested **Table 1**. There are 250 receptors in each version, with 50 PSD slots available. All receptors start outside the synaptic zone.

Composition Name	#GluA1/2s	#GluA2/3s
<i>All 1/2</i>	250	0
<i>All 2/3</i>	0	250
<i>Most 1/2</i>	200	50
<i>Most 2/3</i>	50	200
<i>Equal</i>	125	125

Table 1 AMPAR compositions in the model.

For the LTP length and slot introduction simulations, the *Most 1/2* composition was used.

Results

Synaptic AMPARs exhibit stable basal population numbers

In each composition variant we observe a rapid uptake of receptors into PSD slots, which stabilises to an average of 23 receptors across all five (**Figure 38**). It is only in populations which are primarily GluA1/2 AMPARs which exhibit any instability in this population, and a slower incorporation of receptors. In the *All 1/2* case, the average across the simulation is only 15 synaptic AMPARs.


Sustained HFS saturates GluA1/2 AMPAR insertion into the synapse

The overdriven potentiation of the model for 30 s drives an increase in synaptic receptor number in all GluA1-containing variants of the model, with the average number of receptors shifting from 23 to 37 (**Figure 38**). To compare, the *All 2/3* variant averages at 29 for the duration of the simulation. The *most 1/2* variant is the only one to show a stable potentiation for the full simulation run, with all other variants returning to pre-induction levels.


Short periods of HFS can trigger stable synaptic potentiation

The model continues to show activity-driven incorporation of GluA1/2 AMPARs, even with the new HFS protocol (**Figure 38**). There is an almost complete filling of all slots, with average synaptic receptor numbers across all protocols increasing from 18 to 38. If the 1 s HFS is excluded, which was the only one to fail to maintain a potentiated amount, the averages move to 19 and 42 respectively.

Basal



LTP



LTP length





Figure 38 Model responses to different AMPAR compositions in basal and LTP conditions

The basal model shows no major change in synaptic receptors between all five groups, although the All 1/2 population is more unstable over the course of the simulation. GluA1-containing AMPARs are driven through LTP induction into the synaptic zone, but are unstable in all but the most 1/2 case. HFS > 1 s were all sufficient drive stable increases in synaptic AMPARs. N=1 for each simulation, HFS at 2500s.

LTP Slots



LTP + Delayed Slots




Figure 39 Slot addition during LTP induction stabilises higher synaptic AMPAR populations only when added concurrently with the induction protocol.

When slots are added simultaneously to LTP induction at 2500s, most are filled by free AMPARs. Low numbers of new slots seem to destabilise the potentiation, with 50 slots being the most stable. When the addition of slots is delayed by 30 minutes, the system becomes unstable. LTP induction was for 30 s in all simulations. N=1 for each simulation.

Small increases in slot capacity can destabilise synaptic AMPAR populations

After new PSD slots were created at 2500 s in conjunction with the LTP stimulus, all variants showed increases in AMPAR numbers which exceeded those in the LTP-only protocol (**Figure 39**). The number of slots added strongly correlated with the synaptic AMPAR numbers present by the end of the simulation. The lower slot values (10 and 25) were not stable, and returned to basal levels by the end of their respective runs.

Adding slots 30 mins post-LTP fails to sustain enhanced synaptic receptor numbers

Adding slots 30 minutes after the LTP stimulus is designed to represent the evidence of CaMKII-mediated synaptic remodelling of the PSD occurring within 30 minutes after a strong stimulus (Bosch et al., 2014; Murakoshi, Wang, & Yasuda, 2011; Okamoto, Narayanan, Lee, Murata, & Hayashi, 2007). These runs act similarly upon addition of new slots to the simultaneous LTP and slots, with the magnitude in potentiation correlating positively with the number of slots added (**Figure 39**). However, there is no maintenance of this increased potentiation in any variant, with the 10, 50 and 100 slot variants returning to basal levels. Only the 25-slot variant retained receptor levels that were the same as post-LTP, pre-slot (39 receptors).

Discussion

The model successfully shows divergent responses between basal and potentiated scenarios. All the different AMPAR population compositions showed stable levels of synaptic AMPARs under basal conditions. This suggests the model does not experience major shifts in outputs due to the naturalistic calcium module added previously.

When driven by HFS, all population compositions show increases from baseline. The All1/2 variant however shows increased instability, failing to maintain the potentiation for > 1500 s. Similarly, the Equal, Most2/3 and All2/3 variants return to basal levels before the end of the simulation. That the Most1/2 variant is the most stable supports experimental evidence that AMPAR composition at hippocampal synapses is roughly 80% GluA1-containing AMPARs to 20% GluA3-containing AMPARs (W. Lu et al., 2009; Wenthold et al., 1996). Given the similarity to experimental evidence, this population is used for all future experiments.

Altering the length of the LTP-inducing stimulus indicated that only a small window of heightened calcium levels is necessary to trigger stable increases in

synaptic AMPAR numbers. Only the 1 s stimulus duration failed to show sustained potentiation, suggesting a lack of sustained recruitment of kinase activity in this run. Given that values past this did not seem to quickly saturate the system, only short (< 5s) LTP induction protocols were considered for future experiments.

AMPARs increase in number during LTP events, and may be due to the generation or uncovering of “slots” on PSD proteins that stabilise receptors within the synaptic zone. Concurrent addition of new slots with HFS showed rapid saturation of these new slots in all instances. Only large slot increases were stable however, with the +10 and +25 runs falling to basal levels. These runs failed because of PP1 recovering from being saturated by CaMKII and PKA activity, which quickly destabilises the synaptic AMPAR population through dephosphorylation at S845. These findings are surprising in part because there is no distinction between synaptic and extra-synaptic AMPARs in the model. All kinases and phosphatases can access the receptors regardless of their location.

What may be occurring is that the number of S845 phosphorylated receptors are rapidly detected and inserted by the trafficking system, which offers protection from being dephosphorylated. They are additionally protected from ubiquitination, and can engage in glutamate binding events. All these possible events, coupled with the autophosphorylation of CaMKII, prevent PP1 from being able to act quickly enough to return the system to basal conditions. In this manner, the model agrees with the findings of Hayer and Bhalla (2005), in that overwhelming of the PP1 system, opposed to any increase in direct inactivation via I1, is sufficient to switch the synapse to a “high” state.

Unlike Hayer and Bhalla however, my model shows these switches with much shorter increases in calcium (and indirectly, cAMP), which is much more representative of calcium dynamics at the synapse. To explore this relationship further, I decoupled the addition of slots from the HFS. Returning to the slot hypothesis, synaptic remodelling takes some minutes to be observed post-LTP induction (Bosch et al., 2014; Murakoshi et al., 2011; Okamoto et al., 2007). Whilst this does not rule out slots emerging prior to this remodelling, this simulation assumed that slots were present only after this was complete.

In this scenario the addition of the slots destabilised all simulations. Prior to slot addition, all runs show a stable almost saturated increase in synaptic AMPAR numbers. Like in the concurrent scenario, the larger slot increases generate the largest change in AMPAR number. Both rapidly decay however, with no significant difference between runs by the end of the simulation. In all these runs, free PP1 levels recover after the slots are added into the model.

Whilst investigating this decay, I noticed the rate of ubiquitination in the model caused receptor loss at levels much higher than desired based on HEK-cell derived data, where GluA1-containing AMPARs were reduced by 40 % after 4 hours (A. Lin et al., 2011). In this version, the rate stands at a 40 % loss after 0.4 hours, suggesting the current kinetics of the model are 10 times too high. This may be interfering with interpretation of the maintenance phase of these outputs, as the readily available pool of non-synaptic AMPARs shrinks rapidly and therefore makes any AMPARs trafficked out of the synaptic zone more susceptible to being degraded fully and removed from the system. The model's ubiquitination rate was therefore reduced to a tenth of the original values in subsequent experiments.

Summary

- The model shows stable baseline responses and potentiation of the synapse under a strong calcium input.
- Different AMPAR compositions show different and expected responses to the LTP induction protocol.
- The model shows a sensitivity to strong calcium inputs which extend past 1s that saturates quickly.
- Altering the synaptic slot number during LTP induction destabilises the heightened response in low (< 25) additions.
- Adding slots 30 minutes after the LTP induction protocol destabilises the potentiated response when both small and large numbers of slots are added.

Experiment #5 CaMKII Parameter Set Testing

Introduction

The initial model tried where possible to maintain the parameter sets from AMPS and the imported models, to maintain a consistency and to identify differences that may have been due to the model language and implementation. For some components of the model however, there has been a much greater deal of both experimental and mathematical work conducted to both identify and refine kinetic rate parameters.

One of the largest bodies of work is for CaMKII, and encompasses subunit phosphorylation, autophosphorylation, actin binding and AMPAR interactions. The data in the first model instantiation originates from literature analysis conducted prior to 2000, and subsequently used in several models hosted in DOQCS (Sivakumaran et al., 2003).

Since the publication of these models, there has been several publications that have estimated CaMKII interactions on a much finer grained scale, making use of more recent experimental findings. Some of these parameters show significant differences to those used in the DOQCS models, which may have an impact on model outputs.

Given that veridical parameters are difficult to acquire, and sensitive to a wide number of methodological assumptions and environmental factors, ignoring these new estimates risks the model predicting hypotheses that may be both biochemically and physiologically implausible. To assess the extent to which these changes may impact the model, I altered the rate of CaMKII-PP1 binding and dephosphorylation to those in Coomber et al 1998, and CaMKII-CaM binding, to match those found in Stefan et al.(2012).

Parallel to this investigation, I also assessed the impact of assuming steric hinderance in the model between PP1 and CaM binding to CaMKII. First postulated in (Pharris et al., 2019), they hypothesise that the binding of PP1 blocks the CaM binding to the CaMKII regulatory domain, which is close to Thr286 (Hanson, Meyer, Stryer, & Schulman, 1994; Waxham, Tsai, & Putkey, 1998).

One aim of the model is to accurately describe AMPAR trafficking events, and the impact the kinase and phosphatase balance has. To get a better sense as to how these different binding rates and steric hindrance assumptions impacted the model, runs were tested with a series of high and low frequency patterns of calcium stimulation.

Methodology Changes

The rules changes concerned only CaMKII interactions with its major activation and dephosphorylation partners, CaM and PP1 respectively. The changes are summarised below in **Table 2**.

<i>Variable</i>	AMPS Values	Stefan <i>et al</i> (2012) & Coomber <i>et al</i> (1998) Values
<i>CaM Binding</i>	7.47 molecule ⁻¹ s ⁻¹	0.0697 molecule ⁻¹ s ⁻¹
<i>CaM Unbinding</i>	4.5 s ⁻¹	0.126 s ⁻¹
<i>PP1 Binding</i>	0.005696 molecule ⁻¹ s ⁻¹	0.008303 molecule ⁻¹ s ⁻¹
<i>PP1 Unbinding</i>	1.4 s ⁻¹	50 s ⁻¹
<i>CaMKII dephosphorylation</i>	0.35 s ⁻¹	1 s ⁻¹

Table 2 Variable testing sets for CaMKII interactions

Both sets of values were simulated with the same conditions across all three trafficking hypotheses and compared. Additionally, in the steric variants, PP1 may only dephosphorylate the Thr286 site when CaM is explicitly not bound.

The number of AMPARs present in the previous models was felt to be too high for the volume of the spine being modelled (David Wyllie, personal communication). The numbers of both AMPAR compositions were therefore halved from their values in Experiment #4, split evenly between synaptic and extrasynaptic “compartments” (which are just state flags on the receptors). The number of subunits to allow for very gradual increases in these populations over the simulation run remained unchanged.

The maximum number of slots available in the model was increased from 50 to 75. This better reflects the values calculated for the capacity of a synapse to possess 75 AMPARs at any one time (Nair *et al.*, 2013).

Induction protocols

Tetanic Protocols

The testing here uses a series of tetanic-like protocols of LTP induction, which are not physiologically impossible but are not (except for compressed and spaced theta burst trains) used as induction protocols in the literature. I have kept them here because they show the system being driven by heightened patterns of activity, still show a variation in plasticity induction. They also help demonstrate the difference in response by these two sets of parameters much more clearly than the correct protocols (see Experiment #6). Note that the names of these protocols match their intended patterning, which is explained in Experiment #6.

HFS – 20 s 100 Hz tetanus

Theta – 5 s 100 Hz tetanus

Delta – 10 s 100 Hz tetanus

weakLTP – 20 trains of 10 bursts (a burst defined as four 100 Hz Calcium pulses). Trains were spaced 5 s apart.

Compressed Theta Burst Train (cTheta)– A set of three trains (25 bursts with a 200 ms gap between bursts), with a 10 s gap between trains (P. Park et al., 2016).

Spaced Theta Burst Train (sTheta) – A set of three trains, as detailed above, but with a 600 s gap between trains (P. Park et al., 2016).

These runs were also used to assess the use of two different parameter sets for CaMKII interactions with CaM and PP1.

Results

Both sets of values without steric hinderance allow for synaptic potentiation

Simulation results show that in each parameter set there is an increase in the number of synaptic AMPARs present above basal levels under LTP induction protocols. The original AMPS values show a much greater stability of this increase across all protocols, with only cTheta showing a noticeable failure to maintain beyond 1000 s (**Figure 40, 41**). This is seen to correlate with a loss of CaMKII phosphorylation, which allows for PP1-mediated dephosphorylation of GluA1 residues and subsequent loss of synaptic AMPARs. sTheta, whilst failing to induce any sustained level of CaMKII phosphorylation across the simulation runtime, does show a small additive effect of these spaced trains in reaching a level of synaptic AMPARs similar to the other protocols.

The Stefan and Coomber set has less diversity in the different protocol responses, especially by the end of the simulation runtimes, with GluA1/2 AMPARs returning to or just below the initial numbers (**Figure 44**). Kinase activity is much more tightly regulated, with S831 and S845 phosphorylation unsustainable post-protocol (**Figure 45**).

By the end of the simulation, no significant difference in synaptic receptors was found between the groups for AMPS ($F(5,24) = 1.257, p = 0.315$) or the Stefan and Coomber values ($F(5,22) = 0.965, p = 0.460$).

LFS responses diverge between the two sets of values

The Stefan and Coomber values show a much greater increase in the level of phosphorylated CaMKII (and subsequent S831 phosphorylation) than AMPS

during the three LTD-associated induction protocols (**Figure 46, Figure 47**). In particular, the response to paired-pulse stimulation is noticeably closer to a weak potentiation of the CaMKII associated effects on AMPARs. This response fails however to result in any wider change to the system response, which exhibits a transient LTD that closely matches the AMPS data (**Figure 41, Figure 43**).

The simple LFS protocol exhibits a steady rundown of synaptic AMPARs, reducing the number by over half in the AMPS ($n = 5$, mean = 20.20, SD = 6.76, 95% CI [11.80, 28.59]) and Stefan and Coomber ($n = 5$, mean = 23.20, SD = 5.44, 95% CI [16.43, 29.97]) variants. There was a significant difference between the LFS and PP groups in both AMPS ($F(2,11) = 13.867$, $p = 0.001$) and Stefan and Coomber ($F(2,12) = 10.226$, $p = 0.003$) variants.

AMPS values show unstable CaMKII-mediated saturation of PPI activity

Excluding the cTheta and sTheta protocols, CaMKII quickly became fully phosphorylated in the AMPS LTP model (**Figure 40**). The decay for the tetanic protocols takes much of the simulation runtime, which in turn maintains a saturation of PP1 in the system. This saturation ensures both S831 and S845 residues remain highly phosphorylated until activated CaMKII levels fall sufficiently for free PP1 levels to recover (**Figure 41**).

Steric hinderance does not alter the broad response profiles of both parameter sets

The incorporation of steric hindrance had no significant impact on the outputs of the model, showing only a slight increase in improvement in CaMKII switching in the AMPS set (**Figure 40, 42**), and no effect on the Stefan and Coomber values (**Figure 44, 46**). This property was therefore incorporated into all future versions of the model.

By the end of the LTP protocol simulations, no significant difference in synaptic receptors was found between the groups for the Stefan and Coomber values ($F(5,18) = 1.773$, $p = 0.169$). There was however a significant difference in the AMPS data ($F(5,23) = 3.223$, $p = 0.024$), which was between the HFS and theta protocols ($p = 0.017$, determined by Tukey HSD post hoc analysis).

The LTD induction protocols show the same behaviour, with a significant difference between the LFS and PP groups in both AMPS ($F(2,12) = 18.092$, $p = 0.000239$) and Stefan and Coomber ($F(2,12) = 22.793$, $p = 0.000082$) variants.

The simple LFS protocol exhibits a steady rundown of synaptic AMPARs, reducing the number by over half in the AMPS ($n = 5$, M = 24.60, SD = 4.03, 95% CI

[19.58, 29.61]) and Stefan and Coomber (n = 5, M = 20.60, SD = 3.64, 95% CI [16.07, 25.12]) variants.







- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP

Figure 40 Original (AMPS) values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTP)

All protocols in both variants show activity driven transient increases in synaptic GluA1/2 AMPARs, but also a complete loss of GluA2/3 AMPARs.

The introduction of steric hinderance makes no change to the broad progress of the simulations, but does indicate some ability for HFS to maintain enhanced levels of synaptic GluA1/2 AMPARs.

n=5 for each protocol except steric hinderance sTheta (n=4).





- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP



No Steric Hinderance

Steric Hinderance

Receptor Phosphorylation
S831



S845



Free PP1





Figure 41 Original (AMPS) values with and without steric hinderance: AMPAR phosphorylation and PP1. (LTP)

Both variants show similar profiles, although HFS produces a larger level of receptor phosphorylation and PP1 saturation in the steric hinderance variant.


n=5 for each protocol except steric hinderance sTheta (n=4).

— LFS
 — LFSPP200
 — LFSPP50

Figure 42 Original (AMPS) values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTD)

LFS shows a gradual decay in AMPARs across the simulation run, whereas PP-LFS shows small to no change in GluA1/2 AMPARs, seemingly stabilising them in the membrane. GluA2/3 AMPARs show the same complete loss as in LTP. Steric hinderance values show little change from the non-steric hinderance results

n=5 for all protocols, except PP-LFS50 without steric hinderance (n=4).









- LFS
- LFSP200
- LFSP50

Figure 43 Original (AMPS) values with and without steric hinderance: AMPAR phosphorylation and PP1. (LTD)

Almost no AMPAR phosphorylation exists in all protocols across both variants. Only the PP-LFS protocols can reduce the availability of PP1 in the system.


n=5 for all protocols, except PP-LFS50 without steric hinderance (n=4).



- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP

Figure 44 Stefan and Coomber values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTP)

All protocols show similar but less diverse increases in synaptic GluA1/2 AMPARs compared to AMPs, but shows identical loss of GluA2/3 AMPARs. Enzymatic dynamics are much more tightly regulated with these values, showing much more transient CaMKII activity. Except for sTheta, steric hinderance only shows minor effects on the maintenance of elevated synaptic GluA1/2 AMPARs. Non-steric: n=5 except delta and theta (n=4). Steric: n=5 except for theta (n=2), stheta (n=3), and ctheta (n=4).









- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP

Figure 45 Stefan and Coomber values with and without steric hinderance: AMPAR Phosphorylation and PP1. (LTP)

Both the non-steric and steric variants behave similarly in AMPAR phosphorylation and PP1 saturation, with only Steric sTheta exhibiting a slight difference due to increased membrane AMPAR numbers.

Non-steric: n=5 except delta and theta (n=4).
Steric: n=5 except for theta (n=2), stheta (n=3), and ctheta (n=4).



— LFS
 — LFSPP200
 — LFSPP50


Figure 46 Stefan and Coomber values for CaM, PP1 and CaMKII interactions with and without steric hinderance: AMPARs and CaMKII. (LTD)

Like the AMPS values seen in Figure 16, LFS is the only one to show sustained loss of AMPARs over the simulation. PP-LFS shows a slight increase in synaptic GluA1/2s and a greater, stepped activation of CaMKII subunits, which is rapidly lost post-protocol.


No major change is seen between the steric and non-steric simulations, although the activation of CaMKII subunits is enhanced in the PP-LFS protocols. $n=5$ for all protocols.

No Steric Hinderance


Receptor Population
GluA1/2






GluA2/3



CaMKII phosphorylation





Steric Hinderance



— LFS
 — LFSPP200
 — LFSPP50

Figure 47 Stefan and Coomber values with and without steric hinderance: AMPAR phosphorylation and PP1. (LTD)

Unlike in the original AMPS model variants, PP-LFS protocols here show a considerable amount of phosphorylation of AMPARs at both S831 and S845 residues. Steric variants show a slightly larger level of S831 phosphorylation and slightly greater reduction in PP1 availability. n=5 for all protocols.



Discussion

The Stefan and Coomber dataset is arguably a much more conservative model of CaMKII activation under LTP inducing protocols, with synaptic AMPAR numbers showing a smaller and shorter potentiation compared to the AMPS dataset. Even the strongest protocols failed to saturate the number of phosphorylated CaMKII subunits, only achieving 50% activation. This decayed much more rapidly than in AMPS for all induction protocols, except for sTheta.

Despite this, PP-LFS triggered higher levels of active CaMKII subunits, with over ten times as many activated subunits as AMPS. This does not alter the GluA1/2 AMPAR profile across the simulation, which shows only transient reductions in this AMPAR population, compared to the gradual rundown seen in LFS.

Neither dataset shows any capture of GluA2/3 AMPARs in the synaptic zone. This may be because of the stimulation protocols destabilising the synapse in the case of LTP. Activity-driven GluA1/2 AMPARs outcompete the small GluA2/3 AMPAR population for both synaptic space and trafficking machinery, making them less likely to be maintained synaptically. However, their continued loss within the LFS protocols suggests further investigation as to why these receptors are unable to maintain a stable population within the synaptic zone.

The outcomes of the steric hindrance simulations suggests that the incorporation of this hypothesis do not materially change the broad outcomes of different induction events. For the Stefan and Coomber parameter set, the effects are largely in a greater stability of the sTheta induction protocol, sustaining a potentiated synapse for the duration of the simulation.

For the AMPS parameter values, the effect of steric hinderance is to widen the variance in the levels of CaMKII activation, PP1 saturation and GluA1/2 AMPAR incorporation between the LTP induction protocols. This translates into decreased stability of synaptic GluA1/2 AMPARs, and differences in the active CaMKII subunit decay profile, which is now distinct for each LTP protocol.

The duration of CaMKII phosphorylation has been shown to have very different results experimentally, ranging from seconds through to days and possibly much longer (Singh & Bhalla, 2018). Much of the recent evidence presents a case where this is modulated by localisation of both CaMKII and PP1 into effectively separate PSD and cytosol fractions.

The constraints of Kappa make exploring this within this model at this time impractical, but do emphasise that the system as whole needs to balance switching from low to high states with the spontaneous activity and ion fluctuations which

may insert too much noise into any information transfer if this sensitivity is set too high. AMPS values have a rapid accumulation of phosphorylated CaMKII subunits, and can switch into a high state. The Stefan and Coomber values show no propensity towards switching even under high calcium level protocols. They do however show a marked increase in sensitivity to LFS compared to AMPS values, gradually ramping up for the duration of the LFS protocol.

Given that the Stefan and Coomber values are derived at least in part from more recent experimental evidence, and still show broadly the expected response profiles in high and low frequency stimulation states, this parameter set was used in all future simulations found in this chapter.

Summary

- AMPS and Stefan and Coomber values show similar model outputs regarding synaptic AMPAR expression during LTP and LTD induction protocols.
- AMPS values show a much greater likelihood of a CaMKII switch even in sub-optimal induction contexts.
- Stefan and Coomber values show a divergent synaptic response to paired-pulse LFS protocols, suggesting a failure to capture LTD.

Experiment #6 Modelling AMPAR Trafficking during Plasticity Events

The past five preliminary explorations with the model emphasise that the movement of AMPARs is controlled by several different systems whose interactions can greatly change the end result. This is further compounded by the number of different hypotheses of AMPAR trafficking that currently exist. To gain a clear understanding of how these different mechanisms may interact, I ran the model under several different trafficking paradigms. These can be split into the three major hypotheses detailed in Hugarir & Nicoll (2013): Slot-centric, Receptor-centric and Exocytosis-centric. The final 4 experiments of this Chapter constitute the investigation in how much the model can tell us about the extent to which these different hypotheses can explain the current understanding of how AMPARs are trafficked during plasticity inducing events.

Slot-centric




Figure 48 The Slot-centric model variant.

AMPA entry into the synaptic zone is limited by the number of free slots (or stable nanodomains). Slots are promiscuous to any glutamatergic receptor, and are stabilised by CaMKII activity. The mechanism by which this works is not explicitly modelled due to a lack of definitive information.

This approach assumes that the number of receptors within the synaptic zone is limited by binding sites in the PSD. Increasing the kinase activity makes available or creates more of these “slots” for receptors to bind to, with phosphatase activity achieving the opposite (**Figure 48**). I modelled this by assuming that there is a positive correlation between CaMKII activation and available slots. For every 100 active CaMKII subunits, the number of available slots in the model increases by 10, with a limit of 60. These slots are unstable until the level of active CaMKII subunits reaches 500. Slot instability is reinstated when active CaMKII subunit

numbers begin to fall (< 100 subunits). At this point an unoccupied slot may decay from the system before being occupied again.

This is an assumption based upon the idea that active CaMKII phosphorylates the C-tail of Stargazin, a TARP strongly associated with AMPAR trafficking (Tomita et al., 2005). LTP events might recruit AMPAR-Stargazin through this phosphorylation release the C-tail from being membrane bound and facilitating its binding to PSD-95 (Opazo et al., 2012; Sumioka, Yan, & Tomita, 2010). The “slot” therefore is not necessarily the uncovering of a new binding location, but rather a relational property between diffusing AMPAR complexes and PSD scaffolding proteins, which is modulated through phosphorylation (Earnshaw & Bressloff, 2006; S. H. Shi et al., 2001).

Receptor Phosphorylation-centric

I focus on three different residues in this iteration: S818, S831 and S845. The assumption of this model is that phosphorylation at one or more of these sites is necessary for insertion of AMPARs into the synaptic zone (**Figure 49**). Unlike the PSD-centric model, there is no limit to the number of receptors that may enter (the limit is the overall receptor population in the model). There were five versions of this model made, to assess different relationships between phosphorylation and trafficking. The first three assume that either S818, S831 or S845 phosphorylation respectively were necessary and sufficient for receptor trafficking into the synaptic zone. The phosphorylation of the other residues had no effect on trafficking. Receptor removal only occurred if the residue was unphosphorylated.




Figure 49 The Receptor Phosphorylation model variant.

AMPA entry is controlled by residue modifications, which here correspond to serine residues on the GluA1 subunit. AMPAR trafficking in this variant is sensitive to changes in the balance of kinase and phosphatase activity.




Figure 50 The Exocytosis model variant.

The variant assumes extrasynaptic receptors are a minor transient population of AMPARs in basal conditions. AMPARs are transported to perisynaptic sites by CaMKII and PKC-mediated mechanisms.

For the final two versions, all three residues were important for trafficking. In the *all_sufficient* model, any single phosphorylation event would allow for membrane insertion. The addition of other phosphate groups would increase the rate of entry. All 3 residues need to be dephosphorylated for the removal of the AMPAR from the synaptic zone. The *all_necessary* model assumes all three residues must be phosphorylated for receptor insertion into the synaptic zone. Dephosphorylation of any single residue will result in a chance of removal from the synaptic zone, scaling with the number of phosphorylated residues remaining.

Exocytosis-centric

This variant assumes that AMPARs are primarily either in endocytic vesicle pools or at the synaptic zone (**Figure 50**). Endocytic AMPARs may be trafficked to the membrane if they are phosphorylated by CaMKII and bind to NSF, a molecule related to vesicle trafficking in dendritic spines. These may then enter the synaptic zone if they are further phosphorylated at S845 by PKA. S818 phosphorylation increases the speed of exocytosis. In this variant, we assume that this endo- and exocytosis occurs within the spine volume, with a reserve pool being maintained within the spine.

Plasticity Protocols

Multiple different mechanisms of LTP have been discovered that appear to require different molecular networks in spines and dendrites. As we are not fully modelling the spine biochemistry, there may be plasticity protocols which will not function as they do *ex silico*. To explore this and the three trafficking hypotheses, I have used nine different protocols here, for both LTP (six) and LTD (three).

LTP Protocols

Single Theta Burst Train – Four 100 Hz Calcium pulses (hence a “burst”) are delivered 25 times with a 200 ms gap between bursts. This constitutes a single “train” (based on P. Park et al., 2016)

Compressed Theta Burst Train (cTheta) – A set of three trains, as detailed above, but with a 10 s gap between trains (P. Park et al., 2016).

Spaced Theta Burst Train (sTheta) – A set of three trains, as detailed above, but with a 600 s gap between trains (P. Park et al., 2016).

Delta Burst Train – A set of three trains, as detailed above, but with a 10 s gap between trains and the inter-burst interval increased to 500 ms (Grover et al., 2009).

weakLTP Burst Train – A set of three trains, as detailed above, but with a 10 s gap between trains and the inter-burst interval increased to 5000 ms (Grover et al., 2009).

High Frequency Stimulation (HFS) – Calcium is delivered at 100 Hz for 4 s (T. V. P. Bliss & Lømo, 1973).

LTD Protocols

Low Frequency Stimulation (LFS) – A 1 Hz pulse of Calcium was delivered for the first 900 s of the simulation (Gonzalez, Morales, Villarreal, & Derrick, 2014; Kemp et al., 2000).

50 ms Paired Pulse (PP) Stimulation (LFSP50) – Two pulses of Calcium are delivered per second, with an intra-pulse gap of 50 ms (Gonzalez et al., 2014; Kemp et al., 2000).

200 ms Paired Pulse Stimulation (LFSP200) – Two pulses of Calcium are delivered per second, with an intra-pulse gap of 200 ms (Gonzalez et al., 2014; Kemp et al., 2000).

In each instance, a single pulse of calcium added 60 ($\sim 1 \mu\text{M}$). This is derived from work suggesting an influx of $10 \mu\text{M}$ into spines during signal propagation, which is reduced by 90 % via sequestering mechanisms (Higley & Sabatini, 2012). This gives close to $1 \mu\text{M}$ registered as “free” in the model.

Model Changes

The rate of endocytosis was found to be too high for the current version of the model. In a previous iteration, the variable for ‘ampa_endo’ in the receptor phosphorylation variant of the model had been adjusted from a rate of 3 s^{-1} to a rate of 40 s^{-1} . No evidence as to this being a deliberate change has been located,

and it has been determined to be because of an error in variable editing which went undetected. As no well justified reason for this change could be found, both 'ampa_endo' and 'ampa_exo' were changed to their Gallimore (2016) values (0.03 s^{-1} and 0.045 s^{-1} respectively). This better standardises this part of the model, by making all the kinetic parameters derived from the same source.

I further determined that a previous decrease in sensitivity of cAMP activation (which is a simplification of actual Calcium-mediated cAMP activation networks) was too conservative and was preventing any input-mediated shifts in system state. The model modification was therefore adjusted down to a previously used lower threshold (from $> 15 \text{ Ca ions}$ to > 10). This is still above the average basal Calcium levels but allows for cAMP activation to occur sufficiently to trigger PKA activation under experimentally derived induction protocols.

Results

The Slot variant allows for a transient increase of synaptic receptors

The cTheta, Delta, and HFS protocols all trigger a rapid incorporation of both GluA1/2 and GluA2/3 AMPARs, with the peak levels occurring within the first 500 s of the simulation (**Figure 51**). These all however decay throughout the remainder, returning to (or falling just below) the initial basal level. The remaining three LTP protocols show only small, rapidly decaying increases in GluA1/2 AMPARs, or in the case of weakLTP, just a gradual rundown of synaptic receptors.

By the end of the simulation, a significant difference in synaptic receptors was found between the groups ($F(5,24) = 22.483$, $p = 2.4455\text{E-}8$). Significance ($p < 0.05$), determined by Tukey HSD post hoc analysis, was between the cTheta, Delta, and HFS runs and the sTheta, Theta, and weakLTP protocols.

An additive sufficient Receptor Phosphorylation variant allows for a slight increase in synaptic receptors

S831-only, S845-only, S818-only, and the *all_necessary* model variants showed no ability to traffic AMPARs into the synaptic zone (data not shown). *All_sufficient* was the only version to show any insertion of AMPARs in response to induction protocols (**Figure 52**).

Both GluA1/2 and GluA2/3 synaptic populations are relatively stable across all LTP protocols. Whilst initial increases in GluA1/2 numbers are related to the strength of the underlying induction, by the end of the simulation, no significant

difference in synaptic receptors was found between the groups ($F(5,24) = 1.1$, $p = 0.386$). GluA2/3 receptor numbers are extremely similar in all protocols, averaging between 3 – 4 across all six.

Exocytosis is insufficient for potentiation of the synapse

Whilst some activity-driven trafficking is seen immediately following the induction protocols, all synaptic receptors are driven out by the end of every one (**Figure 53**). GluA2/3 receptor incorporation is seen to gradually increase across the course of the simulation for each protocol.

By the end of the simulation, a significant difference in synaptic receptors was found between the groups ($F(5,24) = 6.458$, $p = 0.001$). This was between HFS and all other protocols ($p < 0.05$), except for Delta ($p = 0.242$, determined by Tukey HSD post hoc analysis).

All variants show defects in LTD induction

Responses to LFS in all three protocols once again show only some LTD-like responses (**Figure 54**). PP-LFS always initiates a gradual potentiation of the synapse which remain above basal levels by the end of the simulation. In the receptor phosphorylation variant, a very large potentiation is also observed in the LFS protocol. The Slot and Exocytosis variants otherwise shows a gradual decay in synaptic receptor number.

There are significant differences between PP-LFS and LFS in the Exocytosis ($F(2,12) = 29.621$, $p = 0.000023$) and Slot ($F(2,12) = 99.028$, $p = 3.476E-8$) variants, but no difference in the Receptor Phosphorylation ($F(2,12) = 1.835$, $p = 0.202$).


- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP

Figure 51 Slot model variant shows activity-driven insertion of AMPARs in the synapse


Both GluA1/2 and GluA2/3 AMPARs show activity dependent insertion into the synaptic zone. No protocol is able to sustain stable populations of either AMPAR type for the duration of the simulation, but cTheta, HFS and Delta were able to show large transient increases in synaptic receptors. $n=5$ for all protocols.

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation




Free PP1




Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation



Free PP1



- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP

Figure 52 Receptor Phosphorylation model variant shows modest stable potentiation of AMPARs in the synapse.

Only GluA1/2 AMPARs are affected by the LTP induction protocols, showing small increases in all protocols. All protocols remain above the basal level by the end of the simulation. n=5 for all protocols.


- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP

Figure 53 Exocytosis model variant fails to maintain initial AMPAR numbers.


Whilst some small transient increases in GluA1/2 AMPARs are seen, all protocols show no sustained AMPAR population by the end of the simulation. GluA2/3 AMPARs are slow to enter the synaptic zone in any protocol. $n=5$ for all protocols.

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation



Free PP1




- LFS
 - LFSPP200
 - LFSPP50


Figure 54 LFS responses differ from expected LTD in all variants.

All variants show strong transient potentiation in the PP-LFS protocols. The Receptor variant also shows potentiation in the LFS protocol. n=5 for all protocols.


Slot
GluA1/2




GluA2/3




Receptor Phosphorylation
GluA1/2




GluA2/3



Exocytosis
GluA1/2



GluA2/3



Discussion

This attempt to dissociate the different trafficking hypotheses into three variants suggests no variant alone is sufficient to stabilise synaptic AMPARs at increased levels after an LTP-inducing stimulus. The most promising variant here is Slot, which exhibits both expected divergent response profiles depending on induction protocol strength, but also at least transient potentiation through additional GluA1/2 AMPARs incorporated post-induction. The Receptor Phosphorylation variant, whilst less potentiated, and with less variation between protocols, is more stable than the other variants for both AMPAR types.

Exocytosis, whilst still unable to show any real stability to activity-driven AMPARs in the synaptic zone, does show the ability to gradually introduce GluA2/3 AMPARs. The small influx of activity-driven AMPARs also points again to this variant as not being sufficient to explain AMPAR trafficking in LTP conditions. It does not discount the ability of this method of injecting a small but significant number of AMPARs, perhaps GluA1 homomers, quickly enough to induce long-lasting changes to PSD morphology and molecular makeup.

The observed failure in PP-LFS to induce LTD suggests that adding metabotropic signalling in the model may be required for LTD expression at the model synapse. mGluR signalling has strong evidence to support it being necessary for successful hippocampal and cerebellar LTD. PP-LFS has repeatedly been shown to activate mGluR1 and mGluR5 receptors in CA1 and induce LTD (Huber, Kayser, & Bear, 2000; H. H. Kim et al., 2015; H. H. Kim, Park, Lee, & Ho, 2019; Reyes-Harde & Stanton, 1998). The mechanism by which mGluRs act to trigger AMPAR endocytosis appears to differ between cerebellar and hippocampal neurons, requiring PKC-mediated phosphorylation and tyrosine residue dephosphorylation, respectively (Lüscher & Huber, 2010).

NMDAR-dependent LTD has long thought to be mediated through the flow of low levels of calcium into the dendritic spine through NMDARs (Artola & Singer, 1993; Kemp & Bashir, 2001). The model, which lacks explicit NMDARs, runs on the evidence that the pattern of low calcium activity preferentially activates PP1 over CaMKII (Huang, Liang, & Hsu, 2001), allowing for dephosphorylation of synaptic AMPARs and their vulnerability to subsequent removal from the synapse (Tomita et al., 2005).

A small but growing body of evidence however supports a metabotropic action of NMDARs on triggering LTD (Dore, Aow, & Malinow, 2016; Mayford, Wang, Kandel, & O'Dell, 1995; Scanziani, Malenka, & Nicoll, 1996). Blockade of the

NMDA channel paired with repeated HFS triggers LTD in hippocampal slices (Nabavi et al., 2013).

The repeated failure of the model to induce LTD through LFS suggests that any investigation of this phenomenon will require the addition of metabotropic signalling through both mGluR and NMDARs. CaMKII and PP1, the key interacting partners for LTP induction, are both localised close to NMDARs. Indeed, the PP1-PSD population appears to be unable to dephosphorylate CaMKII bound to NMDARs until structural conformational changes occur as a result of glutamate binding, allowing CaMKII to reach saturated phosphorylation levels easier (Aow, Dore, & Malinow, 2015; Dore et al., 2016; J. Lisman & Raghavachari, 2015). This all points to the model requiring a major reappraisal to faithfully capture LTD. Because of this, subsequent model runs used LFS protocols more as a diagnostic of the model's function, rather than capturing any physiological behaviour.

Summary

- Both Receptor Phosphorylation and Slot variants show potentiation and the ability to maintain a synaptic population of AMPARs above basal levels.
- Exocytosis, whilst not sufficient for potentiation, does exhibit activity-dependent trafficking which may be important in driving sufficient AMPARs to trigger a synaptic switch into a "high" state.
- LFS protocols in this scenario show unphysiological behaviour in all hypotheses and LFS types, supporting non-ion flux hypotheses of NMDAR-dependent LTD, and the requirement for more metabotropic signalling to be incorporated in future model iterations.

Experiment #7 A Merged Model of AMPAR Trafficking

These findings suggest that two hypotheses, Slot and Receptor phosphorylation, are likely the drivers of early trafficking in a basally active spine. Of the two, Slot has the highest capacity for synaptic zone entry, but if slots are assumed unstable this approach shows a gradual persistent downscaling of AMPAR numbers. Receptor Phosphorylation has a much more modest increase in AMPAR numbers, but these are persistent for up to an hour before a slow return to baseline. The latter also sees a stable maintenance of GluA2/3 AMPARs throughout the simulation and has only a modest fall below the baseline start values by the end of the simulation run.

From these findings I constructed a merged model which still assumes that all three hypotheses play a role in AMPAR trafficking during plasticity.

The slot hypothesis I will unpack here into two separate mechanisms. One is the immediate stabilisation of synaptic zone nanodomains triggered by a fast-acting molecule already present in the PSD. This means AMPARs remain in the synaptic zone for longer, increasing the number present. However, there is not more space available, but the rate of “decay” of a nanodomain is reduced. If activity is persistent however, the synaptic zone will increase in size, leading to an increasing in the number of nanodomains. These again can be made more stable through some unknown molecular stabilisation mechanism. This second mechanism presumably takes some time to manifest.

I assume CaMKII is responsible for both mechanisms, but due to two distinct populations. CaMKII kinetics appear to be modified in different spine locations, with PSD-CaMKII possessing robust stability and cytosolic-CaMKII a much quicker activity decay rate (Singh & Bhalla, 2018). Taken together, PSD-CaMKII ensures initial nanodomain stability, which is eventually supported by actin remodelling mediated by transient CaMKII activation in the cytosol (Bosch et al., 2014; Okamoto et al., 2007). The balancing act here is the relationship between CaMKII autophosphorylation and its tolerance of PP1 activity. A low ratio of PP1:CaMKII will ensure much more stable populations of activated CaMKII holoenzymes, with instability increasing as this ratio increases. Destabilisation of a potentiated synapse can be triggered therefore by increasing the level of PP1. In the cytosol, this will prevent CaMKII-actin disassembly, which whilst stabilising the structure of the spine, will prevent any movement of CaMKII into PSD zones and limit the number of free active CaMKII holoenzymes. In the PSD, this increase might increase the probability of destabilising any bistable switch, which would lead to nanodomain collapse and thus a loss of synaptic AMPARs.

This PP1 increase does not have to be due to changes in the PP1 copy number. I1 in its phosphorylated form is an effective sequester of PP1 (Hayer & Bhalla, 2005). This phosphorylation, mediated by PKA, is not particularly high post-induction. However, strong induction protocols that allow for sufficient numbers of S831 and S845 phosphorylated AMPARs to enter the synaptic zone, and a coupled high level of activated PKA and CaMKII, allow for a sustained higher level of phosphorylated I1. Targeted reductions in PKA activity, perhaps through restricting cAMP synthesis or increasing levels of PKI, could cause the release of inactivated PP1 and tip the balance in favour of synapse destabilisation.

The receptor phosphorylation variant findings tease out the importance of phosphorylation for the maintenance of GluA1/2 AMPARs in the synaptic zone. GluA2/3 AMPARs may persist long term in the synapse simply because they are not internalised and/or broken down as readily as GluA1-containing AMPARs (Widagdo et al., 2015). After plasticity inducing stimuli, they passively replace activity dependent AMPARs through simple binding to the promiscuous proteins that generate a stable nanodomain (Tanaka & Hirano, 2012). Their numbers are lower overall (but see (Kessels, Kopec, Klein, & Malinow, 2009)) and so can be outcompeted when necessary, but they otherwise maintain a stable level of minimal AMPAR function at the synapse (Diering & Huganir, 2018; Pandya et al., 2017). Their half-life from recent work appears to be considerably longer than active GluA1-AMPAARs, at 3.5 days to GluA1's 0.5 days (Dörrbaum, Kochen, Langer, & Schuman, 2018; but see Mathieson et al., 2018), which may reflect the lack of interaction with plasticity mechanisms. Their much reduced conductance under basal conditions (Oh & Derkach, 2005; Renner et al., 2017) protects the PSD and spine at large from any major accidental disruption of the kinase:phosphatase balance established through LTP/D. They are an extra layer of stabilisation against change that is importantly easily overridden.

GluA1 subunits appear to be preferentially inserted into the synaptic zone under unstable synaptic conditions (T. Benke & Traynelis, 2019)⁶⁹. There is growing evidence for homomeric GluA1 AMPARs being held perisynaptically or intracellularly close by to respond and reinforce plasticity inducing stimuli (K. He et al., 2009; Mattison et al., 2014; Y. Yang et al., 2008). Residue modifications enable the mobilisation to, and stabilisation in, the membrane necessary for an increased or decreased sensitivity of the system to incoming inputs (T. Benke & Traynelis, 2019; Diering & Huganir, 2018). I argue change is *necessarily* unstable- any stable change in conductivity could risk driving the system into either cytotoxic or silent

⁶⁹ Alternatively, they *cause* the unstable state. Maybe a bit of both.

states, both ultimately causing synaptic dysfunction (Jurado, 2018; Whitcomb et al., 2015).

GluA1 subunits are thus tightly regulated through a “code”, whose interactions are complex but whose consequences are simple – balancing regulation of GluA1-containing AMPARs (Diering & Huganir, 2018). The longer GluA1-containing receptors remain in the synaptic zone, the more likely they are to be activated by glutamate and induce changes to synaptic plasticity. In this hypothesis, there are no slots because this is not a limiting factor – what matters is the balance of diffusion between extrasynaptic and synaptic zones, and the function of many of the AMPAR binding partners and residue modifications is to alter the rate of this diffusion. Complexity in trafficking is then the result of a dynamic competition between tens of modifiers, of which phosphorylation plays an oversized role.

The Exocytosis variant is a lesser but important narrative addition to the above. The maintenance of potentiation requires a constant supply of AMPARs, and the composition of these will determine whether the synapse maintains any changes and determine future responses to input patterns. This might be in turn driven by previous patterns of activity, both at this and previous synapses. The model’s kinetics here assume that initial changes to synaptic receptors necessitates at least some extrasynaptic membrane-bound receptors. It does not however rule out much faster exocytosis mechanisms from working in parallel. Given recent evidence for much lower concentrations of extrasynaptic AMPAR than previously assumed (Helm, 2019; S. H. Lee et al., 2017), such quick exocytosis mechanisms may play a greater role than anticipated.

The Merged Model

The merged model assumes that slots are present but promiscuous and are only stable in small numbers initially. CaMKII activity stabilises extra slots in an activity-dependent manner. Slots are unstable and may disappear in low CaMKII conditions. Entry and exit for GluA1/2 AMPARs is determined by their phosphorylation profile. S845 phosphorylation enhances entry, with S818 also sufficient and together offering an additive effect. A GluA1/2 AMPAR is thus only removed if both S831 and S845 are unphosphorylated – activity of either kinase is sufficient for maintenance.

Exocytosis works as two distinct mechanisms – a constitutive one dependent only on GRIP, PICK and NSF, and an activity dependent one. The activity dependent one allows rapid transfer of intracellular AMPARs to the extrasynaptic zone when phosphorylated at S845.

Results

The merged model fails to show any AMPAR trafficking

The combined model is unable to drive receptors into the membrane of any composition, with all protocols leading to a complete loss of all receptors by the end of the simulated time (**Figure 55**). This failure is not because of a failure in calcium signalling, with both kinase and phosphatase activity proportional to the strength of the induction protocol used. By the end of the simulation, no significant difference in synaptic receptors was found between the groups ($F(5,20) = 0.355$, $p = 0.873$).

The failure in the PP-LFS protocols, which previously showed strong potentiation, suggests the rules and rates to enter the synaptic zone shift the system to prefer the net movement of receptors out of the synapse (**Figure 56**). Unlike in the previous experiments, there were no significant differences in synaptic receptors between the LTD protocols by the end of the simulations ($F(2,6) = 0.970$, $p = 0.431$).


- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP

Figure 55 Merged Model (LTP) shows no AMPAR trafficking to the membrane.


Despite calcium-mediated activation of CaMKII, AMPARs of all types are unsustainable in within the synaptic zone in the merged model. n=5 except for Delta (n=3), sTheta (n=4), HFS (n=4).

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation




Free PP1




- LFS
 - LFSP200
 - LFSP50

Figure 56 Merged Model (LTD)
 suggests all protocols are too weak to maintain synaptic AMPARs. Whilst all protocols show a decay of synaptic AMPARs, this appears to be not as a result of the LTD stimulus regime, but a non-activity dependent decay mechanism. PP-LFS interestingly slows this decay and shows a large gradual increase in activated CaMKII, which decays rapidly post-protocol. n=4 LFS, n=3 LFSP50, n=2 LFSP200.


Receptor Population GluA1/2




GluA2/3




Receptor Phosphorylation S831




S845



CaMKII phosphorylation



Free PP1



Discussion

It is apparent from the above data that the process of merging the three variants removed much of the expected trafficking phenomena. The data shows that whilst enzymatic activity was reduced, it does not show behaviours that are suggestive of a failure in these parts of the model rule structure. This suggests that the assumptions made during the merge process on trafficking, as well as the adjustments made during that process to variables and initial conditions, is adversely impacting the viability of the model. In particular, the fact that all protocols result in an almost complete loss of AMPARs over the simulation duration suggest that the rules and variables here are no longer representative of physiological behaviour. A full reassessment of the model components is therefore necessary to determine what elements are incorrect, and to keep each variant separate until such a reassessment is complete.

Summary

- The merged model assumes a greater role for Slot and Receptor hypotheses over Exocytosis, but still assumes that receptors are also slowly trafficking between cytosol and membrane.
- The subsequent merged model fails to output any potentiation of AMPAR numbers in the synaptic zone, even when under strong induction protocols.
- The model tends towards a complete removal of AMPARs from the membrane, suggesting a failure in the trafficking rules or variables.

Experiment #8 Modelling AMPAR Trafficking during Plasticity Events, Part II

Model Reappraisal and Refinement

Given the number of changes made during the switch to both the trafficking hypotheses testing and the merged model, and the subsequent failure of the latter, it was necessary at this point to reassess those changes and audit them.

The rules and assumptions made in the merged model whilst justified were judged to be too hasty whilst the individual variants were still not fully explored. As such the model was broken back down into the three trafficking hypotheses for the remaining experiments.

The only variable change found initially was that at around the time of Experiment 5, the initialisation of the model had resulted in two populations of PP2A. This doubled the number of PP2A agents, and resulted in a possibly crucial balance change in the model, where PP1 sequestering was reduced, and S818 phosphorylation more severely limited than expected.

Further exploration also indicated that the implementation of the Stefan & Coomber values (See Experiment 5) was only partially achieved. The PP1 binding, unbinding, and dephosphorylation of CaMKII values were not called in the *Rules.ka* file, where the old values were written in as actual values rather than with a variable name.

These changes were significant enough to require repeat runs of the trafficking hypotheses to investigate how much of the previous findings were because of these errors, or in spite of them.

In previous model iterations, the particle numbers from (Gallimore et al., 2016) for GRIP, PICK1 etc. were all scaled proportional to the compartment size to accommodate this model's use of $0.01 \mu\text{m}^3$. In the original Gallimore model, there is a note that the published data assumes that GRIP and PICK1 are never saturated (based on Czöndör et al., 2012), leading to a constant concentration of free GRIP and PICK1 in the model via diffusion from a bulk cytosol compartment. However, from further investigation I found that recent GRIP copy number estimates for a dendritic spine of a similar volume was taken to be just 200 (to stabilise 100 synaptic AMPARs) (Antunes et al., 2016) and experimentally derived values were even lower, at only 17 molecules for spines which were on average ten times larger

than the current one (Helm, 2019)⁷⁰. The value was therefore changed from 1330 to 80 molecules, and the PP2A-value adjusted models rerun to see if this too made any effect on the model output.

Results

The Exocytosis variant accommodates transient activity-driven receptor insertion

Unlike previous attempts, exocytosis now sustains a population of both GluA1/2 and GluA2/3 AMPARs in the synapse (**Figure 57**). As in the prior versions of receptor and slot variants, the higher frequency protocols (cTheta, Delta, HFS) show a distinct increase in GluA1/2 receptors compared to the lower frequency protocols. In this instance however, sTheta also ends with the same level of synaptic receptors. Whilst stable, the final receptor levels are not much higher than the beginning 13 receptors, with weakLTP being the lowest ($M = 6.8$, $SD = 3.27$, $CI = [2.74, 10.86]$) and Delta the highest ($M = 16$, $SD = 3.67$, $CI = [11.44, 20.56]$). GluA2/3 receptors show a slow but persistent incorporation into the synaptic zone, which is the same for all protocols. This change is between 1 – 3 receptors.

By the end of the simulation, a significant difference in synaptic receptors was found between the groups ($F(5,24) = 3.647$, $p = 0.014$). Tukey HSD post hoc analysis shows significance differences between weakLTP and cTheta ($p = 0.026$) and Delta ($p = 0.015$).

The Receptor Phosphorylation variant shows stable synaptic AMPAR populations and activity-dependent changes in total receptor number

All protocols show a stable and increased number of synaptic AMPARs for the duration of the simulation (**Figure 58**). cTheta, Delta and sTheta, and HFS, ,Theta and weakLTP are grouped in their response profile regarding GluA1/2 synaptic expression. All are above initial basal levels, ranging from 18 - 42 receptors above basal levels by the end of the simulation, or an increase of between 242 - 422 %. GluA2/3 AMPARs averaged by the simulation's end 8 receptors across all protocols, increasing rapidly post-induction in all cases.

cTheta, Delta and sTheta all show persistence in CaMKII activation, with the remaining protocols showing a decay of active subunits post-protocol. By the end of the simulation, a significant difference was found in synaptic receptors between the groups ($F(5,24) = 2.875$, $p = 0.036$). A Tukey HSD post hoc analysis shows only

⁷⁰ In addition, both the original, and an alternative model available in Gallimore et al., (2016), which makes use of finite GRIP and PICK1, use a small number of molecules. A disconnect between the published model and the COPASI model file however makes it hard to ascertain if the values in these model versions are correct.

a significant difference between cTheta and Theta ($p = 0.029$). The Slot variant shows distinct activity-dependent profiles and capacity for a stable potentiation of the synapse

Half the LTP protocols (cTheta, Delta, HFS) show robust increases in both AMPAR types, coming close to saturating the available slots ($M = 68.0, 65.4, 72.0$, $SD = 1.73, 2.60, 0$, $CI [65.84, 70.15; 62.16, 68.64; 72.0, 72.0]$ respectively) (**Figure 59**). These are stable for the duration of the simulation. sTheta shows similar stability but at a slightly lower level ($M = 53$, $SD = 13.1$, $CI = [36.69, 69.30]$). Theta shows a large range of values between the five simulations ($M = 21.4$, $SD = 23.00$, $CI = [-7.16, 49.97]$), ranging between 7 and 62 final synaptic AMPARs.

This persistence in cTheta, Delta, and HFS protocols appears to be a consequence of CaMKII autophosphorylation overcoming PP1 dephosphorylation capacity, indicated by the considerably greater depression of free PP1 levels compared to the Theta and weakLTP protocols. Looking at individual runs (data not shown) cTheta, Delta and HFS show a 100% success rate of this CaMKII-mediated stabilisation of synaptic AMPARs. sTheta stands at 80%, and Theta 20%. Only the weakLTP protocol failed to ever reach a potentiated state that lasted the duration of the simulation.

By the end of the simulation, a significant difference was found in synaptic receptors between the groups ($F(5,24) = 30.671$, $p = 1.1194E-9$). *Levene's F* is significant ($p = 0.018$), but *Welch's F* test not possible due to HFS having zero variance. Therefore a post hoc Game's Howell analysis was used. HFS mean difference was significantly different ($p > 0.05$) than all except sTheta ($p = 0.160$). weakLTP mean difference was significantly different ($p < 0.05$) than all other protocols except Theta ($p = 0.751$).

Reducing GRIP values allows for greater chance of CaMKII-mediated stabilised synaptic potentiation

Reducing GRIP levels allows for an increased probability of receptor trafficking. All three variants show that a reduction in GRIP causes a greater number receptor numbers localised to the membrane. Exocytosis shows a much stronger transient incorporation of GluA1/2 AMPARs, although this is not sustainable, and they are all removed by the end of the simulation (**Figure 60**). The GluA2/3 receptors however are inserted much more quickly and robustly, averaging to 10 receptors by the end of the simulation. No significant difference is found between the protocols ($F(5,24) = 0.223$, $p = 0.949$).

The Receptor Phosphorylation variant shows more pronounced and stable maintenance of GluA1/2 AMPARs across their respective simulations, whilst

GluA2/3 incorporation and maintenance are broadly identical to the higher GRIP model (**Figure 61**). All protocols except Theta and weakLTP display CaMKII-mediated stabilisation and maintenance. Looking at the individual runs (data not shown), Theta could generate a similar “high” state of AMPARs and CaMKII, but this was only 20% of the time. WeakLTP did not achieve this in any of the simulations. At the final simulation time, a significant difference in synaptic receptors was found between the protocols ($F(5,24) = 7.704$, $p = 0.000193$), with weakLTP significantly ($p < 0.05$) lower than all protocols except Theta ($p = 0.926$). Theta was significantly lower than only the cTheta ($p = 0.012$) and Delta ($p = 0.014$) protocols.

The slot variant shows greater GluA1/2 and GluA2/3 receptor stability, without any change to the magnitude of the initial response (**Figure 62**). Once again, all protocols except Theta and weakLTP show a CaMKII-mediated stabilisation of GluA1/2 receptors, but also of GluA2/3 receptors. Theta can succeed to a high state, but only occurs in 40% of the simulations, with the remaining simulations failing to show any sustained potentiation of the synapse. At the final simulation time, a significant difference in synaptic receptors was found between the protocols (Lavene’s F test was not met ($p = 1.2323\text{E-}12$); *Welch’s* $F(5,10.676) = 1206.593$, $p = 2.5572\text{E-}14$). These differences through post hoc Games-Howell are between weakLTP and all other protocols except Theta.

Removing Slot decay shows extremely stable and distinct protocol responses


The assumption that slots can disappear presumes CaMKII activity is necessary to maintain a PSD environment, and that its loss triggers synaptic remodelling. If the decay of these slots is removed, the model only retains the rule that increasing CaMKII activation led to more slots being present. In these simulations, weakLTP, which is unable to trigger CaMKII phosphorylation, shows no change to the synapse (**Figure 63**). The remaining protocols are grouped based upon their activation level of CaMKII, with cTheta, sTheta, Delta and HFS triggering the maximum number of possible slots. In all these runs, the number of synaptic receptors post induction remains largely unchanged throughout the simulations.

At the final simulation time, a significant difference in synaptic receptors was found between the protocols. *Levene’s F Test* for the homogeneity of variance is significant ($p = 0.000037$). Therefore, *Welch’s F* test was used, which still shows significance (*Welch’s* $F(5, 10.743) = 10.374$, $p = 0.000782$). Post hoc Games-Howell analysis shows significance ($p < 0.05$) between weakLTP and all other protocols.


- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP

Figure 57 Exocytosis variant with updated PP2A values and PP1-CaMKII interactions: activity-driven AMPAR insertion now possible.
The model shows a clear distinction between activity driven and constitutive trafficking of different AMPAR types. GluA1/2 AMPARs can be incorporated from internal stores in the higher frequency stimulation induction protocols, although this is not sustained over the simulation. n=5 for all protocols.


Receptor Population GluA1/2




GluA2/3




Receptor Phosphorylation S831




S845



CaMKII phosphorylation



Free PP1




- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP


Figure 58 Receptor variant with updated PP2A values and PP1-CaMKII interactions: more CaMKII activation and protocol differentiation in AMPAR numbers.

The model shows distinct and stable profiles for both AMPAR populations, with all protocols showing enhanced levels of AMPARs by the end of the simulation compared to initial conditions. This maintenance is not dependent on sustained CaMKII activation, although the magnitude of the final synaptic AMPAR population is. $n=5$ for all protocols.


Receptor Population GluA1/2




GluA2/3




Receptor Phosphorylation S831




S845



CaMKII phosphorylation



Free PP1




- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP


Figure 59 Slot variant with updated PP2A values and PP1-CaMKII interactions: strong CaMKII activation and synaptic AMPAR numbers.
LTP protocols show a distinct hierarchy in magnitude of response, with higher frequency ones showing greater enhancement of total synaptic AMPAR numbers. HFS, Delta and cTheta's notable maintenance above baseline levels is attributable to CaMKII autophosphorylation reducing levels of PP1 available to dephosphorylate AMPARs. n=5 for all protocols.

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation



Free PP1




- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP


Figure 60 Exocytosis variant with reduced GRIP numbers: greater AMPAR insertion that lasts for longer. Fewer GRIP numbers allow for GluA2/3 AMPAR trafficking to gradually increase across the simulation runtime. Activity-driven AMPAR insertion is larger in magnitude but fails to maintain a population of GluA1/2 AMPARs by the end of the simulation for all protocols. n=5 for all protocols.

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation




Free PP1




- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP

Figure 61 Receptor with reduced GRIP numbers: greater CaMKII activation and stabler GluA1/2 synaptic numbers. Fewer GRIP molecules show no effect on constitutive AMPAR trafficking, but considerably destabilises synaptic GluA1/2 AMPARs, preventing maintenance in all but the cTheta protocol. In this protocol, CaMKII-mediated suppression of PP1 activity maintains AMPARs through persistent kinase activity. n=5 for all protocols.


Receptor Population GluA1/2




GluA2/3




Receptor Phosphorylation S831




S845



CaMKII phosphorylation




Free PP1




- cTheta
 - Delta
 - HFS
 - sTheta
 - Theta
 - weakLTP

Figure 62 Slot variant with reduced GRIP numbers: synaptic AMPAR numbers more stable.
 weakLTP is the only protocol to show no temporary enhancement of AMPAR numbers during the simulation. Whilst cTheta shows a stable maintenance of AMPAR numbers, all other protocols decay quicker than when higher GRIP numbers are used. n=5 for all protocols.


Receptor Population
GluA1/2




GluA2/3




Receptor Phosphorylation
S831




S845



CaMKII phosphorylation



Free PP1




- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP


Figure 63 Reduced GRIP Slot variant with no decay: even greater stability of synaptic AMPARs.
 Allowing slots to exist permanently once generated shows a clear stratification of induction protocols and decouples the need for CaMKII-mediated saturation of PP1 to allow for stable synaptic GluA1/2 AMPARs. n=5 for all protocols.

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation



Free PP1



Discussion

The resetting of the PP2A values and the incorporation of the Coomber et al (1998) PP1-CaMKII interaction set allows for stronger and more stable potentiation in all three model variants. In particular, the sTheta protocol become much more viable, with the third and final train of calcium leading to comparable final receptor numbers to the massed protocols.

It is arguable that spaced protocols like sTheta may be more likely to succeed because they offer a greater number of opportunities for PKA activity and CaMKII phosphorylation to correlate sufficiently to saturate PP1. Whilst not the most successful protocol here, given the uncertainty in parameters and localisation of molecules this may be a feasible explanation for *in vivo* observations that use this protocol.

There was much more successful phosphorylation of CaMKII, PP1 saturation, and PKA activation than seen in Experiment 6. This was exacerbated in the reduced GRIP versions, exemplifying the complex and delicate balance between AMPAR localisation and their main enzymatic interacting partners. Causally, the model is no closer to ascertaining the exact structure of this interaction, and will require more contextual and spatial richness to work more on this aspect. What is clearer from these findings is that all three trafficking hypotheses are viable, and that different patterns of calcium flux can have unique fingerprints on the system state. It will be interesting to start using the model to explore these in much richer detail.

A significant finding from using the reduced level of GRIP was it shows that the Exocytosis variant could now show a pronounced activity-driven increase in the number of GluA1-containing AMPARs for a brief period post-induction. This is consistent with evidence that argues AMPARs are not maintained as an extrasynaptic population, but are predominantly either synaptic, peri-synaptic, or intracellular (S. H. Lee et al., 2017). This mechanism could therefore provide a temporary increase in the number of available AMPARs sufficiently to maintain synaptic sensitivity long enough to induce a state change driven by CaMKII-mediated saturation of PP1, and the recruitment of AMPARs from nearby spines.

Summary

- Adjusted PP2A values and PP1-CaMKII interactions allow for greater probability of receptor maintenance in the synaptic zone post-LTP induction.
- These rates also allow for a brief but large incorporation of intracellularly sourced AMPARs into the synaptic zone, which is consistent with experimental evidence.

- Reduction in GRIP numbers in line with evidence from Helm (2019) improves the stability of AMPAR numbers in the Receptor Phosphorylation and Slot variants, and improves the transient delivery of receptors in the Exocytosis variant.

Experiment #9 A Monomeric Model of Synaptic AMPAR Trafficking

Introduction

In Chapter 1 we alluded to the most recent findings that suggest AMPARs are transient structures that exist for less than a second. In this theory, monomer and dimer forms diffuse within membranes and transiently form receptor complexes (or trimers). The existence of these transient receptors is brief, but the sheer number of available components results in a stable observable population of receptors at synaptic sites (Morise et al., 2019).

This approach has several trafficking advantages, as monomers would be less impeded within membranes, and be able to be more tightly packed within vesicles for activity-driven release into membranes. It also allows for simpler regulation of population dynamics within a spine by focussed movement of monomer/dimer pairs to reshape the synaptic AMPAR profile. The timescale of existing is also practical, as receptor open states occur within milliseconds, which allows for multiple feasible open/close events for the lifetime of a receptor (Morise et al., 2019).

To assess the feasibility of this theory, I have reconstructed the existing rule-based model to shift AMPAR properties and states to individual monomers. The aim was to assess whether the outputs of the model would be qualitatively similar to those seen in previous experiments within this chapter. To be feasible, subunit exchange must be able to occur in all compartments, and the population of AMPARs must not drastically vary within the second timeframe with which the model records outputs from that seen in previous model iterations.

Methodology Changes

All AMPAR properties were reduced to single monomers as opposed to a monolithic receptor species. As the open state of the receptor is based upon Monod-Wyman-Changeux interaction rules, I have not allowed for monomers to possess the small medium or large open states unless as part of a full tetramer. This contrasts with other allosteric theories, which allow for individual monomer conformation changes (Koshland, Némethy, & Filmer, 1966). In this version, glutamate binding is allowed, but does not cause any changes unless the subunit is part of a tetramer.

To avoid fixing monomers in open states, new rules adjust any monomer or dimer pair so that the subunit(s) revert to a basal (closed) state. In some instances, the number of binding partners for a tetramer has increased. For example, GRIP and

PICK would only have one binding site for a tetramer in previous model iterations, whereas here it is necessary to allow for binding to any subunit.

Additional rules were written to simulate dimer and tetramer formation. Trimer formation, whilst possible, was omitted here in the interests of simplicity. Rule parameters were estimated to fit with the lifespans of tetramers and dimers published in (Morise et al., 2019). They do not take into account the difference in affinities between subunits however at this stage.

Unlike the typical model, we include the possibility of GluA1 homomeric receptors here. The nature of allowing for dimer and tetramer formation, and this perhaps being dictated by the relative abundance of subunits, made its exclusion in this instance a possible large confound to any observations.

Model observations were edited to show total monomer populations within compartments as well as tetramers. The number of membrane-bound monomers was set much higher than previous experiments' tetrameric numbers, to assess the feasibility, as we assume that the number of monomers and dimers may be much higher than previously thought. We therefore moved the monomer numbers (500 GluA1 and GluA2 monomers, and 200 GluA3 monomers) which were normally within the cytosol, to the membrane. These were situated in the extrasynaptic region (the Exocytosis variant was not tested in this experiment).

The model was given theta burst stimulation during initialisation to populate the synaptic zone with a population of AMPARs and allowed to run for 300 s. This did not result in any excessive number of synaptic AMPARs, with the Receptor variant initialisation resulting in 22 synaptic receptors (17 GluA1/2, 5 GluA2/3), which is comparable with the typical model.

The Slot variant of the model had a high failure rate during simulations. This led to almost no runs reaching the 5000s which was standard for the rest of this Chapter. As most runs failed after 4000s, this was used as the end time for simulations of this variant⁷¹.

⁷¹ The reason behind this failure is unclear. Of the three successful runs, two show only that GluA2/3 receptor numbers plateau at roughly their 4000 s values. The final run shows a sudden drop of observable receptors of both populations, which may be due to the recovery of PP1 from saturation. Further work is required to explore this further.

Results

The monomeric hypothesis does not prevent AMPAR formation and detection

The model is able to form populations of GluA1, GluA1/2 and GluA2/3 tetramers, which are constantly being made, but due to the numbers present, form stable population levels (**Figure 64**). At an observation rate of 1/s, almost all monomers are part of a tetramer complex. Whilst not a part of the following graphs, 48 GluA1 homomeric AMPARs are present in each variant, along with 147 GluA1/2 AMPARs and 97 GluA2/3 AMPARs. Calcium entry during induction protocols has no effect on receptor formation or composition ratios.

GluA2/3, and GluA1/2 AMPARs are observable post LTP induction protocols within the synaptic zone

Despite being able to form GluA1/2 AMPARs, maintaining these as activity-driven, significantly increased populations within the synaptic zone is not possible in most of the protocols or variants tested (**Figure 65**, **Figure 67**). HFS in the Receptor Phosphorylation variant is the only LTP protocol to show any considerable activity-driven incorporation of GluA1/2 AMPARs which is sustained for the whole simulation time ($M = 39.2$, $SD = 4.76$, $CI = [33.28, 45.11]$) (**Figure 65**). The other protocols in Receptor Phosphorylation variant do show an ability to sustain GluA1/2 AMPARs, but at a lower level ($M=14.96$).

GluA2/3 AMPARs on the other hand are observed in stable, maintained populations in all protocols ($M = 7.13$) by the end of the simulations.

In the Receptor Phosphorylation variant there is a significant difference between the groups ($F(5,24) = 3.275$, $p = 0.021$). A Tukey HSD post hoc analysis indicates HFS has significantly higher ($p < 0.05$) levels of synaptic receptors than all protocols except Delta ($p = 0.069$) and weakLTP ($p = 0.09$).

Protocols in the Slot variant show no impact on the behaviour of AMPARs in the model. (**Figure 67**). In all cases a gradual increase in GluA1/2 and GluA2/3 AMPARs occurs at the same rate, with an average of 139.83 AMPARs across all protocols ($SD = 16.39$). As such, total synaptic numbers at the end of the stimulation show no significant difference between the groups in the Slot ($F(5,24) = 1.765$, $p = 0.158$) variant.

LFS induces a potentiation of the synapse through the formation of different AMPAR tetramer types

As in the standard model variants, LFS potentiates rather than depotentiates the synapse. Unlike those model variants however, much of this potentiation occurs very differently in the Receptor Phosphorylation variant compared to the Slot

variant. The receptor variant acts much like the standard monolithic tetramer model, with activity-driven trafficking of GluA1 and GluA2 monomers to form GluA1/2 AMPARs, tied closely to the activation level of CaMKII (**Figure 66**). GluA2/3 receptors are present, but in very low numbers, and are unaffected by the LFS protocol.

The slot variant however shows the same large increase in the number of both GluA1/2 and GluA2/3 tetramers, which reaches an averaged peak of 83 and 74 receptors respectively (**Figure 68**). This is the same behaviour seen in the LTP protocols, where the actions of the model are unaffected by the protocols themselves.




Figure 64 AMPA-tetramer complexes form and show stable populations in the monomeric model.


The monomeric model shows clearly that even with constant degradation and reforming of AMPAR tetramers, there is a stable population observable on the seconds level. The proportions of monomers used produces distributions of AMPARs comparable to estimates of AMPAR population distributions (W. Lu et al., 2009; Schwenk et al., 2014; Wenthold et al., 1996). Examples here from Receptor variant (n=5 for each protocol).

- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP


Figure 65
Monomeric Model:
Receptor Variant
(LTP): HFS can
trigger formation of
GluA1/2 synaptic
AMPArs.

Only the HFS protocol was able to elicit both activity-driven synaptic incorporation of AMPARs and CaMKII activation. GluA2/3 population are low but persistent in all protocols. n=5 for all protocols.


Receptor Population GluA1/2




GluA/3




Receptor Phosphorylation S831




S845



CaMKII phosphorylation



Free PP1




- LFS

Figure 66
Monomeric Model:
Receptor Variant
(LTD): Enhanced
stable populations
of GluA1/2
receptors.


LFS shows gradual increases in GluA1/2 AMPARs which decays post-protocol but still shows potentiation over the initial levels. This increase is considerable and is associated with sustained CaMKII activation. $n=5$ for all protocols.

Receptor Population

GluA1/2




GluA2/3




Receptor Phosphorylation


S831




S845



CaMKII phosphorylation




Free PP1




- cTheta
- Delta
- HFS
- sTheta
- Theta
- weakLTP

Figure 67
Monomeric Model:
Slot Variant (LTP):
AMPA trafficking
insensitive to
induction protocols.
 No change in GluA1/2
 AMPARs is seen in
 any protocols, despite
 CaMKII activation
 seen in the HFS
 protocol. However
 both AMPAR types
 increase gradually
 over the run. n=5 for
 all protocols.


Receptor Population
GluA1/2




GluA2/3




Receptor Phosphorylation
S831




S845



CaMKII phosphorylation



Free PP1




- LFS


Figure 68 Monomeric Model: Slot Variant (LTD): AMPAR trafficking insensitive to LFS but triggers high levels of CaMKII activation.

A slight temporary increase in GluA1/2 AMPARs is seen, correlated with CaMKII-mediated PP1 suppression. A large increase in this AMPAR population is present at the end of the simulation, as in the LTP runs. GluA2/3 AMPARs are suppressed until after the LFS protocol, after which they increase at a comparable rate to the LTP runs. $n=5$ for all protocols.


**Receptor Population
GluA1/2**




GluA2/3




**Receptor Phosphorylation
S831**




S845



CaMKII phosphorylation



Free PP1



Discussion

The monomeric model, whilst unable to show any protocol-driven change in 1/2-containing AMPARs at the synapse like those seen in Experiment 8 (outside of HFS in the Receptor Phosphorylation variant) does indicate that the hypothesis is viable in observing a stable population of synaptic AMPARs, even though the lifespan of each tetramer is only milliseconds in length.

The failure of some of the LTP protocols to show activity-driven stabilisation of GluA1/2 AMPARs is at this stage likely due to the paucity in data on how the monomeric populations work at spines. The model has a very large pool of monomers, but does not have any major diffusion of monomers from extra-spine regions, for example. It may be that this and the initial AMPAR subunit numbers chosen are too low for the spine modelled. Another strong possibility is that the monomeric model may depend more upon spatial localisation of molecular components. Tetramers may seed around complexes localised to extrasynaptic or synaptic positions, with complex composition influencing both tetramer lifespan but also tetramer composition itself.

The HFS and LFS protocols in the Receptor Phosphorylation variant show that certain patterns of calcium flux are able to cause large numbers of observable synaptic GluA1/2 AMPARs, but not GluA2/3 AMPARs. This is driven as usual by sufficient activation of CaMKII depressing PP1 activity, sustained due to a switch of CaMKII to the high state.

The most striking outcome of this experiment was the major disparity in the trafficking hypotheses. Whilst the Receptor Phosphorylation variant behaved comparably to the model seen in Experiment 8, the Slot variant showed a drastically different behaviour. None of the protocols have an effect on the trafficking of GluA1/2 or GluA2/3 AMPARs, with the latter increasing gradually and the former exponentially by the end of the simulation. As with the Receptor Phosphorylation variant, HFS protocols induce a stable high level of activated CaMKII, with the other protocols seeing similar low-level activation. In this case however it has no noticeable effect on receptor trafficking or phosphorylation profiles. A similar decoupling of trafficking and enzyme activation is seen in the LFS version. All of this requires careful examination, of both the nature of the rules underpinning the model, and what it can say of the validity of the monomeric hypothesis.

As the monomeric model makes no changes to the enzymatic action beyond the identity of their targets, that fact that receptors still may form, be localised synaptically, and do so differentially depending on their composition makes the monomeric hypothesis a viable and potentially quite powerful explanation of AMPAR dynamics at the synapse. Both experimental work and adjustments to the current model are necessary however to confirm both the validity of the prior investigations, as well as the lifespan and composition of the AMPAR complexes formed in the simulations.

Further analysis of the internal state of the model during simulations highlighted that some ambiguity in the rule syntax allows for unwanted binding events to occur, clumping molecular agents together in a manner not expected to occur *in vivo*. This is primarily within the trafficking machinery, where the PICK/GRIP/PKC complexes are shared in part between subunits on different receptor tetramers. Future simulations of this model should remove this ambiguity and prevent such bridging complexes to form. However, this result may be worth pursuing in any spatial variant of the model. As previously stated, these membrane complexes may be what act as modifiers to AMPARs location and velocity through the membrane. Sharing complexes through binding to different parts may act to localise AMPARs to particular points above the PSD, but be transient bindings such that AMPARs are still free to move over longer timescales. This highly dynamic interaction hypothesis, if true, would likely benefit the monomeric model of AMPARs.

Summary

- A model where AMPARs are modelled explicitly as transient shows stable AMPAR populations on the seconds timescale.
- These receptors can be modified and trafficked in the same way as when treated as stable tetramers.
- Both GluA1 and GluA2-containing receptors can be inserted into the synaptic zone through calcium driven activity and through constitutive transfer of subunits.
- The Slot model variant shows no impact of the protocols on AMPAR trafficking.
- The Receptor Phosphorylation variant shows stable increased AMPAR populations in all protocols, with HFS showing consistently significantly higher levels across simulation runs.

- LFS shows a substantial increase in AMPARs present synaptically post-protocol which persists for the remainder of the simulation, similar to the original model.

General Discussion

This chapter has covered the conception and development of the first functioning rule-based model of AMPAR function at the synapse. I have shown that the rule-based model environment is capable of replicating existing models accurately and reliably, and can combine these models easily. I have also shown through the course of testing and building the model the capacities for investigation afforded by this method that outstrip those possible with more traditional biochemical modelling approaches. The model comprises of many the key molecular interactions involved in synaptic plasticity, at a range of scales which integrate to form a naturalistic model of AMPAR responses to changes in calcium which has until recently not been possible to model easily.

Both from the volume of investigation, and the importance of synaptic plasticity in information transfer, AMPARs have many known modifications and binding partners, as we saw in Chapter 1. This model has been developed to be the foundation of a comprehensive modelling project that investigates these known changes both *in toto* and in a naturalistic context across multiple time scales. I have shown here that Kappa can cope with the volume of interactions generated by such a complex model, which has more than 2 billion events during a 5000 s run, and allow for a much more comprehensive investigation of AMPAR properties in relation to its function *in silico*.

This final point is important because it is increasingly more difficult for *in vitro* and *in vivo* experiments to be designed that can consider all the possible off-target, indirect or emergent properties being identified as potentially part of typical AMPAR function. The possible state profiles that a PSD-anchored GluA1/2 AMPAR may possess in the model is 8,398,080. The level of interactions and modifications make identifying the causal pathways accurately more difficult and require a much greater number of control experiments to be conducted to be determined. *In silico* approaches were in part designed to allow for these types of controls and preliminary investigations to be done without the considerable cost in time, resources, and animals, but until recently fell short in the level of complexity they could afford to an experimenter. This rule-based model provides a much more complex picture of AMPARs and their contextual environment than previously available, and can be rapidly implemented to predict the likely outcomes of a given experimental manipulations on AMPAR function.

The model also helps to highlight several areas where our experimental understanding misses out some of the mechanistic detail necessary to replicate *in vivo*

function faithfully. For example, the modifications to residue S818 are known to require PKC, but no phosphatase has been conclusively shown to be responsible for removal of the phosphate. Additionally we lack evidence for just which isoforms of PKC work on this residue, which is especially important given the different activity driven and constitutive isoforms, and their differential localisation (Chu, Fioravante, Leitges, & Regehr, 2014; Sossin, 2007).

PKA signalling, which here is tightly controlled and crucial to any activity-driven trafficking of AMPARs, is also shown by the model to be absolutely necessary for accurate synaptic modelling. Whilst some recent work has begun to unpick some of the fine filtering this cAMP-mediated pathway performs on incoming signals (Ohadi et al., 2019), more is necessary to understand the key elements of this pathway to be incorporated into future simulations. At present, I believe the current model would benefit from a more accurate cAMP activation module, taking into account adenylate cyclase activity and the influence of intracellular calcium release outwith times of external input.

Most dramatically, we have seen how LFS of the synapse can lead to potentiation when kinetic conditions match those in the model. Whilst we cannot rule out the possibility this is due to a paucity in the calcium dynamics of the model, it also lends support to several findings that show LFS and LTD to be complexly related, and likely very dependent on the spine's molecular composition and local environment for the resultant output profile (Kemp et al., 2000; Pinar et al., 2017). The monomeric model outcome, where GluA2/3 AMPARs not only form but stabilise in the synaptic zone only after an LFS input, points similarly in this direction.

The model does present some more immediately practical headaches regarding LFS and LTD. Is the pronounced potentiation triggered by PP-LFS a symptom of improper calcium signalling, or a response normally prevented by the interaction with related mGluR signalling pathways, which have long been shown to be involved in at least some LTD types (Bashir & Collingridge, 1994; Delgado & O'Dell, 2005; Zho, You, Huang, & Hsu, 2002)? Does the action of LFS alter depending on the composition and total numbers of AMPARs present upon induction?

Despite some of these gaps, the model has also proven insightful in assessing the veracity of several AMPAR-related hypotheses of synaptic function. Intracellular pools of AMPARs expressed to the surface after strong activity is potentially viable as a method of adjusting synaptic strength, although it does appear to be insufficient

for any maintenance on its own. This insertion hypothesis also suggests possibilities such as the transportation of synaptic states to neighbouring spines through vesicular transfer to their pools, communicating the local context and therefore priming a new or nascent spine for expected input.

Exocytosis may also be a crucial player in certain forms of plasticity induction or maintenance, and whilst not always viable here, the fact it *is* feasible renders this a potentially fruitful avenue of research. Equally intriguing is the preliminary findings here that a monomeric model of AMPARs is also kinetically feasible and should be seriously considered as a powerful and persuasive alternative to monolithic stable tetrameric hypotheses of both AMPAR structure and function.

Rule based models meet with many of the desired requirements of maximal models laid out in Chapter 3. Rather than immediately supplant traditional approaches however, I believe the Kappa models should act in partnership with new highly sophisticated synaptic models, such as that of Mäki-Marttunen et al. (2020). They use a statistical model to estimate the properties of AMPAR subunits which depends on the composition and number present in the membrane, which is related to the activity of the rest of the biochemical network. Much like in this rule-based model, their statistical approach allows them to have the possibility of many different feasible AMPAR tetramer profiles, with the advantage of maintaining an ODE-based model network elsewhere.

This allows for much more rapid simulation run-times and parameter scanning, which are the biggest pitfalls to the rule-based approach. My model currently takes between 36 – 48 hours to run a single 5000 s simulation. If the two model types were to be calibrated to one another, the quicker simulation predictions can then be tested in the rule-based approach as a kind of verification step, ensuring that assumptions made in the statistical model for example are valid.

As the models become more complex, the rule-based approach becomes the more sensible (and likely tractable) approach. But by being calibrated and used as part of a verification pair with quicker models, we can be surer of its accuracy and therefore accept more willingly the punitive simulation times⁷².

⁷² If this obstacle could be overcome, the rule-based approach supersedes the ODE one almost immediately, even for small models.

The model I present is a faithfully rendered foundational model of AMPARs at a synapse. It is possible with this rule-based approach to not only modularise a model, but to combine it with existing ones. The model shows both the complexity of the network and its response under multiple different scenarios, showing a general fit with experimental literature even whilst highlighting several gaps in mechanistic process knowledge that have considerable implications on our understanding of AMPAR trafficking, and AMPAR functionality in general.

CHAPTER 5

Biological Models in Teaching

Both this and Chapter 7 explore the use of models as teaching tools. As I will argue, the use of modelling in teaching is underdeveloped and can provide a different way of comprehending complex biological concepts. Furthermore, models can be used in a creative manner by both learners and educators, which is not often afforded by more traditional teaching methods in biology.

This chapter outlines several models developed by students under my direct supervision. The students include STEM-oriented high school students from the Nuffield Programme and Honours year undergraduate students. Despite the large differences in biological background and experience with computational approaches, all of these projects successfully developed working initial models which could be expanded upon for academic research. I will begin by presenting each model in turn, showcasing the strengths of each and how well it addresses the brief that was set. I will continue with an analysis of the pedagogical benefits of learning using such modelling approaches, emphasising how the construction of a model from scratch encourages rapid acquisition of necessary knowledge and understanding of a biological phenomenon. This approach will be further exemplified in Chapter 7, which explores a specific method of model construction.

The Nuffield Scholar Programme

As part of the Scholar programme, the Nuffield Foundation offer post-16 students from disadvantaged backgrounds an opportunity to work in a STEM environment for up to six weeks (Cilauro & Paull, 2019). In this time they are expected to plan, conduct and evaluate their own investigation based upon the host's specialism and available materials. These students are paired to hosts based upon their interests and a brief personal statement.

I outline here the work of two students on this programme, Lewis Strachan and Kyle Robertson, who worked with me on their projects in 2017 and 2018 respectively.

A Biochemical Model of FMRP

Fragile X Syndrome (FXS), as we saw in Chapter 2, is a mono-genetic disorder arising from an abnormal amount of CGG repeats in the gene coding for Fragile X Mental Retardation Protein (FMRP). Most research and modelling of FXS focusses on the downstream effects of this loss, and the impact it has on the larger biochemical network. I noticed there was a gap in the modelling literature for a model which took FMRP to be its primary focus. The creation of such a model would be necessary for any serious investigation of FXS within AMPS. I thus conceived of a small summer project exploring the construction of a model of the immediate interactions FMRP has. This complemented my original modelling objectives as well as being an opportunity for a student to construct something novel and potentially useful in a research setting with minimal setup required.

Local FMRP Networks

FMRP activity is controlled by the phosphorylation of its S499 residue by ribosomal protein S6 kinase (S6K). This appears to be a necessary requirement for its function as a translational repressor, with S6K KO mice showing no FMRP activity (Narayanan et al., 2008). The phosphorylated FMRP is protected from ubiquitin tagging and thus proteasomal breakdown, which occurs rapidly after mGluR-LTD (Hou et al., 2006). The phosphorylation can be reversed through the activity of protein phosphatase 2A (PP2A) (Nalavadi, Muddashetty, Gross, & Bassell, 2012; Prieto, Folci, & Martin, 2020). Both S6K and PP2A are activity dependent enzymes modulated respectively by mGluR and TrkB signalling pathways (Inoki, Li, Zhu, Wu, & Guan, 2002; Jain & Bhalla, 2009; M. Li, Makkinje, & Damuni, 1996; Z.-H. Wang et al., 2019).

FRMP has two homologues: FXRP1 and FXRP2. These reliably form heterodimers with FMRP. Why these heterodimers form and for what purpose is currently not known (Dolzhanskaya, Merz, Aletta, & Denman, 2006). FMRP additionally complexes with CYFIP1 and eIF4E, which when present correlates with translational repression (Fernández, Rajan, & Bagni, 2013; Napoli et al., 2008). The CYFIP1-eIF4E bond is disrupted during mGluR and TrkB pathway activity, possibly to allow access to FMRP's phosphorylation site (DeRubeis et al., 2013; Napoli et al., 2008).

Model Construction

The model was built in COPASI (Hoops et al., 2006) by Lewis, after completing a self-directed tutorial exercise developed in the lab by Dr Melanie Stefan. Lewis was responsible for choosing the key molecules associated with FMRP mentioned above and generating a basic model of their interactions. This was done through literature searching and the use of the UniProt Database (UniProt Consortium, 2019). The final model consisted of 23 species and 27 reactions based on the interactions detailed above (**Table 3, Table 4**⁷³).

⁷³ DOQCS-derived reaction values are taken from similar binding reactions, and so can be treated as estimates.

<i>Species</i>	<i>Initial Concentration ($\mu\text{mol/l}$)</i>	<i>Source</i>
<i>P-FMRP</i>	0.4	Estimate
<i>FXRP1</i>	0.4	Estimate
<i>FXRP2</i>	0.4	Estimate
<i>FMRP-FXRP1</i>	0	-
<i>FMRP-FXRP2</i>	0	-
<i>CYFIP1</i>	0.6	Estimate
<i>eIF4E</i>	0.4	(Rau, Ohlmann, Morley, & Pain, 1996)
<i>CYFIP1-FMRP</i>	0	-
<i>CYFIP1-eIF4E</i>	0	-
<i>FMRP-eIF4E</i>	0	-
<i>CEF</i>	0	-
<i>mGluR</i>	0.3	DOQCS
<i>Glu</i>	*	AMPS
<i>mGluR-Glu</i>	0	-
<i>TrkB</i>	0.25	AMPS
<i>BDNF</i>	*	Fitzpatrick, 2017
<i>TrkB-BDNF</i>	0	-
<i>DP-FMRP</i>	0	-
<i>U-FMRP</i>	0	-
<i>PP2A</i>	1.11	Fitzpatrick, 2017
<i>S6K</i>	1.25	AMPS
<i>UE</i>	0.12	Estimate
<i>null</i>	1 (Fixed)	-

* Glu Expression: $\text{if}(\{\text{Values}[\text{Time}]\} \geq 100, 0, \text{if}(\{\text{Values}[\text{Time}]\} \geq 50, 1, 0))$

* BDNF Expression: $\text{if}(\{\text{Values}[\text{Time}]\} > 300, 0.1, \text{if}(\{\text{Values}[\text{Time}]\} > 200, 3.7, 0.1))$

Table 3 Species and concentrations of Lewis' model.

<i>Reaction</i>	<i>Equation</i>	<i>Rate</i>	<i>Source</i>
<i>FMRP-FXRP1 (forward)</i>	$P\text{-FMRP} + \text{FXRP1} \rightarrow \text{FMRP-FXRP1}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>FMRP-FXRP1 (backward)</i>	$\text{FMRP-FXRP1} \rightarrow P\text{-FMRP} + \text{FXRP1}$	0.15s^{-1}	DOQCS
<i>FMRP-FXRP2 (forward)</i>	$P\text{-FMRP} + \text{FXRP2} \rightarrow \text{FMRP-FXRP2}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>FMRP-FXRP2 (backward)</i>	$\text{FMRP-FXRP2} \rightarrow P\text{-FMRP} + \text{FXRP2}$	0.15s^{-1}	DOQCS
<i>CYFIP1-FMRP (forward)</i>	$P\text{-FMRP} + \text{CYFIP1} \rightarrow \text{CYFIP1-FMRP}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>CYFIP1-FMRP (backward)</i>	$\text{CYFIP1-FMRP} \rightarrow P\text{-FMRP} + \text{CYFIP1}$	0.15s^{-1}	DOQCS
<i>CYFIP1-eIF4E (forward)</i>	$\text{CYFIP1} + \text{eIF4E} \rightarrow \text{CYFIP1-eIF4E}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>CYFIP1-eIF4E (backward)</i>	$\text{CYFIP1-eIF4E} \rightarrow \text{CYFIP1} + \text{eIF4E}$	0.15s^{-1}	DOQCS
<i>FMRP-eIF4E (forward)</i>	$P\text{-FMRP} + \text{eIF4E} \rightarrow \text{FMRP-eIF4E}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>FMRP-eIF4E (backward)</i>	$\text{FMRP-eIF4E} \rightarrow P\text{-FMRP} + \text{eIF4E}$	0.15s^{-1}	DOQCS
<i>CEF1 (forward)</i>	$\text{CYFIP1-eIF4E} + P\text{-FMRP} \rightarrow \text{CEF}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>CEF1 (backward)</i>	$\text{CEF} \rightarrow \text{CYFIP1-eIF4E} + P\text{-FMRP}$	0.15s^{-1}	DOQCS
<i>CEF2 (forward)</i>	$\text{CYFIP1-FMRP} + \text{eIF4E} \rightarrow \text{CEF}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>CEF2 (backward)</i>	$\text{CEF} \rightarrow \text{CYFIP1-FMRP} + \text{eIF4E}$	0.15s^{-1}	DOQCS
<i>CEF3 (forward)</i>	$\text{FMRP-eIF4E} + \text{CYFIP1} \rightarrow \text{CEF}$	$49.999\text{l}/(\mu\text{mol*s})$	DOQCS
<i>CEF3 (background)</i>	$\text{CEF} \rightarrow \text{FMRP-eIF4E} + \text{CYFIP1}$	0.15s^{-1}	DOQCS
<i>mGluR-Glu (forward)</i>	$\text{mGluR} + \text{Glu} \rightarrow \text{mGluR-Glu}$	$16.8\text{l}/(\mu\text{mol*s})$	AMPS
<i>mGluR-Glu (backward)</i>	$\text{mGluR-Glu} \rightarrow \text{mGluR} + \text{Glu}$	10s^{-1}	AMPS
<i>TrkB-BDNF (forward)</i>	$\text{TrkB} + \text{BDNF} \rightarrow \text{TrkB-BDNF}$	$0.99998\text{l}/(\mu\text{mol*s})$	AMPS
<i>TrkB-BDNF (backward)</i>	$\text{TrkB-BDNF} \rightarrow \text{TrkB} + \text{BDNF}$	0.0499967s^{-1}	AMPS
<i>mGluR-Glu CYFIP1-eIF4E Dissociation</i>	$\text{CYFIP1-eIF4E} \rightarrow \text{CYFIP1} + \text{eIF4E}; \text{mGluR-Glu}$	$k_{\text{cat}} = 75\text{s}^{-1}$ $K_M = 40\text{ }\mu\text{mol/l}$	Estimate
<i>TrkB-BDNF CYFIP1-eIF4E Dissociation</i>	$\text{CYFIP1-eIF4E} \rightarrow \text{CYFIP1} + \text{eIF4E}; \text{TrkB-BDNF}$	$k_{\text{cat}} = 75\text{s}^{-1}$ $K_M = 40\text{ }\mu\text{mol/l}$	Estimate
<i>P-FMRP Dephosphorylation</i>	$P\text{-FMRP} + \text{PP2A} \rightarrow \text{DP-FMRP} + \text{PP2A}$	PP2A rate	AMPS
<i>DP-FMRP Phosphorylation</i>	$\text{DP-FMRP} + \text{S6K} \rightarrow P\text{-FMRP} + \text{S6K}$	S6K rate	AMPS
<i>DP-FMRP Ubiquitination</i>	$\text{DP-FMRP} + \text{UE} \rightarrow \text{U-FMRP} + \text{UE}$	$0.5\text{l}/(\mu\text{mol*s})$	Estimate
<i>U-FMRP Degradation</i>	$\text{U-FMRP} \rightarrow \text{null}$	1s^{-1}	Estimate
<i>FMRP Synthesis</i>	$\text{null} \rightarrow P\text{-FMRP}$	0.001s^{-1}	Estimate

Table 4 Reactions and parameters of Lewis' model.

The model volume was set at 0.09 fL, with concentrations and parameters derived from DOQCS (Sivakumaran et al., 2003) and other sources (Rau et al., 1996). When no source was available, estimates were made based on similar reactions and molecules. The rate of phosphorylation and dephosphorylation was written into the model as two rules. A threshold concentration of Glu-mGluR and BDNF-TrkB complexes respectively increased the rates of complex formation for FMRP-S6K and FMRP-PP2A. This stood in for the whole signalling cascades of these pathways which lay beyond the scope of this model.

Model Findings

Lewis tested three hypotheses with this model: Whether FMRP was centrally important to this system, whether the system is modulated significantly by FMRP regulation, and to what extent does activating S6K and PP2A alter the system. **Figure 69** shows the composition of the final model and an example output.

Lewis used their model to determine that FMRP was the rate-limiting factor in their biochemical network, that its loss prevents other complexes from being able to form, and that upstream signalling events from ligand binding episodes can have observable effects on system function. They identified their model was limited by a lack of experimental data on the biochemical network being modelled, but that the model's findings were still useful and informative in showcasing that the local reaction network surrounding FMRP is disrupted heavily when the molecule is less prevalent.

Learning Outcomes

Lewis was initially provided with some introductory material to Fragile X and what FMRP is, and tasked with designing a model based on the regulation of FMRP. From this slight sketch they set about researching the local binding partners of FMRP, how these were regulated by enzyme activity, and how these enzymes were modulated by ligand-receptor binding events upstream of their FMRP network. The model constructed requires a considerable breadth of research to characterise, especially when no prior knowledge was applicable. In part this aspect benefitted from the novelty of the situation – all the material was new, and there was no exam-related requirement to comprehend the theory fully.

Lewis' use of the model however shows a mature and conscious effort to engage with primary research articles to ensure their model features key players in the local regulation of FMRP. They set a series of hypotheses they wished to test with the

model themselves, and were able to critically analyse their results and place them into context, particularly in how they might help direct further research into FMRP regulation.

Lewis gained a practical and self-constructed experience in the conception, development and testing of a scientific idea from start to finish in an unpressured environment. That this made use of unfamiliar tools, with unfamiliar molecules, showed to them that even when derived from sparse data, models can have value. It also shows however how a lack of prior experience or engagement with something need not be a barrier to self-constructed learning, so long as there is guidance available. This is reminiscent of Bruner's scaffolding theory of learning (D. Wood, Bruner, & Ross, 1976), but I feel is more closely attuned in this case with Vygotsky's Zone of Proximal Development (Vygotsky, 1930-34/1978). In this theory, the learner has actions they are expert in, actions they can perform with the guidance of a teacher or fellow learner, and actions they are unable to perform.

In this instance, I worked to guide Lewis to using their skills in questioning, hypothesis generation, and searching for specific data that they already possessed from their education. Through unstructured, daily communication I was able to affirm their approaches, direct them to relevant research material, and troubleshoot issues with the software being used (COPASI). By the end of the project Lewis was much more confident in their own skills. They were able to demonstrate this to others in the form of their model and its findings, culminating in a poster presentation held for Nuffield Scholars to present their work.





Figure 69 Lewis' model of FMRP molecular interactions.

(A) A model plan of Lewis' model, showing the interaction of FMRP with immediate binding partners, and signalling pathways which affect the response of these interactions. (B) An example of one of Lewis' graphs, showing how glutamate binding to mGluR triggers a reduction in dephosphorylated FMRP and increase in CEF (CYFIP1-eIF4E-FMRP) complex concentrations.

Sonic Hedgehog Signalling

A member of the Hedgehog signalling pathway, Sonic Hedgehog (Shh) is a signalling molecule most closely associated with neural development in vertebrates. In this capacity it acts as a morphogen, controlling the genetic activation profile of neural cells across the neural plate in conjunction with Bmp signals acting at the opposite end (Ingham & McMahon, 2001).

Shh acts on the Patched receptor, which is commonly associated with another membrane bound protein, Smoothened. Patched in the absence of Shh inhibits Smoothened, meaning the binding of Shh to Patched is an inhibitory reaction which releases Smoothened. Free Smoothened then works to increase the activation of three Gli transcription factors (Gli 1 – 3) (Y. Chen et al., 2011). This is the canonical pathway, with many more pathways of action uncovered (Carballo, Honorato, De Lopes, & Spohr, 2018).

The underlying mechanisms and localisation are still poorly understood and are only just coming to light. Our interest in the molecule comes from the growing body of work exploring the role it plays post-development. In recent decades, the Shh signalling pathway has become implicated within a number of brain functions, with mRNA for Shh, Patched and Smoothened present in adult brain tissue (Coulombe, Traiffort, Loulier, Faure, & Ruat, 2004; Traiffort, Charytoniuk, Faure, & Ruat, 2002; Traiffort et al., 1999).

Of particular interest to my own work was the localisation of Shh signalling pathway components in both developing and mature hippocampal neurons (Masdeu, Bernard, Faure, Traiffort, & Ruat, 2007; Petralia, Schwartz, Wang, Mattson, & Yao, 2011), and some emerging evidence of a non-canonical interaction with glutamate receptors during synaptic plasticity (Sundberg, 2018).

Given these links to synaptic pathways, and a wealth of experimental work, I felt the development of a basic model of canonical Shh signalling was appropriate for a summer project. Unlike the FMRP model, this model underwent further development during an Honours Project. The Honours Programme offers final year undergraduates a chance to spend twelve weeks conducting a research project.

I will now detail both models and their development.

First Model

The initial model was developed by Kyle, a Nuffield Scholar who spent six weeks constructing the model. As previously with Lewis, I provided a basic level of understanding, training in COPASI and advice on how to search literature for data and information.

Their investigations into existing models found that all models were all focussed on single sections of the pathway, and no accessible functional model of the whole had yet been made. Using this as a starting point, Kyle focussed much of their time on determining the interaction partners and localisation of each of the major pathway components of the canonical model to construct an accurate usable model of canonical Shh signal transduction.

During the creation of this descriptive model form, Kyle identified the primary cilium as being a crucial structure in the pathway's function. This ubiquitous microtubular-based cellular structure acts as a vital semi-isolated environment necessary to allow signal transduction to occur (Wheway et al., 2018). Mutated primary cilia are associated with major disruptions to canonical Shh signalling. Kyle therefore endeavoured to include this area as a compartment in the model to assess whether its absence would also result in a failure of signalling in the final model. This was helped by discussions with faculty members (Dr Thomas Theil and Mr Kai Boon Tan) with research interests in the role of Gli3 and the primary cilium in relation to Shh, who were able to direct Kyle to the key mechanistic elements and to papers that supported this.

The final model consists of 21 species and 28 reactions (**Table 5, Table 6**), moving from Shh binding to the Ptc receptor through to gene transcription (**Figure 70**). Through the course of researching the pathway, it became apparent that the scarcity of accessible reaction parameters and concentrations required a relative model rather than an absolute one, and so parameters were scaled by Kyle to match with evidence from experimental and modelling outcomes.

Shh input is determined by a cyclic set of reactions that release a regular pulse of Shh. Each pulse led to a similarly pulse-like changes in the concentration of Smoothened, Patched, and active Gli, which in turn led to pulsatile changes in gene transcription. Increasing the size of the primary cilium reduces the maximal rate of gene transcription (**Figure 70**).

Learning Outcomes

Viewing this model from a pedagogical perspective, the model shows a good understanding of the currently accepted mechanism of canonical Shh signalling. Its output is qualitatively consistent with the canonical pathway dynamics, and it makes a solid case for the necessity of the primary cilium for the pathway to function at all. It also demonstrates a solid grasp of the key important molecular components of the pathway.

Unlike in Lewis' model, the model encompasses a major signalling pathway which has been extensively researched and characterised. Much like Lewis however, Kyle found that the ability to model the network biochemically was restricted by a lack of kinetic information in the available experimental literature. This may be in part why Kyle's attempts to find pre-existing models only uncovered purely descriptive or narrowly focussed models (e.g. Cohen, Page, Perez-Carrasco, Barnes, & Briscoe, 2014).

Kyle's model is therefore one of the first readily accessible computational models of canonical Shh signalling made (**Figure 70**). This is both surprising and valuable, as it shows model construction of scientifically-relevant models is both achievable by learners but also necessary for our characterisation of these pathways. As in Lewis's investigation, Kyle had to search through many different primary sources to discover concentrations, kinetic rates, or at least data which would help to parameterise the model. Kyle focussed much more on the model construction than its potential use, wanting to ensure the model reflected the Shh pathway behaviour as much as possible.

Kyle adapted some of the COPASI tutorial material provided to generate a cyclic entry of Shh into the system. This allowed them to explore the temporal relationship of the pathway, as well as a way of adjusting Shh signalling frequency for other users. Kyle also showed that changing the primary cilium compartment volume impacted on the ability of the pathway to trigger gene transcription, and that active recruitment of cAMP to this area is necessary for the pathway to function effectively.

This area was of particular interest to Kyle, and in meeting with faculty was able to discuss as a researcher, not a student. This became apparent to Kyle when they expected the PhD students and professors to know the answers, and found that their own knowledge was in places superior, and that they were interested in the




Figure 70 Kyle's model of Shh signalling.

(A) Kyle's model focusses on canonical Shh signalling via the Ptc1 receptor, and highlights the necessity of the primary cilium in allowing this pathway to function. There are two compartments: cytosol (orange) and cilium (pink). (B) Some examples of Kyle's chosen graphical outputs, highlighting the cyclic pulsatile feature of the model, as well as the effect of cilia size on gene synthesis.

possibility of using this model to help them understand the biochemical mechanics of the pathway themselves.

<i>Species</i>	<i>Compartment</i>	<i>Initial concentration</i>
<i>Shh</i>	Cytoplasm	0
<i>Ptc1</i>	Cytoplasm	0.75
<i>Ptc1(l)</i>	Cytoplasm	0
<i>Smo</i>	Cytoplasm	0
<i>Smo(l)</i>	Cytoplasm	0.35
<i>Smo(p)</i>	Primary Cilia	0
<i>SuFu</i>	Primary Cilia	0.8
<i>SuFu(A)</i>	Cytoplasm	0
<i>cAMP</i>	Cytoplasm	0
<i>PIP3</i>	Primary Cilia	0.2
<i>PIP3(l)</i>	Primary Cilia	0
<i>PKA</i>	Cytoplasm	1
<i>PKA(c)</i>	Cytoplasm	0
<i>β-arrestin</i>	Primary Cilia	0
<i>Kif3a</i>	Primary Cilia	0
<i>Gli(FL)</i>	Cytoplasm	1
<i>Gli(A)</i>	Cytoplasm	0
<i>Gli(R)</i>	Cytoplasm	0
<i>Shh Target Genes</i>	Cytoplasm	0
<i>A</i>	Cytoplasm	1
<i>B</i>	Cytoplasm	9

Table 5 A list of species in Kyle's model. Species are split across two compartments.

<i>Reaction Name</i>	<i>Reactions</i>	<i>Kinetic Rate</i>
<i>Smo Inhibition</i>	$\text{Smo} + \text{Ptc1} \rightarrow \text{Smo(l)}$	0.1
<i>Ptc1 Inhibition</i>	$\text{Shh} + \text{Smo(l)} \rightarrow \text{Ptc1(l)} + \text{Smo}$	0.1
<i>Repression of Gli</i>	$\text{SuFu(A)} + \text{Gli(FL)} \rightarrow \text{Gli(R)} + \text{SuFu(A)}$	0.1
<i>Activation of SuFu</i>	$\text{SuFu} + \text{PKA(c)} \rightarrow \text{SuFu(A)}$	0.1
<i>Localisation of cAMP</i>	$\text{PIP3} \rightarrow \text{cAMP} + \text{PIP3}$	0.1
<i>Localisation of β-arrestin & Kif3a</i>	$\text{Smo(p)} \rightarrow \beta\text{-arrestin} + \text{Kif3a} + \text{Smo}$	0.1
<i>Suppression of PIP3</i>	$\beta\text{-arrestin} + \text{PIP3} \rightarrow \text{PIP3(l)}$	0.1
<i>Activation of Genes</i>	$\text{Gli(A)} \rightarrow \text{Shh-Target-Genes}$	0.1
<i>cAMP</i>	$\text{cAMP} \rightarrow$	0.1
<i>Removal/Degradation</i>		
<i>Repression of Genes</i>	$\text{Gli(R)} + \text{Shh-Target-Genes} \rightarrow$	1
<i>Activation of Gli</i>	$\text{Gli(FL)} + \text{Kif3a} \rightarrow \text{Gli(A)}$	0.1
<i>Ptc1(l) Degradation</i>	$\text{Ptc1(l)} \rightarrow$	0.1
<i>SuFu Degradation down</i>	$\text{SuFu} + \text{Gli(A)} \rightarrow \text{Gli(A)}$	0.1
<i>Active Path</i>		
<i>Transcription of protein</i>	$\text{Shh-Target-Genes} \rightarrow \text{Gli(FL)} + \text{Ptc1}$	$\text{if}(\{\text{[Shh]}\} < 0.2, 2, 0)$
<i>β-arrestin reuptake</i>	$\beta\text{-arrestin} \rightarrow$	0.1
<i>Repression of SuFu(A)</i>	$\text{SuFu(A)} + \text{Smo(p)} \rightarrow \text{SuFu} + \text{Smo(p)}$	0.1
<i>PIP3 reactivation</i>	$\text{PIP3(l)} \rightarrow \text{PIP3}$	0.02
<i>Cyclic_System-1</i>	$\text{A} + \text{B} \rightarrow 2 * \text{A} + \text{B}$	0.008
<i>Cyclic_System-2</i>	$\text{A} \rightarrow$	0.15
<i>Cyclic_System-3</i>	$\text{B} \rightarrow 2 * \text{B}$	0.05
<i>Cyclic_System-4</i>	$\text{A} + \text{B} \rightarrow \text{A}$	0.008
<i>Cyclic_Release</i>	$\text{A} \rightarrow \text{A} + \text{Shh}$	0.19
<i>Shh Reuptake</i>	$\text{Shh} \rightarrow$	1.8
<i>Activation of PKA</i>	$\text{PKA} + \text{cAMP} \rightarrow \text{PKA(c)}$	0.35
<i>Deactivation of the PKA complex</i>	$\text{PKA(c)} \rightarrow \text{PKA} + \text{cAMP}$	0.27
<i>Phosphorylation of Smo</i>	$\text{PKA} + \text{Smo} \rightarrow \text{PKA} + \text{Smo(p)}$	0.1
<i>Dephosphorylation of Smo</i>	$\text{PKA} + \text{Smo(p)} \rightarrow \text{PKA} + \text{Smo}$	0.1
<i>Kif3a Relocation</i>	$\text{Kif3a} \rightarrow$	$\text{if}(\{\text{[Kif3a]}\} \geq 5, 3, 0.1)$

Table 6 Reaction rate parameters for Kyle's model.

Second Model

Given that the initial model construction had uncovered a lack of biochemical models on Shh, I felt it would be useful to keep building the same pathway and further develop it into a well-characterised format usable by others. In addition, I had come across research that pointed to Shh influencing the expression of glutamatergic receptors in the brain, and wished to see these findings could be modelled (Sundberg, 2018). My supervisor (Dr Stefan) and myself offered an Honours project in 2019 titled “Computational Modelling of the Sonic Hedgehog (SHH) Pathway”.

The idea was to use Kyle’s model as a springboard to develop a more rigorous model that may be added to future maximal models. James Schofield, a 4th year Anatomy and Development Honours student, was chosen to undertake this project.

James’ first aim was to get up to speed with Shh literature, something which was already familiar from courses taken during their undergraduate studies. James then began two assessments in parallel: one of the literature around the claim of Shh’s role in neural plasticity, and a critical analysis of Kyle’s model. From this, James decided that the best approach was to rebuild Kyle’s model using it as a reference only. This was partially due to a lack of written rationale behind some of the decisions about molecular function made by Kyle, and partially to allow for practice in generating a model using COPASI.

Communication with some of the original researchers of the neural plasticity findings suggested that the work in this area was too nascent to make up a major part of the model (Dr Gonzalo Sanchez, personal communication). It was clear however from related literature that the canonical pathway plays a not insignificant role in multiple brain regions, and that one of the major hindrances to further understanding was a lack of knowledge surrounding the biochemical dynamics.

It is with this problem in mind that James began to construct the model. Their personal aim was to generate a fully characterised, working model of canonical Shh biochemical signalling which had parameters fully derived from experimental data. This was ambitious, but not something I discouraged, as this would be the ultimate aim of the model and would also help to confirm the presence of any ‘gaps’ in the experimental literature more quickly and readily.




Figure 71 James' Shh Model

(A) An overview of James' model. An upgrade of Kyle's model, it includes 3 different compartments (soma (pink), cilium (orange), nucleus (green)), and many more of the regulatory elements of the signalling pathway.

(B) Example figure from James' dissertation. In the absence of Shh, Ptch1 suppresses Smo activity and ciliary localisation (Tukachinsky et al., 2010). Suppressor of Fused (SuFu) translocates to the primary cilium, where it is phosphorylated by both protein kinase A (PKA) and glycogen synthase kinase-3 beta (GSK3β; Chen et al., 2011). Phosphorylated SuFu binds to Gli2 or Gli3 transcription factors (Tukachinsky et al., 2010).

<i>Biochemical Species</i>	<i>Compartment</i>	<i>Concentration (nM)</i>	<i>Source</i>
<i>GPR161</i>	Soma	0.1 – 5.0	Estimate
<i>Ptch1-Smo(I)</i>	Soma	0	Estimate
<i>Shh</i>	Soma	3	(Saha & Schaffer, 2006)
<i>Smo</i>	Soma	0.1 – 15	Estimate
<i>SPOP/Cul3</i>	Soma	0.1	Estimate
<i>SuFu</i>	Soma	100	(Tukachinsky, Lopez, & Salic, 2010)
<i>βTrCP/Cul1</i>	Soma	0.1	Estimate
<i>Proteins</i>	Soma	0	
<i>Gli2/3(A)</i>	Nucleus	0	
<i>Gli2/3(R)</i>	Nucleus	0	
<i>AC</i>	Primary Cilium	0.1 – 5.0	Estimate
<i>ATP</i>	Primary Cilium	525	(Surin et al., 2014)
<i>cAMP</i>	Primary Cilium	0	
<i>CK1</i>	Primary Cilium	0.1	Estimate
<i>EvC-EvC2</i>	Primary Cilium	0.1 – 10	Estimate
<i>Gli2/3</i>	Primary Cilium	5.8	(Saha & Schaffer, 2006)
<i>Gli2/3(P)</i>	Primary Cilium	0	Estimate
<i>GPR161(A)</i>	Primary Cilium	0	Estimate
<i>GRK2</i>	Primary Cilium	50	(Tesmer, Lennarz, Mayer, & Tesmer, 2012)
<i>GSK3β</i>	Primary Cilium	49.1	(Padala et al., 2017)
<i>Kif7(P)</i>	Primary Cilium	500	(M. He et al., 2014)
<i>PKA</i>	Primary Cilium	590	(Saucerman, Brunton, Michailova, & McCulloch, 2003)
<i>PKA(I)</i>	Primary Cilium	590	(Saucerman et al., 2003)
<i>PKA(I/C)</i>	Primary Cilium	0	
<i>PKA(S)</i>	Primary Cilium	0	
<i>PP2A/PPFIA1</i>	Primary Cilium	240	(Padala et al., 2017)
<i>Ptch1</i>	Primary Cilium	0.3	(Saha & Schaffer, 2006)
<i>Smo(A)-EvC-EvC2</i>	Primary Cilium	0	
<i>Smo(P)</i>	Primary Cilium	0	
<i>SuFu(P)</i>	Primary Cilium	0	
<i>SuFu(P)-Gli2/3</i>	Primary Cilium	0	
<i>SuFu(P)-Gli2/3-Kif7</i>	Primary Cilium	0	
<i>SuFu(P)-Gli2/3-Kif7(P)</i>	Primary Cilium	0	

Table 7 Species list for James' model.

<i>Name</i>	<i>Reaction</i>	<i>Global Quantity</i>	<i>Source</i>
<i>Smoothered Inhibition</i>	Ptch1 + Smo → Ptch1-Smo(I)	1.0	-
<i>GPR161 Activation</i>	GPR161 → GPR161(A)	1.0	-
<i>AC Activation</i>	ATP + AC → cAMP + AC	0.5 – 1.0	-
<i>PKA Activation 1</i>	PKA(I) + 2 * cAMP → PKA(I/C)	1.0×10^{-3} – 1.0	-
<i>PKA Activation 2</i>	PKA(I/C) + 2 * cAMP → PKA	1.0×10^{-3} – 1.0	-
<i>SuFu Phosphorylation</i>	SuFu + GSK3 β + PKA(S) → SuFu(P) + GSK3 β + PKA(S)	0.5	(Vertommen et al., 2000)
<i>SuFu Binding</i>	SuFu(P) + Gli2/3 → SuFu(P)-Gli2/3	1.0	-
<i>SuFu Loss</i>	SuFu →	1.0	-
<i>Gli2/3 Phosphorylation</i>	SuFu(P)-Gli2/3 + PKA + CK1 + GSK3 β → SuFu + Gli2/3(P) + PKA + CK1 + GSK3 β	0.5	(Vertommen et al., 2000)
<i>Gli2/3(R) Production</i>	Gli2/3(P) + β TrCP/Cul1 → Gli2/3(R) + β TrCP/Cul1	2.0×10^{-4}	(Lai, Robertson, & Schaffer, 2004)
<i>Gli2/3(R) Degradation</i>	Gli2/3(R) →	1.5×10^{-4}	(Lai et al., 2004)
<i>Shh Binding</i>	Shh + 2 * Ptch1 →	1.5×10^{-3}	(French & Lauffenburger, 1996)
<i>Smoothered Phosphorylation</i>	Smo + CK1 + GRK2 → Smo(P) + CK1 + GRK2	0.5	(Vertommen et al., 2000)
<i>Smoothered Localisation</i>	Smo(P) + EvC-EvC2 → Smo(A)-EvC-EvC2	1.0	-
<i>Kif7 Binding</i>	SuFu(P)-Gli2/3 + Kif7(P) → SuFu(P)-Gli2/3-Kif7(P)	1.0	-
<i>Kif7 Translocation</i>	SuFu(P)-Gli2/3-Kif7(P) + PP2A/PPFIA1 → SuFu(P)-Gli2/3-Kif7 + PP2A/PPFIA1	1.0	-
<i>Gli2/3 Activation</i>	SuFu(P)-Gli2/3-Kif7 → SuFu + Gli2/3(A) + Kif7(P)	1.0	-
<i>Gli2/3(A) Degradation</i>	Gli2/3(A) + SPOP/Cul3 → SPOP/Cul3	1.5×10^{-4}	(Lai et al., 2004)
<i>Gene Expression</i>	Gli2/3(A) → Gli2/3(A) + Proteins	8.0×10^{-3}	(Lai et al., 2004)
<i>Protein Loss</i>	Proteins →	1.0×10^{-5} – 1.0	-

Table 8 Reaction list for James' model. – indicates estimated values.

James' final model consisted of 33 species and 20 reactions (**Table 7, Table 8**) across three compartments. The model is a logical one, focussing on the qualitative aspect of Shh as a switch for the signalling network (**Figure 71**). Whilst the model did not meet with the expectations James had from the start, it did offer some valuable lessons in the progression of scientific discovery. Modelling is quick⁷⁴, and able to investigate multiple aspects of a system rapidly. This however requires the system to be well-characterised kinetically. Despite the large experimental body of work, Shh is poorly characterised in this respect, which was highlighted quite clearly by James' exhaustive search through the literature. In this situation, modelling acts just like any other methodology, in that it generates a small amount of new information that can be slowly iterated on over time.

In this regard, James' model is a success. It points out that as theoretically rich and experimentally wide-reaching as investigations on Shh as a signalling pathway and its outcome are, they are less able to detail the exact mechanisms by which the pathway functions molecularly. As a result, understanding how Shh works in tissues where it is less well-studied becomes harder to do, as researchers cannot easily see linkages with localised biochemical networks and have to do a lot of extra work to see how Shh acts on phenomena.

Better kinetic characterisations would allow the model to better pinpoint the likely important (or even causal) interactions the pathway has within different biochemical environments, or the modifications necessary to match the observed outcomes. This model shows exactly which areas are deficient, providing fruitful avenues for experimental research which will help generate a model with greater predictive power.

James' experience with modelling overall was mixed. Whilst they were able to complete a model, and use it to generate useful information, it was not the model they set out to complete. The frustration of this became more apparent towards the end of the project, reflecting the heightened worries of failure common in undergraduate students exposed to research. In James's case, this may be compounded by their decision to not use Kyle's model, but build one which in some ways was less characteristic of how we have seen biochemical models function thus far in this thesis.

This element of frustration that can come with computational modelling is to my knowledge not something that has been investigated specifically. However, the

⁷⁴ Well, quickish.

problem of frustration in computer programming courses offers some insights into how in the future this experience can be avoided. The most noticeable one is that in programming students, the pairing with a peer seems to help considerably in reducing the level of frustration within learners (Mcdowell, Werner, Bullock, & Fernald, 2003; Nagappan et al., 2003; Williams, Yang, Wiebe, Ferzli, & Miller, 2002). Only Lewis had a working context where a similarly inexperienced peer was working on model building, which formed an accidental paired programming environment.

The second element is structure. The modelling projects all possess a broad aim (build a model), and a specific area of interest prior to the student becoming involved. Once the student takes ownership of the project however, skills or techniques were taught on an 'as-and-when necessary' basis. Whilst this gave the student more control over the direction of their modelling, it may have also isolated them from asking for advice or knowing what tools are available to solve a problem they had.

Thinking back at how these models may offer a space for a student to understand their expertise however, this having to fail and understand that failure before asking for help conforms to the ethos behind zones of proximal development (Vygotsky, 1930-34/1978). I do wonder however if it becomes too high of a cognitive load, and therefore does more harm than good (Kirschner, Sweller, & Clark, 2006). Coupled with situations where no peers are present working on computer modelling, this is likely to be exacerbated. This is a challenge that needs to be addressed in future analysis of this project-based modelling approach.

Blackwell (2002) notes that one key difficulty when introduced to programming is the "loss of direct manipulation". Writing code is not visual in the sense that changes made have abstract, not immediately noticeable effects. The feedback to the programmer is temporally stretched, and being able to link feedback when it does appear is a heavy cognitive load. Similarly, in modelling there are many settings and model components that make retaining a sense of all the changes made difficult without implementing a strong practice of version control, code notation, and record keeping.

Being able to receive immediate feedback on program writing has been a large part of new methods of teaching code, such as CodeCombat (<https://codecombat.com/>), and the Scratch programming language (scratch.mit.edu; Maloney, Resnick, Rusk, Silverman, & Eastmond, 2010; Resnick et al., 2009). Many programming distributions also now come with integrated development environments that trace and flag errors

in real-time, such as Spyder (<https://www.spyder-ide.org/>) or even do much of the coding automatically, such as Jedi (<https://jedi.readthedocs.io/en/latest/>).

COPASI has, in its most recent versions, begun to flag potentially incorrect and model-breaking user edits with visual flags (http://copasi.org/Support/Change_History/COPASI_4_23_Build_184/) but is still not equivalent to a fully reactive system. One development then, would be an introductory modelling software in which both progress and mistakes were visually rendered in real-time, reducing cognitive load through affording direct manipulation. Whilst not the reason for Chapter 7, the software used there is quite close to this idea.

Constructing Knowledge Through Modelling

A modelling project as we have seen can take on different forms that are both subject and learner constrained. In each instance a broad biomedical concept was identified by myself and then constrained by the learner's skillset, tools, and interests. None of these constraints were demonstrably restrictive, and arguably were a necessary feature to enable the level of understanding acquired during these projects.

Model construction is I believe a powerful example of constructivist practice, where the learner builds their own understanding of a concept based on what they can do, as well as what they find personally interesting. In these scenarios the educator acts as a mentor or guide, directing away from extraneous or ambiguous scholarship and ensuring the learner has enough support with the tools or skills being used during knowledge acquisition and dissemination (Palincsar, 1998).

Virginia Richardson's insightful article on constructivist pedagogy (Richardson, 2003) lists five characteristics to constructivist pedagogy (paraphrased here):

1. *Attention to the individual*
2. *Facilitation of group dialogue*
3. *Introduction of formal domain knowledge into conversation*
4. *Provision of opportunities for students to determine, challenge, change or add to existing understanding*
5. *Development of students' metawareness [sic] of their own understandings and learning processes.*

(Richardson, 2003; p1626)

I would argue that the development of models by all the students described here (and indeed in Chapter 7) met these characteristics. The shape and direction of the model was determined almost wholly by the student, but still within a structured task format (to build a model of a "thing"). This student-led shaping was encouraged through dialogue and an emphasis on their expertise within the "thing" they were modelling. Furthermore, the writing of final reports and presentations helped to reinforce the amount of personal understanding and skills gained from the experience. This was achieved by the emphasising through these methods of explaining the model that each finding was discovered through their own decisions and experimental design.

Each learner was left largely to develop their own way in working, including the location of their work and the duration of any work on the project in a day. Furthermore, I made it clear from the onset that my role would be to supervise the project, but not to direct it nor determine its end points. Aside from the initial introduction to both concept and tools, the direction was entirely up to the learner.

In this way each learner was able to take ownership of their project from the onset and be explicitly aware that there was no consequence to any failure or setback. Additionally, and in part by accident, every project involved concepts in which neither the lab nor myself were experts in. We are experts in modelling though. This may have played a bigger role than anticipated in ensuring each learner felt comfortable in shaping their projects as much as they did. By being aware that their expertise would be equal, if not surpass, their supervisor's about the concept, the learners gained a greater confidence in their own knowledge gathering skills and subsequent understanding.

This does not remove the pressure of needing to understand things correctly. Having to impart the reasons why they have made the choices they have, or why a particular model outcome is interesting is still necessary. But the learner can now no longer presume any implicit understanding by the educator, and must present things clearly and with evidence to persuade the educator of the same perspective. This was especially useful for the various written reports each had to carry out at the end of their projects. In all cases, these assessments are not marked by someone knowledgeable about biochemical modelling, which required each learner to therefore work to both introduce the biological concepts as well as the methodology in a clear manner. Both had to be understood by the learner in order to successfully show the experimental work they conducted, and to thus make inferences from the data collected.

We have seen up to this point models being used for the purposes of academic knowledge enhancement. The student models outlined here⁷⁵ however have shown that the capacity for individual growth of understanding to be accelerated in the process of model construction and subsequent interpretation of its outputs. This requires multiple or prolonged interaction, however. A single experience of biochemical modelling software is not sufficient to understand how it works, or to

⁷⁵ And indeed, my own.

really appreciate how the model is designed, why it is that way, and what would be *possible* with some adjustments.

I set out to see if there was a way to introduce the idea of model generation or use in a way which was longer (both in total time and number of times), less constrained, and more creative. Importantly however, it still needed to be something which could be used in a wide range of pedagogical settings, and be appropriate for tertiary teaching. In Chapter 7 I outline my investigations in this area, specifically with the use of Minecraft, and how the gamification of modelling amply meets the above requirements.

CHAPTER 6

Validation through Collaboration

Before we get to Minecraft however, I want to present an idea which I think ties in with the discovery of how students engage with model construction and development. As we saw in Chapter 3, modelling requires careful verification to ensure that what is built is representative of the systems being modelled. A model is built from the works of many other people, and a modeller must curate this work to form a coherent model.

During this process, a modeller may neglect to record the source of a parameter, or the edits needed to reproduce published graphical outputs of the model. Over time, the experimental evidence available about a system may change. So too can the software available to run models become obsolete. These all require models to be tested to check the original claims of the model and its modeller still hold true.

Standardisation initiatives such as SBML (Hucka et al., 2003), SED-ML (Waltemath et al., 2011), MIRIAM (Novère et al., 2005) begin to make models easier to reproduce and integrate with one another, databases such as BioModels (Malik-Sheriff et al., 2020) and ModelDB (McDougal et al., 2017) allow models to be shared and reused more easily, and kinetic databases like DOQCS (Sivakumaran et al., 2003) and SABIO-RK (Wittig et al., 2012) offer curated resources to source and update model parameters.

Despite these excellent initiatives however, the availability of absolute values of molecular concentrations, kinetic rate parameters, and fine-grained localisation data are sparse. Whilst some of this data is difficult to acquire, lots can be approximated quite easily. These experiments are largely uninteresting or irrelevant to the investigations of experimenters however, and so are rarely published.

The following paper addresses ways of resolving both horns of this dilemma, through the development respectively of reproducibility audits and an incentivisation database of modeller-requested experiments.

The paper is a submitted article to *Neuroinformatics* (currently under review). I am co-author of the paper, and contributed equally to the writing of it.

Validation through Collaboration

Encouraging team efforts to ensure internal and external validity of computational models of biochemical pathways

Richard Fitzpatrick · Melanie I. Stefan

Received: date / Accepted: date

Abstract Computational modelling of biochemical reaction pathways is an increasingly important part of neuroscience research. In order to be useful, computational models need to be valid in two senses: First, they need to be consistent with experimental data and able to make testable predictions (external validity). Second, they need to be internally consistent and independently reproducible (internal validity). Here, we discuss both types of validity and provide a brief overview of tools and technologies used to ensure they are met. We also suggest the introduction of new collaborative technologies to ensure model validity: an incentivised experimental database for external validity and reproducibility audits for internal validity. Both rely on FAIR principles and on collaborative science practices.

Keywords Systems biology, reproducibility, computational modelling, model validity

PACS 82.39.-k · 82.39.Rt

Mathematics Subject Classification (2010) 92-08 · 92C20 · 92C42 · 92C45

R. Fitzpatrick, M.I Stefan
Centre for Discovery Brain Sciences
University of Edinburgh
Edinburgh, UK
E-mail: melanie.stefan@ed.ac.uk

M.I. Stefan
ZJU-UoE Institute, Zhejiang University
Haining, China

1 Introduction

The large number of molecules and interactions underpinning most biological phenomena call for *in silico* approaches to understand biochemical networks [1]. This is especially true for neuroscience, where the interpretation of a molecular signalling network can have major implications on translational approaches to diseases and disorders. A good computational model makes testable predictions which can be used to narrow down the number of experimental investigations required to reach an understanding of a given phenomenon [2]. With the newer multistate computational models and tools [3–9] we can see the impact of modifying one aspect of a molecule's function on all others, and how that affects the biochemical network as a whole, without needing to construct many different computational models, or run multiple *in vivo* or *in vitro* experiments.

These powerful aspects of modelling have not reached their full potential within neuroscience. This may stem in part from a lack of clarity on how our modelling approaches represent the biological mechanisms they claim to simulate [2,10] and the soundness of the models themselves. Two major questions can be asked about the validity of a computational model:

1. How can we be sure that the model is representative of *in vivo* states?
2. How do we know the model is reliable?

The first question relates to the external validity of a model (how well the model fits with experimentally knowable data), the second relates to its internal validity (whether the model is soundly and consistently constructed).

This commentary provides two possible pathways to answer these questions about model validity. Both require a greater level of collaboration, both between biochemical modellers and between modellers and their experimental counterparts. We believe by fostering such connections models will be better utilised, better parameterised, and embedded more into the driving of neuroscientific inquiry.

2 External validity: Comparing a computational model to experimental data

Computational models of biological systems are important tools: They can synthesise the current state of knowledge about a biological process into a coherent system. Using models, we can explore overall behaviours of a biological system that would be impossible to predict from just examining its component parts [11]. Using computational models, we can quickly test a large number of possible scenarios. They are therefore especially useful for generating hypotheses about a system, and making testable predictions about its behaviour.

The predictive power of a computational model relies on the model being an accurate (enough) representation of biological reality. Modellers rely on experimental data to construct and constrain the model. Once a model is completed, experiments are needed to validate models, test model predictions, or select from competing models of the same process.

A biochemical modeller wants their model parameters to closely resemble the situation in vivo. This requires binding constants and concentrations specific to a specific cell type or functional component such as a dendritic spine. There are databases of biochemical parameters [12–15], but at this point they suffer from incomplete coverage, especially when it comes to data on signalling pathways.

Compounding this data scarcity is that much of these experimentally derived sources for reaction constants and concentrations come from decades-old research. This work is often of excellent quality, but does not cover many more recently discovered molecules and interactions. The urgent need for new experimental data for models is being approached in interesting ways, with frameworks such as FindSim [16] encouraging the integration of multiscale models with experimental datasets, and the FAIR initiative improving the extraction of data from published studies to improve discovery, standardisation, and enable the re-use of this data [17]. These cannot generate data which is not there though. As our in vivo techniques have improved, there has been a major shift in analysing the function of molecules in situ. The actual mechanisms that underpin function are often included quite late in how we currently construct and perceive biological theory [18,19]. This contrasts with biochemical modellers, who are almost always concerned with mechanisms [20,21].

Modellers do have an array of tools to work around this problem. We can run parameter sensitivity analyses to pinpoint the parameters that matter to a reaction network [22], and then estimate values that fit with experimental outcomes. We can assess the robustness a reaction network, determining sloppy parameters, whose 'true' value does not matter much to the behaviour of the model [23]. Indeed it may be enough to focus experimental efforts on a few parameters that the model is most sensitive to, instead of measuring every single model parameter, which would be both experimentally costly and, for some parameters, unnecessary [23,24].

And yet, as we move to larger and more complex models, the question of how well these sensitivity and parameter identification analyses scale to larger contexts is still open [25]. Parameterisation techniques for large models have evolved rapidly to reduce the intractable computational load and to accommodate the sparsity of absolute datasets [26], but they do not eliminate the need to experimentally determine some key parameters.

Once built, a model can be tested against experimental data to establish its validity. A good model can make predictions that go beyond the currently available data, and may even go beyond currently possible experimental techniques - this is indeed quite common in fields such as physics (see for instance [27–31].) This means that a modeller may predict the outcome of an experiment, but may never be able to conduct the experiment itself, or even see it conducted.

Taken together, experiments can provide useful information at of computational modelling, from early on in model development to years after a model has been completed and distributed.

We suspect all modellers have a “wish list” of experiments that would improve and accelerate their model development, or test model predictions. But not all modellers have access to the necessary infrastructure and skills. Finding experimental collaborators is not always easy: There is not necessarily complete overlap between the experiments that would be informative to a computational modeller and the experiments that interest the experimentalist.

3 Incentivised experimental database

As we have seen, there is a gap between computational models and the experimental data needed to both constrain those models and test their outcomes. This is partly because existing data is not always published and shared, but partly also because some experiments have just never been conducted.

What is needed is a way of incentivising these experiments, to persuade our experimentalist colleagues that there is some benefit to them for carrying them out. We propose one such way is to take some lessons from the past. Could we not present our experimental wishlist, specifying the data we need to complete or check our models, and offer a cash incentive for providing this data?

Offering cash rewards for solving scientific problems is not without precedent. Historically, Challenge Prizes drove major advancements in problems of navigation (e.g. The UK Longitude Act of 1714 [32]) and aviation (e.g. the Orteig Prize [33]). The Millennium Prize offers \$1 million for the solution of any of seven stated mathematical problems [34]. Even more recently, foundations such as Nesta offer considerable sums of money for the solving of defined problems within a range of different fields [35].

What we are suggesting is not quite financially on the same scale as this but captures some of that spirit; of incentivising innovation to accelerate improvements in biochemical modelling. Our “problem” is that we have limited access to specific biochemical data necessary to accelerate the construction and testing of complex dynamical models of biochemical systems. Our solution is to reward experimenters who “solve” parts of this problem through the provision of this data.

We envisage this working thusly. Modellers submit a wishlist of experiments to a database, with explicit instructions on the biological background, the model, the data needed, and (if available) a suggested experimental design. These experiments are sorted into categories related to difficulty and the experimental methodology required to implement them. The relative difficulty or complexity of the experiment is linked to a cash reward, which not only compensates for the time and resources used but provides extra income for the lab to continue their own research. These “microgrants” would be split into two components – money up front for the experiment, with the bonus provided upon submission of raw data and documentation following FAIR principles (regardless of the nature of the outcome). The dataset publication would also include authorship and contribution information of all experimental collaborators, as well as a link to the original data request and the model it stemmed from, thus giving credit and facilitating provenance tracking for model parameters.

These ‘arranged’ collaborations may even prove more fruitful in connecting researchers exploring the same phenomena through different approaches. This leads to our second parallel intention for this database: prediction testing.

As initiatives like FindSim [16] emphasise, models ideally exist in a dynamic cycle with experimental research. Model outcomes produce predictions that are tested experimentally, which provides data to update the model to drive further predictions. This ideal scenario is rarely achieved however, and what we have is a mostly decoupled system where model predictions are not seen by experimental researchers, or only discovered after convergently reaching the same

outcome. Our database would encourage modellers to post the major predictions of their models which can then be validated by experimental work. This approach allows for a gradual and visible increase in the utility of modelling alongside experimental work. As models receive higher fidelity parameter sets, the predictions made will have more weight and power to guide real conceptual breakthroughs.

This provides another use of the modeller “wish list”, supplying experimenters with readily available predictions and a clearly defined direction for potentially fruitful future research. The results of these investigations are importantly just as valuable if they contradict the model as when they agree, leading in each case to publishable findings and model enhancement.

Thus, the incentivised experiment database provides a mechanism for longterm mutually beneficial cycles of models and experiments to arrive at a deeper understanding of biological questions. Importantly, the cycle does not involve fixed teams of researchers. At any point, another experimentalist can claim an open experiment and contribute their data. And a modeller can pick up and refine an existing model. This brings us to the second issue around model validity: The “internal validity” of a computational model, i.e. its ability to be reproduced.

4 Internal validity: Ensuring reproducibility of computational model

The importance of reproducible research has received much attention in recent years across all areas of science [36], including computational modelling of biological systems [37]. Reproducibility is an important condition for model sharing and reuse [38,39].

Much work has been done on how to ensure the reproducibility of computational work. Standards for reproducible computational research have been formulated both as stand-alone guidelines [40,41], and within the FAIR framework [17].

Community efforts to ensure reproducible modelling of biochemical reaction systems include efforts to standardise model specification (e.g. [42–48], model databases [49,50], and standards for model annotation and documentation [51,52,48].

There is now also a journal specifically designed for replication studies of previous computational work [53].

There are, however, persistent problems with model reproducibility. Not all modellers use available standards and share their code. Even for those who attempt to, there are often additional assumptions, e.g. about simulation parameters, model interfacing, or model data analysis that are not made explicit and that hinder future reproducibility [48].

To carefully annotate and document a model and ensure its reproducibility involves time and work. There is as yet little incentive or reward to doing this - the benefits of a model being reproducible often become clear years after it is first published, and scientific career paths are not currently structured to invite this level of foresight.

In theory, pre-publication peer review could pick up problems around reproducibility, but peer reviewers do not always have access to the code, computational resources, and time it would take to reproduce a model that they are reviewing. Post-publication review or replication studies [53] are valuable, but may come too late to salvage the original model.

There is thus room for a new robust process to ascertain model reproducibility pre-publication, and even pre-submission, so that any reproducibility gaps can be caught and addressed early.

5 Reproducibility audits

In order to increase model reproducibility, we suggest the introduction of prepublication reproducibility audits for modelling projects. This means that each project should involve a reproducibility auditor, whose role is to ensure the model is reproducible before it is published.

A reproducibility auditor would be a person familiar with the biological framework and modelling methodologies used, but who was not involved in the original model development. They would ideally not be part of the same research group as the original model developers, so that they do not share the workflows and implicit assumptions prevalent in that group.

When a computational model is ready for publication, the model developers send the model and write-up to their reproducibility auditor. The auditor attempts to run the model and reproduce the figures in the paper based on the information given to them, and identifies gaps in reproducibility. Both parties then work together to improve the documentation and ensure model reproducibility.

Once this is achieved, they submit the manuscript describing the model for publication together, with a short report of the reproducibility audit (steps taken, results obtained) in the appendix, and an author contribution statement specifying the role of the auditor.

The benefit for modellers is an external confirmation of reproducibility prior to publishing the model: possible gaps in documentation can thus be caught and fixed early.

For the field as such, the benefit is that there is an extra quality control step pre-publication. Journals could highlight this by introducing a "reproducibility audited" badge, similar to existing open science badges [54].

For the reproducibility auditor, the benefit is an opportunity to establish a collaboration with another research group in their field and learn first-hand about their modelling methods and process. This could be especially useful for early-career scientists just starting in a particular field. A simple reproducibility audit could even be an exercise assigned in a computational biology, biomedical informatics, or neuroscience class.

Standardisation efforts would benefit from reproducibility audits in two ways: First, reproducibility audits provide a good incentive for adhering to standards and good practices, thereby popularising the standards. Second, reproducibility audits will generate feedback on both

the usefulness and usability of the standards used, and can thus feed back into standards development.

How will modellers find their reproducibility auditors? This could be done fairly informally within existing collaborative network. Some institutions or consortia may also create the role of a reproducibility auditor or make reproducibility audits part of the role of their research integrity advisors [55].

For researchers not having access to these channels, there could also be a centralised online forum where people can post a short description of the project, techniques used, skills expected of auditor, expected workload associated with auditing and an indicative time frame. This may provide especially valuable opportunities for scientists who may not have access to mainstream scientific networks to act as auditors and thereby gain experience and establish collaborative ties.

6 Conclusions

We show here that the fostering of collaborative practices has potential for improving the validity of biochemical models, both external and internal. These collaborations are designed to be mutually beneficial, bringing researchers in similar fields closer together whilst also addressing the challenge of model validity. Furthermore, they have the potential to both accelerate model development and increase the number of biologically-derived parameters available to all researchers. The specifics of how to exactly implement these ideas should also be collaboratively decided. By setting out a possible framework we now invite readers to discuss and debate how we can best turn these ideas into reality.

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. T.D. Pollard, PLoS Biol **11**(12), e1001734 (2013)
2. J. Berro, Biophysical reviews **10**(6), 1637 (2018)
3. H. Bazzazi, Y. Zhang, M. Jafarnejad, A.S. Popel, Journal of theoretical biology **455**, 212 (2018)
4. P. Boutillier, M. Maasha, X. Li, H.F. Medina-Abarca, J. Krivine, J. Feret, I. Cristescu, A.G. Forbes, W. Fontana, Bioinformatics **34**(13), i583 (2018)
5. L.A. Harris, J.S. Hogg, J.J. Tapia, J.A. Sekar, S. Gupta, I. Korsunsky, A. Arora, D. Barua, R.P. Sheehan, J.R. Faeder, Bioinformatics **32**(21), 3366 (2016)
6. M.I. Stefan, T.M. Bartol, T.J. Sejnowski, M.B. Kennedy, PLoS Comput Biol **10**(9), e1003844 (2014)
7. E.C. Stites, M. Aziz, M.S. Creamer, D.D. Von Hoff, R.G. Posner, W.S. Hlavacek, Biophysical journal **108**(7), 1819 (2015)
8. M.I. Stefan, D.P. Marshall, N. Le Nov`ere, PLoS One **7**(1), e29406 (2012)
9. M.C. Pharris, N.M. Patel, T.G. VanDyk, T.M. Bartol, T.J. Sejnowski, M.B. Kennedy, M.I. Stefan, T.L. Kinzer-Ursem, PLoS computational biology **15**(12), e1006941 (2019)
10. A. Mogilner, R. Wollman, W.F. Marshall, Developmental cell **11**(3), 279 (2006)
11. N. Le Nov`ere, Nature reviews. Genetics **16**, 146 (2015). DOI 10.1038/nrg3885
12. M. Glont, C. Arankalle, K. Tiwari, T.V. Nguyen, H. Hermjakob, R.S. Malik Sheriff, Bioinformatics (2020)
13. L. Jeske, S. Placzek, I. Schomburg, A. Chang, D. Schomburg, Nucleic acids research **47**(D1), D542 (2019)
14. S. Sivakumaran, S. Hariharaputran, J. Mishra, U.S. Bhalla, Bioinformatics **19**(3), 408 (2003)
15. U. Wittig, R. Kania, M. Golebiewski, M. Rey, L. Shi, L. Jong, E. Algaa, A. Weidemann, H. Sauer-Danzwith, S. Mir, et al., Nucleic acids research **40**(D1), D790 (2012)
16. N.A. Viswan, G.V. HarshaRani, M.I. Stefan, U.S. Bhalla, Frontiers in neuroinformatics **12**, 38 (2018)
17. M.D. Wilkinson, M. Dumontier, I.J.J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.W. Boiten, L.B. da Silva Santos, P.E. Bourne, J. Bouwman, A.J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C.T. Evelo, R. Finkers, A. Gonzalez-Beltran, A.J.G. Gray, P. Groth, C. Goble, J.S. Grethe, J. Heringa, P.A.C. 't Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S.J. Lusher, M.E. Martone, A. Mons, A.L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S.A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M.A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao, B. Mons, Scientific data **3**, 160018 (2016). DOI 10.1038/sdata.2016.18
18. Y. Lazebnik, Cancer cell **2**(3), 179 (2002)
19. M.B. Kennedy. Biochemistry and neuroscience: the twain need to meet (2017). DOI 10.1016/j.conb.2017.01.004. URL [https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5447485/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5447485/?report=abstract)
20. W.W. Chen, M. Niepel, P.K. Sorger, Genes & development **24**(17), 1861 (2010)
21. N.A. van Riel, Briefings in bioinformatics **7**(4), 364 (2006)
22. Z. Zi, Y. Zheng, A.E. Rundell, E. Klipp, BMC bioinformatics **9**, 342 (2008). DOI 10.1186/1471-2105-9-342
23. R.N. Gutenkunst, J.J. Waterfall, F.P. Casey, K.S. Brown, C.R. Myers, J.P. Sethna, PLoS Comput Biol **3**(10), e189 (2007)
24. M.K. Transtrum, B.B. Machta, K.S. Brown, B.C. Daniels, C.R. Myers, J.P. Sethna, The Journal of chemical physics **143**(1), 07B201 1 (2015)
25. A.C. Babbie, M.P. Stumpf, Journal of The Royal Society Interface **14**(133), 20170237 (2017)
26. L. Schmiester, Y. Sch`alte, F. Fr`ohlich, J. Hasenauer, D. Weindl, Bioinformatics **36**(2), 594 (2020)
27. F. Englert, R. Brout, Physical Review Letters **13**(9), 321 (1964). DOI 10.1103/physrevlett.13.321
28. P.W. Higgs, Physical Review Letters **13**(16), 508 (1964). DOI 10.1103/physrevlett.13.508
29. G.S. Guralnik, C.R. Hagen, T.W.B. Kibble, Physical Review Letters **13**(20), 585 (1964). DOI 10.1103/physrevlett.13.585
30. ATLAS Collaboration, Physics Letters B **716**(1), 1 (2012). DOI

- 10.1016/j.physletb.2012.08.020
31. CMS Collaboration, *Physics Letters B* **716**(1), 30 (2012). DOI 10.1016/j.physletb.2012.08.021
 32. The discovery of longitude at sea act 1714 (13 anne, c. 14) (1714)
 33. T. Brady, *Journal of Aviation/Aerospace Education & Research* **12**(1), 9 (2002)
 34. A.M. Jaffe, *Notices of the AMS* **53**(6) (2006)
 35. R. Puttick, P. Baeck, P. Colligan, Londres: Nesta & Bloomberg Philantropies (2014)
 36. M. Baker, *Nature* **533**, 452 (2016). DOI 10.1038/533452a
 37. P. Mendes, *Bulletin of mathematical biology* **80**, 3081 (2018). DOI 10.1007/s11538-018-0498-z
 38. L. Cucurull-Sanchez, M.J. Chappell, V. Chelliah, S.Y.A. Cheung, G. Derks, M. Penney, A. Phipps, R.S. Malik-Sheriff, J. Timmis, M.J. Tindall, P.H. Graaf, P. Vicini, J.W.T. Yates, *CPT: Pharmacometrics & Systems Pharmacology* **8**(5), 259 (2019). DOI 10.1002/psp4.12381
 39. M. Scharm, T. Gebhardt, V. Tour'e, A. Bagnacani, A. Salehzadeh-Yazdi, O. Wolkenhauer, D. Waltemath, *BMC Systems Biology* **12**(1) (2018). DOI 10.1186/s12918-0180553-2
 40. G.K. Sandve, A. Nekrutenko, J. Taylor, E. Hovig, *PLoS computational biology* **9**, e1003285 (2013). DOI 10.1371/journal.pcbi.1003285
 41. A. Elofsson, B. Hess, E. Lindahl, A. Onufriev, D. van der Spoel, A. Wallqvist, *PLoS computational biology* **15**, e1006649 (2019). DOI 10.1371/journal.pcbi.1006649
 42. M. Hucka, A. Finney, H.M. Sauro, H. Bolouri, J.C. Doyle, H. Kitano, J. Doyle, A.P. Arkin, B.J. Bornstein, D. Bray, A. Cornish-Bowden, A. Cuellar, S. Dronov, E.D. Gilles, M. Ginkel, V. Gor, I.I. Goryanin, W.J. Hedley, T.C. Hodgman, J.H. Hofmeyr, P.J. Hunter, N.S. Juty, J. Kasberger, A. Kremling, U. Kummer, N. Le Nov'ere, L.M. Loew, D. Lucio, P. Mendes, E. Minch, E.D. Mjolsness, Y. Nakayama, M.R. Nelson, P. F. Nielsen, T. Sakurada, J.C. Schaff, B.E. Shapiro, T.S. Shimizu, H. D. Spence, J. Stelling, K. Takahashi, M. Tomita, J. Wagner, J. Wang, *Bioinformatics* **19**(4), 524 (2003)
 43. F. Zhang, L.P. Smith, M.L. Blinov, J. Faeder, W.S. Hlavacek, J.J. Tapia, S.M. Keating, N. Rodriguez, A. Dr'ager, L.A. Harris, A. Finney, B. Hu, M. Hucka, M. MeierSchellersheim, *Journal of Integrative Bioinformatics* **1**(ahead-of-print) (2020)
 44. W.J. Hedley, M.R. Nelson, D.P. Bellivant, P.F. Nielsen, *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **359**(1783), 1073 (2001). DOI 10.1098/rsta.2001.0817
 45. N. Le Nov'ere, M. Hucka, H. Mi, S. Moodie, F. Schreiber, A. Sorokin, E. Demir, K. Wegner, M.I. Aladjem, S.M. Wimalaratne, F.T. Bergman, R. Gauges, P. Ghazal, H. Kawaji, L. Li, Y. Matsuoka, A. Vill'eger, S.E. Boyd, L. Calzone, M. Courtot, U. Dogrusoz, T.C. Freeman, A. Funahashi, S. Ghosh, A. Jouraku, S. Kim, F. Kolpakov, A. Luna, S. Sahle, E. Schmidt, S. Watterson, G. Wu, I. Goryanin, D.B. Kell, C. Sander, H. Sauro, J.L. Snoep, K. Kohn, H. Kitano, *Nat Biotechnol* **27**(8), 735 (2009). DOI 10.1038/nbt.1558. URL <http://dx.doi.org/10.1038/nbt.1558>
 46. V. Tour'e, A. Dr'ager, A. Luna, U. Dogrusoz, A. Rougny, in *Reference Module in Biomedical Sciences* (Elsevier, 2020). DOI 10.1016/b978-0-12-801238-3.11515-6
 47. F. Schreiber, B. Sommer, T. Czauderna, M. Golebiewski, T.E. Gorochowski, M. Hucka, S.M. Keating, M. K'onig, C. Myers, D. Nickerson, D. Waltemath, *Journal of Integrative Bioinformatics* **0**(0) (2020). DOI 10.1515/jib-2020-0022
 48. D. Waltemath, M. Golebiewski, M.L. Blinov, P. Gleeson, H. Hermjakob, M. Hucka, E.T. Inau, S.M. Keating, M. K'onig, O. Krebs, R.S. Malik-Sheriff, D. Nickerson, E. Oberortner, H.M. Sauro, F. Schreiber, L. Smith, M.I. Stefan, U. Wittig, C.J. Myers, *Journal of Integrative Bioinformatics* **0**(0) (2020). DOI 10.1515/jib-2020-0005
 49. M. Glont, T.V.N. Nguyen, M. Graesslin, R. H'alke, R. Ali, J. Schramm, S.M. Wimalaratne, V.B. Kothamachu, N. Rodriguez, M.J. Swat, J. Eils, R. Eils, C. Laibe, R.S. Malik-Sheriff, V. Chelliah, N. Le Nov'ere, H. Hermjakob, *Nucleic Acids Research* **46**(D1), D1248–D1253 (2018). DOI 10.1093/nar/gkx1023. URL <https://dx.doi.org/10.1093/nar/gkx1023>
 50. R.S. Malik-Sheriff, M. Glont, T.V.N. Nguyen, K. Tiwari, M.G. Roberts, A. Xavier, M.T. Vu, J. Men, M. Maire, S. Kananathan, E.L. Fairbanks, J.P. Meyer, C. Arankalle, T.M. Varusai, V. Knight-Schrijver, L. Li, C. Duen'as-Roca, G. Dass, S.M. Keating, Y.M. Park, N. Buso, N. Rodriguez, M. Hucka, H. Hermjakob, *Nucleic Acids Research* (2019). DOI 10.1093/nar/gkz1055. URL <https://doi.org/10.1093/nar/gkz1055>. Gkz1055 51.

51. D. Waltemath, R. Adams, F.T. Bergmann, M. Hucka, F. Kolpakov, A.K. Miller, I.I. Moraru, D. Nickerson, S. Sahle, J.L. Snoep, N. Le Nov`ere, *BMC systems biology* **5**, 198 (2011). DOI 10.1186/1752-0509-5-198
52. F.T. Bergmann, R. Adams, S. Moodie, J. Cooper, M. Glont, M. Golebiewski, M. Hucka, C. Laibe, A.K. Miller, D.P. Nickerson, et al., *BMC bioinformatics* **15**(1), 369 (2014)
53. K. Hinsén, N.P. Rougier, (2019). DOI 10.5281/ZENODO.3069619
54. M.C. Kidwell, L.B. Lazarević, E. Baranski, T.E. Hardwicke, S. Piechowski, L.S. Falkenberg, C. Kennett, A. Slowik, C. Sonnleitner, C. Hess-Holden, T.M. Errington, S. Fiedler, B.A. Nosek, *PLoS biology* **14**, e1002456 (2016). DOI 10.1371/journal.pbio.1002456
55. C. Winchester, *Nature* **557**, 281 (2018). DOI 10.1038/d41586-018-05140-x

Conclusion: Student involvement in validation

So where do students fit into this idea? As we saw in the previous chapter, even secondary students can harness computer models and generate predictions, or at least testable outcomes. Similarly, many of the experiments that a model requires for coarse kinetic data are well within the capacity of an undergraduate student to perform. It is therefore entirely plausible to have at least part of this collaborative endeavour be conducted solely by students. Laboratories provide the materials and space to conduct the experiments, or the computing equipment to generate the models. They receive the microgrants to cover the costs of this, and students can take part in work which has an immediate impact for a specific person, group, or field.

Furthermore, this approach gives a much clearer insight to younger students on the collaborative nature of science, that data should be shared and generated not just for your own specific perceived needs. It also shows the slow, iterative way knowledge is gained and shared. For researchers, this idea provides easy, funded internship and Honours Project ideas, and helps foster collaborations.

The student examples I have shown in this thesis indicate that a lot of useful, foundational pilot models can be achieved readily by non-specialists. These models are impactful on much more than just the lab they were carried out in, and have highlighted that we underestimate the amount of value such “basic” experiments have on challenging our understanding, and on making sense of complex concepts.

CHAPTER 7

MinecraftED

What I hope I have successfully put across in the previous chapters is that modelling can play a much greater role in biomedical education than has been currently realised. In this vein I have developed a novel way of adapting a pre-existing game environment to construct both structural and functional models of biomedical concepts. In this chapter I will demonstrate this work in both academic and public outreach settings and look at what more can be done with this idea by myself and other interested parties.⁷⁶ All Minecraft worlds in this chapter can be downloaded from https://github.com/yishusing/Thesis_Repository⁷⁷ and work for Minecraft Java Edition⁷⁸⁷⁹.

⁷⁶ Many of these projects are ongoing at the time of writing; interested readers can email the author at neuronsafari@gmail.com about any of the projects listed here for post-publication updates.

⁷⁷ The file names are listed at the start of each relevant section in this Chapter.

⁷⁸ Additionally, summaries of both the Pilot Study and Project Millport have been written up for the Teaching Matters blog (<https://www.teaching-matters-blog.ed.ac.uk/generating-immersive-large-scale-virtual-field-work-for-zoology-students/> ; <https://www.teaching-matters-blog.ed.ac.uk/projects-peers-passion-and-play-how-minecraft-transforms-hybrid-teaching/>)

⁷⁹ This chapter has several videos throughout. These *should* work in the pdf format by just clicking the image. I have given URLs in case where necessary, and all are available in the repository. All screenshots were made using Minecraft Java Edition, with the Optifine addon <https://www.optifine.net/home> and Sildur's vibrant shaders <https://sildurs-shaders.github.io/> active.

Learning Through Games and Play

The use of games to help convey knowledge and understanding to learners is well established. The first critical analysis of this came about in 1980 by Thomas Malone (Malone, 1980a, 1980b), who identified challenge, fantasy and curiosity as three key elements of early games that sustain intrinsic motivation – the continuing of an activity for no appreciable reward that is still perceived to be fun. Since then, many different elements have been identified to as defining attributes of games (see (DeRouin-Jessen, 2008; Wilson et al., 2009), but these three were the primary attributes by which the work of this chapter was guided by.

Some readers may query the relationship here with gamification of learning. This is the rather widespread use of activities in teaching which are normally associated with recreational events in non-game settings (Deterding, Dixon, Khaled, & Nacke, 2011). Whilst undoubtedly related, what we are addressing here is game-based learning, which is the use of games themselves as vehicles of learning (Plass, Homer, & Kinzer, 2015)⁸⁰. This is more involved, more personalised, and affords more agency to the learner than gamified approaches. Rather than enhancing learning in a traditional context, game-based learning extends learning and cognition into a new physical or digital context (Gee, 2005).

Projects, Peers, Passion, Play

A useful exemplar of this approach can be found in the work of Mitchel Resnick, who runs the Lifelong Kindergarten group at MIT. His group's work revolves around the generation of ways to get learners engaged creatively with their learning, through new technologies and methods. To encapsulate this, he uses the Four P's: Projects, Peers, Passion and Play (Resnick, 2014).

I believe it is useful here to go into this paper in a little more detail, as this approach mirrors closely what I have both intended for this project and where it appears to be heading. Projects are only as useful to learning as they are meaningful – they must be something the learner directs, creates and takes charge of. Supplying a ready-made project to a learner provides a framework and grounding, but it too often removes the learner's agency. An active project is not just one which is being conducted, but

⁸⁰ The two terms (gamification and game-based learning) are often conflated or poorly distinguished in the current literature.

importantly one which is being constantly updated and modified to what its participants feel should be the project's goal.

Peers are vital because all of us are much more willing to take risks and fail with those we perceive to be at the same level. Learning from one another is often through highly applied and appropriately contextually framed examples and explanations that allow for rapid acquisition. Taking the idea further, it exemplifies that learners *should* share ideas and work to improve others, and that this should not be seen as some optional component of learning. As academics, we build on the work of others, taking elements and enhancing or refuting it to drive knowledge on. Learning should not be any different.

The final two Ps, Passion and Play, are the heart of the method sketched out in this chapter. Anyone with a hobby knows that acquiring new skills and information about something you enjoy is often effortless, and usually perceived differently to work. When we care about the thing we are learning or teaching about, and feel those around us do also, we shrug off many of the distractors, dissuaders and failures that would normally stop us (Ballantyne, Bain, & Packer, 1999; Carbonneau, Vallerand, Fernet, & Guay, 2008; Kahu, Nelson, & Picton, 2017; Lyubomirsky, King, & Diener, 2005). This I see links strongly to Projects, and that taking ownership of the content you choose to deepen your understanding of.

Play will be the most obvious of the four to track through all the examples detailed here. Video games are a success precisely because they tap into our desire for the new, for a safe environment to fail and do the same things over and over simply because we like doing it. 'Playing' a video game can have a depth to it that is not fully captured by how the word is used by those with minimal engagement with the medium, and there are many examples now beginning to show just how much power this medium has in teaching skills and shaping thought (Barr, 2017, 2018; Marín-Díaz, Sampedro-Requena, & López-Pérez, 2020; Nebel, Schneider, Schledjewski, & Rey, 2017).

Deepening the complexity of play is the relationship to what a 'game' is. Something can be regarded as play without being a game, such as manipulating toys or objects. These may have transient rules, but they can change at will. Likewise, a game does not have to be playful. A game can be treated a structured space with rules which may have arisen out of playful experiences, and arguably once made it is a separate category that should not be conflated with play (Walther, 2003). Roger Caillois (1961) places these on a spectrum, with play (*paidia*) and game (*ludus*) at opposite ends.

Play has no or unclear goals, little consequence of the final outcome, and is often spontaneous (Lucero, Karapanos, Arrasvuori, & Korhonen, 2014). Games, or gamification, assigns rules and goals – it structures play, adds motivations and challenge through requiring skill to progress (Barr, 2017). Games can be fun, they can be escapes from reality, but unlike play this may not stem from the imagination of the user, and they may have no real choice in the game world itself.

This chapter will focus almost exclusively on the use of Minecraft, and how I implemented this in an extremely diverse number of ways quite quickly into pedagogical practices⁸¹. On the spectrum of *paidia* and *ludus*, Minecraft has a rather unique position in that it can be used almost wholly playfully, wholly gamefully, or anywhere in the middle⁸². Indeed, some of the work in the chapter begins on the *paidia* side, some on the *ludus* side. All meet somewhere in between these two points.

Minecraft

Minecraft is a sandbox gaming environment where players predominantly create, interact with, and destroy blocks. Since its initial release in 2009, it has sold over 200 million copies (Chiang, 2020), with an additional 200 million users of the free Chinese version (Ahmad, 2019). There is no “correct” way of playing the game, although an end-goal was added in 2011⁸³. Players may choose between several game modes to interact with the game environment. Survival Mode is the default, where players construct bases and collect materials, avoiding hostile entities and exploring the procedurally generated world. Creative Mode allows users unlimited access to all the game’s resources, removes dying and attack from hostile entities, and allows the user to fly. These combined changes allow the user to quickly construct structures and machines that would otherwise take many hours to achieve in Survival Mode.

Creative Mode has been used to create of a diverse range of structures, from simple hobbyist houses, through to detailed scale replicas of famous monuments, cities, and in one case, a whole country (Magnussen & Elming, 2015). The game world is

⁸¹ We will initially explore this via a thematic analysis of the reflections from students and myself, and go on to suggest potential avenues for larger scale implementation into existing educational structures.

⁸² Indeed, this may be why it has been as successful and enduring as it has been thus far.

⁸³ This is never explicitly presented to the player however, and no directions as to how to get there. Players have to stumble upon it, and the solution. There is no punishment for not doing so, and only a token reward for completing it.

procedurally generated⁸⁴ in chunks. These are 16 blocks wide x 16 blocks long and 256 blocks tall (a total of 65, 326 blocks), and are generated as and when required as the user moves through the game space. This just-in-time processing of chunks is necessary due to the size of the Minecraft map, which is limited normally to 30 million blocks in the X and Z planes. The Y plane is restricted in vanilla game versions to 256 blocks. Users may build anywhere in this map, with changes made saved to the relevant chunks.

Tiers of Model Construction

A user-generated structure (or world) is not limited to just the complex arrangement of different structural blocks. There are three tiers of build complexity I have identified which all build on the previous tier to enhance the capabilities any user-generated structure has. Purely structural builds are the base tier, and the most common amongst general users. This build approach shares much of the same reasons for sustained engagement and creativity as Lego does, given a similar mechanic of additive construction.⁸⁵ Models are constructed either to reproduce real-world structures or objects (as seen above) or to create new structures based on the user's imagination. Models may also evolve organically, changing shape to fit the user's changing functional requirements for the space.

The next tier makes use of a specific type of in-game mineral, redstone. Depending on its configuration, it may act as a power source or conduit, and functions as an in-game form of electricity. Users may construct circuits powered by redstone that may activate several specialised blocks (which themselves are usually constructed in part from redstone). For example, a lamp block will switch on and emit light in the presence of an active redstone circuit, remaining on so long as this source is intact. Repeater and comparator blocks further enhance the complexity of these circuits, adding in the ability to set "clocks" and control the pattern and magnitude of redstone signals. These properties allow a user to construct machines, from simple piston operated doors through to working computers.

The previous two tiers combined allow a user to generate in theory almost anything and are available in both Creative and Survival Modes. Constructing highly complex machines however requires vast amounts of space and planning to realise, and so can be impractical for some of its theoretical uses. Command blocks allow users to greatly

⁸⁴ This can be controlled through adjustments to a few parameters prior to world creation.

⁸⁵ Carving is also entirely feasible, but unnecessary and rarely used.

simplify these builds by granting access to a syntactically constrained but still powerful coding language⁸⁶. They also go beyond mere simplification and offer users the capacity to manipulate the game's properties in ways impossible through the previous two tiers, such as editing block properties, controlled creation and deletion of blocks, and the creation of trigger conditions needed for a code sequence to fire.

Each tier can be used alone or in conjunction with any or all the others. Proficiency in using these different construction methods is not strictly related to amount of time spent using the software, but as most users spend their time constructing in Survival Mode, they only slowly develop redstone structures and are unable to utilise command blocks, which would "break" the mechanics of this mode. The mechanics of these two higher tiers is also more intricate and are more powerful as the user acquires more knowledge of the game properties.

But Why Minecraft?

Much of the work that follows uses a blend of all three tiers to varying degrees. This still however does not explain quite *why* I have used Minecraft to investigate game-based learning.

The decision to explore this avenue of work came from work carried out by Philippa Sheail, Noreen Dunnett and Stuart Nichol at the University of Edinburgh. In a talk titled: 'HELP PLZ': *Navigating Minecraft as a digital education space with postgraduate students* (2018), they outlined their experiences with using Minecraft as an environment to engage with distance-learning postgraduate students in Digital Education, and give them space to experiment with being creative and communicating in a digital platform.

After this talk I reflected on the value Minecraft might have as a tool for creatively engaging students with biomedical concepts. In particular, the immediate interest was in how it may map to the biochemical modelling research done in the lab and how it might enhance understanding of pathways and networks in a classroom setting. Our models are often visually depicted as box and arrow diagrams, which simplify complex networks into a causal chain of events (in most cases, from ligand binding through to gene transcription). These allow a reader to see the important

⁸⁶ The details of these blocks will be dealt with later in this Chapter.

molecules, sometimes their localisation, and how many other elements they interact with.

When we compress the network into this state however we remove the mechanisms by which the network functions. A → B and B → C might look the same in the diagram, but what the arrow means is not always clear, and not always the same. Looking back at the diagrams in this thesis, arrows can mean: translocation, activation, complex formation, state changes. For a learner trying to make sense of a functional network, knowing what these arrows mean is a hidden task for them to resolve. A study by Kate Wright and colleagues (2018) found that biology textbooks fail to show consistency in arrow use, down to differences between figures in the same publication, and are only vaguely comprehensible to undergraduates (see also Wright, Fisk, & Newman, (2014). There are several attempts to general universal graphical notation formats, such as SBGN (Mi et al., 2015; Rougny et al., 2019; Sorokin et al., 2015) and PDN (Kitano, Funahashi, Matsuoka, & Oda, 2005). This will, eventually, remove this confusion, as they work on principles similar to flow diagram notations. Right now though, efforts are focussed on correcting the use of notation in research, not necessarily teaching material.

What is needed then is not so much an overhaul of these diagrams, but a complementary way of exploring these networks that has the capacity to introduce these mechanistic concepts and thus make sense of these arrows in relation to the boxes they link together (Newman, Stefkovich, Clasen, Franzen, & Wright, 2018). It is from this that I conceived of an exploratory study into the use of Minecraft in giving students an insight into the mechanisms that underpin biological phenomena, and how that might find itself implemented within teaching.

Pilot Study

World File: BioMedEd (CaMKII)

I applied for a small grant with the Principal's Teaching Award Scheme, which is designed to allow educators to test out novel ways of enhancing learning and teaching for both undergraduates and postgraduates at the University of Edinburgh. The money was used to buy copies of Minecraft: Java Edition, a Realm subscription and primarily to pay students for their time conducting the study.

Students were recruited through an email advertisement to all biomedical science undergraduates. Four positions were offered, with the requirement that students be

entering their third or final year of study in the Autumn (2019). Each position was to be paid at a rate of 12 hours a week for one month and require the write up of a reflective report to aid in determining how the project would progress. Entries were shortlisted from their CVs and interviewed by myself both in-person and via Skype.

The four students chosen (Okuk Takon, Simon Parker, Connor Schlemmer, and Srivatsan (Chikku) Rajagopalan)⁸⁷ represented a diverse range of prior Minecraft experience, educational backgrounds, and biomedical interests. An initial meeting laid out the plans for the study, which was to be conducted in the students' own time, with their own computer devices. Whilst they were to be paid for 12 hours a week, no checking of this would be carried out and they were discouraged from working past this limit. Check-ups were to be done through weekly meetings, and that anything built should be available for all to see on the Realm. The Realm is a multiplayer Minecraft world that allows everyone to work in the same environment.

Two tasks were given in the month, each totalling two weeks of work. These will be explained in more detail below, but the broad aims were to see how students responded to constructing models when given a specific brief, and when they had free choice in what to build. From these two tasks we intended to glean what would be the most promising avenues for further research in using this software.

Task #1: CaMKII Model

All team members were presented with a very brief description of CaMKII, showcasing three diagrams that pointed out its holoenzyme and subunit structures. They were told as little as possible, as I wanted to assess how much research would be done, and how much was necessary, to construct a model of some aspect of this molecule without any bias. All team members (including myself) were tasked with spending much of the first week researching and choosing an interesting aspect. The second week would involve the construction of a model based on this work.

Task #2: Free Choice Model

This task was designed to explore both the creativity of students and how Minecraft might be used in ways which were not anticipated. It was also however optional – team members could choose to continue developing their first model if they wished instead. The only constraint was that the model had to be of something biomedical

⁸⁷ Names have been obscured with letter attributions from here onwards (Students A - D).

that was taught at a university level and had to be understandable without the modeller being present.

Only Student D decided to continue with building the first model. This was explained to me as being out of frustration at not having a model to show to the team, and a desire to work out how to translate the model in their head into reality. Student A began construction of a separate but related model to the first, with Student B and Student C deciding on completely independent models (**Figure 72**). This may have been partially to do with the latter two's greater experience with Minecraft and some of its more advanced mechanics.



Figure 72 Example of student work in Minecraft.

Student C's two models are shown above. He constructed a colour-coded model of the CaMKII holoenzyme structure (left), and an animation of the direction of blood flow through the heart (right). The animation can be seen [here](#).

Reflective Reports

Each student was asked to write a reflective report on their experiences during the study. I provided a list of possible questions grouped into six categories:

- Ease of use
- Value of using a shared workspace
- Time
- Concepts built
- Teaching Potential
- Personal Reflection

Questions are available from https://github.com/yishusing/Thesis_Repository. Whilst not required, most of the students decided to adhere to these categories as they structured their reports. Rather than paraphrase their contributions, I will be presenting key extracts in the relevant sections below and discussing them⁸⁸.

To carry out this analysis, I uploaded all four reports into Dedoose (<https://www.dedoose.com/>; Version 8.0.35) and performed a content analysis. I used 15 codes, which were mostly derived from the guidance notes given to the students. Each code was applied at least once across all four reports. I allowed for multiple codes for a given extract, and did not limit extract length. A total of 124 extracts were identified, with a total of 171 code applications logged (**Figure 73**).

Given the prior guidance for these reports, most extracts as can be inferred from the above were focussed to particular points and so made analysis much easier. Where code co-occurrence did occur, the strongest pairing was between *Teaching Potential x Engagement*, with 16 occurrences. This was followed by *Teaching Potential x Concept Realisation* (9) and *Ease of Use x Concept Realisation* (7) (**Figure 74**).

⁸⁸ Emphasis in quotes are my own, unless stated.

Name	Biomedical Content	Concept Realisation	Ease of Use	Engagement	Minecraft Knowledge	Prior Experience	Project Description	Thoughts on Project feasibility	Shared Workspace	Teaching Potential	Assessment	Outreach	Personal Learning	Time Management	Time taken	Totals
Student C	0	4	6	6	5	1	0	1	1	6	2	0	1	3	3	39
Student B	2	1	2	4	0	1	0	1	1	5	5	1	1	1	3	28
Student A	1	12	5	11	3	4	3	1	2	16	3	0	2	2	1	66
Student D	2	4	1	8	3	2	2	1	1	5	2	0	5	1	1	38
Totals	5	21	14	29	11	8	5	4	5	32	12	1	9	7	8	

Figure 73 Code quantities for each student report on the MinecraftED pilot.

Almost all codes were found in each student report, with teaching potential, engagement and concept realisation being the most common. Only outreach was almost never mentioned.

Code	Biomedical Content	Concept Realisation	Ease of Use	Engagement	Minecraft Knowledge	Prior Experience	Project Description	Thoughts on Project feasibility	Shared Workspace	Teaching Potential	Assessment	Outreach	Personal Learning	Time Management	Time taken	Totals
Biomedical Content		2	0	0	0	0	1	0	0	0	0	0	0	0	0	3
Concept Realisation	2		7	5	3	3	0	0	1	9	0	0	2	4	2	38
Ease of Use	0	7		3	3	4	0	0	1	3	0	0	0	0	1	22
Engagement	0	5	3		2	0	1	0	1	16	5	0	5	2	1	41
Minecraft Knowledge	0	3	3	2		1	0	0	0	3	0	0	0	0	0	12
Prior Experience	0	3	4	0	1		0	1	0	2	0	0	0	1	0	12
Project Description	1	0	0	1	0	0		0	0	1	0	0	0	0	0	3
Thoughts on Project feasibility	0	0	0	0	0	1	0		0	0	0	0	0	0	0	1
Shared Workspace	0	1	1	1	0	0	0	0		1	0	0	0	0	0	4
Teaching Potential	0	9	3	16	3	2	1	0	1		5	1	5	1	0	47
Assessment	0	0	0	5	0	0	0	0	0	5		0	1	0	0	11
Outreach	0	0	0	0	0	0	0	0	0	1	0		0	0	0	1
Personal Learning	0	2	0	5	0	0	0	0	0	5	1	0		0	0	13
Time Management	0	4	0	2	0	1	0	0	0	1	0	0	0		4	12
Time taken	0	2	1	1	0	0	0	0	0	0	0	0	0	4		8
Totals	3	38	22	41	12	12	3	1	4	47	11	1	13	12	8	

Figure 74 Code Co-occurrence across all reports on the MinecraftED pilot.

Values for all co-occurring codes found during the process of analysing student reports. Where multiple codes were present, they were more likely to occur when one code was Teaching Potential, Engagement or Concept Realisation. The most common co-occurrence was Teaching Potential x Engagement, highlighting student understanding of the pilot aims and its potential.

I will now outline the main findings from the report, using the six categories as a grouping factor, and two additional ones (prior experience and engagement). Student responses are identified by letter and colour. Any emphasis (in bold) is my own unless stated.

Full reports can be found at: https://github.com/yishusing/Thesis_Repository

Prior Experience

The team consisted of a range of prior experiences with Minecraft, which was intentional. The feasibility of this approach depends on how quickly a student can pick up using the software, both in terms of the control mechanisms used and its features. Furthermore, it was necessary to see how easy it was for players of the game normally to adapt to using the software for an educational purpose, as this may contrast too much with their expectations of using the software.

Student A had the least experience prior to the project. They were aware of the game, and had supervised children in the past who were playing it, but had had no hands-on experience with the software. Student D had a lot of play-time with friends over the years, but almost exclusively in the game's default 'Survival' mode. This meant the creative aspect of the game was constrained by the additional needs to stay alive, explore and collect resources, and so their prior experience had been much more about utility rather than design or creativity.

Both Student C and Student B had more advanced experiences with the software. Both had made use of the Creative mode, as well as Survival mode, and were therefore more comfortable with its advanced features. They therefore acted here as useful benchmarks to how far the software could be manipulated into expressing biological concepts, and offer insights into software functionality which I was unaware of⁸⁹.

Concept construction: what did you build, and how?

Student D's experience with the software was tempered by a frustration of ideas not mapping easily into implementation. This is a recurring theme within their reflections on construction:

I went about seeking inspiration for the models online, using various diagrams as a guide. After sifting through enough, I was able to pull up some simple diagrams that I could mimic with my limited Minecraft skills. Again, **I was frustrated by my lack of skills [...]**

⁸⁹ For context, at the time of this pilot, I was at the same experience level as Student A.

I was thinking about how I could make models my own; I did not want to transpose a diagram I saw online and make a copy. [...] I had many ideas and concepts; however, I was limited. It was a great deal harder, perhaps due to the expectations I had set for myself.

Progress was slow and difficult. Yet with the construction of the molecule in mind, I believe that my personal understanding of the molecule was growing. I was starting to understand the potential benefits and applications of such a task. It was difficult to gauge progress.

Was progress in making a brilliant model that others could use? Or would the project be deemed more successful if the activity improved and furthered my own understanding?

Student D's frustration was borne partially out of their own high expectations of what *could* be achieved. Knowing a lot about the capacity of the software led to an impasse where much of their time was spent trying to build a complex model all at once. Their ideas outstripped their in-game ability, and with the short time frame and independent nature of the study, there was insufficient time to direct them out of this frustration.

Tellingly however, Student D's comments show that despite these setbacks, they were gleaning a lot of understanding about the molecule from them. The fixation on one aspect of CaMKII (structural changes due to individual subunit activation) led to further research to see if a different model could be done on a closely related phenomenon. So whilst their ambitions of an extravagant model were thwarted, they identified that one positive outcome would still be the enhancing of their understanding.

This reflection points again to the undervaluing of failure in biology learning. With creativity constrained by the subject's nature and pedagogical structure, failure becomes internalised as a metric of no understanding, which is not always true. In this instance, the failure to construct a working model does not invalidate all the research and learning done in the attempt. The wider gulf between what is known and what was presented however may be here what helps a student and an educator tease out the difference more explicitly. A poor mark can demoralise and persuade learners that they do not understand the topic, when the actual failing was in how they presented it, and is very difficult to write into a feedback form coherently. If Student D's ability to disconnect the two maps to other students too, this may be one way in which we can exemplify the differences

between understanding of a topic, and effective dissemination of that understanding.

Whilst Student D was in some ways restricted by their prior experience with Minecraft, Student A was able to more easily overcome their lack of experience by identifying earlier the *purpose* of their model:

I elected to create a dynamic functional model of the initial entry of calcium into the E/F hand domains of calmodulin. As this followed the rule of attraction, I was led to the use of zombies and villagers as these are two components in the game that are closely linked together. I ultimately chose this model **because it fulfilled my aim of creating an educational model** to represent processes that may be hard to visualise otherwise.

Student A was more concerned about how others would understand their model, rather than the complexity possible using Minecraft for modelling. As such, what was important was more its value as an educational tool, than how it looked:

My design progressed from a simple structural model with no moving parts to one that used Redstone, repeaters, pistons etc. I also managed to encase the model in a structure that resembled the shape of the calmodulin domains. On completion of the model, I felt satisfied that it would be an effective visual representation of the calcium-binding the E/F hand domains. However, I had greatly simplified the content I had initially researched in order to develop an idea I could model.

Whilst the finished model was not perceived by Student A to be any different to perhaps a simpler method of visualisation, they overlook benefit on themselves. To have to greatly simplify means tacitly that the level of understanding of the E/F hand domains modelled was quite high, and that they were able to reduce that down to a model which was more comprehensible. Student A also unfairly calls the model 'simple' – to have a working model which shows the movement of calcium ions into their binding sites on calmodulin, however simplified, is an impressive achievement considering that they had no knowledge of Minecraft or calmodulin prior to the study.

Both Student C and Student B also had challenges with their model building, but overall they found the experience much easier because of their backgrounds. Unlike Student D and Student A, who kept to refining models related to the initial brief, they made a separate model based on their own interests in the second half of the pilot. Interestingly, their approach to what their models should be doing

was different, which allows us to explore two different advanced ways of approaching using the software.

For Student C, the aim for the first model was similar to Student A – to be educational and primarily visual:

Essentially my intention with this model was to help teach and make others think about the physical nature of the protein and thus in order to explain my model it would be easy with me present to fly around the model and each element of the model through my viewpoint. If I was to convert the molecule to be self-explanatory, I would add large labels using in-game banners and then would transcribe by thoughts about the model into books on lecterns for user to read through for additional information.

Like Student A, Student C wanted their model to be usable by someone who hadn't done the same amount of research, and that it should be static. Their final model (visible in **Figure 72**) is very visual, and conveys the full structure of the CaMKII holoenzyme, colour-coded to represent major functional parts determined through their research. This kind of structural modelling appears from this pilot to be a common way of what students mean by a model.

Student B however had a different way of thinking about how their models should operate:

My CaMKII models were highly functional, with no attention at all paid to the structural intricacies of the CaMKII molecule. The models represented CaMKII as one 'thing', ignoring the fact that adjacent subunits in real CaMKII molecules can interact with one-another to affect the overall catalytic activity of the complex. I decided to build functional models to begin with as I don't hold my artistic skills in very high regard. **I wanted to prove I could make something that represented concepts in their function before trying to represent concepts visually, combining both structure and function.**

As a result, Student B's CaMKII model is much more abstracted away from any structural components of the molecule. The model consists of a complex set of modular redstone circuitry, each of which is representative of a molecular reaction or interaction related to CaMKII activation by calmodulin, and its ability to autophosphorylate.

This model was much more complex and powerful than I had envisaged prior to the pilot. It showed that the idea of rendering actual mechanisms that functioned over time, independent of user input, was possible. It also indicated that students

could, from having no knowledge of a system, construct a sophisticated model in a short time frame, much as we saw with the traditional modelling approaches in Chapter 5. Even then, Student B was not necessarily content with the final model:

I built my concepts piece by piece, testing as I went. For the CaMKII molecule, this meant building my timers first and working to achieve that 'autoactivation/central circuit', then trial-erroring until I found a solution which would sometimes deactivate the 'central circuit'. I should note here that I don't believe my final model is even a very good representation of CaMKII. This is due to the lack of precision with which I made the timings and the lacklustre output – I never worked out a good way to represent CaMKII's enzymatic nature within the cell, hence the cop-out "My name's CaMKII, and I love you!" output. As for the timings, they work because that's the arrangement in which I found that they work, as opposed to careful planning on my part during building.

This personal feeling on the model is not reflected with the reality of it – whilst it does not perfectly capture the actual timings and show some physical manifestation of CaMKII, it is a very good and powerful representation of the network actually functioning. This was one of the ideas behind the pilot – whether biochemical network processes could be rendered somehow into Minecraft. Student B demonstrates the capacity to do so here, but they lack the context that up until this point there were very few accessible ways of demonstrating this without specialist knowledge or software.




Figure 75 Student B's redstone-based CaMKII model.

This is an example of a functional model within Minecraft, where the CaMKII activation reaction network is represented by a series of redstone circuits that “fire” at different rates. The stochastic circuitry allows the network to more faithfully represent molecular dynamics.

This discovery of the potential depth to modelling in Minecraft was further exemplified in both Student C and Student B's second models. Student C, continuing the theme of the model being primarily visual, created a small animation of blood flow through a human heart (**Figure 72**):

[...] due to a personal interest in cardiology I decided to try modelling my version of the heart in the game. I began with looking at pixel-based art of a heart and through with other reference models of the heart, I recreated cartoon model of a heart with a decent reflection of reality and then incorporated a very basic redstone circuit in addition to somewhat basic use of the coding system in the game brining [sic] everything together with the foundational concept of animation – all together to produce a relatively simple animated heart.

Here we see Student C engaging with the advanced features of the software and being enthusiastic about how far they could push themselves in creating something that was more functional. Much like with the rest of the group, there is an

underselling of their achievements here – timing the animation, generating the model frames and coding that correctly is not a ‘simple’ task. In attempting this approach, Student C show that even more visual models do not have to be static, and can be animated to help demonstrate a mechanism or process.

Student B’s second model is intended to be much more useful as a teaching tool – they felt that the first model was not as accessible as it could be. Like Student C, they turned their attention to something which was animated. Unlike Student C however, they wanted to build something which was more interactive and 3D in nature:

My second model was intended to be more akin to a teaching tool than the CaMKII models. Hence the ‘bank’ of reference stills that I could aim to recreate in animation and walls of ‘which block represents which biological molecule’. This pre-preparation provided a clear basis from which I could get to making animated representations of protein synthesis. To further achieve this ‘teaching tool’ aim, the designs for the ribosome/tRNAs were based on common textbook representations of the process. This is so that, in theory, the animation would be easier understood by students.

To put the quote into context, Student B’s second model was one to help explain protein synthesis (**Figure 76**). The whole model area acts as a representation of the model’s development, model components which are used in the animation, and three interactive spaces where users allow the model to function. This ‘DNA Factory’ is very complex and sophisticated, as it allows for a user to input their own DNA sequence. This requires much more work, as the model has to allow for every possible user input, which can generate up to 20 different amino acids or a STOP codon⁹⁰ at any point in the inputted sequence.

⁹⁰ Some of the original coding for this can be seen in the Github repository.

The blend of complexity, interactivity and visual engagement present in this model I think shows the true potential of using this software, and others like it, for assisting with the explanation of complex concepts. This is a student-built model, created in just over a month, which acts as an abstracted replication of protein synthesis that is usable by almost anyone, and visually appealing and distinct enough to help an educator use to explain the concept accurately. It is from this model that the rest of this Chapter has sprung, and we shall see it again in different guises later.




Figure 76 Video demonstration of the DNA Factory.

Student B created a video guide to their model, which helps further introduce the concept of protein synthesis, and demonstrates how the model operates in detail. Video link: <https://www.youtube.com/embed/OfkCSAgr35I?feature=oembed> .

How easy was it to use Minecraft?

The most important voice here is Student A's, who came into this study with no prior Minecraft experience:

[...] I was interested in determining the usability for someone with similarly minimal experience. My initial impressions were that it was generally very user-friendly. Within my first hour of trialling the programme, I was familiar with the relevant keyboard controls, the location of the tools I would use for building and generally how to navigate the realm.

The other team members all had a range of prior interaction with the software. Student D did not mention in their report, although during interviewing for the position they did explain they were a regular player of the game, albeit not in its Creative mode. Student C had a similar level of experience, but with more experience in the creative possibilities of the game:

[...] the software was incredibly user friendly and working with the program and the realm were easy – however I have had significant Minecraft experience before and thus am not a good metric to measure ease of use really.

Student B already had a high level of experience with Minecraft, particularly in the use of redstone and general construction. Their reflections are thus focussed on the more specialised features of the software, namely redstone and command blocks:

I entered the project as someone used to the mechanics involved in modding and simple redstone usage, as well as in Minecraft's control scheme and 'vibe'. Even then, devising methods with which to implement real-world phenomena in the game was difficult. [...]

I found working with command blocks more fun, as I could see the code changing as I edited and wrote it. I felt I could almost feel it being correct before I ran it within Minecraft. That said, building what was necessary for the redstone and command block models was simple – but obviously could be (and still can be) optimised [...]

All the students reported little difficulty with acquiring the know-how to build structurally in Minecraft. Student A however mentions that for a beginner the sheer number of possibilities can be overwhelming:

Initially, I wondered how limiting exclusively modelling with cube-shaped building blocks would be, however, I found it

reasonably easy to build much larger, alternative shapes by simply layering the blocks.

[...] the vast array of design potential was hindered by my inexperience. I felt that my design ideas were limited by what I could see in my toolkit and without scrolling through every available tool, I found it difficult brainstorm. Essentially, the foundation on which to build my ideas upon was present but I lacked the knowledge to represent these ideas in Minecraft.

This frustration however was mitigated with time – the more time it was possible to spend understanding a feature, the easier it was to then use effectively. Student A also notes that this could be accelerated through the availability of pre-built structures, which would be helpful for beginners in particular:

An easy way to circumvent this will be to predesign models and allow pupils to learn via exploration. However, this will remove the aforementioned benefit of interactive learning. Despite this challenge, I generally found that as I developed my model, I better understood how to use each feature and my general confidence using the software increased.

Much of the difficulty arose from the use of more functional features in the game, such as redstone. This contrasts with Student C's comments, who built models which were overall much more structural than their colleague's approaches:

However, the physical building of the model once a clear picture in my head had been conceived didn't take a long time, at least in comparison to the time I had to take to understand CaMKII's physicality.

Student D however did run into problems in translating their ideas into the game, especially due to the difficulty in making circular structures:

Minecraft blocks are cubes, and therefore creating a circular model poses it is on geometrical issues. For instance, I wanted to explore a concept; however, it was difficult to make anything quickly. If it were to look around, it took a bigger structure, which took quite long.

This issue is largely personal however, as we did not require the models to include circles, and approximations are possible. Part of the issue was with how Student D wanted to explain CaMKII function (see Concept construction).

Overall, no one reported any intrinsic difficulties with using the software from a controls angle, just the initial issue of distinguishing 'destroy' and 'add/use' block

control buttons (normally left and right click respectively). The challenges of implementing ideas into Minecraft we will discuss further in Concept construction.

Was the Realm a valuable space?

One aspect I was keen to maintain in using this software was the ability to work within the same world-space. I wanted to encourage the sharing of ideas and spark inspiration from being able to observe other team members building, or at least visit their constructions and explore them. I strongly encourage the use of the Realm, although I did not enforce it. Of the team members, only myself, Student B and Student C chose to build in the Realm more than in a single-player world. However, the existence of the Realm was seen as useful even when it was not used:

[...] I generally found it quite daunting. The open plan nature of the realm we were given to work on, whilst sufficiently spaced, did make me feel slightly self-conscious.

In the early stages of the project, I knew I was not yet in a place to produce the impressively complex models that I could see my colleagues building. Moreover, I had previously determined that the best way to increase proficiency with Minecraft was to experiment with a range of structures, yet I could not bring myself to do so in front of my other teammates.

Student A here demonstrates that my intention for sharing ideas was appreciated but perhaps not implemented as well as it could have been. To aid in organising the Realm, I had laid out a zone with specific workspaces for each group member. These were connected by a central zone where a user could take a minecart to any other group members' workspace. These individual workspaces were reasonably close together however, and travel to each one did not take any appreciable time. As such, the self-conscious aspect of building in public was perhaps heightened. If travelling to another's workspace took more effort, or perhaps required requesting of co-ordinates to reach, the shared space may have worked better.

It seems that for more experience/confident users, the Realm setting was much more immediately beneficial:

It was extremely useful being able to build in a semi-public setting as it was wonderful being able to look at everyone else's projects through their iterations and take inspiration from them to improve my own projects.

Additionally, due to the realm permanently being up our schedules could be very flexible within a simple deadline and thus

I never really build anything on the realm when anyone else was present, hence the potential pressure of being watched was never there for me.

Seeing the creations other people had made additionally pushed me to work harder in order to keep up with everyone and meet deadlines and aim higher for more interesting projects

I used the Realm for everything I made/did. I found its features helpful and unhelpful in different ways. For example, it was helpful that I could join the Realm on any PC with Minecraft installed and have the world ready to work within. This meant that at home I could flip between using my laptop and my home PC, always able to carry on from where I started. [...]

Student B and Student C were the only students to make considerable use of the Realm alongside myself, and they were able to take inspiration from seeing these other structures much like Student A. Student B notes that the Realm setting made them less likely to redo elements, which was in part my intention. The failed versions show a logical progression of their ideas, and also help to direct others away from methods of implementing that didn't work as expected.

The benefit to the ability to see your own as well as others development and growth of ideas is best exemplified by Student D:

We worked on a realm where everyone participating in the investigation worked in the same space. Everyone could look at the others' models. **This was both an advantage and a disadvantage.**

It was advantageous to seek inspiration or when someone had utilized a different technique or method to be able to adapt that to my own model. It was also interesting to just see what everyone was up to.

However, it made me more conscious of my own work. I ended up deleting my past failures and tried to keep my workspace as 'tidy' as I could. An advantage of leaving models that did not work is just to be able to look back and avoid past mistakes, but as everyone could see it – I felt more aware that anyone could have insight into it. That made me feel more conscious about my failure and made me delete it.

I decided to experiment and try new things in a single player world and try to transfer only the working models onto a shared realm. I felt more comfortable to experiment and try new models in my own realm. I preferred having a private workspace as I prefer not having a feeling that someone is 'looking over my shoulder' as I

worked. Although I did like having access to both areas of work and I interchanged for the reasons mentioned above.

Like Student B and Student A, Student D finds the Realm experience mixed. Unlike Student B, the exposure made them remove things much more readily, and there is only a small element of their work in the Realm. They do identify the potential benefits of having work there to be looked at by others, but also the uncomfortable feeling of being judged by their peers. I suspect this is partially because we only met as a full group once before commencing on the project, and that knowing each other more may have helped with feeling more able to ask for help or keep failed attempts up. On the other hand, unfamiliarity with classmates is common at the University of Edinburgh, where large cohorts and multiple module choices mean people often are in groups where they do not know anyone well.

The mixed feelings about the Realm present something of a paradox. If we hadn't included it, people could have built independently but perhaps would not have been inspired to go as far and complex as they did. But the presence of the Realm was also detrimental to the focus and confidence of some group members. This paradox was interesting from my own perspective as I also would have preferred to work outside the Realm before putting more polished work into the Realm. As the team leader however I needed to lead by example, and thus held to two rules: to always work in the Realm, and to never delete bad, incorrect or inconsequential structures. This self-imposed forcing did help me to feel more comfortable with failure, as I found myself writing notes in-game as to why it failed, and also giving me a history of attempts I could move through in space.

This physical history the Realm (and Minecraft more generally) was meant to provide students uniquely with sense of their progression (and possibly that of others) through a problem, pointing to the failures, the nascent ideas and early model prototypes as well as the successes. It can be thought a little bit as "showing your working".

This physical presence also has interesting links to the strengthening of memory, not only for the concepts engaged with, but the process of how that understanding evolved and developed. Early mnemonic devices such as the method of loci (or "memory palace") explicitly link the navigation through a mental space to the enhancement of encoding and recall (Yates, 1966). The spatial environment helps to anchor imagery by making use of multiple senses, and aids retrieval through the evolutionary link between episodic memory and spatial navigation (Buzsáki & Moser, 2013; Buzsáki, Peyrache, & Kubie, 2014; Fitzpatrick, 2016; O'Keefe & Nadel, 1978). Minecraft worlds are an externalisation of this mnemonic device,

reducing the cognitive load necessary to initially generate these environments and minimising the risk of extinction or noise by being immutable except via direct and sustained conscious modification.

These advantages however are not tied directly to the Realm idea and could be implemented in a single-player environment. The Realm reduces the risk of permanently losing building attempts through accident or wilful destruction thanks to automatic backups of the space, but as the students have expressed can be difficult to get used to. How to blend the use of individual and shared spaces is outside the scope of this Chapter, but it is likely to be one which requires time to become more comfortable with displaying of work-in-progress and failed structures to peers and educators, and teaching in how moving designs between individual to shared spaces can be beneficial.

Time spent and time management

The success of this pilot required that the time taken to achieve a working model did not become excessive. Students have many time pressures, both within and outwith academia, which this method should not add additional strain to. To help explore this, we asked students to track their time spent and try to keep to 12 hours a week, which could be done at any time.

I felt like I spent a great deal of time just trying to figure out what I was doing. So, in essence, I would work for longer than the 12 hours a week, but it would not directly translate into a model or much progress. [...] I liked the flexibility of being able to work when I wanted. Yet, with a fixed 'work' time, I probably would have worked on it more productively.

I found that 12 hours was a reasonably lengthy amount of time when researching CAMK2 [sic]. I personally elected to work in 2-hour bursts, as beyond that point looking through papers became tiresome. In the end, I likely spent 6-8 hours on the research stage. Spending less time on this stage was beneficial in providing me with more time to work on the designing and making of my model.

The flexible nature of the study allowed me to procrastinate in the week, which meant I'd binge on the work, doing it in 4-6 hour shunts as opposed to (the reasonable strategy of) setting aside a few hours a day to work on my projects.

The 12 hours a week was a very manageable and appropriate timescale and although almost every week I worked more than the prescribed it was more out of enthusiasm and taking on relatively ambitious projects.

For all four students, the time was not seen as onerous. Each also shows very different approaches to how they occupied this time, benefiting from the freedom to do the tasks at their own pace. Indeed, in most cases they all happily mention they exceed the recommended time, but out of personal choice and enthusiasm to make the best possible model. There was no obligation to *finish* any of the models in this pilot, but all four students felt keenly that they wanted to do so.

These findings are constrained by the fact the pilot study took place during the summer, when no lessons or assessment were ongoing. Structured or planned times for working may help with providing a scheduled time for students to work which would be supervised by educators either via email or within a Realm environment. This importantly however would not remove the ability to work outwith those hours but may be useful to those who find time-management or motivation a struggle.

Another important factor in this pilot was that the work was only marginally constrained by myself, and is more representative of a long-form in-course assessment or dissertation project. More constrained projects or pre-built worlds may vastly reduce the number of required hours to see benefits from this approach, as too would group work.

The fact that many students exceeded these hours showed that they had set themselves personal challenges and aims for the models, and wanted to show to myself and the group a model which met those personal goals. This I think arises from the fact each model was personally developed, with all research, planning and building being solely theirs, even if on the same overarching concept. There was pride in being able to show how they mapped theory into a practical, virtual object.

What teaching potential do you see this having, if any?

The team had a diverse range of opinions on where they would implement using Minecraft in teaching, which tended to focus on what area would be most strengthened by the approach, rather than its likely failure in others. Student D and Student C both focus on its use for improving the knowledge of the model creator:

In my experience with the whole research project I had to deal a great deal of research from peer-reviewed papers and journals and other university level resources and build a thorough, in-depth, scientifically relevant understanding of the basis behind my individual projects [...]

[...] even if the model is not very complex that a person with not much experience makes, it would help for their own personal learning. This is as they are spending time critically thinking about the various aspects they need to address or show in their model. So to others, it may not be complicated or understandable; to the creator, it is a way of reinforcing knowledge.

Student A and Student B gave multiple possible ways of incorporation of this approach. Student B in particular wrote at length on their own feelings for how well Minecraft could be used in a range of pedagogical settings (one shown below):

Tutorials:

With the amount of content typically covered in a tutorial, I can see some use in using Minecraft to take students (in groups, or as a large group perhaps) through a set of Minecraft-contained representations of concepts over the course of an hour. Here the students could experience the concepts dynamically, engaging with animated content and maybe editing bits of it themselves to see the results. This might make them respond more positively to concepts they'd usually feel divorced from. The students could then be provided with a link to an unlisted YouTube video that shows the Minecraft representations that they'd been working with. That way, if they wished to go over the material again, they could in video format. Map files could also be provided on a VLE for more Minecraft-savvy students.

Both show a mature and sensible approach to this question, and highlight that whilst flexible, this approach is not necessarily ideal for all types of teaching. In particular, the question of familiarity and being able to take time with the approach is mentioned by both throughout. There was some more critical analysis of where and how this approach could be implemented with this in mind:

[...] students with limited experience will need time to improve their confidence with the software and familiarise themselves with the available resources. Given this compromise, a project such as this would exceed the allocated time for a tutorial or a lab and if this were to span multiple sessions, a student would then need to be credited for this effort.

From the perspective of a learner, the approach is potentially quite useful, but not as a one-off, and appears to be better suited to much longer-form or continuous activities. It would need to therefore become embedded into courses across years to reach its full potential. This is not necessarily new – many approaches which

are unfamiliar to students, such as facilitations⁹¹, require being exposed to multiple times in order to match to how educators expect them to function and work for students (Biley & Smith, 1999; Cooke & Moyle, 2002; Lekalakala-Mokgele, 2010). This is in part because a learner needs to see the benefit of the approach, which for them would be reflected in better understanding, which leads to better grades. What is promising from these replies is that from just a small exposure, students can identify freely several possible uses that could be both passive (videos of models), active (projects, tutorials), and where it may not (practicals).

Engagement

The students were not prompted on how the approach might work in terms of engagement, both personally and for students generally. All however wrote in multiple instances about this aspect, in a mostly positive manner. From the personal perspective, we have already seen that all the students wanted to succeed in building the models they planned to, regardless of the time needed to do so. There was also a strong sense that the structure of the first model brief led to them researching much more effectively and comfortably than they perhaps would for a course:

I did find that my understanding of any related topic was reaffirmed as I attempted to build my model. I think that having to creatively represent numerous aspects of a design forces you to better understand what you are designing.

In my experience with the whole research project I had to deal a great deal of research from peer-reviewed papers and journals and other university level resources and build a thorough, in-depth, scientifically relevant understanding of the basis behind my individual projects and with the right facilitation I believe Minecraft could be used in addition to self-study in order to build creations which as a side-effect require and build an extensive base of scientific knowledge.

I believe that my personal understanding of the molecule was growing. I was starting to understand the potential benefits and applications of such a task

⁹¹ Facilitations at Edinburgh area variant of problem-based learning scenarios (D. F. Wood, 2003) where students take on an equal role in shaping discussions on both application of learnt skills and content (e.g critical analysis of research papers) and on skills for learning (essay writing, using and giving feedback, etc.).

All the students also mentioned how the approach might be useful to their peers. Student A focussed on how Minecraft modelling might act as a way both introducing modelling and incorporating more active learning into undergraduate courses:

Generally, modelling opportunities are minimal in my degree, most likely because the software is too expensive or too difficult to use. Traditional learning styles are not ideal for learning the increasingly complex structures and processes we are faced with. We need a more active learning style and I feel that Minecraft may be the refreshing learning experience we need to bring in this new learning style.

Student D's comments also reflected this desire for more interactive and modern ways of learning to be better integrated into undergraduate teaching:

Often web diagrams are used to portray the workings of this molecule; however, these commonly are hard to interpret and make the understanding of the model quite challenging. [...] This may be a more engaging platform and method of knowledge creation. It is a creative outlet that provides an opportunity for the user to engage with molecules in a comfortable environment directly.

When mentioning this project to anyone – there was a certain aura of intrigue and interest. Everyone was always curious about how Minecraft can be utilized at an undergraduate level and to see if it was actually useful. As aforementioned, everyone wants to learn in exciting ways rather than just learning off PowerPoints. This provides a uniquely interactive way to learn.

Student B has a much more measured and cautious take on how Minecraft might be received.

Minecraft can't replace tailor-made protein visualisation software like PyMOL or be a perfect stand-in for static diagrams or molecular animations. However, it can provide easy access to ideas and represent them in a different way to normal, in ways that could engage visual and interactive learners more than their less visually inclined colleagues.

I told you at the time that I really wanted to see if we could make this work, and I believe Minecraft does have potential in a limited range of applications. especially in letting younger students access biological concepts usually kept away from them due to their complexity.

I think Student B's take highlights that the most immediate benefit of this approach would be not as a replacement, but a parallel enhancement of pre-existing methods of teaching biology. It offers a creative and accessible space to either generate or interact with abstracted versions of biological concepts that need not sacrifice complexity, and so help in providing a much more robust theoretical basis for late undergraduate work.

It harnesses much more directly the fact that younger people are used to technology which is extremely responsive, user-friendly, and multi-platform. Educational technology does not always meet these requirements, and the time a student spends with it is likely to be insufficient to make the most use of it, given these are usually from experience tightly regulated either by time available to interact with them, or by requiring students to work on specific computers to access.

Comparisons to Game-Based Learning Literature

In this project, I used a game-based learning approach to address a problem within biomedical sciences. This was at the time however, still quite exploratory, and so no formal analysis of the literature was conducted prior to the pilot. In placing these findings into the wider field, I have identified several related investigations that help with interpreting my findings.

Self-Directed Learning

Self-directed learning was a crucial element of the MinecraftED pilot. Undergraduates are expected upon completing their degree to have the skills necessary to identify their needs, set goals, and know how to go about acquiring the information necessary to learn new expertise (Barr, 2017; Barrie, 2006; Green, Hammer, & Star, 2009). Video games have been identified by several researchers (collected in Coleman & Money, 2020) as offering a familiar context in which self-directed learning occurs during the course of playing the game.

Using extensive coded analysis of interviews and video reflections of experienced video game players, Toh & Kirschner (2020) identify five factors that are important in facilitating self-directed learning in a game environment:

1. Learning analytics and data visualisations
2. Gradual release of new information over time
3. A safe space to try multiple approaches and learn from failures
4. Defamiliarisation mechanics
5. Scaffolded learning

Reflecting on the use of Minecraft, points 1, 3, and 5 are the most important. The group received constant visual feedback on how their models were working/the level of completion, and were in a risk-free environment to try many different methods of reaching their self-set goals. This risk-free perception was tempered by the multiplayer space, but students did identify offline areas to trial different ideas.

Scaffolding is the one area not sufficiently appreciated during this pilot. A scaffolded learning environment is one where a learner is helped to identify their own expertise, and to build on this with the aid of someone with a greater level of experience, or as part of a group. This approach requires dialogue, an appropriate context for the learning, and an “artifact” – an object, piece of knowledge or method that is interacted with to enable the creation or reinforcement of knowledge in the learner(s) (Wells, 1999). In this pilot, the capacity for dialogue between members of the group was possible, but rarely used. Interactions often were weekly, and only between myself and one other group member at a time. As we had no set times to work simultaneously, the chances of working within the shared environment with other group members was low.

In instances where members were online together, dialogue did occur, which was only between myself and Student B. If present, I would talk using the in-game text feature and ask about progress, what had been done, and what was not working yet. Student B could then demonstrate their models and receive feedback immediately, then we would reverse roles. For Student A and Student D in particular, the absence of this dialogue likely contributed to the feelings of insecurity or frustration respectively with their work, and discouraged asking for advice or even direct help from other group members.

Points 2 and 4 may still have roles to play in future uses of this approach. Defamiliarisation is in Toh and Kirschner (2020) a way in which games force players into atypical situations where they have to reappraise their situation and how they interact with the game to progress. As a sandbox game, it is difficult currently for this writer to determine just how such an approach is not actually an inherent part of the experience of playing Minecraft normally.

Point 2 refers to a “just-in-time” deployment of information to users that I reflected as potentially detrimental in Chapter 5. Toh and Kirschner’s argument is that this promotes self-directed learning, as it reduces the risk of overwhelming the player and allows them to work up to new skills, concepts etc, at their own pace. At the onset of this pilot I was not in a position to consider this, given I was a beginner with the software. But having in place a written resource pack to introduce new methods and skills in-game when requested would especially

benefit less experienced users of Minecraft. The information a user needs to solve an in-game problem is available online, but these are often fragmented across various websites. Being able to gradually increase knowledge without the extra load of sourcing that knowledge may improve self-directed learning of the biological content through reducing time taken solving implantation issues.

Another, international qualitative analysis of self-directed learning specifically in Minecraft in secondary education identified three specific themes: context manipulation, storytelling, and defining roles within group work (Faas & Lin, 2017). The potentially overwhelming freedom of the software was quickly compartmentalised by students taking on specific jobs, such as builders, miners, and explorers. Within each role, a student is still free to modify the world (the 'context'), but with more manageable constraints. Storytelling was identified as a powerful motivator, especially when done in conjunction with modifications to the base game to tailor the game world for a specific task.

In the process of designing the Realm world layout, I did introduce non-essential elements that gave a sense of something hidden or already present in the world when the students began the study. This included a temple-like structure, with hidden rooms and passages, hidden homes both on land and in the air, and sculptures. These were in part my learning of the software, but also were on reflection nascent forms of storytelling. They offered a sense that the world space could evolve into some parallel narrative that might in some way relate to the structures we were building⁹².

Digital Game-Based Learning

A small body of literature has accumulated on the use of video games as a method of delivering self-directed learning⁹³. A recent systematic analysis of this literature (Coleman & Money, 2020) identified that much of the work in this field is at the primary level, within classroom environments, and qualitative in their analysis. Games are often specifically designed for educational purposes, and run on desktop/laptop computer systems.

Ideally, these analyses of game-based learning are longitudinal, taking onboard the reflections of multiple student cohorts. In the one paper that did so, they tracked two classes of secondary school students over three years (Bate, MacNish, & Males, 2014). Games for these two groups were never used for complex or new

⁹² This is something I return to later in this Chapter with Project Millport.

⁹³ The reader should note I take the terms 'self-directed' and 'student-centred' to be identical in meaning for the purposes of this Chapter. The theoretical literature does not to my knowledge make any meaningful distinction between the two.

learning, instead focussing on consolidation or reward for task completion. Another key message was the element of risk in bringing in games – that educators would lose control over the pedagogical setting. I will not read into this too much here, but the reluctance to adopt digital game-based learning at secondary and tertiary stages may be partially generated from educators being unfamiliar how to make these approaches fit with desired learning outcomes, and being overwhelmed by the design choices available (Lameras et al., 2017).

Interestingly, many of the papers were assessed to focus on single student learning missing the opportunity for collaboration and co-operation. This occurred even in multiplayer game studies, making it difficult to assess how well the digital approach works in group environments. Whilst my pilot study did not prevent collaboration, and anticipated it occurring, it did not promote it. Doing so, for example through a small problem-solving exercise, may have helped with bringing out this dimension in the pilot.

Perception of Using Video Games for Learning

The ability of video games to facilitate knowledge and skill acquisition is only useful if learners are receptive to this atypical use (Barr, 2018). An exploratory study with Flemish secondary school students found that they would be amenable to the idea of game-based learning, so long as these were easy to use and were ‘useful’ (Bourgonjon, Valcke, Soetaert, & Schellens, 2010). This ‘usefulness’ metric was probed as being able to identify as a learner if the game was improving their exam grades, or their ability to learn. Others (collected in Ak & Kutlu, 2017) have identified that likability or engagement with the digital environment is also important – it has to be well constructed. Using established commercial games benefits from accommodating most of this likability factor, and can focus more on ensuring learning goals are met.

Interviews with 20 Scottish undergraduate students on their perceptions of multiplayer games being played as intended in a controlled environment showed an appreciation of how the experience could reinforce communication, adaptability and resourcefulness attributes (Barr, 2018). Students also notably enjoyed the experiment, even though they were diverse in terms of gaming experience, although one should note the retention rate was not specified⁹⁴.

A higher level example of the success of this approach with undergraduate students involves using the Second Life (secondlife.com) platform to educate 3rd year medical students in radiology (Lorenzo-Alvarez, Rudolphi-Solero, Ruiz-Gomez, & Sendra-Portero, 2020), and is built on a series of previous uses of the

⁹⁴ It is likely to be about 50 % however; see (Barr, 2017).

software for medical educational purposes by the same group (Lorenzo-Alvarez, Pavia-Molina, & Sendra-Portero, 2018; Lorenzo-Alvarez, Rudolphi-Solero, Ruiz-Gomez, & Sendra-Portero, 2019; Lorenzo-Alvarez, Ruiz-Gomez, & Sendra-Portero, 2019). The authors created a competitive game within Second Life, which ran for one month. Students could interact with a bespoke environment featuring information panels containing radiological content. These were present for four days, before being replaced with a series of test panels. Students were ranked after two days, and the lowest scoring participants eliminated. This continued for a total of six rounds, leaving 31 winners out of 90 participants.

Participants scored significantly higher on subsequent tests than a control group, although there was no significant impact on their final grades. Users did find the idea engaging and perceived it as helpful however, especially the competitive gamified nature. Given the material used in-game was identical to what would be normally presented in lecture or in online course material, this enthusiasm has to be attributed to both the competitive game aspect, and the game environment itself.

Minecraft-Specific Studies

Unlike many other commercial games, Minecraft has had a dedicated *Education Edition* since 2016. This version of the game is based on the Bedrock Edition and offers a series of education-specific features that benefit educators using the platform to teach. This format is targeted largely at primary and secondary school educators who wish to use the software in the classroom, with a dedicated website, community, and training programme⁹⁵.

The changes to the game's function emphasise using the software for collaboration, data collection, and in-game information transfer. There is also much greater control possible by the educator over the student's ability to manipulate the environment and explore, providing in-game controls that normally require knowledge of programming or hosting a server to achieve. The most recent iteration of the software is also compatible with Chromebooks, increasing the accessibility and flexibility of using the software for teaching.

A notable study into the potential benefits of Minecraft for tertiary teaching used an earlier, educator-built version of Minecraft called MinecraftEdu⁹⁶. The researchers built a virtual "Atlantis", which featured various buildings, zones, and circuitry puzzles to help in understanding logic gates and binary counters

⁹⁵ <https://education.minecraft.net/>

⁹⁶ This is the precursor to Minecraft: Education Edition, and many of its features were carried over into the official version of the software.

(Nebel et al., 2017). Students were split into 3 groups – goal free, learning goal, and performance goal. Students with defined learning goals expressed more fun and a more manageable cognitive load after 3 hours of play, although no measurable change in learning was observed. The goal-free group, which Minecraft is by default, were much less likely to remember the instructions they were given, and found the experience less fun, suggesting constraints are a necessary feature for maintaining engagement.

Tertiary-level use of Minecraft within courses is much less visible, which may in part be due to a lack of awareness of the availability of the software in Institutions' Campus agreements with Microsoft. A search of "Minecraft" and "undergraduate" in the PubMed, ERIC (Education Resources Information Center) and the British Education Index found no published reports of teaching with Minecraft within an undergraduate course. Searching in Twitter for the same terms found a surprisingly large number of campus replications in Minecraft, but only a few instances of evidence of use in tertiary-level teaching.⁹⁷

⁹⁷<http://www.ilearninguk.com/molcraft/>; <https://createandexpand.co.za/2019/09/03/using-minecraft-education-edition-in-the-undergraduate-occupational-therapy-curriculum-at-the-university-of-the-witwatersrand-in-johannesburg-south-africa/>

Epidemiological Modelling in Minecraft

The evidence from pilot studies into using video games in teaching suggested that exploring the ability to use Minecraft in biological education would benefit from a longer, more focussed project. Along with my supervisor (Dr Stefan) I offered an Honours project titled “Building Biological Models in Minecraft”. The description detailed our broad intentions to explore whether Minecraft could be used to bridge the gap between computational biology and more traditional biological investigations. For the student, the aim would be to construct and reproduce a biological concept or experiment that could be used for both teaching or research. For myself, I wanted to assess if the success of the other student modelling projects (Chapter 5) translated into Minecraft, and compare the two for their potential in future student projects.

Unlike the previous models, there was no specific area of interest provided to the students. Instead, the concept would be chosen by the student through initial discussions with myself and my supervisor. As the project required an unorthodox methodology and a high chance of failure, allowing the student to choose allowed them to select based on their specific field of interest, and on topics which they may have already have a solid understanding. This allowed us to separate out the difficulty of the topic more easily from the difficulty of implementation.

Our project was chosen by one student, Siobhan Vickerstaff, who was specialising in Infectious Diseases. From our initial discussions Siobhan settled on the reproduction of the Susceptible, Infectious, Recovered (SIR) model of disease spread (Hethcote, 1976). An SIR model attempts to explore how a disease spreads through a population mathematically. In its original form, each entity in the model may be Susceptible (not infected but not immune), Infected (able to pass the infection onto others) or Recovered (no longer infected and immune to future infections of the same type).

Given the limited time of the project, Siobhan’s aim was to reproduce the original version of the SIR model in several different forms. The first was to construct the SIR model in a 2D grid pattern. This ‘lattice’ model represents each member of a population as a square on a regular grid (Rhodes & Anderson, 1996). Each square may possess the susceptible, infected, or recovered attributes, represented by a colour.

The model is initialised with all squares susceptible except for one or two, which are infected. The rate of spread of this infection is determined by an infection rate and a set of rules about when such an infection may occur. For example, a


susceptible square may only become infected when at least two adjacent squares are infected. The infection time is also then determined by a rate, upon which a square adopts the recovered attribute. In this form no reinfection is possible.

A simple version of the SIR model was available from Dr Melanie Stefan in Python, which demonstrated this approach, and was used as the template for the Minecraft version. This instantiation involves an 18 x 18 array of command blocks constructed 18 blocks from ground level. Each command block contains code which uniquely identifies it with a scoreboard objective.

Objectives are a way of storing information about in-game actions, which are assigned to entities. These are normally to a player name but can also be for any entity so long as its universally unique identifier (UUID) is used. The player name does not have to be any 'real' player and is therefore useful for assigning objectives that require careful controlling or use across multiple players. Objectives are numerical and are incrementally added to when the proposed trigger for an objective is met.

Standard built-in objectives match to expected uses of the software – mining, breaking, using, or crafting objects, for example. There is however a special class of objective called 'dummy'. This class is only modifiable using commands, and therefore fully controllable by the developer. This allows for the value to be set to any value once certain conditions are met, which may then be used to trigger a command sequence.

Above the 18 x 18 array of command blocks is a layer of diorite blocks, which represent "susceptible". Other layers of command blocks set these blocks to "recovered" (grey concrete) or "infected" (red wool). To change from susceptible to infected, two conditions need to be met. The first is that a neighbouring block is detected that is infected, and that there is power available for the command blocks to initiate a change. Power was provided stochastically through the integrity value of a constantly updated layer of blocks. The existence of a block activated the command blocks once. If the right conditions were met, then the command would execute, and the topmost block of the whole model would change.



B



Figure 77 Siobhan's Minecraft SIR Models.

(A) A replication of a simple SIR model in Minecraft shows a similar but significantly different profile over time when compared to an existing Python version. Each model was run 10 times for 30 evolutions. (B) Two videos showing both the lattice (left) and villager (right) models in action. The lattice video also shows the Python model.

Videos are available at

<https://www.youtube.com/channel/UC9-gPbLo8aUT8RPpw3eUp0A>

or Github https://github.com/yishusing/Thesis_Repository

Two versions of this were made for the infected and recovery parameters. An observer looking top-down therefore sees a spread of infection through the lattice over time, with recovery happening in a delayed wave behind this (**Figure 77**). A user can also modify these coarsely by adjusting the integrity of the block layer used as a coincidence detector, with higher values corresponding with a higher rate of infection/recovery.

Villager Model

The second model variant was built to explore using the in-game entities to generate a more particle-based model. Villagers have several properties which lend themselves well to being used for this. Game entities have a rudimentary artificial intelligence (AI) which respond to environmental changes, nearby entities, and player interaction. The level of this AI varies, but some entities (including villagers) have a quite complex AI that can be manipulated.

Siobhan used this feature as a way of creating a spatial infection model using a human-like population. To do so, villagers were summoned wearing different hats – the colour of the hat determined if the villager was susceptible, infected or recovered. Villagers do not commonly interact with one another on short timescales however, so additional changes were necessary to show how a disease may spread through close contact.

To do so, Siobhan made use of the zombie villager entity. Upon infection by a zombie “bite” (really just contact of the two entities), a villager may turn into a villager zombie. These entities will then chase other villagers, and upon contact with a healthy villager will turn them into a zombie villager also. These entities have an attribute which determines the time they remain as a zombie villager, before turning back into a villager.

Siobhan used these features to allow for the transitions between susceptible and infected, infected and recovered. All villagers begin as susceptible. Upon activating the model, one of these villagers (the one closest to the centre of the arena) is “infected” and changed into a zombie villager (**Figure 77**). Upon contact with a susceptible villager, they too become infected. All infected villagers have the same decay rate of infection. Once this ends, the villager is deleted and immediately replaced with a villager wearing a green turtle hat, therefore distinguishing them from the susceptible (no hat) and infected (straw hat) variants. These villagers are immune to infection.

This end point may be varied however by activating one of three command blocks on a control platform. The villager can be killed by the infection (illustrated by an immobile skeleton) or return to a susceptible state. In this manner we can

introduce a few simple variants to the SIR model, which in turn exemplify different diseases and their progression through a population.

Siobhan's models show two things. The first is that mathematical modelling can be achieved within Minecraft without any major changes to how the software operates normally. These models are buildable without any specific programming knowledge, have error messages, and visual feedback. Importantly, they also resemble a "true" mathematical instantiation, albeit statistically different to the original model. The second is that these models can be generated in very different forms through clever manipulation of pre-existing code and AI behaviours. This is important because it allows for the same software to present models that are abstracted (lattice) and more concrete (villager). This not only allows for a user to compare the model outputs, but also see how abstracted conceptualisation can be made more visually comprehensible, and how this changes both *what* and *how* data is collected.

Computer models can therefore be presented in visual ways in a familiar environment to many young learners, in ways which allow them to dissect and understand how the models operate both theoretically and mechanically (i.e., the code). This approach renders the same models useful to a diversity of audiences at different levels of analysis, which can be supplemented either in-game or externally through text, lectures and discussions.

Neuron Safari

World File: Neuron Safari

Website: www.neuronsafari.com

Using public engagement or outreach was not something that the pilot study set out to explore. However, given the obvious success of the pilot in bringing complex processes into the software, I decided to see if I could develop an experience in Minecraft related to my own interests in neuroscience⁹⁸. I conceived the notion of a virtual tour of a neuron. This would involve the construction of a cartoon neuron in-game, and three interactive stations exploring intra and inter-cellular functions. The DNA Factory would be one station, with the others covering intracellular transport and neuronal communication. Users would be allowed to investigate each station and interact with a game or animation that would offer a visual representation of the phenomena. This would then be further explained by an academic researcher sat next to the user.

This guided tour aspect led to the name “Neuron Safari” being adopted, as it pointed towards discovery, highly visual content and the focus upon neurons and neuroscience. This idea was subsequently proposed in September 2019 for the 2020 Edinburgh Science Festival. This was accepted as a “mini-workshop”: a 20-minute ticketed activity that should be aimed at children 7 – 12 years old.

Initial construction of the Safari envisioned all the material taking place in a giant neuron. The DNA Factory would be central within a nucleus, with a complex network of minecart tracks surrounding and extending to the neuronal membrane. These would represent the protein transport network. Users would then travel down a large axon to watch an animation of synaptic vesicle exocytosis whilst being informed about how information transfer occurs through electrochemical means.

This first idea however did not fit well with the Scottish curriculum guidelines (“Curriculum for Excellence: Experiences and Outcomes,” 2008). Our target audience ran from Primary 3 through to Secondary 1. The concept of cells is not introduced until Level 3 (S1 upwards), so some of our learners would have required a quite dense exposition to have even begun engaging with the content.

To overcome this, I reconceived of the Safari to take place in different “zones”. The beginning three zones would give the user a simple introduction to the idea of a neuron by gradually moving down a size scale from head to cell. Once complete, the user is ready to explore three interactive intracellular-focussed zones (the

⁹⁸ One student (Student B) did mention the possibility of using Minecraft for outreach, but their report came after the conception of the Safari.

DNA Factory, protein transport and cell communication). This idea had the added benefit of turning the activity construction into something more modular. This meant each zone could be kept physically distinct and grow to what was required, not what was limited by space constraints.

General Construction

From the onset, the aim was to keep the model as close to biological reality as possible. Whilst the activity was requested to be aimed at 7 – 12 years, there was no need for the activity to be necessarily restricted to this for future events. In addition, I wanted to allow users to take the activity home with them if they had Minecraft and share it. Thus the more that the activity was close to being veridical, the better someone using the activity unsupervised could take away important information and avoid misconstruing elements.

The first three zones were therefore designed to incorporate real structural data of skulls, brains and neurons respectively. This is feasible because most scientific visualization data is stored as voxels. The term voxel derives itself from “volumetric” and “pixel” – it is a single point data value within a regularly spaced 3D grid. Structures are determined by the relative relationship between voxels. Minecraft’s “chunks” work in a similar manner, and thus opens the possibility of transferring scientific data in this format into the software.

This data is not directly insertable into the software however; it requires conversion into a *.schematic* file that utilizes the Named Binary Tag (NBT) format of data storage⁹⁹. Data is stored in a tree structure made up of various unique tags. This data format was created for Minecraft and is therefore rarely used elsewhere.

A conversion pipeline was therefore required to allow for the transfer of biological structural data. Tinkercad (<https://www.tinkercad.com/>) is a free CAD software made by Autodesk for use in education and hobbyist circles for 3D printing. It allows for the import and export of *.stl*, *.obj* and *.svg* files, as well as the export of *.schematic* files. This covers most common structural 3D model types, and most specialized software will allow for conversion into these types. Some data required a more complex conversion process (see **Figure 78**).

⁹⁹ <https://web.archive.org/web/20110723210920/http://www.minecraft.net/docs/NBT.txt>




Figure 78 Data Conversion Pipeline for importing neuroscientific models into Minecraft.

The process by which neuron morphological data from the NeuroMorpho database is converted into a readable format for Minecraft. TinkerCAD is a free to use online simplified CAD environment created by AutoDesk (<https://www.tinkercad.com/>).

These structures were then insertable in-game using World Edit, a modification allowing for large scale command line changes to the game world. I used the *schematic/load* and *paste -a* functions to insert the major structural files. The DNA Factory was inserted using an in-game capability known as structure blocks. These can save a maximum 32 x 32 x 32 game region as an *.nbt* file that may be transferred between game worlds. These files are placed into the `\.minecraft\saves\[world_name]\generated\minecraft\structures` folder.

Other structures were built by hand using the software's Creative Mode. This provides an unlimited resource of all in-game blocks for easy and quick construction times. It also allows the use of structure, barrier, and command blocks, which were necessary for some of the more complex builds.

The world for the Safari is in the Void biome, which lacks any procedurally generated terrain. This served three purposes. The first is it maintains user focus on only those structures necessary for the activity, reducing distractions. The second is it improves loading times across large viewing distances by reducing the number of blocks needed to be generated in each zone. Finally, this biome does

not generate entities automatically, which might otherwise have interfered with the activity¹⁰⁰.

Further *gamerules* were specified to prevent any random weather events (*gamerule doWeatherCycle false*) and command block event outputs (*gamerule CommandBlockOutput false*). The world spawn point (where the player would be reset to upon any character death) was set to the start co-ordinates of the skull zone (*setworldspawn -854 137 30; spawnradius 0*). As each zone was quite far apart in space, some were beyond the chunk loading distance. Users teleporting between these sections of the world would sometimes become stuck as the world did not render in time for the player to be correctly placed. To overcome this, the starting position of each zone was forced to load upon game startup using *forceload add [from] [to]*.

Coding

The command block allows for a more controlled use of the software coding language, executing commands if conditions are met and not requiring manual input. There are three “flavours” of command block: impulse, chain and repeat. Blocks may be further adjusted to be *conditional* (firing only if a preceding command block command was activated) or *unconditional*, and to either *require redstone* input or be *always active*.

Impulse blocks fire once upon activation. Activation can occur from directly active power components such as levers or redstone block, or a block adjacent to one of these active components. Redstone circuit components such as dust, repeaters and comparators may also activate a block. Always active blocks do not require these inputs. Activation allows for the single execution of an impulse block; if the input remains ON the command block will not fire again.

Chain blocks are by default *always active*. They only fire if they are triggered. If a command block facing the chain block has fired, regardless of whether this was successful, the chain block will be triggered. If *conditional*, the command will not execute unless the preceding command was successful. Regardless of success of this or activation, it will trigger the next chain command block if present. A line of joined chain commands will all execute simultaneously in the order placed, acting similar to a loop structure in more traditional code.

¹⁰⁰ This may also be done through the *gamerule* command as *gamerule doMobSpawning false*, but this can be easily switched back by inquisitive and knowledgeable users



Figure 79 Neuron Safari Zones.

Each zone has its own specific style and concept to be explained, with the exception of the Library Zone, which acts as an information hub and a place to teleport to each other Zone. This area is hidden in the Safari, with access possible from the Cell Communication Lab.

A repeat block's default is always *active* and *unconditional*. This block will fire every 0.05 seconds, corresponding to the software game tick. If all the command conditions are met, the command may successfully be executed twenty times a second. *Conditional* variants fire only if the previous block was successful. *Needs redstone* variants will fire repeatedly only as long as powered by redstone.

World Structures

Given the scope of this activity, I will break down the description of specific structures into each of the zones present in the world (**Figure 79**). The world file is available at

https://github.com/yishusing/Thesis_Repository.

Skull Zone

The skull used was taken from the National Institutes of Health (NIH) 3D Print Exchange (Coakley et al., 2014) and was submitted by kserralde¹⁰¹. The skull is that of a healthy child and was chosen to better fit with our target audience. Users enter the skull by choosing a minecart rail track from three options. Choosing an option by pressing a button summons a minecart on the correct track (and deletes any others currently present). The tracks take the user into the cranial cavity through either of the eye sockets or the nose. All three converge towards the back of the skull, where the user is teleported whilst within the minecart to the Brain zone.

Brain Zone

Like the Skull zone, this zone consists of just one major structure. The brain was taken from cadnav.com, submitted by an unknown author¹⁰². The right hemisphere was modified in WorldEdit using the *replacenear* command to colour-code the four major lobes and the cerebellum. Users begin already in a minecart at the back of the brain between the two hemispheres. They travel out and around the exterior brain surface before entering inside the brain at the temporal lobe. Here they are deposited at the top of the brainstem, where there is some reading material and a transporter to the Neuron zone.

Neuron Zone

The neurons used were all taken from NeuroMorpho.org (Ascoli, Donohue, & Halavi, 2007). There are four human pyramidal neurons [NMO_01052 (K. K. Watson, Jones, & Allman, 2006); NMO_01074 (K. K. Watson et al., 2006); NMO_03476 (Jacobs, 2001); NMO_01706 (Hayes, 1996)] and one mouse retinal ganglion neuron [NMO_75925 (El-Danaf & Huberman, 2015)]. NMO_01052 was

¹⁰¹ <https://3dprint.nih.gov/discover/3DPX-009510>

¹⁰² <https://www.cadnav.com/3d-models/model-44903.html>

modified to include a teleportation area to the next zone and a minecart track running down the inside of the axon, which acted as a transient zone between the Protein Transport Network and Cell Communication Laboratory. NMO_01706 was modified to have repeaters and redstone dust on all dendritic and axonal branches to show the direction of information flow. Users move again by minecart on rails around these structures, centered around NMO_03476 before entering NMO_01074 at the cell soma.

DNA Factory

This zone is primarily comprised of the DNA Factory, a model designed by Simon Parker as part of the MinecraftED pilot study. The factory allows users to input a short DNA sequence, which is subsequently transcribed and translated at two other sites. This was imported via three *.nbt* files. There are three distinct DNA structures which surrounded the DNA Factory. Two are from the NIH 3D Print Exchange authored by fmontecinosfran and depict coiled DNA with and without a histone¹⁰³. They are adapted from Tachiwana et al., 2011. There is also a more traditional representation taken from the University of Illinois Library¹⁰⁴. The latter was used twice, once vertically to complement the Factory representations and once horizontally to transport users into the zone via a short minecart track.

A recent addition to this zone has been three amino acid sequences taken from the UniProt database (The UniProt Consortium, 2019) of Human CER1, DACH1, and WFIKK1. These relatively small (< 800 a.a) sequences were chosen because they fit into the game world, they have eye-catching 'common' names, and are not related to any known disease type. Their sequence was converted into the block representations shown in the DNA Factory.

Users enter the zone, take the short track to a teleporter, which then takes them to the DNA generator. Here a user can input their own short (15 bp) DNA sequence. Users then travel forward to reach the Transcription (inputted DNA -> mRNA animation) and Translation (mRNA -> correct amino acid sequence animation) machines. Once done, users can travel to the next zone via a walkway which takes them past the three example protein sequences described above. The final Translation machine was not fully working in the first two public runs of the activity, but as of May 2020 is fully functional for the Java Edition.

¹⁰³ <https://3dprint.nih.gov/discover/3DPX-012095>; <https://3dprint.nih.gov/discover/3DPX-012092>

¹⁰⁴ <https://sketchfab.com/3d-models/dna-molecule-11209f199eb54e58a3701060196488d1>
(has been deleted since download)

Protein Transport Network

Once proteins have been made, they require transporting to locations in the cell where they can function within biochemical networks. This zone is designed as a small game where users must sort proteins and transport them to the correct 'station'. Proteins must be located from three large chests, matched to a coloured coded chart, and then placed into a minecart to be sent to the end station. This is meant to represent the transporting of newly synthesised proteins along microtubule networks within the cell by motor proteins (Goldstein & Yang, 2000) through reimagining this process as a rail network.

The central hub station of this zone was an adaptation of an existing automatic rail station (Frilioth, 2019). Users can select one of nine different destinations which alters the rail network upon the detection of a minecart on the track. The original structure was modified to include teleportation to destinations and the next zone, and a reference wall for the protein sorting carried out in this zone. The departure track was also modified to allow both standard minecarts and those with hoppers, which allow for the transportation of blocks.

Each destination consists of the same structure which sorts incoming items based on the block type (Avomance, 2018). These were all configured to accept only three specific block types. Incorrect blocks are placed into a junk chest for collection by the user. These structures are all linked to the central hub via colour-coded rail networks and teleporters.

Users select a station with a button situated underneath a coloured block representation. This lights up a redstone lamp and primes the track network to change if a minecart is detected. Users may then summon a minecart to either transport themselves or blocks situated in three chests within the station. Once moving, the minecart will be sent on the correct rail.

Cell Communication Laboratory

The entry point to this zone is an adaptation of a door by Oliver Brotherhood (MumboJumbo, 2020). Three dendritic branches, of varying spine densities, were taken from the NIH 3D Print Exchange and are examples derived from Kasthuri et al., 2015¹⁰⁵. A nicotinic acetylcholine receptor was also taken from the NIH 3D Print Exchange, authored by catherineelee¹⁰⁶ and derived from Unwin, 2005.

The remaining structures were generated by myself, depicting the synapse and receptors in various interactive forms within various rooms of the laboratory.

¹⁰⁵ <https://3dprint.nih.gov/users/bobby-kasthuri>

¹⁰⁶ <https://3dprint.nih.gov/discover/3dpx-002494>

Users enter into a reception area before having the option to explore four different experiments. One allows users to place gold blocks onto a 2D grid representation of the synaptic zone. This creates generic receptors in a larger version of this grid. The user may then press buttons that show neurotransmitter release and ion flow through open receptors based on their previous input.

The second room shows the effect of activity on receptor populations in the synaptic zone. Neurotransmitter release is represented as falling red sand, which hits a platform below with two types of pressure pads. Gold pads represent receptors, and stone pads slots where receptors may be present. Users may press buttons to increase or slow the rate of red sand summoning. The gold pressure pads add a value of one to a scoreboard objective which is being constantly depleted by a redstone clock activating a command block. With no user input, the score is maintained within a small range, representing basal activity. If activity increases and the score reaches a defined threshold, all stone pads are replaced with gold pads to represent potentiation of the synapse. If activity is reduced sufficiently the gold pads reset to basal values. If the score falls even further, the number of gold pads reduces. This system can be reset at any time using a labelled button.

The third room has a representation of an ion channel receptor in a membrane. Users can open and close the channel by the addition of neurotransmitter blocks via a button press. Another button causes the summoning of ions represented by blue wool. These travel through open channels or are destroyed upon contact with a closed channel. A teleporter in this room takes the user to the top of the acetylcholine receptor structure. Users can move to be centred directly above the central pore and drop through the receptor to simulate being an ion travelling through the channel. Teleporters at the bottom return the user to the top or back to the main laboratory.

The Library

A hidden room in the Cell Communication Laboratory takes users to a library derived from the Endonian Library by ToadieOdie (2014)¹⁰⁷. This contains several books which have facts about each zone, and teleporters to get to each zone. Furthermore, there are full credits for all structures used and links to further educational resources. This structure is large enough to contain additional books and zones as the activity grows over time. There are also several command blocks present which allow users to make changes to the gamerule, adjusting the game mode, mob spawning and weather.

¹⁰⁷ <https://www.planetminecraft.com/project/endonian-library-17-3-types-of-schematics/>

Activity Run #1: Leith Labs

The initial version of the Neuron Safari was tested at a public engagement event in January 2020 at Leith Labs. This is a community project focused on bridging the gap between the public and researchers focused on the city of Leith community. The activity was advertised online and was a drop-in session over the course of one day. This additionally acted as a test run for the activity for its use in the Edinburgh Science Festival, to help confirm the average time spent by each user, the areas which might be too difficult to understand, and how easy to use the software was for both volunteers and users.

There were five zones present for this trial (no Cell Communication Laboratory or Library), and was run like a guided tour, using volunteer biomedical science students and staff to help users navigate their way through the activity. Users could spend as long as they wished in each area and were encouraged to take control of navigation and ask questions. To help with the task, users had access to physical 3D brain models and soft toys of neurons and brains. Volunteers were equipped with factsheets and guidance as to how to run the activity. These were provided beforehand and coupled with a one-to-one walkthrough of the activity.

On the day, visitors to Ocean Terminal (an indoor shopping centre) in Leith could come between 11 am and 4 pm to try out the activity. We were full from the outset and throughout the whole day, with users tending to spend closer to 30 minutes with the activity. Four stations were available, with at least one volunteer at each station. An additional two to three volunteers were present to organise waiting groups, and talk to them about neuroscience using 3D models and posters.

Many visitors chose to control and navigate through the world themselves, although many were unfamiliar with the controls. Whilst Minecraft is incredibly popular, it is also used across a diverse range of platforms. From discussions with users, parents and other volunteers, it appears that phone and tablet use are prominent in younger users.

Whilst every attempt has been made to “child-proof” the activity (invisible barriers, limited navigational choices and clear instructions), some volunteers had difficulties if issues did arise. This stemmed more from a lack of Minecraft experience than from activity-breaking errors, but were noted to rectify in future iterations. This did in part lead to the longer than expected time taken per child on the activity.

Despite these hiccups, users were engaged throughout their time with the activity. The initial introduction zones, being mostly passive visual experiences, gave users a chance to familiarise themselves with the equipment and controls, and gave the

guides a chance to gauge the user's current knowledge level and begin to introduce some neural and cellular concepts. This meant that by the time the more interactive zones were reached, guides and users were comfortable in discussing what was on screen and users were not overwhelmed with too much information.

One anticipated downside was that the content level would be too hard to grasp and result in a loss of engagement. We did not find this to be the case. Whilst the DNA Factory was hard to explain to younger users, the concepts had been translated into Minecraft sufficiently to allow at least a surface level understanding of the material. The Protein Transport Network, with its obvious train network analogy, was much easier to grasp, with no reported difficulties in understanding. In the first three zones, facts were kept simple and centred around numbers and comparisons. This aided in quicker comprehension and gave users some easy to remember material to take home.

Users were given a flyer which acted as an advert for future events using the Safari and provided some zone details for interested users. It also had a QR code which linked to a download of the world for Java Edition users. This version was defaulted to Creative mode and gave users the freedom to use the activity for their own ends. Before acquisition, a short feedback form was presented asking for a quick evaluation of the activity. There were four questions, which were to evaluate the enjoyment of the activity, how each zone was perceived, and how well people understood the information contained within the activity.

We received nine responses to the form. 75 % of respondents gave the Safari a rating of ten (out of a maximum of ten), with all scores being nine or ten. On a scale of one (bad) to five (excellent), all zones scored at least 3 (**Figure 80**). The brain zone was seen as particularly interesting, followed by the neuron zone and skull zone. Almost half the respondents gave a N/A answer for the DNA Factory and Protein Transport Network. This may have been that these regions were not reached due to time constraints. Where scored however most responses were of five (excellent).

44.4 % of respondents felt the difficulty level of the Safari material was average (three on a scale of one to five, with one representing "too difficult" and five "too easy"), with results skewed more towards easy than difficult. An open feedback question elicited five responses, four emphasizing enjoyment of the activity and one with a suggestion of how to add to the activity in future.




Figure 80 Leith Labs
Feedback: strong engagement and content comprehension

Eight families responded to the request for feedback on the flyer given to every attendee (top). All responses indicate a strong enjoyment for each Zone, although comprehension of some of the content is found to be difficult.

At the event, participants were asked to write three words on their experience, and their ages. A word cloud (generated on wordclouds.co.uk) shows the comments received (right). Ages ranged from 5 – 10 years.



Activity Run #2: Edinburgh Science Festival

Changes to the Safari

As a result of the feedback from volunteers, parents and users, both major and minor changes were made to the Safari. To reduce the duration of the activity, the speed of movement during the first three zones was increased by adding more redstone-powered rails. We had found that users did not need to spend a long time in these zones to receive all the information necessary to understand and for any questions they might have had. One of the major delays previously was summoned but unused minecarts, which are impossible to remove in Adventure Mode and required knowledge of the command line interface to fix. These would clutter the tracks and prevent progress through the Safari. To remove this issue, all summon commands for minecarts were preceded by a deletion of any minecart in a specified radius.

Two extra zones, the Cell Communication Laboratory and Library, were added. As we had found the time limit to make fitting in such zones potentially difficult, I made these sections not immediately noticeable and hid the teleport from the Protein Transport Zone behind a banner block. This allowed for the use of the zone if time permitted but could be ignored otherwise and not be expected by the user during their time with their guide. These extra zones were importantly still available for any independent use of the activity and did not require any major structural changes to the Protein Transport Network to access.

SARS-CoV-2 Changes

The increasing prevalence of SARS-CoV-2 in the UK led to the cancellation of the physical instantiation of the Festival. Rather than cancelling the whole event however, activity leaders were asked to reconfigure if possible their activities to work online. For the Neuron Safari, this meant adjusting the activity in two major ways. The first was to add in-game text commentary which would stand-in for the tour guides that would normally be present. This was added in the form of minecart-activated command blocks situated throughout the minecart rail portions of the Safari. Material is brief and conveys the major aspects of each zone. The text automatically disappears after a short period, so longer explanations are instead stored in books in the world, placed on lecterns. Books are much more prevalent in the final interactive zones, where user movement and positioning are not precisely controlled.

An exception is in the Cell Communication zone, where information is conveyed by information points that require a button press. This is more in keeping with the design of the zone as well as acting as a trial for use of this format in future zones. Information is displayed and colour coded into short sections. This allows the user

to return quickly to a section to read if the text disappears from the chat window before they finish.

The second major change necessary was the conversion of the world from Java to Bedrock Edition. This is the more popular version amongst our target audience, and so the reach of the Safari would be stymied if only the older, Java Edition format was available. Whilst this had been successfully tested with the Leith Labs version, that iteration had much fewer command blocks and redstone structures. The two versions are related but have key differences in coding structure and block properties. For example, concatenation of multiple commands in one command block is extremely limited in Bedrock¹⁰⁸, leading to an increase in the number of command blocks used and adjustments to the code to accommodate this. The coding language is also slightly different, with block IDs for different colours of blocks sharing a base name and then a numerical code for the variant. Redstone is also slightly different in Bedrock, with components moving slower, timings being different and sometimes stochastic in nature. This makes many Java-built redstone constructs difficult or impossible to replicate.

Conversion of the Safari world required the use of the MCCToolchest tool, developed by *cynodontA*¹⁰⁹. This allows for quick conversion of a world file from Java to Bedrock and vice versa. Not all block types are shared however, and orientations of certain blocks are reset as a result. Conversion thus required a considerable amount of time to convert all the code, redstone and blocks to match as closely as possible to the Java version. I was able to reproduce almost all the elements in the Bedrock version, with the exception of the redstone entrance to the Cell Communication Laboratory and falling blocks in the same zone. Workarounds were implemented which do not impact adversely on user experience.

#ElementsofEdSciFest

The Edinburgh Science Festival website was reconfigured to link to online versions of events and was rebranded as *Elements of EdSciFest*. Much of the advertising and promotion of this virtual festival was carried out on Twitter, using specialized hashtags such as #EdSciFest. For Edinburgh University specific content, there were the additional hashtags #EdUniEdSciFest and #EdUniSciFest.

¹⁰⁸ To provide an example: *execute if x1 y1 z1 red_wool run setblock x2 y2 z2 green_wool replace* would require splitting into two blocks. One impulse command block with *testforblock x1 y1 z1 wool 14* and a chained conditional block with *setblock x2 y2 z2 wool 13*.

¹⁰⁹ <http://www.mcctoolchest.com/>

This virtual festival ran officially from the 4th to the 20th April 2020, although most of the content is still available online and linked to from the Festival website.

The new version of the Safari was uploaded to its own webpage (www.neuronsafari.com). This provides a description of the activity and its creators. There is a dedicated download page which has both Java and Bedrock versions of the activity, with links to videos and webpages with instructions on how to install the activity on a prospective user's device. There is also a link to an optional survey to provide feedback on the activity after playing.

The advantage of an external host for the activity was the ability to record anonymous statistical data from incoming visitors. I used StatCounter¹¹⁰, which records many different metrics that can be used to assess visitor engagement and demographics¹¹¹.

Neuron Safari reaches a global audience As of May 2020, the website was visited by 523 unique visitors from 26 different countries. The majority of visitors came from the UK (77.7 %), which was anticipated given the nature of the promotion of the activity and the language it uses. Of these UK visitors, many were from Scotland and specifically the Edinburgh area, suggesting much of the interest came from people aware of the Edinburgh Science Festival.

Most visitors are on mobile devices 47.94 % of all visitors to the website used a mobile device, increasing to 53.77 % if tablets are included. For Bedrock Edition users, this theoretically can allow for a short time between initial engagement and activity use, as users only need to download the correct file and open it in the app version of Minecraft. Difficulties with file type, inexperience installing files like this, and whether visitors have the mobile/tablet version of the software will all be factors however that might impact on this.

Download conversion rate is about average IP logs showed 16 unique IP address downloads of the Safari, all of which were of the Java edition. This would lead to a conversion rate (visits that result in a "purchase") of 3.06 %, which is within the average website conversion rate (McDowell, Wilson, & Kile, 2016). Rates for educational websites specifically were not found. Our score is likely to be much higher however, as downloads did not always register upon testing on mobile devices, especially with the Bedrock *.mcworld* file. Given the prevalence of both this device type and software version, the conversion rate may be greater.

¹¹⁰ <https://statcounter.com/>

¹¹¹ The website and StatCounter utility have been maintained, and at the time of publication of this thesis are still running.

Optional feedback survey was not visited The website has a “What Do You Think?” page which links to a feedback form made using Jisc Online Surveys¹¹². This was an adaptation of a form which would have been used to garner feedback at the end of the Festival mini-workshops. The survey is split into two sections. The first is intended to be filled in by the user of the Safari. Given our target audience, the questions were written to promote answering and to avoid miscomprehension. There is evidence that shorter questions with simple syntax and straightforward language are necessary to elicit accurate responses (Bell, 2007). They also state that questions to children should avoid being negatively constructed (i.e., an answer of no indicates a positive response).

There are eight questions, which cover the enjoyment of the activity, the ease of use, and if the respondent learnt anything. All but two questions kept to a maximum of three possible responses, which has been regarded as more likely to lead to a response (De Leeuw, 2011, but see Borgers & Hox, 2001).

The second section is designed for parents or guardians to complete and is aimed at reinforcing some of the answers from children and assessing the level of engagement with the material. There are five questions, with two open-ended ones which ask for general positive and negative feedback on the activity. The first question in this section is designed as a multi-part Likert scale, exploring the levels of enjoyment, engagement, ease of use and content.

A final section requests some basic demographic information. To avoid the collection of identifiable data, ages were placed into broad categories and gender and country requests were open questions. All questions were optional, with an introduction to the survey pointing this out and encouraging no response if any doubt of the question’s meaning.

The survey was kept online until August 2020. No responses were forthcoming to this survey. Analysis of the website visits indicates only 2.48 % of visitors accessing the “What Do You Think” page, with no click-throughs to the survey. Because of data protection considerations, there was no opportunity to request an email address from visitors prior to downloading the Safari, and there is only one reference to this in the Safari itself towards the end.

This situation highlights the difficulty in acquiring feedback when an activity is unsupervised and available for anonymous download. For future online events that make use of the Safari, the request for feedback will feature more prominently in the activity. I will also request that organisers email participants with a

¹¹² <https://www.onlinesurveys.ac.uk>

reminder if contact information has been collected and this request meets with ethical approval.¹¹³

¹¹³ The current questionnaire, as it is both optional and acquires no identifiable personal information, does not require ethical approval at the University level. The short time frame of the switch to an online event meant ethical approval was not possible to acquire beforehand and is why email addresses were not collected when downloading the Safari.

Activity Run #3: Glasgow & Midlothian Science Festivals and Organic Visits

For the remainder of 2020, there were no major content changes made to the Safari. Instead, the focus was to provide it as a ready-made activity to the autumn and winter Scottish science festivals, which were all online. My contribution to the Glasgow Science Festival (subtitled 'Science on the Sofa') ran from the 9th September to the 9th November, and was placed in the 'Hands On(line)' theme. Visitors to the site were provided with brief description and a link to the main Safari website. I added to the website a special Glasgow Science Festival-themed page, which was designed to provide an experience of the Safari to those visitors who did not have a copy of Minecraft, and provide an extended preview of the activity.

The page displays several facts about each zone, most of which are not present within the Safari itself, but were intended for use during in-person use. These were further supported by the listing of many of the weblinks found in the Safari's Library Zone, which enhance many of the concepts in written and interactive formats. The main element of the page is an embedded video which is a guided tour of the Safari, commented by myself. This was intended to enhance the understanding of some of the more complicated or less obvious concepts and components of the Safari, as well as provide a chance for those lacking a copy of the software to experience some of the Safari and learn some facts about Neuroscience and basic cell biology. The video does not show all parts of the Safari, nor does it interact fully with some of the activities within the Zones. This was in part due to requiring the video to be at around the 20 minute mark (something suggested by Glasgow Science Festival), and also to retain some of the curiosity and mystery of the activity for those that could then go on and play through the Safari (**Figure 81**).

To tailor the video somewhat to Glasgow Science Festival, I also included in-game a version of the Glasgow Science Festival logo for 2020. This was imported using SpriteCraft (v 1.17), which converts image pixels into Minecraft blocks, exporting a *.schematic* file which can be imported into Minecraft using WorldEdit. This was positioned at the start of the Safari.



Figure 81 Neuron Safari @ Glasgow Science Festival.

To accommodate Festival visitors without Minecraft, a video was produced to take a viewer through six of the seven Zones, providing some extra biological information and explanation of the activities and structures (left). A reproduction of the Festival 2020 logo in Minecraft (right).

Video: <https://www.youtube.com/embed/U7QDvOcHwgc>

The Safari also had a modification of the DNA Factory Zone. For the previous iterations, the translation machinery was either not present or ran into errors when running sequences that featured premature stop codons. After consultation with Student B, by the time of Glasgow Science Festival a new version was available which functioned correctly, although only in Java Edition. To accommodate this new version, I reconfigured the zone to be more visually linear, reducing the number of steps required to reach each zone. I also added the example amino acid sequences for of Human CER1, DACH1, and WFIKKN1 as mentioned earlier at this point, which helps visually link the smaller user-made sequences and the amino acid library to real-world examples.

This new version was also later used in the Midlothian Science Festival, and has up until the time of writing still being accessed by new visitors. By December 2020, the visitor count stands at 1016 unique visitors (almost doubling since May 2020), 1458 sessions and 2396 page views. Many of these visitors came from the Festivals and from a link placed on the Federation of European Neuroscience Societies (FENS) website, but there were also noticeable spikes from Twitter posts and a poster presentation at the FENS 2020 conference, which was one of the 10% most viewed posters at the conference. Visitors have also shared the world via messenger apps, Google Classroom and social media. The world file has been downloaded at least 50 times, although as mentioned earlier this is the minimum value for the Java Edition, and fails to account for any Bedrock downloads, or those for Education Edition (which was added to the website in November 2020).

As **Figure 82** shows, visitors came from 52 different countries, with many of these still coming from the UK. There is however a broadening of international interest, as this represents double the number of countries since May, and a reduction of the UK share of traffic from 77.1 % to 67.6 %. The data for what sites these visitors came from is unreliable, with some websites noted as referrals but most being identified as direct links to the website. Visitors are split almost evenly between mobile and desktop/laptop.

Future Iterations

Given the relative novelty of neuroscience concepts to a general populace, and the increasingly global reach of the Safari, there are several options being explored to how best to take the activity further in the future. Considerably lacking from the current uses by people has been feedback, which has limited the ability to know which elements to change or adapt to help with user engagement and understanding. The gradual return to in-person outreach and public engagement will facilitate this, but for now it must be assumed that the Safari has successfully met with my own and user expectations, is easy to download and install, and requires no major changes.

The now obvious international potential of the Safari also suggests that a useful method of expanding its reach is to offer some translations of the embedded information. This is outside the immediate scope of the Safari, but given the relative sparsity of embedded text, not insurmountable. What will be key is being able to source accurate translations, ideally that can be quickly called upon for any future updates. Doing so will allow me to target a wider audience, and reach communities which have comparatively fewer science communication opportunities than English speakers, here in the UK as well as abroad.

The immediate focus is on how to expand the experience further. The modular aspect of the Safari makes construction and integration of new Zones easy. Conceptualisation and realisation are the only early hurdles to overcome. Currently in development is the addition of a model of the blood-brain barrier, a crucial mechanism that isolates the brain from the rest of the body whilst still allowing for the active transfer of essential molecules (Daneman & Prat, 2015). To summarise, capillaries supplying blood to the brain surface are surrounded by pericytes and astrocytic 'feet', which help to tightly regulate the transfer of






Figure 82 Neuron Safari Website Statistics.

Visitors to the website are mainly from Europe, although there is a sizable number from the US, Australia, and Canada. Many of the visitors come from direct links, suggesting an organic sharing of the website between interested visitors. Visitors are split almost evenly between desktop and mobile, with a slight preference for Windows OS over Apple OS in the former category.

substances from blood to brain and ensure that passive transfer of unwanted substances is close to zero. This is achieved through a series of different membrane transfer mechanisms (Abbott, Rönnbäck, & Hansson, 2006).

Users will have to help sort molecules from a central blood room and ensure they are correctly navigated into the brain by the correct transfer mechanisms. In this module they will learn about different types of molecules, and how their properties change how they can navigate through the barrier (if at all). We therefore gently introduce specificity, compartmentalisation, and enhance the emphasis on molecular transport introduced in the Protein Transport Zone.

In addition to this new zone, I will be introducing some replications of animal behavioural study methods, which I began exploring as a way of generating outreach and research opportunities using Minecraft. I will explain the process of this development in the next section.

Reproduction of Animal Behavioural Experiments in Minecraft

World File: Animal Behaviour Experiments

Comparative Behavioural Experiments

Translating animal behavioural work to humans requires a careful consideration of the cognitive differences between species. Any animal work to some extent assumes that cognitive phenomena in the model animal, and any molecular and cellular components linked to this through experiment, are applicable to human cognitive phenomena. Animal studies give us access to the mechanisms that would be highly unethical to investigate in a human study, allowing for direct inference and genetic manipulation to assess how a phenomenon is realised.

To ensure these experiments are useful comparisons, human researchers can create proxies that aim to look at the same cognitive phenomena and assess if behavioural strategies differ between species. In some cases, indirect measures are possible (i.e., eye tracking or brain imaging), which can help to test the correlation with animal work and provide suggestions for further study.

Approximations of behavioural experiments however run the risk of appearing homologous but being merely analogous. A gold standard for a behavioural study would be one in which humans and animals can be participants in exactly the same experiment. Often there are logistical issues which prevent this from being possible. If they are possible however, they would arguably allow for much more robust conclusions to be made and strengthen the truth of the assumption that work in the animal model translates well into humans.

In this section I will show the adaptation of two well-established animal memory experiments into Minecraft. Each already has an existing human correlate, which I will compare with this possible approach, that would make use of standard video game as well as Virtual Reality (VR) technology.

Comparative studies to date

Our understanding of memory and spatial navigation¹¹⁴ has advanced rapidly over the last century largely thanks to well-designed animal behavioural experiments (Eichenbaum, 2017; Grieves & Jeffery, 2017). These experimental setups are now limited by the strict controlling of the environment that has aided in uncovering so much. It is difficult to tell if the strategies implemented by an

¹¹⁴ Spatial navigation often being used as a stand in for human declarative, or episodic, memory.

animal in an experiment match to what would happen for the same premise in a naturalistic setting (Genzel, 2021). By extension, it is difficult to claim the findings map well to human behaviour (M. Zhao, 2018).

One way of checking for this would be to run the same experiment in humans as in an animal model. This is often quite a difficult thing to implement, however. For example, the current Morris water maze for rats has a diameter of 200 cm¹¹⁵ (e.g., Genzel et al., 2017). To recreate the same experience for human participants would require a pool of 11 m diameter, with perfectly smooth interior walls to remove proximal cues. Such a replication would be highly impractical and cost prohibitive.

We tend to make use of analogues instead that capture the same idea as the animal study, but in a different manner. The Morris Water Maze for example has been achieved virtually, where subjects are tasked to find a hidden treasure chest on a featureless island (Goodrich-Hunsaker, Livingstone, Skelton, & Hopkins, 2009; Moffat & Resnick, 2002; Samanta et al., 2020; Schoenfeld, Schiffelholz, Beyer, Leplow, & Foreman, 2017). Distal cues are situated in the water surrounding the island. This commonly used setup however has major discrepancies from the animal version, so much so that the comparative aspect should be questioned.

Firstly, the human participants are told the aim of the task. This means a behavioural strategy is primed and instigated from the beginning, and is further reinforced through in-world text that tells the individual that there will be subsequent trials. Second, the human task involves no atypical movement or mild stressor. Water is not the default environment for a rodent, nor usually is the inability to get out of it. Third, the cues used are familiar and recognisable ones, with the rodent ones being abstract shapes. In some studies with both animal and human comparisons, the maze water is also transparent, with a transparent hidden platform. From the images used in Schoenfeld et al. (2017) however, the platform apparatus is clearly distinguishable from the rest of the pool, adding a proximal cue to the experiment.

The final issue is that these studies are not clear with the use of the term ‘virtual reality’¹¹⁶. Virtual reality is an immersive experience where the user has no other field of view other than the virtual environment. The user is thus experiencing the virtual world in place of the real one, with at least one sensory modality fully

¹¹⁵ An increase of 70 cm, for those interested (Morris, 1981)

¹¹⁶ Some researchers use “Desktop” or “non-immersive” VR for the method of interacting with the “virtual environment” by sitting in front of a computer. Other terminologies include “flat”, “pancake”, or “screen”. I use ‘flat’ for the remainder of this Chapter.

interacting with the virtual one. Movement through the environment can be done through actual movement of the subject also, which has been improved through the development of wireless VR technology.

The methodology of these papers make no mention of the use of VR headsets, and the original methodology can be traced back to Moffat & Resnick (2002), which describes participants sitting in front of a screen and controlling their movement with a joystick in a virtual *environment*. They are thus making use of a virtual world, but their experience is detached and fails to match that of the animal experiment. The documented similarities in results, and more recent studies on the utility of VR over ‘flat’ virtual experiences suggest that this analogous method might be sufficient (P. Srivastava, Rimzhim, Vijay, Singh, & Chandra, 2019; Stevens, Kincaid, & Sottolare, 2015). The lack of physical movement in these studies however poses a potential confound, with work allowing for motion in high-fidelity VR showing significant improvements in object location memory (Murcia-López & Steed, 2016).

The Human Event Arena

As I argued earlier, a truly comparative study of a behavioural task should have the human and animal subjects performing the exact same task. VR allows us to come very close to this ideal without a prohibitive cost. To exemplify this, I first constructed a VR version of the event arena. In its rodent form, the arena is a 1.6 x 1.6 m clear-walled open space with two 3D intramaze cues towards the centre (Tse et al., 2007, 2011). The floor has a 7 x 7 grid of interchangeable white plastic squares. These can be switched out for panels with a central circular hole in which sandwells can be placed. Sandwells are sand-filled containers which can be dug in for a potential reward.

Most of the current users of this setup use a delayed-matching-to-place paradigm to explore memory. An animal is placed within a “startbox”: an opaque container situated on the outer middle of each of the four sides of the arena. Within this startbox is a water source and a food pellet of a specific flavour. The arena contains 49 possible sandwell locations¹¹⁷ of which normally between one and six contain sandwells. One of these sandwells will contain the same flavoured food pellets as in the startbox, the rest will be empty. The aim for the animal is to find this correct well, take the food, and return to the startbox. They can then run out again and find the same sandwell to collect more food.

¹¹⁷ Realistically only 43 once the positions directly in front of the startbox and where the two intra-maze cues are located are discounted.

Variations of this method have been utilised to provide evidence of long term consolidation of episodic memory into pre-frontal cortices (Bethus, Tse, & Morris, 2010; Tse et al., 2007, 2011), and as a test of everyday memory showing that spaced learning protocols encourage memory persistence (Nonaka et al., 2017). The latter complements a comparative study that compares with human memory studies and trials of cognitive enhancers. The former has been complemented with human imaging studies (van Kesteren, Fernández, Norris, & Hermans, 2010; van Kesteren, Rijpkema, Ruiter, Morris, & Fernández, 2014; van Kesteren, Ruiter, Fernandez, & Henson, 2012).

Virtual setup

In all formats, the event arena requires an animal to navigate through predominantly visual cues to find the food reward. Given the food is merely being used as a motivator, it allows us to fully reconstruct the methodology into VR with the only change being how the cue is given in the startbox, and how the “reward” is discovered.

I reconstructed the event arena in Minecraft to closely resemble one of the animal experiment rooms used by Nonaka et al. (2017) (**Figure 83**). The walls are white with black geometric cues. The entrance to the arena zone is filled with black banners to represent the curtain that is present in animal experiments. The arena is modelled exactly on the animal version, with four startboxes made of black glass blocks, containing a water source and a container for the cue block. A narrow entrance gives access to the arena, which has 49 possible sandwell positions.

Sandwells are represented by a 2 x 2 x 3 structure that lies flush with the arena floor. On top of this are removable trapdoor lids which are removed upon user interaction. In the non-cued sandwells, the underlying blocks are all red sand. For the cued sandwell, one block will be identical to the cue block.

A user must “mine” the cue block, run into the arena and search the sandwells to find the corresponding reward block. Once found, this needs to be mined and the user return to the startbox. From here they may be asked to go out and find the same sandwell again, or finish the session.

During the time the user is in the startbox, the arena may be reset. To do this, the preperation zone of the animal rooms is recreated with additional tools to allow for rapid arena resetting. An experimenter has access to a 7 x 7 representation of the potential sandwell locations. Placement of a gold block on one of these locations opens up the corresponding arena location (or replaces a sandwell if present with empty space). The experimenter can then identify which sandwell type (cued or non-cued) is to be placed by pressing a button in front of the correct

sandwell from a list of six (five different cue blocks, one uncued). The map can be reset at anytime, returning the arena floor to having no sandwells.

Upon completion of placement, the experimenter can press a button corresponding to the startbox which the user is in, which will place a cue block (if the beginning of a trial) and reset two scoreboard objectives which measure time taken to find the reward block and the number of errors (measured as one error for every four trapdoors opened). These act as rough in-game versions of the major two behavioural output metrics for the task in animal settings.

The detailed metrics that can normally be derived from the animal data, such as tracking paths and more fine-grained time data have been shown to be possible by using Spigot¹¹⁸ and a dedicated Minecraft server. The Spigot software allows for the insertion of tailored Java code to track many different aspects of user interaction with the game, and has been used to evaluate the effects of Minecraft environments on hippocampal memory (Clemenson, Henningfield, & Stark, 2019). The use of VR headsets offers the added benefit of tracking eye movements in humans too, allowing for further analysis for how visual cues are being used.

In this VR environment, users must walk around to find the reward block, and can only make use of visual cues in game to do so. There is no need to inform the user of the goal of the experiment beforehand. The same experiment can therefore be run and offer ways of exploring if both animal and human participants use the same strategies when solving the arena, and if they have similar memory consolidation using probe tests (no reward pellet/block in the cued sandwell). This is a much more direct comparison to the animal studies and requires only a decently sized room and a VR setup.

Controls will entail the use of the event arena as a flat virtual experience and a computer-based but non-Minecraft delayed matching to place task. These will again require no explanation as to the task or how to complete it to be given, beyond an explanation of what buttons or keys to press to register choices and to move, and a neutral space to practice these motions (akin to habituation in rodents). Ideally, only a small number of trials should be allowed on a given day, and then repeated each day for about a week. We expect that in all three cases the determining of how the task should be performed will be acquired, but that the additional complexity of the event arena may take longer for users to reduce the number of errors made. We also predict that the VR experience will enable quicker and more accurate development of a schema that maps block colour to specific

¹¹⁸ www.spigotmc.org

locations, as users have no distracting or potentially changing cues from the real world during the task.




Figure 83 Virtual Event Arena.

The event arena room is built to represent one used in Nonaka et al. (2017), split into an experimental set up area (top left) and experimental area. The arena itself (top right, bottom left) consists of four identical startboxes surrounding a central open space which has 43 possible spaces for sandwells to be placed (seen bottom right).




Figure 84 Virtual Water Maze.

Room design is similar to that in the virtual event arena (Figure 6). A setup zone (top left) allows a user to set platform position and player start location within the water maze. Platforms may be set to be hidden until the player is within an adjustable block radius, or after a set period. Users can only use extramaze cues to navigate.

Human Water Maze

The Morris Water Maze is one of the most widely used behavioural tests used to investigate memory in rodents. As previously mentioned, a human version is necessarily different given the difficulties of scaling the method. The virtual experiences used to date however are much closer to the event arena method than the water maze, removing the crucial component of discomfort and sub-optimal navigation that rodents experience when being placed into a giant pool of water¹¹⁹. Strategies and behaviours under pressure or in unfamiliar conditions may diverge from those made whilst comfortable or in familiar settings (Brunyé, Gardony, Holmes, & Taylor, 2018; Gouirand & Matuszewich, 2005; Zimmermann, Stauffacher, Langhans, & Würbel, 2001), and so current comparisons in this area may suffer from unwanted confounds.

Virtual environments, or the interactions with them by users, can be made to generate similar feelings of discomfort, frustration, even stress (Boulton, Hourizi, Jefferies, & Guy, 2017; Gilleade & Dix, 2004). This is something many games utilise as a way of immersing a player, fostering a sense of realism, challenge and subsequently a sense of achievement when the obstacle is overcome (Canossa, Drachen, & Sørensen, 2011; Jennett et al., 2008; Nylund & Landfors, 2015; Sweetser & Wyeth, 2005). In these instances players have to spend time generating a strategy to overcome the obstacle whilst under the stressful conditions, test that, and refine. This refinement can occur during the task, or outside the game after a period of time or sleep.

This mild peril is relatively absent from current virtual instantiations. I set out to attempt to develop a prototype in Minecraft which achieves this and has a similar environmental setup to the Morris Water Maze. The water maze is represented by a hollow cylinder of rendered in black concrete blocks and with the topmost side removed. This cylinder is filled with water. Surrounding this water maze is a room with both 2D and 3D cues present on the walls. One wall is made of banners, which are mostly opaque but allow for a player to pass through (akin to the curtain used as one wall in the Morris Water Maze setups I have used in the past).

In the Java version of the game, water cannot be dyed to be opaque. This creates a problem for the generation of a hidden platform, as they are potentially viewable from a distance by the user. The hidden platform also runs into problems surrounding block choice. Barrier blocks are the obvious choice for a hidden platform, as they are invisible to the player but can be stood on. Currently

¹¹⁹ Interestingly, some of the earliest forms of the virtual water maze were more faithful replications of the task (Astur, Ortiz, & Sutherland, 1998; Sandstrom, Kaufman, & A. Huettel, 1998).

however these blocks cannot be waterlogged in Java versions of Minecraft. Waterlogging allows for a water block to exist in the same space as another block – otherwise a block will be deemed outside water, and so any surrounding water will treat it as an obstacle and move around it. This renders the platform visible, as to the player the block's location appears to be filled by air instead of water.

To circumvent these issues, the platform is designed to be generated only when the player is within two blocks of the centre of the platform's designated location. This prevents navigation by unwanted detection of the platform but still allows for the participant to locate the platform when in the correct area. This mechanism can, with a minor edit, be also used for experiments where the platform is absent for the first minute of the rodent in the maze. In this instance the platform is tied to a timer instead of the current position of the participant.

The platform can be placed in any part of the pool, at any depth. Normally the pool is split into four quadrants, and the platform is placed in a particular location once in each quadrant. Participants can be started again from any location, in the pool, but the standard procedure of placement according to relative north/south/east/west with regards to the pool is the one coded in the examples shown here.

The participant when placed in the pool must search to find the platform. There is no other way out of the pool, and any swimming under the water results in a player losing oxygen and potentially dying. A player must therefore continually keep their heads above the water line whilst navigating the pool to find the platform. This requires multitasking - keeping afloat and being aware of their surroundings. The difficulty of the task is therefore restored to a level similar to that for the rodent, which must also keep its head above water and locate a platform through visual navigation (**Figure 84**).

Data Extraction

The main metrics of the water maze in rodent studies are: time spent in quadrant (correct vs incorrect), distance travelled, time taken to reach platform. The final two are quite easy to accommodate within the current simulation, using pre-existing scoreboard objectives that allow for detection of distance travelled in water and a timer which operates only if the participant is in the pool and not on a platform. The first metric is somewhat harder to accommodate, as it requires careful mapping of each quadrant, but is possible. Each quadrant has a command block associated, which checks for the presence of a player in a defined radius

from the centre of the quadrant¹²⁰. If a player is located, a scoreboard timer can update. When the player moves to another quadrant, the timer is paused. Currently there is a necessary overlap point at the edge of each quadrant, therefore avoiding any dead space. This can be modified if required.

With this setup the main metrics of a water maze trial are thusly observable and recordable. This still requires some manual extraction of the data from the software, but this is also how researchers use the existing rodent software, extracting time values and recording them in a data sheet to take out of the arena for data entry at a later time point.

A missing and potentially valuable metric is the detection of the path of a participant in the water maze, not just the distance travelled. An overhead camera is often used in the Morris Water Maze that can track a rodents' movement and save this for the generation of heat maps or density diagrams (Genzel et al., 2017; Samanta et al., 2020). In client-side versions of this virtual water maze, there is not the ability to extract this data accurately. It would be possible to record an overhead view of the pool, which may be possibly used to extract the same detail. A better option however would be to place the world onto a server and utilise the data package developed by Clemenson (Clemenson et al., 2019), which extracts the co-ordinates of the player at defined time steps and is a potentially easier way of generating accurate maps.

VR, as mentioned in relation to the event arena, also provides an opportunity to make use of eye-tracking, which allows us to understand how cues are being used in both immersive and flat VR. This gives us the ability to assess cue quality and compare to rodent studies and assess how these may differ.

Implementation and Development

Unfortunately, due to the SARS-CoV-2 pandemic, conducting tests of these environments has not been possible. The intention was to run these behavioural tests in both 'flat' and immersive VR, both to assess task performance as related to rodent studies, but also to the pre-existing virtual water maze. A working prototype of both methodologies can be seen in **Figure 85**.

I also intend in the future to trial these experiences as part of public engagement programs to both assess their utility for comparative behavioural analysis, and the possibility of using this as a method of data collection of less commonly accessible demographics. Testing within public settings offers multiple benefits to both

¹²⁰ Radii here are representative of a square from the centre, not a circle, thus avoiding any gaps.

participants and researchers. Both the event arena and water maze demonstrate animal studies which can help challenge some of the public perceptions of this form of experimentation, as they can be conducted without any insult or chronic stress on the animal. This form of comparative analysis is also arguably a way of showcasing how digital approaches can help reduce the level of animal experimentation, and show how we may someday be able to replace it.

Testing in outreach and public engagement environments also provides a much greater number of participants of varying degrees of familiarity with computers and virtual reality, allowing for a wealth of data collection on how the methodology should be run and adapted to suit different groups. This reduces the level of calibration or the risk of failed experiments if these methods were used in a formal investigation framework. It also importantly provides a chance to assess how well this method may be developed into a citizen science project. The world save can be downloaded by an international audience and be used at home to generate data about spatial navigation in virtual environments. Whilst not quite as rich as the data from a formal in-person set up, this method opens up a much wider range of participants that can provide a much more accurate picture of spatial navigation across different demographics.

Human studies are still quite heavily dependent on university member demographics and that of the surrounding residential areas. This creates a so-called WEIRD demographic (Western, Educated, Industrial, Rich, and Democratic) (Henrich, Heine, & Norenzayan, 2010) which distorts experimental outcomes and generates broad assumptions about human behaviour which may not hold up as strongly in a larger more diverse cohort, or even be unique to the population tested. Whilst primarily used as an argument for social behaviour, reasoning and other so-called 'higher' level cognitive functions, the argument I believe also holds for processes such as spatial navigation. It is well reported now that different people employ different navigational strategies, and that this can differ by gender, culture and age (C.-H. Chen, Chang, & Chang, 2009; Goeke et al., 2015; Wolbers & Hegarty, 2010). Moreover, early work in the 20th century showed that assumptions about cognition of seemingly simple or basal concepts are greatly shaped by educational exposure and environment (Luria, 1976).

These behavioural experiments are complex in design and interpretation, but relatively easy to conduct as a participant. By automating much of the experimental setup and data collection, I hope that these may be introduced as a way of collecting a rich and diverse dataset of spatial navigation techniques, strategies and ability that can help to better understand the translational elements

of current animal research, and in so doing also define how much has diverged and in what manner.



Figure 85 Development Version of Water Maze and Event Arena.

The above video shows example runs of both behavioural experiments detailed in this chapter.

Video: <https://www.youtube.com/embed/UWEI0C6JoXQ?feature=oembed>

Planned Pilot Studies

The immediate uses of these experimental simulations are within public engagement and citizen science. The event arena and Morris Water Maze are being incorporated as an extra zone into the Neuron Safari world, to be released in March 2021 as part of Brain Awareness Week

(<https://www.brainawareness.org/>). They will be enhanced with some in-game information about how these experiments work, and how and why animals are used currently in these investigations.

My next step is to additionally develop virtual environments in Minecraft that mimic rodent object exploration tasks to complement work by Dr Antonios Asiminas (Asiminas, 2016; Asiminas et al., 2019). These will be accompanied by educational videos explaining the science behind the tasks. This project has two distinct objectives:

- To educate the public on non-invasive animal behavioural testing.
- Collect human behavioural data as part of a unique citizen science project that will look at object memory across age points.

Given the popularity of Minecraft amongst neurodivergent individuals (Ringland, Boyd, Faucett, Cullen, & Hayes, 2017; Ringland, Wolf, Faucett, Dombrowski, & Hayes, 2016; Zolyomi & Schmalz, 2017), this project will also allow us to engage this demographic in citizen science. Collecting behavioural data from neurodivergent individuals, especially those who struggle with anxiety, is challenging. Our project allows us to engage collaborative data collection with these individuals from the comfort of their homes, using a platform they love.

MinecraftEdAR: Integrating Minecraft into the Classroom

A key aim of MinecraftEd is to harness both learner and educator creativity, and to rethink how we teach certain concepts in biology. Whilst the previous examples cover a broad range of possible uses for this idea, they all still require Minecraft itself. As mentioned earlier, our use of Minecraft stems from its sandbox capacity, shallow learning curve and multi-platform support. What this picture lacks is how to extend beyond this, a broader interpretation of this idea of model construction and creativity.

Since February 2020 I have been slowly developing a collaborative project with the University of Edinburgh's School of Engineering. This project combines two PTAS projects, MinecraftEd and Educational Apps in Augmented (and Mixed) Reality (EdAR; <https://edar.online/>)¹²¹. The latter was established in 2017 and has been developing pilot Augmented Reality (AR) applications within a multi-disciplinary collaboration. Importantly, this idea features a large amount of student involvement, with undergraduate Engineering students adapting CAD expertise into AR environments. The EdAR project has successfully developed applications for use in Engineering and other disciplines, at the time of writing all within the University of Edinburgh.

With MinecraftEd, we have established the potential of Minecraft as a modelling software that requires little to no training that can be used by staff and students. The in-built versatility lends itself to much quicker development of AR models, with the pipelines for importing of many different 3D file formats established in this thesis. Whilst these lack the usual accuracy of CAD-built models, they often do not need to be so exact for pedagogical purposes.

This collaboration is as of time of writing ongoing, and has several aims:

- To test how well Minecraft models translate into AR versions
- If these models lend themselves to a physical version through 3D printing
- If successful, to what extent can this pipeline be student-led, and
- Do they enhance learning in an established teaching environment (e.g a tutorial)?

¹²¹ EdAR is run by Prof Andrew Sherlock at the University of Ediburgh.



Figure 86 Development Version of AR DNA Factory App.

The mechanisms, structures and aesthetics of the Minecraft DNA Factory have been adapted into Unity assets and animated in a similar manner in AR. Users may interact with any element to bring up extra information, structures, and animations. Video courtesy of Simon Parker.

Video: https://www.youtube.com/embed/_XAoCoUTIYk?feature=oembed

I will present here the initial findings of this collaboration, which has been the product of myself and two undergraduate colleagues from both the Deanery of Biomedical Sciences (Simon Parker) and the School of Engineering (Justin Savage).

The Unity Engine

The AR applications were written in the Unity Engine (Unity Technologies), which is built on the C# programming language. All team members made use of AR assets developed by Dr Andrew Sherlock and used the Unity Student Plan.

AR Project: DNA Factory

The DNA Factory has proven to be a versatile and valuable tool for testing out the possibilities afforded by digital education. The model's shape and size lends itself well to this project, particularly for the construction of a physical 3D version to be used to anchor the AR model in the user's local environment. A minor modification was necessary for 3D printing, as the tRNA model design was too fragile and caused errors.

This aside, the focus of this project was to assess how quickly a pre-existing model can be translated into AR. Additionally, we sought to assess if such a translation could be done by students working collaboratively, with each student an expert in either the biomedical or AR modelling. This sharing of skills was predicted to speed up the construction and refinement of the AR models using Minecraft, and to allow for the drafting of a documented pipeline that faculty and students could use for their own purposes.

To successfully implement the model in an AR environment, the model elements were translated into a series of animations written in C#. The user is presented with two prefabricated DNA strands which are separated by a gap measuring 21 nucleotide bases. The user can input their own sequence of nucleotides using on-screen buttons up to the 21-base limit. Mistakes are removable through a separate button. The user may then animate the transcription of their chosen sequence through another button press. If a user fails to provide a sequence divisible by three at the time of pressing this button, the app randomly generates enough bases to compensate.

The transcription animation begins by the fusing of the prefabricated and user-inputted DNA. The DNA splits in two lengthways in the y-axis. A script is called that creates mRNA base pairings complementary to the upper DNA strand, as well as a prefabricated poly-A tail. Once complete, the DNA objects move away and are deleted. The user may then press an additional button to translate the mRNA.

Translation summons the 40S ribosomal subunit and an initiator tRNA to a fixed point on the prefabricated portion of the mRNA. This codes for Methionine, which is the initial amino acid for all proteins at translation. The rest of the ribosomal machine then enters the scene and begins to read the mRNA, with the mRNA moving from right to left in this process. Each codon summons the correct tRNA holding the correct amino acid, which is added to the existing amino acids to slowly form a polypeptide. If none of the user inputted sequence included a STOP codon, this continues into the second prefabricated mRNA section which contains the poly-A tail. Here a STOP codon is present, which runs a termination animation sequence that removes all but the generated polypeptide sequence. The whole process can be restarted through another button at this point.

During every step of the application, the user may touch any on-screen model component. This will stop any ongoing animation and bring up a reactive UI with further information about the pressed object. Within this UI are extra buttons which summon 3D chemical structures related to the object (**Figure 86**). The user may spend an unlimited time within this “Touch and Explain” interface and

remove it by pressing the close button (“X”) in the upper right of the UI. This reactivates any paused animation sequences and allows the user to select another model component.

A non-AR version of the app was also made to enable those without compatible phones and tablets to still learn using the app materials. All structures are instead held in specific locations with a fixed viewpoint. The user can still make use of the “Touch and Explain” and animation elements, but in a fully virtual environment.

Implementation

The app has undergone considerable refinement and adjustment under EdAR but has retained the block structure and simplicity of the original model. The 3D model element is not currently needed for the app to function, although this was largely due to a shift in methodology of AR tagging and may be returned to any future educational deployment of the application. Some of the changes have come about through consultation with the School of Biological Sciences at the University of Edinburgh, who have expressed an interest in the deployment of this and other AR applications in teaching over the next few years. The application is designed to combine with a hands-on workshop that gets students to construct models of DNA, RNA and ribosomes, which may end up being tagged to act as components of the app also. True structural models are thus paired with both real and virtual simplifications, as well as interactive demonstrations of the key biological mechanisms occurring. The aim is to determine if this blending of real and virtual, simple and veridical will improve or accelerate student learning of protein synthesis, and then offer these applications to other interested parties.

At the time of writing, the restrictions surrounding SARS-CoV-2 and the teaching calendar has meant implementation is yet to be trialled. Interested readers may wish to enquire with EdAR or the author for updates on this venture.

Project Millport: Virtual Ecological Fieldwork

World File: Project Millport

At the time of writing, an additional use-case for Minecraft has arisen out of the requirement for virtual fieldwork enforced by the current SARS-CoV-2 pandemic. Ecological fieldwork often involves the active sampling of an environment, using a variety of tools to demarcate sample areas and observe species of plants and animals within related environmental contexts. From this sampling, a researcher may extract several different types of data that can be used to explore variations in populations that may be because of the environmental context a species finds itself in.

Student fieldwork is often conducted at sites such as beaches and forests, with students learning about how to sample an environment just as much as skills in data management and relation of ecological principles to a real environment. Students not only perceive field work to be more informative than simply receiving data to analyse, but also perform better on applying the knowledge acquired in related assessment (Boyle et al., 2007; Rickinson et al., 2004; Scott et al., 2012).

At Edinburgh, Field Zoology 3 (a core Zoology course) students normally undertake five days of fieldwork in April based at FSC Millport on the island of Great Cumbrae. Students are tasked with laying down transects and quadrats on the island's beaches to sample the natural fauna and flora. Typical questions posed to students include assessing the relationship between limpet size and distance from the low tide mark, or to measure the relation between dog whelk populations, barnacle coverage, and height above sea level.

Importantly, the species investigated, and the questions asked, are not of any novel biological or ecological significance. They are designed to assess students in their ability to sample and record data, and to learn how transects, quadrats, and other ecological tools work in practice. I was invited to work with Prof Tom Little on rapidly developing a replacement for this fieldwork. This gave me an opportunity to develop a Minecraft environment which would capture these learning outcomes and still provide an immersive experience that required active exploration of the environment and crucially provide users with the freedom to choose how they sampled the environment. This was important because it allows for a much greater variation in the data sets collected and reported as outcomes of this fieldwork, whilst allowing for constraints by educators to ensure they can tell when data has been collected appropriately or not.

Building the Island

This project first entailed the attempts to recreate the actual island used normally. Whilst Minecraft procedurally generates environments, including islands, the scale, topography and biomes¹²² are often not that resemblant to the real world. Creating an island is also possible in game, but even with the tools such as WorldEdit¹²³, this would take a considerably long time to achieve on the scale required. A tool called Minecraft Earth Tiles¹²⁴ allows for the conversion of parts of the world into Minecraft world files. The world is segmented into 'tiles' of 1° latitude and longitude. The tiles can be scaled horizontally between 1:10 up to 1:200, and vertically from 1:5 and 1:200. Selected tiles are then processed into working Minecraft world files through a batch process that involves extraction of OpenStreetMap¹²⁵ geographical data and bathymetry data via QGIS¹²⁶, heightmap data from the Advanced Land Observing Satellite Digital Surface Model¹²⁷, and filtering of this data into geographical features. These are then 'painted' into a single file using ImageMagik¹²⁸ and converted into a working Minecraft world file using WorldPainter¹²⁹.

I wanted to achieve a sense of geographical realism through using Puerto Rico as a template. Puerto Rico requires 6 EarthTiles to fully render, and exhibits a range of elevations, biomes, and locations suitable for adapting into beaches (**Figure 87**). The island is much more rectangular however, which reduced the number of diagonal beaches, which are difficult to work with if elements such as tides or transects were needed to be generated. The island was generated at a 1:25 scale horizontally and 1:5 scale vertically. The vertical scale does lead to some clipping of the higher elevations (Minecraft worlds unmodded are limited to a height of 255 blocks), but as these were not part of the primary brief of the project this was not felt to be an issue.

The key elements of the project were the need for users to travel to distinct beaches and collect information both visually (by identifying species and counting population sizes), and through some form of measurement (equivalent to

¹²² Biomes are distinct geographical parts of a Minecraft world that have distinct flora and fauna compositions, as well as properties that affect weather patterns, block colouration, and sky colour.

¹²³ <https://worldedit.enginehub.org/en/latest/>

¹²⁴ <https://earthtiles.motfe.net/>

¹²⁵ <https://www.openstreetmap.org/>

¹²⁶ <https://qgis.org/en/site/index.html>

¹²⁷ <https://www.eorc.jaxa.jp/ALOS/en/aw3d30/index.htm>

¹²⁸ <https://imagemagick.org/index.php>

¹²⁹ <https://www.worldpainter.net/>

dimensional or weight data). It was important therefore to determine an appropriate size for the beaches, how to show the gradual increase in height above sea level, and how to insert custom data which a user could extract. It was also important to determine which in-game entities would be suitable for representing beach fauna.

I identified at this point that the key requirements for virtual fieldwork- to be engaging, accessible and informative – matched closely to Malone's (1980b) challenging/fantasy/curiosity identification of what makes a game self-motivating. First and most important is to be engaging – this is a task which students should want to do and explore and perform without requiring educator prompting. This ties to one of the learning outcomes, which is that students should begin to independently develop their own experimental design. On Cumbrae, this is surrounding the times they sample, the beaches they sample from, the species chosen etc. I wanted to ensure a similar level of choice was on offer.

This requires striking a careful balance within the game world. We want the user to have the freedom to choose their own path to data collection, but there is also the need for the data to be comparable between users and to real-world data collection. Drawing from real-world instantiations, the easiest way of doing this is to frame the task with a handbook or other instructional material. These scaffold the users thinking and expectations of the task, and allow for influence to be exerted on what are good avenues of research to be explored.

As I discovered in the development of Neuron Safari, this works best virtually when you design from the beginning for the world to be more than the intended aims. In essence, you approach the task as being about the creation of a game, not just a representation of concepts or tasks. This takes Minecraft away from being a mere vehicle for educational content, and instead about generating a whole identity which has this content as a part of it.

Unlike Neuron Safari, there were fewer constraints surrounding both audience and time. The primary users are 19 – 21 years old, have already an understanding of the concepts exemplified, and are allowed over a week to interact with the world. This allowed me to explore much more deeply how to make the world greater than just being about collecting data on various beaches. The brief could have been met with a simple square island with four beaches, all with the required diversity of flora and fauna, and simply a small bit of flat land separating each one. This however would promote rapid completion of the task, and offer no opportunities for students to be creative. Fieldwork for this class normally takes place over the course of a week, and I wanted to keep this element of time in the virtual version.

The final island is approximately 6,700 blocks in length and 2,300 in width. Travelling in a straight line with no obstacles would take 10 minutes in real-time, roughly equivalent to the in-game daytime portion of a day. Straight line travel however is impossible with the current geography, making any exploration of the island considerably longer¹³⁰.

Once generated, beach generation and environmental morphing were conducted in parallel. Several sites were identified for beaches, classified according to the surrounding geography, biomes, and the presence of any offshore islands. From this list four sites were selected that were sufficiently diverse in the first two of these categories, and offered a diversity in maximum size, depth, and potential layout. No sufficiently sized offshore landmasses visible from a listed shoreline were identified during this survey, so the decision was made to move some existing smaller islands to the necessary beach locations at a later date.

All four sites were then modified using WorldPainter. First, each location was changed to a beach biome, which would remove the possibility of entities or other biome-associated elements of the game being generated on beaches. Secondly, each location's topmost 3 blocks were replaced with sand blocks so as to better represent a beach environment. The boundaries of each location were then modified to be in most cases at a higher elevation to the beach zone (which ranged from 61 to 71 in the y-axis) at the sides and at the supralittorial zone. The beach zone itself was then adjusted to show a graded increase in elevation from sea level up to either 70 or 71 on the y-axis.

¹³⁰ Great Cumbrae, incidentally, is about an hour to walk from one end to the other, making the Minecraft island somewhat comparable in this area.



Figure 87 Final Island Design for Project Millport.

The geographical data for Puerto Rico, imported from OSM via EarthTiles was further adjusted in WorldPainter to move offshore islands, build beaches, and modify some of the island's shape and interior region (top). Beaches were initially prototyped to determine suitable sizes, rockpool distributions etc. (bottom).

This beach shaping was not done with any deliberate decisions on the size of the beach or the width of each level. This led to a much more natural-looking and diverse set of beaches that can be identified by these properties alone. Once complete, each new beach required a considerable effort to remove pre-existing flora such as trees, vines, and grasses, which are not removable in WorldPainter. Furthermore, modification of a Minecraft file in this manner is prone to several bugs which require manual adjustment in-game. Most notably this is the presence of water at random heights or locations, and artifact blocks (usually rendered as mossy cobblestone) in locations where elevations are considerably lower than their original levels.

These manual adjustments were also necessary during the reshaping of the island. I did not wish to keep the island as Puerto Rico, as this may have generated confusion about the island content, and would have potentially restricted its future use and development. As such, large parts of the island had changes to elevation profiles, and some of the island was submerged. A central mountainous zone was flattened to make space for the proposed virtual Millport, which would be roughly equidistant from each beach. Millport is the main population centre on the real island of Great Cumbrae, acting as the hub for student projects in normal times. I felt it important to maintain this hub space, and the name, to keep the experience thematically similar for students. Nearby, a considerable area was turned into a mountain lake, with a single river that runs out to sea and generates a delta at its mouth.

Beaches were populated by a mixture of manual and automated methods. Turtle data regarding height above sea level, width, shell length and sex was provided by Prof Tom Little for each beach. Each level above sea level for the beaches was then mapped for its width and length. Each turtle was then assigned a co-ordinate position, randomly selected from the ranges calculated using Excel's `RANDBETWEEN` function. These were then written into a summon command and saved as *.mcfunction* files and packaged into a datapack, which is saved into the world's *datapacks* folder. This allows for the */function* command to be called in-game. Each *.mcfunction* file contains a list of commands in the form of:

```
summon turtle -272675 63 -73546 {Tags:["width 19.31071",  
"shell length 19.83665","sex female"], Team:Shore,  
NoAI:1,Invulnerable:1,PersistenceRequired:1}
```

Calling the function therefore summons all the turtles to the desired coordinates at once. As beach levels were not of equal widths or lengths, some locations resulted in buried or floating turtles. These were rectified by a mixture of beach modification and repositioning of the affected entities.

The other necessary entities (tropical fish, sea pickles, coral, and silverfish) were placed manually in game according again to data provided by Prof Tom Little.




Figure 88 Beach Maps.

Millport Station has maps laying out the topology and rockpool distribution of each beach. Users can make decisions on where to visit and how to segment a beach prior to visiting, allowing for more careful planning of data sampling excursions.




Figure 89 Fintray Beach.

Example views of the rockpools and turtle distributions on one of the island's beaches. This is a sheltered beach, with the offshore island shown bottom right.

Task Components

The students are required during the fieldwork to spend up to five days working on projects within Dumbrae (the island's current working title). The field handbook gives a map of the island, as well as the locations of Millport and each named beach. There are six proposed projects listed for students to undertake:

Turtle size

Students must design an investigation on how sea turtle size varies either intra- or inter-beach (or both). Sea turtles are distributed at each beach differently, with differences in total population, density, and numbers at each co-ordinate point above sea level. They additionally have extractable data of their width, shell length and sex.

Turtle Cannibalism

Baby sea turtles are hypothesised to avoid adult turtles as they are in danger of being cannibalised. Students must determine if this hypothesis is validated through sampling.

Silverfish-Sea Pickle relationships

The primary food source for turtles on Dumbrae is reported to be sea pickles. Silverfish, which are semi-aquatic and invasive also eat sea pickles. Students must determine both the relationship between silverfish populations and sea pickle abundance, but also whether this has any impact on the turtle numbers found on any given beach.

Rockpool diversity

Each rockpool has up to six different species of fish and four different species of coral, in varying distributions. Furthermore, rockpools possess different depths, sizes, and heights above sea level. Students must choose how they wish to investigate diversity in rockpools through the assessment of one or more of these variables.

Rockpool predation

The dottyback fish, present in many rockpools, is predated upon by the native bird population. Students must test the hypothesis that the dottyback is safer when it lives near many other fish of different species.

Clownfish-coral relationships

Clownfish are reported to eat organisms only found on fire coral. Students must determine if clownfish abundance is linked to the number of fire coral present in a rockpool.

EcoLab

In addition to fieldwork, students may also optionally wish to conduct ‘experiments’ at an Eco Lab, situated in the main Millport zone. The lab consists of two rooms (*Pool* and *Arena*), each with four experimental stations. As **Figure 90** shows, each tank or arena is identical, and has two areas at each end where test items (both plant and sessile animals) can be added. These items can be used to investigate the behavioural response of three different species found on the island. The user summons the required animal from an observation platform, which allows them to track manually the dwell time in each pool/arena half, the path taken, and therefore any discernible preference of the species for one of the test items.

Students may wish to tie these experiments into their field projects, as most of the species in the lab are taken from the beach population. However, a student is also free to investigate using the lab independently of these projects, or not use the lab at all. This provides an extra layer of choice as well as complexity to the student-defined elements of the task, giving greater ownership and independence that helps to better resemble real-world data collection.

Tutorial Island

Given the size and scope of the experience, and that about 60 students would be taking part in the course, we discussed a way of introducing how the game worked, as well as how the specifics of the fieldwork would work also. I decided the best way would be to generate a separate area which could act as an in-game tutorial zone, where students would need to meet a series of requirements before being allowed to reach the main island (**Figure 91**)

The island features six different zones, designed to test the student in the game controls, object identification, tool use and the biological concept of zonation. All students start on this island, which is a considerable distance from the main island. The first area (or Zone 0) is an optional space for those unfamiliar with the game and the controls. There are instructions in the form of books and signs in this zone, which if accidentally deleted are immediately regenerated. This allows a user to get used to which controls are for use and which for destroy, which they can also manually change in the settings. The area also introduces buttons, chests and minecarts, which are all used prominently in the main task.

Once satisfied, a student can pull a lever to enter the tutorial proper. The first zone, *Rockpools*, is designed to get students to think about blocks as units of measurement, and to get used to using buttons to activate in-game elements.




Figure 90 EcoLab.

The two rooms offer an optional manner of investigating the island's animal species by observing preference to up to three different native plant and sessile animal species. Users have the option of dry land and water species, each with different options. An observation platform is situated above each pool to allow a user to make observations.




Figure 91 Tutorial Island Zone.

The main 'rooms' of the tutorial area are show above, consisting of (clockwise, from top left): Rockpools, Turtles, Beach, and Museum. Apart from Beach, all others are required to be completed in order to reach the next room and therefore the main island.

A student has to count the longest length of three rockpools, divide by two, and then input this in a control panel for each rockpool. When correct, a light is lit and in the case of three correct answers a door opens to the next zone. The door is not openable in any other manner even using commands. A reset button is available in case a user goes beyond the correct rockpool length and needs to start again.

The second zone, *Turtles*, introduces the extraction of data from in-game entities using a tool, and the calculation of averages from sampled data. The room resembles a turtle hatchery, with several baby turtles and some adult turtles present. Some of these turtles are tagged with randomised data on their width. The user must pick up a “Turtle Measurer” and nametags from a chest, using the former to extract data and the latter to prevent double counting. Once complete, the student must work out the average width of the turtles. The door to the next zone is surrounded by a series of signs with different feasible averages (taken from a random distribution around the correct value). Right clicking an incorrect answer teleports the user back to Zone 0, with only the correct answer opening the door. There are 48 signs, making brute-forcing this room much longer than conducting the required task.

Museum, the next zone, is a way of ensuring students can correctly identify the in-game animals they will be investigating. Students must observe a different entity at each station, and identify it based on its colouration or markings using their handbook. Each station has two species names, with levers above them. The student must select the correct species name for each station in order to open the door to the next zone. Upon pressing a lever, some information is provided about the species, regardless of if correct or not. To check if they are correct, a row of red blocks is next to the door. These change colour to green if the correct response is chosen for each station. This is not signposted or explained – the user must work this out for themselves. The handbook does explain that the images will help with this zone, and therefore helps to identify those students who have not read the handbook closely.

The next two zones *Methods* and *Beach* are optional in-game, having no restrictions to being bypassed. *Methods* is an information-rich zone, with signposted buttons providing some background to the methods of sampling being used in the fieldwork, as well as other contextual information of interest to a keen learner. *Beach* is the final zone, and features an idealised graded beach. There are six countable species on the beach, represented by bees, baby turtles, slimes, sea pickles, warped nylium and podzol. Students are asked in their handbook and in an in-game book to count the number of each species at each level of the beach, and graph the results. All six have been distributed to show zonation between

them, pairing bees, slimes and baby turtles with a preference for one of the remaining three species. Students are asked to submit their results of their whole-beach sampling as proof of completing the tutorial zone.

At the end of the beach is a small platform which teleports the player to Millport station, and thus allows them to commence their main fieldwork projects. Return to this area is possible at any time through a teleportation device in Millport station.

A “lived-in” environment: non-essential additions

An aim I had was to make the island feel alive, or at least “lived-in”. A world of this scale would feel empty and lack engagement if it was strictly just the beaches and the automatically generated environments of the island. I therefore set out to create a structure to the world that would blend elements of the real Millport and Great Cumbrae. This was reflected somewhat in the directives I received, such as the naming of the beaches, island, and central zone after those found on Great Cumbrae. Running with this idea, I began to create Millport. This is the central “hub” of both the island and the users progress through the fieldwork. From this zone, users can access all four beaches directly through two means of transport.

These are housed in a large building designed to resemble a Victorian train station (**Figure 92**). The station is likely to be the most seen aspect by the majority of users, as it is necessary to return to when trying to access distant beaches, the lab, or any of the additional elements of Millport itself. Given the high footfall, it was necessary to make the space large, easy to navigate, and to contain many of the necessary materials for students to carry out their projects.

The station contains at the entrance a ‘tool collection point’. This provides the user with the Turtle Measurer, nametags, and materials for laying down quadrats. This point can be accessed again and again if a user loses any of the items. Nearby are four maps of each beach, which allows a user to see the topology, rockpool distribution, and size of each beach before visiting. This may help with students planning their trips and with their write ups of the project itself.

The path to these and each mode of transportation is both signposted and directed through a gold block pathway inserted into the floor. This is paired with large open spaces which help to clarify the important areas within the station currently, as well as allowing for easier expansion of station elements later. In the centre of the station is a board that directs the user to each of the train networks for each beach, which are also clearly colour coded. Behind this board is a teleportation station, which allows for instantaneous travel to any beach and the Tutorial Zone. A user can choose to use one or both these methods as they wish or require.

Train journeys have been designed to enhance the element of immersion and to emphasise the size of the island. The journeys range from two through to almost nine minutes in length and take different routes and approaches to reach their destinations (**Figure 93**). Most outbound journeys are overground, and show to the user some of the topology of the rest of the island. Return journeys are underground in a subway system, which is designed to be more creative in its design. Outbound journeys also have a series of announcements and audio-visual elements that are both project specific and are meant to represent slightly a train journey whilst listening to the radio. These elements are again optional but are available to those who wish for a more immersive experience.

Millport, in addition to the station and lab, also has several other buildings and places which are much more closely tied to the “lived-in” aim. These include gardens, a house, library, pub, and café (**Figure 94**). The latter two are modelled on actual locations in the real Millport, and have been replicated as close as possible given available photographs, with the scale somewhat modified. The library is a location where staff can add information in the form of in-game books, as well as hyperlinks to extra material on the course VLE and other websites. This allows a user to remain in the game and still access additional material that can run at the side of the game in windowed mode.

The island also possesses several custom achievements that cover a range of likely and unlikely actions a user may do in their time on the island. Achievements have become a staple of video games in recent years, and have been shown to amplify engagement and time spent in game environments, as well as promote a sense of reward in the user when triggered.

Users receive awards that scale with the number of slabs laid down for quadrat and transect formation, and for how many times they use the Turtle Measurer and nametags. There are also rewards for completing the tutorial zone, exploration, visiting all the buildings in Millport, taking every possible train journey and for finding hidden structures and entities scattered around the island.



Figure 92 Millport Station. The station is the central hub of most user's fieldwork projects, connecting to each beach by train or teleportation. Users can also pick up their equipment, view maps, and access the rest of Millport from this location. All aspects of the station are well signposted throughout.




Figure 93 Island Train Network.

Each beach takes a different length of time to reach by train, and is linked locally to at least one other beach to reduce the level of redundant travelling necessary. Trains operate both overground and as part of a subway network. A copy of this map is also present within the station.




Figure 94 Millport Aesthetical Additions.

The Millport zone is enhanced through the construction of several non-essential elements that provide areas for multiplayer socialisation, further information, and exploration. Some of the structures are references to establishments found in the real Millport (top left, top right).

Much of the latter elements are entirely hidden, and not expected to be found by users. They are at most only weakly visible in a task-focussed run-through of the island, but are thematically linked and carefully constructed. There is no at present utility for these elements, except in adding a sense of fantasy and depth to the environment which may be capitalised on as the island grows in content. Their addition also acts as a coarse proxy of a user's exploration and curiosity, as they require additional effort to locate for no tangible reward.

Once more, I have strived to develop a world which pairs educational content with a user's natural curiosity and inquisitiveness that allows for a wide range of uses to be easily applied in future. This also has a long-term advantage. The world is large enough to add extra elements without needing to remove existing material, and is diverse enough in geography to minimise landscape reshaping. The world building elements can be used to link together new additions such that they blend into the existing world, and further targets a user's curiosity and inquiry into *why* these elements are there, what they mean, and in so doing promote a greater time spent in the environment and potentially on the tasks set by the educator. These act too to inject the 'fantasy' element, which is arguably most useful for any variations of the task which might be deemed suitable for outreach and public engagement settings.

A Scaffolded Learning Environment

Project Millport is gentle on the constraints it places on the user. In Neuron Safari, it was necessary to offer only the illusion of choice, and to restrict any chance of users moving outside of the zones. For this activity however we have much more time allocated for a user to spend interacting, and an audience who are meant to be learning to learn independently.

Whilst there are outcomes that we want the students to meet, we do not want to too closely manage the route any single student takes to reach it. This is always a careful balancing act in any setting. Provide just the island with no material, and users must spend a considerable amount of time learning how to navigate, where they should be going, and what tools they have available to them. If we provide too much guidance however, all students will conduct the same investigations, find the same data, resulting in a lower chance of engaging and understanding both the concepts and skills being taught.

The level of scaffolding here contrasts with that used in Neuron Safari. There the level of information and guidance in-game is tailored to new, young learners. This level of guidance can be altered through in-person discussions between the user and educator, facilitating more complex knowledge acquisition and relating concepts to other previously known information. In the Safari, the external (tour guides) and internal (in-game written material) scaffolding structures can be adapted, but the external scaffolds are more responsive and immediate.

In Project Millport, there will be a chance to have external scaffolding (handbooks, synchronous video calls and asynchronous videos) and internal scaffolds (books, signs, images). The users here however are much more able to learn independently, and already have some proficiency in the concepts and topics that make up the learning outcomes of the activity. Scaffolding here becomes less about relaying of new information, but rather encouraging users to *apply* their knowledge, and to facilitate them developing their own interpretations of the in-game data. This is what they would do in the field: they can choose the beaches they sample from, where they sample, and what they sample. They can also choose the times, the number of methods etc., all of which are advised upon by educators, but are not chosen for them.

Collaborative Learning Opportunities

As we saw in Chapter 5, a learner may proactively reflect on their weaknesses and use collaboration to improve their skills through asking peers to show how to do something, or to explain a concept. In the real-world fieldwork, students can work in pairs or small groups to assist with planning investigations and sampling.

In our digital setting, collaborative work is possible, and may be arranged to work in a similar manner. There are however some challenges to digital learning environments for this collaborative approach. One is the relative skill levels of a group. Whilst we have made efforts to ensure users all have the same base level of ability using the software, those with poor computer literacy may choose to disengage from the activity if group members are seen to be comfortable with the software. The other is communication channels – without an established, accessible online space to communicate both synchronously and asynchronously, group work becomes difficult, slow, and takes away from the task being performed. For the former, we will have to see the outcomes of using this activity. For the latter, we are exploring setting up Discord (<https://discord.com/>) servers, which can be set up with minimal effort and offer distinct group-only communication spaces accessible on any internet-enabled device. These servers allow for live-streaming and video calling, offering a single space for the synchronous and asynchronous group activities within a well-designed accessible user interface.

This may be enhanced through setting up a multiplayer server for the Minecraft island file itself. This would allow for software issues to be resolved more easily, ease the computational load on older laptops, and allow users to work together in the same island space simultaneously. Users can then be placed into in-game teams and communicate privately using the in-game chat function. Both this and the Discord servers can be implemented independently of one another, but both allow for students to work together to plan their investigations and collect the necessary data.

Students will also be able to communicate with educators in similar ways. Demonstrators, who would usually assist with field work questions, can contact individuals or groups via their Discord servers, or in-game if a multiplayer version is implemented. Being able to provide immediate live (or at least rapid) feedback may prove crucial in the success of this activity, which lacks the physical displacement and sense of authority of real-world field work.

Implementation

The island will be distributed to all 60 students on the Field Zoology 3 undergraduate course at the University of Edinburgh. This will occur in the first week of April 2021, with a preview of the activity in March. The implementation structure, as mentioned, is still being finalised, but will likely involve a mix of server and client (downloaded) versions. Evaluation of the experience is anticipated to be incorporated into the debrief at the end of the field work week, to be submitted along with the field reports the students write. The evaluation will

take the form of a short user experience quiz, and will especially target the perceived value of the experience by students, and their perceptions generally of the use of Minecraft in their education.

Other Universities have also expressed interest in similar versions of the island for their own Zoology courses – I intend to evaluate those in a similar manner, focussing on the perception of its use and the amount of Zoology knowledge and skills it offers.

General Reflections

Minecraft has shown itself to be a remarkably powerful and creative vehicle for student-centred learning, pedagogical development, and public engagement. It pairs the relative creative freedom of coding languages with a practical but highly visual style, and a simple user interface. As an educational tool, it has had some consideration, but never seriously at the tertiary level. I hope that with much of the above examples I have shown that this, and other game-based learning approaches, should be seen as a viable way of engaging with tertiary learners.

In creating pedagogical tools, we are often constrained by budgets, feasibility, and time. With MinecraftED, often the challenge has less been about if something is possible, but more in how to go about realising it. This has been aided immeasurably by the wide player base who have got there first, constructing applications, modifications and programs that allowed for a much greater range of possibilities to be realised in much of what I produced.

Time is admittedly still a major challenge that plays into how game-based learning may progress in the future. Each of the above worlds was hundreds of hours in the development and deployment. For an educator without a background in the software, the thought of having to learn how to use the software may impact their likelihood of attempting. Familiar fears of software become outdated or incompatible are also no doubt worries that some readers may have.

I'd argue that this can be reduced through the design of training and templates to help accelerate educator skills in use of Minecraft. Microsoft already have such a program for Minecraft Education Edition (<https://education.microsoft.com/en-us/learningPath/3eede2ae>). Furthermore, the pilot study here shows that 24 hours of time is more than enough for even inexperienced students to know enough to generate informative constructs and hone their skills in learning and communication in the process.

The medium here is Minecraft, but the discoveries can apply to any other current or future software which has similar freedoms in user access and manipulation to how it operates. This freedom is key, as at the heart of why all of the projects listed here have been successful is because they have not been constrained too much creatively. Given the tools, and a goal, learners can make extraordinary things in the process of learning.

In this chapter I have only been able to allude to the more social or in-game collaborative facets to this idea. Project Millport is being tested as a multiplayer environment in the upcoming year (2021) and should offer some insights into how well this method of virtual fieldwork can accommodate the social aspects of

fieldwork. The behavioural experiments will offer another angle into how users might use these experiments, how their use is communicated, and whether this route is viable for citizen science approaches. I am particularly keen to explore how we may engage with the design of future experiments with neurodivergent communities, using in-game non-verbal communication outlets.

I began this project to test the idea that Minecraft might be a creative way of learning about complex biomedical concepts and models. What I have discovered is an untapped enthusiasm and curiosity from both educators and learners in the adaptation of existing digital environments. It has spawned multidisciplinary collaborations, attracted an international audience, and continues to grow. Most encouraging of all, it appears to be inspiring those who encounter it to think about how we can harness student's digital knowledge, experience, and talents in the design of pedagogy, not just as a method of deploying it.

CHAPTER 8

Epilogue

This thesis has laid out two ways in which computational modelling may be used more effectively in academic and pedagogic settings. These are not a radical rethinking, nor a major reworking of current practice, but rather a tailoring and synthesis of disparate pre-existing elements. I have tried to both highlight and persuade the reader that our current utilisation of computational modelling approaches in biomedical science is too restricted to a few well-worn methods.

The double-edged sword of computational work is always that you can try something in a different way. There will always be an alternative implementation or solution to any problem, one that may be demonstrably better than current practice. The key is always a mixture of ease-of-use and dissemination. I believe that both directions outlined in this thesis meet at least partially these two points. Minecraft meets both very highly, and it is my intention to pursue this approach avidly over the coming years.

Modelling and visualising biomedical concepts can only be improved by the increase in the number and range of people able to take part in doing so. From an academic standpoint, it increases the knowledge of a given concept and across different levels of explanatory difficulty. From a pedagogical stance, it allows for a physically constructivist take on biomedical concepts that would normally be impractical to anyone below postgraduate level. It also allows for much greater agency in how a learner expresses their learning or understanding, enhancing written explanations, and requiring a much deeper level of understanding to construct than is minimally required for a traditional essay.

This of course is not without some caveats, which we have discussed already. But the fear of obsolescence is dampened by the fact that the *idea* being offered here transcends any singular platform. Constructing rather than merely describing or explaining a concept is for this writer an overlooked concept at the tertiary level of teaching, further constrained by budgets and time. Computational approaches free us from these constraints and enable us to give back to students the autonomy of their learning without removing existing practice or structures. It harnesses the strengths of a learner's imagination and requires them to adapt their understanding to be understandable by others.

Rule-based modelling has already begun to impact both computational modelling generally and specifically the development of scientific theory. Crucially it opens up the ability to model single molecules in great detail. In this alone it will allow

for investigators to develop highly detailed explanations of molecular function. At the same time it offers an ability to explore emergent dynamics in both multiscale and multistate paradigms. Furthermore, it allows for models to be run constantly and updated on the fly, allowing for the possibility of a naturalistic modelling approach.

This naturalistic modelling ideal is perhaps the most important aspect of rule-based modelling, and for computational biochemical models generally. Our common stance is to try and make models bigger – to encompass more of a mechanism sketch. This was the position seen in Chapter 3 with the AMPS model. And whilst it is true that biochemical models have not grown in size over the years, and that more component parts should be and need to be included into these models, this requires context.

Deterministic well-mixed models are an idealised system in a largely undefined context. Stochastic models are an improvement, but still lack fine-grained detail and contextual information. As modellers we need to remember from the start that the systems we explore are noisy, and riddled with assumptions and unknowns. We need to embrace this noisiness wholeheartedly if we are to truly delineate model-based investigations as being something distinctly different to traditional experimental work. My sketch of maximal models is a first stab at how we might approach this.

In Chapter 7 I have shown how concepts quite often requiring abstracted simplification in traditional visual representations can be rendered more naturally and complexly in virtual environments. Importantly, these types of models are accessible across multiple age groups, regardless of the inherent complexity, and can be constructed by learners and educators alike. We have also seen how modelling of this type can blur the lines between experiment and simulation and allow for novel approaches to both human cognitive science and tertiary level pedagogy without losing any scientific rigour.

A subtler thread running through all these findings has been an emphasis on collaboration and reuse. Models have long tended to be tightly tied to their creator's interests and career, which tightens their focus but loses some of their ability to be used in other investigations, or to be added to in meaningful ways. The need for models to be iterative has been interpreted in such a way that leads us to quite often reinvent the wheel, with multiple models of the same molecules in slightly different contexts. Whilst admittedly my attempts to hew more closely to the notion of iteration in part failed, they did highlight that existing models can be merged successfully, and that they can be used to glean new information about a biochemical system.

In Chapters 5 and 7, the models generated were intended to be widely shared and used by others and constructed in such a manner which allows them to grow organically to suit the needs of the user. They provide springboards to encourage collaboration between educator and learner, even experimenter and learner, and require more in the way of imagination and creativity than technical skill. They are not restricted to my intentions as a builder, and can be freely and easily adapted by those who want to do so. The spectre of software obsolescence still looms on the distant horizon, but the wider advantage of promoting the use off-the-shelf software to generate pedagogical experiences provides a longer-term advantage that hopefully inspires others.

The future of modelling needs to be collaborative, iterative, and maximal if it is to meet the ambitions it has as a discipline in driving our understanding forward in ways comparable to traditional approaches. In this spirit, this thesis provides the foundational work for others to take this forward and to persuade others of the potential impact of implementing these ideas would have.

REFERENCES

- ✓ Minecraft Java 1.13 & 1.14 Automatic Minecart Rail Station Tutorial Railway Selector - YouTube. (n.d.). Retrieved May 27, 2020, from <https://www.youtube.com/watch?v=HHerGRJhq4Y&feature=youtu.be>
- Abbott, N. J., Rönnbäck, L., & Hansson, E. (2006, January). Astrocyte-endothelial interactions at the blood-brain barrier. *Nature Reviews Neuroscience*, Vol. 7, pp. 41–53. <https://doi.org/10.1038/nrn1824>
- Adamczyk, A., Mejias, R., Takamiya, K., Yocum, J., Krasnova, I. N., Calderon, J., ... Wang, T. (2012). GluA3-deficiency in mice is associated with increased social and aggressive behavior and elevated dopamine in striatum. *Behavioural Brain Research*, 229(1), 265–272. <https://doi.org/10.1016/j.bbr.2012.01.007>
- Ahmad, D. (2019). "In China, Minecraft is now approaching 300 million registered accounts. The game in China is published by NetEase on PC/Mobile and is free to play. Outside China, the game has sold over 180 million paid copies. Minecraft has 112. Retrieved February 18, 2021, from <https://twitter.com/ZhugeEX/status/1193253634609500160>
- Ait-Haddou, R., & Herzog, W. (2003, May). Brownian ratchet models of molecular motors. *Cell Biochemistry and Biophysics*, Vol. 38, pp. 191–213. <https://doi.org/10.1385/CBB:38:2:191>
- Aitken, S., Alexander, R. D., & Beggs, J. D. (2013). A rule-based kinetic model of RNA polymerase II C-terminal domain phosphorylation. *Journal of The Royal Society Interface*, 10(86), 20130438. <https://doi.org/10.1098/rsif.2013.0438>
- Ak, O., & Kutlu, B. (2017). Comparing 2D and 3D game-based learning environments in terms of learning gains and student perceptions. *British Journal of Educational Technology*, 48(1), 129–144. <https://doi.org/10.1111/bjet.12346>
- Akgül, G., & McBain, C. J. (2016). Diverse roles for ionotropic glutamate receptors on inhibitory interneurons in developing and adult brain. *The Journal of Physiology*, 594(19), 5471–5490. <https://doi.org/10.1113/JP271764>
- Aldridge, B. B., Burke, J. M., Lauffenburger, D. A., & Sorger, P. K. (2006). Physicochemical modelling of cell signalling pathways. TL - 8. *Nature Cell Biology*, 8 VN-re(11), 1195–1203. <https://doi.org/10.1038/ncb1497>
- Altschuler, S. J., & Wu, L. F. (2010, May). Cellular Heterogeneity: Do Differences Make a Difference? *Cell*, Vol. 141, pp. 559–563. <https://doi.org/10.1016/j.cell.2010.04.033>
- American Psychiatric Association. (2013). *Diagnostic and Statistical Manual of Mental Disorders*. <https://doi.org/10.1176/appi.books.9780890425596>
- Anggono, V., Clem, R. L., & Huganir, R. L. (2011). PICK1 loss of function occludes homeostatic synaptic scaling. *Journal of Neuroscience*, 31(6), 2188–2196. <https://doi.org/10.1523/JNEUROSCI.5633-10.2011>
- Anggono, V., & Huganir, R. L. (2012). Regulation of AMPA receptor trafficking and synaptic plasticity. *Current Opinion in Neurobiology*, 22(3), 461–469. <https://doi.org/10.1016/j.conb.2011.12.006>
- Antunes, G., Roque, A. C., & Simoes-de-Souza, F. M. (2016). Stochastic Induction of Long-Term Potentiation and Long-Term Depression. *Scientific Reports*, 6(1), 30899. <https://doi.org/10.1038/srep30899>
- Aow, J., Dore, K., & Malinow, R. (2015). Conformational signaling required for synaptic plasticity by the NMDA receptor complex. *Proceedings of the National Academy of Sciences of the United States of America*, 112(47), 14711–14716. <https://doi.org/10.1073/pnas.1520029112>
- Artola, A., & Singer, W. (1993). Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends in Neurosciences*, 16(11), 480–487. [https://doi.org/10.1016/0166-2236\(93\)90081-V](https://doi.org/10.1016/0166-2236(93)90081-V)
- Ascoli, G. A., Donohue, D. E., & Halavi, M. (2007, August 29). NeuroMorpho.Org: A central resource for neuronal morphologies. *Journal of Neuroscience*, Vol. 27, pp. 9247–9251. <https://doi.org/10.1523/JNEUROSCI.2055-07.2007>
- Asiminas, A. (2016). *Modelling Fragile X Syndrome in rats: New directions in translational research*.
- Asiminas, A., Jackson, A. D., Louros, S. R., Till, S. M., Spano, T., Dando, O., ... Kind, P. C. (2019). Sustained correction

- of associative learning deficits after brief, early treatment in a rat model of Fragile X Syndrome. *Sci. Transl. Med.*, 11. Retrieved from <http://stm.sciencemag.org/>
- Astur, R. S., Ortiz, M. L., & Sutherland, R. J. (1998). A characterization of performance by men and women in a virtual morris water task: A large and reliable sex difference. *Behavioural Brain Research*, 93(1–2), 185–190. [https://doi.org/10.1016/S0166-4328\(98\)00019-9](https://doi.org/10.1016/S0166-4328(98)00019-9)
- Auerbach, B. D., Osterweil, E. K., & Bear, M. F. (2011). Mutations causing syndromic autism define an axis of synaptic pathophysiology. *Nature*, 480(7375), 63–68. <https://doi.org/10.1038/nature10658>
- Ballantyne, R., Bain, J. D., & Packer, J. (1999). Researching university teaching in Australia: Themes and issues in academics' reflections. *Studies in Higher Education*, 24(2), 237–257. <https://doi.org/10.1080/03075079912331379918>
- Banke, T. G., Bowie, D., Lee, H. K., Huganir, R. L., Schousboe, A., & Traynelis, S. F. (2000). Control of GluR1 AMPA receptor function by cAMP-dependent protein kinase. *Journal of Neuroscience*, 20(1), 89–102. <https://doi.org/10.1523/jneurosci.20-01-00089.2000>
- Baral, C. (2000). Abductive reasoning through filtering. *Artificial Intelligence*, 120(1), 1–28. [https://doi.org/10.1016/S0004-3702\(00\)00020-5](https://doi.org/10.1016/S0004-3702(00)00020-5)
- Barr, M. (2017). Video games can develop graduate skills in higher education students: A randomised trial. *Computers and Education*, 113, 86–97. <https://doi.org/10.1016/j.compedu.2017.05.016>
- Barr, M. (2018). Student attitudes to games-based skills development: Learning from video games in higher education. *Computers in Human Behavior*, 80, 283–294. <https://doi.org/10.1016/j.chb.2017.11.030>
- Barria, A., Muller, D., Derkach, V., Griffith, L. C., & Soderling, T. R. (1997). Regulatory phosphorylation of AMPA-type glutamate receptors by CaM-KII during long-term potentiation. *Science*, 276(5321), 2042–2045. <https://doi.org/10.1126/science.276.5321.2042>
- Barrie, S. C. (2006). Understanding what we mean by the generic attributes of graduates. *Higher Education*, 51(2), 215–241. <https://doi.org/10.1007/s10734-004-6384-7>
- Bashir, Z. I., & Collingridge, G. L. (1994). An investigation of depotentiation of long-term potentiation in the CA1 region of the hippocampus. In *Exp Brain Res* (Vol. 100). Retrieved from Springer-Verlag website: https://idp.springer.com/authorize/casa?redirect_uri=https://link.springer.com/content/pdf/10.1007/BF00229183.pdf&casa_token=5RIuzNhFVzEAAAAA:tCBw9GG0FvRQUjKRPv5kvscXXzkKYqkqSpMhb-TEnle3ZhgWY_1us59q78kNc2iM-daj6M1BxVh-58
- Bate, F., MacNish, J., & Males, S. (2014). The politics of gaming in schools: a sociocultural perspective from Western Australia. *Learning, Media and Technology*, 39(3), 306–327. <https://doi.org/10.1080/17439884.2013.872655>
- Bear, M. F., Dölen, G., Osterweil, E., & Nagarajan, N. (2008). Fragile X: translation in action. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, 33(1), 84–87. <https://doi.org/10.1038/sj.npp.1301610>
- Bear, M. F., Huber, K. M., & Warren, S. T. (2004). The mGluR theory of fragile X mental retardation. *Trends in Neurosciences*, 27(7), 370–377. <https://doi.org/10.1016/j.tins.2004.04.009>
- Bell, A. (2007). Designing and testing questionnaires for children. *Journal of Research in Nursing Los Angeles*, 12(5), 461–469. <https://doi.org/10.1177/17449871079616>
- Bender, V. A., Bender, K. J., Brasier, D. J., & Feldman, D. E. (2006). Two coincidence detectors for spike timing-dependent plasticity in somatosensory cortex. *Journal of Neuroscience*, 26(16), 4166–4177. <https://doi.org/10.1523/JNEUROSCI.0176-06.2006>
- Benke, T. A., Lüthi, A., Isaac, J. T. R., & Collingridge, G. L. (1998). Modulation of AMPA receptor unitary conductance by synaptic activity. *Nature*, 393(6687), 793–797. <https://doi.org/10.1038/31709>
- Benke, T., & Traynelis, S. F. (2019). AMPA-Type Glutamate Receptor Conductance Changes and Plasticity: Still a Lot of Noise. *Neurochemical Research*, 44(3), 539–548. <https://doi.org/10.1007/s11064-018-2491-1>
- Berro, J. (2018). “Essentially, all models are wrong, but some are useful”—a cross-disciplinary agenda for building useful models in cell biology and biophysics. *Biophysical Reviews*, 10(6), 1637–1647. <https://doi.org/10.1007/s12551-018-0478-4>

- Bethus, I., Tse, D., & Morris, R. G. M. (2010). Dopamine and Memory: Modulation of the Persistence of Memory for Novel Hippocampal NMDA Receptor-Dependent Paired Associates. *Journal of Neuroscience*, *30*(5), 1610–1618. <https://doi.org/10.1523/JNEUROSCI.2721-09.2010>
- Bezanilla, F. (2008). Ion Channels: From Conductance to Structure. *Neuron*, *60*(3), 456–468. <https://doi.org/10.1016/j.neuron.2008.10.035>
- Bhalla, U. S. (2014). Molecular computation in neurons: A modeling perspective. *Current Opinion in Neurobiology*, *25*, 31–37. <https://doi.org/10.1016/j.conb.2013.11.006>
- Bhalla, U. S., & Iyengar, R. (1999). Emergent properties of networks of biological signaling pathways. *Science*, *283*(5400), 381–387. <https://doi.org/10.1126/science.283.5400.381>
- Biley, F. C., & Smith, K. L. (1999). Making sense of problem-based learning: the perceptions and experiences of undergraduate nursing students. *Journal of Advanced Nursing*, *30*(5), 1205–1212. <https://doi.org/10.1046/j.1365-2648.1999.01188.x>
- Bissen, D., Foss, F., & Acker-Palmer, A. (2019). AMPA receptors and their minions: auxiliary proteins in AMPA receptor trafficking. *Cellular and Molecular Life Sciences*, *76*(11), 2133–2169. <https://doi.org/10.1007/s00018-019-03068-7>
- Blackwell, A. F. (2002). *What is Programming?* Retrieved from www.ppig.org
- Bladt, F., Tafuri, A., Gelkop, S., Langille, L., & Pawson, T. (2002). Epidermolysis bullosa and embryonic lethality in mice lacking the multi-PDZ domain protein GRIP1. *Proceedings of the National Academy of Sciences*, *99*(10), 6816–6821. <https://doi.org/10.1073/pnas.092130099>
- Bliss, T. V. P., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, *232*(2), 331–356. <https://doi.org/10.1113/jphysiol.1973.sp010273>
- Bliss, Tim V.P., & Collingridge, G. L. (2013, January 22). Expression of NMDA receptor-dependent LTP in the hippocampus: bridging the divide. *Molecular Brain*, Vol. 6, p. 5. <https://doi.org/10.1186/1756-6606-6-5>
- Bliss, Tim V.P., Collingridge, G. L., Morris, R. G. M., & Reymann, K. G. (2018). Long-term potentiation in the hippocampus: discovery, mechanisms and function. *Neuroforum*, *24*(3), A103–A120. <https://doi.org/10.1515/nf-2017-A059>
- Blitzer, R. D., Wong, T., Nouranifar, R., Iyengar, R., & Landau, E. M. (1995). Postsynaptic CAMP pathway gates early LTP in hippocampal CA1 region. *Neuron*, *15*(6), 1403–1414. [https://doi.org/10.1016/0896-6273\(95\)90018-7](https://doi.org/10.1016/0896-6273(95)90018-7)
- Boehm, J., Kang, M. G., Johnson, R. C., Esteban, J., Huganir, R. L., & Malinow, R. (2006). Synaptic Incorporation of AMPA Receptors during LTP Is Controlled by a PKC Phosphorylation Site on GluR1. *Neuron*, *51*(2), 213–225. <https://doi.org/10.1016/j.neuron.2006.06.013>
- Borgdorff, A. J., & Choquet, D. (2002). Regulation of AMPA receptor lateral movements. *Nature*, *417*(6889), 649–653. <https://doi.org/10.1038/nature00780>
- Borgers, N., & Hox, J. (2001). Item Nonresponse in Questionnaire Research With Children. In *Journal of Official Statistics* (Vol. 17).
- Borst, J. G. G. (2010). The low synaptic release probability in vivo. *Trends in Neurosciences*, *33*(6), 259–266. <https://doi.org/10.1016/j.tins.2010.03.003>
- Bortolotto, Z. ., & Collingridge, G. . (1998). Involvement of calcium/calmodulin-dependent protein kinases in the setting of a molecular switch involved in hippocampal LTP. *Neuropharmacology*, *37*(4–5), 535–544. [https://doi.org/10.1016/S0028-3908\(98\)00058-6](https://doi.org/10.1016/S0028-3908(98)00058-6)
- Bosch, M., Castro, J., Saneyoshi, T., Matsuno, H., Sur, M., & Hayashi, Y. (2014). Structural and molecular remodeling of dendritic spine substructures during long-term potentiation. *Neuron*, *82*(2), 444–459. <https://doi.org/10.1016/j.neuron.2014.03.021>
- Bosch, M., & Hayashi, Y. (2012). Structural plasticity of dendritic spines. *Current Opinion in Neurobiology*, *22*(3), 383–388. <https://doi.org/10.1016/J.CONB.2011.09.002>
- Boudkkazi, S., Brechet, A., Schwenk, J., & Fakler, B. (2014). Cornichon2 Dictates the Time Course of Excitatory Transmission at Individual Hippocampal Synapses. *Neuron*, *82*(4), 848–858.

- <https://doi.org/10.1016/j.neuron.2014.03.031>
- Boulton, A., Hourizi, R., Jefferies, D., & Guy, A. (2017). A Little Bit of Frustration Can Go a Long Way. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 10664 LNCS (pp. 188–200). https://doi.org/10.1007/978-3-319-71649-7_16
- Bourgonjon, J., Valcke, M., Soetaert, R., & Schellens, T. (2010). Students' perceptions about the use of video games in the classroom. *Computers & Education*, 54(4), 1145–1156. <https://doi.org/10.1016/j.compedu.2009.10.022>
- Boutillier, P., Maasha, M., Li, X., Medina-Abarca, H. F., Krivine, J., Feret, J., ... Fontana, W. (2018). The Kappa platform for rule-based modeling. *Bioinformatics*, 34(13), i583–i592. <https://doi.org/10.1093/bioinformatics/bty272>
- Boyle, A., Maguire, S., Martin, A., Milsom, C., Nash, R., Rawlinson, S., ... Conchie, S. (2007). Fieldwork is Good: the Student Perception and the Affective Domain. *Journal of Geography in Higher Education*, 31(2), 299–317. <https://doi.org/10.1080/03098260601063628>
- Bradshaw, J. M., Hudmon, A., & Schulman, H. (2002). Chemical Quenched Flow Kinetic Studies Indicate an Intraholoenzyme Autophosphorylation Mechanism for Ca²⁺/Calmodulin-dependent Protein Kinase II. *Journal of Biological Chemistry*, 277(23), 20991–20998. <https://doi.org/10.1074/jbc.M202154200>
- Branco, T., & Staras, K. (2009). The probability of neurotransmitter release: Variability and feedback control at single synapses. *Nature Reviews Neuroscience*, 10(5), 373–383. <https://doi.org/10.1038/nrn2634>
- Broadhead, M. J., Horrocks, M. H., Zhu, F., Muresan, L., Benavides-Piccione, R., DeFelipe, J., ... Grant, S. G. N. (2016). PSD95 nanoclusters are postsynaptic building blocks in hippocampus circuits. *Scientific Reports*, 6. <https://doi.org/10.1038/srep24626>
- Brockie, P. J., Jensen, M., Mellem, J. E., Jensen, E., Yamasaki, T., Wang, R., ... Maricq, A. V. (2013). Cornichons Control ER Export of AMPA Receptors to Regulate Synaptic Excitability. *Neuron*, 80(1), 129–142. <https://doi.org/10.1016/j.neuron.2013.07.028>
- Brunyé, T. T., Gardony, A. L., Holmes, A., & Taylor, H. A. (2018). Spatial decision dynamics during wayfinding: intersections prompt the decision-making process. *Cognitive Research: Principles and Implications* 2018 3:1, 3(1), 1–19. <https://doi.org/10.1186/S41235-018-0098-3>
- Büchel, F., Rodriguez, N., Swainston, N., Wrzodek, C., Czauderna, T., Keller, R., ... Le Novère, N. (2013). Path2Models: Large-scale generation of computational models from biochemical pathway maps. *BMC Systems Biology*, 7. <https://doi.org/10.1186/1752-0509-7-116>
- Buonarati, O. R., Hammes, E. A., Watson, J. F., Greger, I. H., & Hell, J. W. (2019). Mechanisms of postsynaptic localization of AMPA-type glutamate receptors and their regulation during long-term potentiation. *Science Signaling*, 12(562), 6889. <https://doi.org/10.1126/scisignal.aar6889>
- Buzsáki, G., & Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, 16(2), 130–138. <https://doi.org/10.1038/nn.3304>
- Buzsáki, G., Peyrache, A., & Kubie, J. (2014). Emergence of cognition from action. *Cold Spring Harbor Symposia on Quantitative Biology*, 79, 41–50. <https://doi.org/10.1101/sqb.2014.79.024679>
- Caillois, R. (1961). *Man, Play, Games*. Champaign, IL: University of Illinois Press.
- Cais, O., Herguedas, B., Krol, K., Cull-Candy, S. G., Farrant, M., & Greger, I. H. (2014). Mapping the interaction sites between AMPA receptors and TARPs reveals a role for the receptor N-terminal domain in channel gating. *Cell Reports*, 9(2), 728–740. <https://doi.org/10.1016/j.celrep.2014.09.029>
- Canossa, A., Drachen, A., & Sørensen, J. R. M. (2011). Arrrgghh!!! *Proceedings of the 6th International Conference on Foundations of Digital Games - FDG '11*, 61–68. <https://doi.org/10.1145/2159365.2159374>
- Carballo, G. B., Honorato, J. R., De Lopes, G. P. F., & Spohr, T. C. L. D. S. E. (2018, March 20). A highlight on Sonic hedgehog pathway. *Cell Communication and Signaling*, Vol. 16, pp. 1–15. <https://doi.org/10.1186/s12964-018-0220-7>
- Carbonneau, N., Vallerand, R. J., Fernet, C., & Guay, F. (2008). *The Role of Passion for Teaching in Intrapersonal and Interpersonal Outcomes*. <https://doi.org/10.1037/a0012545>
- Carvalho, A. L., Kameyama, K., & Huganir, R. L. (1999). Characterization of phosphorylation sites on the glutamate receptor 4 subunit of the AMPA receptors. *Journal of Neuroscience*, 19(12), 4748–4754.

- <https://doi.org/10.1523/jneurosci.19-12-04748.1999>
- Cerpa, W., Gambrell, A., Inestrosa, N. C., & Barria, A. (2011). Regulation of NMDA-receptor synaptic transmission by Wnt signaling. *Journal of Neuroscience*, 31(26), 9466–9471. <https://doi.org/10.1523/JNEUROSCI.6311-10.2011>
- Chen, C.-H., Chang, W.-C., & Chang, W.-T. (2009). Gender differences in relation to wayfinding strategies, navigational support design, and wayfinding task difficulty. *Journal of Environmental Psychology*, 29(2), 220–226. <https://doi.org/10.1016/j.jenvp.2008.07.003>
- Chen, L., & Toth, M. (2001). Fragile X mice develop sensory hyperreactivity to auditory stimuli. *Neuroscience*, 103(4), 1043–1050.
- Chen, Lu, Chetkovich, D. M., Petralia, R. S., Sweeney, N. T., Kawasaki, Y., Wenthold, R. J., ... Nicoll, R. A. (2000). Stargazin regulates synaptic targeting of AMPA receptors by two distinct mechanisms. *Nature*, 408(6815), 936–943. <https://doi.org/10.1038/35050030>
- Chen, W. W., Niepel, M., & Sorger, P. K. (2010). Classic and contemporary approaches to modeling biochemical reactions. *Genes and Development*, 24(17), 1861–1875. <https://doi.org/10.1101/gad.1945410>
- Chen, Y., Sasai, N., Ma, G., Yue, T., Jia, J., Briscoe, J., & Jiang, J. (2011). Sonic Hedgehog Dependent Phosphorylation by CK1 α and GRK2 Is Required for Ciliary Accumulation and Activation of Smoothened. *PLoS Biology*, 9(6), e1001083. <https://doi.org/10.1371/journal.pbio.1001083>
- Chetkovich, D. M., Chen, L., Stocker, T. J., Nicoll, R. A., & Brecht, D. S. (2002). Phosphorylation of the postsynaptic density-95 (PSD-95)/discs large/zona occludens-1 binding site of stargazin regulates binding to PSD-95 and synaptic targeting of AMPA receptors. *Journal of Neuroscience*, 22(14), 5791–5796. <https://doi.org/10.1523/jneurosci.22-14-05791.2002>
- Chiang, H. (2020). Minecraft: Connecting More Players Than Ever Before - Xbox Wire. Retrieved January 27, 2021, from <https://news.xbox.com/en-us/2020/05/18/minecraft-connecting-more-players-than-ever-before/>
- Chirimuuta, M. (2018, April 3). Marr, Mayr, and MR: What functionalism should now be about. *Philosophical Psychology*, Vol. 31, pp. 403–418. <https://doi.org/10.1080/09515089.2017.1381679>
- Chiu, S. L., Diering, G. H., Ye, B., Takamiya, K., Chen, C. M., Jiang, Y., ... Huganir, R. L. (2017). GRASP1 Regulates Synaptic Plasticity and Learning through Endosomal Recycling of AMPA Receptors. *Neuron*, 93(6), 1405–1419.e8. <https://doi.org/10.1016/j.neuron.2017.02.031>
- Chowdhury, D. (2013, August 1). Stochastic mechano-chemical kinetics of molecular motors: A multidisciplinary enterprise from a physicist's perspective. *Physics Reports*, Vol. 529, pp. 1–197. <https://doi.org/10.1016/j.physrep.2013.03.005>
- Chu, Y. X., Fioravante, D., Leitges, M., & Regehr, W. G. (2014). Calcium-Dependent PKC Isoforms Have Specialized Roles in Short-Term Synaptic Plasticity. *Neuron*, 82(4), 859–871. <https://doi.org/10.1016/j.neuron.2014.04.003>
- Cilaurio, F., & Paull, G. (2019). *Evaluation of Nuffield Research Placements: interim report*. London.
- Ciliberto, A., Capuani, F., & Tyson, J. J. (2007). Modeling Networks of Coupled Enzymatic Reactions Using the Total Quasi-Steady State Approximation. *PLoS Computational Biology*, 3(3), e45. <https://doi.org/10.1371/journal.pcbi.0030045>
- Citri, A., Bhattacharyya, S., Ma, C., Morishita, W., Fang, S., Rizo, J., & Malenka, R. C. (2010). Calcium binding to PICK1 is essential for the intracellular retention of AMPA receptors underlying long-term depression. *Journal of Neuroscience*, 30(49), 16437–16452. <https://doi.org/10.1523/JNEUROSCI.4478-10.2010>
- Clague, M. J., & Urbé, S. (2010). Ubiquitin: Same molecule, different degradation pathways. *Cell*, 143(5), 682–685. <https://doi.org/10.1016/j.cell.2010.11.012>
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. <https://doi.org/10.1093/acprof:oso/9780190217013.001.0001>
- Clemenson, G. D., Henningfield, C. M., & Stark, C. E. L. (2019). Improving hippocampal memory through the experience of a rich minecraft environment. *Frontiers in Behavioral Neuroscience*, 13. <https://doi.org/10.3389/fnbeh.2019.00057>
- Coakley, M. F., Hurt, D. E., Weber, N., Mtingwa, M., Fincher, E. C., Alekseyev, V., ... Huyen, Y. (2014). The NIH 3D

- print exchange: A public resource for bioscientific and biomedical 3D prints. *3D Printing and Additive Manufacturing*, 1(3), 137–140. <https://doi.org/10.1089/3dp.2014.1503>
- Cohen, L. D., & Ziv, N. E. (2019, August 1). Neuronal and synaptic protein lifetimes. *Current Opinion in Neurobiology*, Vol. 57, pp. 9–16. <https://doi.org/10.1016/j.conb.2018.12.007>
- Cohen, M., Page, K. M., Perez-Carrasco, R., Barnes, C. P., & Briscoe, J. (2014). A theoretical framework for the regulation of shh morphogen-controlled gene expression. *Development (Cambridge)*, 141(20), 3868–3878. <https://doi.org/10.1242/dev.112573>
- Coleman, T. E., & Money, A. G. (2020). Student-centred digital game-based learning: a conceptual framework and survey of the state of the art. *Higher Education*, 79(3), 415–457. <https://doi.org/10.1007/s10734-019-00417-0>
- Collingridge, G. L., Peineau, S., Howland, J. G., & Wang, Y. T. (2010). Long-term depression in the CNS. *Nature Reviews Neuroscience*, 11(7), 459–473. <https://doi.org/10.1038/nrn2867>
- Compans, B., Choquet, D., & Hosy, E. (2016). Review on the role of AMPA receptor nano-organization and dynamic in the properties of synaptic transmission. *Neurophotonics*, 3(4), 041811. <https://doi.org/10.1117/1.nph.3.4.041811>
- Cooke, M., & Moyle, K. (2002). Students' evaluation of problem-based learning. *Nurse Education Today*, 22(4), 330–339. <https://doi.org/10.1054/nedt.2001.0713>
- Cooney, J. R., Hurlburt, J. L., Selig, D. K., Harris, K. M., & Fiala, J. C. (2002). Endosomal Compartments Serve Multiple Hippocampal Dendritic Spines from a Widespread Rather Than a Local Store of Recycling Membrane. *Journal of Neuroscience*, 22(6), 2215–2224. <https://doi.org/10.1523/JNEUROSCI.22-06-02215.2002>
- Cooper, L. N., & Bear, M. F. (2012). The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nature Reviews Neuroscience*, 13(11), 798–810. <https://doi.org/10.1038/nrn3353>
- Cornish, K. M., Munir, F., & Cross, G. (2001). Differential Impact of the FMR-1 Full Mutation on Memory and Attention Functioning: A Neuropsychological Perspective. *Journal of Cognitive Neuroscience*, 13(1), 144–150. <https://doi.org/10.1162/089892901564126>
- Coulombe, J., Traiffort, E., Loulier, K., Faure, H., & Ruat, M. (2004). Hedgehog interacting protein in the mature brain: membrane-associated and soluble forms. *Molecular and Cellular Neuroscience*, 25(2), 323–333. <https://doi.org/10.1016/j.mcn.2003.10.024>
- Coultrap, S. J., Freund, R. K., O'Leary, H., Sanderson, J. L., Roche, K. W., Dell'Acqua, M. L., & Bayer, K. U. (2014). Autonomous CaMKII mediates both LTP and LTD using a mechanism for differential substrate site selection. *Cell Reports*, 8(3), 431–437. <https://doi.org/10.1016/j.celrep.2014.01.005>
- Coveney, P. V., Dougherty, E. R., & Highfeld, R. R. (2016). Big data need big theory too. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2080). <https://doi.org/10.1098/rsta.2016.0153>
- Craver, C. F. (2002). Interlevel Experiments and Multilevel Mechanisms in the Neuroscience of Memory. *Philosophy of Science*, 69(S3), S83–S97. <https://doi.org/10.1086/341836>
- Craver, C. F. (2007). Explaining the Brain. In *Explaining the Brain*. <https://doi.org/10.1093/acprof:oso/9780199299317.001.0001>
- Craver, C. F. (2013). Functions and mechanisms: a perspectivalist view. In P. Huneman (Ed.), *Functions: Selection and Mechanisms*. Springer.
- Curriculum for Excellence: Experiences and Outcomes. (2008). In *Education Scotland*. Retrieved from <https://education.gov.scot/Documents/All-experiencesoutcomes18.pdf>
- Czöndör, K., Mondin, M., Garcia, M., Heine, M., Frischknecht, R., Choquet, D., ... Thoumine, O. R. (2012). Unified quantitative model of AMPA receptor trafficking at synapses. *Proceedings of the National Academy of Sciences of the United States of America*, 109(9), 3522–3527. <https://doi.org/10.1073/pnas.1109818109>
- D506–D515. (2019). UniProt: a worldwide hub of protein knowledge. *Nucleic Acids Research*, 47(D1), D506–D515. <https://doi.org/10.1093/nar/gky1049>
- Dahlhaus, R. (2018). Of Men and Mice: Modeling the Fragile X Syndrome. *Frontiers in Molecular Neuroscience*, 11, 41. <https://doi.org/10.3389/fnmol.2018.00041>

- Dakoji, S., Tomita, S., Karimzadegan, S., Nicoll, R. A., & Bredt, D. S. (2003). Interaction of transmembrane AMPA receptor regulatory proteins with multiple membrane associated guanylate kinases. *Neuropharmacology*, 45(6), 849–856. [https://doi.org/10.1016/S0028-3908\(03\)00267-3](https://doi.org/10.1016/S0028-3908(03)00267-3)
- Damasch, I. E. (1989). Structural realization of a Hebb-type learning rule [Book]. In R. Cotterill (Ed.), *Models of brain function* (First, pp. 539–553). Cambridge: Cambridge University Press.
- Daneman, R., & Prat, A. (2015). The blood–brain barrier. *Cold Spring Harbor Perspectives in Biology*, 7(1). <https://doi.org/10.1101/cshperspect.a020412>
- Danos, V., Feret, J., Fontana, W., Harmer, R., & Krivine, J. (2007). Rule-Based Modelling of Cellular Signalling. In L. Caires & V. T. Vasconcelos (Eds.), *CONCUR 2007 -- Concurrency Theory. Vol. 4703 LNCS* (pp. 17–41). https://doi.org/10.1007/978-3-540-74407-8_3
- Danos, V., & Laneve, C. (2004). Formal molecular biology. *Theoretical Computer Science*, 325(1), 69–110. <https://doi.org/10.1016/j.tcs.2004.03.065>
- De Brigard, F. (2014). The Nature of Memory Traces. *Philosophy Compass*, 9(6), 402–414. <https://doi.org/10.1111/phc3.12133>
- De Leeuw, E. D. (2011). *Improving Data Quality when Surveying Children and Adolescents: Cognitive and Social Development and its Role in Questionnaire Construction and Pretesting Report prepared for the Annual Meeting of the Academy of Finland: Research Programs Public Health Ch.*
- Dedoose Version 8.0.35, web application for managing, analyzing, and presenting qualitative and mixed method research data. (2018). Retrieved from <https://app.dedoose.com/App/?Version=8.3.41%0A%0A>
- Delgado, J. Y., & O'Dell, T. J. (2005). Long-term potentiation persists in an occult state following mGluR-dependent depotentiation. *Neuropharmacology*, 48(7), 936–948. <https://doi.org/10.1016/j.neuropharm.2005.01.008>
- Derkach, V., Barria, A., & Soderling, T. R. (1999). Ca²⁺/calmodulin-kinase II enhances channel conductance of α -amino-3-hydroxy-5-methyl-4-isoxazolepropionate type glutamate receptors. *Proceedings of the National Academy of Sciences of the United States of America*, 96(6), 3269–3274. <https://doi.org/10.1073/pnas.96.6.3269>
- DeRouin-Jessen, R. E. (2008). Game on: The Impact of Game Features in Computer-Based Training (ProQuest LLC). Retrieved from <http://search.ebscohost.com/login.aspx?direct=true&db=eric&AN=ED527400&site=ehost-live>
- DeRubeis, S., Pasciuto, E., Li, K. W., Fernández, E., DiMarino, D., Buzzi, A., ... Bagni, C. (2013). CYFIP1 coordinates mRNA translation and cytoskeleton remodeling to ensure proper dendritic Spine formation. *Neuron*, 79(6), 1169–1182. <https://doi.org/10.1016/j.neuron.2013.06.039>
- Desmond, N. L., & Levy, W. B. (1990). Morphological correlates of long-term potentiation imply the modification of existing synapses, not synaptogenesis, in the hippocampal dentate gyrus. *Synapse*, 5(2), 139–143. <https://doi.org/10.1002/syn.890050208>
- Deterding, S., Dixon, D., Khaled, R., & Nacke, L. (2011). From game design elements to gamefulness: Defining “gamification.” *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments, MindTrek 2011*, 9–15. <https://doi.org/10.1145/2181037.2181040>
- Devonshire, I. M., Grandy, T. H., Dommert, E. J., & Greenfield, S. A. (2010). Effects of urethane anaesthesia on sensory processing in the rat barrel cortex revealed by combined optical imaging and electrophysiology. *European Journal of Neuroscience*, 32(5), 786–797. <https://doi.org/10.1111/j.1460-9568.2010.07322.x>
- Diering, G. H., Heo, S., Hussain, N. K., Liu, B., & Huganir, R. L. (2016). Extensive phosphorylation of AMPA receptors in neurons. *Proceedings of the National Academy of Sciences*, 113(33), E4920–E4927. <https://doi.org/10.1073/pnas.1610631113>
- Diering, G. H., & Huganir, R. L. (2018). The AMPA Receptor Code of Synaptic Plasticity. *Neuron*, 100(2), 314–329. <https://doi.org/10.1016/j.neuron.2018.10.018>
- Djakovic, S. N., Schwarz, L. A., Barylko, B., DeMartino, G. N., & Patrick, G. N. (2009). Regulation of the proteasome by neuronal activity and calcium/calmodulin-dependent protein kinase II. *Journal of Biological Chemistry*, 284(39), 26655–26665. <https://doi.org/10.1074/jbc.M109.021956>
- Dölen, G., Osterweil, E., Rao, B. S. S., Smith, G. B., Auerbach, B. D., Chattarji, S., & Bear, M. F. (2007). Correction of Fragile X Syndrome in Mice. *Neuron*, 56(6), 955–962. <https://doi.org/10.1016/j.neuron.2007.12.001>

- Dolzanskaya, N., Merz, G., Aletta, J. M., & Denman, R. B. (2006). Methylation regulates the intracellular protein-protein and protein-RNA interactions of FMRP. *Journal of Cell Science*, 119(9), 1933–1946. <https://doi.org/10.1242/jcs.02882>
- Dong, H., O'Brien, R. J., Fung, E. T., Lanahan, A. A., Worley, P. F., & Huganir, R. L. (1997). GRIP: A synaptic PDZ domain-containing protein that interacts with AMPA receptors. *Nature*, 386(6622), 279–284. <https://doi.org/10.1038/386279a0>
- Dore, K., Aow, J., & Malinow, R. (2016). The Emergence of NMDA Receptor Metabotropic Function: Insights from Imaging. *Frontiers in Synaptic Neuroscience*, 8, 20. <https://doi.org/10.3389/fnsyn.2016.00020>
- Dörbaum, A. R., Kochen, L., Langer, J. D., & Schuman, E. M. (2018). Local and global influences on protein turnover in neurons and glia. *eLife*, 7. <https://doi.org/10.7554/eLife.34202>
- Dudani, N., Bhalla, U. S., & Ray, S. (2014). MOOSE, the Multiscale Object-Oriented Simulation Environment. In *Encyclopedia of Computational Neuroscience* (pp. 1–4). https://doi.org/10.1007/978-1-4614-7320-6_257-1
- Dutta-Roy, R., Rosenmund, C., Edelstein, S. J., & Le Novère, N. (2015). Ligand-Dependent Opening of the Multiple AMPA Receptor Conductance States: A Concerted Model. *PLOS ONE*, 10(1), e0116616. <https://doi.org/10.1371/journal.pone.0116616>
- Earnshaw, B. A., & Bressloff, P. C. (2006). Biophysical Model of AMPA Receptor Trafficking and Its Regulation during Long-Term Potentiation / Long-Term Depression. *Receptor*, 26(47), 12362–12373. <https://doi.org/10.1523/JNEUROSCI.3601-06.2006>
- Edelstein, S. J., Stefan, M. I., & Le Novère, N. (2010). Ligand depletion in vivo modulates the dynamic range and cooperativity of signal transduction. *PLoS ONE*, 5(1), e8449. <https://doi.org/10.1371/journal.pone.0008449>
- Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. *Journal of Neurophysiology*, 117(4), 1785–1796. <https://doi.org/10.1152/jn.00005.2017>
- El-Danaf, R. N., & Huberman, A. D. (2015). Characteristic patterns of dendritic remodeling in early-stage glaucoma: Evidence from genetically identified retinal ganglion cell types. *Journal of Neuroscience*, 35(6), 2329–2343. <https://doi.org/10.1523/JNEUROSCI.1419-14.2015>
- Engert, F., & Bonhoeffer, T. (1999). Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature*, 399(6731), 66–70. <https://doi.org/10.1038/19978>
- Erickson, C. A., Davenport, M. H., Schaefer, T. L., Wink, L. K., Pedapati, E. V., Sweeney, J. A., ... Berry-Kravis, E. (2017, June 12). Fragile X targeted pharmacotherapy: Lessons learned and future directions. *Journal of Neurodevelopmental Disorders*, Vol. 9. <https://doi.org/10.1186/s11689-017-9186-9>
- Eronen, M. (2008). Explaining the Brain: Ruthless Reductionism or Multilevel Mechanisms? 1. In *Papers of the 31st International Wittgenstein Symposium, Vol. XVI* (pp. 86–88).
- Esteban, J. A., Shi, S. H., Wilson, C., Nuriya, M., Huganir, R. L., & Malinow, R. (2003). PKA phosphorylation of AMPA receptor subunits controls synaptic trafficking underlying plasticity. *Nature Neuroscience*, 6(2), 136–143. <https://doi.org/10.1038/nn997>
- Faas, T., & Lin, C. (2017). Self-directed learning in teacher-lead Minecraft classrooms. *Conference on Human Factors in Computing Systems - Proceedings, Part F127655*, 2569–2575. <https://doi.org/10.1145/3027063.3053269>
- Fariás, G. G., Alfaro, I. E., Cerpa, W., Grabowski, C. P., Godoy, J. A., Bonansco, C., & Inestrosa, N. C. (2009). Wnt-5a/JNK Signaling Promotes the Clustering of PSD-95 in Hippocampal Neurons. *Journal of Biological Chemistry*, 284(23), 15857–15866. <https://doi.org/10.1074/jbc.M808986200>
- Fernández, E., Rajan, N., & Bagni, C. (2013). The FMRP regulon: from targets to disease convergence. *Frontiers in Neuroscience*, 7(October), 191. <https://doi.org/10.3389/fnins.2013.00191>
- Fitzpatrick, R. (2016). *ENACTING MEMORY Rethinking our Capacity to Remember* (The University of Edinburgh). Retrieved from https://project-archive.inf.ed.ac.uk/msc/20161829/msc_proj.pdf
- Fiuza, M., Rostovsky, C. M., Parkinson, G. T., Bygrave, A. M., Halemani, N., Baptista, M., ... Hanley, J. G. (2017). PICK1 regulates AMPA receptor endocytosis via direct interactions with AP2 α -appendage and dynamin. *Journal of Cell Biology*, 216(10), 3323–3338. <https://doi.org/10.1083/jcb.201701034>

- Fox, K., & Stryker, M. (2017). Integrating Hebbian and homeostatic plasticity: Introduction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715). <https://doi.org/10.1098/rstb.2016.0413>
- French, A. R., & Lauffenburger, D. A. (1996). Intracellular receptor/ligand sorting based on endosomal retention components. *Biotechnology and Bioengineering*, 51(3), 281–297. [https://doi.org/10.1002/\(SICI\)1097-0290\(19960805\)51:3<281::AID-BIT4>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1097-0290(19960805)51:3<281::AID-BIT4>3.0.CO;2-9)
- Fu, A. K. Y., Hung, K. W., Fu, W. Y., Shen, C., Chen, Y., Xia, J., ... Ip, N. Y. (2011). APCCdh1 mediates EphA4-dependent downregulation of AMPA receptors in homeostatic plasticity. *Nature Neuroscience*, 14(2), 181–191. <https://doi.org/10.1038/nn.2715>
- Galanis, C., & Vlachos, A. (2020). Hebbian and Homeostatic Synaptic Plasticity—Do Alterations of One Reflect Enhancement of the Other? *Frontiers in Cellular Neuroscience*, 14, 50. <https://doi.org/10.3389/fncel.2020.00050>
- Gallimore, A. R., Aricescu, A. R., Yuzaki, M., & Calinescu, R. (2016). A Computational Model for the AMPA Receptor Phosphorylation Master Switch Regulating Cerebellar Long-Term Depression. *PLoS Computational Biology*, 12(1), 1–23. <https://doi.org/10.1371/journal.pcbi.1004664>
- Gallimore, A. R., Kim, T., Tanaka-Yamamoto, K., & De Schutter, E. (2018). Switching On Depression and Potentiation in the Cerebellum. *Cell Reports*, 22(3), 722–733. <https://doi.org/10.1016/j.celrep.2017.12.084>
- Gallinaro, J. V., & Rotter, S. (2018). Associative properties of structural plasticity based on firing rate homeostasis in recurrent neuronal networks. *Scientific Reports*, 8(1). <https://doi.org/10.1038/s41598-018-22077-3>
- Gan, Q., Salussolia, C. L., & Wollmuth, L. P. (2015). Assembly of AMPA receptors: mechanisms and regulation. *The Journal of Physiology*, 593(1), 39–48. <https://doi.org/10.1113/jphysiol.2014.273755>
- García-Nafria, J., Herguedas, B., Watson, J. F., & Greger, I. H. (2016). The dynamic AMPA receptor extracellular region: a platform for synaptic protein interactions. *The Journal of Physiology*, 594(19), 5449–5458. <https://doi.org/10.1113/JP271844>
- Garnier, S. (2018). *viridis: Default Color Maps from "matplotlib"*. Retrieved from <https://cran.r-project.org/package=viridis%0A>
- Gee, J. P. (2005). Learning by Design: Good Video Games as Learning Machines. *E-Learning and Digital Media*, 2(1), 5–16. <https://doi.org/10.2304/elea.2005.2.1.5>
- Genzel, L. (2021). How to control behavioral studies for rodents— don't project human thoughts onto them. *ENeuro*, 8(1), 1–7. <https://doi.org/10.1523/ENEURO.0456-20.2021>
- Genzel, L., Rossato, J. I., Jacobse, J., Grieves, R. M., Spooner, P. A., Battaglia, F. P., ... Morris, R. G. M. (2017). The Yin and Yang of Memory Consolidation: Hippocampal and Neocortical. *PLOS Biology*, 15(1), e2000531. <https://doi.org/10.1371/journal.pbio.2000531>
- Gill, M. B., Kato, A. S., Roberts, M. F., Yu, H., Wang, H., Tomita, S., & Bredt, D. S. (2011). Cornichon-2 modulates AMPA receptor-transmembrane AMPA receptor regulatory protein assembly to dictate gating and pharmacology. *Journal of Neuroscience*, 31(18), 6928–6938. <https://doi.org/10.1523/JNEUROSCI.6271-10.2011>
- Gilleade, K. M., & Dix, A. (2004). *Using Frustration in the Design of Adaptive Videogames*.
- Glont, M., Arankalle, C., Tiwari, K., Nguyen, T. V. N., Hermjakob, H., & Malik-Sheriff, R. S. (2020). BioModels parameters: A treasure trove of parameter values from published systems biology models. *Bioinformatics*, 36(17), 4649–4654. <https://doi.org/10.1093/bioinformatics/btaa560>
- Goeke, C., Kornpetpanee, S., Köster, M., Fernández-Revelles, A. B., Gramann, K., & König, P. (2015). Cultural background shapes spatial reference frame proclivity. *Scientific Reports*, 5(1), 1–13. <https://doi.org/10.1038/srep11426>
- Goldstein, L. S. B., & Yang, Z. (2000). Microtubule-based transport systems in neurons: The roles of kinesins and dyneins. *Annual Review of Neuroscience*, Vol. 23, pp. 39–71. <https://doi.org/10.1146/annurev.neuro.23.1.39>
- Gonzalez, J., Morales, I. S., Villarreal, D. M., & Derrick, B. E. (2014). Low-frequency stimulation induces long-term depression and slow onset long-term potentiation at perforant path-dentate gyrus synapses in vivo. *Journal of Neurophysiology*, 111(6), 1259–1273. <https://doi.org/10.1152/jn.00941.2012>
- Goodrich-Hunsaker, N. J., Livingstone, S. A., Skelton, R. W., & Hopkins, R. O. (2009). Spatial deficits in a virtual

- water maze in amnesic participants with hippocampal damage. *Hippocampus*, 20(4), NA-NA. <https://doi.org/10.1002/hipo.20651>
- Gouirand, A. M., & Matuszewich, L. (2005). The effects of chronic unpredictable stress on male rats in the water maze. *Physiology & Behavior*, 86(1–2), 21–31. <https://doi.org/10.1016/J.PHYSBEH.2005.06.027>
- Granger, A. J., & Nicoll, R. A. (2014). Expression mechanisms underlying long-term potentiation: a postsynaptic view, 10 years on. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633), 20130136. <https://doi.org/10.1098/rstb.2013.0136>
- Granger, A. J., Shi, Y., Lu, W., Cerpas, M., & Nicoll, R. A. (2013). LTP requires a reserve pool of glutamate receptors independent of subunit type. *Nature*, 493(7433), 495–500. <https://doi.org/10.1038/nature11775>
- Greca, I. M., & Moreira, M. A. (2000). Mental models, conceptual models, and modelling. *International Journal of Science Education*, 22(1), 1–11. <https://doi.org/10.1080/095006900289976>
- Green, W., Hammer, S., & Star, C. (2009). Facing up to the challenge: why is it so hard to develop graduate attributes? *Higher Education Research & Development*, 28(1), 17–29. <https://doi.org/10.1080/07294360802444339>
- Greengard, P., Jen, J., Nairn, A. C., & Stevens, C. F. (1991). Enhancement of the glutamate response by cAMP-dependent protein kinase in hippocampal neurons. *Science*, 253(5024), 1135–1138. <https://doi.org/10.1126/science.1716001>
- Greger, I. H., & Esteban, J. A. (2007). AMPA receptor biogenesis and trafficking. *Current Opinion in Neurobiology*, 17(3), 289–297. <https://doi.org/10.1016/j.conb.2007.04.007>
- Grieves, R. M., & Jeffery, K. J. (2017). The representation of space in the brain. *Behavioural Processes*, 135, 113–131. <https://doi.org/10.1016/j.beproc.2016.12.012>
- Grover, L. M., Kim, E., Cooke, J. D., & Holmes, W. R. (2009). LTP in hippocampal area CA1 is induced by burst stimulation over a broad frequency range centered around delta. *Learning & Memory (Cold Spring Harbor, N.Y.)*, 16(1), 69–81. <https://doi.org/10.1101/lm.1179109>
- Guéniot, L., Lepere, V., De Medeiros, G. F., Danckaert, A., Flamant, P., Le Dudal, M., ... Jouvion, G. (2020). Muscle Injury Induces Postoperative Cognitive Dysfunction. *Scientific Reports*, 10(1), 1–11. <https://doi.org/10.1038/s41598-020-59639-3>
- Guest, O., & Martin, A. (2020). *How computational modeling can force theory building in psychological science*. <https://doi.org/10.31234/osf.io/rybh9>
- Guntupalli, S., Jang, S. E., Zhu, T., Hugarir, R. L., Widagdo, J., & Anggono, V. (2017). GluA1 subunit ubiquitination mediates amyloid- β -induced loss of surface α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) receptors. *Journal of Biological Chemistry*, 292(20), 8186–8194. <https://doi.org/10.1074/jbc.M116.774554>
- Hagerman, R., Hoem, G., & Hagerman, P. (2010, September 21). Fragile X and autism: Intertwined at the molecular level leading to targeted treatments. *Molecular Autism*, Vol. 1, pp. 1–14. <https://doi.org/10.1186/2040-2392-1-12>
- Hanson, P. I., Meyer, T., Stryer, L., & Schulman, H. (1994). Dual role of calmodulin in autophosphorylation of multifunctional cam kinase may underlie decoding of calcium signals. *Neuron*, 12(5), 943–956. [https://doi.org/10.1016/0896-6273\(94\)90306-9](https://doi.org/10.1016/0896-6273(94)90306-9)
- Harris, J. J., Jolivet, R., & Attwell, D. (2012). Synaptic Energy Use and Supply. *Neuron*, 75(5), 762–777. <https://doi.org/10.1016/J.NEURON.2012.08.019>
- Hayer, A., & Bhalla, U. S. (2005). Molecular Switches at the Synapse Emerge from Receptor and Kinase Traffic. *PLoS Computational Biology*, 1(2), e20. <https://doi.org/10.1371/journal.pcbi.0010020>
- Hayes, T. L. (1996). Magnopyramidal Neurons in the Anterior Motor Speech Region. *Archives of Neurology*, 53(12), 1277. <https://doi.org/10.1001/archneur.1996.00550120089021>
- He, K., Song, L., Cummings, L. W., Goldman, J., Hugarir, R. L., & Lee, H. K. (2009). Stabilization of Ca²⁺-permeable AMPA receptors at perisynaptic sites by GluR1-S845 phosphorylation. *Proceedings of the National Academy of Sciences of the United States of America*, 106(47), 20033–20038. <https://doi.org/10.1073/pnas.0910338106>

- He, M., Subramanian, R., Bangs, F., Omelchenko, T., Liem, K. F., Kapoor, T. M., & Anderson, K. V. (2014). The kinesin-4 protein Kif7 regulates mammalian Hedgehog signalling by organizing the cilium tip compartment. *Nature Cell Biology*, 16(7), 663–672. <https://doi.org/10.1038/ncb2988>
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley & Sons, Inc.
- Heil, K. F., Wysocka, E. M., Sorokina, O., Kotaleski, J. H., Simpson, T. I., Armstrong, J. D., & Sterratt, D. C. (2018). Analysis of proteins in computational models of synaptic plasticity. *BioRxiv*, 254094. <https://doi.org/10.1101/254094>
- Helm, M. (2019). *A quantitative analysis of the molecular organization of dendritic spines from hippocampal neurons submitted by*. Georg-August-Universität Göttingen.
- Henley, J. M., & Wilkinson, K. a. (2016). Synaptic AMPA receptor composition in development, plasticity and disease. *Nature Reviews. Neuroscience*, 17(6), 337–350. <https://doi.org/10.1038/nrn.2016.37>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature*, 466(7302). <https://doi.org/10.1038/466029a>
- Herreras, O. (2016). Local Field Potentials: Myths and Misunderstandings. *Frontiers in Neural Circuits*, 10(DEC), 101. <https://doi.org/10.3389/fncir.2016.00101>
- Herring, B. E., Shi, Y., Suh, Y. H., Zheng, C. Y., Blankenship, S. M., Roche, K. W., & Nicoll, R. A. (2013). Cornichon Proteins Determine the Subunit Composition of Synaptic AMPA Receptors. *Neuron*, 77(6), 1083–1096. <https://doi.org/10.1016/j.neuron.2013.01.017>
- Hethcote, H. W. (1976). Qualitative analyses of communicable disease models. *Mathematical Biosciences*, 28(3–4), 335–356. [https://doi.org/10.1016/0025-5564\(76\)90132-2](https://doi.org/10.1016/0025-5564(76)90132-2)
- Heynen, A. J., Abraham, W. C., & Bear, M. F. (1996). Bidirectional modification of CA1 synapses in the adult hippocampus in vivo. *Nature*, 381(6578), 163–166. <https://doi.org/10.1038/381163a0>
- Higley, M. J., & Sabatini, B. L. (2012). Calcium signaling in dendritic spines. *Cold Spring Harbor Perspectives in Biology*, 4(4). <https://doi.org/10.1101/cshperspect.a005686>
- Higuchi, M., Maas, S., Single, F. N., Hartner, J., Rozov, A., Burnashev, N., ... Seeburg, P. H. (2000). Point mutation in an AMPA receptor gene rescues lethality in mice deficient in the RNA-editing enzyme ADAR2. *Nature*, 406(6791), 78–81. <https://doi.org/10.1038/35017558>
- Holtmaat, A., & Svoboda, K. (2009, September). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*, Vol. 10, pp. 647–658. <https://doi.org/10.1038/nrn2699>
- Hoops, S., Gauges, R., Lee, C., Pahle, J., Simus, N., Singhal, M., ... Kummer, U. (2006). COPASI - A COMplex PATHway Simulator. *Bioinformatics*, 22(24), 3067–3074. <https://doi.org/10.1093/bioinformatics/btl485>
- Hosokawa, T., Mitsushima, D., Kaneko, R., & Hayashi, Y. (2015). Stoichiometry and Phosphoisotypes of Hippocampal AMPA-Type Glutamate Receptor Phosphorylation. *Neuron*, 85(1), 60–67. <https://doi.org/10.1016/j.neuron.2014.11.026>
- Hou, L., Antion, M. D., Hu, D., Spencer, C. M., Paylor, R., & Klann, E. (2006). Dynamic Translational and Proteasomal Regulation of Fragile X Mental Retardation Protein Controls mGluR-Dependent Long-Term Depression. *Neuron*, 51(4), 441–454. <https://doi.org/10.1016/j.neuron.2006.07.005>
- How to make an Item Sorter in Minecraft with Multiple delivery points in Survival Minecraft - YouTube. (n.d.). Retrieved May 27, 2020, from https://www.youtube.com/watch?v=xQn_PllhgmQ
- Hruska, M., Henderson, N., Le Marchand, S. J., Jafri, H., & Dalva, M. B. (2018). Synaptic nanomodules underlie the organization and plasticity of spine synapses. *Nature Neuroscience*, 21(5), 671–682. <https://doi.org/10.1038/s41593-018-0138-9>
- Hu, M., Li, P., Li, M., Li, W., Yao, T., Wu, J. W., ... Shi, Y. (2002). Crystal structure of a UBP-family deubiquitinating enzyme in isolation and in complex with ubiquitin aldehyde. *Cell*, 111(7), 1041–1054. [https://doi.org/10.1016/S0092-8674\(02\)01199-6](https://doi.org/10.1016/S0092-8674(02)01199-6)
- Huang, C. C., Liang, Y. C., & Hsu, K. Sen. (2001). Characterization of the Mechanism Underlying the Reversal of Long Term Potentiation by Low Frequency Stimulation at Hippocampal CA1 Synapses. *Journal of Biological Chemistry*, 276(51), 48108–48117. <https://doi.org/10.1074/jbc.M106388200>

- Huber, K. M., Kayser, M. S., & Bear, M. F. (2000). Role for rapid dendritic protein synthesis in hippocampal mGluR-dependent long-term depression. *Science*, 288(5469), 1254–1256. <https://doi.org/10.1126/science.288.5469.1254>
- Hucka, M., Finney, A., Sauro, H. M., Bolouri, H., Doyle, J. C., Kitano, H., ... Wang, J. (2003). The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models. *Bioinformatics*, 19(4), 524–531. <https://doi.org/10.1093/bioinformatics/btg015>
- Huganir, R. L., & Nicoll, R. A. (2013). AMPARs and synaptic plasticity: The last 25 years. *Neuron*, 80(3), 704–717. <https://doi.org/10.1016/j.neuron.2013.10.025>
- Huguet, G., Ey, E., & Bourgeron, T. (2013). The Genetic Landscapes of Autism Spectrum Disorders. *Annual Review of Genomics and Human Genetics*, 14(1), 191–213. <https://doi.org/10.1146/annurev-genom-091212-153431>
- Incontro, S., Díaz-Alonso, J., Iafrati, J., Vieira, M., Asensio, C. S., Sohal, V. S., ... Nicoll, R. A. (2018). The CaMKII/NMDA receptor complex controls hippocampal synaptic transmission by kinase-dependent and independent mechanisms. *Nature Communications*, 9(1), 1–21. <https://doi.org/10.1038/s41467-018-04439-7>
- Ingham, P. W., & McMahon, A. P. (2001, December 1). Hedgehog signaling in animal development: Paradigms and principles. *Genes and Development*, Vol. 15, pp. 3059–3087. <https://doi.org/10.1101/gad.938601>
- Inoki, K., Li, Y., Zhu, T., Wu, J., & Guan, K.-L. (2002). TSC2 is phosphorylated and inhibited by Akt and suppresses mTOR signalling. *Nature Cell Biology*, 4(9), 648–657. <https://doi.org/10.1038/ncb839>
- Irwin, S. A., Patel, B., Idupulapati, M., Harris, J. B., Crisostomo, R. A., Larsen, B. P., ... Greenough, W. T. (2001). Abnormal dendritic spine characteristics in the temporal and visual cortices of patients with fragile-X syndrome: A quantitative examination. *American Journal of Medical Genetics*, 98(2), 161–167. [https://doi.org/10.1002/1096-8628\(20010115\)98:2<161::AID-AJMG1025>3.0.CO;2-B](https://doi.org/10.1002/1096-8628(20010115)98:2<161::AID-AJMG1025>3.0.CO;2-B)
- Isaac, J. T. R., Nicoll, R. A., & Malenka, R. C. (1995). Evidence for silent synapses: Implications for the expression of LTP. *Neuron*, 15(2), 427–434. [https://doi.org/10.1016/0896-6273\(95\)90046-2](https://doi.org/10.1016/0896-6273(95)90046-2)
- Iwakura, Y., Nagano, T., Kawamura, M., Horikawa, H., Ibaraki, K., Takei, N., & Nawa, H. (2001). N-Methyl-D-aspartate-induced α -Amino-3-hydroxy-5-methyl-4-isoxazolepropionic Acid (AMPA) Receptor Down-regulation Involves Interaction of the Carboxyl Terminus of GluR2/3 with Pick1. *Journal of Biological Chemistry*, 276(43), 40025–40032. <https://doi.org/10.1074/jbc.M103125200>
- Jacobs, B. (2001). Regional Dendritic and Spine Variation in Human Cerebral Cortex: a Quantitative Golgi Study. *Cerebral Cortex*, 11(6), 558–571. <https://doi.org/10.1093/cercor/11.6.558>
- Jain, P., & Bhalla, U. S. (2009). Signaling Logic of Activity-Triggered Dendritic Protein Synthesis: An mTOR Gate But Not a Feedback Switch. *PLoS Comput Biol*, 5(2), 1000287. <https://doi.org/10.1371/journal.pcbi.1000287>
- Jędrzejewska-Szmek, J., & Blackwell, K. T. (2019). From membrane receptors to protein synthesis and actin cytoskeleton: Mechanisms underlying long lasting forms of synaptic plasticity. *Seminars in Cell & Developmental Biology*, 95, 120–129. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1084952118300582>
- Jennett, C., Cox, A. L., Cairns, P., Dhoparee, S., Epps, A., Tijs, T., & Walton, A. (2008). Measuring and defining the experience of immersion in games. *International Journal of Human-Computer Studies*, 66(9), 641–661. <https://doi.org/10.1016/j.ijhcs.2008.04.004>
- Jeske, L., Placzek, S., Schomburg, I., Chang, A., & Schomburg, D. (2019). BRENDA in 2019: a European ELIXIR core data resource. *Nucleic Acids Research*, 47(D1), D542–D549. <https://doi.org/10.1093/nar/gky1048>
- Jeste, S. S., & Geschwind, D. H. (2016). Clinical trials for neurodevelopmental disorders: At a therapeutic frontier. *Science Translational Medicine*, 8(321), 321fs1–321fs1. <https://doi.org/10.1126/scitranslmed.aad9874>
- Jurado, S. (2018, January 9). AMPA receptor trafficking in natural and pathological aging. *Frontiers in Molecular Neuroscience*, Vol. 10. <https://doi.org/10.3389/fnmol.2017.00446>
- Kahu, E., Nelson, K., & Picton, C. (2017). Student interest as a key driver of engagement for first year students. *Student Success*, 8(2), 55–66. <https://doi.org/10.5204/ssj.v8i2.379>
- Kang, M., Day, C. A., Kenworthy, A. K., & DiBenedetto, E. (2012). Simplified Equation to Extract Diffusion Coefficients from Confocal FRAP Data. *Traffic*, 13(12), 1589–1600. <https://doi.org/10.1111/tra.12008>

- Karr, J. R., Sanghvi, J. C., MacKlin, D. N., Gutschow, M. V., Jacobs, J. M., Bolival, B., ... Covert, M. W. (2012). A whole-cell computational model predicts phenotype from genotype. *Cell*, 150(2), 389–401. <https://doi.org/10.1016/j.cell.2012.05.044>
- Karr, J. R., Takahashi, K., & Funahashi, A. (2015). The principles of whole-cell modeling. *Current Opinion in Microbiology*, 27, 18–24. <https://doi.org/10.1016/j.mib.2015.06.004>
- Kasthuri, N., Hayworth, K. J., Berger, D. R., Schalek, R. L., Conchello, J. A., Knowles-Barley, S., ... Lichtman, J. W. (2015). Saturated Reconstruction of a Volume of Neocortex. *Cell*, 162(3), 648–661. <https://doi.org/10.1016/j.cell.2015.06.054>
- Kastning, K., Kukhtina, V., Kittler, J. T., Chen, G., Pechstein, A., Enders, S., ... Haucke, V. (2007). Molecular determinants for the interaction between AMPA receptors and the clathrin adaptor complex AP-2. *Proceedings of the National Academy of Sciences of the United States of America*, 104(8), 2991–2996. <https://doi.org/10.1073/pnas.0611170104>
- Kato, A. S., Gill, M. B., Ho, M. T., Yu, H., Tu, Y., Siuda, E. R., ... Bredt, D. S. (2010). Hippocampal AMPA Receptor Gating Controlled by Both TARP and Cornichon Proteins. *Neuron*, 68(6), 1082–1096. <https://doi.org/10.1016/j.neuron.2010.11.026>
- Kato, A. S., Zhou, W., Milstein, A. D., Knierman, M. D., Siuda, E. R., Dotzla, J. E., ... Bredt, D. S. (2007). New transmembrane AMPA receptor regulatory protein isoform, γ -7, differentially regulates AMPA receptors. *Journal of Neuroscience*, 27(18), 4969–4977. <https://doi.org/10.1523/JNEUROSCI.5561-06.2007>
- Keating, S. M., Waltemath, D., König, M., Zhang, F., Dräger, A., Chaouiya, C., ... Hucka, M. (2020). SBML Level 3: an extensible format for the exchange and reuse of biological models. *Molecular Systems Biology*, 16(8), e9110. <https://doi.org/10.15252/MSB.20199110>
- Kemp, N., & Bashir, Z. I. (2001). Long-term depression: a cascade of induction and expression mechanisms. *Progress in Neurobiology*, 65(4), 339–365. [https://doi.org/10.1016/S0301-0082\(01\)00013-2](https://doi.org/10.1016/S0301-0082(01)00013-2)
- Kemp, N., McQueen, J., Faulkes, S., & Bashir, Z. I. (2000). Different forms of LTD in the CA1 region of the hippocampus: role of age and stimulus protocol. *European Journal of Neuroscience*, 12(1), 360–366. <https://doi.org/10.1046/j.1460-9568.2000.00903.x>
- Kerr, R. A., Bartol, T. M., Kaminsky, B., Dittrich, M., Chang, J.-C. J., Baden, S. B., ... Stiles, J. R. (2008). Fast Monte Carlo Simulation Methods for Biological Reaction-Diffusion Systems in Solution and on Surfaces. *SIAM Journal on Scientific Computing*, 30(6), 3126–3149. <https://doi.org/10.1137/070692017>
- Kessels, H. W., Kopec, C. D., Klein, M. E., & Malinow, R. (2009). Roles of stargazin and phosphorylation in the control of AMPA receptor subcellular distribution. *Nature Neuroscience*, 12(7), 888–896. <https://doi.org/10.1038/nn.2340>
- Kim, E., Niethammer, M., Rothschild, A., Nung Jan, Y., & Sheng, M. (1995). Clustering of Shaker-type K⁺ channels by interaction with a family of membrane-associated guanylate kinases. *Nature*, 378(6552), 85–88. <https://doi.org/10.1038/378085a0>
- Kim, H. H., Lee, K. H., Lee, D., Han, Y. E., Lee, S. H., Sohn, J. W., & Ho, W. K. (2015). Costimulation of AMPA and metabotropic glutamate receptors underlies phospholipase C activation by glutamate in hippocampus. *Journal of Neuroscience*, 35(16), 6401–6412. <https://doi.org/10.1523/JNEUROSCI.4208-14.2015>
- Kim, H. H., Park, J. M., Lee, S. H., & Ho, W. K. (2019). Association of mGluR-dependent LTD of excitatory synapses with endocannabinoid-dependent LTD of inhibitory synapses leads to EPSP to spike potentiation in CA1 pyramidal neurons. *Journal of Neuroscience*, 39(2), 224–237. <https://doi.org/10.1523/JNEUROSCI.2935-17.2018>
- Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist*, 41(2), 75–86. https://doi.org/10.1207/s15326985ep4102_1
- Kitano, H., Funahashi, A., Matsuoka, Y., & Oda, K. (2005). Using process diagrams for the graphical representation of biological networks. *Nature Biotechnology*, 23(8), 961–966. <https://doi.org/10.1038/nbt1111>
- Köhler, A., Krivine, J., & Vidmar, J. (2014). A Rule-Based Model of Base Excision Repair. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Vol. 8859, pp. 173–195). https://doi.org/10.1007/978-3-319-12982-2_13

- Kornau, H. C., Schenker, L. T., Kennedy, M. B., & Seeburg, P. H. (1995). Domain interaction between NMDA receptor subunits and the postsynaptic density protein PSD-95. *Science*, 269(5231), 1737–1740. <https://doi.org/10.1126/science.7569905>
- Koshland, D. E., Némethy, G., & Filmer, D. (1966). Comparison of Experimental Binding Data and Theoretical Models in Proteins Containing Subunits *. *Biochemistry*, 5(1), 365–385. <https://doi.org/10.1021/bi00865a047>
- Kotaleski, J. H., & Blackwell, K. T. (2010). Modelling the molecular mechanisms of synaptic plasticity using systems biology approaches. *Nature Reviews Neuroscience*, 11(4), 239–251. <https://doi.org/10.1038/nrn2807>
- Kristensen, A. S., Jenkins, M. A., Banke, T. G., Schousboe, A., Makino, Y., Johnson, R. C., ... Traynelis, S. F. (2011). Mechanism of Ca²⁺/calmodulin-dependent kinase II regulation of AMPA receptor gating. *Nature Neuroscience*, 14(6), 727–735. <https://doi.org/10.1038/nn.2804>
- Kwak, Y. D., Wang, B., Li, J. J., Wang, R., Deng, Q., Diao, S., ... Liao, F. F. (2012). Upregulation of the E3 ligase NEDD4-1 by oxidative stress degrades IGF-1 receptor protein in neurodegeneration. *Journal of Neuroscience*, 32(32), 10971–10981. <https://doi.org/10.1523/JNEUROSCI.1836-12.2012>
- Lai, K., Robertson, M. J., & Schaffer, D. V. (2004). The Sonic Hedgehog Signaling System as a Bistable Genetic Switch. *Biophysical Journal*, 86(5), 2748–2757. [https://doi.org/10.1016/S0006-3495\(04\)74328-3](https://doi.org/10.1016/S0006-3495(04)74328-3)
- Lameras, P., Arnab, S., Dunwell, I., Stewart, C., Clarke, S., & Petridis, P. (2017). Essential features of serious games design in higher education: Linking learning attributes to game mechanics. *British Journal of Educational Technology*, 48(4), 972–994. <https://doi.org/10.1111/bjet.12467>
- Lee, B. R., Ma, Y.-Y., Huang, Y. H., Wang, X., Otaka, M., Ishikawa, M., ... Dong, Y. (2013). Maturation of silent synapses in amygdala-accumbens projection contributes to incubation of cocaine craving. *Nature Neuroscience*, 16(11), 1644–1651. <https://doi.org/10.1038/nn.3533>
- Lee, H.-K., & Huganir, R. L. (2008). AMPA Receptor Regulation and the Reversal of Synaptic Plasticity – LTP, LTD, Depotentiation, and Dedepression. In *Learning and Memory: A Comprehensive Reference* (pp. 633–648). <https://doi.org/10.1016/B978-012370509-9.00035-8>
- Lee, H. K., & Kirkwood, A. (2011, July 1). AMPA receptor regulation during synaptic plasticity in hippocampus and neocortex. *Seminars in Cell and Developmental Biology*, Vol. 22, pp. 514–520. <https://doi.org/10.1016/j.semcdb.2011.06.007>
- Lee, H. K., Takamiya, K., He, K., Song, L., & Huganir, R. L. (2010). Specific roles of AMPA receptor subunit GluR1 (GluA1) phosphorylation sites in regulating synaptic plasticity in the CA1 region of hippocampus. *Journal of Neurophysiology*, 103(1), 479–489. <https://doi.org/10.1152/jn.00835.2009>
- Lee, Hey-Kyoung, & Kirkwood, A. (2011). AMPA receptor regulation during synaptic plasticity in hippocampus and neocortex. *Seminars in Cell & Developmental Biology*, 22(5), 514–520. <https://doi.org/10.1016/j.semcdb.2011.06.007>
- Lee, Hey-Kyoung, & Kirkwood, A. (2019). Mechanisms of Homeostatic Synaptic Plasticity in vivo. *Frontiers in Cellular Neuroscience*, 13. <https://doi.org/10.3389/fncel.2019.00520>
- Lee, S. H., Jin, C., Cai, E., Ge, P., Ishitsuka, Y., Teng, K. W., ... Selvin, P. R. (2017). Super-resolution imaging of synaptic and extra-synaptic AMPA receptors with Different-Sized fluorescent probes. *ELife*, 6. <https://doi.org/10.7554/eLife.27744>
- Lekalakala-Mokgele, E. (2010). Facilitation in problem-based learning: Experiencing the locus of control. *Nurse Education Today*, 30(7), 638–642. <https://doi.org/10.1016/j.nedt.2009.12.017>
- Lengler, J., Jug, F., & Steger, A. (2013). Reliable neuronal systems: The importance of heterogeneity. *PLoS ONE*, 8(12), 80694. <https://doi.org/10.1371/journal.pone.0080694>
- Lesperance, L. S., Yang, Y. M., & Wang, L. Y. (2020). Delayed expression of activity-dependent gating switch in synaptic AMPARs at a central synapse. *Molecular Brain*, 13(1), 6. <https://doi.org/10.1186/s13041-019-0536-2>
- Li, L., Stefan, M. I., & Le Novère, N. (2012). Calcium Input Frequency, Duration and Amplitude Differentially Modulate the Relative Activation of Calcineurin and CaMKII. *PLoS ONE*, 7(9). <https://doi.org/10.1371/journal.pone.0043810>
- Li, M., Makkinje, A., & Damuni, Z. (1996). The Myeloid Leukemia-associated Protein SET Is a Potent Inhibitor of

- Protein Phosphatase 2A. *Journal of Biological Chemistry*, 271(19), 11059–11062. <https://doi.org/10.1074/jbc.271.19.11059>
- Li, W., & Ye, Y. (2008). Polyubiquitin chains: Functions, structures, and mechanisms. *Cellular and Molecular Life Sciences*, 65(15), 2397–2406. <https://doi.org/10.1007/s00018-008-8090-6>
- Li, X., Mao, Z., Wu, M., & Xia, J. (2013). Rescuing infertility of pick1 knockout mice by generating testis-specific transgenic mice via testicular infection. *Scientific Reports*, 3(1), 1–8. <https://doi.org/10.1038/srep02842>
- Lin, A., Hou, Q., Jarzylo, L., Amato, S., Gilbert, J., Shang, F., & Man, H.-Y. (2011). Nedd4-mediated AMPA receptor ubiquitination regulates receptor turnover and trafficking. *Journal of Neurochemistry*, 119(1), 27–39. <https://doi.org/10.1111/j.1471-4159.2011.07221.x>
- Lin, D.-T., Makino, Y., Sharma, K., Hayashi, T., Neve, R., Takamiya, K., & Huganir, R. L. (2009). Regulation of AMPA receptor extrasynaptic insertion by 4.1N, phosphorylation and palmitoylation. *Nature Neuroscience*, 12(7), 879–887. <https://doi.org/10.1038/nn.2351>
- Lipp, P., & Reither, G. (2011). Protein kinase C: The “Masters” of calcium and lipid. *Cold Spring Harbor Perspectives in Biology*, 3(7), 1–17. <https://doi.org/10.1101/cshperspect.a004556>
- Lisman, J. E. (1985). A mechanism for memory storage insensitive to molecular turnover: a bistable autophosphorylating kinase. *Proceedings of the National Academy of Sciences*, 82(9), 3055–3057. <https://doi.org/10.1073/pnas.82.9.3055>
- Lisman, J., & Raghavachari, S. (2015). Biochemical principles underlying the stable maintenance of LTP by the CaMKII/NMDAR complex. *Brain Research*, 1621, 51–61. <https://doi.org/10.1016/j.brainres.2014.12.010>
- Liu, S., Rao, Y., & Daw, N. W. (2003). Roles of protein kinase A and protein kinase G in synaptic plasticity in the visual cortex. *Cerebral Cortex*, 13(8), 864–869. <https://doi.org/10.1093/cercor/13.8.864>
- Lodge, D. (2009). The history of the pharmacology and cloning of ionotropic glutamate receptors and the development of idiosyncratic nomenclature. *Neuropharmacology*, 59(1), 6–21. <https://doi.org/10.1016/j.neuropharm.2008.08.006>
- Long, J., Wei, Z., Feng, W., Yu, C., Zhao, Y., & Zhang, M. (2008). Supramodular Nature of GRIP1 Revealed by the Structure of Its PDZ12 Tandem in Complex with the Carboxyl Tail of Fras1. *Journal of Molecular Biology*, 375(5), 1457–1468. <https://doi.org/10.1016/j.jmb.2007.11.088>
- Lorenzo-Alvarez, R., Pavia-Molina, J., & Sendra-Portero, F. (2018). Exploring the Potential of Undergraduate Radiology Education in the Virtual World Second Life with First-cycle and Second-cycle Medical Students. *Academic Radiology*, 25(8), 1087–1096. <https://doi.org/10.1016/j.acra.2018.02.026>
- Lorenzo-Alvarez, R., Rudolphi-Solero, T., Ruiz-Gomez, M. J., & Sendra-Portero, F. (2019). Medical Student Education for Abdominal Radiographs in a 3D Virtual Classroom Versus Traditional Classroom: A Randomized Controlled Trial. *American Journal of Roentgenology*, 213(3), 644–650. <https://doi.org/10.2214/AJR.19.21131>
- Lorenzo-Alvarez, R., Ruiz-Gomez, M. J., & Sendra-Portero, F. (2019). Medical Students' and Family Physicians' Attitudes and Perceptions Toward Radiology Learning in the Second Life Virtual World. *American Journal of Roentgenology*, 212(6), 1295–1302. <https://doi.org/10.2214/AJR.18.20381>
- Lorenzo-Alvarez, R., Rudolphi-Solero, T., Ruiz-Gomez, M. J., & Sendra-Portero, F. (2020). Game-Based Learning in Virtual Worlds: A Multiuser Online Game for Medical Undergraduate Radiology Education within Second Life. *Anatomical Sciences Education*, 13(5), 602–617. <https://doi.org/10.1002/ase.1927>
- Louros, S. R., & Osterweil, E. K. (2016). Perturbed proteostasis in autism spectrum disorders. *Journal of Neurochemistry*, 139(6), 1081–1092. <https://doi.org/10.1111/jnc.13723>
- Löwel, S., & Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255(5041), 209–212. <https://doi.org/10.1126/science.1372754>
- Lu, J., Helton, T. D., Blanpied, T. A., Rácz, B., Newpher, T. M., Weinberg, R. J., & Ehlers, M. D. (2007). Postsynaptic Positioning of Endocytic Zones and AMPA Receptor Cycling by Physical Coupling of Dynamin-3 to Homer. *Neuron*, 55(6), 874–889. <https://doi.org/10.1016/j.neuron.2007.06.041>
- Lu, W. H., Yeh, N. H., & Huang, Y. S. (2017). CPEB2 Activates GRASP1 mRNA Translation and Promotes AMPA Receptor Surface Expression, Long-Term Potentiation, and Memory. *Cell Reports*, 21(7), 1783–1794. <https://doi.org/10.1016/j.celrep.2017.10.073>

- Lu, W., Shi, Y., Jackson, A. C., Bjorgan, K., During, M. J., Sprengel, R., ... Nicoll, R. A. (2009). Subunit Composition of Synaptic AMPA Receptors Revealed by a Single-Cell Genetic Approach. *Neuron*, 62(2), 254–268. <https://doi.org/10.1016/j.neuron.2009.02.027>
- Lu, W. Y., Man, H. Y., Ju, W., Trimble, W. S., MacDonald, J. F., & Wang, Y. T. (2001). Activation of synaptic NMDA receptors induces membrane insertion of new AMPA receptors and LTP in cultured hippocampal neurons. *Neuron*, 29(1), 243–254. [https://doi.org/10.1016/S0896-6273\(01\)00194-5](https://doi.org/10.1016/S0896-6273(01)00194-5)
- Lu, W., & Ziff, E. B. (2005). PICK1 interacts with ABP/GRIP to regulate AMPA receptor trafficking. *Neuron*, 47(3), 407–421. <https://doi.org/10.1016/j.neuron.2005.07.006>
- Lu, Y., Allen, M., Halt, A. R., Weisenhaus, M., Dallapiazza, R. F., Hall, D. D., ... Hell, J. W. (2007). Age-dependent requirement of AKAP150-anchored PKA and GluR2-lacking AMPA receptors in LTP. *The EMBO Journal*, 26(23), 4879–4890. <https://doi.org/10.1038/sj.emboj.7601884>
- Lucero, A., Karapanos, E., Arrasvuori, J., & Korhonen, H. (2014, May 1). Playful or Gameful? *Interactions*, 21(3), 34–39. <https://doi.org/10.1145/2590973>
- Luchkina, N. V., Coleman, S. K., Huupponen, J., Cai, C., Kivistö, A., Taira, T., ... Lauri, S. E. (2017). Molecular mechanisms controlling synaptic recruitment of GluA4 subunit-containing AMPA-receptors critical for functional maturation of CA1 glutamatergic synapses. *Neuropharmacology*, 112, 46–56. <https://doi.org/10.1016/j.neuropharm.2016.04.049>
- Luria, A. R. (1976). *Cognitive Development: It's Cultural and Social Foundations*. Harvard University Press.
- Lüscher, C., & Huber, K. M. (2010). Group 1 mGluR-Dependent Synaptic Long-Term Depression: Mechanisms and Implications for Circuitry and Disease. *Neuron*, 65(4), 445–459. <https://doi.org/10.1016/j.neuron.2010.01.016>
- Lussier, M. P., Herring, B. E., Nasu-Nishimura, Y., Neutzner, A., Karbowski, M., Youle, R. J., ... Roche, K. W. (2012). Ubiquitin ligase RNF167 regulates AMPA receptor-mediated synaptic transmission. *Proceedings of the National Academy of Sciences of the United States of America*, 109(47), 19426–19431. <https://doi.org/10.1073/pnas.1217477109>
- Lussier, M. P., Nasu-Nishimura, Y., & Roche, K. W. (2011). Activity-dependent ubiquitination of the AMPA receptor subunit GluA2. *Journal of Neuroscience*, 31(8), 3077–3081. <https://doi.org/10.1523/JNEUROSCI.5944-10.2011>
- Lyall, K., Croen, L., Daniels, J., Fallin, M. D., Ladd-Acosta, C., Lee, B. K., ... Newschaffer, C. (2017). The Changing Epidemiology of Autism Spectrum Disorders. *Annual Review of Public Health*, 38, 81–102. <https://doi.org/10.1146/annurev-publhealth-031816-044318>
- Lyubomirsky, S., King, L., & Diener, E. (2005). *The Benefits of Frequent Positive Affect: Does Happiness Lead to Success?* <https://doi.org/10.1037/0033-2909.131.6.803>
- Magnani, L. (2001). Abduction, Reason and Science. In *Abduction, Reason and Science*. <https://doi.org/10.1007/978-1-4419-8562-0>
- Magnussen, R., & Elming, A. (2015, October 8). *Cities at Play: Children's redesign of deprived neighbourhoods in Minecraft* (pp. 331–337). pp. 331–337. Retrieved from <https://vbn.aau.dk/en/publications/cities-at-play-childrens-redesign-of-deprived-neighbourhoods-in-m>
- Mahajan, G., & Nadkarni, S. (2019). Intracellular calcium stores mediate metaplasticity at hippocampal dendritic spines. *The Journal of Physiology*, 597(13), 3473–3502. <https://doi.org/10.1113/JP277726>
- Mäki-Marttunen, T., Iannella, N., Edwards, A. G., Einevoll, G. T., & Blackwell, K. T. (2020). A unified computational model for cortical post-synaptic plasticity. *eLife*, 9, 1–42. <https://doi.org/10.7554/eLife.55714>
- Making the Most Secure Door in Minecraft! - YouTube. (n.d.). Retrieved May 27, 2020, from <https://www.youtube.com/watch?v=xz0xGDrXUBQ>
- Maletic-Savatic, M., Malinow, R., & Svoboda, K. (1999). Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity. *Science*, 283(5409), 1923–1927. <https://doi.org/10.1126/science.283.5409.1923>
- Malik-Sheriff, R. S., Glont, M., Nguyen, T. V. N., Tiwari, K., Roberts, M. G., Xavier, A., ... Hermjakob, H. (2020). BioModels-15 years of sharing computational models in life science. *Nucleic Acids Research*, 48(D1), D407–D415. <https://doi.org/10.1093/nar/gkz1055>

- Malinow, R., Mainen, Z. F., & Hayashi, Y. (2000). LTP mechanisms: From silence to four-lane traffic. *Current Opinion in Neurobiology*, 10(3), 352–357. [https://doi.org/10.1016/S0959-4388\(00\)00099-4](https://doi.org/10.1016/S0959-4388(00)00099-4)
- Malone, T. W. (1980a). *What makes things fun to learn? A study of intrinsically motivating computer games*. ProQuest Information & Learning.
- Malone, T. W. (1980b). What makes things fun to learn? heuristics for designing instructional computer games. *Proceedings of the 3rd ACM SIGSMALL Symposium and the First SIGPC Symposium on Small Systems - SIGSMALL '80*, 162–169. <https://doi.org/10.1145/800088.802839>
- Maloney, J., Resnick, M., Rusk, N., Silverman, B., & Eastmond, E. (2010). The scratch program-ming language and environment. *ACM Trans. Comput. Educ*, 10. <https://doi.org/10.1145/1868358.1868363>
- Man, H. Y., Sekine-Aizawa, Y., & Huganir, R. L. (2007). Regulation of α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptor trafficking through PKA phosphorylation of the Glu receptor 1 subunit. *Proceedings of the National Academy of Sciences of the United States of America*, 104(9), 3579–3584. <https://doi.org/10.1073/pnas.0611698104>
- Marín-Díaz, V., Sampedro-Requena, B. E., & López-Pérez, M. (2020). Students' perceptions about the use the videogames in secondary education. *Education and Information Technologies*, 25(4), 3251–3273. <https://doi.org/10.1007/s10639-020-10122-6>
- Marrink, S. J., Corradi, V., Souza, P. C. T., Ingólfsson, H. I., Tieleman, D. P., & Sansom, M. S. P. (2019). Computational Modeling of Realistic Cell Membranes. *Chemical Reviews*, 119(9), 6184–6226. <https://doi.org/10.1021/acs.chemrev.8b00460>
- Martin, S. J., Grimwood, P. D., & Morris, R. G. M. (2000). Synaptic Plasticity and Memory: An Evaluation of the Hypothesis. *Annual Review of Neuroscience*, 23(1), 649–711. <https://doi.org/10.1146/annurev.neuro.23.1.649>
- Masdeu, C., Bernard, V., Faure, H., Traffort, E., & Ruat, M. (2007). Distribution of Smoothed at hippocampal mossy fiber synapses. *NeuroReport*, 18(4), 395–399. <https://doi.org/10.1097/WNR.0b013e32801421ce>
- Mathieson, T., Franken, H., Kosinski, J., Kurzawa, N., Zinn, N., Sweetman, G., ... Savitski, M. M. (2018). Systematic analysis of protein turnover in primary cells. *Nature Communications*, 9(1), 1–10. <https://doi.org/10.1038/s41467-018-03106-1>
- Matsuda, S., Launey, T., Mikawa, S., & Hirai, H. (2000). Disruption of AMPA receptor GluR2 clusters following long-term depression induction in cerebellar Purkinje neurons. *EMBO Journal*, 19(12), 2765–2774. <https://doi.org/10.1093/emboj/19.12.2765>
- Matsuda, S., Mikawa, S., & Hirai, H. (2002). Phosphorylation of Serine-880 in GluR2 by Protein Kinase C Prevents Its C Terminus from Binding with Glutamate Receptor-Interacting Protein. *Journal of Neurochemistry*, 73(4), 1765–1768. <https://doi.org/10.1046/j.1471-4159.1999.731765.x>
- Matsuzaki, M., Honkura, N., Ellis-Davies, G. C. R., & Kasai, H. (2004). Structural basis of long-term potentiation in single dendritic spines. *Nature*, 429(6993), 761–766. <https://doi.org/10.1038/nature02617>
- Mattison, H. A., Bagal, A. A., Mohammadi, M., Pulimood, N. S., Reich, C. G., Alger, B. E., ... Thompson, S. M. (2014). Evidence of calcium-permeable AMPA receptors in dendritic spines of CA1 pyramidal neurons. *Journal of Neurophysiology*, 112(2), 263–275. <https://doi.org/10.1152/jn.00578.2013>
- Maurin, T., Zongaro, S., & Bardoni, B. (2014). Fragile X Syndrome: From molecular pathology to therapy. *Neuroscience and Biobehavioral Reviews*.
- Mayer, M. L. (2011). Structure and mechanism of glutamate receptor ion channel assembly, activation and modulation. *Current Opinion in Neurobiology*, 21(2), 283–290. <https://doi.org/10.1016/J.CONB.2011.02.001>
- Mayford, M., Wang, J., Kandel, E. R., & O'Dell, T. J. (1995). CaMKII regulates the frequency-response function of hippocampal synapses for the production of both LTD and LTP. *Cell*, 81(6), 891–904. [https://doi.org/10.1016/0092-8674\(95\)90009-8](https://doi.org/10.1016/0092-8674(95)90009-8)
- McDougal, R. A., Morse, T. M., Carnevale, T., Marengo, L., Wang, R., Migliore, M., ... Hines, M. L. (2017). Twenty years of ModelDB and beyond: building essential modeling tools for the future of neuroscience. *Journal of Computational Neuroscience*, 42(1), 1–10. <https://doi.org/10.1007/s10827-016-0623-7>
- Mcdowell, C., Werner, L., Bullock, H. E., & Fernald, J. (2003). The Impact of Pair Programming on Student

- Performance, Perception and Persistence. *ICSE '03: Proceedings of the 25th International Conference on Software Engineering*. <https://doi.org/10.5555/776816.776899>
- McDowell, W. C., Wilson, R. C., & Kile, C. O. (2016). An examination of retail website design and conversion rate. *Journal of Business Research*, 69(11), 4837–4842. <https://doi.org/10.1016/j.jbusres.2016.04.040>
- Mejias, R., Chiu, S. L., Han, M., Rose, R., Gil-Infante, A., Zhao, Y., ... Wang, T. (2019). Purkinje cell-specific Grip1/2 knockout mice show increased repetitive self-grooming and enhanced mGluR5 signaling in cerebellum. *Neurobiology of Disease*, 132, 104602. <https://doi.org/10.1016/j.nbd.2019.104602>
- Meng, Y., Zhang, Y., & Jia, Z. (2003). Synaptic transmission and plasticity in the absence of AMPA glutamate receptor GluR2 and GluR3. *Neuron*, 39(1), 163–176. [https://doi.org/10.1016/S0896-6273\(03\)00368-4](https://doi.org/10.1016/S0896-6273(03)00368-4)
- Merenstein, S. A., Sobesky, W. E., Taylor, A. K., Riddle, J. E., Tran, H. X., & Hagerman, R. J. (1996). Molecular-clinical correlations in males with an expanded FMR1 mutation. *American Journal of Medical Genetics*, 64(2), 388–394. [https://doi.org/10.1002/\(SICI\)1096-8628\(19960809\)64:2<388::AID-AJMG31>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1096-8628(19960809)64:2<388::AID-AJMG31>3.0.CO;2-9)
- Mi, H., Schreiber, F., Moodie, S., Czauderna, T., Demir, E., Haw, R., ... Villéger, A. (2015). Systems Biology Graphical Notation: Activity Flow language Level 1 Version 1.2. *Journal of Integrative Bioinformatics*, 12(2), 340–381. <https://doi.org/10.1515/jib-2015-265>
- Michel, J. J. C., & Scott, J. D. (2002). AKAP mediated signal transduction. *Annual Review of Pharmacology and Toxicology*, 42(1), 235–257. <https://doi.org/10.1146/annurev.pharmtox.42.083101.135801>
- Milstein, A., Li, Y., Bittner, K., Grienberger, C., Soltesz, I., Magee, J., & Romani, S. (2020). Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity. *BioRxiv*, 2020.02.04.934182. <https://doi.org/10.1101/2020.02.04.934182>
- Minecraft: Education Edition: Teacher Academy - Microsoft Educator Center. (n.d.). Retrieved February 19, 2021, from <https://education.microsoft.com/en-us/learningPath/3eede2ae>
- Mizuseki, K., & Buzsáki, G. (2013). Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell Reports*, 4(5), 1010–1021. <https://doi.org/10.1016/j.celrep.2013.07.039>
- Moffat, S. D., & Resnick, S. M. (2002). Effects of age on virtual environment place navigation and allocentric cognitive mapping. *Behavioral Neuroscience*, 115(5), 851–859. <https://doi.org/10.1037/0735-7044.116.5.851>
- Mogilner, A., Wollman, R., & Marshall, W. F. (2006). Quantitative Modeling in Cell Biology: What Is It Good for? *Developmental Cell*, 11(3), 279–287. <https://doi.org/10.1016/j.devcel.2006.08.004>
- Monod, J., Wyman, J., & Changeux, J. P. (1965). On the nature of allosteric transitions: A plausible model. *Journal of Molecular Biology*, 12(1), 88–118. [https://doi.org/10.1016/S0022-2836\(65\)80285-6](https://doi.org/10.1016/S0022-2836(65)80285-6)
- Moretto, E., & Passafaro, M. (2018). Recent findings on AMPA receptor recycling. *Frontiers in Cellular Neuroscience*, 12, 286. <https://doi.org/10.3389/fncel.2018.00286>
- Morise, J., Suzuki, K. G. N., Kitagawa, A., Wakazono, Y., Takamiya, K., Tsunoyama, T. A., ... Oka, S. (2019). AMPA receptors in the synapse turnover by monomer diffusion. *Nature Communications*, 10(1), 1–18. <https://doi.org/10.1038/s41467-019-13229-8>
- Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12(2), 239–260. [https://doi.org/10.1016/0023-9690\(81\)90020-5](https://doi.org/10.1016/0023-9690(81)90020-5)
- Moser, E. I., & Moser, M. B. (1999). Is learning blocked by saturation of synaptic weights in the hippocampus? *Neuroscience and Biobehavioral Reviews*, 23(5), 661–672. [https://doi.org/10.1016/S0149-7634\(98\)00060-8](https://doi.org/10.1016/S0149-7634(98)00060-8)
- Mukherjee, S., & Manahan-Vaughan, D. (2013). Role of metabotropic glutamate receptors in persistent forms of hippocampal plasticity and learning. *Neuropharmacology*, 66, 65–81. <https://doi.org/10.1016/J.NEUROPHARM.2012.06.005>
- Mullard, A. (2015). Fragile X disappointments upset autism ambitions. In *Nature Reviews Drug Discovery* 2015 14:3 (Vol. 14). <https://doi.org/10.1038/nrd4555>
- Murakoshi, H., Wang, H., & Yasuda, R. (2011). Local, persistent activation of Rho GTPases during plasticity of single dendritic spines. *Nature*, 472(7341), 100–104. <https://doi.org/10.1038/nature09823>

- Murcia-López, M., & Steed, A. (2016). The Effect of Environmental Features, Self-Avatar, and Immersion on Object Location Memory in Virtual Environments. *Frontiers in ICT*, 3(NOV), 3. <https://doi.org/10.3389/fict.2016.00024>
- Musumeci, S. A., Bosco, P., Calabrese, G., Bakker, C., De Sarro, G. B., Elia, M., ... Oostra, B. A. (2000). Audiogenic seizures susceptibility in transgenic mice with fragile X syndrome. *Epilepsia*, 41(1), 19–23. <https://doi.org/10.1111/j.1528-1157.2000.tb01499.x>
- Musumeci, S. A., Hagerman, R. J., Ferri, R., Bosco, P., Dalla Bernardina, B., Tassinari, C. A., ... Elia, M. (1999). Epilepsy and EEG findings in males with fragile X syndrome. *Epilepsia*, 40(8), 1092–1099. <https://doi.org/10.1111/j.1528-1157.1999.tb00824.x>
- Nabavi, S., Kessels, H. W., Alfonso, S., Aow, J., Fox, R., & Malinow, R. (2013). Metabotropic NMDA receptor function is required for NMDA receptor-dependent long-term depression. *Proceedings of the National Academy of Sciences of the United States of America*, 110(10), 4027–4032. <https://doi.org/10.1073/pnas.1219454110>
- Nagappan, N., Williams, L., Ferzli, M., Wiebe, E., Yang, K., Miller, C., & Balik, S. (2003). *Improving the CS1 Experience with Pair Programming*.
- Nair, D., Hosy, E., Petersen, J. D., Constals, A., Giannone, G., Choquet, D., & Sibarita, J. B. (2013). Super-resolution imaging reveals that AMPA receptors inside synapses are dynamically organized in nanodomains regulated by PSD95. *Journal of Neuroscience*, 33(32), 13204–13224. <https://doi.org/10.1523/JNEUROSCI.2381-12.2013>
- Nalavadi, V. C., Muddashetty, R. S., Gross, C., & Bassell, G. J. (2012). Dephosphorylation-induced ubiquitination and degradation of FMRP in dendrites: A role in immediate early mGluR-stimulated translation. *Journal of Neuroscience*, 32(8), 2582–2587. <https://doi.org/10.1523/JNEUROSCI.5057-11.2012>
- Napoli, I., Mercaldo, V., Boyle, P. P., Eleuteri, B., Zalfa, F., De Rubeis, S., ... Bagni, C. (2008). The Fragile X Syndrome Protein Represses Activity-Dependent Translation through CYFIP1, a New 4E-BP. *Cell*, 134(6), 1042–1054. <https://doi.org/10.1016/j.cell.2008.07.031>
- Narayanan, U., Nalavadi, V., Nakamoto, M., Thomas, G., Ceman, S., Bassell, G. J., & Warren, S. T. (2008). S6K1 phosphorylates and regulates fragile X mental retardation protein (FMRP) with the neuronal protein synthesis-dependent mammalian target of rapamycin (mTOR) signaling cascade. *Journal of Biological Chemistry*, 283(27), 18478–18482. <https://doi.org/10.1074/jbc.C800055200>
- Nebel, S., Schneider, S., Schledjewski, J., & Rey, G. D. (2017). Goal-Setting in Educational Video Games. *Simulation & Gaming*, 48(1), 98–130. <https://doi.org/10.1177/1046878116680869>
- Newman, D. L., Stefkovich, M., Clasen, C., Franzen, M. A., & Wright, L. K. (2018). Physical models can provide superior learning opportunities beyond the benefits of active engagements. *Biochemistry and Molecular Biology Education*, 46(5), 435–444. <https://doi.org/10.1002/bmb.21159>
- Nicholson, D. J. (2012). The concept of mechanism in biology. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(1), 152–163. <https://doi.org/10.1016/j.shpsc.2011.05.014>
- Nicholson, D. J. (2019). Is the cell really a machine? *Journal of Theoretical Biology*, 477, 108–126. <https://doi.org/10.1016/j.jtbi.2019.06.002>
- Nicoletti, F., Bruno, V., Ngomba, R. T., Gradini, R., & Battaglia, G. (2015). Metabotropic glutamate receptors as drug targets: what's new? *Current Opinion in Pharmacology*, 20, 89–94. <https://doi.org/10.1016/j.coph.2014.12.002>
- Nicoll, R. A., & Roche, K. W. (2013). Long-term potentiation: Peeling the onion. *Neuropharmacology*, 74, 18–22. <https://doi.org/10.1016/j.neuropharm.2013.02.010>
- Niethammer, M., Kim, E., & Sheng, M. (1996). Interaction between the C terminus of NMDA receptor subunits and multiple members of the PSD-95 family of membrane-associated guanylate kinases. *Journal of Neuroscience*, 16(7), 2157–2163. <https://doi.org/10.1523/jneurosci.16-07-02157.1996>
- Nonaka, M., Fitzpatrick, R., Lapira, J., Wheeler, D., Spooner, P. A., Corcoles-Parada, M., ... Morris, R. G. M. (2017). Everyday memory: towards a translationally effective method of modelling the encoding, forgetting and enhancement of memory. *European Journal of Neuroscience*, 46(4), 1937–1953. <https://doi.org/10.1111/ejn.13637>

- Novère, N. Le, Finney, A., Hucka, M., Bhalla, U. S., Campagne, F., Collado-Vides, J., ... Wanner, B. L. (2005). Minimum information requested in the annotation of biochemical models (MIRIAM). *Nature Biotechnology*, 23(12), 1509–1515. <https://doi.org/10.1038/nbt1156>
- Nylund, A., & Landfors, O. (2015). *Frustration and its effect on immersion in games*. Umea University.
- O'Brien, R. J., Desheng, X., Petralia, R. S., Steward, O., Huguinir, R. L., & Worley, P. (1999). Synaptic clustering of AMPA receptors by the extracellular immediate-early gene product Narp. *Neuron*, 23(2), 309–323. [https://doi.org/10.1016/S0896-6273\(00\)80782-5](https://doi.org/10.1016/S0896-6273(00)80782-5)
- O'Donnell, C., Tiago Gonçalves, J., Portera-Cailliau, C., & Sejnowski, T. J. (2017). Beyond excitation/inhibition imbalance in multidimensional models of neural circuit changes in brain disorders. *eLife*, 6. <https://doi.org/10.7554/eLife.26724>
- O'Donnell, W. T., & Warren, S. T. (2002). A Decade of Molecular Studies of Fragile X Syndrome. *Annual Review of Neuroscience*, 25(1), 315–338. <https://doi.org/10.1146/annurev.neuro.25.112701.142909>
- O'Driscoll, T. (2017). *Computational Modelling of CaMKII, a Multi-state Biomolecule Involved in Memory*. University of Edinburgh.
- O'Keefe, J., & Nadel, L. (1978). Hippocampus Physiology. *The Hippocampus as a Cognitive Map*, (1), 141–230.
- Official Site | Second Life - Virtual Worlds, Virtual Reality, VR, Avatars, Free 3D Chat. (n.d.). Retrieved February 23, 2021, from <https://secondlife.com/>
- Oh, M. C., & Derkach, V. A. (2005). Dominant role of the GluR2 subunit in regulation of AMPA receptors by CaMKII. *Nature Neuroscience*, 8(7), 853–854. <https://doi.org/10.1038/nn1476>
- Oh, M. C., Derkach, V. A., Guire, E. S., & Soderling, T. R. (2006). Extrasynaptic membrane trafficking regulated by GluR1 serine 845 phosphorylation primes AMPA receptors for long-term potentiation. *Journal of Biological Chemistry*, 281(2), 752–758. <https://doi.org/10.1074/jbc.M509677200>
- Ohadi, D., Schmitt, D. L., Calabrese, B., Halpain, S., Zhang, J., & Rangamani, P. (2019). Computational Modeling Reveals Frequency Modulation of Calcium-cAMP/PKA Pathway in Dendritic Spines. *Biophysical Journal*, 117(10), 1963–1980. <https://doi.org/10.1016/j.bpj.2019.10.003>
- Okamoto, K.-I., Narayanan, R., Lee, S. H., Murata, K., & Hayashi, Y. (2007). The role of CaMKII as an F-actin-bundling protein crucial for maintenance of dendritic spine structure. *Proceedings of the National Academy of Sciences*, 104(15), 6418–6423. <https://doi.org/10.1073/pnas.0701656104>
- Ooms, J. (2020). *writexl: Export Data Frames to Excel "xlsx" Format*. Retrieved from <https://cran.r-project.org/package=writexl%0A>
- Opazo, P., Sainlos, M., & Choquet, D. (2012). Regulation of AMPA receptor surface diffusion by PSD-95 slots. *Current Opinion in Neurobiology*, 22(3), 453–460. <https://doi.org/10.1016/j.conb.2011.10.010>
- Osterweil, E. K., Chuang, S. C., Chubykin, A. A., Sidorov, M., Bianchi, R., Wong, R. K. S., & Bear, M. F. (2013). Lovastatin Corrects Excess Protein Synthesis and Prevents Epileptogenesis in a Mouse Model of Fragile X Syndrome. *Neuron*, 77(2), 243–250. <https://doi.org/10.1016/j.neuron.2012.01.034>
- Pachernegg, S., Munster, Y., Muth-Kohne, E., Fuhrmann, G., & Hollmann, M. (2015). GluA2 is rapidly edited at the Q/R site during neural differentiation in vitro. *Frontiers in Cellular Neuroscience*, 9, 69. <https://doi.org/10.3389/fncel.2015.00069>
- Padala, R. R., Karnawat, R., Viswanathan, S. B., Thakkar, A. V., & Das, A. B. (2017). Cancerous perturbations within the ERK, PI3K/Akt, and Wnt/ β -catenin signaling network constitutively activate inter-pathway positive feedback loops. *Molecular BioSystems*, 13(5), 830–840. <https://doi.org/10.1039/C6MB00786D>
- Palincsar, A. S. (1998). SOCIAL CONSTRUCTIVIST PERSPECTIVES ON TEACHING AND LEARNING. *Annual Review of Psychology*, 49(1), 345–375. <https://doi.org/10.1146/annurev.psych.49.1.345>
- Pandya, N. J., Koopmans, F., Slotman, J. A., Paliukhovich, I., Houtsmuller, A. B., Smit, A. B., & Li, K. W. (2017). Correlation profiling of brain sub-cellular proteomes reveals co-assembly of synaptic proteins and subcellular distribution. *Scientific Reports*, 7(1), 12107. <https://doi.org/10.1038/s41598-017-11690-3>
- Park, C. W., & Ryu, K. Y. (2014). Cellular ubiquitin pool dynamics and homeostasis. *BMB Reports*, Vol. 47, pp. 475–482. <https://doi.org/10.5483/BMBRep.2014.47.9.128>

- Park, P., Sanderson, T. M., Amici, M., Choi, S.-L., Bortolotto, Z. A., Zhuo, M., ... Collingridge, G. L. (2016). Calcium-Permeable AMPA Receptors Mediate the Induction of the Protein Kinase A-Dependent Component of Long-Term Potentiation in the Hippocampus. *Journal of Neuroscience*, 36(2), 622–631. <https://doi.org/10.1523/JNEUROSCI.3625-15.2016>
- Pasternack, A., Coleman, S. K., Joupila, A., Mottershead, D. G., Lindfors, M., Pasternack, M., & Keinänen, K. (2002). α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) receptor channels lacking the N-terminal domain. *Journal of Biological Chemistry*, 277(51), 49662–49667. <https://doi.org/10.1074/jbc.M208349200>
- Patrick, G. N., Bingol, B., Weld, H. A., & Schuman, E. M. (2003). Ubiquitin-Mediated Proteasome Activity Is Required for Agonist-Induced Endocytosis of GluRs. *Current Biology*, 13(23), 2073–2081. <https://doi.org/10.1016/j.cub.2003.10.028>
- Pedersen, S. W., Albertsen, L., Moran, G. E., Briévesque, B., Pedersen, S. B., Bartels, L., ... Strømgaard, K. (2017). Site-Specific Phosphorylation of PSD-95 PDZ Domains Reveals Fine-Tuned Regulation of Protein–Protein Interactions. *ACS Chem. Biol.*, 12. <https://doi.org/10.1021/acscchembio.7b00361>
- Pelkey, K. A., Chittajallu, R., Craig, M. T., Tricoire, L., Wester, J. C., & McBain, C. J. (2017). Hippocampal GABAergic Inhibitory Interneurons. *Physiological Reviews*, 97(4), 1619–1747. <https://doi.org/10.1152/physrev.00007.2017>
- Penders, B., Horstman, K., & Vos, R. (2008). Walking the Line between Lab and Computation: The “Moist” Zone. *BioScience*, 58(8), 747–755. <https://doi.org/10.1641/B580811>
- Penn, A. C., Zhang, C. L., Georges, F., Royer, L., Breillat, C., Hosy, E., ... Choquet, D. (2017). Hippocampal LTP and contextual learning require surface diffusion of AMPA receptors. *Nature*, 549(7672), 384–388. <https://doi.org/10.1038/nature23658>
- Petralia, R. S., Schwartz, C. M., Wang, Y.-X., Mattson, M. P., & Yao, P. J. (2011). Subcellular localization of patched and smoothened, the receptors for sonic hedgehog signaling, in the hippocampal neuron. *The Journal of Comparative Neurology*, 519(18), 3684–3699. <https://doi.org/10.1002/cne.22681>
- Pharris, M. C., Patel, N. M., VanDyk, T. G., Bartol, T. M., Sejnowski, T. J., Kennedy, M. B., ... Kinzer-Ursem, T. L. (2019). A multi-state model of the CaMKII dodecamer suggests a role for calmodulin in maintenance of autophosphorylation. *PLOS Computational Biology*, 15(12), e1006941. <https://doi.org/10.1371/journal.pcbi.1006941>
- Piard, J., Umanah, G. K. E., Harms, F. L., Abalde-Atristain, L., Amram, D., Chang, M., ... Kutsche, K. (2018). A homozygous ATAD1 mutation impairs postsynaptic AMPA receptor trafficking and causes a lethal encephalopathy. *Brain*, 141(3), 651–661. <https://doi.org/10.1093/brain/awx377>
- Pinar, C., Fontaine, C. J., Triviño-Paredes, J., Lottenberg, C. P., Gil-Mohapel, J., & Christie, B. R. (2017). Revisiting the flip side: Long-term depression of synaptic efficacy in the hippocampus. *Neuroscience & Biobehavioral Reviews*, 80, 394–413. <https://doi.org/10.1016/j.neubiorev.2017.06.001>
- Plant, K., Pelkey, K. A., Bortolotto, Z. A., Morita, D., Terashima, A., McBain, C. J., ... Isaac, J. T. R. (2006). Transient incorporation of native GluR2-lacking AMPA receptors during hippocampal long-term potentiation. *Nature Neuroscience*, 9(5), 602–604. <https://doi.org/10.1038/nn1678>
- Plass, J. L., Homer, B. D., & Kinzer, C. K. (2015). Foundations of Game-Based Learning. *Educational Psychologist*, 50(4), 258–283. <https://doi.org/10.1080/00461520.2015.1122533>
- Popescu, G. K. (2014). Accessories assist AMPA receptors to close pockets. *The Journal of General Physiology*, 145(1), 17–21. <https://doi.org/10.1085/jgp.201411320>
- Preyer, G., & Mans, D. (1999). Introduction: On Contemporary Developments in the Theory of Argumentation. *ProtoSociology*, 13, 3–13. <https://doi.org/10.5840/PROTOSOCIOLOGY1999131>
- Prieto, M., Folci, A., & Martin, S. (2020). Post-translational modifications of the Fragile X Mental Retardation Protein in neuronal function and dysfunction. *Molecular Psychiatry*, 25(8), 1688–1703. <https://doi.org/10.1038/s41380-019-0629-4>
- R: a language and environment for statistical computing. (n.d.). Retrieved February 16, 2021, from <https://www.gbif.org/tool/81287/r-a-language-and-environment-for-statistical-computing>
- Rau, M., Ohlmann, T., Morley, S. J., & Pain, V. M. (1996). A reevaluation of the Cap-binding protein, eIF4E, as a rate-limiting factor for initiation of translation in reticulocyte lysate. *Journal of Biological Chemistry*

- 277(15), 8983–8990. <https://doi.org/10.1074/jbc.271.15.8983>
- Ray, S., Deshpande, R., Dudani, N., & Bhalla, U. S. (2008). A general biological simulator: the multiscale object oriented simulation environment, MOOSE. *BMC Neuroscience*, 9, P93. <https://doi.org/10.1186/1471-2202-9-S1-P93>
- Raz, A., Grady, S. M., Krause, B. M., Uhlich, D. J., Manning, K. A., & Banks, M. I. (2014). Preferential effect of isoflurane on top-down vs. bottom-up pathways in sensory cortex. *Frontiers in Systems Neuroscience*, 8(OCT). <https://doi.org/10.3389/fnsys.2014.00191>
- Redondo, R. L., & Morris, R. G. M. (2011). Making memories last: the synaptic tagging and capture hypothesis. *Nature Reviews Neuroscience*, 12(1), 17–30. <https://doi.org/10.1038/nrn2963>
- Reinders, N. R., Pao, Y., Renner, M. C., Da Silva-Matos, C. M., Lodder, T. R., Malinow, R., & Kessels, H. W. (2016). Amyloid- β effects on synapses and memory require AMPA receptor subunit GluA3. *Proceedings of the National Academy of Sciences of the United States of America*, 113(42), E6526–E6534. <https://doi.org/10.1073/pnas.1614249113>
- Renner, M. C., Albers, E. H., Gutierrez-Castellanos, N., Reinders, N. R., van Huijstee, A. N., Xiong, H., ... Kessels, H. W. (2017). Synaptic plasticity through activation of GluA3-containing AMPA-receptors. *eLife*, 6. <https://doi.org/10.7554/eLife.25462>
- Resnick, M. (2014). Give P's a Chance: Project, Peers, Passion, Play. *Constructionism and Creativity Conference*. Vienna.
- Resnick, M., Maloney, Andrés Monroy-Hernández, J., Rusk, N., Eastmond, E., Brennan, K., Rosenbaum, E., ... Kafai, Y. (2009). "Digital fluency" should mean designing, creating, and remixing, not just browsing, chatting, and interacting. 52(11). <https://doi.org/10.1145/1592761.1592779>
- Reyes-Harde, M., & Stanton, P. K. (1998). Postsynaptic phospholipase C activity is required for the induction of homosynaptic long-term depression in rat hippocampus. *Neuroscience Letters*, 252(3), 155–158. [https://doi.org/10.1016/S0304-3940\(98\)00496-0](https://doi.org/10.1016/S0304-3940(98)00496-0)
- Rhodes, C. J., & Anderson, R. M. (1996). PHYSICS LETTERS A Dynamics in a lattice epidemic model. In *Physics Letters A* (Vol. 210).
- Richardson, V. (2003). Constructivist Pedagogy. *Teachers College Record*, 105(9), 1623–1640. <https://doi.org/10.1046/j.1467-9620.2003.00303.x>
- Rickinson, M., Dillon, J., Teamey, K., Morris, M., Choi, M. Y., Sanders, D., & Benefield, P. (2004). *A Review of Research on Outdoor Learning*. Unknown Publisher.
- Ringland, K. E., Boyd, L., Faucett, H., Cullen, A. L. L., & Hayes, G. R. (2017). Making in Minecraft. *Proceedings of the 2017 Conference on Interaction Design and Children*, 340–345. <https://doi.org/10.1145/3078072.3079749>
- Ringland, K. E., Wolf, C. T., Faucett, H., Dombrowski, L., & Hayes, G. R. (2016). "Will I always be not social?": Re-conceptualizing sociality in the context of a minecraft community for autism. *Conference on Human Factors in Computing Systems - Proceedings*, 1256–1269. <https://doi.org/10.1145/2858036.2858038>
- Rodrigues, E. M., Scudder, S. L., Goo, M. S., & Patrick, G. N. (2016). A β -induced synaptic alterations require the E3 ubiquitin ligase Nedd4-1. *Journal of Neuroscience*, 36(5), 1590–1595. <https://doi.org/10.1523/JNEUROSCI.2964-15.2016>
- Romano, D. R., Pharris, M. C., Patel, N. M., & Kinzer-Ursem, T. L. (2017). Competitive tuning: Competition's role in setting the frequency-dependence of Ca²⁺-dependent proteins. *PLOS Computational Biology*, 13(11), e1005820. <https://doi.org/10.1371/journal.pcbi.1005820>
- Romero, F. (2015). Why there isn't inter-level causation in mechanisms. *Synthese*, 192(11), 3731–3755. <https://doi.org/10.1007/s11229-015-0718-0>
- Rosenmund, C., Stern-Bach, Y., & Stevens, C. F. (1998). The tetrameric structure of a glutamate receptor channel. *Science*, 280(5369), 1596–1599. <https://doi.org/10.1126/science.280.5369.1596>
- Rotschafer, S. E., & Razak, K. a. (2014). Auditory Processing in Fragile X Syndrome. *Frontiers in Cellular Neuroscience*, 8(February), 19. <https://doi.org/10.3389/fncel.2014.00019>
- Rouach, N., Byrd, K., Petralia, R. S., Elias, G. M., Adesnik, H., Tomita, S., ... Nicoll, R. A. (2005). TARP γ -8 controls hippocampal AMPA receptor number, distribution and synaptic plasticity. *Nature Neuroscience*, 8(11),

- 1525–1533. <https://doi.org/10.1038/nn1551>
- Rougny, A., Touré, V., Moodie, S., Balaur, I., Czauderna, T., Borlinghaus, H., ... Luna, A. (2019). Systems Biology Graphical Notation: Process Description language Level 1 Version 2.0. *Journal of Integrative Bioinformatics*, 16(2). <https://doi.org/10.1515/jib-2019-0022>
- Roy, M., Sorokina, O., Skene, N., Simonnet, C., Mazzo, F., Zwart, R., ... Grant, S. G. N. (2018). Proteomic analysis of postsynaptic proteins in regions of the human neocortex. *Nature Neuroscience*, 21(1), 130–141. <https://doi.org/10.1038/s41593-017-0025-9>
- Roy, R. D., Rosenmund, C., & Stefan, M. I. (2017). Cooperative binding mitigates the high-dose hook effect. *BMC Systems Biology*, 11(1), 74. <https://doi.org/10.1186/s12918-017-0447-8>
- Saglietti, L., Dequidt, C., Kamieniarz, K., Rousset, M. C., Valnegri, P., Thoumine, O., ... Passafaro, M. (2007). Extracellular Interactions between GluR2 and N-Cadherin in Spine Regulation. *Neuron*, 54(3), 461–477. <https://doi.org/10.1016/j.neuron.2007.04.012>
- Saha, K., & Schaffer, D. V. (2006). Signal dynamics in Sonic hedgehog tissue patterning. *Development*, 133(5), 889–900. <https://doi.org/10.1242/dev.02254>
- Salussolia, C. L., Gan, Q., Kazi, R., Singh, P., Allopenna, J., Furukawa, H., & Wollmuth, L. P. (2013). A eukaryotic specific transmembrane segment is required for tetramerization in AMPA receptors. *Journal of Neuroscience*, 33(23), 9840–9845. <https://doi.org/10.1523/JNEUROSCI.2626-12.2013>
- Salzer, U., Kostan, J., & Djinić-Carugo, K. (2017). Deciphering the BAR code of membrane modulators. *Cellular and Molecular Life Sciences*, 74, 2413–2438. <https://doi.org/10.1007/s00018-017-2478-0>
- Samanta, A., Rongen, L. S. van, Rossato, J. I., Jacobse, J., Schoenfeld, R., & Genzel, L. (2020). Sleep leads to system-wide neural changes independent of allo- and egocentric spatial training in humans and rats. *BioRxiv*, 2020.12.19.423580. <https://doi.org/10.1101/2020.12.19.423580>
- Sanderson, J. L., Gorski, J. A., & Dell'Acqua, M. L. (2016). NMDA Receptor-Dependent LTD Requires Transient Synaptic Incorporation of Ca²⁺-Permeable AMPARs Mediated by AKAP150-Anchored PKA and Calcineurin. *Neuron*, 89(5), 1000–1015. <https://doi.org/10.1016/j.neuron.2016.01.043>
- Sanderson, J. L., Scott, J. D., & Dell'Acqua, M. L. (2018). Control of homeostatic synaptic plasticity by AKAP-anchored kinase and phosphatase regulation of Ca²⁺-permeable AMPA receptors. *Journal of Neuroscience*, 38(11), 2863–2876. <https://doi.org/10.1523/JNEUROSCI.2362-17.2018>
- Sandstrom, N. J., Kaufman, J., & A. Huettel, S. (1998). Males and females use different distal cues in a virtual environment navigation task. *Cognitive Brain Research*, 6(4), 351–360. [https://doi.org/10.1016/S0926-6410\(98\)00002-0](https://doi.org/10.1016/S0926-6410(98)00002-0)
- Saucerman, J. J., Brunton, L. L., Michailova, A. P., & McCulloch, A. D. (2003). Modeling β -Adrenergic Control of Cardiac Myocyte Contractility in Silico. *Journal of Biological Chemistry*, 278(48), 47997–48003. <https://doi.org/10.1074/jbc.M308362200>
- Sawicka, K., Hale, C. R., Park, C. Y., Fak, J. J., Gresack, J. E., Van Driesche, S. J., ... Darnell, R. B. (2019). FMRP has a cell-type-specific role in CA1 pyramidal neurons to regulate autism-related transcripts and circadian memory. *eLife*, 8. <https://doi.org/10.7554/eLife.46919>
- Saxena, S., & Maze, M. (2018). Impact on the brain of the inflammatory response to surgery. *La Presse Médicale*, 47(4), e73–e81. <https://doi.org/10.1016/j.lpm.2018.03.011>
- Scanziani, M., Malenka, R. C., & Nicoll, R. A. (1996). Role of intercellular interactions in heterosynaptic long-term depression. *Nature*, 380(6573), 446–450. <https://doi.org/10.1038/380446a0>
- Schoenfeld, R., Schifflholz, T., Beyer, C., Leplow, B., & Foreman, N. (2017). Variants of the Morris water maze task to comparatively assess human and rodent place navigation. *Neurobiology of Learning and Memory*, 139, 117–127. <https://doi.org/10.1016/j.nlm.2016.12.022>
- Schwarz, L. A., Hall, B. J., & Patrick, G. N. (2010). Activity-dependent ubiquitination of GluA1 mediates a distinct AMPA receptor endocytosis and sorting pathway. *Journal of Neuroscience*, 30(49), 16718–16729. <https://doi.org/10.1523/JNEUROSCI.3686-10.2010>
- Schwenk, J., Baehrens, D., Haupt, A., Bildl, W., Boudkazi, S., Roeper, J., ... Schulte, U. (2014). Regional diversity and developmental dynamics of the AMPA-receptor proteome in the mammalian brain. *Neuron*, 84(1), 41–54. <https://doi.org/10.1016/j.neuron.2014.08.044>

- Schwenk, J., Harmel, N., Brechet, A., Zolles, G., Berkefeld, H., Müller, C. S., ... Fakler, B. (2012). High-Resolution Proteomics Unravel Architecture and Molecular Diversity of Native AMPA Receptor Complexes. *Neuron*, 74(4), 621–633. <https://doi.org/10.1016/j.neuron.2012.03.034>
- Schwenk, J., Harmel, N., Zolles, G., Bildl, W., Kulik, A., Heimrich, B., ... Klöcker, N. (2009). Functional proteomics identify cornichon proteins as auxiliary subunits of AMPA receptors. *Science*, 323(5919), 1313–1319. <https://doi.org/10.1126/science.1167852>
- Schworer, C. M., Colbran, R. J., Keefer, J. R., & Soderling, T. R. (1988). Ca²⁺/calmodulin-dependent protein kinase II. Identification of a regulatory autophosphorylation site adjacent to the inhibitory and calmodulin-binding domains. *Journal of Biological Chemistry*, 263(27), 13486–13489. [https://doi.org/10.1016/s0021-9258\(18\)68264-x](https://doi.org/10.1016/s0021-9258(18)68264-x)
- Scott, G. W., Goulder, R., Wheeler, P., Scott, L. J., Tobin, M. L., & Marsham, S. (2012). The Value of Fieldwork in Life and Environmental Sciences in the Context of Higher Education: A Case Study in Learning About Biodiversity. *Journal of Science Education and Technology*, 21(1), 11–21. <https://doi.org/10.1007/s10956-010-9276-x>
- Seidenman, K. J., Steinberg, J. P., Huganir, R., & Malinow, R. (2003). Glutamate receptor subunit 2 serine 880 phosphorylation modulates synaptic transmission and mediates plasticity in CA1 pyramidal cells. *Journal of Neuroscience*, 23(27), 9220–9228. <https://doi.org/10.1523/jneurosci.23-27-09220.2003>
- Sellers, K. K., Bennett, D. V., Hutt, A., Williams, J. H., & Fröhlich, F. (2015). Awake vs. anesthetized: Layer-specific sensory processing in visual cortex and functional connectivity between cortical areas. *Journal of Neurophysiology*, 113(10), 3798–3815. <https://doi.org/10.1152/jn.00923.2014>
- Seol, G. H., Ziburkus, J., Huang, S. Y., Song, L., Kim, I. T., Takamiya, K., ... Kirkwood, A. (2007). Neuromodulators Control the Polarity of Spike-Timing-Dependent Synaptic Plasticity. *Neuron*, 55(6), 919–929. <https://doi.org/10.1016/j.neuron.2007.08.013>
- Shanks, N. F., Cais, O., Maruo, T., Savas, X. J. N., Zaika, E. I., Azumaya, C. M., ... Nakagawa, T. (2014). Molecular dissection of the interaction between the AMPA receptor and cornichon homolog-3. *Journal of Neuroscience*, 34(36), 12104–12120. <https://doi.org/10.1523/JNEUROSCI.0595-14.2014>
- Sheail, P., Dunnett, N., & Nicol, S. (2018). 'PLZ HELP': Navigating Minecraft as a digital education space with postgraduate students. *The University of Edinburgh Learning and Teaching Conference*, 30. Edinburgh.
- Sheng, M., & Ertürk, A. (2014). Long-term depression: a cell biological view. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633), 20130138. <https://doi.org/10.1098/rstb.2013.0138>
- Shepherd, J. D., & Huganir, R. L. (2007). The Cell Biology of Synaptic Plasticity: AMPA Receptor Trafficking. *Annual Review of Cell and Developmental Biology*, 23(1), 613–643. <https://doi.org/10.1146/annurev.cellbio.23.090506.123516>
- Shi, S. H., Hayashi, Y., Esteban, J. A., & Malinow, R. (2001). Subunit-specific rules governing AMPA receptor trafficking to synapses in hippocampal pyramidal neurons. *Cell*, 105(3), 331–343. [https://doi.org/10.1016/S0092-8674\(01\)00321-X](https://doi.org/10.1016/S0092-8674(01)00321-X)
- Shi, Y., Suh, Y. H., Milstein, A. D., Isozaki, K., Schmid, S. M., Roche, K. W., & Nicoll, R. A. (2010). Functional comparison of the effects of TARPs and cornichons on AMPA receptor trafficking and gating. *Proceedings of the National Academy of Sciences of the United States of America*, 107(37), 16315–16319. <https://doi.org/10.1073/pnas.1011706107>
- Sia, G. M., Béïque, J. C., Rumbaugh, G., Cho, R., Worley, P. F., & Huganir, R. L. (2007). Interaction of the N-Terminal Domain of the AMPA Receptor GluR4 Subunit with the Neuronal Pentraxin NP1 Mediates GluR4 Synaptic Recruitment. *Neuron*, 55(1), 87–102. <https://doi.org/10.1016/j.neuron.2007.06.020>
- Silverman, J. B., Restituito, S., Lu, W., Lee-Edwards, L., Khatri, L., & Ziff, E. B. (2007). Synaptic anchorage of AMPA receptors by cadherins through neural plakophilin-related arm protein-AMPA receptor-binding protein complexes. *Journal of Neuroscience*, 27(32), 8505–8516. <https://doi.org/10.1523/JNEUROSCI.1395-07.2007>
- Singh, D., & Bhalla, U. S. (2018). Subunit exchange enhances information retention by CaMKII in dendritic spines. *eLife*, 7. <https://doi.org/10.7554/eLife.41412>
- Sivakumaran, S., Hariharaputran, S., Mishra, J., & Bhalla, U. S. (2003). The database of quantitative cellular signaling: Management and analysis of chemical kinetic models of signaling networks. *Bioinformatics*, 19(3), 408–415. <https://doi.org/10.1093/bioinformatics/btf860>

- Sneddon, M. W., Faeder, J. R., & Emonet, T. (2011). Efficient modeling, simulation and coarse-graining of biological complexity with NFsim. *Nature Methods*, 8(2), 177–183. <https://doi.org/10.1038/nmeth.1546>
- Song, H., Smolen, P., Av-Ron, E., Baxter, D. A., & Byrne, J. H. (2006). Bifurcation and Singularity Analysis of a Molecular Network for the Induction of Long-Term Memory. *Biophysical Journal*, 90(7), 2309–2325. <https://doi.org/10.1529/biophysj.105.074500>
- Sorokin, A., Heil, K. F., Armstrong, J. D., & Sorokina, O. (2018). Rule-based modelling provides an extendable framework for comparing candidate mechanisms underpinning clathrin polymerisation. *Scientific Reports*, 8(1), 5658. <https://doi.org/10.1038/s41598-018-23829-x>
- Sorokin, A., Le Novère, N., Luna, A., Czauderna, T., Demir, E., Haw, R., ... Villéger, A. (2015). Systems Biology Graphical Notation: Entity Relationship language Level 1 Version 2. *Journal of Integrative Bioinformatics*, 12(2), 281–339. <https://doi.org/10.1515/jib-2015-264>
- Sorokina, O., Sorokin, A., Armstrong, J. D., & Danos, V. (2013). A simulator for spatially extended kappa models. *Bioinformatics*, 29(23), 3105–3106. <https://doi.org/10.1093/bioinformatics/btt523>
- Sossin, W. S. (2007). Isoform specificity of protein kinase Cs in synaptic plasticity. *Learning & Memory*, 14(4), 236–246. <https://doi.org/10.1101/lm.469707>
- Srivastava, P., Rimzhim, A., Vijay, P., Singh, S., & Chandra, S. (2019). Desktop VR Is Better Than Non-ambulatory HMD VR for Spatial Learning. *Frontiers in Robotics and AI*, 6, 50. <https://doi.org/10.3389/frobt.2019.00050>
- Srivastava, S., Osten, P., Vilim, F. S., Khatri, L., Inman, G., States, B., ... Ziff, E. B. (1998). Novel anchorage of GluR2/3 to the postsynaptic density by the AMPA receptor-binding protein ABP. *Neuron*, 21(3), 581–591. [https://doi.org/10.1016/S0896-6273\(00\)80568-1](https://doi.org/10.1016/S0896-6273(00)80568-1)
- States, B. A., Khatri, L., & Ziff, E. B. (2008). Stable synaptic retention of serine-880-phosphorylated GluR2 in hippocampal neurons. *Molecular and Cellular Neuroscience*, 38(2), 189–202. <https://doi.org/10.1016/j.mcn.2008.02.003>
- Stefan, M. I., Bartol, T. M., Sejnowski, T. J., & Kennedy, M. B. (2014). Multi-state Modeling of Biomolecules. *PLoS Computational Biology*, 10(9), e1003844. <https://doi.org/10.1371/journal.pcbi.1003844>
- Stefan, M. I., Edelstein, S. J., & Novere, N. Le. (2008). An allosteric model of calmodulin explains differential activation of PP2B and CaMKII. *Proceedings of the National Academy of Sciences*, 105(31), 10768–10773. <https://doi.org/10.1073/pnas.0804672105>
- Stein, E. L. A., & Chetkovich, D. M. (2010). Regulation of stargazin synaptic trafficking by C-terminal PDZ ligand phosphorylation in bidirectional synaptic plasticity. *Journal of Neurochemistry*, 113(1), 42–53. <https://doi.org/10.1111/j.1471-4159.2009.06529.x>
- Steinberg, J. P., Takamiya, K., Shen, Y., Xia, J., Rubio, M. E., Yu, S., ... Huganir, R. L. (2006). Targeted in vivo mutations of the AMPA receptor subunit GluR2 and its interacting protein PICK1 eliminate cerebellar long-term depression. *Neuron*, 49(6), 845–860. <https://doi.org/10.1016/j.neuron.2006.02.025>
- Sterratt, D., Graham, B., Gillies, A., & Willshaw, D. (2011). Intracellular mechanisms. In *Principles of Computational Modelling in Neuroscience* (pp. 133–171). <https://doi.org/10.1017/CBO9780511975899.007>
- Stevens, J., Kincaid, P., & Sottolare, R. (2015). Visual modality research in virtual and mixed reality simulation. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 12(4), 519–537. <https://doi.org/10.1177/1548512915569742>
- Stiles, J. R., Van Helden, D., Bartol, T. M., Salpeter, E. E., & Salpeter, M. M. (1996). Miniature endplate current rise times less than 100 microseconds from improved dual recordings can be modeled with passive acetylcholine diffusion from a synaptic vesicle. *Proceedings of the National Academy of Sciences*, 93(12), 5747–5752. <https://doi.org/10.1073/pnas.93.12.5747>
- Suderman, R., & Deeds, E. J. (2013). Machines vs. Ensembles: Effective MAPK Signaling through Heterogeneous Sets of Protein Complexes. *PLoS Computational Biology*, 9(10), e1003278. <https://doi.org/10.1371/journal.pcbi.1003278>
- Sumioka, A., Yan, D., & Tomita, S. (2010). TARP Phosphorylation Regulates Synaptic AMPA Receptors through Lipid Bilayers. *Neuron*, 66(5), 755–767. <https://doi.org/10.1016/j.neuron.2010.04.035>
- Sundberg, S. (2018). *Neuromodulation, Short-Term and Long-Term Plasticity in Corticothalamic and Hippocampal Neuronal Networks*. Linköping University.

- Surin, A. M., Gorbacheva, L. R., Savinkova, I. G., Sharipov, R. R., Khodorov, B. I., & Pinelis, V. G. (2014). Study on ATP concentration changes in cytosol of individual cultured neurons during glutamate-induced deregulation of calcium homeostasis. *Biochemistry (Moscow)*, 79(2), 146–157. <https://doi.org/10.1134/S0006297914020084>
- Sweetser, P., & Wyeth, P. (2005). GameFlow: A Model for Evaluating Player Enjoyment in Games. In *ACM Computers in Entertainment* (Vol. 3).
- Tabet, R., Moutin, E., Becker, J. A. J., Heintz, D., Fouillen, L., Flatter, E., ... Moine, H. (2016). Fragile X mental retardation protein (FMRP) controls diacylglycerol kinase activity in neurons. *Proceedings of the National Academy of Sciences of the United States of America*, 113(26), E3619–E3628. <https://doi.org/10.1073/pnas.1522631113>
- Tabet, R., Vitale, N., & Moine, H. (2016). Fragile X syndrome: Are signaling lipids the missing culprits? *Biochimie*, 1–7. <https://doi.org/10.1016/j.biochi.2016.09.002>
- Tachiwana, H., Osakabe, A., Shiga, T., Miya, Y., Kimura, H., Kagawa, W., & Kurumizaka, H. (2011). Structures of human nucleosomes containing major histone H3 variants. *Acta Crystallographica Section D: Biological Crystallography*, 67(6), 578–583. <https://doi.org/10.1107/S0907444911014818>
- Tai, C.-Y., Kim, S. A., & Schuman, E. M. (2008). Cadherins and synaptic plasticity. *Current Opinion in Cell Biology*, 20(5), 567–575. <https://doi.org/10.1016/j.ceb.2008.06.003>
- Takamiya, K., Kostourou, V., Adams, S., Jadeja, S., Chalepakis, G., Scambler, P. J., ... Adams, R. H. (2004). A direct functional link between the multi-PDZ domain protein GRIP1 and the Fraser syndrome protein Fras1. *Nature Genetics*, 36(2), 172–177. <https://doi.org/10.1038/ng1292>
- Takamiya, K., Mao, L., Haganir, R. L., & Linden, D. J. (2008). The glutamate receptor-interacting protein family of GluR2-binding proteins is required for long-term synaptic depression expression in cerebellar Purkinje cells. *Journal of Neuroscience*, 28(22), 5752–5755. <https://doi.org/10.1523/JNEUROSCI.0654-08.2008>
- Takeuchi, T., Duzkiewicz, A. J., & Morris, R. G. M. (2014). The synaptic plasticity and memory hypothesis: encoding, storage and persistence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633), 20130288. <https://doi.org/10.1098/rstb.2013.0288>
- Takeuchi, Y., Morise, J., Morita, I., Takematsu, H., & Oka, S. (2015). Role of Site-Specific N-Glycans Expressed on GluA2 in the Regulation of Cell Surface Expression of AMPA-Type Glutamate Receptors. *PLOS ONE*, 10(8), e0135644. <https://doi.org/10.1371/journal.pone.0135644>
- Tanaka, H., & Hirano, T. (2012). Visualization of Subunit-Specific Delivery of Glutamate Receptors to Postsynaptic Membrane during Hippocampal Long-Term Potentiation. *Cell Reports*, 1(4), 291–298. <https://doi.org/10.1016/j.celrep.2012.02.004>
- Tang, Y., Ye, M., Du, Y., Qiu, X., Lv, X., Yang, W., & Luo, J. (2015). EGFR signaling upregulates surface expression of the GluN2B-containing NMDA receptor and contributes to long-term potentiation in the hippocampus. *Neuroscience*, 304, 109–121. <https://doi.org/10.1016/j.neuroscience.2015.07.021>
- Tangherloni, A., Nobile, M. S., Besozzi, D., Mauri, G., & Cazzaniga, P. (2017). LASSIE: Simulating large-scale models of biochemical systems on GPUs. *BMC Bioinformatics*, 18(1), 246. <https://doi.org/10.1186/s12859-017-1666-0>
- Telias, M. (2019). Molecular Mechanisms of Synaptic Dysregulation in Fragile X Syndrome and Autism Spectrum Disorders. *Frontiers in Molecular Neuroscience*, 12, 51. <https://doi.org/10.3389/fnmol.2019.00051>
- Terashima, A., Pelkey, K. A., Rah, J. C., Suh, Y. H., Roche, K. W., Collingridge, G. L. L., ... Isaac, J. T. R. (2008). An Essential Role for PICK1 in NMDA Receptor-Dependent Bidirectional Synaptic Plasticity. *Neuron*, 57(6), 872–882. <https://doi.org/10.1016/j.neuron.2008.01.028>
- Tesmer, V. M., Lennarz, S., Mayer, G., & Tesmer, J. J. G. (2012). Molecular mechanism for inhibition of G protein-coupled receptor kinase 2 by a selective RNA aptamer. *Structure*, 20(8), 1300–1309. <https://doi.org/10.1016/j.str.2012.05.002>
- Thomas, A. M., Bui, N., Graham, D., Perkins, J. R., Yuva-Paylor, L. A., & Paylor, R. (2011). Genetic reduction of group 1 metabotropic glutamate receptors alters select behaviors in a mouse model for fragile X syndrome. *Behavioural Brain Research*, 223(2), 310–321. <https://doi.org/10.1016/j.bbr.2011.04.049>
- Tian, H., Xu, Y., Liu, F., Wang, G., & Hu, S. (2015). Effect of acute fentanyl treatment on synaptic plasticity in the

- hippocampal CA1 region in rats. *Frontiers in Pharmacology*, 6(OCT). <https://doi.org/10.3389/fphar.2015.00251>
- Toh, W., & Kirschner, D. (2020). Self-directed learning in video games, affordances and pedagogical implications for teaching and learning. *Computers & Education*, 154, 103912. <https://doi.org/10.1016/j.compedu.2020.103912>
- Tolle, D. P., & Le Novère, N. (2010). Brownian diffusion of AMPA receptors is sufficient to explain fast onset of LTP. *BMC Systems Biology*, 4(1), 25. <https://doi.org/10.1186/1752-0509-4-25>
- Tomita, S., Stein, V., Stocker, T. J., Nicoll, R. A., & Bredt, D. S. (2005). Bidirectional synaptic plasticity regulated by phosphorylation of stargazin-like TARPs. *Neuron*, 45(2), 269–277. <https://doi.org/10.1016/j.neuron.2005.01.009>
- Torrecillas, A., Laynez, J., Menéndez, M., Corbalán-García, S., & Gómez-Fernández, J. C. (2004). Calorimetric Study of the Interaction of the C2 Domains of Classical Protein Kinase C Isoenzymes with Ca²⁺ and Phospholipids †. *Biochemistry*, 43(37), 11727–11739. <https://doi.org/10.1021/bi0489659>
- Traiffort, E., Charytoniuk, D. A., Faure, H., & Ruat, M. (2002). Regional Distribution of Sonic Hedgehog, Patched, and Smoothened mRNA in the Adult Rat Brain. *Journal of Neurochemistry*, 79(3), 1327–1330. <https://doi.org/10.1046/j.1471-4159.1998.70031327.x>
- Traiffort, E., Charytoniuk, D., Watroba, L., Faure, H., Sales, N., & Ruat, M. (1999). Discrete localizations of hedgehog signalling components in the developing and adult rat nervous system. *European Journal of Neuroscience*, 11(9), 3199–3214. <https://doi.org/10.1046/j.1460-9568.1999.00777.x>
- Traynelis, S. F., Wollmuth, L. P., McBain, C. J., Menniti, F. S., Vance, K. M., Ogden, K. K., ... Dingledine, R. (2010). Glutamate Receptor Ion Channels: Structure, Regulation, and Function. *Pharmacological Reviews*, 62(3), 405–496. <https://doi.org/10.1124/pr.109.002451>
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. a, Wood, E. R., ... Morris, R. G. M. (2007). Schemas and memory consolidation. *Science (New York, N.Y.)*, 316(5821), 76–82. <https://doi.org/10.1126/science.1135935>
- Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., ... Morris, R. G. M. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science (New York, N.Y.)*, 333(6044), 891–895. <https://doi.org/10.1126/science.1205274>
- Tukachinsky, H., Lopez, L. V., & Salic, A. (2010). A mechanism for vertebrate Hedgehog signaling: Recruitment to cilia and dissociation of SuFu-Gli protein complexes. *Journal of Cell Biology*, 191(2), 415–428. <https://doi.org/10.1083/jcb.201004108>
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., & Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670), 892–896. <https://doi.org/10.1038/36103>
- Unwin, N. (2005). Refined structure of the nicotinic acetylcholine receptor at 4 Å resolution. *Journal of Molecular Biology*, 346(4), 967–989. <https://doi.org/10.1016/j.jmb.2004.12.031>
- Urban-Ciecko, J., & Barth, A. L. (2016). Somatostatin-expressing neurons in cortical networks. *Nature Reviews Neuroscience*, 17(7), 401–409. <https://doi.org/10.1038/nrn.2016.53>
- Van der Molen, M. J. W., Huizinga, M., Huizenga, H. M., Ridderinkhof, K. R., Van der Molen, M. W., Hamel, B. J. C., ... Ramakers, G. J. A. (2010). Profiling Fragile X Syndrome in males: Strengths and weaknesses in cognitive abilities. *Research in Developmental Disabilities*, 31(2), 426–439. <https://doi.org/10.1016/j.ridd.2009.10.013>
- van der Most, P. J., Dolga, A. M., Nijholt, I. M., Luiten, P. G. M., & Eisel, U. L. M. (2009). Statins: Mechanisms of neuroprotection. *Progress in Neurobiology*, 88(1), 64–75. <https://doi.org/10.1016/j.pneurobio.2009.02.002>
- van Elburg, R. A. J., & van Ooyen, A. (2010). Impact of dendritic size and dendritic topology on burst firing in pyramidal cells. *PLoS Computational Biology*, 6(5), 1–19. <https://doi.org/10.1371/journal.pcbi.1000781>
- van Kesteren, M. T. R., Fernandez, G., Norris, D. G., & Hermans, E. J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences*, 107(16), 7550–7555. <https://doi.org/10.1073/pnas.0914892107>
- van Kesteren, M. T. R., Rijpkema, M., Ruiter, D. J., Morris, R. G. M., & Fernández, G. (2014). Building on Prior Knowledge: Schema-dependent Encoding Processes Relate to Academic Performance. *Journal of Cognitive*

- Neuroscience*, 26(10), 2250–2261. https://doi.org/10.1162/jocn_a_00630
- van Kesteren, M. T. R., Ruiter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in Neurosciences*, 35(4), 211–219. <https://doi.org/10.1016/j.tins.2012.02.001>
- van Riel, N. A. W. (2006). Dynamic modelling and analysis of biochemical networks: Mechanism-based models and model-based experiments. *Briefings in Bioinformatics*, 7(4), 364–374. <https://doi.org/10.1093/bib/bbl040>
- Varela, F. G., Maturana, H. R., & Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *BioSystems*, 5(4), 187–196. [https://doi.org/10.1016/0303-2647\(74\)90031-8](https://doi.org/10.1016/0303-2647(74)90031-8)
- Vargas, J. Y., Fuenzalida, M., & Inestrosa, N. C. (2014). In vivo activation of Wnt signaling pathway enhances cognitive function of adult mice and reverses cognitive deficits in an Alzheimer's disease model. *Journal of Neuroscience*, 34(6), 2191–2202. <https://doi.org/10.1523/JNEUROSCI.0862-13.2014>
- Vayttaden, S. J., & Bhalla, U. S. (2004). Developing Complex Signaling Models Using GENESIS/Kinetikit. *Science Signaling*, 2004(219), pl4–pl4. <https://doi.org/10.1126/stke.2192004pl4>
- Vertommen, D., Rider, M., Ni, Y., Waelkens, E., Merlevede, W., Vandenheede, J. R., & Van Lint, J. (2000). Regulation of Protein Kinase D by Multisite Phosphorylation. *Journal of Biological Chemistry*, 275(26), 19567–19576. <https://doi.org/10.1074/jbc.M001357200>
- Vianna Gallinaro, J. (2020). *Neuronal assembly formation and non-random recurrent connectivity induced by homeostatic structural plasticity* (Albert-Ludwigs Universität Freiburg im Breisgau; KTH, Royal Institute of Technology, Stockholm). Retrieved from <https://www.diva-portal.org/smash/get/diva2:1412927/FULLTEXT01.pdf%23page%3D61&hl=fr&sa=X&d=328487950777889279&scisig=AAGBfm0wkJvaqXw2QNNnVEuzHz6R3INXA&nossl=1&oi=scholar&rt&hist=aBs5f48AAA AJ:12411142020873097189:AAGBfm1L4nNWxlQz3QzSFyD3Nick5Mhakg#page=>
- Villalobos, M., & Ward, D. (2015). Living Systems: Autonomy, Autopoiesis and Enaction. *Philosophy & Technology*, 28(2), 225–239. <https://doi.org/10.1007/s13347-014-0154-y>
- Viswan, N. A., Vani HarshaRani, G., Stefan, M. I., Bhalla, U. S., Ekkehard Plesser, H., Daniel Meliza, C., & Phillip Nickerson, D. (2018). *FindSim: A Framework for Integrating Neuronal Data and Signaling Models*. <https://doi.org/10.3389/fninf.2018.00038>
- Volk, L., Kim, C. H., Takamiya, K., Yu, Y., & Huganir, R. L. (2010). Developmental regulation of protein interacting with C kinase 1 (PICK1) function in hippocampal synaptic plasticity and learning. *Proceedings of the National Academy of Sciences of the United States of America*, 107(50), 21784–21789. <https://doi.org/10.1073/pnas.1016103107>
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Cambridge, Mass.: Harvard University Press.
- Waltemath, D., Adams, R., Bergmann, F. T., Hucka, M., Kolpakov, F., Miller, A. K., ... Le Novère, N. (2011). Reproducible computational biology experiments with SED-ML - The Simulation Experiment Description Markup Language. *BMC Systems Biology*, 5(1), 198. <https://doi.org/10.1186/1752-0509-5-198>
- Walther, B. K. (2003). Playing and Gaming: Reflections and Classifications. *Game Studies*, 3(1). Retrieved from <http://www.gamestudies.org/0301/walther/>
- Walton, D. (2001). Abductive, presumptive and plausible arguments. *Informal Logic*, 21(2). Retrieved from <https://philpapers.org/rec/WALAPA-5>
- Wang, H. (2014). Lipid rafts: a signaling platform linking cholesterol metabolism to synaptic deficits in autism spectrum disorders. *Frontiers in Behavioral Neuroscience*, 8(MAR), 104. <https://doi.org/10.3389/fnbeh.2014.00104>
- Wang, H., Pati, S., Pozzo-Miller, L., & Doering, L. C. (2015). Targeted pharmacological treatment of autism spectrum disorders: fragile X and Rett syndromes. *Frontiers in Cellular Neuroscience*, 9(February), 55. <https://doi.org/10.3389/fncel.2015.00055>
- Wang, L., Chiang, H. C., Wu, W., Liang, B., Xie, Z., Yao, X., ... Zhong, Y. (2012). Epidermal growth factor receptor is a preferred target for treating Amyloid- β -induced memory loss. *Proceedings of the National Academy of Sciences of the United States of America*, 109(41), 16743–16748. <https://doi.org/10.1073/pnas.1208011109>
- Wang, Y., Rowan, M. J., & Anwyl, R. (1997). Induction of LTD in the Dentate Gyrus In Vitro Is NMDA Receptor

- Independent, but Dependent on Ca²⁺ Influx via Low-Voltage-Activated Ca²⁺ Channels and Release of Ca²⁺ From Intracellular Stores. *Journal of Neurophysiology*, 77(2), 812–825. <https://doi.org/10.1152/jn.1997.77.2.812>
- Wang, Y. T., & Linden, D. J. (2000). Expression of cerebellar long-term depression requires postsynaptic clathrin-mediated endocytosis. *Neuron*, 25(3), 635–647. [https://doi.org/10.1016/S0896-6273\(00\)81066-1](https://doi.org/10.1016/S0896-6273(00)81066-1)
- Wang, Z.-H., Xiang, J., Liu, X., Yu, S. P., Manfredsson, F. P., Sandoval, I. M., ... Ye, K. (2019). Deficiency in BDNF/TrkB Neurotrophic Activity Stimulates δ -Secretase by Upregulating C/EBP β in Alzheimer's Disease. *Cell Reports*, 28(3), 655–669.e5. <https://doi.org/10.1016/j.celrep.2019.06.054>
- Waters, J., & Smith, S. J. (2002). Vesicle pool partitioning influences presynaptic diversity and weighting in rat hippocampal synapses. *The Journal of Physiology*, 541(3), 811–823. <https://doi.org/10.1113/jphysiol.2001.013485>
- Watson, J. F., Ho, H., & Greger, I. H. (2017). Synaptic transmission and plasticity require AMPA receptor anchoring via its N-terminal domain. *eLife*, 6. <https://doi.org/10.7554/eLife.23024>
- Watson, K. K., Jones, T. K., & Allman, J. M. (2006). Dendritic architecture of the von Economo neurons. *Neuroscience*, 141(3), 1107–1112. <https://doi.org/10.1016/j.neuroscience.2006.04.084>
- Waxham, M. N., Tsai, A., & Putkey, J. A. (1998). A Mechanism for Calmodulin (CaM) Trapping by CaM-kinase II Defined by a Family of CaM-binding Peptides. *Journal of Biological Chemistry*, 273(28), 17579–17584. <https://doi.org/10.1074/jbc.273.28.17579>
- Weber, A., & Varela, F. (2002). Life after Kant: Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences*, 1(2), 97–125. <https://doi.org/10.1023/A:1020368120174>
- Wei, J., Xiong, Z., Lee, J. B., Cheng, J., Duffney, L. J., Matas, E., & Yan, Z. (2016). Histone modification of Nedd4 ubiquitin ligase controls the loss of AMPA receptors and cognitive impairment induced by repeated stress. *Journal of Neuroscience*, 36(7), 2119–2130. <https://doi.org/10.1523/JNEUROSCI.3056-15.2016>
- Wells, G. (1999). Discourse and Knowing in the Classroom. In *Dialogic Inquiry* (pp. 98–132). <https://doi.org/10.1017/CBO9780511605895.005>
- Wentholt, R. J., Petralia, R. S., Blahos, J., & Niedzielski, A. S. (1996). Evidence for multiple AMPA receptor complexes in hippocampal CA1/CA2 neurons. *Journal of Neuroscience*, 16(6), 1982–1989. <https://doi.org/10.1523/jneurosci.16-06-01982.1996>
- Whewey, G., Nazlamova, L., & Hancock, J. T. (2018). Signaling through the Primary Cilium. *Frontiers in Cell and Developmental Biology*, 6(FEB). <https://doi.org/10.3389/fcell.2018.00008>
- Whitcomb, D. J., Hogg, E. L., Regan, P., Piers, T., Narayan, P., Whitehead, G., ... Cho, K. (2015). Intracellular oligomeric amyloid-beta rapidly regulates GluA1 subunit of AMPA receptor in the hippocampus. *Scientific Reports*, 5(1), 10934. <https://doi.org/10.1038/srep10934>
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Retrieved from <https://ggplot2.tidyverse.org>
- Wickham, H., & Bryan, J. (2019). *readxl: Read Excel Files*. Retrieved from <https://cran.r-project.org/package=readxl%0A>
- Widagdo, J., Chai, Y., Ridder, M. C., Chau, Y., Johnson, R. C., Sah, P., ... Anggono, V. (2015). Activity-Dependent ubiquitination of GluA1 and GluA2 regulates AMPA receptor intracellular sorting and degradation. *Cell Reports*, 10(5), 783–795. <https://doi.org/10.1016/j.celrep.2015.01.015>
- Widagdo, J., Guntupalli, S., Jang, S. E., & Anggono, V. (2017). Regulation of AMPA receptor trafficking by protein ubiquitination. *Frontiers in Molecular Neuroscience*, 10, 347. <https://doi.org/10.3389/fnmol.2017.00347>
- Wikström, M. A., Matthews, P., Roberts, D., Collingridge, G. L., & Bortolotto, Z. A. (2003). Parallel kinase cascades are involved in the induction of LTP at hippocampal CA1 synapses. *Neuropharmacology*, 45(6), 828–836. [https://doi.org/10.1016/S0028-3908\(03\)00336-8](https://doi.org/10.1016/S0028-3908(03)00336-8)
- Williams, L., Yang, K., Wiebe, E., Ferzli, M., & Miller, C. (2002). Pair Programming in an Introductory Computer Science Course: Initial Results and Recommendations. *Proceedings of the AAAI Workshop on Approximation and Abstraction of Computational Theories*.
- Wilson, K. A., Bedwell, W. L., Lazzara, E. H., Salas, E., Burke, S., Estock, J. L., ... Conkey, C. (2009). Relationships

- Between Game Attributes and Learning Outcomes Review and Research Proposals. *Simulation & Gaming*, 40, 217–266. <https://doi.org/10.1177/1046878108321866>
- Wimsatt, W. C. (2006). Reductionism and its heuristics: Making methodological reductionism honest. *Synthese*, 151(3), 445–475. <https://doi.org/10.1007/s11229-006-9017-0>
- Wittig, U., Kania, R., Golebiewski, M., Rey, M., Shi, L., Jong, L., ... Muller, W. (2012). SABIO-RK--database for biochemical reaction kinetics. *Nucleic Acids Research*, 40(D1), D790–D796. <https://doi.org/10.1093/nar/gkr1046>
- Wolbers, T., & Hegarty, M. (2010). What determines our navigational abilities? *Trends in Cognitive Sciences*, 14(3), 138–146. <https://doi.org/10.1016/j.tics.2010.01.001>
- Wood, D., Bruner, J. S., & Ross, G. (1976). THE ROLE OF TUTORING IN PROBLEM SOLVING*. In *J. Child Psychol. Psychiat* (Vol. 17). Pergamon Press.
- Wood, D. F. (2003). ABC of learning and teaching in medicine: Problem based learning. *BMJ*, 326(7384), 328–330. <https://doi.org/10.1136/bmj.326.7384.328>
- Wright, L. K., Cardenas, J. J., Liang, P., & Newman, D. L. (2018). Arrows in biology: Lack of clarity and consistency points to confusion for learners. *CBE Life Sciences Education*, 17(1). <https://doi.org/10.1187/cbe.17-04-0069>
- Wright, L. K., Fisk, J. N., & Newman, D. L. (2014). DNA → RNA: What Do Students Think the Arrow Means? *CBE—Life Sciences Education*, 13(2), 338–348. <https://doi.org/10.1187/cbe.cbe-13-09-0188>
- Xia, J., Zhang, X., Staudinger, J., & Huganir, R. L. (1999). Clustering of AMPA receptors by the synaptic PD domain-containing protein PICK1. *Neuron*, 22(1), 179–187. [https://doi.org/10.1016/S0896-6273\(00\)80689-3](https://doi.org/10.1016/S0896-6273(00)80689-3)
- Xiao, M. Y., Wasling, P., Hanse, E., & Gustafsson, B. (2004). Creation of AMPA-silent synapses in the neonatal hippocampus. *Nature Neuroscience*, 7(3), 236–243. <https://doi.org/10.1038/nn1196>
- Xu, D., Hopf, C., Reddy, R., Cho, R. W., Guo, L., Lanahan, A., ... Worley, P. (2003). Narp and NP1 form heterocomplexes that function in developmental and activity-dependent synaptic plasticity. *Neuron*, 39(3), 513–528. [https://doi.org/10.1016/S0896-6273\(03\)00463-X](https://doi.org/10.1016/S0896-6273(03)00463-X)
- Yan, D., Yamasaki, M., Straub, C., Watanabe, M., & Tomita, S. (2013). Homeostatic Control of Synaptic Transmission by Distinct Glutamate Receptors. *Neuron*, 78(4), 687–699. <https://doi.org/10.1016/J.NEURON.2013.02.031>
- Yan, Q. J., Asafo-Adjei, P. K., Arnold, H. M., Brown, R. E., & Bauchwitz, R. P. (2004). A phenotypic and molecular characterization of the fmr1-tm1Cgr Fragile X mouse. *Genes, Brain and Behavior*, 3(6), 337–359. <https://doi.org/10.1111/j.1601-183X.2004.00087.x>
- Yang, B., & Kumar, S. (2010). Nedd4 and Nedd4-2: Closely related ubiquitin-protein ligases with distinct physiological functions. *Cell Death and Differentiation*, 17(1), 68–77. <https://doi.org/10.1038/cdd.2009.84>
- Yang, Y., Wang, X. Bin, Frerking, M., & Zhou, Q. (2008). Delivery of AMPA receptors to perisynaptic sites precedes the full expression of long-term potentiation. *Proceedings of the National Academy of Sciences of the United States of America*, 105(32), 11388–11393. <https://doi.org/10.1073/pnas.0802978105>
- Yasaka, T., Tiong, S. Y. X., Hughes, D. I., Riddell, J. S., & Todd, A. J. (2010). Populations of inhibitory and excitatory interneurons in lamina II of the adult rat spinal dorsal horn revealed by a combined electrophysiological and anatomical approach. *Pain*, 151(2), 475–488. <https://doi.org/10.1016/j.pain.2010.08.008>
- Yates, F. A. (Frances A. (1966). *The art of memory*. University of Chicago Press.
- Yu, C. (2018). *Stoichiometry of AMPA receptors measured by single molecule imaging* (Albert-Ludwigs-Universität Freiburg im Breisgau). Retrieved from <https://d-nb.info/1213244684/34>
- Zhang, Y., Guo, O., Huo, Y., Wang, G., & Man, H. Y. (2018). Amyloid-β Induces AMPA Receptor Ubiquitination and Degradation in Primary Neurons and Human Brains of Alzheimer's Disease. *Journal of Alzheimer's Disease: JAD*, 62(4), 1789–1801. <https://doi.org/10.3233/JAD-170879>
- Zhao, H., Berger, A. J., Brown, P. H., Kumar, J., Balbo, A., May, C. A., ... Schuck, P. (2012). Analysis of high-affinity assembly for AMPA receptor amino-terminal domains. *Journal of General Physiology*, 139(5), 371–388. <https://doi.org/10.1085/jgp.201210770>
- Zhao, H., Lomash, S., Chittori, S., Glasser, C., Mayer, M. L., & Schuck, P. (2017). Preferential assembly of heteromeric

- kainate and AMPA receptor amino terminal domains. *ELife*, 6. <https://doi.org/10.7554/eLife.32056>
- Zhao, M. (2018). Human spatial representation: what we cannot learn from the studies of rodent navigation. *Journal of Neurophysiology*, 120(5), 2453–2465. <https://doi.org/10.1152/jn.00781.2017>
- Zhao, Y., Chen, S., Swensen, A. C., Qian, W. J., & Gouaux, E. (2019). Architecture and subunit arrangement of native AMPA receptors elucidated by cryo-EM. *Science*, 364(6438), 355–362. <https://doi.org/10.1126/science.aaw8250>
- Zho, W. M., You, J. L., Huang, C. C., & Hsu, K. Sen. (2002). The group I metabotropic glutamate receptor agonist (S)-3,5-dihydroxyphenylglycine induces a novel form of depotentiation in the CA1 region of the hippocampus. *Journal of Neuroscience*, 22(20), 8838–8849. <https://doi.org/10.1523/jneurosci.22-20-08838.2002>
- Zhou, Z., Liu, A., Xia, S., Leung, C., Qi, J., Meng, Y., ... Jia, Z. (2018). The C-terminal tails of endogenous GluA1 and GluA2 differentially contribute to hippocampal synaptic plasticity and learning. *Nature Neuroscience*, 21(1), 50–65. <https://doi.org/10.1038/s41593-017-0030-z>
- Zhu, J. J., Esteban, J. A., Hayashi, Y., & Malinow, R. (2000). Postnatal synaptic potentiation: Delivery of GluR4-containing AMPA receptors by spontaneous activity. *Nature Neuroscience*, 3(11), 1098–1106. <https://doi.org/10.1038/80614>
- Zhu, J., Lee, K. Y., Jewett, K. A., Man, H.-Y., Chung, H. J., & Tsai, N.-P. (2017). Epilepsy-associated gene Nedd4-2 mediates neuronal activity and seizure susceptibility through AMPA receptors. *PLOS Genetics*, 13(2), e1006634. <https://doi.org/10.1371/journal.pgen.1006634>
- Zi, Z., Zheng, Y., Rundell, A. E., & Klipp, E. (2008). SBML-SAT: A systems biology markup language (SBML) based sensitivity analysis tool. *BMC Bioinformatics*, 9(1), 342. <https://doi.org/10.1186/1471-2105-9-342>
- Zimmermann, A., Stauffacher, M., Langhans, W., & Würbel, H. (2001). Enrichment-dependent differences in novelty exploration in rats can be explained by habituation. *Behavioural Brain Research*, 121(1–2), 11–20. [https://doi.org/10.1016/S0166-4328\(00\)00377-6](https://doi.org/10.1016/S0166-4328(00)00377-6)
- Zolyomi, A., & Schmalz, M. (2017). Mining for Social Skills: Minecraft in Home and Therapy for Neurodiverse Youth. *Proceedings of the 50th Hawaii International Conference on System Sciences*. Retrieved from <http://hdl.handle.net/10125/41569>