



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.


Doctoral Thesis in Computer Science

Models of Corticostriatal Synaptic Plasticity and Plateau Potentials in Striatal Projection Neurons

DANIEL TRPEVSKI

KTH ROYAL INSTITUTE OF TECHNOLOGY

Stockholm, Sweden 2024



Models of Corticostriatal Synaptic Plasticity and Plateau Potentials in Striatal Projection Neurons

DANIEL TRPEVSKI

Academic Dissertation which, with due permission of the KTH Royal Institute of Technology and University of Edinburgh, is submitted for public defence for the Degree of Doctor of Philosophy on Wednesday the 25th September 2024, at 10:00 a.m. in F3, Lindstedtsvägen 26, Stockholm.

Doctoral Thesis in Computer Science
KTH Royal Institute of Technology
Stockholm, Sweden 2024
University of Edinburgh
Edinburgh, United Kingdom 2024

© Daniel Trpevski

ISBN 978-91-8040-984-1
TRITA-EECS-AVL-2024:59

Printed by: Universitetservice US-AB, Sweden 2024

To progress

Abstract

In this thesis we studied synaptic plasticity and neuronal computation in single striatal projection neurons (SPNs), which have a major role in goal-directed learning. Goal-directed or reward learning means to learn, based on sensory information from the body and the environment, to select actions out of all the behavioral repertoire that lead to obtaining a goal or reward (such as food or water). In mammals, all the behavioral motor repertoire is under constant, tonic inhibition, and the direct-pathway SPNs (dSPNs) select (disinhibit) goal-obtaining actions. The learning process is guided by the neuromodulator dopamine which signals the positive or negative outcome of an action. The synapses from cortical neurons onto the dSPNs, called corticostriatal synapses, are responsive to dopamine signals, and can strengthen and weaken based on the (positive or negative) action outcome. This promotes or discourages future actions in the same or similar sensory context.

Within a collaborative computational modeling effort, we studied the biochemical circuitry in the corticostriatal synapses with multiscale modeling and simulations. This circuitry in the corticostriatal synapses responds to neuromodulatory signals and controls the expression of synaptic plasticity. Multiscale modeling and simulations enable studying a system at multiple temporal and spatial scales, and integrating the results across the different scales. Based on molecular dynamics simulations of the enzyme which transduces extracellular neuromodulatory signals into an intracellular second messenger molecule, and Brownian dynamics simulations of regulator molecules binding to the enzyme, we constructed a kinetic model of the enzyme-based signal transduction network. The kinetic model showed that two co-occurring neuromodulatory signals, a dopamine peak and an

acetylcholine pause, are required to produce the second messenger and thus enable strengthening of corticostriatal synapses onto dSPNs, and that only the dopamine signal is not enough.

Next, we developed a local, calcium- and reward-dependent learning rule based on what is known about the biochemical circuitry of corticostriatal synapses onto dSPNs. We show that with this biologically-based learning rule, single SPNs can learn to solve the nonlinear feature binding problem (NFBP), a computationally hard problem representing the class of linearly nonseparable tasks. This result suggests that different, unrelated or partially related stimuli that require executing the same action to obtain a goal, can use the same SPNs responsible for selecting that action, and that a single SPN can reliably distinguish between similar stimuli.

The solution of the NFBP with the aforementioned learning rule relies on supralinear dendritic voltage elevations called plateau potentials. Experimentally, plateau potentials are all-or-none events, a property crucial for performing nonlinear computations required to solve the NFBP. However, computational models of plateau potentials often produce graded voltage elevations. We analyzed and compared existing plateau potential models, and found that long-lasting glutamate spillover in the extrasynaptic space robustly produces all-or-none plateau potentials by activating extrasynaptic N-methyl-D-aspartate (NMDA) glutamate receptors. This suggests that glutamate spillover may be a mechanism for generating all-or-none plateau potentials *in vivo*, as well.

In summary, the findings presented in this thesis advance our understanding of the role of single dSPNs in goal-directed learning, the biophysical mechanisms involved in performing their nonlinear computations, and the neuromodulatory signals necessary to produce synaptic strengthening and thus implement goal-directed learning.

Sammanfattning

I denna avhandling studerades synaptisk plasticitet samt förmågan till lokal beräkning i de striatala projektionsneuronens (SPNs) dendriter. Detta har direkt relevans för förståelsen av s.k. målstyrd inlärning. Målstyrd inlärning och belöningsinlärning innebär att man, baserat på sensorisk information från kroppen och omgivningen, lär sig selektera hur man skall agera/handla så att man uppnår ett mål eller erhåller en belöning (t.ex. mat eller vatten). Hos alla däggdjur är motoriska/exekutiva centra i hjärnstammen och thalamus under konstant tonisk inhibition via basala ganglierna, men när man aktiverar SPNs i den s.k. direkta vägen genom basala ganglierna (dSPN) så disinhiberas de motoriska centra som behövs för att initiera specifika målstyrda beteenden. Inlärningsprocessen för detta guidas av bl.a. dopamin, en neuromodulator som signalerar huruvida ett beteende ger positivt eller negativt resultat. Synapserna från kortikala projektionsneuron till dSPNs, s.k. kortikostriatala synapser, påverkas av dopamin och de kan förstärkas eller försvagas baserat på om dopaminsignalen signalerar ett positivt eller negativt utfall. Detta antingen främjar eller motverkar att man väljer samma beteende/handlingar under en liknande situation i framtiden.

Via samarbete med andra beräkningsbiologigrupper kunde vi studera de biokemiska signaleringsnätverken i de kortikostriatala synapserna med hjälp av multiskal modeller och simuleringar. Synapsens intracellulära biokemiska signaleringsnätverk kontrollerar synaptisk plasticitet och påverkas av neuromodulering. Genom att använda multiskalsimuleringar kunde vi studera systemet över multipla temporala och spatiala skalor, och integrera resultaten över de olika skalorna. Baserat på molekylärdynamiska simuleringar av det enzym som överför de extracellulära neuromodulatoriska signalerna till en intracellulär s.k. 2nd messenger molekyl, och Brownianska dynamiksimu-

leringar av de regulatoriska molekyler som binder till enzymet, kunde vi konstruera en kinetisk modell av detta signaleringsnätverk. Denna kinetiska modell predicerade att två samtidigt förekommande neuromodulatoriska signalförändringar, nämligen en pik i dopamin och en paus i acetylkolin, behövs för att aktivera 2nd messengersignaleringen på ett optimalt sätt, vilket i sin tur leder till att synapsen på dSPN förstärks. Simuleringarna predicerade också att en förändring i dopaminsignalen inte var tillräckligt.

Vi utvecklade därefter en förenklad lokal synaptisk plasticitetsregel baserat på vad som är känt om hur plasticiteten styrs i de kortikostriatala synapserna på dSPNs. Vi kunde visa att med denna inlärningsmodell för synaptisk plasticitet så kan en enskild SPN lära sig att lösa problemet med olinjär funktionsbindning (NFBP), ett beräkningsmässigt svårt problem som representerar klassen av linjärt icke separerbara problem. Detta resultat pekar på att olika, orelaterade eller delvis relaterade stimuli som kräver att man utför samma handling för att uppnå ett mål, kan använda samma SPN för att välja den handlingen, och att en enstaka SPN kan på ett tillförlitligt sätt skilja mellan liknande stimuli.

Lösningen av NFBP med den ovan nämnda inlärningsregeln bygger på supralinjära dendritiska förändringar av membranspänningen vilket kallas platåpotentialer. Experimentellt uppvisar platåpotentialer ett allt-eller-inget beteende, en egenskap som är avgörande för att utföra icke-linjära beräkningar som krävs för att lösa NFBP. Dock producerar beräkningsmodeller av platåpotentialer ofta graderade spänningshöjningar. Vi analyserade och jämförde befintliga modeller för platåpotentialer och fann att glutamatspillover (läckage) i det extrasynaptiska utrymmet producerar, på ett robust sätt, allt-eller-inget platåpotentialer genom att aktivera extrasynaptiska N-methyl-D-aspartat (NMDA) glutamatreceptorer. Detta tyder på att glutamatspillover kan vara en mekanism för att generera dessa allt-eller-inget platåpotentialer även *in vivo*.

Sammanfattningsvis fördjupar de resultat som presenteras i denna avhandling vår förståelse av enskilda dSPNs roll i målinriktat lärande, de biofysiska mekanismer som är involverade i att utföra neuronens icke-linjära beräkningar, och de neuromodulerande signaler som krävs för att producera synaptisk plasticitet och därmed implementera målinriktat lärande.

Contents

Abstract	v
Sammanfattning	vii
List of articles	xiii
List of abbreviations	xvi
Acknowledgements	xvii
Declaration	xix
 Part I:	
1 Introduction	1
1.1 What is the function of the nervous system?	1
1.2 Why does the nervous system need to be plastic?	2
1.3 Why study the nervous system? The state of scientific progress about the nervous system	4
1.4 The questions we chose to study in this thesis	5
1.4.1 Which (neuromodulatory) signals cause corticostriatal synaptic plasticity?	5
1.4.2 Solution to the nonlinear feature binding problem . . .	5
1.4.3 Glutamate spillover and its role in nonlinear compu- tation	6

2	Background Material	7
2.1	What is a neuron?	7
2.2	Ionic battery and action potentials	7
2.3	Synaptic potentials	12
2.4	Plateau potentials	14
2.5	Signal transduction through AC5	16
2.6	Synaptic plasticity	19
2.7	LTP	19
2.8	LTD	20
2.9	Goal-directed (reward) learning	22
2.10	Basal ganglia and goal-directed learning	25
2.11	Reward-related midbrain dopamine signals	27
2.12	Project-specific background material	29
2.12.1	AC5 signalling network	29
2.12.2	Synaptic plasticity rule	36
2.12.3	Synaptic clustering	40
2.12.4	The binding problem	40
2.12.5	The nonlinear feature binding problem	41
2.12.6	Glutamate spillover	42
3	Methods	47
3.1	Kinetic model of the AC5 signaling network	47
3.2	Measures of coincidence detection	50
3.2.1	Average catalytic rate	51
3.2.2	Synergy	52
3.3	Membrane voltage and Hodgkin-Huxley models of ion channels	54
3.4	Multicompartmental neuron models	57
3.5	Calcium dynamics	60
3.6	Synaptic models	62
3.6.1	Dual exponential synaptic model (difference of two ex- ponentials)	62
3.6.2	Saturating synapse models	64
3.7	Extrasynaptic glutamate receptors	65
3.8	Evoking plateau potentials	66
3.9	Synaptic plasticity	68
4	Present Investigation	73
4.1	Kinetic model of the AC5 signaling network	73

4.1.1	Mathematical analysis of the null model's regulatory motif	73
4.1.2	Comparison of the kinetic models with and without ternary complex	75
4.1.3	Discussion	79
4.2	Plasticity rule for solving the NFBP	81
4.2.1	Characterization of plateau potentials in the SPN model	82
4.2.2	A single run of the plasticity rule	84
4.2.3	Performance on the NFBP	88
4.2.4	The role of metaplasticity	90
4.2.5	Discussion	94
4.3	Glutamate Spillover	97
4.3.1	Thresholded and accumulative spillover	98
4.3.2	Long glutamate duration drives all-or-none plateau potentials	101
4.3.3	Low extrasynaptic glutamate concentration can produce all-or-none plateau potentials	102
4.3.4	Discussion	102
5	Future Outlook	107
	References	109

Part II: Research Articles

Article A

Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals

Article B

Local, calcium- and reward-based synaptic learning rule that enhances dendritic nonlinearities can solve the nonlinear feature binding problem

Article C

Glutamate spillover drives robust all-or-none dendritic plateau potentials—an in silico investigation using models of striatal projection neurons

List of articles

Articles included in the thesis:

- **Article A:** Bruce, N. J.*, Narzi, D.*, **Trpevski, D.***, van Keulen, S. C.*, Nair, A. G., R othlisberger, U., Wade, R. C., Carloni, P., and Hellgren Kotaleski, J. (2019). “Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals”. *PLOS Computational Biology* 15(10):e1007382. doi: 10.1371/journal.pcbi.1007382.
- **Article B:** Khodadadi, Z., **Trpevski, D.**, Lindroos, R., Hellgren Kotaleski, J. (2024) “Local, calcium- and reward-based synaptic learning rule that enhances dendritic nonlinearities can solve the nonlinear feature binding problem”. *eLife* 13:RP97274; doi: 10.7554/eLife.97274.1
- **Article C:** **Trpevski, D.**, Khodadadi, Z., Carannante, I., and Hellgren Kotaleski, J. (2023). “Glutamate spillover drives robust all-or-none dendritic plateau potentials—an *in silico* investigation using models of striatal projection neurons”. *Front. Cell. Neurosci.* 17, 1196182. doi: 10.3389/fncel.2023.1196182

* – equal contribution

Articles not included in the thesis:

- **Article D:** Santos, J. P. G., Pajo, K., **Trpevski, D.**, Stepaniuk, A., Eriksson, O., Nair, A. G., Keller, D., Hellgren Kotaleski, J., Kramer, A. (2022). “A modular workflow for model building, analysis, and parameter estimation in systems biology and neuroscience”. *Neuroinformatics* 20: 241–259. doi: 10.1007/s12021-021-09546-3
- **Article E:** van Keulen, S. C., Martin, J., Colizzi, F., Frezza, E., **Trpevski, D.**, Cirauqui Diaz, N., Vidossich, P., Röthlisberger, U., Hellgren Kotaleski, J., Wade, R. C., Carloni, P. (2023). “Multiscale molecular simulations to investigate adenylyl cyclase-based signaling in the brain”. *WIREs Computational Science* 13(1): e1623. doi: 10.1002/wcms.1623
- **Article F:** **Trpevski, D.***, Khodadadi, Z.*, Hellgren Kotaleski, J. (2024). “An inhibitory plasticity rule for control of neuronal firing rate and dendritic compartmentalization”. Manuscript

* – equal contribution

List of abbreviations

2-AG	2-arachidonoylglycerol
AC	adenylyl cyclase
ACh	acetylcholine
AMPA	α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid
AMPA R	AMPA receptor
ATP	adenosine triphosphate
cAMP	cyclic adenosine monophosphate
CB1R	type 1 cannabinoid receptor
Da	dopamine
DGL α	diacylglycerol lipase α
dSPN	direct-pathway striatal projection neuron
GDP	guanosine diphosphate
GPCR	G-protein-coupled receptor
GTP	guanosine triphosphate
IP ₃	inositol triphosphate
iSPN	indirect-pathway striatal projection neuron
LTD	long-term depression
LTP	long-term potentiation
mGluR	metabotropic glutamate receptor
NFBP	nonlinear feature binding problem
NMDA	N-methyl D-aspartate
NMDAR	NMDA receptor
PIP2	phosphatidylinositol 4,5-bisphosphate
PKA	protein kinase A
PP1	protein phosphatase 1
PP2A	protein phosphatase 2A
PP2B	protein phosphatase 2B
PPi	pyrophosphate
RGS	regulator of G protein signaling
RyR	ryanodine receptor
SPN	striatal projection neuron

Acknowledgements

I would like to thank Jeanette, my supervisor, for her enthusiastic encouragement, unending patience, and wise pragmatism. I would also like to thank Matthias Hennig, my co-supervisor, for his enthusiasm and guidance.

Alex Kozlov, thank you for your support in all questions technical, all your provided models and scripts for running them on clusters. Olivia Eriksson, thank you for all the discussions and sharing your knowledge into many aspects of model building. Jovana Belic, I would simply not have managed to find my way through many of KTH's requirements for orderly pursuing the PhD programme, and would have missed out on many things had it not been for you! Your advice and experience contributed to my PhD being a much smoother ride!

Much of the research was done in collaboration with my colleagues Anu, Robert, Mohammad, Ilaria and Zahra. It was great that through joint efforts we reached new conclusions and findings! There were many other attempts at pursuing more projects, and much discussion with all other members of the lab: Kadri, João, Parul, Andrei, Mikael, Jan, Johannes, and Kai. Ilaria, your help with navigating the thesis submission process was tremendously valuable!

One of the studies in this thesis was a multidisciplinary collaboration between several project partners, achieving a synergistic outcome that would have been unattainable if the individual research teams had worked separately. I would like to thank Neil Bruce, Siri Camee van Keulen, Rebecca Wade, Daniele Narzi, and Paolo Carloni for all their help, discussions and sharing specialized knowledge which shaped the research carried out in our branch of the project. I would also like to thank Juliette Martin, Francesco

Colizzi, Elisa Frezza, Nuria Cirauqui Diaz, Pietro Vidossich and Ursula Röthlisberger.

I thank Ana Kalajdjieva for professionally illustrating the mouse brain. The illustration appeared in the article reporting the results of one of the studies, and is also used in this thesis.

A lot of the research done in this thesis would have remained as ideas, concepts, or demonstrations of working examples without the computing clusters at PDC KTH and the Fenix Infrastructure resources.

SciLifeLab, KTH Royal Institute of Technology, and the University of Edinburgh provided a sound environment for research with excellent potential for collaboration.

Finally, the establishing of the Erasmus Mundus PhD programmes by the European Commission has been an extraordinary initiative. EuroSPIN, the European Study Programme in Neuroinformatics, enabled a terrific environment across the participating European and Indian institutions to pursue PhD studies in.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified. The included publications are joint collaborations, the contributions to which are listed on page xiii.

Stockholm, 2024

1 Introduction

1.1 What is the function of the nervous system?

The nervous system is an organ system specialized in information processing, which is present only in animals. From the fact that the outgoing fibers of the nervous system terminate on muscles and glands, one can reason that its role is to control movement and regulate internal bodily processes. Indeed, it is thought that nervous systems have evolved in heterotrophic multicellular organisms (those that do not produce their own food through, for example, photosynthesis) to control opposing contractile tissue/cells and to coordinate the movement of a multicellular body (Bucher and Anderson 2015; Keijzer et al. 2013), perhaps out of the need to move in the environment to obtain food. Movement then implies the animal will experience a changing environment, hence a need arises for senses that can track changes in the world, as well as for circuitry that processes those sensory stimuli (Greenspan 2007). Whatever the evolutionary history of nervous systems turns out to be, their function is to sense the environment, both external and internal, and based on the information obtained by processing these sensory stimuli, to control movement and regulate bodily processes as necessary. At least in humans, the nervous system is the seat of consciousness and awareness, creates feelings and thoughts, creates abstract models of the world and based on them and on memories and experiences, is able to make predictions or imagine the outcome of future events.

The basic functional unit in the nervous system that processes information and triggers a behavioral response is a neural circuit. One of the simplest examples for a neural circuit is the myotactic (“knee-jerk”) reflex (Fig. 1.1). This circuit, and all others, consists of three basic components:

afferent neurons, efferent neurons, and local circuit neurons, or interneurons (Purves et al. 2012). Afferent neurons receive sensory information from the sensory organs in the muscles and carry it to the central nervous system, in this case to several types of central neurons in the spinal cord. One type are the motor neurons, efferent neurons which carry information towards the effectors, in this case the leg muscles. Another type are the interneurons, or local circuit neurons, which make connections with the motor neurons projecting to the flexor muscles of the leg, and can thus modulate their output. The function of the myotactic reflex is to quickly extend the leg after tripping on a stone, for example, and thus prevent a fall. This is achieved by simultaneously contracting the extensor muscle and relaxing the flexor muscle, and the processing required for the neural circuit to achieve this is very simple. The synapses, or connections, from the sensory afferent neurons to the effector motor neurons projecting to the extensor muscles and to the interneurons are excitatory, while the synapses from the interneurons to the flexor motor neurons are inhibitory. Tripping on a stone while walking causes a stretch in the tendon which, in turn stretches the sensory organs in the extensor muscle, stimulating the sensory neurons. They, in turn, will simultaneously stimulate the motor neurons to the extensor muscles, causing them to contract, and stimulate the inhibitory interneurons, which will inhibit the flexor motor neurons, relaxing the flexors. Thus, with one monosynaptic and one disynaptic processing route, a circuit for automatic (reflexive) and innate behavior is implemented, which helps to prevent falls.

1.2 Why does the nervous system need to be plastic?

Some important parts of behavior are wired in the brain architecture. However, to adapt to a changing environment, the brain needs to be able to remodel, to be malleable. The most drastic example of brain remodeling can be seen in complete metamorphosis, present in 88% of insects, the most diverse group of multicellular organisms in terms of number of species. The larva and the adult frequently have completely different bodies and behaviors, requiring appropriate changes in the nervous system to control the new body and its behavior in the environment (Truman et al. 2023).

In humans and other vertebrates, despite lacking such striking developmental changes, the nervous system is still highly plastic, enabling the

1.2. WHY DOES THE NERVOUS SYSTEM NEED TO BE PLASTIC? 3




Figure 1.1: The neural circuit for the myotactic reflex. Figure reused from Purves et al. 2012 with permission from Oxford University Press. © 2012

organism to learn, meaning to store information about the world and to associate cause and effect, especially one's own actions to outcomes in the world.

The latter type of learning is important for obtaining desired “goals”, such as food, water, or finding a mate. Neural circuits, much larger than the one for the myotactic reflex, exist in the brain for such goal-directed, or reward learning. These circuits enable an organism to learn, based on sensory information from the environment and the body, to select actions out of all the behavioral repertoire (all possible actions) which lead to obtaining the goal (reward). In the real world this might mean learning the association between the appearance of a plant and the fruits it produces, and that foods with certain color and smell are edible, while others should be avoided. In simplified experimental settings, whose aim is to be simple enough to isolate

and identify the mechanisms underlying a phenomenon, this might mean learning that pressing a button only when a green light and a sound is on, will lead to food delivery, but not when a red light with a different sound is on. Also, the appropriate action towards the button might be important, i.e. that it is important to press, instead of, for example, bite or smell the button.

1.3 Why study the nervous system? The state of scientific progress about the nervous system

The human nervous system determines so much of the human condition, making it one of the most intriguing subjects to study for providing at least a mechanistic explanation of how it achieves its feats.

A lot about the nervous system is known. Its structure and detailed anatomy, organ composition, and direction of information flow, as well as the roles of various brain regions have been established. Additionally, many mechanistic details such as molecules responsible for sensory signal transduction, action potential generation, and chemical signal transmission between neurons have been elucidated, and an astounding diversity of cells has been documented. The functionality of local microcircuits in some brain regions is beginning to be understood.

Much more about the nervous system is unknown. One of the biggest questions in neuroscience is how events at the subcellular and cellular levels interact and give rise to organismal behaviors. To answer this question, a huge amount of the molecular machinery awaits to be discovered, the question of how transduced sensory information is integrated in early nuclei, as well as in layers beyond sensory cortices is still far from completely understood, and the way neurons encode information with their firing activity is not fully answered. Further, how precise neuronal activity affects, and is integrated into, functional neural circuits is an active area of research.

With regard to nervous system plasticity, the foremost candidate mechanism for implementing learning and memory is synaptic plasticity. It is hypothesised that information is stored in neural circuits in the pattern of synaptic strengths. In the neural circuits for goal-directed learning, the striatal projection neurons (SPNs) are considered to play a major role: to select appropriate actions for a goal and inhibit inappropriate ones. They receive processed sensory and contextual information through synapses from

cortical neurons, and modifying these synapses is thought to represent the linking of relevant sensory information to appropriate actions¹. Synaptic plasticity of the corticostriatal synapses is the subject of this thesis.

1.4 The questions we chose to study in this thesis

Although neural circuits are the unit for expressing behavior, we focus on understanding the function of a single SPN within a circuit. Phrased generally, we studied what kind of functionality and computations a single SPN can perform with its corticostriatal synaptic plasticity mechanisms. The thesis addresses three research questions:

1.4.1 Which (neuromodulatory) signals cause corticostriatal synaptic plasticity?

Synaptic plasticity in corticostriatal synapses relies, among other things, on the levels of the modulatory neurotransmitter dopamine, which control the activity of the enzyme adenylyl cyclase 5 (AC5) expressed in these synapses. We studied the regulation of the enzyme AC5 by transient neuromodulatory signals. The results suggested that two coincident neuromodulatory signals are required for strengthening corticostriatal synapses, and that only dopamine signaling is not enough. In the type of SPNs that we studied – the direct-pathway SPNs (dSPNs) – these signals are a dopamine peak, and an acetylcholine pause.

1.4.2 Solution to the nonlinear feature binding problem

We chose a well-known, computationally hard problem to study the computational capability of single SPNs. We show that dSPNs can solve the nonlinear feature binding problem with a learning rule based on the available synaptic plasticity mechanisms in their corticostriatal synapses. This implies that the same SPNs involved in selecting an action can be used in different situations that warrant the execution of that same action, and moreover, that single SPNs can reliably distinguish between similar situations in which the action should not be executed.

¹Around 60 % of the synaptic inputs to the SPNs are from the cortex. SPNs also receive inputs from thalamic neurons, but these are not treated in this thesis in any specific detail.

1.4.3 Glutamate spillover and its role in nonlinear computation

The solution of the nonlinear feature binding problem that we proposed relies on long-lasting excitations of the neuron, called plateau potentials. We studied how plateau potentials can be robustly evoked by excitatory synapses and found that employing glutamate spillover could be a key requirement for robustly producing plateau potentials, even when the necessary receptors to evoke them might have different characteristics.

2 Background Material

2.1 What is a neuron?

Neurons are the information-processing cells in the nervous system. Apart from neurons, the nervous system contains glial cells of several types, of which astrocytes are increasingly being recognized as important for information transmission at synapses and information processing. In contrast to the classical image of a spherical eukaryotic cell, neurons are usually highly branched, with a starlike appearance (Fig. 2.1). They are not the only cells in the body with this shape, but they are the only starlike cells that use electrical signals to rapidly communicate between themselves. Their shape allows them to make connections called synapses with many individual neurons, which allow for point-to-point unidirectional communication.

In this chapter we will describe the physical basis for the electrical excitability of neurons, and the generation of the fast electrical signals for interneuronal communication, the action potentials (APs). We will further describe the molecular machinery which controls the strength of synaptic connections between neurons, specifically the corticostriatal synapses. Lastly, we will give a brief description of reward learning, the type of learning process which is most frequently associated with these synapses.

2.2 Ionic battery and action potentials

All cells have electrical voltage across their membranes. This voltage is a result of the different concentration of ions, which in neurons are mainly Na^+ and K^+ , in the cytoplasm and in the extracellular environment (Fig. 2.2). Regardless of the concentration difference, both sides of the membrane

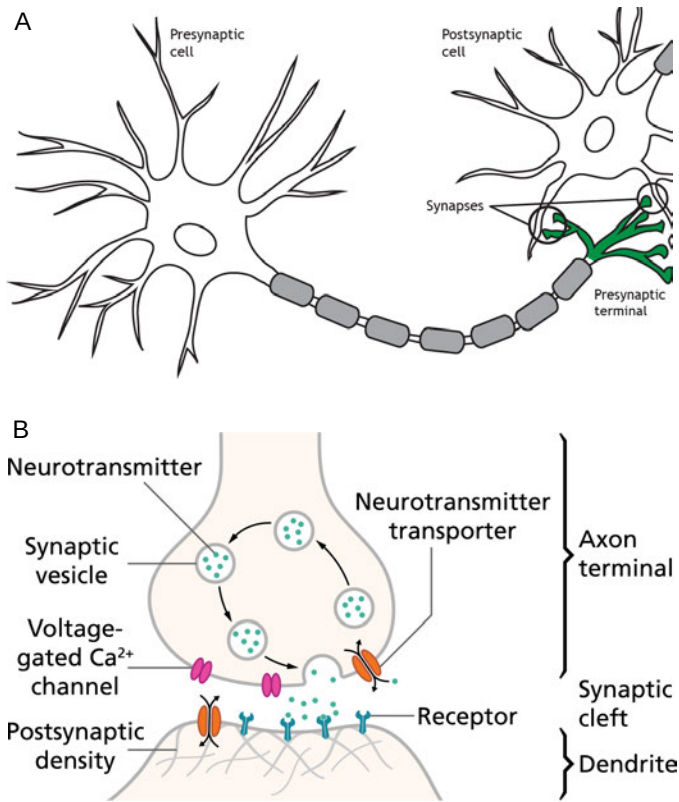


Figure 2.1: Illustration of a neuron and a synapse. (A) A neuron makes synapses with another neuron. Synapses are typically made on thin cylindrical structures called dendrites, which branch out from the soma, the spherical part which contains the nucleus and major cellular organelles. (B) A synapse consists of a presynaptic and postsynaptic terminal, and between them is a small space called the synaptic cleft. Synaptic vesicles filled with neurotransmitter sit in a vesicle pool in the presynaptic terminal. The AP, arriving from the axon to the presynaptic terminal, triggers the release of neurotransmitter from the vesicles into the synaptic cleft. Neurotransmitters bind to receptors on the postsynaptic terminal, which then either let current in the postsynaptic cell or initiate intracellular signaling. The effects of both processes are described later in this chapter. Image in (A) by Casey Henley. Image in (B) by Thomas Splettstösser. CC BY-SA 4.0 licence. <https://creativecommons.org/licenses/by-nc-sa/4.0/>

are electrically neutral. With regard to the diffusing ions, the lipid cellular membrane has two properties. On one hand it acts as a capacitor: ions can associate closely to its surface, leading to an accumulation of charge on the membrane. On the other hand, it is not a complete barrier, i.e. it is permeable to some ions (we will come back to the nature of the membrane permeability at the end of this section). When ions from the more concentrated side of the membrane permeate to the less concentrated side due to the concentration gradient, they associate with the membrane on the less concentrated side very quickly (i.e. charge the membrane capacitor) causing a local excess of electric charge (Fig. 2.2B, left). At the same time, their rapid membrane permeation causes an opposite local shortage of electric charge on the more concentrated side of the membrane, due to the sudden local lack of ions which had permeated the membrane. This local charge difference at the permeation site creates a difference in electric potentials between both membrane sides, thus giving rise to an electric field across the membrane, whose strength corresponds to some characteristics of the ion (Fig. 2.2B, right). The difference in electric potentials is called the ionic battery. This is how the electric voltage measured across a cell's membrane arises. At rest, when no stimuli arrive to the neuron, the electric voltage is called the resting membrane voltage.

As the ions flow across the membrane down their concentration gradient, the developing electric field opposes their movement because it is directed oppositely from the direction of positive current flow. At the point of electrochemical equilibrium, the electric field is strong enough to completely counter the concentration gradient, and no net flow of ions across the membrane occurs. The value of the membrane voltage in this situation is called the reversal voltage (or also, reversal or Nernst potential), and can be calculated based on the resting intra- and extracellular concentrations of the ion. So, both the ionic concentrations and the membrane voltage determine the amount of ionic current flow through the membrane. The direction of ionic current is generally determined by the ionic concentrations, but experimentally applying voltage with a magnitude larger than an ion's reversal voltage can change the direction of that ionic current through across the membrane.

As mentioned above, all cells exhibit electric voltage across their membranes due to the different concentration of ions on the two sides of the membrane and due to the membrane permeability. However, not all cells are electrically excitable. Electrical excitability is the property of cells to exhibit transient variations in the membrane voltage due to an electrical or




Figure 2.2: The ionic battery – basis of electric voltage across the membrane. (A) A membrane permeable only to K^+ separates two compartments of equal K^+ concentration. There is no net flux of K^+ across the membrane and no voltage. (B) One compartment initially contains 10 times higher K^+ concentration than the other. Initially there is no voltage. After K^+ starts diffusing through the membrane to the less concentrated compartment, the local excess of K^+ on the less concentrated side together with the local lack of K^+ on the more concentrated side creates an electric field across the membrane. Figure reused from Purves et al. 2012 with permission from Oxford University Press. © 2012

other stimulus (such as light, mechanical or chemical stimulation). Electrical excitability is due to proteins in the membrane which have opening and closing pores that are selectively permeable to some ions. There is an enormous family of such proteins, called ion channels, which are frequently selective for only one major ion (Fig. 2.3). Let's consider the situation when the membrane possesses a K^+ channel. When this protein is in the membrane, the mechanism of the ionic battery we described above will occur only for the K^+ ion flowing across the channel pore, i.e. this channel will

create an ionic battery for K^+ since it is only permeable to K^+ . The more different ion channel proteins are present in the membrane, the more ionic batteries appear across that membrane. The ion channel proteins most frequently have an additional property: they are not always open. For a huge number of them, the opening of the pore depends on the current value of the membrane voltage; because of this they are called voltage-gated ion channels (Fig. 2.3A-D). Each of them has a region of voltage where it will open, and each of them opens and closes with its own speed, which also depends on the voltage. If we apply an electrical stimulus to the membrane, causing some ion channels to open, this means that the ionic batteries “appear and disappear” across the membrane as the ion channels open and close. With their transient presence, the corresponding membrane currents are also transiently present, creating fast and transient voltage changes across the membrane. In this way, various kinds of electrical signals are possible across the membrane, depending on which ion channels are expressed. If the same ion channels as those at the site of electrical stimulation are distributed in the dendritic and axonal membrane, the electrical signals can propagate unattenuated through the neuron.




Figure 2.3: Voltage-gated and ligand-gated channels. Figure reused from Purves et al. 2012 with permission from Oxford University Press. © 2012

The most famous electrical signal is the action potential (Fig. 2.4A). It is a fast and short spike in the membrane voltage, is created by Na^+ and K^+ channels in the axon hillock, the origin of the axon, and is essential for interneuronal communication. It is an all-or-none electrical event, meaning that it appears after a threshold level of the somatic voltage is reached, and it has an almost constant amplitude irrespective of the strength of the

evoking stimulus (i.e., it is not graded in amplitude). Because Na^+ and K^+ channels are distributed in the axon, the action potential propagates down the axon to all the outgoing synapses the neuron makes. Its role is to initiate the chemical signaling at the synapses to the postsynaptic neurons to which the axon projects by triggering the release of neurotransmitter. Neurotransmitter at the synapse requires elevated membrane voltage, provided usually by trains of action potentials, which opens voltage-gated calcium channels on the presynaptic membrane, after which calcium starts the process of neurotransmitter release (Fig. 2.1B).

The membrane permeability which, together with the different ionic concentrations on each side of the membrane, gives rise to the resting membrane voltage, has several origins. Firstly, some K^+ ion channels are open at resting membrane voltages, and K^+ current flowing outward keeps the membrane voltage low (these channels close at elevated membrane potentials). Secondly, there is a large family of K^+ ion channels which are voltage-independent, and whose pore opening is governed by intracellular ligands (Fig. 2.3G). Their opening and closing dynamics is much slower than voltage-gated channels, and they may remain open for very long periods of time, because of which they are called leak channels (Goldstein et al. 2001). Finally, pure lipid bilayer membranes exhibit thermal fluctuations which also form ion-conducting pores through the bilayer membrane, simply as a consequence of the second law of thermodynamics. Current passed through these pores resembles current through ion channel proteins (Blicher and Heimburg 2013; Kaufmann and Silman 1983; Hanke et al. 1989). The exact contributions of each of these sources to the resting membrane voltage vary for different neuron types, and typically, only the contribution of protein ion channels is considered.

2.3 Synaptic potentials

Not all ion channels are voltage-gated. Many ion channels require a molecule to be bound to them, a ligand, in order to open (Fig. 2.3E–F). A typical example are the ion channels in the membrane of the postsynaptic site whose ligands are frequently neurotransmitters released from the presynaptic terminal. After its presynaptic release, a neurotransmitter binds to its ion channel receptors on the postsynaptic side, opening them and allowing ion influx to raise the postsynaptic membrane voltage locally at the spine or




Figure 2.4: Electrical signals in neurons. (A) A train of action potentials recorded in vivo from a rat SPN. The action potential is the fast upward spike triggered after the somatic voltage threshold is reached. Figure is reused from Mahon et al. (2006). © 2006 Society for Neuroscience. (B) Synaptic potentials evoked with glutamate uncaging and recorded from the soma of an SPN in a slice preparation from a rat. The red trace is the synaptic potential evoked from a single synaptic input, and the black trace is evoked from two synaptic inputs spaced apart 50 ms. Figure is reused from Carter et al. (2007). © 2007 Society for Neuroscience. (C) Dendritic plateau potentials evoked with glutamate iontophoresis (extracellular current delivering glutamate) in a single dendrite, and measured at the soma. Current steps increase in equal increments between traces. Green trace indicates the last response before plateaus appear and the red trace is the first supralinear, plateau response. Figure reused from Oikonomou et al. (2014). © 2014 Oikonomou et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

dendrite where the synapse is positioned. This generates a local synaptic potential, which is called excitatory if it raises the voltage, and inhibitory if it lowers it (Fig. 2.4B). Such local voltage changes propagate in all directions away from their point of origin as the ions that caused them travel through the dendritic cytoplasm. The changes in voltage attenuate with distance

because the ion currents that generate them leak out of the membrane.

We will consider three main ligand-gated neurotransmitter receptors in this thesis. Two are receptors for the excitatory neurotransmitter glutamate. One of them, the AMPA receptor (AMPA) is responsible for fast and short excitations, and the other, the NMDA receptor (NMDAR) gives rise to slow and long excitations. The third receptor is the GABA receptor, for the inhibitory neurotransmitter GABA.

When enough of the synapses get activated, the propagating voltage changes will raise the somatic membrane voltage to the threshold level required to initiate the action potential at the axon hillock. The threshold for action potential initiation and the need to integrate excitatory potentials from many synapses give the opportunity to implement a computational logic deciding when a neuron gets activated, i.e. when an action potential is discharged. Moreover, synapses can excite the postsynaptic neuron with varying strength, giving the possibility for adaptable computational logic. Stronger synapses cause larger synaptic potentials in the postsynaptic neuron, achieved by passing more current, and weaker synapses cause smaller synaptic potentials. In SPNs, whose resting voltage is very low, typically hundreds of synapses distributed across the dendrites need to be coactive to cause somatic spiking. However, when 10–20 synapses located on a small dendritic segment of 10–20 μm are coactivated within a short interval (for example, up to 50 ms), large dendritic excitations occur which can cause somatic spiking. These supralinear dendritic voltage events are described in the next section.

2.4 Plateau potentials

Dendrites also have supralinear, all-or-none voltage events which are called dendritic spikes. They can be generated by several mechanisms, and of interest to us are the dendritic spikes generated by NMDARs (Antic et al. 2010). They are evoked by clustered synaptic inputs, localized within a dendritic stretch of approximately 20 μm . Sometimes a distinction is made in the literature between an NMDA spike and a plateau potential (Oikonomou et al. 2014). Both events have supralinear voltage amplitudes and require synchronous activation of more than one synapse containing NMDARs. NMDA spikes are shorter (last 15–50 ms), their amplitude is smaller, can be elicited with smaller stimulus intensities, and evoke NMDAR Ca^{2+} influx only lo-

cally, where the stimulation occurs. Plateau potentials, on the other hand, are longer-lasting (above 100 ms, typically 200–500 ms), their amplitude is larger and may drive somatic spiking by themselves. They require the activation of a larger cluster of synapses and evoke Ca^{2+} influx not only from the engaged NMDARs that mediate them, but also from voltage-gated $^{2+}$ channels along the dendritic branch, since the voltage elevation spreads throughout much of the originating branch. Plateau potentials occur in neurons of multiple brain regions, SPNs among them, and because of their all-or-none nature, also allow the implementation of specific computations, most notably to detect the coincident activation of a synaptic cluster. Additionally, NMDA spikes have been found to be both necessary and sufficient to trigger synaptic plasticity, at least in the collateral synapses between CA3 hippocampal neurons (Brandalise et al. 2016).

NMDA spikes and plateau potentials are generated through a positive feedback loop in the activation of NMDARs. The NMDARs are ligand-gated ion channels activated by glutamate and glycine that are additionally blocked by Mg^{2+} ions at resting membrane voltages. At a given extracellular Mg^{2+} concentration, $[\text{Mg}^{2+}]$, the amount of blockage depends on the local voltage. This voltage dependence is a sigmoidal nonlinearity, where higher voltage provides rapid removal of Mg^{2+} from the NMDAR pore. The positive feedback loop is initiated when the dendritic voltage reaches a threshold value, usually around -50 mV, which partially alleviates the Mg^{2+} block on NMDARs. This causes the flow of an inward current through the unblocked NMDARs, raising the voltage in a cycle that progressively alleviates more of the Mg^{2+} block in remaining blocked NMDARs and allows for more inward current. In this way, a sustained voltage elevation is produced, which lasts as long as there is glutamate and glycine to bind to the NMDARs. The sigmoidal shape of the voltage nonlinearity is a crucial ingredient for the all-or-none behavior of the plateau potentials since, for linear increases in voltage, it provides a “sudden” relief of Mg^{2+} ions from the NMDARs. The Mg^{2+} block is described with a gating function, which describes both the probability of a single unblocked NMDAR at a given voltage and $[\text{Mg}^{2+}]$, as well as the macroscopic NMDA current sustained by the fraction of an unblocked NMDAR population:

$$g(v) = \frac{1}{1 + \eta [\text{Mg}^{2+}] e^{-\alpha v}}, \quad (2.1)$$




Figure 2.5: The gating function in Eq. 2.1. (A) The parameter α determines the steepness of the gating function, with higher values representing steeper curves. Plots corresponding to the values $\alpha \in \{0.04, 0.05, 0.062, 0.07, 0.08\}$ are shown, and darker colors indicate higher values. $\eta = 0.38$. (B) The parameter η determines the shift of the gating function along the x axis. Plots corresponding to the values $\eta \in \{e^{-2}, 0.38, 1, e, e^2\}$ are shown, with darker colors indicating higher values. $\alpha = 0.062$. In both panels, the thick lines indicate the gating function with the parameters originally reported by Jahr and Stevens (1990). $[\text{Mg}^{2+}] = 1 \text{ mM}$.

where α is the steepness of the curve, η is related to the voltage value $v_{1/2}$ where half of the receptors are free of Mg^{2+} according to $v_{1/2} = \alpha \ln(\eta [\text{Mg}^{2+}])$, $[\text{Mg}^{2+}]$ is the concentration of extracellular Mg^{2+} in millimolar (mM) and v is the membrane potential in millivolts (mV) (Jahr and Stevens 1990). Figure 2.5 shows examples of the gating function where the α and η parameters are varied.

2.5 Signal transduction through AC5

The strength of a synapse is controlled by molecular circuitry within the synapse, called a signaling network or a signal transduction network. In fact, signaling networks are present in all cells. They represent the cell's sensory system for extracellular chemicals, light, temperature, mechanical force and voltage, and also implement the decision functions which integrate the extracellular signals and trigger the appropriate cellular response to an environmental event. A now classical example of signal transduction is bacterial chemotaxis. It allows bacteria to respond to changes in the chemical composition of their environment by moving toward increased concentrations of an attractant molecule or away from a repellent molecule. The signaling

network in *E. coli* starts with glucose receptors in the cell membrane, and effects a response on the molecular motors with which a bacterium moves. In synapses, the signaling networks start with neurotransmitter receptors, ligand-gated and/or voltage-gated ion channels in the membrane of the post-synaptic site, and the responses of the signaling networks trigger synaptic plasticity. In this section we provide the background to detecting and transducing neuromodulatory signals in corticostriatal synapses. Neuromodulators are a type of neurotransmitter that bind to receptors which are not ion channels, i.e. they do not pass current inside the cell, but they activate various signaling molecules intracellularly (they are called metabotropic receptors). This is the most common form of intercellular signaling, described next.

The cellular membrane isolates the inside of the cell from the outside, so for the cell to be able to receive information about extracellular events, outside signals need to somehow “cross” the cellular membrane. Chemical signals, which are molecules interacting with the cell, do this in two ways. Some are able to diffuse through the membrane and bind to their intracellular receptor, activating a signaling network. A vast majority of chemical signals bind to membrane receptors, which transduce the signal by interacting with intracellular molecules when the outside chemical ligand is bound. Membrane receptors are highly specific, meaning they can only bind a single molecular species, or several very closely related molecular species. A large class of such receptors are the G protein-coupled receptors (GPCR), which when activated by a ligand can interact with intracellular G proteins (or GTP-binding proteins) (Fig. 2.6). In the corticostriatal synapse, the G proteins are heterotrimers composed of three subunits, α , β and γ . When the GPCR activates a G protein, the α subunit binds to GTP and dissociates from the heterotrimer, and represents the active form of the G protein. It can interact with various molecules, most commonly enzymes, further relaying the extracellular signal within the signaling network. In summary, the binding of a ligand to a GPCR releases an active $G\alpha$ subunit, which represents the transduced intracellular signal corresponding to the extracellular ligand.

The length of the transduced intracellular signal in the form of active $G\alpha$ subunits is regulated by proteins which deactivate the $G\alpha$ subunits, called regulators of G protein signaling, or RGS proteins. When RGS proteins bind to a $G\alpha$ subunit, they promote its intrinsic GTPase activity (GTP hydrolysis to GDP), and GTP hydrolysis to GDP is what inactivates an




Figure 2.6: Signal transduction networks through GPCRs. Figure reused from Purves et al. (2012) with permission from Oxford University Press. © 2012

active $G\alpha$ subunit. Many RGS proteins have been identified in the mouse, of which RGS2, RGS4, RGS9 and RGS14 have been found to be expressed in the striatum (Gold et al. 1997; Lerner and Kreitzer 2012; Xie et al. 2012).

In the corticostriatal synapses, the $G\alpha$ subunit binds to adenylyl cyclase 5 (AC5), an enzyme which catalyzes the conversion of ATP (adenosine triphosphate) to cAMP (cyclic adenosine monophosphate), a common intracellular signaling molecule (a second messenger). cAMP then activates protein kinase A (PKA), which in the SPNs can effect changes on a variety of molecules, such as increasing the conductance of ion channels (thus strengthening a synapse), modulating the activity of further downstream signaling molecules, or activating transcription factors, molecules that initiate gene transcription. Controlling the levels of the second messenger cAMP is the primary checkpoint in controlling whether an active corticostriatal synapse will undergo LTP.

2.6 Synaptic plasticity

In this section we present the signaling networks (or molecular circuits) which control synaptic long-term potentiation (LTP) and long-term depression (LTD), which refer to strengthening and weakening of the synapse which lasts for 30 minutes to several hours.

As mentioned in section 2.3, for an excitatory synapse to cause a larger synaptic potential, more current needs to flow in the postsynaptic neuron. This can happen in several ways, and in our case synaptic strengthening means that more glutamate receptors have been inserted in the postsynaptic membrane, or that the existing receptors are modified so that they can pass through more current (usually via receptor phosphorylation). Conversely, to cause a smaller synaptic potential, less current needs to flow in the postsynaptic neuron. This can happen through removal of glutamate receptors from the postsynaptic membrane or through dephosphorylating them, but also by reducing the amount of glutamate that is released from the presynaptic side. In the corticostriatal synapses, synaptic weakening is expressed presynaptically.

2.7 LTP

In corticostriatal synapses, the signaling networks that decide whether to strengthen the synapse can be viewed as having two branches. One branch responds to synaptic activity and the other responds to neuromodulatory signals, such as dopamine. Synaptic activity is usually indicated intracellularly by the amount of Ca^{2+} ion influx through ligand- and/or voltage-gated ion channels. The Ca^{2+} ions affect one branch of the signaling network in corticostriatal synapses. Neuromodulators bind to metabotropic receptors and activate a separate branch of the signaling network.

The signaling network which controls LTP in dSPNs is shown in Fig. 2.7. Ca^{2+} influx from NMDA and voltage-gated Ca^{2+} channels activates the enzyme calmodulin-dependent kinase II (CaMKII), critical for phosphorylating the AMPA receptors and potentiating the synapse (Herring and Nicoll 2016). Dopamine, on the other hand, activates the enzyme AC5 via the D_1 dopamine receptor (a GPCR), and ultimately results in activated PKA, which is necessary for LTP (Yagishita et al. 2014; Cui et al. 2016). At the time we conducted the research projects, it had been proposed that acetyl-

choline signaling might also be involved in controlling the activation of AC5. We studied the regulation of AC5 by both neuromodulatory signals, and for this reason the acetylcholine M_4 receptor is also included in the network in Fig. 2.7. The coincident activation of the neuromodulatory and Ca^{2+} branch is detected by the protein DARPP-32, which controls the balance of the proteins which phosphorylate and dephosphorylate the AMPARs, and can thus promote increased synaptic strength or LTP (Svenningsson et al. 2004; Centonze et al. 2001). An additional protein, ARPP-21, has been hypothesised to detect the order of the glutamatergic and neuromodulatory inputs, allowing LTP only if dopamine shortly follows synaptic activation by glutamate but not if it precedes it (Nair et al. 2016). Thus, this network only strengthens the synapse if synaptic activity is followed by neuromodulatory input activation.

2.8 LTD

The signaling network that controls LTD in corticostriatal synapses is shown in Fig. 2.8. LTD is expressed presynaptically, meaning that the synapse is weakened by reducing the probability of presynaptic neurotransmitter release. However, the signaling events to reduce the release probability are initiated postsynaptically. This signaling network also has two branches. One branch responds to elevated membrane voltage through Ca^{2+} influx from $Ca_v1.3$ channels. Ca^{2+} ions from $Ca_v1.3$ channels and from ryanodine intracellular stores are necessary to activate the enzyme $DGL\alpha$ which can produce endocannabinoids (Plotkin et al. 2013). The $Ca_v1.3$ channel is active at voltage levels above -50 mV and any local dendritic voltage elevation can activate it, irrespectively of whether the synapse is active. Under regular circumstances, this is not enough to trigger LTD (Fino et al. 2010). Activation of mGluR by glutamate is in addition necessary, thus providing synaptic specificity and initiating LTD only at the active synapses which experience elevated voltage in the operating range of $Ca_v1.3$. The mGluR signaling branch produces the substrate for the enzyme $DGL\alpha$, PIP₂, from which endocannabinoids are produced. Moreover, this also produces IP₃ intracellularly, which together with the Ca^{2+} from $Ca_v1.3$ will open the internal IP₃ calcium stores, adding to the rise of Ca^{2+} levels for activation of $DGL\alpha$ (Fino et al. 2010). Endocannabinoids diffuse to the presynaptic terminal and bind to the CB_1 receptors that initiate the signaling cascade

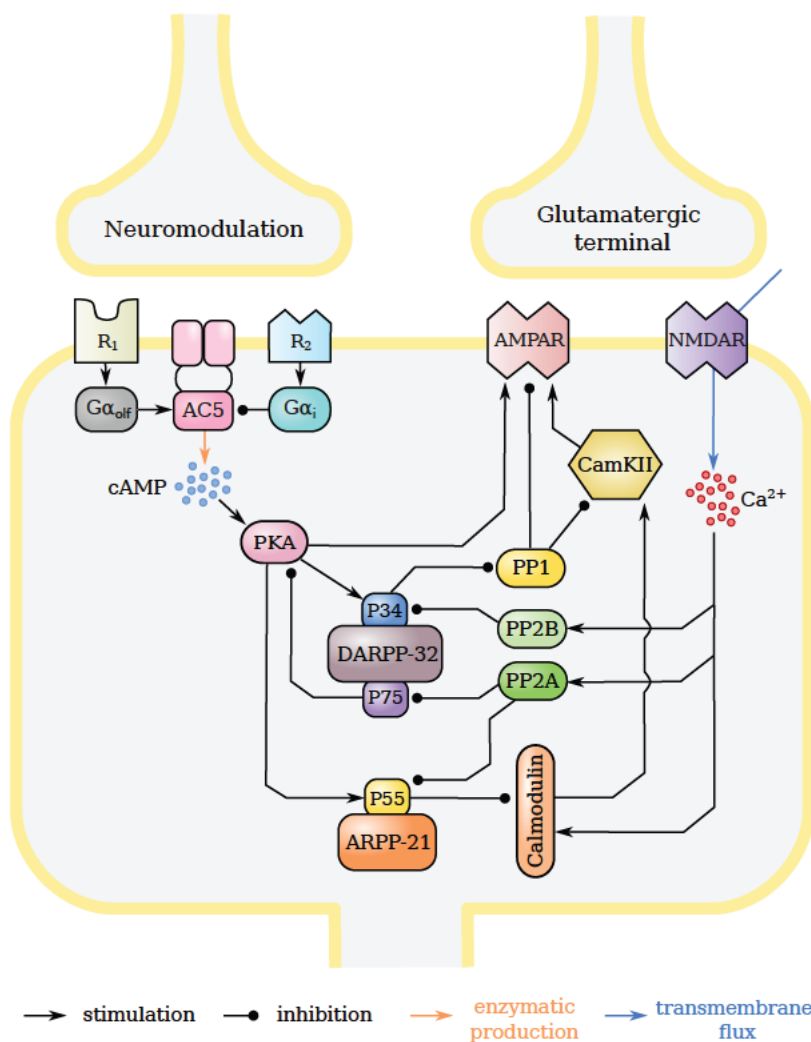


Figure 2.7: Signaling network (biochemical circuitry) for LTP in corticostriatal synapses.

for reducing neurotransmitter release. Lastly, a minimal $[Ca^{2+}]$ threshold for inducing LTD has been reported (Shindou et al. 2011).




Figure 2.8: Signaling network (biochemical circuitry) for LTD in corticostriatal synapses.

2.9 Goal-directed (reward) learning

Behaviors that are not innate need to be learned. In scientific terms, a reward is an object or event that causes approach behavior from the organism and consumption/interaction, produces learning of such behavior, and

is an outcome of decision making (Schultz 2016; Schultz 2007). Through goal-directed or reward learning, a behavior or action that leads to a subjectively valuable outcome (reward) is “reinforced”, i.e. is more likely to be performed again in the same circumstances in order to obtain that outcome. Conversely, actions that lead to a subjectively unvaluable outcome or no outcome are suppressed. Goal-directed learning consists of (at least) two processes: a cognitive process involving expecting a reward based on past experience (memory), and an associative learning process through which sensory stimuli (cues) which precede the reward become future motivators to make reward-obtaining actions (Berridge 2000; O’Doherty et al. 2017). The prefrontal and insular cortices are involved in the cognitive process of reward expectation, while the basal ganglia, more particularly the striatum and nucleus accumbens, are involved in the associative learning process (Berridge 2000; O’Doherty et al. 2017). Both processes operate in parallel and serve to guide behavior in the same direction.

The neural circuit(s) thought to be involved in goal-directed learning are shown in Fig. 2.9. To describe them we will use a simple experiment, where rats are trained that stimulating one of their whiskers is followed by a water reward. This is a typical stimulus-response learning experiment, where after repeated trials, the rats learn the association between whisker stimulation and a water reward, and as a result start licking in anticipation of the water reward, i.e. *before* the water is provided. It is also called associative learning, or Pavlovian conditioning. During the course of the experiment rats are head-restrained, so that as little brain activity is triggered through the other senses. The description of this experiment and the neural circuits for goal-directed learning are presented according to Esmaeili et al. (2020).

Sensory information from stimulating a single whisker first arrives in the trigeminal nucleus, where the cell bodies of the whisker sensory neurons are located (Fig. 2.9A). Signals from the trigeminal nucleus are transmitted via the brainstem to the sensory thalamus, which in turn passes the information to the primary whisker sensory cortex (wS1)¹. The activation of wS1 from even a single whisker causes rapid further activation of several brain areas thought to be relevant for goal-directed learning. In fact, projections from wS1 to at least 18 brain regions have been documented, although not all of them might be activated by single whisker stimulation. To begin with, wS1 projects to the secondary whisker sensory cortex (wS2) and the pri-

¹Also called barrel cortex.

mary whisker motor cortex (wM1) (Fig. 2.9B). These two pathways are thought to represent two parallel processing streams, the first of which processes ‘where’ and ‘when’ the stimulus occurred, and the second is about determining ‘what’ the stimulus is. How this information reaches the motor cortex area responsible for moving the tongue and jaw muscles (tjM1) is not known yet, since neither wS1 and wS2 project directly to tjM1. However, two pathways are thought to be involved. The first is a pathway from wS1 and wS2 via the secondary whisker motor cortex (wM2) which connects to the secondary motor cortex for the tongue and jaw (tjM2), and tjM2 is connected with tjM1. In this pathway the medial prefrontal cortex and the hippocampus are thought to provide decision-, planning- and context-related information. The second is a pathway from wS1 to the dorsolateral striatum of the basal ganglia, which forms the corticostriatal synapses, the topic of a large part of this thesis (Fig. 2.9B). The output of the basal ganglia disinhibits the motor thalamus, which through its projections to tjM1 significantly aids in activating motor commands from tjM1 to the premotor brainstem nuclei controlling the tongue and jaw muscles (Lalive et al. 2018). The brainstem nuclei control muscle movement via projections to the motor neuron nuclei for the tongue and jaw muscles (Fig. 2.9C).

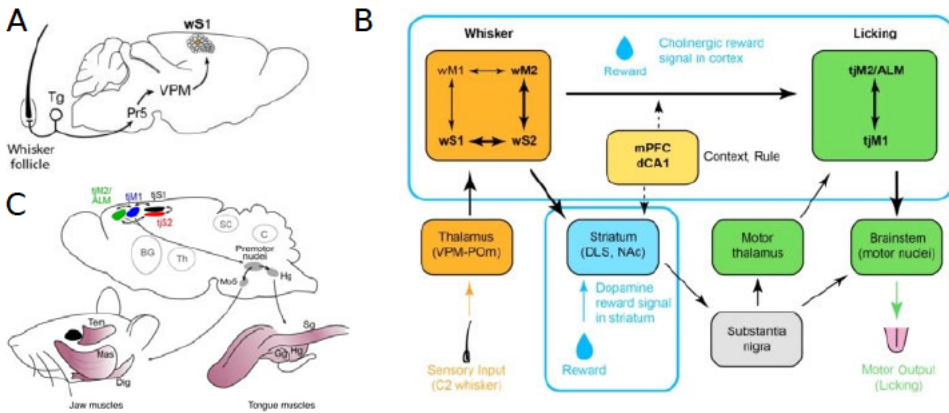


Figure 2.9: Neural circuits for goal-directed learning. Figure panels taken from figures in Esmaeili et al. 2020. © Esmaeili et al. 2020. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

Reward learning is thought to change the synaptic connections in both

the cortical and basal ganglia pathway and thus encodes new behaviors, and learning in the two pathways might correspond to the two cognitive and associative learning processes (Berridge 2000). Synaptic plasticity in the cortical pathway is modulated by the neuromodulator acetylcholine, while corticostriatal synaptic plasticity is modulated by dopamine. Our focus in this thesis is on corticostriatal synaptic plasticity, dealing, therefore, with the associative learning process. It is hypothesised that reward-related cholinergic signals also influence the cortical pathway through disinhibition, affecting synaptic plasticity between regions wS1 and wS2, for example. We do not address this aspect of reward-learning in this thesis.

2.10 Basal ganglia and goal-directed learning

The basal ganglia are a group of nuclei at the base of the cerebrum, under the cortex (they are some of the many subcortical nuclei). The nucleus receiving the cortical inputs is the striatum, and these inputs onto the SPNs form the corticostriatal synapses. Its projection neurons are inhibitory, and they project to the output nuclei of the basal ganglia, which control the activity of the thalamus and brainstem². Motor nuclei in the thalamus and brainstem are under constant inhibition by the basal ganglia output nuclei (Fig. 2.10).

Two parallel output pathways in the basal ganglia act oppositely on the motor nuclei in the thalamus and brainstem. When the SPNs from the direct pathway (dSPNs) are active, they disinhibit the motor nuclei. When the SPNs from the indirect pathway (iSPNs) are active, they increase the constant inhibition to the motor nuclei even more. By controlling the output of the thalamic and brainstem motor nuclei, the basal ganglia disinhibit required motor programs and suppress unwanted ones in the motor cortex and brainstem in order to initiate the appropriate goal-achieving actions.

The basal ganglia are necessary for proper movement initiation. Original theories about their function proposed that they select appropriate actions based on sensory information and motor plans. However, recent findings suggest action selection is done in the cortex, and the role of basal ganglia is to initiate the selected actions and modify how vigorously they are per-

²Besides the indirect route from the basal ganglia via the thalamus and motor cortex to the brainstem, mentioned in section 2.9, there are also direct connections from the basal ganglia to the brainstem (Arber and Costa 2022)




Figure 2.10: A disinhibitory pathway from sensory cortex to motor cortex via the basal ganglia and the thalamus. Figure reused from Purves et al. 2012 with permission from Oxford University Press. © 2012

formed (Thura and Cisek 2017). The corticostriatal synapses are believed to implement associations between sensory stimuli and the required motor output for achieving a goal, thus representing an important site in the gating of proper actions. Midbrain dopamine neurons densely innervate the striatum, and, as described in the previous sections, dopamine modulates when corticostriatal synaptic plasticity occurs.

2.11 Reward-related midbrain dopamine signals

The activity of dopamine neurons differs widely across the midbrain, and it also exhibits variations on three different time-scales, directly affecting striatal dopamine concentrations on those time-scales. The short review presented in this section is mainly according to the information in Schultz (2007) and Schultz (2016). Tonic dopamine neuron activity, which determines the levels of striatal dopamine over many hours or days, is necessary for proper movement, cognition, attention and motivation. Slow activity on the time-scale of seconds to tens of minutes is related to behaviors such as self-initiated limb movements, as well as the uncertainty with which the sensory stimulus predicts that a reward will occur. Finally, fast, sub-second changes in striatal dopamine are reward-related signals, which consist of two components (Fig. 2.11). The first is an initial peak signaling the occurrence of a conspicuous/salient, potentially rewarding, or otherwise attention-requiring sensory event, such as a loud noise, bright or fast moving objects, and liquids touching the mouth. It is an unselective response to a wide range of unpredicted sensory stimuli from any sensory modality. The second, main response, which evolves from the initial response, is related to identifying and valuing the sensory event. It is called a reward prediction error, and its occurrence in time depends on whether the organism is already familiar with the stimulus and its reward properties. A stimulus can itself be a reward (for example, food), or, as in the case of the mouse's whisker stimulation in section 2.9, it can predict that a reward will follow.

If the organism has experienced the stimulus already and knows/remembers that the stimulus is rewarding, or that it predicts a reward, the second dopaminergic response happens immediately after the first and reflects the reward value (Fig. 2.11, top). At the time of interaction with the stimulus or when the reward is delivered, if the reward had been correctly predicted (by the second dopaminergic response), there is no additional dopaminergic

response. If the reward is greater than predicted, dopamine neuron activity increases, resulting in a positive prediction error. If the reward is smaller or nonexistent (omitted), a dopamine pause occurs, representing a negative prediction error.

On the other hand, if an organism has never experienced the stimulus and does not know yet whether it is a reward, or predicts that a reward will follow, there is no second dopaminergic response when the stimulus occurs, and the initial response is followed by a dopamine pause (Fig. 2.11, bottom). The dopamine neurons respond only after interacting with the stimulus or after reward delivery, and their activity increases in proportion to the perceived reward value, resulting in a positive reward prediction error. If the stimulus is not a reward or is not followed by a reward, there is no dopamine neuron activity. The magnitude of the second dopaminergic response reflects the subjective reward value in both the learned (rewarded) and the novel (unrewarded) stimulus scenarios.




Figure 2.11: A schematic depiction of the fast, sub-second, reward-related dopamine signals. (Top) The appearance of the signals for a previously rewarded stimulus, i.e. a stimulus that the organism has encountered before and knows its reward value. (Bottom) The appearance of the signals for an unrewarded stimulus, i.e. a stimulus which the organism has not experienced previously and does not know its reward value. Figure reproduced with permission from Schultz (2016). © 2016 Springer Nature.

Because of this reward-related activity of midbrain dopaminergic neu-

rons, and the dense dopaminergic innervation to the striatum, corticostriatal synapses are thought to be a major site for reward learning.

The dSPNs mainly express D_1 dopamine receptors, and the iSPNs mainly express D_2 dopamine receptors. Because the two receptors have different affinity for dopamine, the two SPN types are differently responsive to dopamine. The tonic levels of dopamine are enough to activate D_2 receptors, while D_1 receptors can be activated with the dopamine peaks, emitted as reward prediction errors.

2.12 Project-specific background material

In this thesis we study questions related to the role of single SPNs in the brain's reward learning neural circuits. We study how the transduced neuromodulatory signals are processed in the corticostriatal synaptic signaling network for LTP by the enzyme AC5. We also study whether a single SPN can solve a computationally difficult reward learning task. It is related to the computational capacity of a single SPN, or how much the same neuron, responsible for gating/initiating the same action, can be engaged in different contexts. The solution of the task requires dendritic plateau potentials, because of which we also studied the how plateau potentials are generated in SPNs, focusing on the role of glutamate spillover.

2.12.1 AC5 signalling network

As mentioned in section 2.5, the levels of cAMP decide whether an active corticostriatal synapse will undergo LTP. Hence, regulating AC5, the enzyme which produces cAMP, can be seen as equivalent to implementing this decision function within the synaptic circuitry.

ACs are a family of membrane-bound enzymes, which contain two dimerized transmembrane domains that anchor them to the membrane, and a dimer catalytic core. The relative positioning of the catalytic monomers C1 and C2 determine whether the catalytic center is properly aligned to perform the conversion of ATP to cAMP. Each of the catalytic monomers has a binding groove for a G protein. The binding of a G protein affects the relative positioning of the catalytic monomers, thus regulating AC activity (Linder 2006). There are four different families of $G\alpha$ subunits that can bind to ACs. The $G\alpha$ subunits that bind to AC5 are $G\alpha_{olf}$, which belongs to the family of AC-stimulating subunits, and $G\alpha_i$, which belongs to the family of

AC-inhibiting subunits. On AC5, $G\alpha_{\text{olf}}$ binds to C2, and $G\alpha_i$ binds to C1 (Fig. 2.12). The binding of $G\alpha_{\text{olf}}$ to C2 causes a conformational change in AC5 such that the catalytic core is aligned to perform catalysis at a high rate, thus having a stimulatory effect on AC5. The binding of $G\alpha_i$ to C1 has the opposite, inhibitory effect: the alignment of the catalytic core prevents binding of ATP to the active site (Keulen and Rothlisberger 2017).

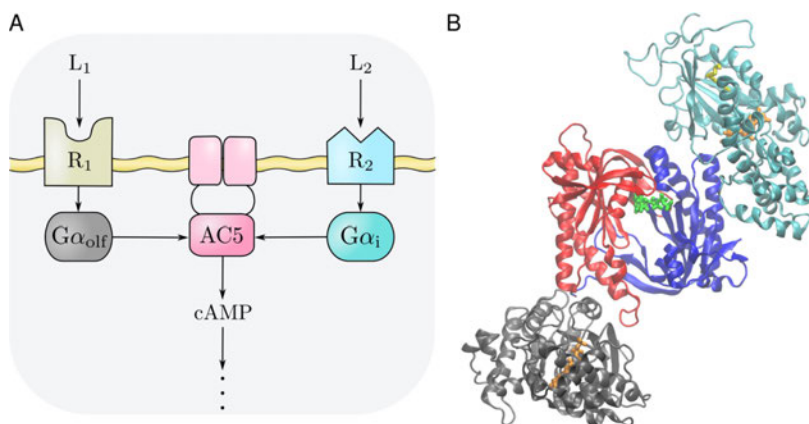


Figure 2.12: (A) General illustration of the AC5 signal transduction network. Two agonists (L_1 and L_2) bind to two GPCRs (R_1 and R_2), releasing the $G\alpha_{\text{olf}}$ and $G\alpha_i$ subunits, respectively. These stimulate and inhibit the conversion of cAMP, respectively. (B) Model of the possible $G\alpha_{\text{olf}} \cdot \text{AC5} \cdot G\alpha_i$ ternary complex. The cytosolic part of the AC5 enzyme consisting of the pseudo-symmetric C1 (blue) and C2 (red) cytoplasmic domains in an ATP-bound (green) conformation are complexed to an active conformation of $G\alpha_{\text{olf}}$ (gray) and to $G\alpha_i$ (cyan). GTP (orange) and a myristoyl moiety in $G\alpha_i$ (yellow) are shown in stick representation. Controlling the relative positions and conformations of C1 and C2 may enhance or inhibit enzymatic function. This is one way in which $G\alpha_{\text{olf}}$ and $G\alpha_i$ exert their regulatory effects: each of them has a separate binding site on the AC5 domain dimer. The model shown is in fact the initial configuration of the ternary complex used in the all-atom molecular dynamics simulations, performed by one of the partners in this study. Figure and caption reused from Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

The structure of a chimeric AC · $G\alpha_s$, comprised of an AC catalytic core (one dimer from AC5 and another from AC2) bound to $G\alpha_s$, a protein

similar to $G\alpha_{\text{olf}}$, has been determined with X-ray crystallography, experimentally confirming its potential formation in vivo (Tesmer et al. 1997). The formation and catalytic activities of the binary complexes $AC5 \cdot G\alpha_{\text{olf}}$ and $AC5 \cdot G\alpha_i$ have been confirmed in preparations with membrane fragments from Sf9 cells containing the entire AC5 enzyme, and with a reconstituted AC5 catalytic core in a solution (Taussig et al. 1994; Dessauer et al. 1998; Chen-Goodspeed et al. 2005). Since $G\alpha_{\text{olf}}$ and $G\alpha_i$ bind to AC5 on separate locations, it has been hypothesized that a ternary $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ complex may form during AC5 regulation (Fig. 2.12B) (Dessauer et al. 1998; Chen-Goodspeed et al. 2005). Nevertheless, it has not been possible to detect this complex experimentally, and it is not known what its catalytic activity would be. The existence of a ternary complex would allow AC5 to act as not just a signal transducer, but also as a signal integrator, since its activity could then be regulated by more than one binding partner, i.e. by two G proteins each activated by a different extracellular signal. Of course, in order for AC5 to bind to both $G\alpha_{\text{olf}}$ and $G\alpha_i$ simultaneously, both G protein subunits need to have been activated very closely in time.

Many ACs have been found to perform coincidence detection – the detection of co-occurring signaling events, which results in significantly increased cAMP production only when more than one signaling event occurs almost simultaneously (Anholt 1994; Bourne and Nicoll 1993; Delmeire et al. 2003; Impey et al. 1994; Lustig et al. 1993; McVey et al. 1999; Mons et al. 1999; Nair et al. 2019; Nair et al. 2015). In addition, a previous study with a large model of the LTP signaling network predicted that AC5 responds most strongly to a simultaneous peak in dopamine (Da \uparrow) and a pause in acetylcholine levels (ACh \downarrow), i.e. both stimulation by increased $G\alpha_{\text{olf}}$ and disinhibition by decreased $G\alpha_i$ are necessary for the enzyme to produce significant amounts of cAMP (Nair et al. 2015). This suggests that AC5 might also function as a coincidence detector. In this previous modeling study, the AC5 regulatory motif contained the ternary complex, but the reaction rates for forming it were small compared to the reaction rates for forming the binary complexes, resulting in small amounts of ternary $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ complex being formed, leaving open questions regarding its role in AC5 regulation.

Because of the difficulty in detecting the ternary complex experimentally and the remaining open questions for its possible role, we resorted to studying the AC5 enzyme and its regulation with computational methods. This was done within the scope of an international project where different

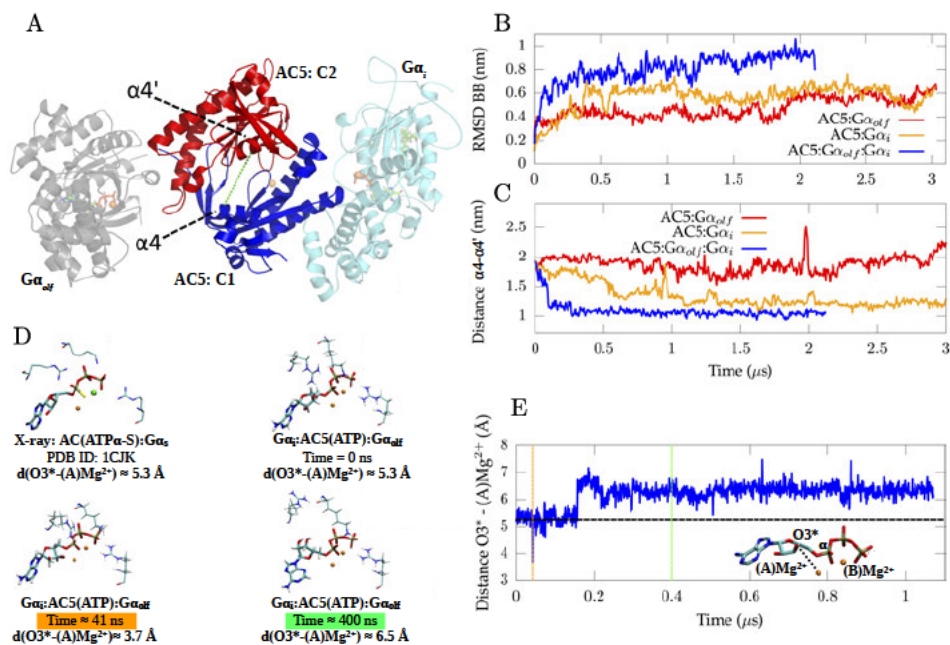


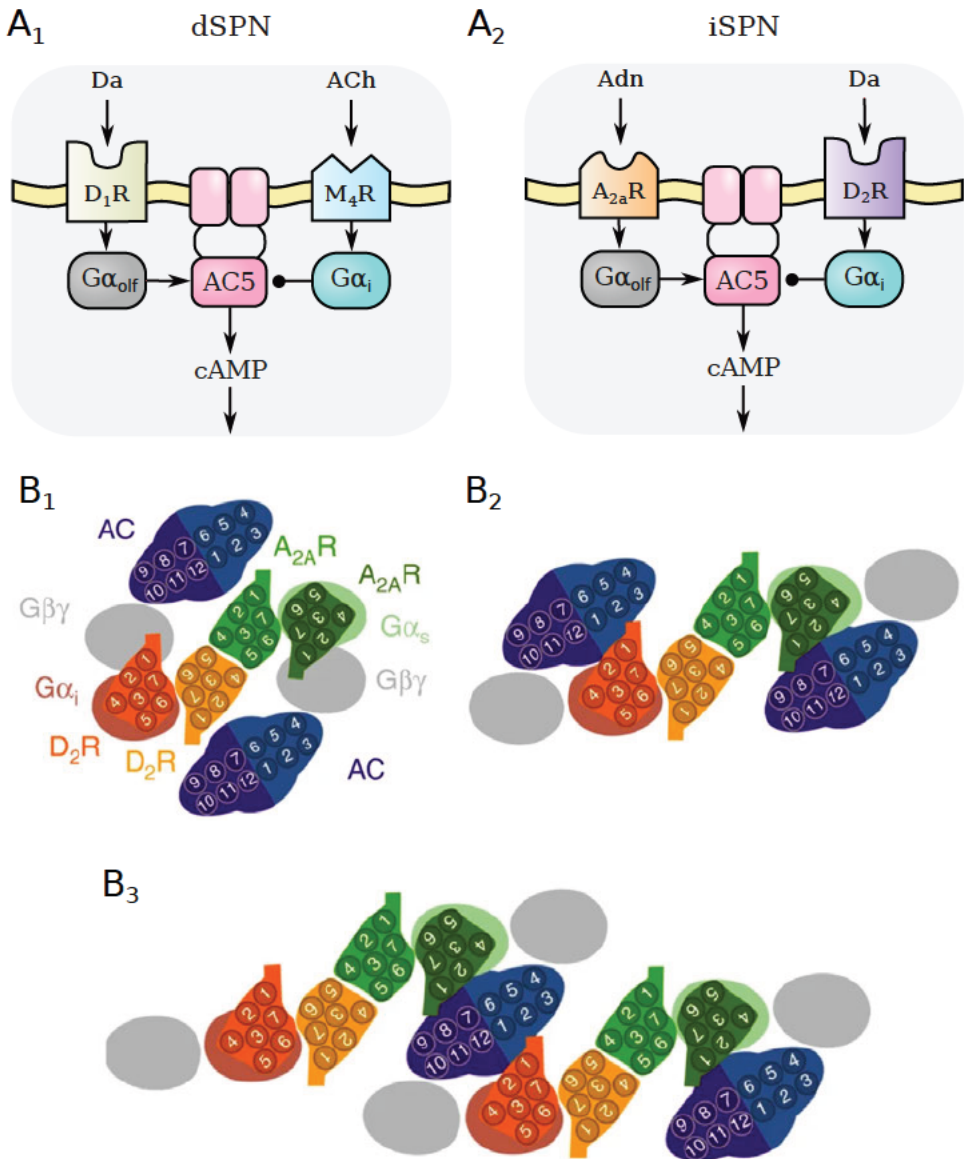
Figure 2.13: Stability analysis of the *apo* ternary complex, conformation of the active site of AC5 in the *apo* state, and conformation of ATP in AC5's active site in the *holo* state. (A) Initial conformation of the *apo* ternary complex including $G\alpha_{olf}$ (gray), $G\alpha_{i1}$ (cyan), C1 (blue), C2 (red). In addition to the protein structures, two GTP molecules as well as the myristoyl moiety of the $G\alpha_{i1}$ subunit and three Mg^{2+} ions are shown. The green dashed line indicates the distance between the $C\alpha$ atom of Thr1007 (belonging to helix $\alpha 4$) and the $C\alpha$ atom of Ser1208 (belonging to helix $\alpha 4'$). (B) Root-mean-square-deviation of the backbone of the protein complexes $G\alpha_{olf} \cdot AC5 \cdot G\alpha_{i1}$, $AC5 \cdot G\alpha_{olf}$ and $AC5 \cdot G\alpha_{i1}$. (C) Time evolution of the distance between Thr1007 and Ser1208 for the three simulated complexes in *apo* form. (D) Conformation of ATP in the $G\alpha_{olf} \cdot AC5(ATP) \cdot G\alpha_{i1}$ system at different time intervals in the trajectory as well as the conformation of ATP α -S in the AC $\cdot G\alpha_s$ X-ray structure (PDB code 1CJK). The color of the time in green and orange corresponds to the coloured lines in image (E). (E) Distance between the oxygen of a hydroxyl on the sugar ring of ATP, $O3^*$, and a Mg^{2+} ion, (A) Mg^{2+} , in the active site of the *holo* ternary complex. The black dashed line shows the distance between the two atoms in the AC $\cdot G\alpha_s$ X-ray structure (PDB ID 1CJK) to which ATP α -S is bound. Figure and caption reproduced from Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

partners studied the enzyme at different spatial and temporal scales, i.e. we employed multiscale simulations. Molecular dynamics (MD) simulations were used to study the conformational stability of the ternary complex on the microsecond timescale. Brownian dynamics (BD) simulations, operating on the millisecond timescale, were used to predict the forward rate constants of $G\alpha$ subunit binding to AC5 and to the binary complexes. The MD and BD simulations were performed by the project partners and we present the obtained results as background material. The MD simulations suggested that if the ternary complex forms, it is stable and catalytically inactive, since the positioning of the catalytic core is unfavorable for performing the conversion from ATP to cAMP (Fig. 2.13). The results from the BD simulations are given in Table 2.1. We incorporated these results to construct a kinetic model of the AC5-based signaling network and study whether AC5 is also a coincidence detector.

Rate constant	<i>apo</i> AC5	<i>holo</i> AC5
k_{f1}	0.02 (0.01)	0.03 (0.01)
k_{f2}	0.02 (0.01)	0.05 (0.01)
k_{f3}	0.01 (0.01)	0.02 (0.01)
k_{f4}	0.02 (0.01)	0.02 (0.01)

Table 2.1: Bimolecular association rate constants ($\text{nM}^{-1} \text{s}^{-1}$) for the forward reactions computed from BD simulations (standard deviation in parentheses). Rate constants were calculated separately for complexes incorporating both *apo* and *holo* AC5 (*apo* AC5 is AC5 without the substrate ATP, and *holo* AC5 is AC5 bound to the substrate ATP). Table and caption reproduced from Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

AC5 is expressed in both dSPNs and iSPNs, and it is regulated by both $G\alpha_{\text{olf}}$ and $G\alpha_i$ in both SPN types (Fig. 2.14A). In dSPNs, $G\alpha_{\text{olf}}$ is released when dopamine binds to the D_1R , while in iSPNs it is released when adenosine binds to the $A_{2a}R$. In dSPNs $G\alpha_i$ is released when acetylcholine binds to the M_4R , while in iSPNs it is released when dopamine binds to the D_2R . So, in dSPNs AC5 is responsive to dopamine and acetylcholine, whereas in iSPNs it is responsive to dopamine and adenosine. Moreover, because of the different properties of the D_1R and D_2R , AC5 responds to dopamine differently in the two SPN types. The D_1R requires high dopamine concentrations



to be activated, such as those produced with the sub-second, reward-related dopamine signals, and it does not respond to basal dopamine levels. The D_2R , on the other hand, can be activated by basal dopamine levels (Yapo et al. 2017). Moreover, experimental studies have suggested that, at least in iSPNs, the D_2Rs and $A_{2a}Rs$ form receptor complexes with each other and with AC5, in a suggested arrangement which allows for the simultaneous binding of the G proteins to AC5, i.e. for the ternary complex to form (Fig. 2.14B₁–B₃) (Navarro et al. 2018). Regarding acetylcholine, experimental data indicates that, at an organism’s resting state, much of AC5 in the dSPNs could be inhibited by $G\alpha_i$ due to a tonic level of acetylcholine produced by the tonic activity of the striatal cholinergic interneurons, which activates $G\alpha_i$ via the M_4 receptor (Nair et al. 2019). Basal levels of cAMP are low. For this reason, both the original large model of the LTP signaling network, and our reduced kinetic model of AC5 regulation have been set up so that in the basal state acetylcholine activates the M_4R and $G\alpha_i$ initially inhibits a large proportion of AC5. As mentioned above, the basal dopamine levels do not activate the D_1R . A $Da \uparrow$ is needed to activate $G\alpha_{olf}$, which can

Figure 2.14 (*preceding page*): (A) In both dSPNs and iSPNs, AC5 is activated by $G\alpha_{olf}$ and inhibited by $G\alpha_i$. The receptors which are coupled to the G proteins are different in the two SPN types. (A, Left) In dSPNs, the dopamine D_1R activates G_{olf} while the acetylcholine M_4R activates G_i . (A, Right) In iSPNs, the adenosine $A_{2a}R$ activates G_{olf} while the dopamine D_2R activates G_i . (B) Suggested spatial arrangement of the $A_{2a}Rs$, D_2Rs and AC5 in iSPNs according to experimental studies (Navarro et al. 2018). G proteins are also added in the model based on the assumption that they form a complex with their GPCRs. The molecules are indicated in a top-down view of the membrane from the extracellular side. The two AC5 transmembrane domains are shown in blue, D_2Rs are shown in shades of orange, and $A_{2a}Rs$ are shown in shades of green. The spatial arrangement is a minimal computational solution to six experimentally determined requirements, among which are those that D_2Rs and $A_{2a}Rs$ form both homodimers and heterodimers, as well as the findings that they interact with different regions of the AC5 transmembrane domains depending on whether they are activated by agonists. (B₁) Suggested precoupled receptor complex when agonists are not bound. (B₂) Suggested precoupled receptor complex when agonists are bound. (B₃) Suggested linear arrangement of precoupled receptor complexes which support the formation of a ternary complex. The numbers on the molecules indicate the transmembrane helices within the molecules. Panels B₁ – B₃ are reproduced from Navarro et al. (2018). © 2018 Navarro et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.




Figure 2.15: A schematic depiction of the hypothesis how coincidence detection by AC5 could work in dSPNs. (A) In the basal state, the tonic level of ACh causes inhibition of AC5 via $G\alpha_i$. (B) A $Da \uparrow$ alone cannot activate AC5, because the formed ternary complex is catalytically inactive (as suggested by the MD simulations). (C) A coincident $ACh \downarrow$ would be needed to disinhibit AC5 from $G\alpha_i$ so that cAMP can be produced.

then activate AC5. In theory, a co-occurring $ACh \downarrow$, which would deactivate $G\alpha_i$, would be needed to disinhibit AC5 from $G\alpha_i$ and produce cAMP (Fig. 2.15).

In order to perform coincidence detection, the signaling network should be able to make a clear distinction between the situation of a simultaneous dopamine peak and acetylcholine pause ($Da \uparrow + ACh \downarrow$) and that of a single signal, i.e. $Da \uparrow$ or $ACh \downarrow$ alone. This distinction is realized by producing different amounts of cAMP, i.e. by differences in the enzyme's catalytic rate. With the kinetic model we can obtain the enzyme's average catalytic rate for the different neuromodulatory signals, and based on that, determine whether it can perform coincidence detection. If the amount of cAMP during $Da \uparrow + ACh \downarrow$ is much greater than during a $Da \uparrow$ or $ACh \downarrow$ alone, then AC5 can perform coincidence detection. If, for example, the amount of cAMP is similar during $Da \uparrow + ACh \downarrow$ and a $Da \uparrow$ alone, then AC5 cannot distinguish whether two or only one neuromodulatory signals occurred extracellularly.

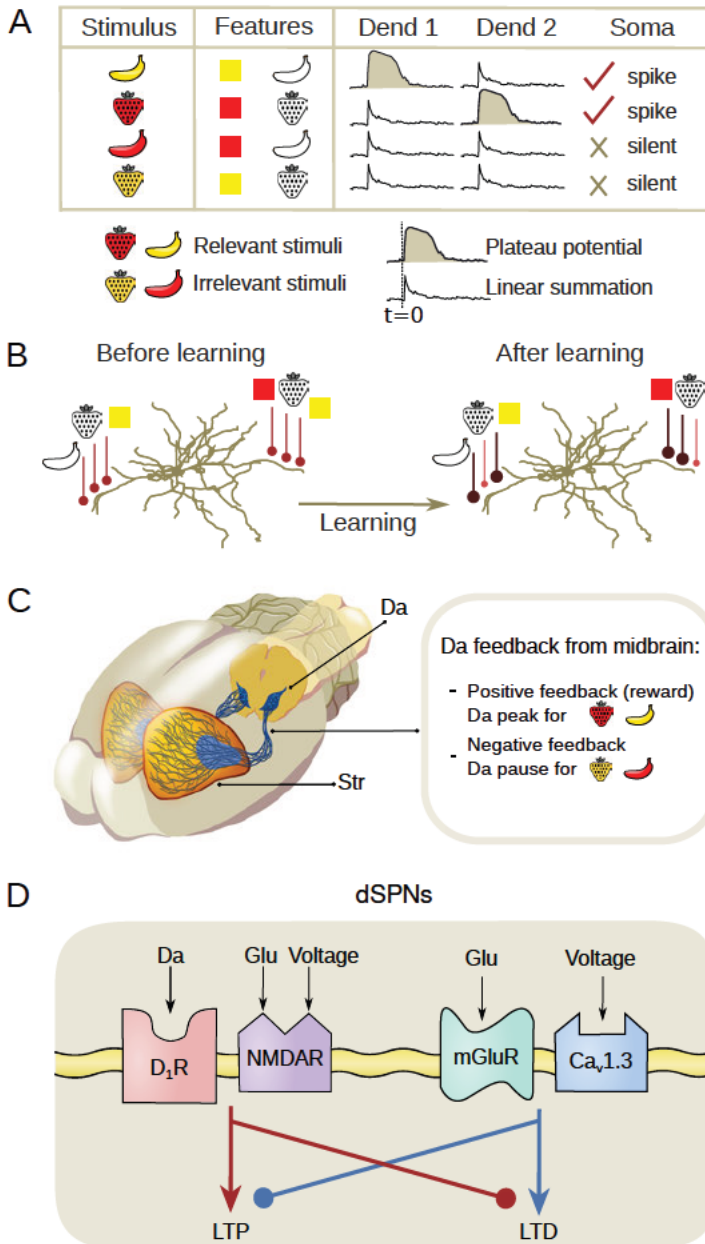
2.12.2 Synaptic plasticity rule

After studying how the regulation of AC5 gates LTP in the striatum, we continued towards formulating a synaptic plasticity rule, also called a learning rule, based on the synaptic signaling networks for LTP and LTD in the corticostriatal synapse, introduced in sections 2.7 and 2.8. We used this

plasticity rule to study the computations that a single SPN can perform, and through this gain insight in the role of single SPNs within the neural circuits for goal-directed learning. Specifically, we used a computationally difficult task, called the nonlinear feature binding problem (NFBP), described in the later sections (Fig. 2.16A). Its solution would imply that the same dSPN involved in initiating a particular action could be recruited in multiple contexts requiring the same action. Moreover, the dSPN can distinguish between similar relevant and irrelevant situations, and learn to spike or be silent accordingly, thus initiating the action only in relevant situations.

We developed a local, calcium- and reward-based learning rule which can utilize plateau potentials to solve the NFBP. Many learning rules use the relative difference between the timing of synaptic activity and the timing of somatic spiking as a basis to update synaptic strength (Feldman 2012). However, in many neurons the action potential does not reliably propagate backwards in the dendrites (backpropagating action potential, bAP). Even if it does, it is unclear how distal synapses can observe bAPs intracellularly and obtain the information about spike timing, since the additional Ca^{2+} influx from the bAPs might be minor compared to the Ca^{2+} influx from the already elevated local voltage at the synaptic site. For this reason, we formulated the learning rule solely based on the local, i.e. synaptic intracellular calcium concentration, $[\text{Ca}^{2+}]_i$. Dopamine reward signals, which represent feedback about whether the initiated action had a positive outcome or not, guide the learning rule as to whether to strengthen or weaken the synapse (Fig. 2.16A).

Experiments have demonstrated that in dSPNs, which express the D_1 dopamine receptor, a $\text{Da} \uparrow$ stimulates the D_1R , which together with significant Ca^{2+} influx through NMDA receptors triggers LTP (Shen et al. 2008; Yagishita et al. 2014; Fisher et al. 2017). Dopamine and Ca^{2+} activate the two branches of the LTP signaling network described in section 2.7. If the D_1Rs are stimulated, the LTD signaling network is inhibited (Shen et al. 2008). Conversely, when little/no dopamine is bound to the D_1 receptors, as during a $\text{Da} \downarrow$, and there is significant Ca^{2+} influx through $\text{Ca}_v1.3$ channels, LTD occurs (Shen et al. 2008; Fino et al. 2010; Plotkin et al. 2013). The D_1R needs to be free of dopamine in order to disinhibit the LTD signaling network and allow the expression of LTD. These experimental findings describe the phenomenological plasticity outcomes from synaptic stimulation without investigating all molecules participating in the signaling networks in sections 2.7 and 2.8. The phenomenological plasticity outcomes from



these experimental studies are graphically summarized in Fig. 2.16D. Also, voltage-gated calcium channels, most notably $\text{Ca}_v3.2$ and $\text{Ca}_v3.3$, have been reported to be necessary for LTP (Fino et al. 2010). We base the learning rule on these findings, without explicitly modeling all molecules in the signaling networks for LTP and LTD.

As we explain in the following sections, we use the learning rule with a setup of pre-existing clustered synaptic inputs which can generate plateau

Figure 2.16 (*preceding page*): The NFBP, a conceptual solution for it, and SPN mechanisms to solve it. (A) The NFBP and a conceptual solution: The NFBP is represented with an example from visual feature binding. In the simplest form of the NFBP, a stimulus has two features, here shape and form, each with two possible values, strawberry and banana, and red and yellow, respectively. The NFBP consists of responding with neuronal spiking to two of the feature combinations, called the relevant stimuli (red strawberry and yellow banana), and remaining silent for the other two feature combinations, called the irrelevant stimuli (yellow strawberry and red banana). If each feature is represented with locally clustered synapses, the NFBP can be solved when the co-active clusters on a single dendrite, corresponding to a relevant stimulus, evoke a plateau potential, thus supralinearly exciting the soma. Conversely, co-activation of synaptic clusters for the irrelevant combinations should not evoke plateau potentials. (B) Illustration of how synaptic plasticity may contribute to solving the NFBP for a pre-existing arrangement of synaptic clusters on two dendrites. A plasticity rule which strengthens only synaptic clusters representing relevant feature combinations, so that they produce robust supralinear dendritic responses, while weakening synapses activated by irrelevant feature combinations, could solve the NFBP. (C) Dopamine (Da) feedback from the midbrain to the striatum (Str) guides the learning process, differentiating between positive feedback for relevant stimuli and negative feedback for irrelevant stimuli. Positive feedback represented by dopamine peaks is necessary for LTP, and negative feedback represented by a dopamine pause is necessary for LTD. (D) Synaptic plasticity outcomes after stimulating the neurotransmitter receptors that initiate plasticity in a dSPN corticostriatal synapse (redrawn from Shen et al. (2008)). NMDA Ca^{2+} influx, followed by stimulation of D_1 Rs, triggers LTP, while inhibiting LTD. Ca^{2+} influx from $\text{Ca}_v1.3$ channels and activation of mGluRs when D_1 Rs are free of Da triggers LTD, while inhibiting LTP. Figure adapted from Khodadadi et al. (2024). © 2024, Khodadadi et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

potentials, since plateaus are crucial for solving the NFBP. Plateau potentials, as well as some synaptic clustering, have been demonstrated in SPNs (Plotkin et al. 2011; Oikonomou et al. 2014; Du et al. 2017; Hwang et al. 2022; Day et al. 2024; Sanabria et al. 2023), and corticostriatal synapses participate in goal-directed learning, which is an additional basis to motivate the study whether a single SPN can solve the NFBP.

2.12.3 Synaptic clustering

Synapses in many brain regions exhibit fine scale organization on the dendrites into functional synaptic clusters. They are spatially localized, or grouped in dendritic stretches approximately 20–30 μm long, and receive correlated inputs or represent a common sensory feature. Synaptic clustering has been observed across brain regions, developmental ages, and different animal species, and frequently forms during development (Kleindienst et al. 2011; Takahashi et al. 2012; Winnubst et al. 2015; Wilson et al. 2016; Winnubst et al. 2015; Iacaruso et al. 2017; Scholl et al. 2017; Niculescu et al. 2018; Kerlin et al. 2019; Ju et al. 2020; Kirchner and Gjorgjieva 2021). Sensory features such as the orientation of a visual stimulus, color, receptive fields, or sound intensity are represented by such functional clusters, and frequently, a single dendrite has clusters for different features positioned adjacently (Chen et al. 2011; Wilson et al. 2016; Iacaruso et al. 2017; Scholl et al. 2017; Ju et al. 2020; Jia et al. 2010; Varga et al. 2011). With this arrangement, it is possible that uniting multiple single features into more complex objects occurs in dendrites by activating several functional clusters together.

2.12.4 The binding problem

The combination of the many single features of an object/entity into a unified representation of that object/entity by the brain is called the binding problem (Ghose and Maunsell 1999; Gray 1999; Reynolds and Desimone 1999; Roskies 1999; Treisman 1999; Malsburg 1999; Wolfe and Cave 1999). The problem is based on the realistic assumption that there are not enough neurons in the brain to encode all possible feature combinations of all the stimuli that exist in the world, so a mechanism must exist which enables the same groups of neurons to dynamically represent objects/entities as they appear to the senses or as evoked from memory (Malsburg 1999). We do

not deal with finding such a mechanism for solving the general binding problem in this thesis. However, large numbers of combination-coding neurons seem to exist at least in the preattentive computational stages of the visual system (Ju et al. 2020). The learning rule we developed is a solution to feature binding, and to the computationally more complex nonlinear feature binding, in such neurons which already receive innervation for multiple features.

2.12.5 The nonlinear feature binding problem

As a task, the NFBP is relevant to brain regions which perform integration of multimodal input signals, or signals representing different features of the same modality (Roskies 1999). It is usually illustrated with examples from the visual system, as in Fig. 2.16A (Roskies 1999; Malsburg 1999; Tran-Van-Minh et al. 2015). It is, nevertheless, also relevant to the striatum, where goal-directed learning might require the processing of several different types of sensory- and motor-related input features. In the most elementary form, the NFBP consists of discriminating between two types of stimuli, each of which are a combination of two features (for example, color and shape). Each feature can have two values (red and yellow for color, and strawberry and banana for shape), giving a total of four possible stimuli (feature combinations). To solve the NFBP, the neuron should respond with somatic spiking to only two stimuli (which we call relevant stimuli), while remaining silent for the other two (irrelevant stimuli). The problem is nonlinear since, on average, the neuron receives the same amount of excitation for all four stimuli, but should learn to respond to them differently. Such nonlinear problems are solved by processing the different stimuli nonlinearly, and SPNs have supralinear processing capabilities in the form of plateau potentials, as we described in Section 2.4. If each feature is represented by a functional synaptic cluster, the SPN should learn to strengthen clusters for the relevant feature combinations so that activated together they will elicit a plateau, while weakening clusters belonging to irrelevant feature combinations so that they do not elicit plateaus (Fig. 2.16B).

To guide the learning process, we set up the task so that an organism's interaction with the relevant stimuli is rewarding and produces a $Da \uparrow$, while the irrelevant stimuli are not rewarding and produce dopamine pauses ($Da \downarrow$) (Fig. 2.16C). We use dSPNs in this study, and we assume that the necessary coincident $ACh \downarrow$ also occurs together with the $Da \uparrow$ to enable LTP.

2.12.6 Glutamate spillover

As introduced in section 2.4, plateau potentials are generated when a synaptic cluster is synchronously activated within a short time window of several tens of milliseconds. Experimentally, plateau potentials have been evoked by activating synaptic clusters by glutamate uncaging or by stimulating stretches of dendrites by glutamate iontophoresis or repeated synaptic stimulation (Schiller et al. 2000; Major et al. 2008; Plotkin et al. 2011; Oikonomou et al. 2012; Du et al. 2017; Kumar et al. 2018; Gao et al. 2021). All three techniques have been used to apply a progressively increasing glutamate stimulus in equal increments. In this case a supralinear, all-or-none response in the amplitude of the somatic voltage is observed, qualifying the plateau potential as an all-or-none voltage event, similarly to the action potential. An example of all-or-none plateau potentials evoked with glutamate iontophoresis is shown in Fig. 2.4C (taken from Oikonomou et al. (2014)). However, when computationally modeling plateau potentials, they often appear graded when evoked by stimuli progressively increasing in equal increments (Fig. 2.17).

As described in section 2.4, the crucial ingredient for an all-or-none plateau potential is the sigmoidal shape of the Mg^{2+} block gating function. For a given $[\text{Mg}^{2+}]$, it is described by the parameters α and η , because of which we decided to study the effects of varying these parameters on the behavior of the plateau potential. Collecting available existing data on the gating function, we found there is a natural variability in the parameters α and η , which depends on the NMDAR isoform and on the concentration of other intra- and extracellular ions (Table 2.2). Moreover, studies which model plateau potentials almost always use steeper gating functions than the ones experimentally reported, without an explicit reason for the change in parameters (Table 2.3). As can be seen in Fig. 2.17, by using a steeper gating function instead, the plateaus become all-or-none events. In light of the large natural variability in gating function steepness, how is it that plateau potentials experimentally appear as robustly all-or-none events? Also, how to incorporate the all-or-none quality in a model without using unrealistically steep gating functions?

A recent modeling study has included glutamate spillover when modeling plateau potentials (Gao et al. 2021). Glutamate spillover is thought to occur when a cluster of synapses is stimulated more strongly or repeatedly, and the amounts of exocytosed glutamate surpass the ability of astrocytes to

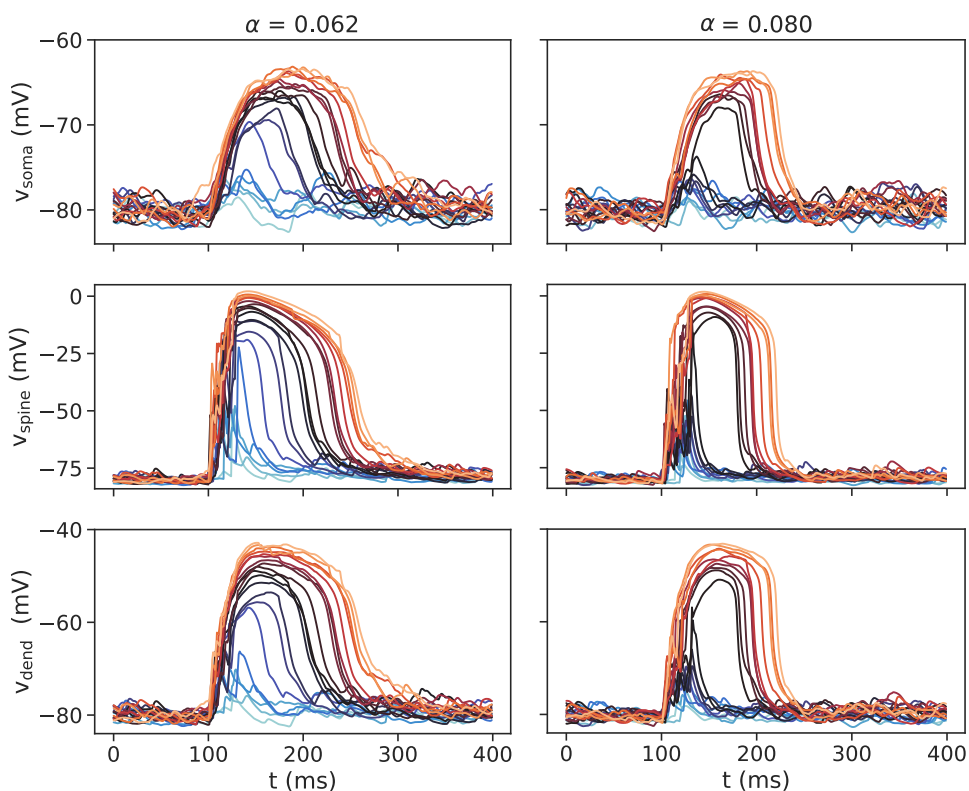


Figure 2.17: Plateau potentials in SPNs evoked by a cluster of varying size from 1 to 20 synapses in one dendrite. (Left) Graded plateau potentials in the soma, spine and dendritic shaft with the experimentally determined value of the steepness parameter α from Jahr and Stevens (1990). (Right) All-or-none plateau potentials in the soma, spine and dendritic shaft with a commonly used larger value for α in computational models (steeper Mg^{2+} block).

	$[\text{Mg}^{2+}]$ (mM)	α (/mV)	η (/mM)
Nowak et al. (1984)	0.5	0.04	1.33
Jahr and Stevens (1990)	1	0.06	0.28
Chen and Huang (1992)	0.03	0.05	0.49
Sharma and Stevens (1996)	3	0.06	0.28
McMenimen et al. (2006)	2	0.06	0.42
	0.2	0.05	3.3
Chiu and Carter (2022)	1	0.07	0.11
	0.7	0.07	0.1
	0.8	0.07	0.12

Table 2.2: Gating function parameter values from experimental data. Parameter values are determined from fits to experimental data, except for Jahr and Stevens (1990) which have reported the parameters themselves. Table reused from Trpevski et al. (2023). Columns on neuron type and ionic concentration have been omitted for ease of comparison with Table 2.3 (for full table see Trpevski et al. (2023)). © 2023 Trpevski, Khodadadi, Carannante and Hellgren Kotaleski. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

take it up, thus causing it to spill over from the synaptic cleft into the space surrounding the spine necks and dendritic shafts. Here it stimulates extrasynaptic NMDARs (eNMDARs), which can aid the generation of plateau potentials (Chalifoux and Carter 2011; Oikonomou et al. 2014; Garcia-Munoz et al. 2015).

We analyzed the plateau potential models with and without glutamate spillover and found that the all-or-none behavior is very sensitive to the steepness of the Mg^{2+} block when no glutamate spillover is included in the model and, conversely, very robust with respect to the steepness of the Mg^{2+} block when glutamate spillover is included. Therefore, in light of the natural variability of the Mg^{2+} block gating function, glutamate spillover might be an important mechanism in generating all-or-none plateau potentials.

	[Mg] (mM)	α (/mV)	η (/mM)
Rhodes (2006)	1 and 2	0.08	0.28
Major et al. (2008)	1.8	0.08	0.11
Farinella et al. (2014)	1	0.08	0.3
Poleg-Polsky (2015)	1	0.08	0.25
Doron et al. (2017)	1	0.08	0.28 and 1.45
Du et al. (2017)	1	0.07	0.33
Dorman et al. (2018)	1.4	0.1	0.06
Kumar et al. (2018)	1	0.08	0.25
Ecker et al. (2020)	1	0.06	0.38
Gao et al. (2021)	1	0.08	0.25

Table 2.3: Examples of studies that have modeled plateau potentials and the parameters used for the gating function. Ecker et al. (2020) have corrected η for the omitted junction potential in Jahr and Stevens (1990), and use the same value for α . Note that Major et al. (2008) have used a different form of the gating function equation, and we have converted their parameter values to match the form in Eq. 2.1. Table reused from Trpevski et al. (2023). © 2023 Trpevski, Khodadadi, Carannante and Hellgren Kotaleski. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

3 Methods

3.1 Kinetic model of the AC5 signaling network

To study the regulation of the enzyme AC5, we constructed a kinetic model describing the interactions of AC5 with the regulatory G proteins. We constructed the model by simplifying the much larger model of the synaptic machinery for LTP in Nair et al. (2015), in order to focus only on the reactions for the regulation of AC5. Kinetic models are ordinary differential equations which quantify how the concentrations of reactants in chemical reactions vary in time as the reaction progresses. This is called the kinetics of the chemical reaction, hence the name of the kinetic models. They describe a homogeneous (well-mixed) solution of reactants and are based on the law of mass action, which states that the reaction rate is proportional to the probability of a collision of the reactants. In a homogeneous solution, the probability of collision is proportional to the concentration of the reactants. For a single reversible chemical reaction in which two compounds (substrates) S_1 and S_2 react to produce m molecules of product P :



the kinetic model that describes the reaction is:

$$\frac{dS_1}{dt} = \frac{dS_2}{dt} = -v; \quad \frac{dP}{dt} = mv$$

$$v = k_f[S_1][S_2] - k_r[P]^m$$

where v is the reaction rate, $[\cdot]$ denotes concentration, and k_f and k_r are the forward and reverse reaction rate constants. The rate constants describe

how often a collision results in a transformation of the reactants, and are dependent on the temperature of the reaction.

Kinetic models are phenomenological and macroscopic descriptions of many, homogeneously distributed, microscopic particles. The synaptic environment, on the other hand, is highly structured, with many molecules assembled into precoupled molecular complexes and with only some molecular species freely diffusing. Kinetic models can be viewed as focusing only on the behavior of a system that is a result of the particular pattern of chemical reactions between the molecules, i.e. the chemical reaction network (biochemical circuitry), without considering any spatial ordering or structure. Nevertheless, kinetic models can be fitted to match experimentally measured concentrations of each molecular species by appropriately fitting the reaction rate constants. As such, they are very valuable in understanding the qualitative behavior of the signal transduction network centered around AC5, located in the structured synaptic environment.

The kinetic model of the AC5-centred signal transduction network is given in Fig. 3.1. To see the effects of the ternary complex on AC5 activity, we compare the results of the kinetic model with the ternary complex to the results from a null model where the ternary complex does not form, i.e. only bimolecular complexes of AC5 with the G proteins occur.

The two versions of the model have 8 and 16 rate constants, respectively. In Fig. 4.2, we have additionally used versions of the model which included cAMP production to illustrate the correspondence between k_c and the levels of cAMP and thus justify the use of k_c as a proxy for the cAMP levels. The reactions and parameters for cAMP production and degradation have been taken from Nair et al. (2015).

Neither Da and ACh nor their receptors are explicitly included in the model. This eliminates the parameters that would have corresponded to the reactions of ligand and receptor binding, and G protein and receptor binding. Instead, when Da \uparrow and ACh \downarrow occur, they affect the conversion of inactive to active G proteins as follows. Firstly, we have omitted the heterotrimeric nature of the G proteins, i.e. we have not included the $G\beta\gamma$ subunit in the model. The G proteins are modeled as simply switching between an active and inactive form, requiring only four parameters: the rates of activation, $k_{fG_{\text{off}}}$ and $k_{fG_{\text{i}}}$, and deactivation of a G protein, $k_{rG_{\text{off}}}$ and $k_{rG_{\text{i}}}$. Basal Da levels, which do not activate the D₁Rs, are modeled with $k_{fG_{\text{off}}} = 0 \text{ s}^{-1}$, during which inactive G_{off} is not converted to its active form. During a Da \uparrow , $k_{fG_{\text{off}}} = 5 \text{ s}^{-1}$, producing active $G\alpha_{\text{off}}$. Conversely, basal

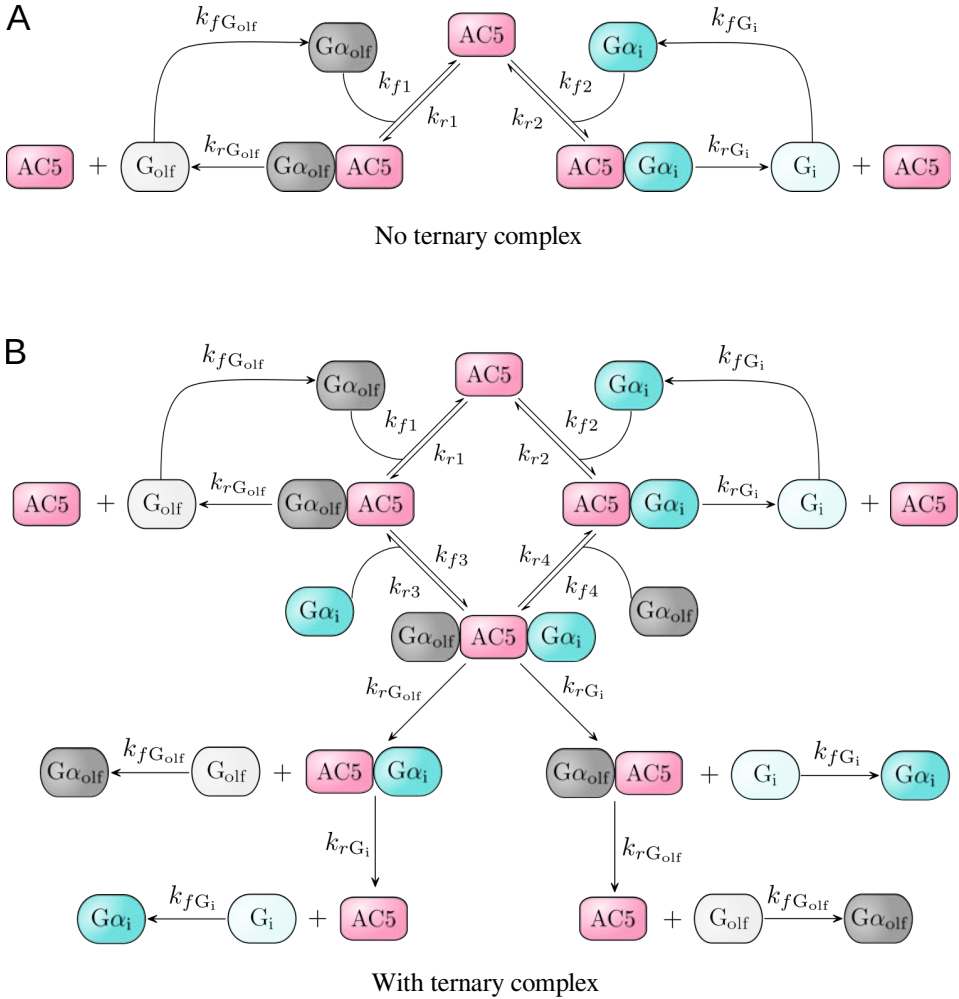


Figure 3.1: The full kinetic models of the AC5 signal transduction network. The model with no ternary complex is also called the null model.

ACh levels activate M_4 R_s, causing conversion of inactive G_i to active G_{α_i} , which is modeled with $k_{fG_i} = 5 \text{ s}^{-1}$. During an ACh \downarrow , $k_{fG_i} = 0 \text{ s}^{-1}$. The value for the rate of G_{α_i} activation is chosen so that there is a high resting-level inhibition of AC5 by the tonic level of ACh, a reasonable assumption based on recent experimental results (Nair et al. 2019). The value for $k_{fG_{\text{olf}}}$

is, similarly, chosen to achieve concentrations of active G_{olf} high enough to drive the binding reactions with AC5 and $AC5 \cdot G_{\alpha_i}$ forward. The total concentrations¹ of AC5 and the two G proteins used for this model are $n_{AC5} = 1500$ nM, $n_{G_{\text{olf}}} = 1500$ nM and $n_{G_i} = 6000$ nM.

The activated $G_{\alpha_{\text{olf}}}$ and G_{α_i} interact with AC5 and can form each of the binary complexes and the ternary complex. Their deactivation is done by the intrinsic GTPase activity of the G proteins themselves, but is increased by AC5 for the case of G_{α_s} (a homologue to $G_{\alpha_{\text{olf}}}$) at least fivefold and the regulator of G protein signaling 9-2 (RGS9-2) for G_{α_i} 20 to 40 times (Xie et al. 2015; Scholich et al. 1999), because of which we have used values of $k_{rG_{\text{olf}}} = 0.5$ s⁻¹ and $k_{rG_i} = 5$ s⁻¹. The slower rate of $G_{\alpha_{\text{olf}}}$ deactivation than that of G_{α_i} deactivation results in a longer-lived intracellular $G_{\alpha_{\text{olf}}}$ than the initiating signal, Da \uparrow , while the drop in G_{α_i} is as brief as the extracellular initiating signal, the ACh \downarrow . If deactivated, the G proteins unbind from AC5.

For the reverse (unbinding) rates of the G proteins from the binary AC5 complexes, we use values 100 times greater than the corresponding forward rate constants, which is the order of magnitude fitted in Chen-Goodspeed et al. (2005). The reverse rates of G protein unbinding from the ternary complex are increased by an order of magnitude compared to the reverse rates for unbinding from the respective binary complexes to qualitatively incorporate possible reduced stability of the ternary complex compared to the binary complexes on longer time scales indicated by the MD simulations (Bruce et al. 2019).

3.2 Measures of coincidence detection

We studied the ability of AC5 to perform coincidence detection in dSPNs, which means to respond with significant amounts of cAMP only when the two incoming signals Da \uparrow and ACh \downarrow coincide in time and in the spatial vicinity of the receptors. Note that there are two aspects of coincidence detection in the definition:

1. Distinguishing between the inputs Da \uparrow + ACh \downarrow and a Da \uparrow or ACh \downarrow alone, and
2. Responding strongly enough with amounts of cAMP that are physiologically relevant.

¹Within the kinetic model, we also refer to concentrations as amounts.

To quantify how well the network distinguishes between the different inputs, we use a quantity called synergy, and to quantify the strength of the response, we use the average catalytic rate, both defined below.

3.2.1 Average catalytic rate

The average catalytic rate k_c is an average of the catalytic rates of each enzyme species. For the null model in Fig. 3.1, where the ternary complex does not form, these are AC5, AC5 · G α_{olf} , and AC5 · G α_i , and k_c is:

$$k_c = p_1 k_{c,\text{AC5}} + p_2 k_{c,\text{AC5}\cdot\text{G}\alpha_{\text{olf}}} + p_3 k_{c,\text{AC5}\cdot\text{G}\alpha_i} \quad (3.2)$$

with the weights

$$p_1 = \frac{[\text{AC5}]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i]}, \quad (3.3)$$

$$p_2 = \frac{[\text{AC5} \cdot \text{G}\alpha_{\text{olf}}]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i]}, \quad (3.4)$$

$$p_3 = \frac{[\text{AC5} \cdot \text{G}\alpha_i]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i]}. \quad (3.5)$$

being the amounts of each enzyme species as a fraction of the total concentration of AC5 in the system. For the kinetic model with the ternary complex, the enzyme species are AC5, AC5 · G α_{olf} , AC5 · G α_i , and G α_{olf} · AC5 · G α_i , and the average catalytic rate is:

$k_c = q_1 k_{c,\text{AC5}} + q_2 k_{c,\text{AC5}\cdot\text{G}\alpha_{\text{olf}}} + q_3 k_{c,\text{AC5}\cdot\text{G}\alpha_i} + q_4 k_{c,\text{G}\alpha_{\text{olf}}\cdot\text{AC5}\cdot\text{G}\alpha_i}$ with the proportions of each enzyme species given with the weights

$$q_1 = \frac{[\text{AC5}]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i] + [\text{G}\alpha_{\text{olf}} \cdot \text{AC5} \cdot \text{G}\alpha_i]}, \quad (3.6)$$

$$q_2 = \frac{[\text{AC5} \cdot \text{G}\alpha_{\text{olf}}]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i] + [\text{G}\alpha_{\text{olf}} \cdot \text{AC5} \cdot \text{G}\alpha_i]}, \quad (3.7)$$

$$q_3 = \frac{[\text{AC5} \cdot \text{G}\alpha_i]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i] + [\text{G}\alpha_{\text{olf}} \cdot \text{AC5} \cdot \text{G}\alpha_i]}, \quad (3.8)$$

$$q_4 = \frac{[\text{G}\alpha_{\text{olf}} \cdot \text{AC5} \cdot \text{G}\alpha_i]}{[\text{AC5}] + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}] + [\text{AC5} \cdot \text{G}\alpha_i] + [\text{G}\alpha_{\text{olf}} \cdot \text{AC5} \cdot \text{G}\alpha_i]}. \quad (3.9)$$

We assume that the catalytic rate of the unbound AC5 is $k_{c,\text{AC5}} = 0.1 \text{ s}^{-1}$, and this rate is scaled by factors of $\alpha_{\text{G}\alpha_{\text{olf}}}$ and $\alpha_{\text{G}\alpha_i}$ when the respective

regulator $G\alpha$ subunit binds:

$$k_{c,AC5 \cdot G\alpha_{olf}} = \alpha_{G\alpha_{olf}} k_{c,AC5} \quad (3.10)$$

$$k_{c,AC5 \cdot G\alpha_i} = \alpha_{G\alpha_i} k_{c,AC5} \quad (3.11)$$

The factors of stimulation and inhibition of AC5 are set to be $\alpha_{G\alpha_{olf}} = 200$ and $\alpha_{G\alpha_i} = 0.01 \text{s}^{-1}$ (Chen-Goodspeed et al. 2005). For the catalytic rate of the ternary complex, we use the result of the MD simulations from section 2.12.1 that the ternary complex is inactive, $\alpha_{G\alpha_{olf},G\alpha_i} = \alpha_{G\alpha_i}$, i.e.:

$$k_{c,G\alpha_{olf} \cdot AC5 \cdot G\alpha_{olf}} = \alpha_{G\alpha_i} k_{c,AC5}, \quad (3.12)$$

except in Figs. 4.2D–F, where the catalytic rate of the ternary complex is varied to investigate its effect on coincidence detection. (A technically nonzero value for α_{G_i} is convenient to vary the ternary complex catalytic rate logarithmically in Figs. 4.2D–F. However, the value $\alpha_{G_i} = 0.01 \text{s}^{-1}$ makes $AC5 \cdot G\alpha_i$ and $G\alpha_{olf} \cdot AC5 \cdot G\alpha_i$ practically inactive.)

3.2.2 Synergy

The synergy of two signals s_1 and s_2 is defined as

$$S(t) = \frac{r(s_1, s_2, t)}{r(s_1, t) + r(s_2, t) - r_{ss}}, \quad (3.13)$$

where r is the response to a given signal and r_{ss} is the response at steady state. This quantity measures how much stronger the response r of the signaling network is for two coincident signals compared to the summed responses for single signals. The synergistic effect of the input signals can be calculated for any molecule in the signaling network that is affected by the inputs. For example, in the large model of the LTP signaling network in Nair et al. (2015), it was calculated for PKA (shown in Fig. 2.7). Not having included PKA or cAMP in the minimal kinetic model, we use the average catalytic rate k_c as a proxy for the amount of cAMP produced, since the latter directly depends on k_c . That is, the synergy of coincident neuromodulatory signals $Da \uparrow + ACh \downarrow$ is:

$$S(t) = \frac{k_c(Da \uparrow + ACh \downarrow, t)}{k_c(Da \uparrow, t) + k_c(ACh \downarrow, t) - k_{c,ss}}. \quad (3.14)$$

$S(t) > 1$ indicates a supralinear response in the presence of the two coincident signals, $S(t) = 1$ indicates a linear response to the coincident signals, and $S(t) < 1$ is a sublinear response. Hence, the case $S(t) > 1$ indicates that the signaling network can perform coincidence detection, whereas $S(t) \leq 1$ indicates an inability to do so.

Example traces for k_c and the corresponding synergy are given in Fig. 3.2. Using

$$\Delta = k_c(\text{Da } \uparrow + \text{ACh } \downarrow, t) - (k_c(\text{Da } \uparrow, t) + k_c(\text{ACh } \downarrow, t) - k_{c,ss}) \quad (3.15)$$

to express the difference between the response of the network for two coincident signals and the response for single signals, the expression for the synergy can also be rewritten as

$$S(t) = \frac{k_c(\text{Da } \uparrow + \text{ACh } \downarrow, t)}{k_c(\text{Da } \uparrow, t) + k_c(\text{ACh } \downarrow, t) - k_{c,ss}} = 1 + \frac{\Delta}{k_c(\text{Da } \uparrow, t) + k_c(\text{ACh } \downarrow, t) - k_{c,ss}}. \quad (3.16)$$

The difference in the responses, Δ , determines how big the synergistic effect of the input signals is (Fig. 3.2).

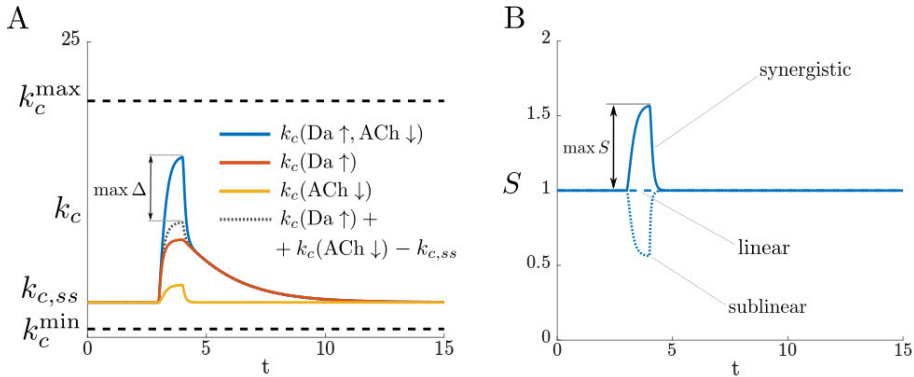


Figure 3.2: (A) A cartoon of the average catalytic rate during each of the input situations $\text{Da } \uparrow + \text{ACh } \downarrow$, $\text{Da } \uparrow$, and $\text{ACh } \downarrow$. (B) The synergy corresponding to panel A. In addition, the dashed and dotted lines illustrate how the synergy for a linear and a sublinear response would look like, respectively. Figure and caption reproduced from Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

Figure 3.2 is an example depicting how k_c relates to the synergy. There are minimum and maximum bounds on k_c : it would attain the minimum value k_c^{\min} if all of AC5 were inhibited by $G\alpha_i$, that is, only the catalytically inactive complex $AC5 \cdot G\alpha_i$ exists, where k_c^{\min} is the catalytic rate of $AC5 \cdot G\alpha_i$ (Eq. 3.11). Analogously, it would attain the maximum value k_c^{\max} if all of AC5 were bound in the catalytically active complex $AC5 \cdot G\alpha_{\text{off}}$, where k_c^{\max} is its catalytic rate (Eq. 3.10). In the models we use in this study, AC5 is never fully occupied by either of the $G\alpha$ subunits, and hence k_c is always between the minimum and maximum bounds. To maximize Δ , one would need to have k_c grow towards the maximum as much as possible during $Da \uparrow + ACh \downarrow$, and to have it as low as possible at steady state, and when each of the signals $Da \uparrow$ and $ACh \downarrow$ comes alone.

A combination of the average catalytic rate and the synergy gives a metric which can be helpful to evaluate whether the network both distinguishes between the different input situations and responds strongly. We use a product of k_c and S for this purpose, $C = Sk_c$. This metric scales the average catalytic rate by the amount of synergy.

3.3 Membrane voltage and Hodgkin-Huxley models of ion channels

Electrical signals are the fastest signals that neurons employ for signal processing/integration and for interneuronal communication, and are the most common subject of experimental and computational study when studying these two neuronal functions. Computational models of the changes in membrane voltage describe a piece of membrane with an equivalent electrical circuit, which contains the membrane capacitance, the membrane permeability, or leak ionic conductance, and the ionic conductances arising from any ion channels present in the membrane (Fig. 3.3).

For the single membrane compartment in Fig. 3.3 where there are no spatial variations in the membrane voltage, the currents flowing in and out of it satisfy the Kirchoff current law, which is a charge conservation law stating that all currents flowing in and out of a single node sum to 0:

$$I_c + I_l + \sum_{\text{ion}} I_{\text{ion}} = 0,$$




Figure 3.3: An equivalent electrical circuit of a piece of neuronal membrane. This electrical circuit contains components describing the membrane capacitance, the leak ionic conductance, and the time-variable, voltage-gated Na^+ and K^+ ionic conductances. The membrane voltage is equal to the difference of the intra- and extracellular electrical potentials, $v = V_i - V_o$. Usually, V_o is taken as a reference potential, or electrical ground, i.e. $V_o = 0$, so v is equal to V_i .

$$I_c = C_m \frac{dv}{dt},$$

$$I_l = g_l(v - E_l),$$

$$I_{\text{ion}} = g_{\text{ion}}(v - E_{\text{ion}}),$$

where v is the membrane voltage, I_c is the capacitive current, I_l is the leak current, I_{ion} is a voltage-gated ionic current, C_m is the membrane capacitance, g_l is the leak conductance and g_{ion} is the voltage-gated ionic conductance for a given ion, E_l is the leak reversal potential, and E_{ion} is the reversal potential for a given ion. The change of membrane voltage with respect to time appears in the capacitive membrane current, and rearranging the Kirchoff current law gives the differential equation which describes the time evolution of membrane voltage, which is numerically solved in computational neuron models to obtain the voltage dynamics as a result of the

various ion channels and leak current present in the membrane.

$$C_m \frac{dv}{dt} = -g_{\text{ion}}(v - E_{\text{ion}}) - g_l(v - E_l) \quad (3.17)$$

The leak current conductance is constant or slow-changing. On the other hand, the conductance of the ion channels is often voltage-dependent, and can be highly variable as the membrane voltage goes in and out of the operational voltage range of the ion channel. The voltage dependence of the ionic conductances is phenomenologically modeled with the Hodgkin-Huxley formalism, which we present through the examples of the Na^+ and K^+ channels involved in the generation of the action potential. Hodgkin-Huxley ionic conductances are expressed in terms of gating variables, which are voltage-dependent and thus time-variable. The gating variables have values from 0 to 1, and scale a maximal conductance, describing the proportion of open channels at a particular voltage value.

$$g_{\text{Na}} = \bar{g}_{\text{Na}} m^3 h \quad (3.18)$$

$$g_{\text{K}} = \bar{g}_{\text{K}} n^4 \quad (3.19)$$

$$\frac{dm}{dt} = \frac{m - m_\infty}{\tau_m} \quad (3.20)$$

$$\frac{dh}{dt} = \frac{h - h_\infty}{\tau_h} \quad (3.21)$$

$$\frac{dn}{dt} = \frac{n - n_\infty}{\tau_n} \quad (3.22)$$

The dynamics of the Na^+ conductance is described by two gating variables, m and h . m increases with increasing voltage, and h decreases, and their product gives a voltage range where the Na^+ channel is open. The dynamics of the K^+ conductance is described by a single gating variable, n . As membrane voltage changes, the gating variables usually do not instantaneously reach their new values, but do so with a delay (τ_m, τ_h, τ_n), which is also voltage-dependent. When the Hodgkin-Huxley formalism was developed, the composition of the neuronal membrane was not known, so the gating variables do not represent any physical quantities. Later interpretations of the Hodgkin-Huxley formalism have suggested that the gating variables can be viewed as a macroscopic description of the probability that an ion channel protein is open for the given value of the membrane voltage, and the time constants describe the time it takes for the protein to change conformation for various voltage values.

Neuron models consisting of a single membrane compartment are used to approximate the electrical behavior of a whole neuron with a single point, usually representing the voltage at the soma.

3.4 Multicompartmental neuron models

Not all parts of a neuron have the same membrane voltage, nor do they contain the same ion channels. Multicompartmental neuron models are used to describe such spatial variations in neurons. They are formed by connecting single membrane compartments, which describe isopotential regions of the neuron, with axial resistances that model the cytoplasmic resistivity (Fig. 3.4). We use multicompartmental neuron models of SPNs in the NEURON simulator developed by Lindroos and Hellgren Kotaleski (2021). These multicompartmental neuron models incorporate a detailed description of the neuronal morphology in order to describe how changes in the membrane voltage propagate through the spatial structure of the neuron.




Figure 3.4: An equivalent electrical circuit of a piece of membrane modeled with three membrane compartments. Each compartment has different ionic conductances. A single compartment describes an isopotential region of the membrane, i.e. where the voltage does not vary (noticeably) with distance. Compartment $i + 1$ contains only the membrane capacitance and a leak conductance, and it is also called a passive membrane. Compartment i in addition contains the voltage-dependent (and thus, time-variable) Na^+ and K^+ conductances, while compartment $i - 1$ contains a voltage-dependent Ca^{2+} conductance. Neighboring compartments are connected through the axial membrane resistance.

In mathematical terms, the propagation of voltage in a one-dimensional

cable is described by the cable equation, which has the same form as the one-dimensional diffusion equation:

$$\lambda^2 \frac{\partial^2 v}{\partial x^2} = \tau \frac{\partial v}{\partial t} + v \quad (3.23)$$

In this equation, both space and time are continuous variables. Multicompartmental neuron models approximate the spatial solution to the cable equation by discretizing space. Having multiple membrane compartments connected to each other is in fact the spatial discretization of the cable equation. For example, if we consider three membrane compartments $i - 1$, i , $i + 1$, in the middle of a dendrite, the voltage in the middle compartment is given by Eq. 3.17 with two additional terms for the axial currents with the adjacent compartments $i - 1$ and $i + 1$:

$$C_{mi} \frac{dv_i}{dt} + I_{ion} + I_l = g_{a,i-1}(v_{i-1} - v_i) + g_{a,i+1}(v_{i+1} - v_i) \quad (3.24)$$

If the compartments have a length of Δx , the membrane capacitance is $C_m = c_m \pi d \Delta x$, the axial resistivity is $R_a = \frac{4r_a \Delta x}{\pi d^2}$, the axial conductances are the reciprocal of the axial resistivity $g_{a,i-1} = g_{a,i+1} = \frac{\pi d^2}{4r_a \Delta x}$, where c_m is the specific membrane capacitance (capacitance per unit of length) and r_a is the specific axial resistivity (resistivity per unit of length). Inserting these in Eq. 3.24, and dividing it by $\pi d \Delta x$ gives:

$$c_m \frac{dv_i}{dt} + i_{ion} + i_l = \frac{d}{4r_a} \frac{v_{i-1} - 2v_i + v_{i+1}}{\Delta x^2}, \quad (3.25)$$

where i_{ion} and i_l are current densities (currents per membrane area). For infinitesimally small compartment sizes, the term on the right hand side of Eq. 3.25 is the second order partial derivative of the voltage with respect to space in the cable equation:

$$c_m \frac{\partial v_i}{\partial t} + i_{ion} + i_l = \frac{d}{4r_a} \frac{\partial^2 v}{\partial x^2}$$

Multiplying the last equation by the membrane resistance r_m (the inverse of the leak membrane conductance) and recognizing that $(i_{ion} + i_l)r_m = v$, as well as using $\tau = c_m r_m$ and $\lambda^2 = \frac{d r_m}{4r_a}$, we obtain the cable equation again:

$$\tau \frac{\partial v}{\partial t} + v = \lambda^2 \frac{\partial^2 v}{\partial x^2}, \quad (3.26)$$

showing that a multicompartment dendrite model is a spatial approximation to the cable equation.

For multicompartment neuron models, NEURON solves the coupled differential equations of the form Eq. 3.24 for each compartment in the model. Numerically integrating Eq. 3.24 for each compartment is in fact discretization of the cable equation with respect to time. NEURON offers several numerical solvers for differential equations, of which we have used the `cnexp` solver, which is second order accurate and computationally efficient for stiff systems.

Multicompartmental neuron models are constructed in the following steps:

1. Digital reconstruction of the neuronal morphology. Several experimental imaging techniques can provide images of the neuronal morphology, which is converted into a 3D model by using specialized software. In NEURON, dendrites are represented by cylindrical sections usually corresponding to dendritic stretches where the diameter is almost constant. Dendritic regions with different diameters are represented with different sections, and sections are connected to each other according to the neuronal morphology. Each section can have multiple electrical compartments. The question of choosing the appropriate compartment size, or equivalently, the number of compartments in each section, is very important in order to minimize spatial discretization errors with respect to the analytical solution of the cable equation.
2. Distribution of ion channels in the soma and dendrites. All of the experimentally detected ion channels in the neuron are described as Hodgkin-Huxley based conductance models and are inserted in the soma and dendrites.
3. Parameter fitting. The maximal conductances of the ion channels are fitted with an optimization method such that the electrical behavior of the neuron matches experimental data.
4. Model validation. The constructed model behavior is validated on experimental data which was not used for parameter fitting.

We have made additions to the SPN models from Lindroos and Hellgren Kotaleski 2021. Spines in the neurons are added as additional compartments consisting of a neck and head with lengths and diameters $l_{\text{neck}} = 0.5 \mu\text{m}$, $l_{\text{head}} = 0.5 \mu\text{m}$, and $d_{\text{neck}} = 0.125 \mu\text{m}$, $d_{\text{head}} = 0.5 \mu\text{m}$, respectively. Axial

resistance in all compartments is $150 \Omega \cdot \text{cm}$, except for the spine neck, where it is $1130 \Omega \cdot \text{cm}$ (Dorman et al. 2018). Spines contain the inwardly rectifying K^+ channel, with a conductance equal to that of the parent dendritic shaft segment. The voltage-gated Ca^{2+} channels $\text{Ca}_v1.2$, $\text{Ca}_v1.3$ (type L), $\text{Ca}_v2.3$ (type R), $\text{Ca}_v3.2$ and $\text{Ca}_v3.3$ (type T), which are present in the dendrites, are added to the spines as well, and their conductances have been manually tuned to match the relative proportions determined in Carter and Sabatini (2004), and Higley and Sabatini (2010), as well as the Ca^{2+} concentration amplitudes arising from backpropagating action potential (bAP) stimulation as in Fig. 2 of Shindou et al. (2011).

3.5 Calcium dynamics

The multicompartment SPN models we have used also track the intracellular concentration of Ca^{2+} ions, $[\text{Ca}^{2+}]_i$, which enter through the voltage-gated Ca^{2+} channels and the NMDARs. As introduced in sections 2.7, 2.8, and 2.12.2, LTP and LTD are triggered by Ca^{2+} from different sources. Because of this, we also tracked $[\text{Ca}^{2+}]_i$ from these sources separately in the model. In the study of the synaptic plasticity rule, we used axial Ca^{2+} diffusion for all Ca^{2+} sources except for Ca^{2+} from $\text{Ca}_v1.2$ and $\text{Ca}_v1.3$ (L-type) channels, which we assumed is localized in a nanodomain (Simms and Zamponi 2014). Axial diffusion was implemented using the rxd module for NEURON, with four Ca^{2+} buffers and a Ca^{2+} pump, according to the detailed Ca^{2+} diffusion model in Dorman et al. (2018). Since fitting a large Ca^{2+} diffusion model is a project in its own right, we have used the buffer parameters from Dorman et al. (2018) with the rationale that the idealized SPN morphology used in that study should describe most biological SPN morphologies, as the one we are using. We modified the Ca^{2+} pump parameters to obtain qualitatively similar results for $[\text{Ca}^{2+}]_i$ as in Dorman et al. (2018). The parameters for Ca^{2+} diffusion are given in Table 3.1. Throughout the study of the synaptic plasticity rule, the diffusing Ca^{2+} is referred to as $[\text{Ca}^{2+}]_{\text{NMDA}}$, since the contribution of Ca^{2+} from NMDARs is largest, and since NMDA Ca^{2+} is important for triggering LTP. $[\text{Ca}^{2+}]_{\text{NMDA}}$ is used to trigger LTP in the synaptic plasticity rule, and since Fino et al. (2010) have reported that voltage-gated Ca^{2+} channels are also important for LTP, they are included in $[\text{Ca}^{2+}]_{\text{NMDA}}$. Nevertheless, with a refitting of the diffusion model it is possible to have only NMDA Ca^{2+} as the source that triggers LTP.

Buffers	Calbindin	CaMN	CaMC	Immobile
k_f (mM/ms)	28	100	6	400
k_r (/ms)	$19.6 \cdot 10^{-6}$	10^{-3}	$9 \cdot 10^{-6}$	40
Concentration (mM)	0.08	0.015	0.015	2.5
D ($\mu\text{m}^2/\text{s}$)	66	66	66	0
Pump	PMCA			
k_t [mM/(μm^2 ms)]	5000			
K_D (mM)	0.3			

Table 3.1: The parameters for the axial diffusion model from Dorman et al. (2018). We have changed only the catalytic activity of the pump.

$[\text{Ca}^{2+}]_{\text{L-type}}$, which triggers LTD, is described by a phenomenological "pool" model, according to which all Ca^{2+} from $\text{Ca}_v1.2$ and $\text{Ca}_v1.3$ currents integrates in a single variable. The model is implemented according to Wolf et al. (2005), and Ca^{2+} dynamics are calculated separately for each compartment according to Eq. 3.27. This model calculates $[\text{Ca}^{2+}]_i$ in a thin cylindrical shell with depth d under the membrane, and $[\text{Ca}^{2+}]_i$ decays with a time constant τ_{Ca} , modeling Ca^{2+} diffusion only phenomenologically, i.e no Ca^{2+} ions cross single dendritic compartments. The pool model differs from the simplest pool models in that it also includes the effect of a Ca^{2+} pump that extrudes intracellular Ca^{2+} to the extracellular space.

$$\frac{d[\text{Ca}^{2+}]_i}{dt} = \frac{[\text{Ca}^{2+}]_\infty - [\text{Ca}^{2+}]_i}{\tau_{\text{Ca}}} + \frac{k}{2Fd} \sum_{x \in \{1.2, 1.3\}} I_{\text{Ca}_v x} - pk_t \frac{[\text{Ca}^{2+}]_i}{[\text{Ca}^{2+}]_i + K_d} \quad (3.27)$$

In Eq. 3.27, F is the Faraday constant, k_t is the catalytic activity of the pump, K_d is the pump dissociation constant for Ca^{2+} , and k and p are phenomenological parameters used in Wolf et al. (2005) to balance Ca^{2+} influx and efflux.

In the study of glutamate spillover, there is also no explicit diffusion of Ca^{2+} between neighboring compartments implemented in the model. Instead, $[\text{Ca}^{2+}]_i$ is modeled with the pool model in Eq. 3.27, with two separate pools for $[\text{Ca}^{2+}]_{\text{NMDA}}$ and $[\text{Ca}^{2+}]_{\text{L-type}}$. In this study, $[\text{Ca}^{2+}]_i$ is not used for any purpose, but we report it since it is implemented in the model. The

parameters for the $[\text{Ca}^{2+}]$ pool models are given in Table 3.2.

Pool	$[\text{Ca}^{2+}]_{\infty}$ (nM)	τ_{Ca} (ms)	k_t (mM/ms)	K_D (mM)	d (μM)	k	p
$[\text{Ca}^{2+}]_{\text{NMDA}}$	50	100	10^{-4}	10^{-4}	0.1	10000	0.02
$[\text{Ca}^{2+}]_{\text{L-type}}$	50	100	10^{-4}	10^{-4}	0.1	10000	0.02

Table 3.2: The parameters for the $[\text{Ca}^{2+}]$ pool models. Parameter values are the same for both models.

Lastly, the default extracellular Ca^{2+} concentration, $[\text{Ca}^{2+}]_o$, is different in NEURON and in the rxd module. We did not change the default values, which is why in the synaptic plasticity study, where the rxd module is used, $[\text{Ca}^{2+}]_o = 1$ mM, while in the glutamate spillover study, where diffusion is not implemented, $[\text{Ca}^{2+}]_o = 2$ mM. Among other things, $[\text{Ca}^{2+}]_o$ also affects the reversal voltage for the NMDA synapses, as will be mentioned in the next section.

3.6 Synaptic models

3.6.1 Dual exponential synaptic model (difference of two exponentials)

To represent background synaptic noise, excitatory and inhibitory synapses are distributed across the dendrites according to the synapse densities reported in Cheng et al. (1997), and they are not located on spines. The synaptic models for both excitatory and inhibitory synapses are represented by a difference of two exponential functions (also called the dual exponential synaptic model), which, for a spike arriving at time t_0 is represented by:

$$g_{syn}(t) = wg_{\max} \frac{1}{C} \frac{\tau_A \tau_B}{\tau_B - \tau_A} \left(\exp\left(-\frac{t-t_0}{\tau_B}\right) - \exp\left(-\frac{t-t_0}{\tau_A}\right) \right)$$

where g_{\max} is the maximal synaptic conductance, and τ_A and τ_B are the rise and decay time constants. C is a normalization factor given with:

$$C = \frac{\tau_A \tau_B}{\tau_B - \tau_A} \left(\exp\left(-\frac{t^*}{\tau_B}\right) - \exp\left(-\frac{t^*}{\tau_A}\right) \right)$$

$$t^* = \frac{\tau_A \tau_B}{\tau_B - \tau_A} \ln \frac{\tau_B}{\tau_A}$$

where t^* is the position of the maximum of the dual exponential function. This normalization ensures that the time varying part of the dual exponential reaches a peak value of 1, which is then scaled by the maximal conductance g_{\max} with units of nanosiemens (nS). The synaptic weight, w , is a non-negative, dimensionless parameter which scales the maximal synaptic conductance. It is modified by synaptic plasticity, and changes the resulting synaptic conductance. We also use the name synaptic strength for the product of the weight and the maximal conductance wg_{\max} , which has units of nS, to differentiate it from the dimensionless synaptic weight.

The synaptic current arising from the conductance g_{syn} is:

$$I_{syn} = g_{syn} (v - E_{syn}) \quad (3.28)$$

where E_{syn} is the synapse's reversal voltage. In NEURON the model is in fact implemented by a scheme integrating two first order differential equations, which allows for convenient integration of spikes arriving at any time.

Excitatory synapses have both AMPA and NMDA components, and NMDA synapses have a Mg^{2+} block described by a sigmoidal function:

$$I_{NMDA} = g_{NMDA} (V - E_{rev}) \sigma; \quad \sigma = \frac{1}{1 + \eta \exp(-\alpha V)}$$

with parameters $\eta = 0.38$ and $\alpha = 0.062$. The parameters of the synaptic models are given in Table 3.3.

	AMPA	NMDA	inhibitory
τ_A (ms)	1.9	2.76	0.1
τ_B (ms)	4.8	115.5	10
E_{rev} (mV)	0	0	-60
g_{\max} (nS)	1.5	3.5	1.5

Table 3.3: The parameters for the dual exponential synaptic models describing the synapses for the background noise.


Because our multicompartment neuron model contains less electrical compartments than the number of synapses reported in Cheng et al. (1997),

this implies that several synapses will arrive in a single compartment. Since their synaptic inputs will be integrated in the same voltage variable, we have represented such multiple synapses with a single synapse instead, whose input frequency is scaled by the number of synapses it represents. This is why we have used a non-saturating synapse model such as the dual exponential model for these synapses.


3.6.2 Saturating synapse models

Clustered excitatory synapses located in particular dendritic branches are represented with saturating synapse models taken from Gao et al. (2021), which are a variation of the saturating synapse models in Destexhe et al. (1994). We chose saturating synapse models in this case since the synaptic clusters receive temporally coincident inputs comprised of three spikes each in one of the studies, and we wanted to avoid single excitatory synapses reaching unrealistically high conductance values.

The saturating synapse models are kinetic models which operate according to two different kinetic schemes depending on the presence of neurotransmitters. When neurotransmitters arrive at the postsynaptic site, receptor dynamics are described by a kinetic scheme that models switching between closed (C) and open (O) states with rates of α and β :



where $[T]$ is the neurotransmitter concentration. After the neurotransmitter has been cleared from the synaptic cleft, receptor dynamics evolves according to a kinetic scheme that describes just the closing of the receptor with a rate β :



When a presynaptic spike arrives, neurotransmitter levels are assumed to always reach a fixed saturating concentration, $[T]_{\max}$, in the synaptic cleft, i.e., are represented by a pulse with amplitude $[T]_{\max}$ and duration of T_{dur} . This “short pulse” representation of glutamate concentration in the synaptic cleft follows from findings in cultured hippocampal synapses showing that glutamate has a very fast time course, rapidly reaching concentrations of 1 mM, with a decay constant of 1.2 ms (Clements et al. 1992). If another

presynaptic spike arrives while the neurotransmitter pulse is still “on”, the pulse duration is lengthened by T_{dur} . In this model, the synaptic conductance is represented by the fraction of glutamate receptors in the open state at a given time. The synaptic current in this model is also given with Eq. 3.28, where g_{syn} is given with

$$g_{\text{syn}} = wg_{\text{max}}[O].$$

The parameters for the saturating synapse models are given in Table 3.4.

	AMPA	NMDA	eNMDA
E_{rev} (mV)	0	0	0
g_{max} (nS)	1.5	3.5	3.5
T_{dur} (ms)	1	1	$50 + 200w$
$[T]_{\text{max}}$ (mM)	1	1	0.2
α [/(ms mM)]	12.5	4	4
β (/ms)	0.25	0.01	0.01
K_d (μM)	20	2.5	2.5

Table 3.4: Parameters for the saturating synapse models used in the clustered synaptic inputs. $[\text{Ca}^{2+}]_o$ affects the reversal voltage of NMDA synapses (Arvanov and Wang 1997). As mentioned in section 3.5, the default $[\text{Ca}^{2+}]_o$ values are different in the models with and without diffusion. Because of this, for the NMDA and eNMDA synapses in the spillover study, we used we used $E_{\text{rev}} = 15$ mV, a value which corresponds to $[\text{Ca}^{2+}]_o = 2$ mM in that study. These discrepancies in the settings between the models with and without diffusion remain to be corrected in future work.

3.7 Extrasynaptic glutamate receptors

To model glutamate spillover, we have placed extrasynaptic NMDA receptors (eNMDARs) in the dendritic shaft directly under each clustered synapse situated on a spine. eNMDARs are modeled by the same kinetic scheme described in the previous section for saturating synapse models, with three differences: (i) T_{dur} is much longer, modeling the spillover effect of glutamate in the extrasynaptic space, (ii) the synaptic spike there arrives with a delay of 1.5 ms after the glutamate threshold has been reached (according to

Szapiro and Barbour (2007)), spillover currents arise after 1–2 ms of synaptic currents), and (iii) T_{\max} is lower for eNMDARs than in the synaptic cleft (lower extrasynaptic glutamate concentrations have been reported, and also spillover-induced currents are very sensitive to low-affinity AMPA antagonists, suggesting lower extrasynaptic glutamate concentrations (Szapiro and Barbour 2007; Okubo et al. 2010). The channel properties of NMDARs and eNMDARs have been reported to be similar (Clark et al. 1997). eNMDARs can be both diffusely and punctually distributed along dendritic shafts, and according to electrophysiological measurements, their conductance is 20–60% of the total NMDAR conductance (Harris and Pettit 2007; Petralia 2012; Papouin and Oliet 2014). Because we place eNMDARs uniformly in the same dendritic segment of 20–30 μm as the clustered spines, we effectively model a diffuse eNMDAR distribution (despite the fact that synapses are represented as punctual inputs in NEURON). We use the same NMDAR and eNMDAR conductance, giving an equal distribution of total NMDAR conductance between the spine and extrasynaptic receptors. The parameters for the eNMDARs are also given in Table 3.4.

3.8 Evoking plateau potentials

Plateau potentials are evoked by clustered synaptic input. In the study for glutamate spillover, where we compared plateau potential models with and without glutamate spillover, clusters ranged from 1 to 20 synapses for the models with spillover and from 1 to 40 synapses for the model without spillover. A cluster is placed on one of the dendrites, approximately 120–140 μm from the soma. The total NMDAR conductance is critical for generating nonlinear dendritic responses, so in order to meaningfully compare the cases with and without spillover, the same total NMDAR conductance needs to be used in both cases. The spine NMDAR conductance in both cases is the same. The case without spillover has no eNMDARs, meaning that, to match the total NMDAR conductance in the spillover scenarios, it was necessary to add more NMDAR conductance. This can be settled by either using a twice larger spine NMDAR conductance or a twice larger cluster size; we opted for the latter, which is why the clusters in the models without spillover range from 1 to 40 synapses. (However, note that this choice caused the total AMPAR conductance to be twice larger in the case without spillover compared to the case with spillover.) The glutamate spillover model that

we studied is proposed by Gao et al. (2021), and the differences in our usage of the model are detailed in Trpevski et al. (2023) (Article C in Part II of this thesis).

Clustered synaptic input consists of one presynaptic spike per synapse, all arriving randomly within a 30 ms interval. When no glutamate spillover is modeled, synaptic inputs arrive only to the spines. For glutamate spillover, the same presynaptic spike is delivered to the extrasynaptic NMDA conductance placed in the dendritic shaft under that spine. When delivering the presynaptic spike to the eNMDAR we explored two scenarios, which we call thresholded spillover and accumulative spillover. In thresholded spillover we assume glutamate spills over suddenly and simultaneously to all eNMDARs once a glutamate threshold is reached, with a delay of 1.5 ms. In accumulative glutamate spillover there is no threshold for glutamate, and a presynaptic spike delivered to a synaptic NMDA synapse is also delivered to the corresponding extrasynaptic NMDA conductance after a delay of 1.5 ms. This means that every synaptic activation causes spillover and activates the corresponding eNMDARs.

The glutamate threshold in thresholded spillover is implemented in the following way. Inspired by the fact that LTP causes withdrawal of the astrocytic processes enveloping a synapse, thus boosting glutamate spillover (Henneberger et al. 2020), we chose to base the threshold on the NMDA synaptic weight, w , which describes how strong the synapse is. The idea is that stronger synapses can cause spillover more easily than weaker synapses. Specifically, when a synapse in a cluster is activated, its weight is added to a variable associated with that cluster which conceptually models the concentration of glutamate. Added weights are normalized in the variable so that a cluster of at least ten synapses with a weight $w > 0.4$ needs to be activated to reach the glutamate threshold at which spillover occurs, and a spike is delivered to the eNMDARs. With this setting, spillover can also occur, for example, if a cluster of five synapses with weights $w > 0.8$ is activated, or a cluster of twenty synapses with weights $w > 0.2$ is activated. To implement the variable where weights are added, we have in fact used a NEURON integrate-and-fire cell with a very long time constant that generates a spike to the eNMDARs in the cluster once it reaches a threshold value of 1.

The two glutamate spillover scenarios we consider are idealized scenarios, and biological glutamate spillover is most likely somewhere in between, varying among brain regions.

In the study of the synaptic plasticity rule, the threshold for glutamate

spillover is set at a different level: 16 synapses with weight $w > 0.5$ need to be activated to evoke a spillover-induced plateau potential.

3.9 Synaptic plasticity

To study whether a single dSPN can solve the feature binding problem, we formulated a local, calcium- and reward-based learning rule. We avoided modeling the entire LTP and LTD signaling networks described in sections 2.7 and 2.8, but we have nevertheless based the learning rule on those networks, which require that the two branches in each network are activated. Supported by the experimental results from Shen et al. (2008), Yagishita et al. (2014), Fisher et al. (2017), and Fino et al. (2010), this means that Ca^{2+} from NMDA channels and voltage-gated Ca^{2+} channels, as well as dopamine activating the D_1Rs , are necessary to trigger LTP. Also, Ca^{2+} from $\text{Ca}_v1.3$ channels and glutamate activating mGluRs, together with D_1Rs free from dopamine, are necessary for LTD (Shen et al. 2008; Fino et al. 2010; Plotkin et al. 2013).

Additionally, not all Ca^{2+} levels can trigger synaptic plasticity. Namely, a lower threshold level of $[\text{Ca}^{2+}]_{\text{L-type}}$ for triggering LTD has been reported (Shindou et al. 2011). Because of this, we have implemented a lower threshold for both $[\text{Ca}^{2+}]_{\text{NMDA}}$ and $[\text{Ca}^{2+}]_{\text{L-type}}$, i.e. for both LTP and LTD the corresponding Ca^{2+} levels need to be above this threshold to trigger plasticity. The LTD requirement that glutamate activates mGluRs is implemented as simply requiring that a presynaptic spike arrives at a synapse, i.e. inactive synapses which have elevated $[\text{Ca}^{2+}]_{\text{L-type}}$ are not weakened.

Moreover, in order to solve the NFBP with plateau potentials as suggested in section 2.12.2, a relevant stimulus needs to trigger a plateau potential, but an irrelevant stimulus should not. Together with the background synaptic input, the plateau potentials should elicit somatic spiking, while the excitations from the irrelevant stimuli should not. Relevant and irrelevant stimuli represented on the same dendrite share synaptic clusters. For example, a yellow banana and a red banana on dendrite 1 in Fig. 2.16B both activate the synapses encoding the shape of the stimulus. Therefore, when strengthening synaptic clusters to elicit plateaus for the relevant stimulus, it is important not to strengthen them too much, as half of those synapses will also be activated by the irrelevant stimulus, also eliciting somatic spiking, even if no plateau is evoked. This means the learning rule should “know”

when to stop strengthening a synapse. The simplest way to endow the rule with such a decision capability, in our view, was to place an upper threshold for $[\text{Ca}^{2+}]_{\text{NMDA}}$ above which no LTP occurs. This threshold should be set to a $[\text{Ca}^{2+}]_{\text{NMDA}}$ level which indicates that plateau potentials have appeared (after potentially several iterations of synaptic strengthening). Assuming that plateaus will cause somatic spiking, there is no further need for synaptic strengthening, and synaptic updates should stop. As we will show in section 4.2, $[\text{Ca}^{2+}]_{\text{NMDA}}$ with and without plateaus differs by an order of magnitude, and also has an all-or-none property, i.e. a pronounced supra-linearity like the plateaus, meaning that $[\text{Ca}^{2+}]_{\text{NMDA}}$ is a good indicator of whether plateaus have appeared. Also, because of the all-or-none quality, a range of values for the upper threshold is possible.

The plasticity rule is given with:

$$\Delta w = \eta_E \Delta t \begin{cases} w_{\max} - w, & \text{Da } \uparrow \text{ and } \theta_{\text{LTP}} < [\text{Ca}^{2+}]_{\text{NMDA}} < \Theta_{\text{LTP}} \\ w - w_{\min}, & \text{Da } \downarrow \text{ and } [\text{Ca}^{2+}]_{\text{L-type}} > \theta_{\text{LTD}} \\ 0, & \text{no Da feedback} \end{cases}$$

The same weight is used in both AMPARs and NMDARs. During LTP the weight can increase to a maximal weight of w_{\max} , and during LTD it can decrease to a minimal weight of w_{\min} , with a learning rate of η_E . Δt is a time step, such as the integration time step in NEURON’s integrators. LTP happens after a $\text{Da } \uparrow$ is delivered and $[\text{Ca}^{2+}]_{\text{NMDA}}$ is between the lower and upper threshold levels θ_{LTP} and Θ_{LTP} , respectively. LTD happens when a $\text{Da } \downarrow$ is delivered and $[\text{Ca}^{2+}]_{\text{L-type}}$ is higher than the threshold level θ_{LTD} . The upper threshold for LTP, Θ_{LTP} , prevents the weight from saturating at the maximal level. Similarly, the threshold for LTD, θ_{LTD} , if it is not set too low, prevents the weight from falling to the minimal level. Each weight update depends on the current weight value, and the update becomes smaller as the weight approaches the maximal or minimal value. This is called a multiplicative weight update model, and is a commonly used formulation in synaptic plasticity rules (Gütig et al. 2003; Rubin et al. 2001; Rossum et al. 2000). The lower Ca^{2+} thresholds for plasticity, θ_{LTP} and θ_{LTD} , are also another commonly used feature in plasticity rules, although they are usually applied to the same Ca^{2+} source (Graupner and Brunel 2012; Jędrzejewska-Szmek et al. 2017). The additional component that we have introduced in the rule is the upper threshold for LTP, Θ_{LTP} .

We have also incorporated a simple metaplasticity rule for the Ca^{2+} thresholds θ_{LTP} and θ_{LTD} :

$$\begin{aligned}\Delta\theta_{\text{LTP}} &= \eta_{\theta}([\text{Ca}^{2+}]_{\text{NMDA}} - \theta_{\text{LTP}})\Delta t, \\ \Delta\theta_{\text{LTD}} &= \eta_{\theta}([\text{Ca}^{2+}]_{\text{L-type}} - \theta_{\text{LTD}})\Delta t,\end{aligned}$$

which are updated after both LTP and LTD, that is, either when $\text{Da} \uparrow$ and $[\text{Ca}^{2+}]_{\text{NMDA}} > \theta_{\text{LTP}}$ or when $\text{Da} \downarrow$ and $[\text{Ca}^{2+}]_{\text{L-type}} > \theta_{\text{LTD}}$. We do not take into account the upper threshold Θ_{LTP} in metaplasticity; it only plays a role in synaptic plasticity (where the condition for synaptic strengthening is $\text{Da} \uparrow$ and $\theta_{\text{LTP}} < [\text{Ca}^{2+}]_{\text{NMDA}} < \Theta_{\text{LTP}}$)². With this metaplasticity rule, each of the thresholds follows the latest $[\text{Ca}^{2+}]_{\text{NMDA}}$ and $[\text{Ca}^{2+}]_{\text{L-type}}$ levels with a delay determined by the rate η_{θ} . As we will show in section 4.2, metaplasticity in the LTP threshold prevents weakened synapses, which thus do not contribute to obtaining future dopamine rewards, to participate in LTP easily. Metaplasticity in the LTD threshold protects strengthened synapses, which are useful for obtaining future dopamine rewards, from easily undergoing LTD. Metaplasticity is very important for solving the NFBP. Without it synaptic strength is very variable and does not stabilize easily, contributing to worsened performance on the NFBP. This formulation of metaplasticity is similar to the one used in the Bienenstock-Cooper-Munro (BCM) rule³, where the threshold follows the average neuronal activity (Bienenstock et al. 1982). In our formulation of metaplasticity, the threshold follows the maximal Ca^{2+} level achieved during stimulus presentation. This implicitly assumes that the maximal $[\text{Ca}^{2+}]_{\text{NMDA}}$ and $[\text{Ca}^{2+}]_{\text{L-type}}$ amplitudes can be remembered by the synaptic machinery. Other quantities such as the average or the integral of $[\text{Ca}^{2+}]$ can be used instead of the maximum; however, we did not try these alternatives. The parameters for the excitatory plasticity rule are given in Table 3.5.

We used two different stimulation protocols for "training" the SPN on the NFBP (Fig. 3.5). One of the protocols was used during in the so-called

²We have also tried two options where Θ_{LTP} influences metaplasticity, as well: i) the condition $\theta_{\text{LTP}} < [\text{Ca}^{2+}]_{\text{NMDA}} < \Theta_{\text{LTP}}$ is used only when updating θ_{LTP} , ii) the condition $\theta_{\text{LTP}} < [\text{Ca}^{2+}]_{\text{NMDA}} < \Theta_{\text{LTP}}$ is used when updating both θ_{LTP} and θ_{LTD} . The first option gave similar results in learning, while the second option gave worse results, which will be commented upon in section 4.2.

³The BCM rule is one of the first synaptic plasticity rules developed to help explain how neuronal orientation selectivity in the primary visual cortex arises (Bienenstock et al. 1982).

Parameter	w_{\max}	w_{\min}	Θ_{LTP}	$\theta_{0, \text{LTP}}$	$\theta_{0, \text{LTD}}$	η_E	η_θ
Value	2	0.01	40 μM	0.4 μM	0.1 μM	0.85	2.5

Table 3.5: The parameters for the plasticity rule. $\theta_{0, \text{LTP}}$ and $\theta_{0, \text{LTD}}$ are the initial values of the thresholds θ_{LTP} and θ_{LTD} .

subthreshold learning mode, in which the neuron initially does not spike for any stimulus, and a plateau potential can elicit somatic spiking together with the background noise (Fig. 3.5A). The other protocol is used during suprathreshold learning mode, in which the neuron initially spikes for all stimuli due to additional feature-unspecific inputs, which are distributed throughout the dendrites and are active during all stimuli (Fig. 3.5B). In this mode the plateau potential cannot elicit somatic spiking without the additional distributed inputs. The two training modes are described in more detail in section 4.2. In both training modes clustered inputs do not elicit plateaus before learning.




Figure 3.5: The stimulation protocol for training the SPN on the NFBP. (A) The protocol used in subthreshold learning mode. (B) The protocol used in suprathreshold learning mode.

In both stimulation protocols, a stimulus activates the synaptic clusters corresponding to two of its features with 3 input spikes per synapse, which arrive randomly within a window of 30 ms. The Da feedback is delivered 400 ms after the beginning of the stimulus, and it lasts 50 ms. During the

Da feedback the synaptic weights and the lower plasticity thresholds are updated. To shorten the simulation running time, after the delivery of the Da feedback, the concentrations of intra- and extracellular Ca^{2+} and the buffer species are reset to their initial values, instead of waiting for them to settle back to baseline values. 600 ms after the beginning of a stimulus, the next stimulus is presented, in random order out of the four possible stimuli. The magnitudes of the Da feedback are always the same, and the Da response never moves from the time of interaction with the stimulus, to the time of stimulus presentation, as occurs after learning in vivo. This stimulation protocol can be thought of learning during which an organism is still unfamiliar with the reward value of the stimulus or the reward it predicts.

The synaptic plasticity rule used in this thesis is different from the one in Article B, with the goal of exploring whether and how different formulations and model assumptions can solve the NFBP.

4 Present Investigation

In this section we present and discuss the key results of the thesis for all three research questions. The detailed results and discussion are contained in the accompanying articles A – C.

4.1 Kinetic model of the AC5 signaling network

The research results indicated that the model with the ternary complex, shown in Fig. 3.1B, is much better at coincidence detection than the null model, shown in Fig. 3.1A, suggesting that indeed two neuromodulatory signals need to co-occur so that AC5 can be activated.

4.1.1 Mathematical analysis of the null model’s regulatory motif

We start with a steady state analysis of the regulatory motif of the null model (Fig. 3.1A), which shows that the motif inherently lacks the ability for coincidence detection. We show this by showing that the synergy for this motif is $S(t) \leq 1$. To isolate the effect of the motif only on coincidence detection we consider a simplified situation where:

1. During a $\text{Da} \uparrow$ and an $\text{ACh} \downarrow$, respectively, the rise in $[\text{G}\alpha_{\text{olf}}]$ and drop in $[\text{G}\alpha_{\text{i}}]$ are assumed to be square-shaped, i.e. that the G proteins interact with AC5 sufficiently fast for $[\text{G}\alpha_{\text{olf}}]$ and $[\text{G}\alpha_{\text{i}}]$ to reach steady states for the duration of their corresponding neuromodulatory signals. This means that the forward and backward rate constants k_{f1} , k_{f2} , k_{r1} and k_{r2} , are large enough so that the binary complexes are quickly

formed and dissolved, making the precise values of the rate constants irrelevant.

2. During the $\text{Da} \uparrow$, enough $\text{G}\alpha_{\text{olf}}$ is produced to occupy all AC5, i.e. $[\text{G}\alpha_{\text{olf}}] \gg [\text{G}\alpha_i]$ and $[\text{G}\alpha_{\text{olf}}]$ saturates AC5. This means that the activation of AC5 is not limited by the amount of $[\text{G}\alpha_{\text{olf}}]$.
3. During the $\text{ACh} \downarrow$, $[\text{G}\alpha_i]$ drops close to 0 and all of AC5 is disinhibited from $\text{G}\alpha_i$.
4. The stimulation and inhibition of AC5 is “perfect”, i.e. the catalytic rate of $\text{AC5} \cdot \text{G}\alpha_{\text{olf}}$ is very large, $\alpha_{\text{G}\alpha_{\text{olf}}} \rightarrow \infty$, and $\text{AC5} \cdot \text{G}\alpha_i$ is catalytically inactive, $\alpha_{\text{G}\alpha_i} = 0$.

With these simplifications only the regulatory motif of the G proteins with AC5 determines the value of S , allowing us to examine its effect on coincidence detection. Evaluating all of the terms in the expression for the synergy in Eq. 3.14, we have the following.

During $\text{Da} \uparrow + \text{ACh} \downarrow$ there is no $\text{G}\alpha_i$ in the system because of the $\text{ACh} \downarrow$ and the assumed $\text{G}\alpha_i \approx 0$ (assumption 3), so all of AC5 is occupied by $\text{G}\alpha_{\text{olf}}$ (assumption 2), so:

$$k_c(\text{Da} \uparrow + \text{ACh} \downarrow, t) = k_{c, \text{AC5} \cdot \text{G}\alpha_{\text{olf}}} = \alpha_{\text{G}\alpha_{\text{olf}}} k_{c, \text{AC5}}. \quad (4.1)$$

During a $\text{Da} \uparrow$ alone, there is again enough $\text{G}\alpha_{\text{olf}}$ to outcompete $\text{G}\alpha_i$ when occupying AC5, and it also saturates AC5 (assumption 2), yielding the same result for the average catalytic rate as for $\text{Da} \uparrow + \text{ACh} \downarrow$,

$$k_c(\text{Da} \uparrow, t) = k_{c, \text{AC5} \cdot \text{G}\alpha_{\text{olf}}} = \alpha_{\text{G}\alpha_{\text{olf}}} k_{c, \text{AC5}}. \quad (4.2)$$

For $\text{ACh} \downarrow$ alone, there is no $\text{G}\alpha_i$ in the system (assumption 3) and AC5 is partly occupied by any resting-state levels of $\text{G}\alpha_{\text{olf}}$:

$$k_{c(\text{ACh}\downarrow, t)} = p_4 k_{c, \text{AC5}} + p_5 k_{c, \text{AC5} \cdot \text{G}\alpha_{\text{olf}}}, \quad (4.3)$$

where p_4 and p_5 are the fractions of AC5 and $\text{AC5} \cdot \text{G}\alpha_{\text{olf}}$ during an $\text{ACh} \downarrow$:

$$p_4 = \frac{[\text{AC5}]_{\text{ACh}\downarrow}}{[\text{AC5}]_{\text{ACh}\downarrow} + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}]_{\text{ACh}\downarrow}} \quad p_5 = \frac{[\text{AC5} \cdot \text{G}\alpha_{\text{olf}}]_{\text{ACh}\downarrow}}{[\text{AC5}]_{\text{ACh}\downarrow} + [\text{AC5} \cdot \text{G}\alpha_{\text{olf}}]_{\text{ACh}\downarrow}} \quad (4.4)$$

and $p_5 = 1 - p_4$. Lastly, during the basal state where no neuromodulatory signals are present, k_c is given with Eq. 3.2, which because of the inactive $\text{AC5} \cdot \text{G}\alpha_i$ (assumption 4), becomes:

$$k_{c,ss} = p_1 k_{c,AC5} + p_2 \alpha_{\text{G}\alpha_{\text{olf}}} k_{c,AC5} \quad (4.5)$$

Substituting these expressions in Eq. 3.14 for the synergy, and taking the limit of $\alpha_{\text{G}\alpha_{\text{olf}}} \rightarrow \infty$, yields:

$$S \xrightarrow{\alpha_{\text{G}\alpha_{\text{olf}}} \rightarrow \infty} \frac{1}{1 + p_5 - p_2} < 1. \quad (4.6)$$

The synergy is always less than 1 since $p_5 > p_2$, i.e. the fraction of $\text{AC5} \cdot \text{G}\alpha_{\text{olf}}$ without any $\text{G}\alpha_i$ in the system is always greater than the fraction of $\text{AC5} \cdot \text{G}\alpha_{\text{olf}}$ in the basal state when $\text{G}\alpha_i$ is present in the system.

An example of this situation is simulated in Fig. 4.1, where the available $\text{G}\alpha_{\text{olf}}$ has been made very high to mimic assumption 2 we made above. Also, the neuromodulatory signals are made very long so that steady state levels of each species can be reached. The upper panels of Fig. 4.1 show each of the enzyme species, AC5 , $\text{AC5} \cdot \text{G}\alpha_{\text{olf}}$, and $\text{AC5} \cdot \text{G}\alpha_i$ (as a fraction of the total amount of AC5) during the three situations of neuromodulatory signals, $\text{Da} \uparrow + \text{ACh} \downarrow$, a $\text{Da} \uparrow$ alone, and an $\text{ACh} \downarrow$ alone. The lower panels show the average catalytic rate k_c and the amount of cAMP (normalized to the basal cAMP value) for the three situations of neuromodulatory signals, and the resulting synergy. It can be seen that for a $\text{Da} \uparrow$ alone, the high amount of $\text{G}\alpha_{\text{olf}}$ indeed binds to almost all of AC5 , producing a similar k_c to the case of $\text{Da} \uparrow + \text{ACh} \downarrow$, thus producing values of $S < 1$.

In reality, the assumptions 3 and 4 are valid, since $\text{AC5} \cdot \text{G}\alpha_i$ is indeed catalytically inactive, and $\text{AC5} \cdot \text{G}\alpha_{\text{olf}}$ is much stronger at catalysis than AC5 (as introduced in sections 2.12.1 and 3.1). Conversely, the assumptions 1 and 2 are not, since the neuromodulatory signals are short-lived, the rate constants estimated with the BD simulations are not very large, and the amount of $\text{G}\alpha_{\text{olf}}$ is not orders of magnitude greater than that of $\text{G}\alpha_i$. This makes coincidence detection possible with the null model, as well.

4.1.2 Comparison of the kinetic models with and without ternary complex

The comparison between the null model and the regulatory motif with the ternary complex is shown in Fig. 4.2. The neuromodulatory inputs modeled

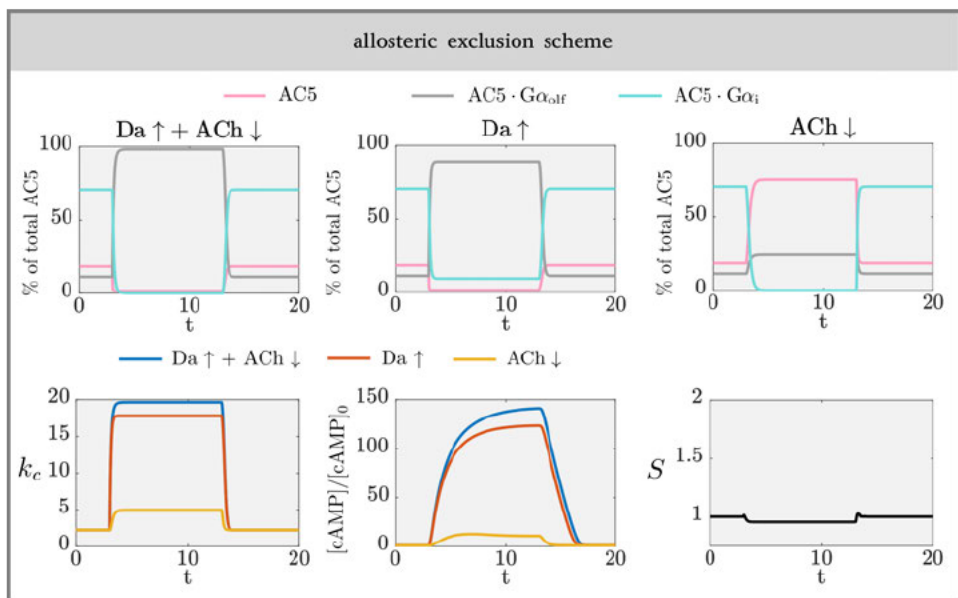
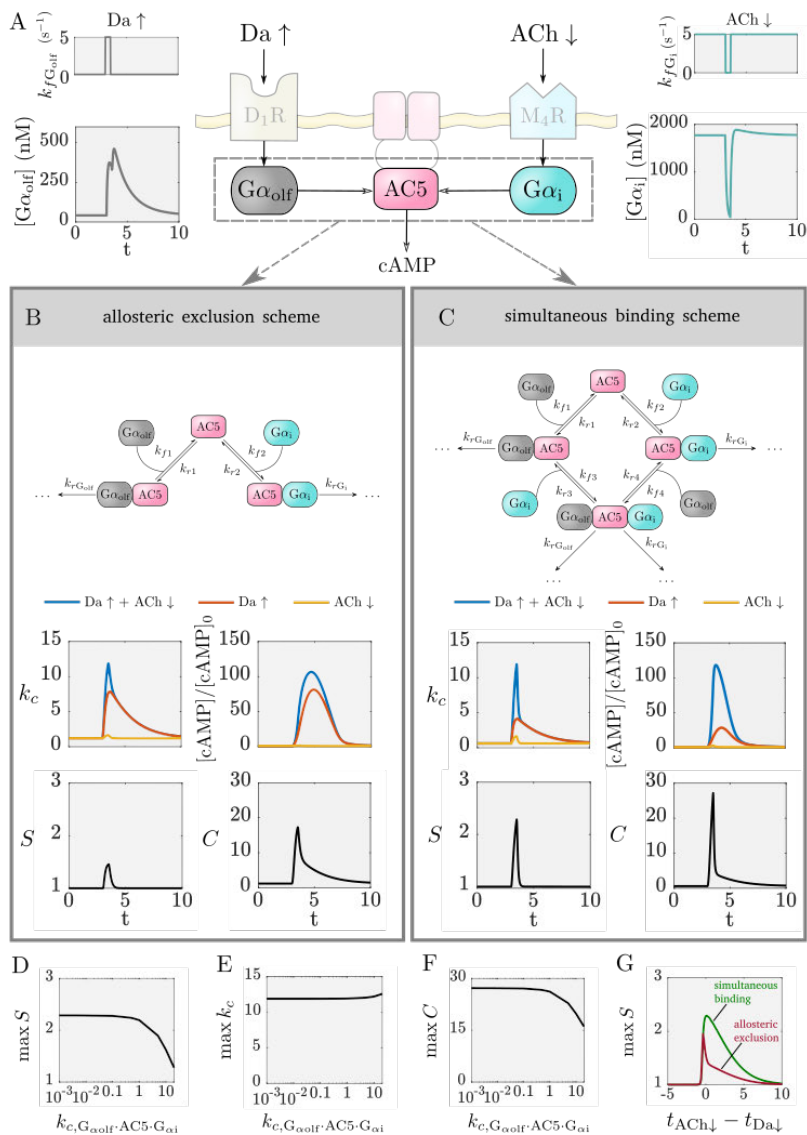


Figure 4.1: The regulatory motif of the null model cannot perform coincidence detection ($S < 1$). Panels above show the concentration of each enzyme species for the different neuromodulatory signals. The panels below show the average catalytic rate, the normalized cAMP response, and the synergy S . Figure reused from the supporting information in Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

as changes to the conversion rates of the G proteins are shown in Fig. 4.2A, as are the concentrations of G α_{olf} and G α_i obtained with the kinetic model with the ternary complex. [G α_{olf}] is a longer-lived signal compared to Da ↑, while [G α_i] follows the length of ACh ↓. Figures 4.2B and 4.2C, indicate that the regulatory motif with the ternary complex is much better at coincidence detection: both the synergy S and the scaled synergy C are larger than those for the null model. Figure 4.2B also shows that during Da ↑ + ACh ↓ and a Da ↑ alone, the null model generates cAMP in the same order of magnitude, which can activate PKA. In the model with the ternary complex, cAMP during Da ↑ + ACh ↓ and a Da ↑ alone differs by one order of magnitude, so that downstream activation of the LTP signaling network should occur only when the two neuromodulatory signals coincide.



In the null model, coincidence detection occurs because the forward and backward rate constants are slow enough such that steady states are not reached in the system while the neuromodulatory inputs last. Furthermore, the model with the ternary complex has a much wider window for coincidence detection, whereas in the null model the two transient neuromodulatory signals must arrive practically simultaneously for significant coincidence detection to occur (Fig. 4.2G). In addition, coincidence detection with the ternary complex is robust in several orders of magnitude of ternary complex catalytic rate (Fig. 4.2D–F). Coincidence detection disappears only for catalytic rates similar to the catalytic rate of the $AC5 \cdot G\alpha_{olf}$ binary complex.

The reason why coincidence detection is better in the motif with the ternary complex can be seen if we inspect the dynamics of each enzyme species during the neuromodulatory signals for both regulatory motifs (Fig. 4.3). Both motifs have similar maximal values for $k_c(Da \uparrow + ACh \downarrow)$ and for $k_c(ACh \downarrow)$. The enzyme species dynamics are only substantially different for a $Da \uparrow$. In the null model, the activated $G\alpha_{olf}$ competes with $G\alpha_i$ in binding to AC5 and creates a large amount of the catalytically active species $AC5 \cdot G\alpha_{olf}$, producing a large value for $k_c(Da \uparrow)$. In the kinetic model with the ternary complex, the activated $G\alpha_{olf}$ can bind to $AC5 \cdot G\alpha_i$ and create the catalytically inactive ternary complex, producing a low value for $k_c(Da \uparrow)$. It is as if the formation of the ternary complex absorbs any activated $G\alpha_{olf}$

Figure 4.2 (*preceding page*): The effect of the different regulatory motifs on coincidence detection. (A) The inputs to the model translate to an elevation in $G\alpha_{olf}$ and a pause in $G\alpha_i$. The plots are results of the model with the ternary complex. The shaded parts of the image are not included in the kinetic model. (B) The null model, also called allosteric exclusion scheme, and the k_c , synergy, and cAMP levels obtained due to this regulatory motif. (C) The regulatory motif with the ternary complex, also called the simultaneous binding scheme, and the k_c , synergy, and cAMP levels obtained due to this regulatory motif. (D, E, F) The effect of the ternary complex' catalytic activity on coincidence detection: the maximum of the synergy (D), the maximum of k_c (E), the maximum of the metric C (F). (G) The detection window for the allosteric occlusion scheme (red) and the simultaneous binding scheme (green). $t_{Da \uparrow}$ and $t_{ACh \downarrow}$ are the times when the $Da \uparrow$ and $ACh \downarrow$ occur, respectively. Note the shared time axes in (A), (B) and (C). Figure and caption reused from Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

and intracellularly there is little sign of a $Da \uparrow$ having occurred unless a coincident $ACh \downarrow$ arrives which will release the inhibition from $G\alpha_i$ and allow the intracellular production of cAMP.

Finally, a critical aspect for coincidence detection to work is to have fast deactivation of the active $G\alpha_i$. Only then can intracellular levels of $G\alpha_i$ decrease significantly to disinhibit AC5 during the short duration of the $ACh \downarrow$ signal. The experimental evidence for this high GTPase rate is listed in the description of the kinetic model in section 3.1. We varied the GTPase rate to see its effect on coincidence detection (Fig. 4.4). There is an optimum value of this rate – it needs to be both high enough to cause a drop in $[G\alpha_i]$ during the $ACh \downarrow$ and low enough so that there is significant inhibition of AC5 in the basal state. This result has also been demonstrated with the more extensive analysis on the robustness of coincidence detection to parameter values in the previous, much larger, kinetic model of the AC5 signal transduction network (Nair et al. 2015).

4.1.3 Discussion

Within this project we used multiscale simulations which showed that an inactive ternary complex between AC5 and its G protein regulators, a molecular-level feature, gives rise to significantly increased coincidence detection, a systems-level function of the signal transduction network that AC5 is embedded in. It suggested that two neuromodulatory signals are necessary to activate AC5 and enable LTP of the corticostriatal synapses. In dSPNs the required signals are a dopamine peak and an acetylcholine pause, and in iSPNs they are an adenosine peak and a dopamine pause. The ternary complex has still not been detected experimentally, neither in reconstituted preparations, nor in the intact system.

However, experimental results are consistent with an inactive ternary complex, of which the most significant are the following. Activating both D_2Rs and $A_{2a}Rs$ with pharmacological agonists in iSPNs from rat striatal primary cultures results in almost baseline cAMP levels (Navarro et al. 2018). Similarly, in striatal brain slices from mice the cAMP response of AC5 is completely prevented in both dSPNs and iSPNs when the activation of the respective G_{olf} -coupled receptors, D_1Rs and $A_{2a}Rs$, is followed by the activation of the respective G_i -coupled receptors, M_4Rs and D_2Rs (Nair et al. 2019, Fig. 3A; Yapo et al. 2017, Fig. 2B). Moreover, after completion of this project, in vivo experiments have shown that coincident

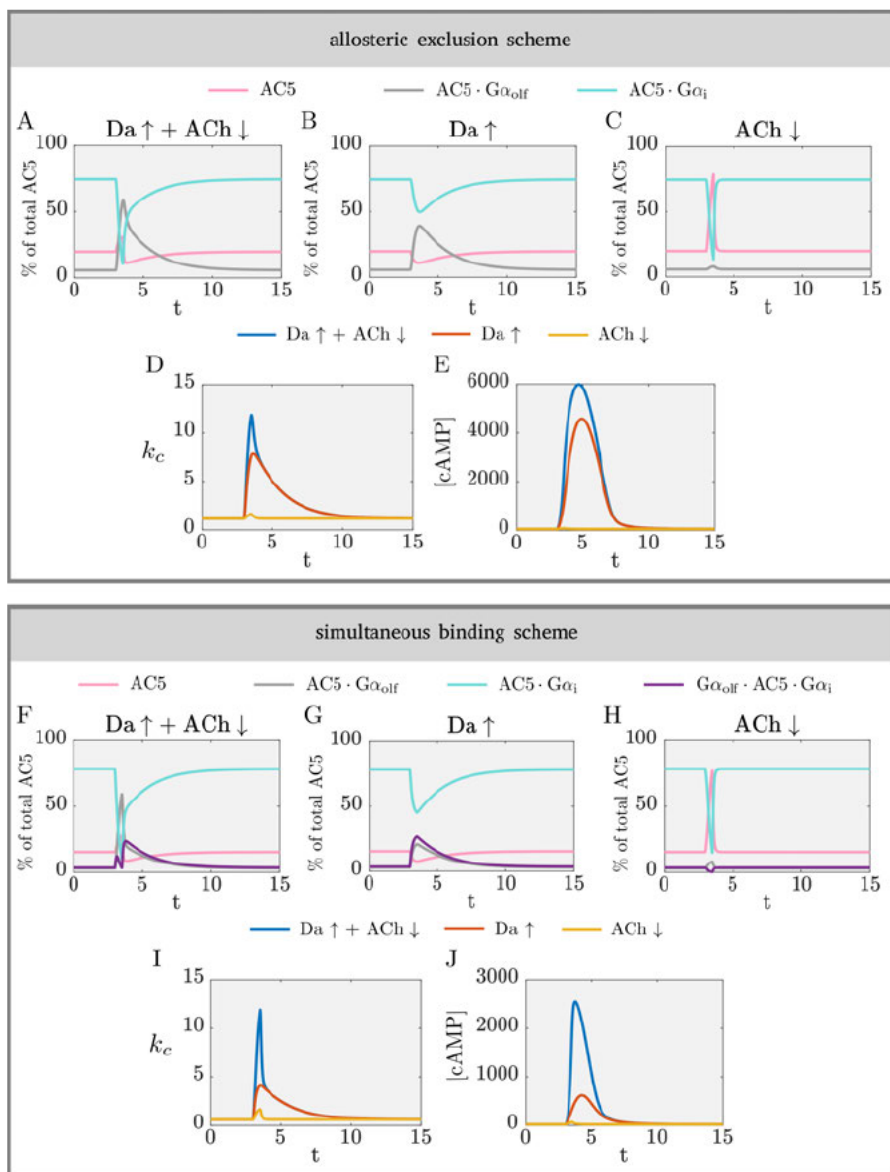


Figure 4.3: The dynamics of each enzyme species during $\text{Da} \uparrow + \text{ACh} \downarrow$, $\text{Da} \uparrow$, and $\text{ACh} \downarrow$, and the resulting k_c and [cAMP] for both kinetic models. Figure reused from the supporting information in Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.




Figure 4.4: Dependence of coincidence detection on the GTPase rate for $G\alpha_i$. Figure reused from the supporting information in Bruce et al. (2019). © 2019 Bruce et al. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

dopamine peaks and acetylcholine pauses are required for LTP in corticostriatal synapses onto dSPNs, consistent with the role of AC5 as a coincidence detector (Reynolds et al. 2022).

The primary regulators of AC5 are the two G proteins, but AC5 is also regulated by numerous other molecules, such as PKA, Ca^{2+} , nitric oxide, and the $G\beta\gamma$ subunit, and, as recently reported, via the transmembrane domains (Linder 2006; Landau et al. 2024). This positions AC5 as a major integration hub in synapses, and a future avenue for research is to investigate how these other regulators modify cAMP signaling and synaptic plasticity.

4.2 Plasticity rule for solving the NFBP

We use two setups to study whether the synaptic plasticity rule can solve the NFBP. In the first setup, the dSPNs initially do not spike upon stimulus presentation, and after learning they can elicit plateau potentials by themselves, without additional synaptic inputs apart from the background noise. This setup is called subthreshold learning, to indicate that the neuron does not spike initially (the somatic firing threshold is not reached). The dopamine feedback is delivered for each stimulus irrespective of whether the dSPN spiked or not. This models a situation where other dSPNs spike at the beginning of the learning process, and are involved in action initiation, which induces behavior and interaction with the stimulus, which in turn induces the dopamine feedback signal. Learning can be thought of as recruiting more dSPNs for reliable action initiation.

Since plateau potentials normally do not elicit somatic spiking in SPNs by themselves, we also used a setup where a plateau potential together with the background noise is not enough for somatic spiking – additional synaptic input distributed across the dendrites and coinciding with the plateaus is necessary. In this setup, the dSPN initially spikes for all stimuli, modeling a neuron that participates in action initiation, triggering interaction with the stimulus before learning occurs. Also, in this setup dopamine feedback is given only if the neuron spikes. Spiking for the relevant stimuli produces a $\text{Da} \uparrow$ and spiking for the irrelevant stimuli produces a $\text{Da} \downarrow$. We call this setup suprathreshold learning, to indicate that the neuron spikes initially. The background noise in this setup is set lower than the background noise in subthreshold learning.

Before showing the results of the learning rule, we first show the characteristics of plateau potentials in SPNs. As mentioned in section 3.9, the learning rule used in this thesis is different from the one in Article B, with the goal of exploring alternative model assumptions that could solve the NFBP.

4.2.1 Characterization of plateau potentials in the SPN model

To characterize plateau potentials in SPNs, we stimulated a single synaptic cluster in one dendrite, with varying size of 1 to 25 synapses (Fig. 4.5). The cluster's location was varied at three different distance ranges from the soma. Figure 4.5 shows the somatic voltage, spine voltage, spine $[\text{Ca}^{2+}]_{\text{NMDA}}$ and $[\text{Ca}^{2+}]_{\text{L-type}}$ for the different cluster sizes. Smaller clusters generate NMDA-dependent nonlinearities (NMDA spikes) which develop into plateau potentials when glutamate spillover occurs. At the threshold for glutamate spillover, a "jump" occurs in all four quantities shown in Fig. 4.5. The jump occurs because glutamate spillover produces a sudden and robust increase in NMDA conductance, caused by the activation of eNMDARs, where the clearance of glutamate is assumed to be slower. In this way an all-or-none plateau potential is evoked.

The supralinearity of the plateau potentials is also shown in Fig. 4.6, where the maximal amplitude of the somatic amplitude, $[\text{Ca}^{2+}]_{\text{NMDA}}$, and $[\text{Ca}^{2+}]_{\text{L-type}}$ are shown for three different values of synaptic weight. The voltage nonlinearity (voltage jump) observed at the soma reduces as the cluster is placed more distally, a result of the attenuation of the electrical

signals as they travel to the soma. Moreover, increased synaptic weight also reduces the size of the voltage nonlinearity in two ways: i) it increases the NMDA conductance of the smaller synaptic clusters, and ii) spillover can occur for smaller synaptic clusters, which evoke plateaus with smaller somatic amplitudes than larger clusters.

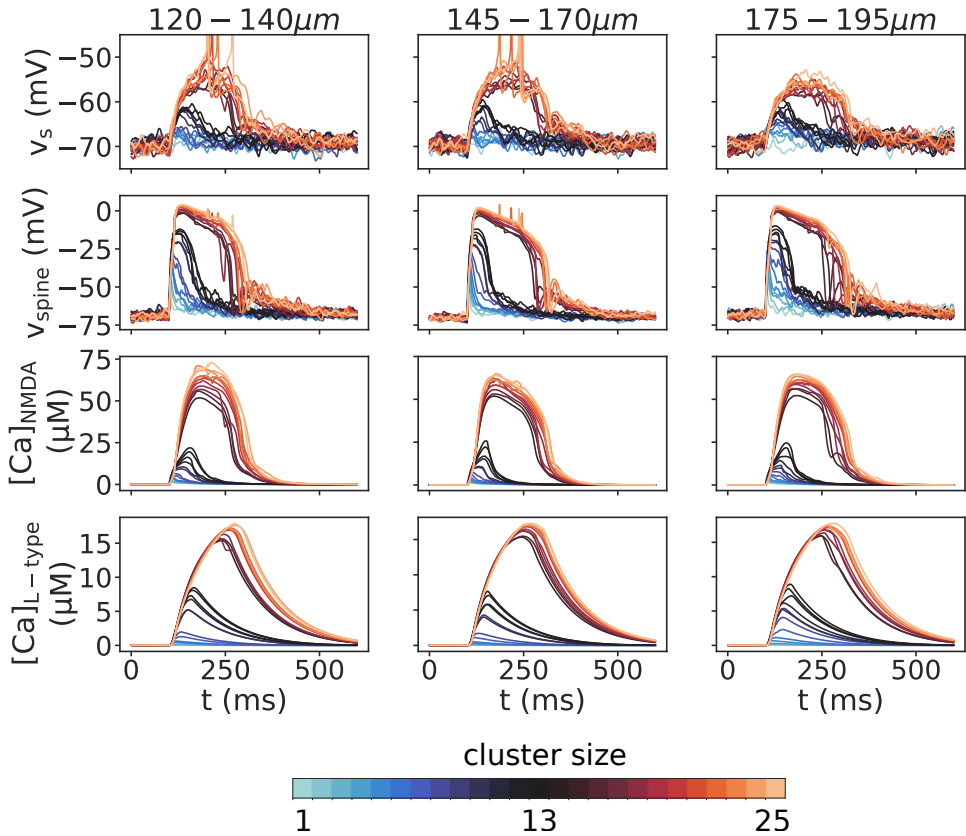


Figure 4.5: Plateau potentials generated by a synaptic cluster in a single dendrite, at three different distances from the soma. The spine voltage, $[\text{Ca}^{2+}]_{\text{NMDA}}$, and $[\text{Ca}^{2+}]_{\text{L-type}}$ represent the mean from all the spines in the cluster. 50 % of the total NMDA conductance is in the spines and 50 % is in the eNMDARs.




Figure 4.6: The maximal amplitude of the somatic voltage, $[\text{Ca}^{2+}]_{\text{NMDA}}$, and $[\text{Ca}^{2+}]_{\text{L-type}}$. Results are averages across 9 dendrites, and 10 trials per dendrite and per cluster size. Three values of synaptic weight were used, $w = 0.25$ (orange), $w = 0.5$ (red), $w = 0.75$ (dark red). 50 % of the total NMDA conductance is in the spines and 50 % is in the eNMDARs.

4.2.2 A single run of the plasticity rule

To illustrate how the plasticity rule works, we use the input configuration in Fig. 2.16A. The four features in the NFBP are distributed in two dendrites such that each dendrite, containing 3 features, can encode one relevant and one irrelevant stimulus, and the two relevant stimuli can be encoded in separate dendrites. Each feature is represented by a synaptic cluster of 10 synapses. We show both a subthreshold and a suprathreshold learning example in Figs. 4.7 and 4.8.

To solve the NFBP with this input configuration, the plasticity rule

should strengthen the synapses for two features in each dendrite and weaken the synapses for the third feature. Specifically, in dendrite 1, the synapses representing yellow and banana should be strengthened, while the synapses for red should be weakened. Similarly, in dendrite 2 the synapses for red and strawberry should be strengthened, while the synapses for yellow should be weakened. Indeed, this is the learning outcome in both the subthreshold and suprathreshold cases (Figs. 4.7D and 4.8D).

As explained in section 3.9, the upper plasticity threshold Θ_{LTP} , which is set at a $[Ca^{2+}]$ value obtainable only with plateau potentials, stops weight updates in synapses representing the relevant stimuli once they are strong enough to evoke plateaus (Figs. 4.7D and 4.8D). In this way, excessive synaptic strengthening is prevented. When an irrelevant stimulus arrives, only half of its synapses are strengthened, while the other half is weakened, and as such they do not generate plateaus. The thresholds θ_{LTP} and θ_{LTD} , together with the maximal $[Ca^{2+}]_{NMDA}$ and $[Ca^{2+}]_{L-type}$ values during a stimulus, are also shown in Figs. 4.7D and 4.8D. Synapses are strengthened only when $[Ca^{2+}]_{NMDA}$ is between θ_{LTP} and Θ_{LTP} , and weakened only when $[Ca^{2+}]_{L-type}$ is above θ_{LTD} . The plasticity thresholds θ_{LTP} and θ_{LTD} are updated to follow the maximal $[Ca^{2+}]_{NMDA}$ and $[Ca^{2+}]_{L-type}$, respectively, only if synaptic strengthening or weakening occurs, and without taking into account the upper threshold Θ_{LTP} . Before learning, the neuron does not spike in the subthreshold setup, and always spikes in the suprathreshold setup, and no plateaus are evoked (Figs. 4.7A and 4.8A, grey lines)¹. After learning, plateaus are evoked only for the relevant stimuli, contributing to somatic spiking, and no plateaus are evoked for the irrelevant stimuli, for which the neuron is silent (Figs. 4.7A and 4.8A, black lines).

In the case of suprathreshold learning, there are additional feature-unspecific synapses, which are distributed across the dendrites and are active during all stimuli. Because initially they contribute to somatic spiking for each stimulus, they should be weakened once the neuron learns to generate plateau potentials through the clustered inputs. They should be weakened to a level which together with the plateau potentials still elicits somatic spiking for the relevant stimuli. Figure 4.8E shows the synaptic weights for the

¹In the suprathreshold setup, prolonged NMDA spikes appear in the dendrites before learning, compared to the subthreshold setup. These are due to the distributed, feature-unspecific inputs that cause somatic spiking. Somatic spiking in turn elevates the voltage in the dendrites, which helps to relieve the Mg^{2+} block of NMDARs, causing prolonged NMDA spikes.

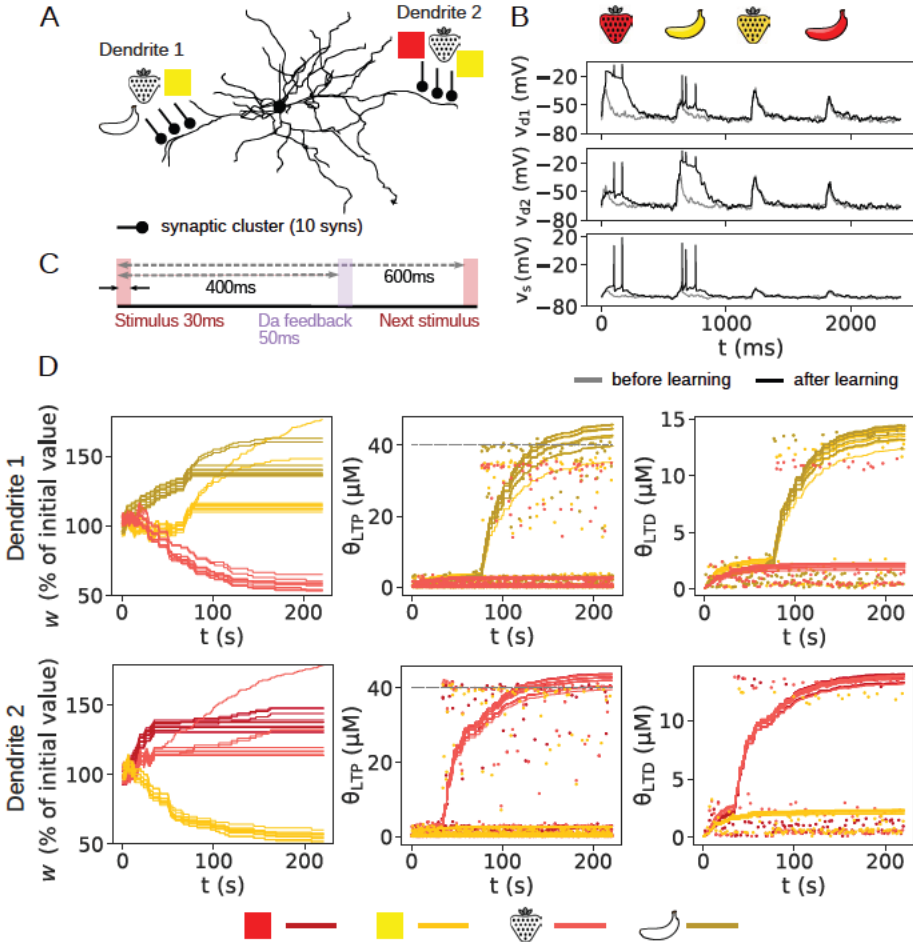
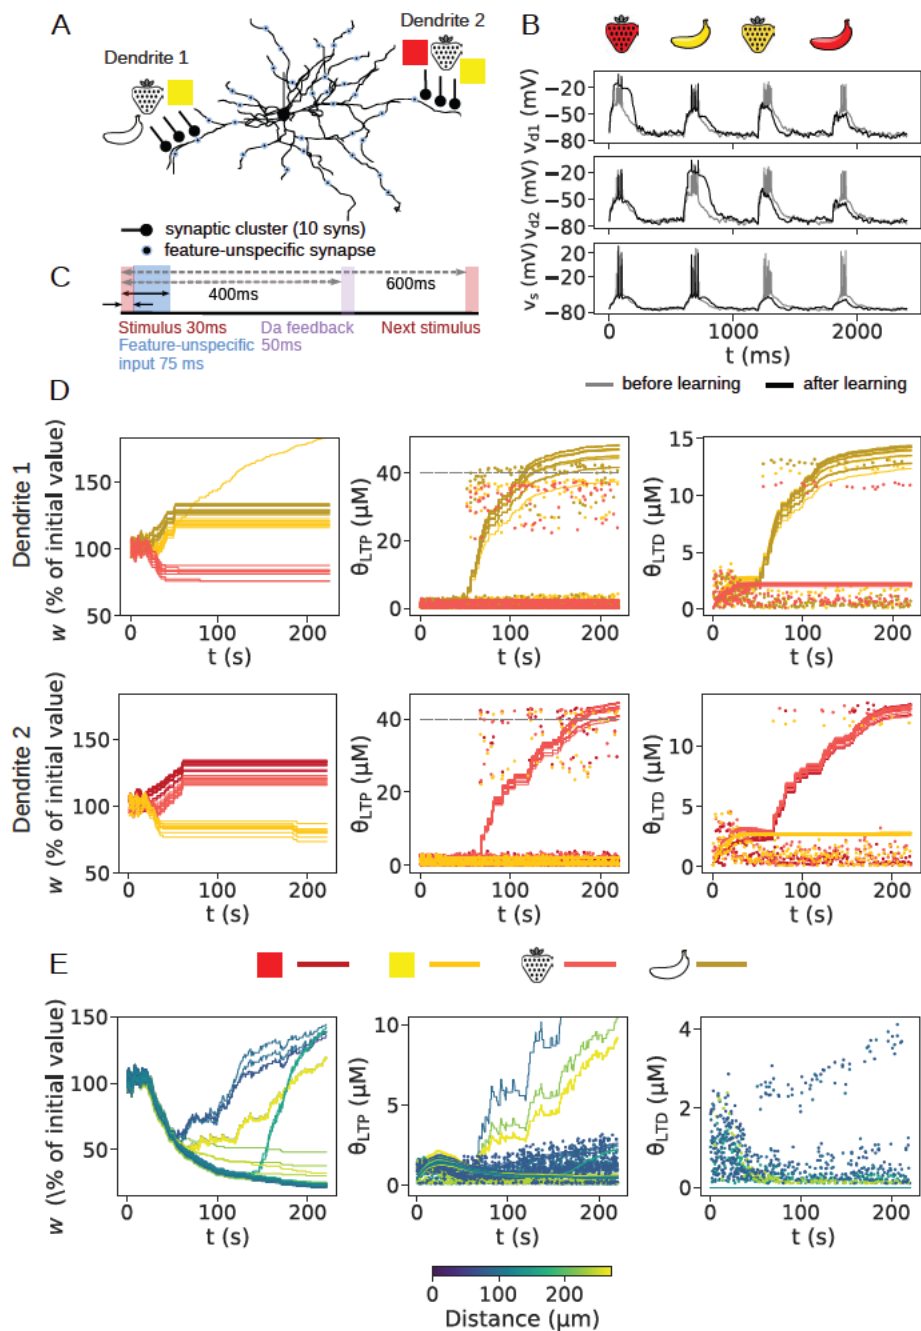


Figure 4.7: Illustration of the learning rule in subthreshold mode. (A) The input configuration contains 3 features per dendrite, so that one relevant and one irrelevant stimulus are represented in each dendrite. A cluster of 10 synapses is used to represent each feature. (B) Voltage traces from the two dendrites and soma before and after learning, for each stimulus. (C) The stimulation protocol (explained in section 3.9). (D) The evolution of synaptic weights (left), θ_{LTP} (middle), and θ_{LTD} (right). Dots show the maximal Ca^{2+} amplitudes for the first synapse in each cluster and for every fourth stimulus (for better legibility). The upper plasticity threshold, θ_{LTP} , is fixed and indicated with grey dashed lines.



feature-unspecific synapses. Initially, they strengthen, as their $[\text{Ca}^{2+}]_{\text{NMDA}}$ levels are higher than their lower plasticity thresholds for LTP, θ_{LTP} . However, because these synapses are not part of clusters, they do not experience plateau potentials and cannot reach $[\text{Ca}^{2+}]_{\text{NMDA}}$ levels as high as the clustered synapses. Therefore, after some iterations of strengthening, the lower plasticity threshold θ_{LTP} is updated to the most current maximal $[\text{Ca}^{2+}]_{\text{NMDA}}$ levels, and prevents them from strengthening further. From that point onwards, feature-unspecific synapses are mostly weakened whenever there is a $\text{Da} \downarrow$ and $[\text{Ca}^{2+}]_{\text{L-type}} > \theta_{\text{LTD}}$. When they are weakened enough so that there is no somatic spiking for the irrelevant stimuli, further weakening stops because without somatic spiking, no Da feedback is elicited. Crucially, the threshold for LTD, θ_{LTD} , is fixed for these synapses. If not, it will protect these synapses from weakening. This will be explained in section 4.2.4 where we will examine the role of the sliding plasticity thresholds, i.e. the role of metaplasticity.

4.2.3 Performance on the NFBP

The performance for both subthreshold and suprathreshold learning is shown in Fig. 4.10. We tested all possible feature combinations on two dendrites which allow for both relevant stimuli to be encoded, i.e. for the NFBP to be solved. There are a total of 31 such feature combinations (Fig. 4.9). We have divided the results in two groups: those describing input configurations with

Figure 4.8 (*preceding page*): Illustration of the learning rule in suprathreshold mode. (A) The same input configuration as in Fig. 4.7A is used, with additional 120 feature-unspecific inputs. (B) Voltage traces from the two dendrites and soma before and after learning, for each stimulus. (C) The stimulation protocol (explained in section 3.9). (D) The evolution of synaptic weights (left), θ_{LTP} (middle), and θ_{LTD} (right). Dots show the maximal Ca^{2+} amplitudes for the first synapse in each cluster and for every fourth stimulus (for better legibility). The upper plasticity threshold, Θ_{LTP} , is fixed and indicated with grey dashed lines. (E) The evolution of synaptic weights (left) and θ_{LTP} (middle) for the feature-unspecific inputs. θ_{LTD} is fixed for these synapses (right). The line color shows the distance of the synapse from the soma. Dots show the maximal Ca^{2+} amplitudes for every fourth synapse. The plot for θ_{LTP} (middle) is clipped to show the dynamics for most of the synapses, i.e. high values are not shown.

up to three features in a dendrite, and those describing input configurations with four features in a dendrite. We interpret the performance results in the

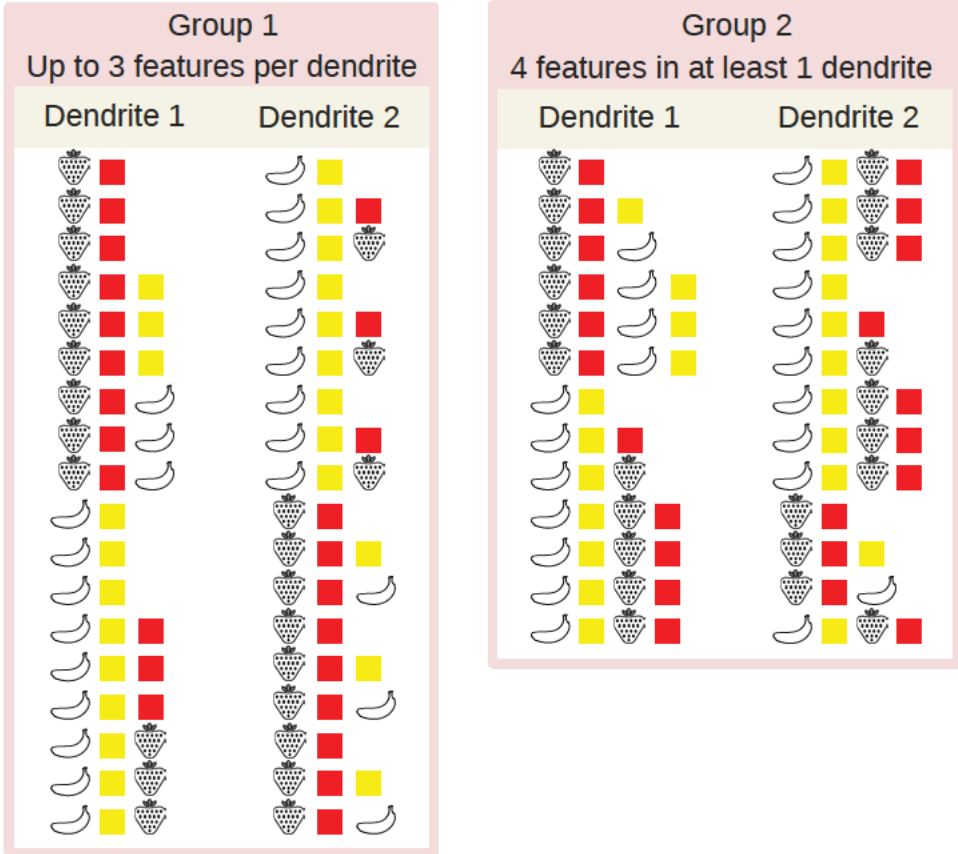


Figure 4.9: A list of all possible feature combinations of the four features on two dendrites, divided in two groups. The first group contains all 18 feature combinations which have up to 3 features per dendrite. The second group contains all 13 feature combinations which have 4 features in at least one dendrite.

following way:

- Performance of 100% indicates that the neuron spikes only for the relevant stimuli and is silent for the irrelevant stimuli.

- Performance of 50% can indicate two situations: i) either the neuron spikes for all four stimuli, or ii) is silent for all four stimuli.
- Performance of 75% indicates spiking behavior between two cases: i) the SPN spikes for only one of the relevant stimuli, remaining silent for the other three, or ii) it is silent for one of the irrelevant stimuli while spiking for the other three.
- We consider the NFBP solved when the performance is greater than 87.5%, exemplified by the situation where the SPN always spikes for one relevant stimulus, and at least half of the time for the other relevant stimulus, while remaining silent for the irrelevant stimuli.

In the group with up to three features per dendrite, the NFBP is successfully solved with close to 100% performance in subthreshold learning and 95% in suprathreshold learning (Figs. 4.10A, and 4.10B). In the group with four features in at least one dendrite, the plasticity rule typically cannot successfully encode both relevant stimuli in separate dendrites. Only one of the relevant stimuli is stored by the neuron and the NFBP is not solved (the mean performance is around 75%). This happens because the dendrite with four features can store either of the relevant stimuli, and the random order of stimulus presentation during learning determines which of the two stimuli will be stored in the dendrite. The NFBP is solved in some trials, which is indicated by the large standard deviation in this group. Once one of the relevant stimuli is stored in the dendrite, the synapses for the other relevant stimulus are weakened when they are activated as part of the irrelevant stimuli.

More distal synaptic clusters produce smaller somatic plateau potential amplitudes (Fig. 4.6, middle and right panels). Nevertheless, in suprathreshold learning, where the additional feature-unspecific inputs can contribute to somatic spiking, the NFBP can still be solved for more distally located clusters when up to 3 features per dendrite are present (Figs. 4.10C and 4.10D).

4.2.4 The role of metaplasticity

Metaplasticity is crucial for the NFBP to be successfully solved. We show the role of metaplasticity on the example of subthreshold learning. Figure 4.11 shows the outcome of learning when the LTD plasticity threshold, θ_{LTD} ,

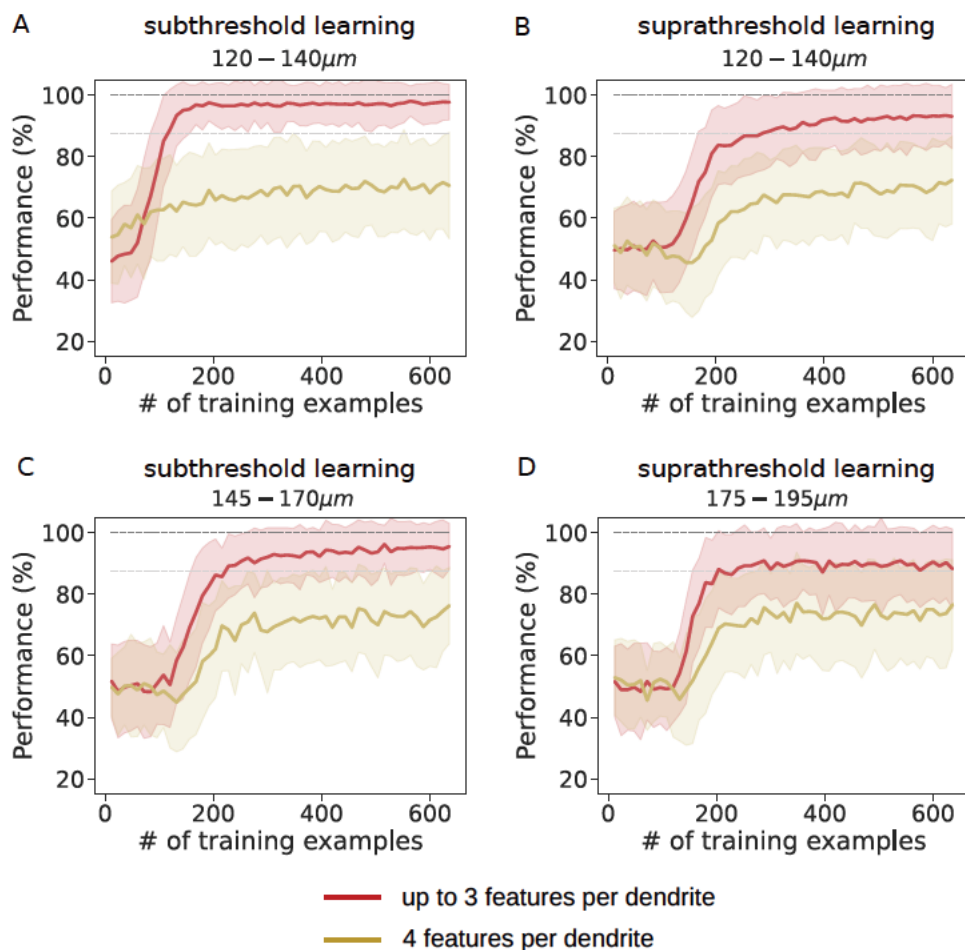


Figure 4.10: The performance on the NFBP for the two groups of input configurations shown in Fig. 4.9. (A) Performance for subthreshold learning. (B) Performance for suprathreshold learning. (C, D) Performance for suprathreshold learning where the cluster is placed more distally (at the same locations as in Fig. 4.6). In all panels, the two gray lines mark the levels 87.5 % and 100 %. Results are averaged over 10 trials per input configuration, and across all input configurations in each group.

is fixed. Without metaplasticity in LTD, the synaptic weights of the relevant stimuli do not stabilize. Comparing with Fig. 4.7D, where LTD metaplasticity is active, we can conclude that once the synaptic weights of the relevant stimuli are strengthened, the rising θ_{LTD} protects them from weakening easily. This is because with a higher θ_{LTD} , strengthened synapses will also require higher $[\text{Ca}^{2+}]_{\text{L-type}}$ to be weakened. Without LTD metaplasticity, there is nothing that prevents strengthened synapses to start weakening, which is the reason why the synaptic weights cannot stabilize (Fig. 4.11C). Moreover, with a fixed θ_{LTD} , synapses that do not belong to the relevant stimuli are weakened until their $[\text{Ca}^{2+}]_{\text{L-type}}$ falls below θ_{LTD} , which in the example of Fig. 4.11 makes them drop close to the minimal weight w_{min} . When LTD metaplasticity is active, it limits the weakening of these synapses (Fig. 4.7D, left and right). Since the neuron does not learn to generate plateaus, the performance remains around 50 % for both groups of input configurations (Fig. 4.11B).

Metaplasticity of LTD is beneficial for clustered synapses. However, in the distributed, feature-unspecific synapses, it is crucial that θ_{LTD} is fixed, in order to allow them to weaken enough so that they do not elicit somatic spiking for all stimuli. Otherwise, if LTD metaplasticity is active for the feature-unspecific synapses, it will prevent them from weakening once θ_{LTD} becomes updated to their $[\text{Ca}^{2+}]_{\text{L-type}}$ levels. Therefore, a necessary component of the learning rule is that the LTD threshold θ_{LTD} is plastic only for the clustered synapses, and not for the distributed, feature-unspecific synapses.

Figure 4.12, on the other hand, shows the outcome of learning when the lower LTP plasticity threshold, θ_{LTP} , is fixed. To make the role of θ_{LTP} more obvious, we have fixed it to a lower value – the initial value used in Fig. 4.8. Compared to Fig. 4.7D, the synapses belonging to the feature in the irrelevant stimulus in each dendrite are not weakened (strawberry in dendrite 1 and yellow in dendrite 2). This occurs because θ_{LTP} is kept fixed at a low value. In this case, the synapses belonging to the irrelevant stimulus can almost always create $[\text{Ca}^{2+}]_{\text{NMDA}}$ levels higher than θ_{LTP} when they are activated by a relevant stimulus (e.g. the synapses for strawberry on dendrite 1 when a red strawberry is presented). This means that they can always get strengthened, despite being weakened when an irrelevant stimulus arrives (e.g yellow strawberry on dendrite 1). Conversely, when LTP metaplasticity is active, θ_{LTP} evolves to match the maximal $[\text{Ca}^{2+}]_{\text{NMDA}}$ levels during learning (Fig. 4.7D, middle panels). In this case, weakening

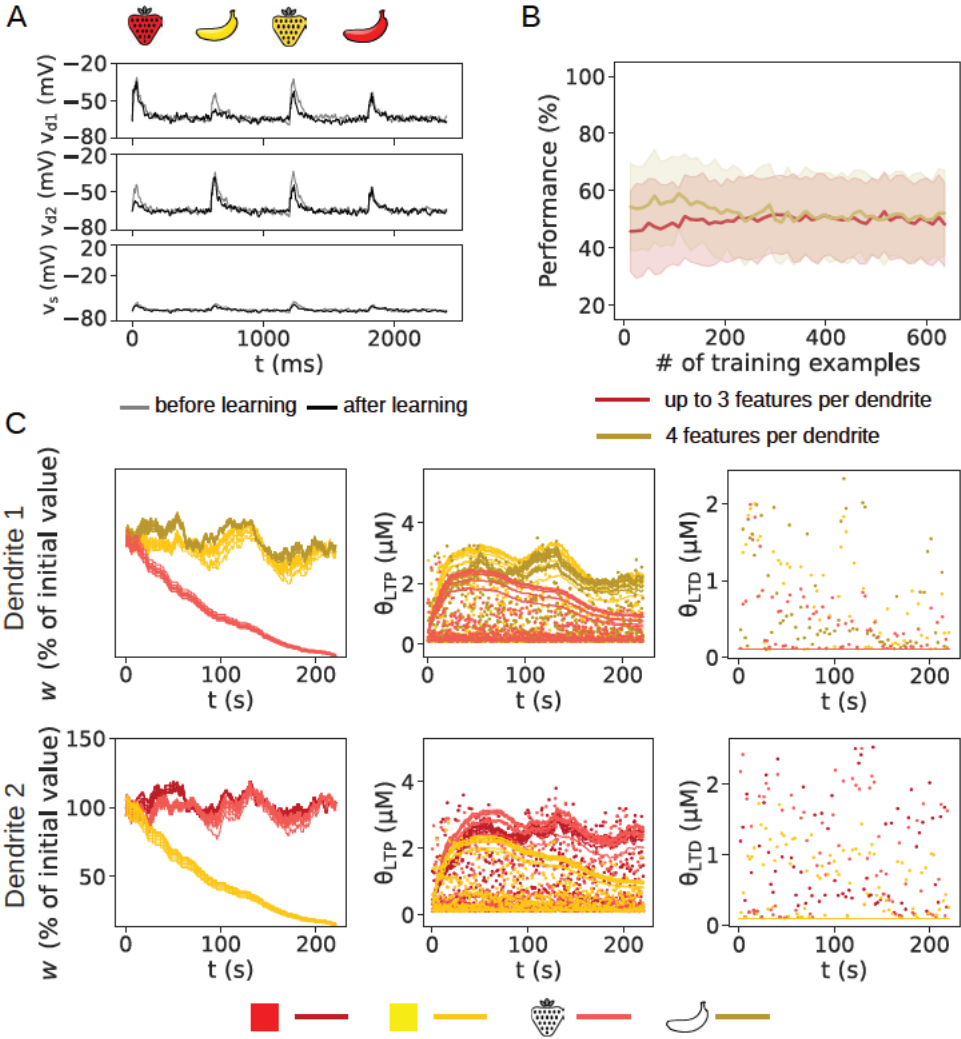


Figure 4.11: The role of the threshold for LTD, θ_{LTD} , in learning. The same input configuration as in Figs. 4.7A and 4.8A is used. $\theta_{LTD} = 0.1 \mu M$, i.e. there is no metaplasticity for LTD. (A) The dendritic and somatic voltage for each stimulus before and after learning. (B) The performance on the NFBP for the two groups of input configurations. Results are averaged over 10 trials per input configuration and over all input configurations in each group. (C) The synaptic weights, θ_{LTP} and θ_{LTD} , with the corresponding maximal $[Ca^{2+}]$ levels.

the synapses for the feature activated by the irrelevant stimulus in each dendrite will generate less $[Ca^{2+}]_{NMDA}$, which will eventually fall below θ_{LTP} , preventing these synapses from strengthening in the future. This is the reason why the synapses for strawberry in dendrite 1 and yellow in dendrite 2 are persistently weakened during learning. The synapses for the relevant stimulus are strengthened, similarly to Fig. 4.7D, until a plateau potential is generated in each dendrite (shown with black lines in Fig. 4.12A and $[Ca^{2+}]_{NMDA}$ amplitudes in Fig. 4.12C, middle panels).

Performance on the NFBP is slightly worse without LTP metaplasticity, and is due to not being able to weaken unnecessary synapses. The neuron correctly generates plateaus and spikes for the relevant stimuli, but also sometimes spikes for at least one of the irrelevant stimuli (Fig. 4.12B).

4.2.5 Discussion

We formulated a synaptic plasticity rule based on the corticostriatal synaptic signaling networks for LTP and LTD, in order to study the computational capacity of SPNs. The plasticity rule is based on the local calcium concentration and does not require that the synapse has access to global neuronal properties such as the timing of somatic spiking. We have used the plasticity rule in dSPNs, where dopamine signaling determines whether an active synapse will be strengthened or weakened. We have assumed that the dopamine signal results from the outcome of an action taken by the organism, so it represents a feedback signal for whether the outcome was rewarding or not. This formulation, together with assuming pre-existing clustered synapses, allows the SPNs to solve the nonlinear feature binding problem. With this learning rule, dSPNs learn to initiate actions towards the “relevant”, or rewarding stimuli in the NFBP, and take no actions towards the “irrelevant” or non-rewarding ones. An analogous learning rule using the same dopamine reward feedback for the stimuli would work in iSPNs, as well. In iSPNs an adenosine peak stimulating the A_{2a} Rs is necessary to activate AC5, while a coincident dopamine pause is necessary to disinhibit it through the D_2 Rs. Corticostriatal synapses onto the iSPNs would thus be strengthened for the irrelevant stimuli, activating iSPNs to suppress actions, while being weakened for the relevant ones, so that iSPNs would not suppress actions. This fits well with the role of the basal ganglia output in initiating actions.

The ability to solve the NFBP can be interpreted as indicating that a

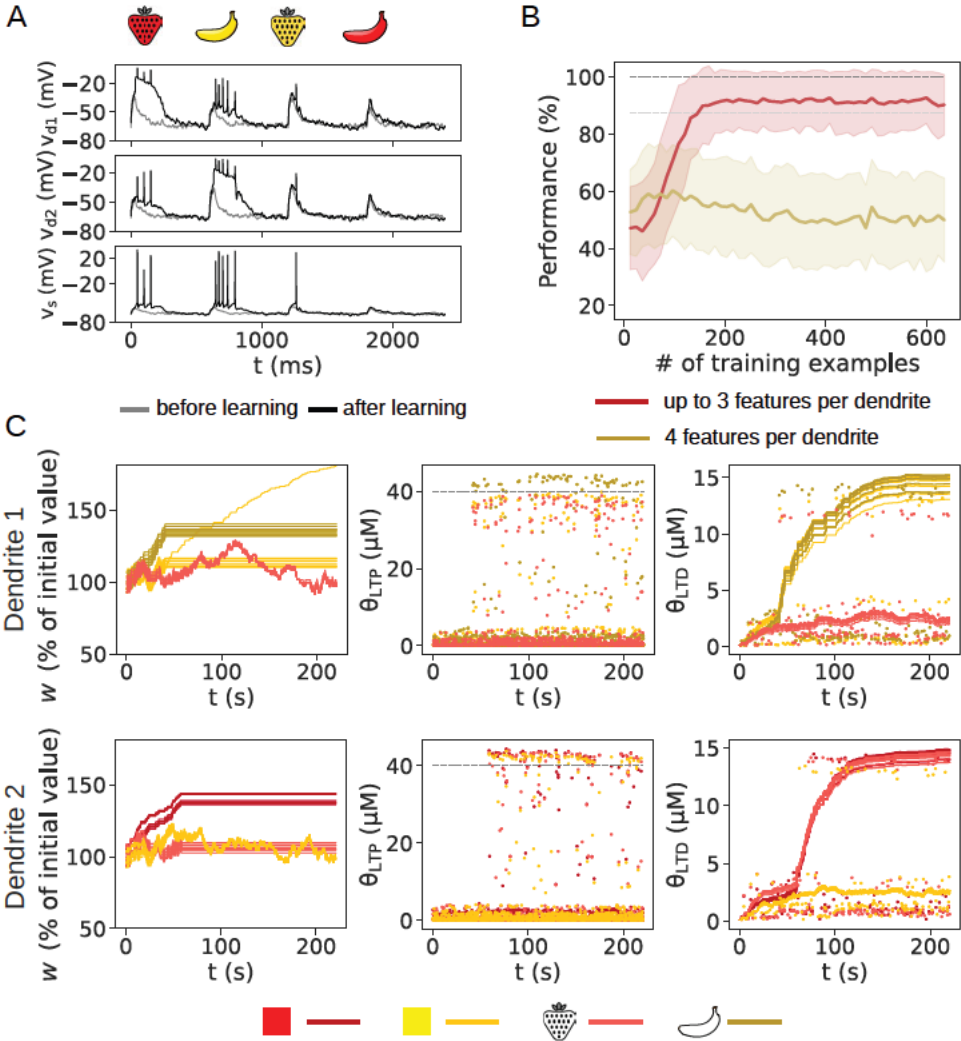


Figure 4.12: The role of the lower threshold for LTP, θ_{LTP} , in learning. The same input configuration as in Figs. 4.7A and 4.8A is used. $\theta_{LTP} = 0.4 \mu M$, i.e. there is no metaplasticity for LTP. (A) The dendritic and somatic voltage for each stimulus before and after learning. (B) The performance on the NFBP for the two groups of input configurations. Results are averaged over 10 trials per input configuration and across all input configurations in each group. (C) The synaptic weights, θ_{LTP} and θ_{LTD} , with the corresponding maximal $[Ca^{2+}]$ levels.

dSPN responsible for initiating a particular action can be engaged in multiple different situations requiring the same action, through, for example, activating a plateau potential on one of the dendrites. Moreover, it means that a single dSPN can learn to discriminate between similar stimuli out of which only some should initiate an action.

In our formulation of the learning rule we have phenomenologically implemented the two branches of the corticostriatal signaling networks belonging to the cortical and neuromodulatory inputs. In addition, we have incorporated three assumptions in the learning rule. The first and major assumption is the upper threshold for synaptic strengthening, i.e. that high $[\text{Ca}^{2+}]_{\text{NMDA}}$ during LTP conditions (dopamine peak and acetylcholine pause) prevents any synaptic plasticity. Whether this is true in biological neurons remains an open question. Nevertheless, LTP has been shown to saturate with more stimulus repetitions in experimental synaptic plasticity protocols (Yger and Gilson 2015). AC5 is inhibited by Ca^{2+} , indicating that preventing synaptic strengthening by high $[\text{Ca}^{2+}]_{\text{NMDA}}$ might occur already very early in the signaling network. Inhibiting AC5 would not activate PKA and would thus leave the balance of phosphorylated and unphosphorylated glutamate receptors unchanged. Additionally, Ca^{2+} activates PP2B through calmodulin, and PP2B inactivates PKA (Church et al. 2021). Due to the complexity of the signaling networks, the influence of high $[\text{Ca}^{2+}]_{\text{NMDA}}$ on plasticity outcomes requires further study, and using a model of the whole signaling network might provide some answers. The second assumption is that metaplasticity operates in the corticostriatal synapses onto dSPNs, which is not a strong assumption, since metaplasticity has been documented for many synapses. The third assumption, critical for suprathreshold learning, is that distributed synapses do not have metaplasticity for synaptic weakening, which is active only in the clustered synapses. This difference in metaplasticity might be implemented by signaling networks that respond to diffusing molecules that only span the short distances within a synaptic cluster, and not the large distances between synapses distributed across many dendrites.

Additionally, in the plasticity rule, updates to the synaptic strength are done in small steps. This resulted in somewhat long training times on the order of 200 training examples to solve the NFBP, although still 2–10 times faster than in other studies (Schiess et al. 2016; Bicknell and Häusser 2021). Future research could use the models of the whole signaling networks to translate the synaptic inputs into a synaptic plasticity outcome, which

might result in different dynamics of the synaptic weights. With a model of the whole signaling network, it would be possible to investigate stimulation protocols where the magnitude of the Da feedback is varied to model subjective reward value. It would also be possible to model scenarios where during learning Da feedback moves to the time of stimulus presentation, becoming a predictor of reward value.

Perhaps more than anything else, solving the NFBP with the proposed plasticity rule highlights the importance of deciding when to stop updating synaptic weights. The two processes/mechanisms that regulate the amount of weight updates in the rule are the upper plasticity threshold for LTP, and metaplasticity. Without the upper plasticity threshold, all strengthened synapses would reach their maximally attainable synaptic strength. Without metaplasticity, the synaptic weights would not stabilize. Important future research directions are to study various such possible mechanisms and their effects on learning. Experimentally, several different forms of metaplasticity have been reported. In the learning rule we considered metaplasticity mechanisms which are activated by the same conditions that activate synaptic plasticity. However, stimulation that does not change synaptic strength can also activate metaplasticity and affect the plasticity thresholds (Abraham 2008). Moreover, we only considered metaplasticity evoked in the same synapse where synaptic plasticity occurs (homosynaptic metaplasticity), yet heterosynaptic metaplasticity mechanisms also exist, where the evoking stimuli originate from different synapses or different neurons (Abraham 2008; Abraham et al. 2019). It is likely that different metaplasticity mechanisms exist in different synapse types, dedicated to support the possibly specific roles of such synapses in various brain regions.

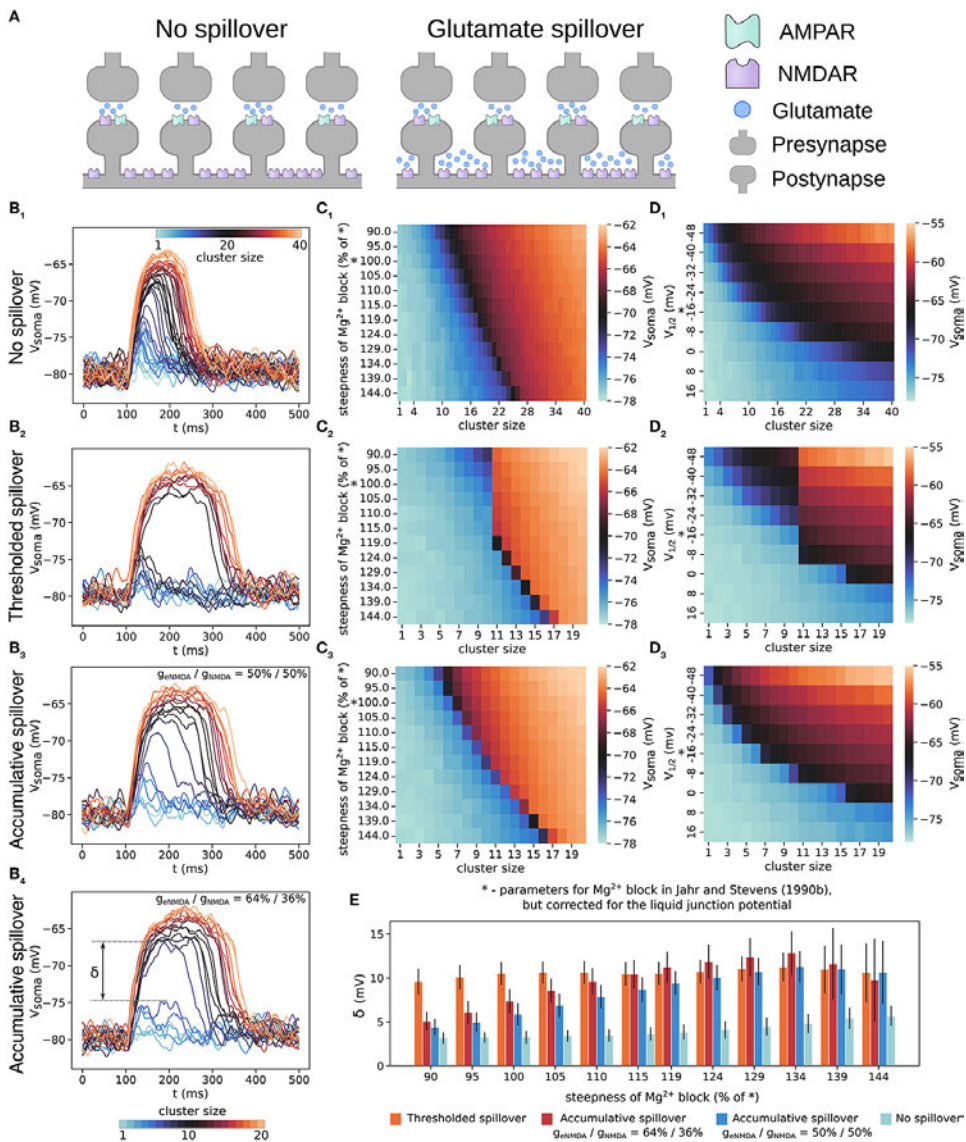
4.3 Glutamate Spillover

The plateau potentials utilized by the plasticity rule are generated with glutamate spillover. Here we present the results that led us to incorporate glutamate spillover in the model for plateau potentials. In the results sections 4.3.1, 4.3.2, and 4.3.3, much of the text is reused and/or paraphrased from the corresponding sections in Trpevski et al. (2023). © 2023 Trpevski, Khodadadi, Carannante and Hellgren Kotaleski. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

4.3.1 Thresholded and accumulative spillover

As introduced in section 2.12.6, most plateau potential models do not capture the all-or-none plateau quality unless unrealistic gating function parameters are used. To demonstrate the role of glutamate spillover, we start by comparing the plateau potentials evoked by a synaptic cluster of increasing size without spillover and with the two spillover models, where in the gating function we used the originally reported parameters for the Mg^{2+} block but corrected for the value of the junction potential (i.e., the parameters found in Jahr and Stevens (1990) but corrected as in Ecker et al. (2020)). The somatic voltage traces in Figs. 4.13B₁, B₂ show that plateau potentials generated without glutamate spillover exhibit a graded increase in amplitude (as in Fig. 2.17 of section 2.12.6), whereas those generated with thresholded glutamate spillover exhibit an all-or-none response. We note again that the model with no spillover has a twice larger cluster size than the models with glutamate spillover, in order to compare the same total NMDAR conductance in the two cases. This means that every second voltage trace in Fig. 4.13B₁ is comparable to the voltage traces in Figs. 4.13B₂ -B₄.

We also varied the steepness of the gating function by varying the parameter α in Eq. 2.1, and the amplitude of the plateau potentials generated in this way is shown as heatmaps in Figs. 4.13C₁, C₂. The all-or-none behavior can be seen as a sharp jump in the colors of the heatmaps for a small increase in cluster size (i.e., stimulus strength for evoking a plateau). Without glutamate spillover the all-or-none behavior is present only for steep sigmoidal curves, whereas with thresholded glutamate spillover, it is always present. In both cases, steeper gating functions require more excitation to overcome the Mg^{2+} block. (As described in section 3.8, thresholded spillover occurs after 10 synapses are activated, hence the sharp jump at 10 synapses for lower gating function steepness.) Comparing with the gating functions estimated from available experimental data (Table 2), most of them are more shallow than the originally reported one by Jahr and Stevens (1990), and result in graded NMDA potentials in the model without spillover, but still produce all-or-none plateau potentials in the model with glutamate spillover. Varying the parameter η , which shifts the gating function along the x-axis, produces the same qualitative results regarding the all-or-none behavior in the models with thresholded spillover and without glutamate spillover as when varying the parameter α (Figs. 4.13D₁, D₂): no glutamate spillover results in graded plateaus and thresholded glutamate spillover produces all-or-none plateaus.



Shifting the gating functions to very depolarized values disables the generation of plateau potentials, as shown in the heatmaps in Figs. 4.13D₁, D₂. In Fig. 4.13E, we compare the maximal difference between the somatic voltage amplitudes of two consecutive voltage traces, termed the “size of the jump”, δ (indicated in Fig. 4.13B₄). The size of the jump, δ , is a measure of the dendrite’s supralinear response. For example, a large sudden jump (≈ 10 mV) indicates all-or-none behavior, whereas small jumps (of a few mV)

Figure 4.13 (*preceding page*): Plateau potentials evoked with and without glutamate spillover. (A) Illustration of glutamate spillover vs. no spillover. (B) Somatic voltage traces of plateau potentials evoked by a cluster of increasing size from 1 to 40 synapses for the model with no spillover (B₁) and from 1 to 20 synapses for the models with: thresholded spillover (B₂), accumulative spillover with equal proportions of NMDAR and eNMDAR conductance (B₃), and accumulative spillover with a larger proportion of eNMDAR conductance (64% of total NMDAR conductance); (B₄). (C) Amplitude of somatic depolarization as a function of cluster size and steepness of the gating function for the models with: no spillover (C₁), thresholded spillover (C₂), accumulative spillover with equal proportions of NMDAR and eNMDAR conductance (C₃). (D) Amplitude of somatic depolarization as a function of cluster size and position of the gating function along the x-axis (determined by varying the parameter η according to $\eta = e^x$, for $x \in \{-3, -2.5, \dots, 2\}$). (D₁–D₃) show results for the same models as in (C₁–C₃). Values in the heatmaps are averages over 50 trials of clustered synaptic inputs elicited in the same dendrite as for the panels in (B), except for (C₁, D₁), which are averages over 30 trials. The “*” on the y-axes in (C, D) indicates the gating function parameters as in Jahr and Stevens (1990), but corrected for the liquid junction potential. (E) The maximal difference between two consecutive voltage traces, δ , for the four models in (B). Large jumps indicate all-or-none plateau potentials, and small jumps are consistent with graded NMDA spikes/plateau potentials. δ for no spillover was calculated taking every second voltage trace (steps of 7 nS in NMDA conductance), so that it is comparable to the δ calculated for (B₃, B₄). In B₁, results are averages over 30 trials and 11 different dendrites and in B₂–B₄ they are averages over 50 trials and 11 dendrites, with plateaus elicited at approximately the same distance from the soma. Figure and caption reused from Trpevski et al. (2023). Error bars represent standard deviation. © 2023 Trpevski, Khodadadi, Carannante and Hellgren Kotaleski. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

would indicate graded potentials. Thresholded glutamate spillover typically produced voltage jumps of 8–12 mV, while no spillover produced jumps of ≈ 3 mV (Fig. 4.13B₁).

In the cerebellum, physiologically relevant parallel fiber stimulation protocols cause gradual accumulation of glutamate in the extrasynaptic space with each stimulation pulse (Okubo et al. 2010). To explore the effects of such more gradual spillover, we used the accumulative spillover model, where each activation of spine NMDARs also activates the corresponding eNMDARs in the dendritic shaft under the spine after a delay, without the need to first reach a glutamate threshold.

The results of these simulations show that accumulative glutamate spillover results in plateau potential behavior which is between that without spillover and with thresholded glutamate spillover (Figs. 4.13B₃, C₃, D₃). Shallow gating functions cause graded plateau potentials, but all-or-none behavior arises much sooner for increasing steepness compared to the case without spillover. With thresholded spillover, clusters of up to 10 synapses activate only the spine NMDAR conductance, and clusters with more than 10 synapses activate both NMDAR and eNMDAR conductances, thus producing the large voltage jump before and after spillover (Fig. 4.13B₂). With accumulative spillover, eNMDARs are activated with a short delay after each NMDAR activation, giving rise to the more graded increase in plateau potential amplitude, as well as duration (Fig. 4.13B₃). Lastly, increasing the proportion of eNMDAR conductance in the total NMDA conductance produced plateaus which are more all-or-none (Fig. 4.13B₄, E). This suggests that there may be particular conditions under which all-or-none plateau potentials appear and that plateau potential behavior might vary among brain regions.

4.3.2 Long glutamate duration drives all-or-none plateau potentials

In Article C, we used dynamical systems analysis to understand how varying the parameters of the gating function influences the generation of plateau potentials. The results of the analysis suggested that the all-or-none quality of plateau potentials is a result of the long duration of glutamate in the extrasynaptic space. We verified this with simulations where we manipulated the duration of the transmitter pulse (the parameter T_{dur} , Fig. 4.14). As explained in section 3.7, T_{dur} for the synapses onto spines is

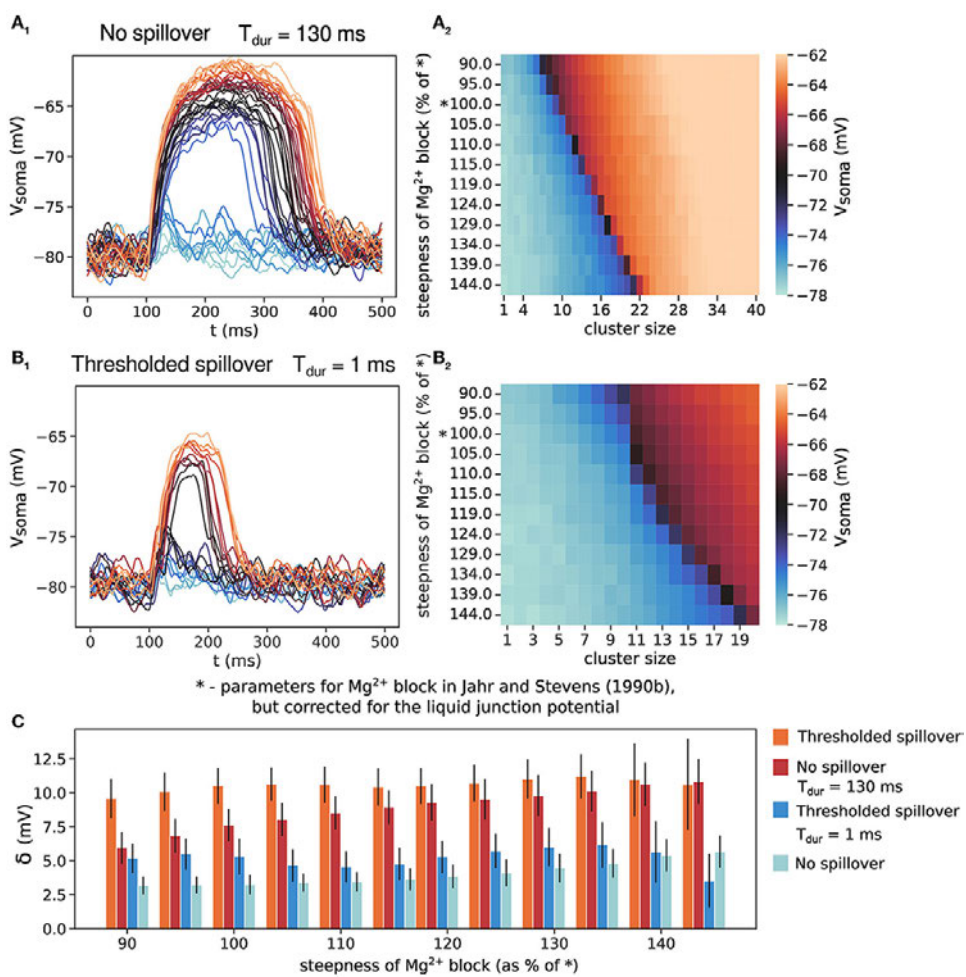
much shorter than T_{dur} for eNMDARs. Increasing the spine T_{dur} in the case without spillover results in all-or-none plateau potential behavior. On the other hand, decreasing the eNMDAR T_{dur} in the case with spillover produces rather graded plateau potentials with much smaller voltage jumps (Fig. 4.14B₁, B₂). The all-or-none behavior in Fig. 4.14A₁ and graded behavior in Fig. 4.14B₁ are also evident in the sharp and gradual color transitions in the respective heatmaps in Figs. 4.14A₂, B₂ as well as in the comparison of the voltage jump, δ in Fig. 4.14C.

4.3.3 Low extrasynaptic glutamate concentration can produce all-or-none plateau potentials

The extrasynaptic glutamate concentration (the parameter T_{max} for eNMDARs) used in the simulations in Figs. 4.13 and 4.14 is 200 μM , an order of magnitude higher than the measurements reported in Okubo et al. (2010). NMDARs have high-affinity for glutamate, in the low-micromolar range, so they should be able to respond to low glutamate concentrations. We have tested the thresholded spillover model with lower values for extracellular glutamate concentration (the parameter T_{max}) and the results are shown in Fig. 4.15. Concentrations of 10 and 20 μM , significantly higher than the NMDAR dissociation constant, $K_d = 2.5 \mu\text{M}$, produce all-or-none plateau potentials (Fig. 4.15A, B). For concentrations of 5 μM , the size of the voltage jump is reduced, although the plateau potentials retain their all-or-none quality. Concentrations comparable to the NMDAR K_d do not activate the eNMDARs fully and produce graded plateau potential amplitude (Fig. 4.15D, E).

4.3.4 Discussion

We studied what causes the all-or-none quality of plateau potentials, and found that glutamate spillover robustly provides this property. Considering the large natural variability in NMDAR isoform properties, which computational modeling suggests will in most cases not produce all-or-none plateaus, glutamate spillover might indeed account for the experimentally observed all-or-none quality. There are several relevant questions for further study in this area. One is the role of astrocytes in glutamate spillover, which have been proposed to even reverse function from glutamate uptake to glutamate release when they have been saturated with glutamate (Malarkey and



Parpura 2008). This could provide even more pronounced supralinear behavior in the plateau potentials through nonlinear glutamate accumulation in the extrasynaptic space. Another is how LTP, with the reported withdrawal of astrocytic processes surrounding the synapse, affects glutamate spillover and its accumulation in the extrasynaptic space (Henneberger et al. 2020). More experimental studies are required to uncover how these processes are regulated, and simulations can help to test the impacts of various mechanisms for controlling glutamate spillover on plateau potential generation.

Lastly, to conclusively study the all-or-none plateau potential behavior, it is necessary to apply types of experimental stimulation which mimic the in vivo processes as closely as possible. Most of the results reporting all-or-none plateau potentials deliver glutamate to the neuron directly through an electrode, or through laser uncaging, which have the advantage of precisely controlling the stimulation strength. However, in vivo glutamate spillover is expected to occur when glutamate escapes the narrow synaptic cleft after synaptic activation. This suggests that techniques like electrical focal synaptic stimulation, which would trigger presynaptic neurotransmitter re-

Figure 4.14 (*preceding page*): The effect of transmitter pulse duration, T_{dur} , on the all-or-none behavior of plateau potentials. (A₁) No glutamate spillover but long transmitter pulse duration produces all-or-none plateau potentials. Cluster size up to 40 synapses. (A₂) Somatic amplitude of plateau potential when varying the steepness for the model in (A₁). (B₁) Thresholded glutamate spillover and short transmitter pulse in the extrasynaptic space produces NMDA potentials with a small all-or-none jump. Cluster size up to 20 synapses. (B₂) Somatic amplitude of plateau potential when varying the steepness for the model in (B₁). The "*" on the y-axes in (A₂, B₂) indicates the gating function parameters as in Jahr and Stevens (1990), but corrected for the liquid junction potential. (C) The size of the voltage jump, δ for the scenarios in (A, B). Results for thresholded spillover and no spillover from Figure 4.13E are added for ease of comparison. Results are averages of 30 trials and 11 different dendrites for (A₁) and 50 trials and 11 different dendrites for (B₁), with clusters positioned at approximately the same distance from the soma. Error bars represent standard deviation. Figure and caption reproduced from Trpevski et al. (2023). © 2023 Trpevski, Khodadadi, Carannante and Hellgren Kotaleski. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

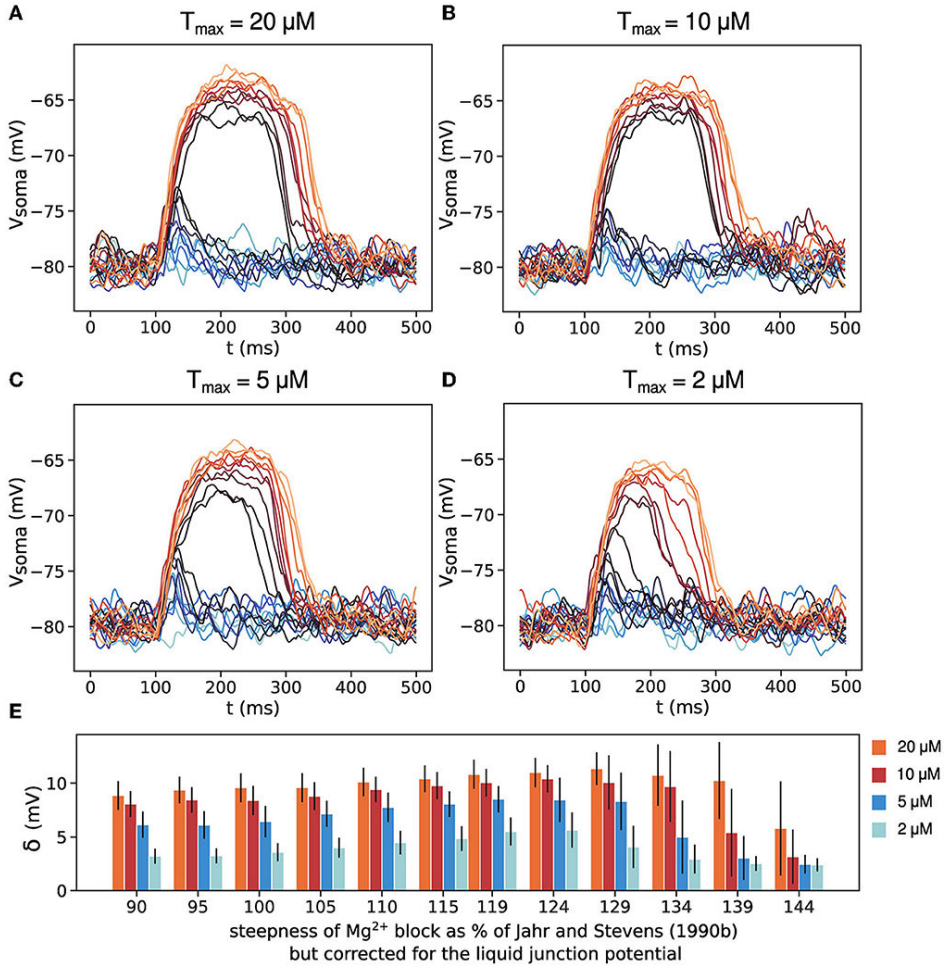


Figure 4.15: The effect of low extrasynaptic glutamate concentration. (A–D) Somatic voltage traces for thresholded spillover and 20, 10, 5, and 2 μM extracellular glutamate concentration, T_{\max} . (E) Size of voltage jump δ for the four values of T_{\max} . Results are averages of 50 trials in 11 different dendrites with clusters positioned at approximately the same distance from the soma. Error bars represent standard deviation. Figure and caption reproduced from Trpevski et al. (2023). © 2023 Trpevski, Khodadadi, Carannante and Hellgren Kotaleski. CC BY 4.0 licence <https://creativecommons.org/licenses/by/4.0/>.

lease, might provide physiologically more realistic information regarding glutamate spillover. All-or-none dendritic spikes have been reported with this technique (Schiller et al. 2000).

5 Future Outlook

In this thesis we studied the computational properties of single striatal projection neurons, provided by molecular mechanisms such as the signaling networks in their synapses, and the neurotransmitter receptors that enable supralinear, all-or-none input integration through dendritic plateau potentials. Building on the results first obtained by studying the signaling network of the corticostriatal synapse and the mechanism of generation of plateau potentials, the results regarding the computational properties of SPNs give insights into their possible role within the neural circuits for goal-directed learning.

One theme of this thesis is multiscale modeling, which uses different computational methods to study systems in which processes occur at multiple temporal and spatial scales, with the purpose of understanding how properties and interactions at one level affect system function at other levels. In the case of neural circuits, we showed how molecular-level interactions give rise to systems-level features such as coincidence detection, which implied that two neuromodulatory signals are required to enable corticostriatal synaptic strengthening for goal-directed learning. The case of AC5 is an example how the two different SPNs, whose synapses have two variations of the same signaling network, have different roles within the neural circuits for goal-directed learning as a result of small differences in molecular composition. These molecular differences enable them to respond differently to the same neuromodulatory signal carrying information about positive and negative action outcomes, and giving them opposite roles in initiating and suppressing actions.

Together with technological advances such as large-scale multi-electrode arrays, which enable monitoring the activity of possibly thousands of neu-

rons, optogenetics and chemogenetics (through designer receptors exclusively activated by designer drugs (DREADDs)) for the control of specific neuronal populations with light and pharmacological agents, multiscale modeling offers the possibility to formulate and test hypotheses about the role and impact of single or multiple molecular components in the function of whole neural circuits for specific behaviors. Large scale brain research initiatives such as the Human Brain Project and the BRAIN initiative are compiling detailed cellular atlases of the brain, have the goal to develop more tools for probing cellular and circuit function, and have provided computational infrastructure for combining experimental data and models, which should substantially accelerate the cycle of hypothesis formulation (in the form of models), generating predictions, and testing them experimentally.

To understand how synaptic plasticity uses and enhances a neuron's computational properties, many more of its governing molecular mechanisms will need to be uncovered and understood. The surrounding glial cells, specifically astrocytes, appear to have a prominent role in both learning and computation, yet glioscience is only at the beginning of its development. Metaplasticity, with its role in determining the outcome of synaptic plasticity, is also a crucial component in how plasticity operates, and most of its underlying molecular implementation awaits discovery.

References

- Abraham, Wickliffe C. (2008). “Metaplasticity: tuning synapses and networks for plasticity”. *Nature Reviews Neuroscience* 9 (5), p. 387. DOI: 10.1038/nrn2356
- Abraham, Wickliffe C., Oliver D. Jones, and David L. Glanzman (2019). “Is plasticity of synapses the mechanism of long-term memory storage?” *npj Science of Learning* 4, p. 9. DOI: 10.1038/s41539-019-0048-y
- Anholt, Robert R. H. (1994). “Signal integration in the nervous system: adenylate cyclases as molecular coincidence detectors.” *Trends Neurosci* 17, pp. 37–41. DOI: 10.1016/0166-2236(94)90033-7
- Antic, Srdjan D., Wen-Liang Zhou, Anna R. Moore, Shaina M. Short, and Katerina D. Ikonomu (2010). “The decade of the dendritic NMDA spike”. *Journal of Neuroscience Research* 88 (14), 2991–3001. DOI: 10.1002/jnr.22444
- Arber, Silvia and Rui M. Costa (2022). “Networking brainstem and basal ganglia circuits for movement”. *Nature Reviews Neuroscience* 23 (6), pp. 342–360. DOI: 10.1038/s41583-022-00581-w
- Arvanov, Viktor L. and Rex Y. Wang (1997). “NMDA-induced response in pyramidal neurons of the rat medial prefrontal cortex slices consists of NMDA and non-NMDA components”. *Brain Research* 768 (1), pp. 361–364. DOI: [https://doi.org/10.1016/S0006-8993\(97\)00842-1](https://doi.org/10.1016/S0006-8993(97)00842-1)
- Berridge, Kent C. (2000). “Reward learning: Reinforcement, incentives, and expectations”. In: *Psychology of Learning and Motivation Volume 40*. Elsevier, 223–278. DOI: 10.1016/s0079-7421(00)80022-5
- Bicknell, Brendan A. and Michael Häusser (2021). “A synaptic learning rule for exploiting nonlinear dendritic computation”. *Neuron* 109 (24), 4001–4017.e10. DOI: 10.1016/j.neuron.2021.09.044

- Bienenstock, Elie L., Leon N. Cooper, and Paul W. Munro (1982). “Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex”. *Journal of Neuroscience* 2 (1), pp. 32–48. DOI: 10.1523/JNEUROSCI.02-01-00032.1982
- Blicher, Andreas and Thomas Heimburg (2013). “Voltage-Gated Lipid Ion Channels”. *PLoS ONE* 8 (6), e65707. DOI: 10.1371/journal.pone.0065707
- Bourne, Henry R. and Roger Nicoll (1993). “Molecular machines integrate coincident synaptic signals”. *Cell* 72, 65–75. DOI: 10.1016/s0092-8674(05)80029-7
- Brandalise, Federico, Stefano Carta, Fritjof Helmchen, John Lisman, and Urs Gerber (2016). “Dendritic NMDA spikes are necessary for timing-dependent associative LTP in CA3 pyramidal cells”. *Nature Communications* 7 (1). DOI: 10.1038/ncomms13480
- Bruce, Neil J., Daniele Narzi, Daniel Trpevski, Siri C. van Keulen, Anu G. Nair, Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, and Jeanette Hellgren Kotaleski (2019). “Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals”. *PLOS Computational Biology* 15 (10), e1007382. DOI: 10.1371/journal.pcbi.1007382
- Bucher, Dirk and Peter A. V. Anderson (2015). “Evolution of the first nervous systems – what can we surmise?” *Journal of Experimental Biology* 218 (4), 501–503. DOI: 10.1242/jeb.111799
- Carter, Adam G. and Bernardo L. Sabatini (2004). “State-Dependent Calcium Signaling in Dendritic Spines of Striatal Medium Spiny Neurons”. *Neuron* 44 (3), pp. 483–493. DOI: 10.1016/j.neuron.2004.10.013
- Carter, Adam G., Gilberto J. Soler-Llavina, and Bernardo L. Sabatini (2007). “Timing and Location of Synaptic Inputs Determine Modes of Sub-threshold Integration in Striatal Medium Spiny Neurons”. *The Journal of Neuroscience* 27 (33), 8967–8977. DOI: 10.1523/jneurosci.2798-07.2007
- Centonze, Diego, Barbara Picconi, Paolo Gubellini, Giorgio Bernardi, and Paolo Calabresi (2001). “Dopaminergic control of synaptic plasticity in the dorsal striatum”. *European Journal of Neuroscience* 13 (6), 1071–1077. DOI: 10.1046/j.0953-816x.2001.01485.x
- Chalifoux, Jason R. and Adam G. Carter (2011). “Glutamate Spillover Promotes the Generation of NMDA Spikes”. *The Journal of Neuroscience* 31 (45), 16435–16446. DOI: 10.1523/jneurosci.2777-11.2011

- Chen, Li and Li-Yen Mae Huang (1992). “Protein kinase C reduces Mg^{2+} block of NMDA-receptor channels as a mechanism of modulation”. *Nature* 356, pp. 521–523. DOI: 10.1038/356521a0
- Chen, Xiaowei, Ulrich Leischner, Nathalie L. Rochefort, Israel Nelken, and Arthur Konnerth (2011). “Functional mapping of single spines in cortical neurons in vivo”. *Nature* 475 (7357), 501–505. DOI: 10.1038/nature10193
- Chen-Goodspeed, Misty, Abolanle N. Lukan, and Carmen W. Dessauer (2005). “Modeling of Galpha(s) and Galpha(i) regulation of human type V and VI adenylyl cyclase.” *J Biol Chem* 280, pp. 1808–16. DOI: 10.1074/jbc.M409172200
- Cheng, Heng-Wei, Jose A. Rafols, Harry G. Goshgarian, Yakir Anavi, Jiang Tong, and Thomas H. McNeill (1997). “Differential Spine Loss and Regrowth of Striatal Neurons Following Multiple Forms of Deafferentation: A Golgi Study”. *Experimental Neurology* 147 (2), pp. 287–298. DOI: 10.1006/exnr.1997.6618
- Chiu, Delia N. and Brett C. Carter (2022). “Synaptic NMDA receptor activity at resting membrane potentials”. *Front. Cell. Neurosci.* 16, p. 916626. DOI: 10.3389/fncel.2022.916626
- Church, Timothy W., Parul Tewatia, Saad Hannan, João Antunes, Olivia Eriksson, Trevor G. Smart, Jeanette Hellgren Kotaleski, and Matthew G. Gold (2021). “AKAP79 enables calcineurin to directly suppress protein kinase A activity”. *eLife* 10, e68164. DOI: 10.7554/eLife.68164
- Clark, Beverley A., Mark Farrant, and Stuart G. Cull-Candy (1997). “A Direct Comparison of the Single-Channel Properties of Synaptic and Extrasynaptic NMDA Receptors”. *The Journal of Neuroscience* 17 (1), 107–116. DOI: 10.1523/jneurosci.17-01-00107.1997
- Clements, John D., Robin A. J. Lester, Gang Tong, Craig E. Jahr, and Gary L. Westbrook (1992). “The Time Course of Glutamate in the Synaptic Cleft”. *Science* 258 (5087), 1498–1501. DOI: 10.1126/science.1359647
- Cui, Yihui, Ilya Prokin, Hao Xu, Bruno Delord, Stephane Genet, Laurent Venance, and Hugues Berry (2016). “Endocannabinoid dynamics gate spike-timing dependent depression and potentiation”. *eLife* 5. DOI: 10.7554/eLife.13185
- Day, Michelle, Marziyeh Belal, William C. Surmeier, Alexandria Melendez, David Wokosin, Tatiana Tkatch, Vernon R. J. Clarke, and D. James Surmeier (2024). “GABAergic regulation of striatal spiny projection neu-

- rons depends upon their activity state”. *PLOS Biology* 22 (1), e3002483. DOI: 10.1371/journal.pbio.3002483
- Delmeire, D., D. Flamez, S. A. Hinke, J. J. Cali, D. Pipeleers, and F. Schuit (2003). “Type VIII adenylyl cyclase in rat beta cells: coincidence signal detector/generator for glucose and GLP-1.” *Diabetologia* 46, pp. 1383–93. DOI: 10.1007/s00125-003-1203-8
- Dessauer, Carmen W., John J. G. Tesmer, Stephen R. Sprang, and Alfred G. Gilman (1998). “Identification of a $G_{i\alpha}$ Binding Site on Type V Adenylyl Cyclase”. *Journal of Biological Chemistry* 273 (40), 25831–25839. DOI: 10.1074/jbc.273.40.25831
- Destexhe, Alain, Zachary F. Mainen, and Terrence J. Sejnowski (1994). “An Efficient Method for Computing Synaptic Conductances Based on a Kinetic Model of Receptor Binding”. *Neural Computation* 6 (1), 14–18. DOI: 10.1162/neco.1994.6.1.14
- Dorman, Daniel B, Joanna Jędrzejewska-Szmek, and Kim T Blackwell (2018). “Inhibition enhances spatially-specific calcium encoding of synaptic input patterns in a biologically constrained model”. *eLife* 7. DOI: 10.7554/eLife.38588
- Doron, Michael, Giuseppe Chindemi, Eilif Muller, Henry Markram, and Idan Segev (2017). “Timed synaptic inhibition shapes NMDA spikes influencing local dendritic processing and global I/O properties of cortical neurons”. *Cell Reports* 21, pp. 1550–1561. DOI: 10.1016/j.celrep.2017.10.035
- Du, Kai, Yu-Wei Wu, Robert Lindroos, Yu Liu, Balázs Rózsa, Gergely Katona, Jun B. Ding, and Jeanette Hellgren Kotaleski (2017). “Cell-type-specific inhibition of the dendritic plateau potential in striatal spiny projection neurons”. *Proceedings of the National Academy of Sciences* 114 (36). DOI: 10.1073/pnas.1704893114
- Ecker, András, Armando Romani, Sára Sárny, Szabolcs Káli, Michele Migliore, Joanne Falck, Sigrun Lange, Audrey Mercer, Alex M. Thomson, Eilif Muller, Michael W. Reimann, and Srikanth Ramaswamy (2020). “Data-driven integration of hippocampal CA1 synaptic physiology in silico”. *Hippocampus* 30 (11), 1129–1145. DOI: 10.1002/hipo.23220
- Esmaili, Vahid, Keita Tamura, Georgios Foustoukos, Anastasiia Oryshchuk, Sylvain Crochet, and Carl CH Petersen (2020). “Cortical circuits for transforming whisker sensation into goal-directed licking”. *Current Opinion in Neurobiology* 65, 38–48. DOI: 10.1016/j.conb.2020.08.003

- Farinella, Matteo, Daniel T. Ruedt, Padraig Gleeson, Frederic Lanore, and R. Angus Silver (2014). “Glutamate-bound NMDARs arising from in vivo-like network activity extend spatio-temporal integration in a L5 cortical pyramidal cell model”. *PLoS Comput. Biol.* 10, e1003590. DOI: 10.1371/journal.pcbi.1003590
- Feldman, Daniel E. (2012). “The Spike-Timing Dependence of Plasticity”. *Neuron* 75 (4), pp. 556–571. DOI: 10.1016/j.neuron.2012.08.001
- Fino, Elodie, Vincent Paille, Yihui Cui, Teresa Morera-Herreras, Jean-Michel Deniau, and Laurent Venance (2010). “Distinct coincidence detectors govern the corticostriatal spike timing-dependent plasticity: Corticostriatal spike timing-dependent plasticity pharmacology”. *The Journal of Physiology* 588 (16), 3045–3062. DOI: 10.1113/jphysiol.2010.188466
- Fisher, Simon D., Paul B. Robertson, Melony J. Black, Peter Redgrave, Mark A. Sagar, Wickliffe C. Abraham, and John N.J. Reynolds (2017). “Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo”. *Nature Communications* 8 (1). DOI: 10.1038/s41467-017-00394-x
- Gao, Peng P., Joseph W. Graham, Wen-Liang Zhou, Jinyoung Jang, Sergio Angulo, Salvador Dura-Bernal, Michael Hines, William W. Lytton, and Srdjan D. Antic (2021). “Local glutamate-mediated dendritic plateau potentials change the state of the cortical pyramidal neuron”. *Journal of Neurophysiology* 125 (1), 23–42. DOI: 10.1152/jn.00734.2019
- Garcia-Munoz, Marianela, Violeta G. Lopez-Huerta, Luis Carrillo-Reid, and Gordon W. Arbuthnott (2015). “Extrasynaptic glutamate NMDA receptors: Key players in striatal function”. *Neuropharmacology* 89, 54–63. DOI: 10.1016/j.neuropharm.2014.09.013
- Ghose, Geoffrey M. and John Maunsell (1999). “Specialized Representations in Visual Cortex”. *Neuron* 24 (1), 79–85. DOI: 10.1016/s0896-6273(00)80823-5
- Gold, Stephen J., Yan G. Ni, Henrik G. Dohlman, and Eric J. Nestler (1997). “Regulators of G-protein signaling (RGS) proteins: region-specific expression of nine subtypes in rat brain.” *J Neurosci* 17, pp. 8024–37. DOI: 10.1523/JNEUROSCI.17-20-08024.1997
- Goldstein, Steve A. N., Detlef Bockenhauer, Ita O’Kelly, and Noam Zilberberg (2001). “Potassium leak channels and the KCNK family of two-domain subunits”. *Nature Reviews Neuroscience* 2 (3), 175–184. DOI: 10.1038/35058574

- Graupner, Michael and Nicolas Brunel (2012). “Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location”. *Proceedings of the National Academy of Sciences* 109 (10), pp. 3991–3996. DOI: 10.1073/pnas.1109359109
- Gray, Charles M. (1999). “The Temporal Correlation Hypothesis of Visual Feature Integration”. *Neuron* 24 (1), 31–47. DOI: 10.1016/s0896-6273(00)80820-x
- Greenspan, Ralph J. (2007). *An Introduction to Nervous Systems*. Cold Spring Harbor Laboratory Press
- Gütig, R., R. Aharonov, S. Rotter, and Haim Sompolinsky (2003). “Learning Input Correlations through Nonlinear Temporally Asymmetric Hebbian Plasticity”. *Journal of Neuroscience* 23 (9), pp. 3697–3714. DOI: 10.1523/JNEUROSCI.23-09-03697.2003
- Hanke, Wolfgang, Konrad Kaufmann, and Ayus Corcia (1989). *Ion Channel Fluctuations in Pure Lipid Bilayer Membranes: Control by Voltage*. Caruaru, Brazil
- Harris, Alexander Z. and Diana L. Pettit (2007). “Extrasynaptic and synaptic NMDA receptors form stable and uniform pools in rat hippocampal slices”. *The Journal of Physiology* 584 (2), 509–519. DOI: 10.1113/jphysiol.2007.137679
- Henneberger, Christian, Lucie Bard, Aude Panatier, James P. Reynolds, Olga Kopach, Nikolay I. Medvedev, Daniel Minge, Michel K. Herde, Stefanie Anders, Igor Kraev, Janosch P. Heller, Sylvain Rama, Kaiyu Zheng, Thomas P. Jensen, Inmaculada Sanchez-Romero, Colin J. Jackson, Harald Janovjak, Ole Petter Ottersen, Erlend Arnulf Nagelhus, Stephane H. R. Oliet, Michael G. Stewart, U. Valentin Nägerl, and Dmitri A. Rusakov (2020). “LTP Induction Boosts Glutamate Spillover by Driving Withdrawal of Perisynaptic Astroglia.” *Neuron* 108, 919–936.e11
- Herring, Bruce E. and Roger A. Nicoll (2016). “Long-Term Potentiation: From CaMKII to AMPA Receptor Trafficking”. *Annual Review of Physiology* 78 (1), 351–365. DOI: 10.1146/annurev-physiol-021014-071753
- Higley, Michael J. and Bernardo L. Sabatini (2010). “Competitive regulation of synaptic Ca^{2+} influx by D_2 dopamine and A_{2A} adenosine receptors.” *Nat Neurosci* 13, pp. 958–66
- Hwang, Fuu-Jiun, Richard H. Roth, Yu-Wei Wu, Yue Sun, Destany K. Kwon, Yu Liu, and Jun B. Ding (2022). “Motor learning selectively strengthens cortical and striatal synapses of motor engram neurons”. *Neuron* 110 (17), 2790–2801.e5. DOI: 10.1016/j.neuron.2022.06.006

- Iacaruso, M. Florencia, Ioana T. Gasler, and Sonja B. Hofer (2017). “Synaptic organization of visual space in primary visual cortex”. *Nature* 547 (7664), 449–452. DOI: 10.1038/nature23019
- Impey, Soren, Gary Wayman, Zhiliang Wu, and Daniel R. Storm (1994). “Type I adenylyl cyclase functions as a coincidence detector for control of cyclic AMP response element-mediated transcription: synergistic regulation of transcription by Ca²⁺ and isoproterenol.” *Mol Cell Biol* 14, pp. 8272–81
- Jahr, Craig E. and Charles F. Stevens (1990). “Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics”. *The Journal of Neuroscience* 10 (9), 3178–3182. DOI: 10.1523/jneurosci.10-09-03178.1990
- Jia, Hongbo, Nathalie L. Rochefort, Xiaowei Chen, and Arthur Konnerth (2010). “Dendritic organization of sensory input to cortical neurons in vivo”. *Nature* 464 (7293), 1307–1312. DOI: 10.1038/nature08947
- Ju, Niansheng, Yang Li, Fang Liu, Hongfei Jiang, Stephen L. Macknik, Susana Martinez-Conde, and Shiming Tang (2020). “Spatiotemporal functional organization of excitatory synaptic inputs onto macaque V1 neurons”. *Nature Communications* 11 (1). DOI: 10.1038/s41467-020-14501-y
- Jędrzejewska-Szmek, Joanna, Sriraman Damodaran, Daniel B. Dorman, and Kim T. Blackwell (2017). “Calcium dynamics predict direction of synaptic plasticity in striatal spiny projection neurons”. *European Journal of Neuroscience* 45 (8), pp. 1044–1056. DOI: <https://doi.org/10.1111/ejn.13287>
- Kaufmann, Konrad and Israel Silman (Sept. 1983). “The induction by protons of ion channels through lipid bilayer membranes”. *Biophysical Chemistry* 18 (2), pp. 89–99. DOI: 10.1016/0301-4622(83)85002-9
- Keijzer, Fred, Marc van Duijn, and Pamela Lyon (2013). “What nervous systems do: early evolution, input–output, and the skin brain thesis”. *Adaptive Behavior* 21 (2), 67–85. DOI: 10.1177/1059712312465330
- Kerlin, Aaron, Boaz Mohar, Daniel Flickinger, Bryan J MacLennan, Matthew B Dean, Courtney Davis, Nelson Spruston, and Karel Svoboda (2019). “Functional clustering of dendritic activity during decision-making”. *eLife* 8. DOI: 10.7554/eLife.46966
- Keulen, Siri Camee van and Ursula Rothlisberger (2017). “Exploring the inhibition mechanism of adenylyl cyclase type 5 by n-terminal myris-

- toyolated Gail”. *PLOS Computational Biology* 13 (9), e1005673. DOI: 10.1371/journal.pcbi.1005673
- Khodadadi, Zahra, Daniel Trpevski, Robert Lindroos, and Jeanette Hellgren Kotaleski (2024). “Local, calcium- and reward-based synaptic learning rule that enhances dendritic nonlinearities can solve the nonlinear feature binding problem”. *eLife* 13 RP97274. DOI: 10.7554/eLife.97274.1
- Kirchner, Jan H. and Julijana Gjorgjieva (2021). “Emergence of local and global synaptic organization on cortical dendrites”. *Nature Communications* 12 (1). DOI: 10.1038/s41467-021-23557-3
- Kleindienst, Thomas, Johan Winnubst, Claudia Roth-Alpermann, Tobias Bonhoeffer, and Christian Lohmann (2011). “Activity-Dependent Clustering of Functional Synaptic Inputs on Developing Hippocampal Dendrites”. *Neuron* 72 (6), 1012–1024. DOI: 10.1016/j.neuron.2011.10.015
- Kumar, Amit, Oded Schiff, Edi Barkai, Bartlett W Mel, Alon Poleg-Polsky, and Jackie Schiller (2018). *Author response: NMDA spikes mediate amplification of inputs in the rat piriform cortex*. DOI: 10.7554/eLife.38446.018
- Lalive, Arnaud L., Anthony D. Lien, Thomas K. Roseberry, Christopher H. Donahue, and Anatol C. Kreitzer (2018). “Motor thalamus supports striatum-driven reinforcement”. *eLife* 7. DOI: 10.7554/eLife.34032
- Landau, Marius, Sherif Elsabbagh, Harald Gross, Adrian Fuchs, Anita C.F. Schultz, and Joachim E. Schultz (2024). “A new class of receptors: Lipids regulate mammalian Gs α -stimulated adenylyl cyclase activities via their membrane anchors”. *bioRxiv*. DOI: 10.1101/2024.08.06.606792
- Lerner, Talia N. and Anatol C. Kreitzer (2012). “RGS4 is required for dopaminergic control of striatal LTD and susceptibility to parkinsonian motor deficits.” *Neuron* 73, pp. 347–59. DOI: 10.1016/j.neuron.2011.11.015
- Linder, Jürgen U. (2006). “Class III adenylyl cyclases: molecular mechanisms of catalysis and regulation”. *Cellular and Molecular Life Sciences* 63 (15), 1736–1751. DOI: 10.1007/s00018-006-6072-0
- Lindroos, Robert and Jeanette Hellgren Kotaleski (2021). “Predicting complex spikes in striatal projection neurons of the direct pathway following neuromodulation by acetylcholine and dopamine.” *Eur J Neurosci* 53, pp. 2117–2134. DOI: 10.1111/ejn.14891
- Lustig, Kevin D., Bruce R. Conklin, Paul Herzmark, Ron Taussig, and Henry R. Bourne (1993). “Type II adenylyl cyclase integrates coincident signals

- from Gs, Gi, and Gq." *J Biol Chem* 268, pp. 13900–5. DOI: 10.1016/S0021-9258(19)85187-6
- Mahon, Séverine, Nicolas Vautrelle, Laurent Pezard, Seán J. Slaght, Jean-Michel Deniau, Guy Chouvet, and Stéphane Charpier (2006). "Distinct Patterns of Striatal Medium Spiny Neuron Activity during the Natural Sleep–Wake Cycle". *The Journal of Neuroscience* 26 (48), 12587–12595. DOI: 10.1523/jneurosci.3987-06.2006
- Major, Guy, Alon Polsky, Winfried Denk, Jackie Schiller, and David W. Tank (2008). "Spatiotemporally Graded NMDA Spike/Plateau Potentials in Basal Dendrites of Neocortical Pyramidal Neurons". *Journal of Neurophysiology* 99 (5), 2584–2601. DOI: 10.1152/jn.00011.2008
- Malarkey, Erik B. and Vladimir Parpura (2008). "Mechanisms of glutamate release from astrocytes." *Neurochem Int* 52, pp. 142–54. DOI: 10.1016/j.neuint.2007.06.005
- Malsburg, Christoph von der (1999). "The What and Why of Binding". *Neuron* 24 (1), 95–104. DOI: 10.1016/S0896-6273(00)80825-9
- McMenimen, Kathryn A., E. James Petersson, Henry A. Lester, and Dennis A. Dougherty (2006). "Probing the Mg²⁺ blockade site of an N-methyl-D-aspartate (NMDA) receptor with unnatural amino acid mutagenesis". *ACS Chem. Biol.* 1, pp. 227–234. DOI: 10.1021/cb6000944
- McVey, Michael, Jennifer Hill, Allyn Howlett, and Claudette Klein (1999). "Adenylyl cyclase, a coincidence detector for nitric oxide." *J Biol Chem* 274, pp. 18887–92. DOI: 10.1074/jbc.274.27.18887
- Mons, N, JL Guillou, and R Jaffard (1999). "The role of Ca²⁺/calmodulin-stimulable adenylyl cyclases as molecular coincidence detectors in memory formation." *Cell Mol Life Sci* 55, pp. 525–33. DOI: 10.1007/s000180050311
- Nair, Anu G., Upinder S. Bhalla, and Jeanette Hellgren Kotaleski (2016). "Role of DARPP-32 and ARPP-21 in the Emergence of Temporal Constraints on Striatal Calcium and Dopamine Integration". *PLOS Computational Biology* 12 (9), e1005080. DOI: 10.1371/journal.pcbi.1005080
- Nair, Anu G., Liliana R.V. Castro, Marianne El Khoury, Victor Gorgievski, Bruno Giros, Eleni T. Tzavara, Jeanette Hellgren-Kotaleski, and Pierre Vincent (2019). "The high efficacy of muscarinic M4 receptor in D1 medium spiny neurons reverses striatal hyperdopaminergia". *Neuropharmacology* 146, 74–83. DOI: 10.1016/j.neuropharm.2018.11.029
- Nair, Anu G., Omar Gutierrez-Arenas, Olivia Eriksson, Pierre Vincent, and Jeanette Hellgren Kotaleski (2015). "Sensing Positive versus Negative Reward Signals through Adenylyl Cyclase-Coupled GPCRs in Direct

- and Indirect Pathway Striatal Medium Spiny Neurons”. *The Journal of Neuroscience* 35 (41), pp. 14017–14030. DOI: 10.1523/jneurosci.0730-15.2015
- Navarro, Gemma, Arnau Cordoní, Verónica Casadó-Anguera, Estefanía Moreno, Ning-Sheng Cai, Antoni Cortés, Enric I. Canela, Carmen W. Dessauer, Vicent Casadó, Leonardo Pardo, Carme Lluís, and Sergi Ferré (2018). “Evidence for functional pre-coupled complexes of receptor heteromers and adenylyl cyclase”. *Nature Communications* 9 (1). DOI: 10.1038/s41467-018-03522-3
- Niculescu, Dragos, Kristin Michaelsen-Preusse, Ülkü Güner, René van Dorland, Corette J. Wierenga, and Christian Lohmann (2018). “A BDNF-Mediated Push-Pull Plasticity Mechanism for Synaptic Clustering”. *Cell Reports* 24 (8), 2063–2074. DOI: 10.1016/j.celrep.2018.07.073
- Nowak, L., P. Bregestovski, P. Ascher, A. Herbet, and A. Prochiantz (1984). “Magnesium gates glutamate-activated channels in mouse central neurons”. *Nature* 307, pp. 462–465. DOI: 10.1038/307462a0
- Oikonomou, Katerina D., Shaina M. Short, Matthew T. Rich, and Srdjan D. Antic (2012). “Extrasynaptic Glutamate Receptor Activation as Cellular Bases for Dynamic Range Compression in Pyramidal Neurons”. *Frontiers in Physiology* 3. DOI: 10.3389/fphys.2012.00334
- Oikonomou, Katerina D., Mandakini B. Singh, Enas V. Sterjanaj, and Srdjan D. Antic (2014). “Spiny neurons of amygdala, striatum, and cortex use dendritic plateau potentials to detect network UP states”. *Frontiers in Cellular Neuroscience* 8. DOI: 10.3389/fncel.2014.00292
- Okubo, Yohei, Hiroshi Sekiya, Shigeyuki Namiki, Hirokazu Sakamoto, Sho Inuma, Miwako Yamasaki, Masahiko Watanabe, Kenzo Hirose, and Masamitsu Iino (2010). “Imaging extrasynaptic glutamate dynamics in the brain”. *Proceedings of the National Academy of Sciences* 107 (14), 6526–6531. DOI: 10.1073/pnas.0913154107
- O’Doherty, John P., Jeffrey Cockburn, and Wolfgang M. Pauli (2017). “Learning, Reward, and Decision Making”. *Annual Review of Psychology* 68 (1), 73–100. DOI: 10.1146/annurev-psych-010416-044216
- Papouin, Thomas and Stéphane H. R. Oliet (2014). “Organization, control and function of extrasynaptic NMDA receptors”. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369 (1654), p. 20130601. DOI: 10.1098/rstb.2013.0601

- Petralia, Ronald S. (2012). “Distribution of Extrasynaptic NMDA Receptors on Neurons”. *The Scientific World Journal* 2012, 1–11. DOI: 10.1100/2012/267120
- Plotkin, Joshua L, Michelle Day, and D James Surmeier (2011). “Synaptically driven state transitions in distal dendrites of striatal spiny neurons”. *Nature Neuroscience* 14 (7), 881–888. DOI: 10.1038/nn.2848
- Plotkin, Joshua L., Weixing Shen, Igor Rafalovich, Luke E. Sebel, Michelle Day, C. Savio Chan, and D. James Surmeier (2013). “Regulation of dendritic calcium release in striatal spiny projection neurons”. *Journal of Neurophysiology* 110 (10), 2325–2336. DOI: 10.1152/jn.00422.2013
- Poleg-Polsky, Alon (2015). “Effects of neural morphology and input distribution on synaptic processing by global and focal NMDA-spikes”. *PLoS ONE* 10, e0140254. DOI: 10.1371/journal.pone.0140254
- Purves, Dale, George J. Augustine, David Fitzpatrick, William C. Hall, Anthony-Samuel LaMantia, and Leonard E. White (2012). *Neuroscience*. 5th. Sinauer Associates
- Reynolds, John H. and Robert Desimone (1999). “The Role of Neural Mechanisms of Attention in Solving the Binding Problem”. *Neuron* 24 (1), 19–29. DOI: 10.1016/s0896-6273(00)80819-3
- Reynolds, John N. J., Riccardo Avvisati, Paul D. Dodson, Simon D. Fisher, Manfred J. Oswald, Jeffery R. Wickens, and Yan-Feng Zhang (2022). “Coincidence of cholinergic pauses, dopaminergic activation and depolarisation of spiny projection neurons drives synaptic plasticity in the striatum”. *Nature Communications* 13 (1). DOI: 10.1038/s41467-022-28950-0
- Rhodes, Paul (2006). “The properties and implications of NMDA spikes in neocortical pyramidal cells”. *J. Neurosci.* 26, pp. 6704–6715. DOI: 10.1523/JNEUROSCI.3791-05.2006
- Roskies, Adina L. (1999). “The binding problem.” *Neuron* 24, pp. 7–9, 111–25
- Rossum, M. C. W. van, G. Q. Bi, and G. G. Turrigiano (2000). “Stable Hebbian Learning from Spike Timing-Dependent Plasticity”. *Journal of Neuroscience* 20 (23), pp. 8812–8821. DOI: 10.1523/JNEUROSCI.20-23-08812.2000
- Rubin, Jonathan, Daniel D. Lee, and H. Sompolinsky (2001). “Equilibrium Properties of Temporally Asymmetric Hebbian Plasticity”. *Phys. Rev. Lett.* 86 (2), pp. 364–367. DOI: 10.1103/PhysRevLett.86.364

- Sanabria, Branden D., Sindhuja S. Baskar, Alex J. Yonk, Iván Linares-García, Victoria E. Abaira, Christian R. Lee, and David J. Margolis (2023). “Cell-Type Specific Connectivity of Whisker-Related Sensory and Motor Cortical Input to Dorsal Striatum”. *eneuro* 11 (1), ENEURO.0503–23.2023. DOI: 10.1523/eneuro.0503–23.2023
- Schiess, Markus, Robert Urbanczik, and Walter Senn (2016). “Somato-dendritic Synaptic Plasticity and Error-backpropagation in Active Dendrites”. *PLoS Computational Biology* 12 (2), e1004638. DOI: 10.1371/journal.pcbi.1004638
- Schiller, Jackie, Guy Major, Helmut J. Koester, and Yitzhak Schiller (2000). “NMDA spikes in basal dendrites of cortical pyramidal neurons”. *Nature* 404 (6775), pp. 285–289
- Scholich, K, JB Mullenix, C Wittpoth, HM Poppleton, SC Pierre, MA Lindorfer, JC Garrison, and TB Patel (1999). “Facilitation of signal onset and termination by adenylyl cyclase.” *Science* 283, pp. 1328–31
- Scholl, Benjamin, Daniel E. Wilson, and David Fitzpatrick (2017). “Local Order within Global Disorder: Synaptic Architecture of Visual Space”. *Neuron* 96 (5), 1127–1138.e4. DOI: 10.1016/j.neuron.2017.10.017
- Schultz, Wolfram (2007). “Multiple Dopamine Functions at Different Time Courses”. *Annual Review of Neuroscience* 30, pp. 259–288. DOI: <https://doi.org/10.1146/annurev.neuro.28.061604.135722>
- (2016). “Dopamine reward prediction-error signalling: a two-component response”. *Nature Reviews Neuroscience* 17 (3), pp. 183–195. DOI: 10.1038/nrn.2015.26
- Sharma, G. and C. F. Stevens (1996). “A mutation that alters magnesium block of N-methyl-D-aspartate receptor channels”. *Proc. Natl. Acad. Sci. U.S.A.* 93, pp. 9259–9263. DOI: 10.1073/pnas.93.17.9259
- Shen, Weixing, Marc Flajolet, Paul Greengard, and D. James Surmeier (2008). “Dichotomous Dopaminergic Control of Striatal Synaptic Plasticity”. *Science* 321 (5890), 848–851. DOI: 10.1126/science.1160575
- Shindou, Tomomi, Mayumi Ochi-Shindou, and Jeffery R. Wickens (2011). “A Ca^{2+} Threshold for Induction of Spike-Timing-Dependent Depression in the Mouse Striatum”. *The Journal of Neuroscience* 31 (36), 13015–13022. DOI: 10.1523/jneurosci.3206–11.2011
- Simms, Brett A. and Gerald W. Zamponi (2014). “Neuronal Voltage-Gated Calcium Channels: Structure, Function, and Dysfunction”. 82 (1)
- Svenningsson, Per, Akinori Nishi, Gilberto Fisone, Jean-Antoine Girault, Angus C. Nairn, and Paul Greengard (2004). “DARPP-32: An Integrator

- of Neurotransmission". *Annual Review of Pharmacology and Toxicology* 44 (1), 269–296. DOI: 10.1146/annurev.pharmtox.44.101802.121415
- Szapiro, Germán and Boris Barbour (2007). "Multiple climbing fibers signal to molecular layer interneurons exclusively via glutamate spillover". *Nature Neuroscience* 10 (6), 735–742. DOI: 10.1038/nn1907
- Takahashi, Naoya, Kazuo Kitamura, Naoki Matsuo, Mark Mayford, Masanobu Kano, Norio Matsuki, and Yuji Ikegaya (2012). "Locally Synchronized Synaptic Inputs". *Science* 335 (6066), 353–356. DOI: 10.1126/science.1210362
- Taussig, Ron, W.J. Tang, J.R. Hepler, and Alfred G. Gilman (1994). "Distinct patterns of bidirectional regulation of mammalian adenylyl cyclases." *Journal of Biological Chemistry* 269 (8), 6093–6100. DOI: 10.1016/s0021-9258(17)37574-9
- Tesmer, John J. G., Roger K. Sunahara, Alfred G. Gilman, and Stephen R. Sprang (1997). "Crystal Structure of the Catalytic Domains of Adenylyl Cyclase in a Complex with $G\alpha \cdot GTP\gamma S$ ". *Science* 278 (5345), 1907–1916. DOI: 10.1126/science.278.5345.1907
- Thura, David and Paul Cisek (2017). "The Basal Ganglia Do Not Select Reach Targets but Control the Urgency of Commitment". *Neuron* 95 (5), 1160–1170.e5. DOI: 10.1016/j.neuron.2017.07.039
- Tran-Van-Minh, Alexandra, Romain D. Cazé, Thérèse Abrahamsson, Laurence Cathala, Boris S. Gutkin, and David A. DiGregorio (2015). "Contribution of sublinear and supralinear dendritic integration to neuronal computations". *Frontiers in Cellular Neuroscience* 9. DOI: 10.3389/fncel.2015.00067
- Treisman, Anne (1999). "Solutions to the Binding Problem". *Neuron* 24 (1), 105–125. DOI: 10.1016/s0896-6273(00)80826-0
- Trpevski, Daniel, Zahra Khodadadi, Ilaria Carannante, and Jeanette Hellgren Kotaleski (2023). "Glutamate spillover drives robust all-or-none dendritic plateau potentials—an *in silico* investigation using models of striatal projection neurons". *Front. Cell. Neurosci.* 17, p. 1196182. DOI: 10.3389/fncel.2023.1196182
- Truman, James W., Jacquelyn Price, Rosa L. Miyares, and Tzumin Lee (2023). "Metamorphosis of memory circuits in *Drosophila* reveals a strategy for evolving a larval brain". *eLife* 12. DOI: 10.7554/elife.80594
- Varga, Zsuzsanna, Hongbo Jia, Bert Sakmann, and Arthur Konnerth (2011). "Dendritic coding of multiple sensory inputs in single cortical neurons

- in vivo". *Proceedings of the National Academy of Sciences* 108 (37), 15420–15425. DOI: 10.1073/pnas.1112355108
- Wilson, Daniel E., David E. Whitney, Benjamin Scholl, and David Fitzpatrick (2016). "Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex". *Nature Neuroscience* 19 (8), 1003–1009. DOI: 10.1038/nn.4323
- Winnubst, Johan, Juliette E. Cheyne, Dragos Niculescu, and Christian Lohmann (2015). "Spontaneous Activity Drives Local Synaptic Plasticity In Vivo". *Neuron* 87 (2), 399–410. DOI: 10.1016/j.neuron.2015.06.029
- Wolf, John A., Jason T. Moyer, Maciej T. Lazarewicz, Diego Contreras, Marianne Benoit-Marand, Patricio O'Donnell, and Leif H. Finkel (2005). "NMDA/AMPA Ratio Impacts State Transitions and Entrainment to Oscillations in a Computational Model of the Nucleus Accumbens Medium Spiny Projection Neuron". *Journal of Neuroscience* 25 (40), pp. 9080–9095. DOI: 10.1523/JNEUROSCI.2220-05.2005
- Wolfe, Jeremy M. and Kyle R. Cave (1999). "The Psychophysical Evidence for a Binding Problem in Human Vision". *Neuron* 24 (1), 11–17. DOI: 10.1016/s0896-6273(00)80818-1
- Xie, Keqiang, Ikuo Masuho, Cameron Brand, Carmen W. Dessauer, and Kirill A. Martemyanov (2012). "The complex of G protein regulator RGS9-2 and G β (5) controls sensitization and signaling kinetics of type 5 adenylyl cyclase in the striatum." *Sci Signal* 5, ra63
- Xie, Keqiang, Ikuo Masuho, Chien-Cheng Shih, Yan Cao, Keita Sasaki, Chun Wan J. Lai, Pyung-Lim Han, Hiroshi Ueda, Carmen W. Dessauer, Michelle E. Ehrlich, Baoji Xu, Barry M. Willardson, and Kirill A. Martemyanov (2015). "Stable G protein-effector complexes in striatal neurons: mechanism of assembly and role in neurotransmitter signaling." *Elife* 4
- Yagishita, Sho, Akiko Hayashi-Takagi, Graham C. Ellis-Davies, Hidetoshi Urakubo, Shin Ishii, and Haruo Kasai (2014). "A critical time window for dopamine actions on the structural plasticity of dendritic spines." *Science* 345, pp. 1616–20. DOI: 10.1126/science.1255514
- Yapo, Cedric, Anu G. Nair, Lorna Clement, Liliana R. Castro, Jeanette Hellgren Kotaleski, and Pierre Vincent (2017). "Detection of phasic dopamine by D1 and D2 striatal medium spiny neurons". *The Journal of Physiology* 595 (24), 7451–7475. DOI: 10.1113/jp274475
- Yger, Pierre and Matthieu Gilson (2015). "Models of Metaplasticity: A Review of Concepts". *Frontiers in Computational Neuroscience* 9. DOI: 10.3389/fncom.2015.00138

Article A



Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals
(2019) PLoS Computational Biology 15(10): e1007382

RESEARCH ARTICLE

Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals

Neil J. Bruce¹*, Daniele Narzi²*, Daniel Trpevski³*, Siri C. van Keulen^{2,4}*, Anu G. Nair⁵, Ursula Röthlisberger^{2*}, Rebecca C. Wade^{1,6,7*}, Paolo Carloni^{8,9,10,11*}, Jeanette Hellgren Kotaleski^{3,12*}

1 Molecular and Cellular Modeling Group, Heidelberg Institute for Theoretical Studies (HITS), Schloss-Heidelberg, Germany, **2** Institut des Sciences et Ingénierie Chimiques, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland, **3** Science for Life Laboratory, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden, **4** Department of Computer Science, Stanford University, Stanford, California, United States of America, **5** Institute of Molecular Life Sciences, University of Zurich, Zurich, Switzerland, **6** Center for Molecular Biology (ZMBH), DKFZ-ZMBH Alliance, Heidelberg University, Heidelberg, Germany, **7** Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Heidelberg, Germany, **8** Department of Physics and Department of Neurobiology, RWTH Aachen University, Aachen, Germany, **9** Institute for Neuroscience and Medicine (INM)-11, Forschungszentrum Jülich, Jülich, Germany, **10** Institute of Neuroscience and Medicine (INM)-9, Forschungszentrum Jülich, Jülich, Germany, **11** Institute for Advanced Simulation (IAS-5), Forschungszentrum Jülich, Jülich, Germany, **12** Department of Neuroscience, Karolinska Institutet, Solna, Sweden

* These authors contributed equally to this work.

* (JHK) (UR); (RCW); (PC);



OPEN ACCESS

Citation: Bruce NJ, Narzi D, Trpevski D, van Keulen SC, Nair AG, Röthlisberger U, et al. (2019) Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals. *PLoS Comput Biol* 15 (10): e1007382. <https://doi.org/10.1371/journal.pcbi.1007382>

Editor: Hugues Berry, Inria, FRANCE

Received: April 1, 2019

Accepted: September 5, 2019

Published: October 30, 2019

Copyright: © 2019 Bruce et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The data and models generated in this work are available for download and visualisation at https://humanbrainproject.github.io/hbp-bsp-live-papers/2019/bruce_et_al_2019/bruce_et_al_2019.html.

Funding: This research has received funding from the European Union Seventh Framework Programme (FP7/2007-2013; <https://ec.europa.eu/research/fp7/>) under grant agreement N° 604102 (HBP) (JHK, PC, RCW), the European Union's Horizon 2020 Framework Programme (<https://ec.europa.eu/horizon/>).

Abstract

Long-term potentiation and depression of synaptic activity in response to stimuli is a key factor in reinforcement learning. Strengthening of the corticostriatal synapses depends on the second messenger cAMP, whose synthesis is catalysed by the enzyme adenylyl cyclase 5 (AC5), which is itself regulated by the stimulatory $G_{\alpha_{olf}}$ and inhibitory G_{α_i} proteins. AC isoforms have been suggested to act as coincidence detectors, promoting cellular responses only when convergent regulatory signals occur close in time. However, the mechanism for this is currently unclear, and seems to lie in their diverse regulation patterns. Despite attempts to isolate the ternary complex, it is not known if $G_{\alpha_{olf}}$ and G_{α_i} can bind to AC5 simultaneously, nor what activity the complex would have. Using protein structure-based molecular dynamics simulations, we show that this complex is stable and inactive. These simulations, along with Brownian dynamics simulations to estimate protein association rates constants, constrain a kinetic model that shows that the presence of this ternary inactive complex is crucial for AC5's ability to detect coincident signals, producing a synergistic increase in cAMP. These results reveal some of the prerequisites for corticostriatal synaptic plasticity, and explain recent experimental data on cAMP concentrations following receptor activation. Moreover, they provide insights into the regulatory mechanisms that control signal processing by different AC isoforms.

europa.eu/programmes/horizon2020/en) for Research and Innovation under the Specific Grant Agreement N°s 720270 (Human Brain Project SGA1) (JHK, PC, RCW, UR, AGN) and 785907 (Human Brain Project SGA2) (JHK, PC, RCW, UR, DT), the Swedish e-Science Research Centre (SeRC; <https://e-science.se/>) (JHK, DT), the Swedish Research Council (<https://www.vr.se/>) (JHK, AGN, DT), the Klaus Tschira Stiftung (www.klaus-tschira-stiftung.de) (RCW, NJB), National Institute on Alcohol Abuse and Alcoholism Grant 2R01AA016022 (www.niaaa.nih.gov) (JHK). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

Adenylyl cyclases (ACs) are enzymes that can translate extracellular signals into the intracellular molecule cAMP, which is thus a 2nd messenger of extracellular events. The brain expresses nine membrane-bound AC variants, and AC5 is the dominant form in the striatum. The striatum is the input stage of the basal ganglia, a brain structure involved in reward learning, i.e. the learning of behaviors that lead to rewarding stimuli (such as food, water, sugar, etc). During reward learning, cAMP production is crucial for strengthening the synapses from cortical neurons onto the striatal principal neurons, and its formation is dependent on several neuromodulatory systems such as dopamine and acetylcholine. It is, however, not understood how AC5 is activated by transient (subsecond) changes in the neuromodulatory signals. Here we combine several computational tools, from molecular dynamics and Brownian dynamics simulations to bioinformatics approaches, to inform and constrain a kinetic model of the AC5-dependent signaling system. We use this model to show how the specific molecular properties of AC5 can detect particular combinations of co-occurring transient changes in the neuromodulatory signals which thus result in a supralinear/synergistic cAMP production. Our results also provide insights into the computational capabilities of the different AC isoforms.

Introduction

Information processing in the brain occurs within circuits of neurons that are interconnected via synapses. The modification of these neuronal circuits, in response to an organism's experiences and interactions with the environment, is crucial for memory and learning, allowing the organism's behaviour to adapt to changing conditions in its environment. One way that neuronal circuits are modified is through the process of synaptic plasticity, in which the strengths of certain synapses are either enhanced or depressed over time in response to neural activity. Insights into when plasticity happens can provide an understanding of the basic functioning of the nervous system, and its ability to learn. A very informative way to gain such insights is through analyzing the molecular circuitry of the synapses - i.e. the networks of biochemical reactions that underlie synaptic modifications. These differ across synapses, and our focus in this study is on the corticostriatal synapse, which is the interface between the cortex and the basal ganglia, a forebrain structure involved in selection of behaviour and reward learning [1,2].

All cells process information from their external and internal environment through signal transduction networks - molecular circuits evolved for producing suitable responses to different stimuli. In neurons, synaptic signal transduction networks determine whether a synapse will be potentiated or depressed. In some cases, even single molecules are able to realize computational abilities within these networks. These molecules are often enzymes, whose activity is allosterically regulated by the binding of other signaling molecules [3]. One such case is the family of mammalian adenylyl cyclase enzymes (ACs). These catalyze the conversion of adenosine triphosphate (ATP) to cyclic adenosine monophosphate (cAMP) - one of the main cellular second messenger signaling molecules.

Mammalian ACs express ten different isoforms [4–6]. Of these, nine are membrane bound, and one is soluble. Their catalytic reaction may be regulated by a variety of interactors, most importantly G protein subunits [6]. These are released in response to extracellular agonists binding to G protein-coupled receptors (GPCRs), the largest superfamily of mammalian

transmembrane receptors. In this way, ACs may function not only as signal transducers but also as *signal integrators*: they perform decision functions that determine the time at which and how much cAMP is produced. One such decision function, attributed to many ACs, is detection of co-occurring signaling events (denoted as *coincidence detection* here), resulting in significantly increased production of cAMP only when more than one signaling event occurs almost simultaneously [7–15].

All nine membrane-bound AC isoforms are expressed in the brain, possibly because this organ is specialized in signal processing. Specific ACs are particularly abundant in specific brain regions, and AC5 is highly expressed in the striatum, the input nucleus of the basal ganglia. It is involved in signal transduction networks that are crucial for synaptic plasticity in the two types of medium spiny neurons (MSNs) of this brain region, which are the direct pathway MSNs that express D₁-type dopamine receptors (D₁ MSNs), and the indirect pathway MSNs expressing D₂-type dopamine receptors (D₂ MSNs).

The same regulatory mechanism exists in both MSN types: AC5 is activated by the stimulatory G α_{olf} protein subunit, and inhibited by the inhibitory G α_i protein subunit. Clearly, the regulation of AC5 is a crucial determinant of the levels of cAMP. This mechanism responds to different extracellular agonists (acting as neuromodulatory signals) associated with the expression of different G protein-coupled receptors (GPCRs) (Fig 1A) [16–18]. In D₁ MSNs, the binding of dopamine to the D₁ GPCR results in the release of G α_{olf} , while binding of acetylcholine to the M₄ GPCR causes the release of G α_i . Conversely, in D₂ MSNs, binding of dopamine to the D₂ GPCRs results in the release of G α_i , and G α_{olf} is released upon binding of adenosine to the A_{2a} GPCR.

Knowing how AC5 processes the neuromodulatory signals would reveal the conditions under which plasticity and learning in the basal ganglia are triggered. In particular, knowing

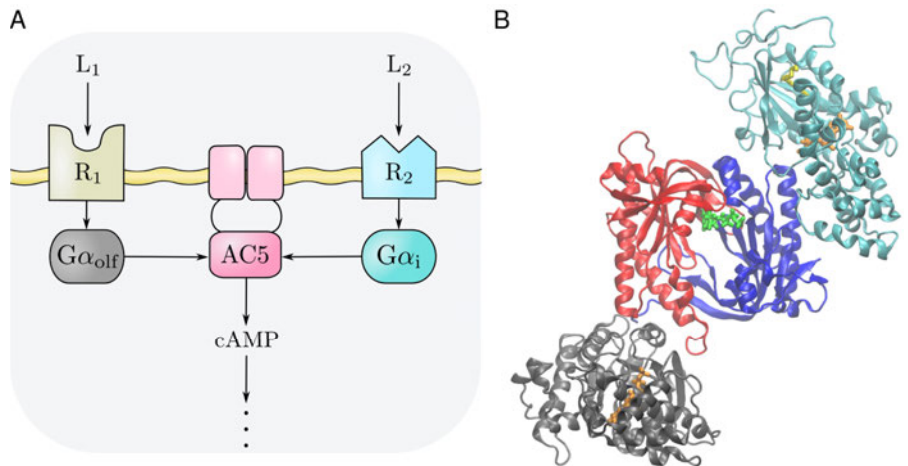


Fig 1. A General scheme of the AC5 signal transduction network. It applies to both the D₁ and D₂ MSNs discussed in the text. Two agonists (L₁ and L₂) bind to two GPCRs (R₁ and R₂), releasing the G α_{olf} and G α_i subunits, respectively. These stimulate and inhibit the conversion of cAMP. **B** Initial modelled configuration of the G α_{olf} · AC5 · G α_i ternary complex, used in the classical MD simulations. The cytosolic part of the AC5 enzyme consisting of the pseudo-symmetric C1 (blue) and C2 (red) cytoplasmic domains in an ATP-bound (green) conformation are complexed to an active conformation of G α_{olf} (gray) and to G α_i (cyan). GTP (orange) and the myristoyl moiety in G α_i (yellow) are shown in stick representation. Controlling the relative positions and conformations of C1 and C2 may enhance or inhibit enzymatic function. This is one way in which G α_{olf} and G α_i exert their regulatory effects: each of them has a separate binding site on the AC5 domain dimer.

<https://doi.org/10.1371/journal.pcbi.1007382.g001>

whether AC5 is a coincidence detector will help us understand if and why changes in more than one of the neuromodulatory signals it receives are necessary to trigger synaptic plasticity. This is what we have set out to determine in this study. We use the neuromodulation of D₁ MSNs as the example here (see [15]).

Recent quantitative estimates of the efficacy of ACh on M₄ receptors in D₁ MSNs suggest that the basal ACh level (even a few tens of nanomolars), resulting from the tonic activity of cholinergic interneurons, may produce G α_i activation [14]. The active G α_i could be sufficient to tonically inhibit significant amounts of AC5 under baseline conditions. In addition, it has been proposed that the dopamine-dependent D₁ activation of AC5 by G α_{olf} is quite low at baseline conditions [19]. This partly inhibited state of AC5 could be considered as the “resting state” of striatal AC5 signaling. Previous modeling studies of this signal transduction network predicted that AC5 responds most strongly to a simultaneous increase in dopamine (Da \uparrow) and a pause in acetylcholine levels (ACh \downarrow), i.e. both stimulation by increased G α_{olf} and disinhibition by decreased G α_i are necessary for the enzyme to produce significant amounts of cAMP [15]. The response to the two neuromodulatory signals was nonlinear and synergistic. This suggests that AC5 might function as a coincidence detector, since the network responds with significant amounts of cAMP only when the two incoming signals Da \uparrow and ACh \downarrow coincide in time and in the spatial vicinity of the receptors. In order to perform strong coincidence detection, the network should be able to make a clear distinction between the situation of a simultaneous dopamine peak and acetylcholine dip (Da \uparrow + ACh \downarrow) and that of a single signal, i.e. Da \uparrow or ACh \downarrow alone. This distinction is realized by producing different amounts of cAMP, i.e. by differences in the enzyme’s catalytic rate.

During the course of our previous kinetic modelling study [15], and follow-up experimental studies on the function of the AC5 signal transduction network [14,19], it became clear that the presence or absence of a ternary G α_{olf} · AC5 · G α_i complex during AC5 regulation, and the level of catalytic activity of such a complex, could significantly affect AC5’s ability to act as a coincidence detector.

While the existence of the binary AC5 · G α_i and AC5 · G α_{olf} complexes, and their catalytic activities, has been confirmed experimentally, a G α_{olf} · AC5 · G α_i ternary complex has not been identified so far [20–22]. However, it has been suggested to exist during AC5 regulation, but it is not known whether it would be catalytically active or inactive (Fig 1B) [21, 23]. So far, we know from molecular dynamics (MD) simulations of the binary complexes that binding of one G α subunit can produce allosteric effects at the binding site for the other [24,25], raising the speculation that allosteric effects influence ternary complex formation.



Resolving the details of AC5 regulation can help to understand whether AC5 acts as a coincidence detector, and hence, whether transient changes in two neuromodulatory signals are necessary for plasticity and learning in the basal ganglia. Here, we take a multiscale modeling approach to address the following question: can the G α_{olf} · AC5 · G α_i ternary complex form in the AC5 signal transduction network? If it does, is it able to catalyse ATP conversion, and does its presence affect the ability of the enzyme to perform coincidence detection? We combine MD simulations, to study the complexation of AC5 with its G α partners, with Brownian dynamics (BD) simulations to estimate the forward (association) rate constants of binding between AC5 and the G α subunits. Snapshots from the MD simulations are used as starting points for BD simulations. We then incorporate the results from these simulations into a kinetic model of the AC5 signal transduction network and quantify the ability of the enzyme to detect coincident extracellular signaling events. Molecular simulation approaches that bridge multiple spatial and temporal scales are well established [26], and here we span the spatial scales by bridging from intra- and inter-molecular dynamics up to the function of a biochemical reaction network [27,28,29]. Multiscale simulation is particularly informative for this system since experimental efforts in studying the ternary complex have not yielded results.

Our MD simulations show that a ternary complex could be present during AC5 regulation, being stable on the μ s timescale, but it appears to be catalytically inactive. The kinetic model, constructed using results from the MD and BD simulations performed here, reveals that: (i) the suggested interaction scheme between the regulatory $G\alpha$ subunits and AC5 strengthens coincidence detection when compared to alternative schemes; (ii) the predicted values of the forward rate constants are favourable for coincidence detection. This suggests that AC5 is indeed a powerful coincidence detector and, as a result of the inactive ternary complex, a rise in dopamine alone does not have an effect on synaptic plasticity; it needs to be accompanied by a pause in the level of acetylcholine. These insights are also discussed with regard to other AC isoforms.

Results

Molecular dynamics simulations indicate that the apo $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ ternary complex is stable on microsecond timescales

The existence of a ternary complex, $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$, during AC5 regulation has been unclear. So far it has not been possible to detect the complex in experiments, possibly due to its unstable nature [21]. However, as $G\alpha_{\text{olf}}$ and $G\alpha_i$ interact with AC5 on different sites on its catalytic domains, the ternary complex could be formed via two reactions:



Previous MD simulation results suggested that upon binding of $G\alpha_i$ to AC5, the $G\alpha_{\text{olf}}$ binding groove adopts a conformation that hinders $G\alpha_{\text{olf}}$ from binding [24], making the first reaction less favourable; however, a stable ternary complex might still be formed via the second route. All-atom MD simulations were employed to investigate the stability of a putative *apo* ternary complex, $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$, in the absence of ATP, on the μ s time scale. We mention here that $G\alpha_{\text{olf}}$ has no post-translational modifications to its protein sequence. In contrast, $G\alpha_i$ is considered in its myristoylated form since a non-myristoylated $G\alpha_i$ subunit is known to be unable to form an $AC5 \cdot G\alpha_i$ complex and, therefore, it is not functional [21, 24].

The *apo* ternary complex appears to be stable over the course of 2.1 μ s of all-atom MD simulation with a root-mean-square deviation (RMSD) of the complex's backbone fluctuating between 0.8 and 1 nm (Fig 2B) compared to the first frame of the trajectory. The RMSD of each individual protein, i.e. AC5, $G\alpha_{\text{olf}}$ or $G\alpha_i$, in the *apo* ternary complex remains below 0.4 nm (S1 Fig). Similar RMSD values have been found for the *apo* forms of the binary complexes $AC5 \cdot G\alpha_{\text{olf}}$ and $AC5 \cdot G\alpha_i$ (S1 Fig). Additionally, analysis of the secondary structures and calculations of the numbers of H-bonds as functions of time for all systems confirm the overall stability suggested by the RMSD analysis (S2–S7 Figs). Root-mean-square-fluctuations (RMSF) per residue have also been calculated (S6 Fig).

Minor differences in RMSF were observed for $G\alpha_{\text{olf}}$ and $G\alpha_i$, independently of the simulated system. Conversely, in the case of AC5, the C2 $\beta 4'$ - $\beta 5'$ region (i.e. residues Val1186 to Trp1200 of the C2 domain) and the C2 $\beta 7'$ - $\beta 8'$ region (i.e. residues Gln1235 to Asn1256 of the C2 domain) show a change in orientation dependent on the presence of the $G\alpha$ subunit. Whereas $\beta 7'$ - $\beta 8'$ appears to be more flexible in the binary $AC5 \cdot G\alpha_{\text{olf}}$ complex, the $\beta 4'$ - $\beta 5'$ region is more flexible in the presence of $G\alpha_i$. Besides RMSD, a measure of compactness of the simulated complexes is provided by the global radius of gyration (R_g), which also provides an indication of the stability of the complex. R_g was calculated as a function of time along the simulated trajectories for all three complexes (S7 Fig). In the case of the *apo* $AC5 \cdot G\alpha_i$ complex, a

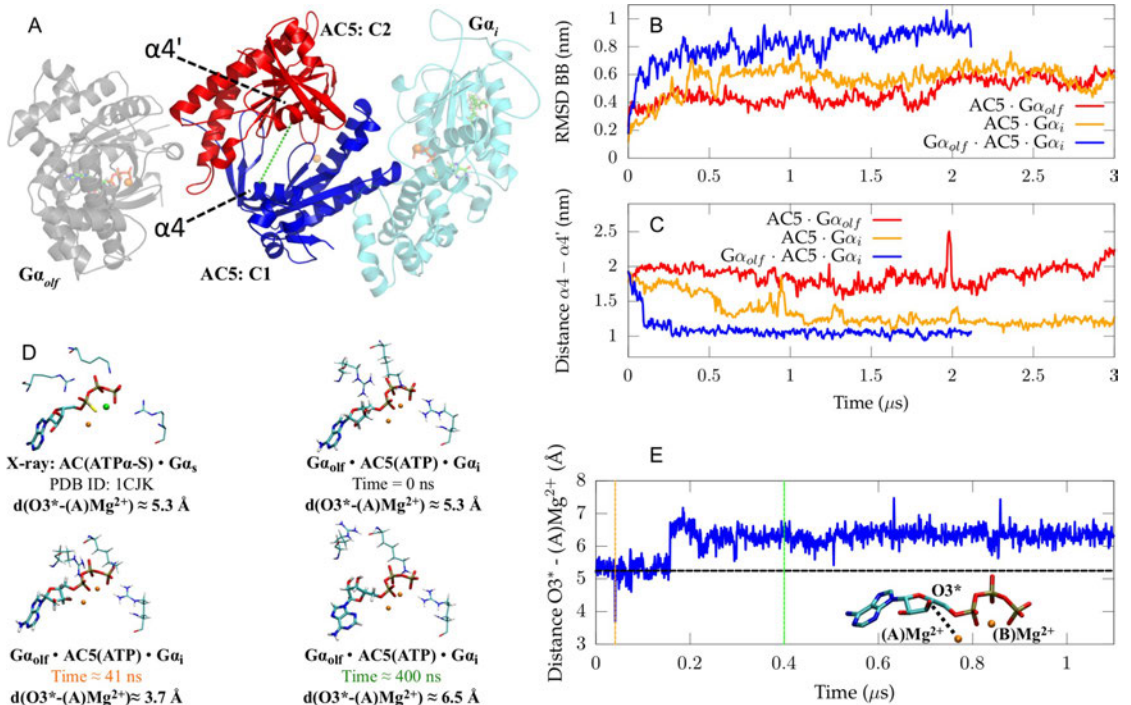


Fig 2. Stability analysis of the *apo* ternary complex, conformation of the active site of AC5 in the *apo* state, and conformation of ATP in AC5's active site in the *holo* state. A Initial conformation of the *apo* ternary complex including $G\alpha_{olf}$ (gray), $G\alpha_{i1}$ (cyan), C1 (blue), C2 (red). In addition to the protein structures, two GTP molecules as well as the myristoyl moiety of the $G\alpha_{i1}$ subunit and three Mg^{2+} ions are shown. The green dashed line indicates the distance between the C α atom of Thr1007 (belonging to helix $\alpha 4$) and the C α atom of Ser1208 (belonging to helix $\alpha 4'$). B RMSD of the backbone of the protein complexes $G\alpha_{olf} \cdot AC5 \cdot G\alpha_{i1}$, $AC5 \cdot G\alpha_{i1}$, and $G\alpha_{olf} \cdot AC5 \cdot G\alpha_{i1}$. C Time evolution of the distance between Thr1007 and Ser1208 for the three simulated complexes in *apo* form. D Conformation of ATP in the $G\alpha_{olf} \cdot AC5(ATP) \cdot G\alpha_{i1}$ system at different time intervals in the trajectory as well as the conformation of ATP- α -S in the AC · $G\alpha_i$ X-ray structure (PDB code 1CJK). The color of the time in green and orange corresponds to the coloured lines in image (E). E Distance between the oxygen of a hydroxyl on the sugar ring of ATP, $O3^*$, and a Mg^{2+} ion, (A) Mg^{2+} , in the active site of the *holo* ternary complex. The black dashed line shows the distance between the two atoms in the AC · $G\alpha_i$ X-ray structure (PDB ID 1CJK) to which ATP- α -S is bound [22].

<https://doi.org/10.1371/journal.pcbi.1007382.g002>

clear reduction of R_g can be observed compared to its initial conformation. The *apo* form of $AC5 \cdot G\alpha_{olf}$ also shows a reduction of R_g over time, while R_g increases for the *apo* ternary complex system with respect to the starting value. In this regard, it is worth pointing out that after the first 500 ns, R_g does not increase further and, instead, oscillates around a constant value, slightly higher than that for the first frame. Such behaviour does not suggest instability of the ternary complex, but is rather an indication of an initial reorientation of the different domains to optimize their interaction. Such structural rearrangement is not surprising considering that the initial structure of the ternary complex was generated by homology modelling.

Simulations of the *apo* and *holo* state of the ternary complex indicate that it is unable to catalyse ATP conversion

To assess the ability of the ternary complex to catalyse the conversion of ATP to cAMP, we investigated the structural dynamics of the ATP binding site during the MD simulation of the

apo ternary complex, and the conformations sampled by ATP in a simulation of the *holo* ternary complex, $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ in which ATP is bound to the active site of AC5. We compare the conformational changes of the proteins in the *apo* ternary complex with those in the previously reported simulation of the *apo* $AC5 \cdot G\alpha_i$ complex [24] and with those of the *apo* $AC5 \cdot G\alpha_{\text{olf}}$ simulation reported here.

The ability of AC5 to convert ATP to cAMP depends on the state of its catalytic domain. A characteristic quantity in this respect is the relative distance between the two helices $\alpha 4'$ and $\alpha 4$ positioned either side of the binding groove (Fig 2A). The distance between the C α atom of Thr1007 in helix $\alpha 4$ and the C α atom of Ser1208 in helix $\alpha 4'$ (highlighted by the green dashed line in Fig 2A) was calculated along the simulated trajectories for the three complexes investigated here and reported as a function of time in Fig 2C.

This distance exhibits a clear decrease in the first 100 ns in the $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ complex, starting from a value of 1.9 nm and reaching a stable value of about 1.1 nm. Similarly, a decrease of the $\alpha 4$ - $\alpha 4'$ distance was also observed in the $AC5 \cdot G\alpha_i$ complex along 3 μ s of simulation. Conversely, when AC5 is bound to $G\alpha_{\text{olf}}$ in a binary complex, the distance between the two helices is characterized by a higher value with respect to the case of the $AC5 \cdot G\alpha_i$ and $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ complexes. A larger value of this distance can be associated with a higher accessibility of the binding groove. On the other hand, a reduced value, as found in the AC5 domain bound to $G\alpha_i$ implies a lower accessibility to the binding groove. The effect of the reduction in $\alpha 4$ - $\alpha 4'$ distance on active site conformation and residue orientation has been described in [24] for the $AC5 \cdot G\alpha_i$ complex. A shorter distance between the $\alpha 4$ and $\alpha 4'$ helices decreases the volume in the active site accessible to ATP as $\alpha 4'$ starts occupying the region in which ATP's adenine moiety docks. Due to the conformational changes in the active site, important residues for ATP stabilization undergo a rearrangement in orientation that appear to significantly differ from the $AC5 \cdot G\alpha_{\text{olf}}$ system, potentially negatively impacting the probability of ATP association. Hence, this partial closure of the active site of AC5 in $AC5 \cdot G\alpha_i$ and $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ reduces the space available for ATP binding and could also lower the probability of stable ATP association, suggesting that the ternary complex is likely to be inactive.

Apart from the ability of the *apo* ternary complex to bind the substrate, we additionally investigated the possible catalytic activity of a *holo* $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$ complex. ATP has to undergo a cyclisation reaction in order to form the products cAMP and pyrophosphate. This reaction is induced by the attack of a deprotonated hydroxyl moiety in ATP's sugar ring, $O3^*$, on the phosphorus atom of the α -phosphate moiety (Fig 2E). Hence, ATP conversion requires oxygen $O3^*$ to be in the proximity of (A)Mg $^{2+}$ to which the phosphorus atom is coordinated in order to obtain a conformation that can undergo the cyclisation reaction [30–32].

A crystal structure of the *holo* AC domain dimer in complex with a stimulatory $G\alpha$ subunit [22] shows the active conformation of ATP α -S, an ATP mimic, in the active site of the enzyme in the presence of a Mg $^{2+}$ ion and a Mn $^{2+}$ ion (Fig 2D), which is assumed to be substituted by a second Mg $^{2+}$ ion under physiological conditions. The $O3^*-(A)Mg^{2+}$ distance in the X-ray structure of the *holo* $AC \cdot G\alpha_s$ complex is 5.25 Å (Fig 2D). In the MD trajectory of the *holo* ternary complex, the $O3^*-(A)Mg^{2+}$ distance starts at a similar value and mainly remains in this state for the first 160 ns. Within the first 160 ns, the $O3^*-(A)Mg^{2+}$ distance even decreases to distances as short as 3.7 Å, closer to the distance which has been suggested to correspond to ATP's reactive state [30–32] (see orange line). However, after 160 ns, the ATP molecule undergoes a drastic conformational change, resulting in an increase of the $O3^*-(A)Mg^{2+}$ distance to \approx 6.5 Å. In this state, the ATP molecule is unable to attain a short $O3^*-(A)Mg^{2+}$ distance, such as observed at 41 ns, which is a feature of the reactive state of ATP. This conformational transition of ATP appears to be irreversible on the μ s time scale, thus suggesting the inhibition of the catalytic reaction.

The rate of diffusional association of each $G\alpha$ protein subunit to AC5 is unaffected by prior binding of the other subunit

To provide initial values for the forward rate constants - the parameters in the kinetic model of AC5 activity which were not constrained experimentally - we performed BD simulations (Table 1). The predicted rate constants suggest both $G\alpha$ subunits form complexes at similar rates, and their association is not greatly affected by the presence of ATP in the AC5 active site, or prior binding of the other subunit. Indeed, the variation in the predicted rate constants for each reaction across the different MD snapshots used in BD simulations is greater than the variation of the mean values for each constant (S1 Table).

It should be noted that the predicted rate constants are for the diffusional approach and initial binding of the $G\alpha$ subunits to AC5. The previously reported MD simulations show that the binding of $G\alpha_{olf}$ to AC5 · $G\alpha_i$ is hindered by a conformational change of its binding groove on AC5, adding an additional conformational gating contribution to the rate of binding that is not described by the BD simulation method used [24].

The presence of an inactive ternary complex improves the ability of the network to detect coincident signals

We incorporated the results from the MD and BD simulations into a kinetic model of the AC5 signal transduction network, the basic feature of which is a regulatory scheme where the ternary complex can form (Fig 3C). We find that this network can perform coincidence detection. To investigate how the ternary complex contributes to the network's ability to perform coincidence detection, we compared the system with a network in which no ternary complex can form.

It is important to note that there are two aspects of coincidence detection: (a) distinguishing between the different inputs, and (b) responding strongly enough with an increase in cAMP concentration that is physiologically relevant. As a proxy for the amount of cAMP produced, we use the average catalytic rate of AC5, k_c , since the amount of cAMP produced is proportional to k_c (see Methods, and also illustrated in Fig 3). The average catalytic rate of AC5 is a weighted average of the catalytic rates of the unbound enzyme and each of the complexes with the $G\alpha$ subunits. Additionally, to measure whether the signal transduction network distinguishes between the different input situations, the synergy quantity, S , is employed. This describes how much greater the average catalytic rate is when both signals coincide compared to the cases when only one of them arrives (see Methods). A synergy value greater than 1 indicates that the network can perform coincidence detection, i.e. the network responds more strongly when the two signals coincide than if each of them occurs individually and are summed. A value of 1, or less than 1, marks a response equal to or weaker than the sum of the individual signals, respectively. In these cases, the average catalytic rate does not produce distinguishable amounts of cAMP that enable the cell to discriminate between the different input

Table 1. Bimolecular association rate constants (nM⁻¹s⁻¹) for the forward reactions computed from BD simulations (standard deviation over all structures of complexes and replica simulations shown in parentheses). Rate constants were calculated separately for complexes incorporating both apo and holo AC5. Data for each individual structure of each complex are shown in S1 Table.

Rate Constant	apo AC5	holo AC5
k_{F1}	0.018 (0.007)	0.026 (0.009)
k_{F2}	0.02 (0.01)	0.05 (0.01)
k_{F3}	0.012 (0.008)	0.017 (0.006)
k_{F4}	0.02 (0.01)	0.020 (0.006)

<https://doi.org/10.1371/journal.pcbi.1007382.t001>

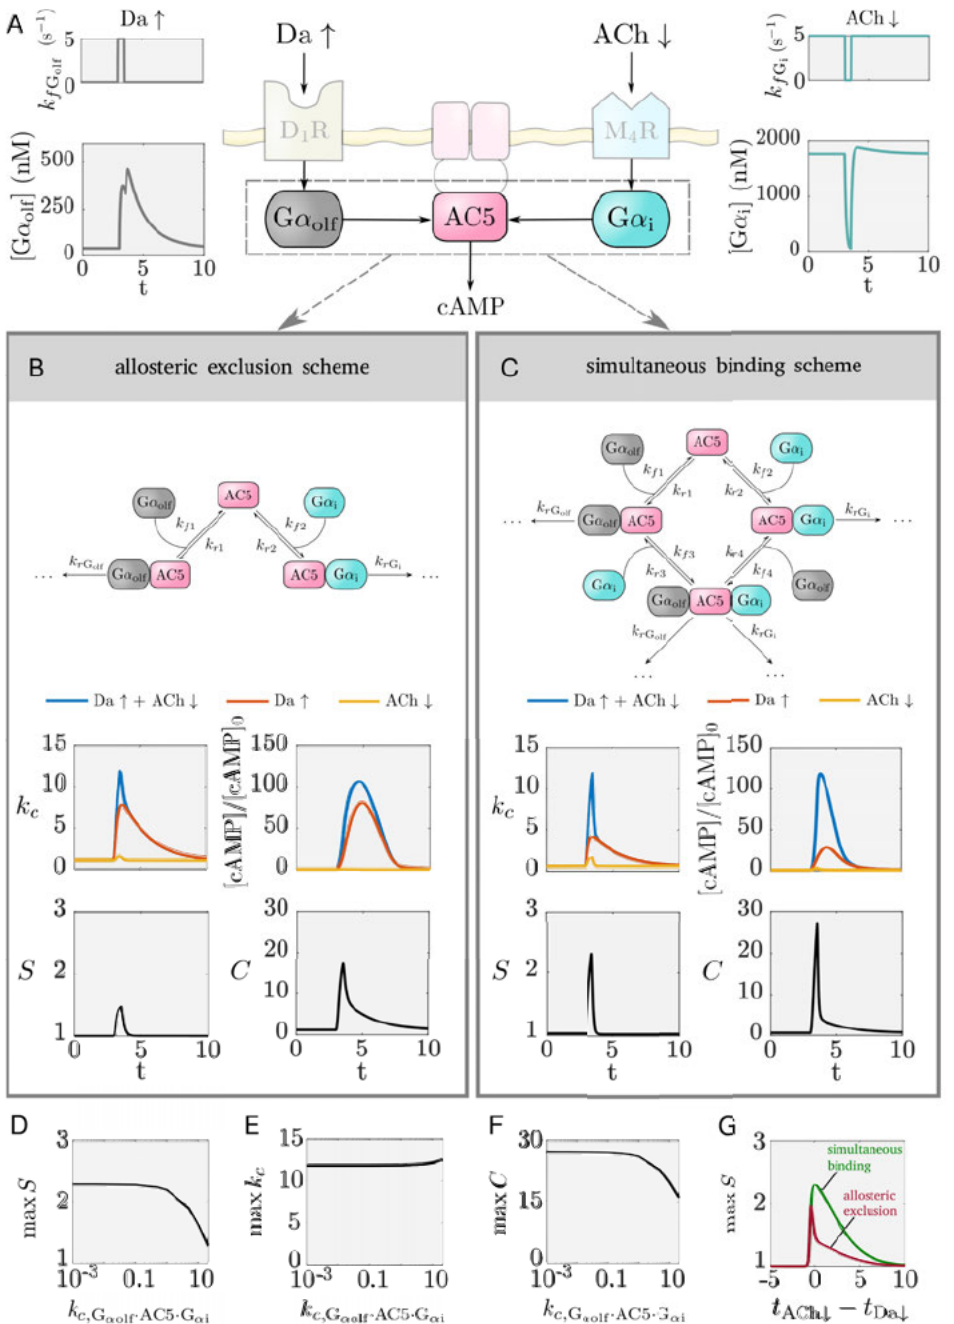


Fig 3. The effect of the different regulatory schemes on coincidence detection. **A** The inputs to the model translate to an elevation in $G\alpha_{olf}$ and pause in $G\alpha_i$. The shaded parts of the scheme are not included in the kinetic model. **B** The allosteric exclusion scheme, and the k_c , synergy, and cAMP levels obtained due to this regulatory scheme. **C** The simultaneous binding scheme, and the k_c , synergy, and cAMP levels obtained due to this regulatory scheme. **D,E,F** The effect of the ternary complex' catalytic activity on coincidence detection: the maximum of the synergy (D), the maximum of k_c (E), the maximum of the metric C (F). **G** The detection window for the allosteric exclusion scheme (red) and the simultaneous binding scheme (green). $t_{Da\uparrow}$ and $t_{ACh\downarrow}$ are the times when the $Da\uparrow$ and $ACh\downarrow$ occur, respectively. Note the shared time axes in (A), (B) and (C).

<https://doi.org/10.1371/journal.pcbi.1007382.g003>

situations. Since a synergy greater than 1 does not necessarily indicate a strong response from the network (i.e. strong cAMP production), a combination of the average catalytic rate and the synergy is applied where needed to quantify coincidence detection (see the metric C in methods).

The regulatory schemes that we compare are given in Fig 3B and 3C. We name the first scheme an *allosteric exclusion scheme* – the binding of one $G\alpha$ subunit excludes the possibility for binding of the other $G\alpha$ subunit. The second scheme is termed the *simultaneous binding scheme*, where both $G\alpha$ subunits can bind to AC5 and the ternary complex can be formed. Simulation results for both schemes are given in Fig 3. The inputs are assumed to be a dopamine peak of 0.5s ($Da\uparrow$) and an acetylcholine dip of 0.5s ($ACh\downarrow$), and the corresponding rise in $G\alpha_{olf}$ and drop in $G\alpha_i$ are shown in Fig 3A. Time courses showing the amounts of all AC5 species (the free enzyme and the complexes with the $G\alpha$) are given in S9 Fig.

The simultaneous binding scheme can better distinguish between $Da\uparrow + ACh\downarrow$ and the individual signals $Da\uparrow$ or $ACh\downarrow$, it has higher synergy. Both the schemes have a similar maximal $k_c(Da\uparrow + ACh\downarrow)$, and, as evident, the increase in synergy in the simultaneous binding scheme versus the allosteric exclusion scheme comes from a reduced $k_c(Da\uparrow)$. This relative difference in the average catalytic rate enables the simultaneous binding scheme to respond differently to the coincident signal, $Da\uparrow + ACh\downarrow$, compared to $Da\uparrow$ alone, as is also visible from the amounts of cAMP produced, and hence to differentiate between the two input situations. The allosteric exclusion scheme, on the other hand, produces similar values for $k_c(Da\uparrow + ACh\downarrow)$ and $k_c(Da\uparrow)$ and thus responds similarly to $Da\uparrow + ACh\downarrow$ and to $Da\uparrow$ alone, being unable to distinguish well between them in terms of cAMP production. The reason for this is the exclusivity of the interaction between each of the $G\alpha$ proteins and AC5: when only $Da\uparrow$ arrives, $G\alpha_{olf}$ is able to compete with $G\alpha_i$ and bind to much of the enzyme (approximately half of it as shown in S9B Fig). This creates the catalytically active complex $AC5 \cdot G\alpha_{olf}$, driving an increase in $k_c(Da\uparrow)$. In this case, reduced inhibition of AC5 by an additional $ACh\downarrow$ does not contribute much to $k_c(Da\uparrow + ACh\downarrow)$. In the simultaneous binding scheme, however, a $Da\uparrow$ causes the formation of the ternary complex S9F Fig), and due to its inactivity $k_c(Da\uparrow)$ does not increase significantly. Only with an additional $ACh\downarrow$ is the inhibition by $G\alpha_i$ relieved and the proportion of the active complex $AC5 \cdot G\alpha_{olf}$ is increased, enabling a high $k_c(Da\uparrow + ACh\downarrow)$. Importantly, $k_c(Da\uparrow)$ is also low for an inactive ternary complex so that little cAMP is produced with a $Da\uparrow$ only and little "stray" activation of downstream signalling would occur. In fact, only for a substantially active ternary complex does the simultaneous binding scheme become comparable to the allosteric exclusion scheme in terms of synergy (Fig 3D, 3E and 3F). For a wide range of low to moderate ternary complex activity, it performs coincidence detection better. The maximum of the catalytic rate is not affected much by the activity of the ternary complex (Fig 3E), and the metric C shows that coincidence detection is most significant for an inactive ternary complex (Fig 3F). An inactive ternary complex enables the lowest catalytic activity of the enzyme at resting state and hence the biggest difference between $k_c(Da\uparrow)$ and $k_c(Da\uparrow + ACh\downarrow)$, and this in turn maximizes the synergy.

In the supplementary material we show that the allosteric exclusion scheme in itself lacks the ability to perform coincidence detection, and this is due to the exclusivity of the regulatory

interaction. Coincidence detection with this scheme, as demonstrated in Fig 3B, is in fact a result of the amounts of $G\alpha_{olf}$ and $G\alpha_i$ and the kinetics determined by the forward rate constants.

An inherent property of a coincidence detector is that there exists a time window over which two signals can be detected as if arriving together. The detector uses some mechanism by which it “remembers” the occurrence of one of the signals for some time interval, and responds when the other signal arrives within this interval. For the AC5 signal transduction network, the existence of the detection window also depends on the regulatory scheme. In fact, the formation of the ternary complex is very important to allow for a broader window of coincidence detection. In the case of only a $Da\uparrow$, a ternary complex that has buffered (or absorbed) the elevated active $G\alpha_{olf}$ provides this memory: allowing the $ACh\downarrow$ to arrive some time later and still elicit a response (Fig 3G). This is potentially relevant, since the cholinergic interneurons responsible for the $ACh\downarrow$ have been found to produce a second $ACh\downarrow$ to certain stimuli [33]. The length of the detection window is determined by the rate of deactivation of the active $G\alpha_{olf}$ (also illustrated in S10D and S10G Fig, and is due to the fact that the GTPase activity for $G\alpha_{olf}$ is lower than the one for $G\alpha_i$). Both schemes technically have the same length of detection windows, but the allosteric exclusion scheme has a high synergistic effect in a very narrow region of the window - the signals need to occur practically simultaneously (Fig 3G). The detection window is asymmetric, i.e. the $Da\uparrow$ needs to arrive first to elicit a response from the network. (Time courses with $ACh\downarrow$ preceding and following a $Da\uparrow$ illustrating the difference between the two schemes are given in S10 Fig.)

Lastly, we note that a critical aspect for coincidence detection to work is to have fast deactivation of the active $G\alpha_i$. Then the dynamics of $G\alpha_i$ inside the cell can follow the short duration of the $ACh\downarrow$ signal. The experimental evidence for this high GTPase rate is listed in the description of the kinetic model (see Methods). The effect of the GTPase rate on coincidence detection is shown in S11 Fig. There is an optimum value of this rate - it needs to be both high enough to cause a drop in $[G\alpha_i]$ during the $ACh\downarrow$ and low enough so that there is significant inhibition of AC5 at steady state. Since the deactivation of $G\alpha_i$ is faster than that for $G\alpha_{olf}$, then, provided there is enough active $G\alpha_{olf}$ to bind to AC5, the duration of the synergistic effect is determined by the duration of the pause (S10E and S10G Fig). More extensive analysis on the robustness of coincidence detection to parameter values has been performed for the previous, much larger kinetic model of the signal transduction network [34].

To summarize, the results of this section show that the formation of the ternary complex aids coincidence detection and prolongs the detection window. Additionally, the less catalytically active the ternary complex is, the better the coincidence detection. As elaborated in the discussion, an inactive ternary complex can also explain recent experimental results on cAMP production due to activation of the implicated GPCRs in the native system [14,19].

Hindrance of $G\alpha_{olf}$ binding to AC5 · $G\alpha_i$ still allows coincidence detection

To see where the predicted values from the BD simulations lie in terms of affecting coincidence detection, we performed parameter scans for the values of the forward rate constants. All of the forward rate constants affect the synergy, since they affect the fractions q_i , $i = 1, \dots, 4$, of each species at any point in time (see equations for q_i in Methods).

To begin with, we address the observation from Van Keulen and R othlisberger [24] that the $G\alpha_{olf}$ binding groove in AC5 · $G\alpha_i$ adopts a conformation that hinders $G\alpha_{olf}$ binding. This could further decrease the value of the rate constant k_{r4} , compared to the value predicted in the BD simulations. We investigated the effect such a decrease has on coincidence detection by altering k_{r4} by varying orders of magnitude (Fig 4A, 4B and 4C). We found that hindered

binding of $G\alpha_{\text{off}}$ to $AC5 \cdot G\alpha_i$ does not affect coincidence detection significantly (the synergy S is increased by 12%). Nevertheless, when performing the parameter scans, we considered two scenarios. In the first one, we used $k_{f1} = k_{f4}$ and $k_{f2} = k_{f3}$ since the BD simulations showed that these values are similar, at least in order of magnitude. In the second scenario, we used $k_{f2} = k_{f3}$ and $k_{f4} = k_{f1}/100$. We call this scheme, in which the reaction corresponding to the rate k_{f4} is slower, the *hindered simultaneous binding scheme*.

For the simultaneous binding scheme, a very wide range of tested values for k_{f1} provides similar k_c and synergy values for a given k_{f2} (Fig 4D and 4E), which can be interpreted as follows. Similarly to the requirement for a high GTPase activity for $G\alpha_i$, it is necessary that $G\alpha_{\text{off}}$ binds to AC5 quickly enough to be able to follow the signal $Da \uparrow$. Since active $G\alpha_{\text{off}}$ has a slower dynamics than the input $Da \uparrow$, i.e. $G\alpha_{\text{off}}$ exists inside the cell for some time after the outside signal has stopped, it is likely that a range of values for the rate constant k_{f1} can enable coincidence detection (not only values that match the length of a $Da \uparrow$, which are around $k_{f1} = 2 \text{ (nMs)}^{-1}$ for a $Da \uparrow$ of 0.5s). The synergy grows with k_{f1} since a higher binding rate of $G\alpha_{\text{off}}$ provides a higher maximum of k_c during the $Da \uparrow + ACh \downarrow$, whereas $k_c(Da \uparrow)$ is not affected much (S12A and S12B Fig).

Increasing the rate constants $k_{f2} = k_{f3}$ for $G\alpha_i$ binding also causes an increase in synergy. The lower these constants, the less inhibited the enzyme will be due to smaller fractions of both $AC5 \cdot G\alpha_i$ and $G\alpha_{\text{off}} \cdot AC5 \cdot G\alpha_i$. This allows for more stimulation by the available $G\alpha_{\text{off}}$ and hence a higher resting-level k_c and a higher $k_c(Da \uparrow)$. The situations $Da \uparrow + ACh \downarrow$ and only $Da \uparrow$ produce a more similar response, and hence show a lower synergy. Higher values for $k_{f2} = k_{f3}$, on the other hand, allow for both stronger resting-level inhibition of AC5 and more ternary complex formed, and hence a lower resting-level k_c and lower $k_c(Da \uparrow)$, thus increasing the synergy. The region of optimal parameters for the simultaneous binding scheme is quite wide in terms of $k_{f1} = k_{f4}$ and $k_{f2} = k_{f3}$, and the most optimal scenario is for the largest values tested, as shown in Fig 4F using the metric C .

Compared to the simultaneous binding scheme, the hindered simultaneous binding scheme does not affect the maximum of the average catalytic rate significantly (Fig 4H), but it moves the region of optimal values for the synergy towards low values for the rate constant k_{f1} and increases it for these values (Fig 4G). Inspecting the dynamics of the model with this regulatory scheme reveals the reason: the hindered reaction $AC5 \cdot G\alpha_i + G\alpha_{\text{off}} \leftrightarrow G\alpha_{\text{off}} \cdot AC5 \cdot G\alpha_i$ causes less ternary complex to be formed and consequently less $AC5 \cdot G\alpha_{\text{off}}$ from the dissociation of $G\alpha_i$ from the ternary complex (the route via k_{f3}) (S12C and S12D Fig). This, in general, lowers $k_c(Da \uparrow)$ compared to the simultaneous binding scheme, creating a larger relative difference between $k_c(Da \uparrow)$ and $k_c(Da \uparrow + ACh \downarrow)$ which results in a high synergy. This difference is largest for low values of k_{f1} and decreases with k_{f1} due to the higher amounts of $AC5 \cdot G\alpha_{\text{off}}$ during the $Da \uparrow$ (higher $k_c(Da \uparrow)$). Importantly, the maximal k_c has the opposite dependence on k_{f1} from the synergy, so that the region of values for k_{f1} optimal for coincidence detection optimizes between a high catalytic rate and a high synergy (Fig 4I). The effect of $k_{f2} = k_{f3}$ on coincidence detection is the same as described for the simultaneous binding scheme.

In both scenarios, the predictions of the BD simulations are favourable for coincidence detection. However, higher binding rates of $G\alpha_i$ than the predicted ones are better suited for coincidence detection. This increase in k_{f2} and k_{f3} could arise if $G\alpha_i$ were part of multiprotein signaling complexes, as has been shown for AC5 and $G\alpha_{\text{off}}$ [35–37].

Discussion

In this study we find that an inactive ternary complex between AC5 and its G protein regulators, a molecular-level feature, gives rise to significantly increased coincidence detection, a

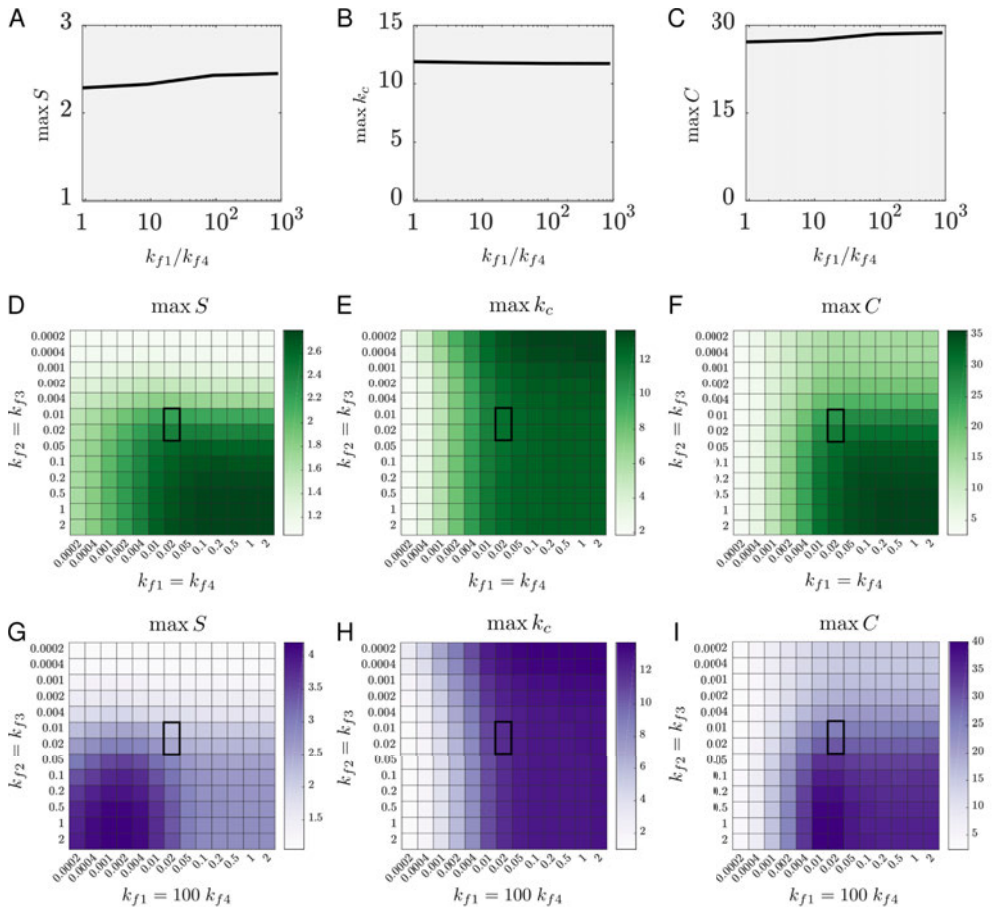


Fig 4. Dependence of coincidence detection on the forward rate constants. A, B, C The maximum of the synergy (A), the maximum of k_c (B) and the maximum of the metric C (C) for a reduced rate constant k_{fi} relative to the prediction from BD simulations. D, E, F Coincidence detection for the simultaneous binding scheme as dependent on the forward rate constants: maximum of the synergy (D), maximum of k_c (E), maximum of the metric C (F). G, H, I Coincidence detection for the hindered simultaneous binding scheme as dependent on the forward rate constants: maximum of the synergy (G), maximum of k_c (H), maximum of the metric C (I). Axes are in units of $(nMs)^{-1}$. Highlighted regions correspond to the forward rate constants predicted from BD simulations.

<https://doi.org/10.1371/journal.pcbi.1007382.g004>

systems-level function of the signal transduction network that AC5 is embedded in. In order to investigate the stability and the activity of the putative $G\alpha_{olf} \cdot AC5 \cdot G\alpha_i$ ternary complex, we carried out all-atom MD simulations. Our results showed that on the μs time scale, the complex seems to be stable independently of the presence or absence of ATP. Additionally, previous MD studies suggest possible pathways for the formation of the ternary complex, showing the possibility of the $G\alpha_i$ protein to bind to the *holo* $AC5 \cdot G\alpha_s$ complex [25], and disfavoring the binding of $G\alpha_{olf}$ to the *apo* $AC5 \cdot G\alpha_s$ complex [24]. Overall, it should be pointed out that MD simulations cannot exclude the instability of the ternary complex on longer time

scales, which cannot be assessed due to computational limits. However, MD simulations can help us to investigate the conformations sampled by the ternary complex at physiological temperature and pressure and, consequently, the activity of the complex. Indeed, the partial closure of the binding groove found in the *apo* ternary complex, analogous to that occurring in the binary *apo* AC5 · G α_i , suggests a lack of catalytic activity in the ternary complex due to the reduced accessibility for the ATP substrate. Additionally, even if a ternary complex could exist with ATP bound to AC5, our results show that the substrate would adopt a conformation not suitable for the subsequent catalytic reaction leading to the formation of cAMP. It is worth mentioning that possible conformations sampled by ATP, prior and during its conversion to cAMP, have been reported by Hahn et al. in a theoretical study [38]. Such conformations, required for an optimal conversion of ATP to cAMP, clearly show an opposite orientation of the oxygen O3^{*} with respect to that sampled in our simulation, thus supporting our hypothesis about the inactivity of the ternary complex.

Experimentally, indirect data exist that are consistent with a significantly inactivated ternary complex. To begin with, the existence of functional oligomeric complexes of AC5 and the A_{2a} and D₂ receptors has been demonstrated, and their proposed spatial arrangement, importantly, supports ternary complex formation [39]. Additionally, experimental results from striatal cultures in this study show that stimulating the D₂ receptor almost entirely counteracts the effects of the applied agonist on the A_{2a} receptor in terms of cAMP production. These results are consistent with an inactive ternary complex since, under these conditions, both active G α_{olf} and active G α_i would exist in the cell, and activation of G α_i stops the activity of the previously formed AC5 · G α_{olf} . Similar experimental studies with the *in vitro* native system have also been performed recently [14,19], although the animals used in these studies are very young (between 7 and 12 days old) and the developmental transition from AC1 to AC5 in the striatum is not complete. However, at P7 AC5 contributes to at least 50% of the cAMP response, and this contribution significantly increases further on, as the transition is reported to be complete by P14 [40]. Comparison with these experimental results is therefore still informative, and in both studies the results are consistent with an inactive ternary complex. In D₁ MSNs, stimulating the D₁ receptors with an agonist followed by activation of the M₄ receptors completely abolishes the cAMP response of AC5 [14, Fig 3A]. The analogous experiment in D₂ MSNs is also consistent with an existing and inactive ternary complex [19, Fig 2B]. Stimulating the A_{2a} receptor with an agonist and then uncaging dopamine, which stimulates the D₂ receptor, also abolishes the cAMP response. Experiments modeling the opposite regulatory effects of G α_s and G α_i on AC5 in living HEK293T cells via two other GPCRs are also consistent with an inactive ternary complex, where activation of the G α_i signaling branch completely reverts cAMP production due to the activated G α_s branch back to baseline levels [41, Fig 6A and 6B]. Earlier *in vitro* experiments with membranes of Sf9 cells expressing AC5, however, have not yielded complete inhibition by G α_i [23,42]. Adding high amounts of both G α_{olf} and G α_i to the assays of AC5-containing membranes did not completely inhibit AC5 - the enzyme still produced significant amounts of cAMP [23,42]. Nevertheless, a catalytically active ternary complex still supports coincidence detection for a wide range of ternary complex activity, as shown in Fig 3D–3F, as long as its catalytic rate is noticeably smaller than that of the active binary complex AC5 · G α_{olf} . Additionally, the data in [23] are interpreted and fitted with an alternative, more elaborate interaction scheme where the ternary complex can form and is not very active, and the production of cAMP is due to higher order, catalytically active complexes of AC5 with more than one G α_{olf} and G α_i subunit. We therefore suggest that the results of our study are plausible and supported by several experimental results, among which are experiments with the native system.

Relevance for corticostriatal synaptic plasticity

Knowing what an intracellular signal transduction network is composed of and the details of how it operates can help to clarify how it responds to extracellular events. The AC5 signal transduction network in D₁ MSNs (as well as in D₂ MSNs) interacts with a calcium-activated signal transduction network to regulate synaptic plasticity. Calcium influx at the synapse is necessary for synaptic potentiation, but exerts its effect only if accompanied by activation of the AC5 signaling module [43–45]. Thus, the AC5 signal transduction network gates plasticity in the corticostriatal synapses onto the MSNs. Now, an existing and inactive ternary complex in AC5 regulation has consequences on how this “gate” would be opened: disinhibition from active G α_i is necessary, accompanied by stimulation from G α_{olf} . That is, our findings suggest that the experimentally observed ACh \downarrow in the striatum is likely physiologically relevant for D₁ MSNs, and both a Da \uparrow and a ACh \downarrow are necessary to enable synaptic potentiation (see also [46]). They can also help in interpreting the functional role of the neuromodulatory signals in the striatum.

The kinetic model of the AC5 signal transduction network, built according to the findings of the MD simulations and with the parameters predicted with the BD simulations, reveals improved coincidence detection when compared to alternative regulatory schemes. This, together with the implications mentioned above arising from the existence and inactivity of the ternary complex, suggests that the regulation of AC5 has indeed evolved to perform coincidence detection of the two neuromodulatory signals.

Comparisons to other AC isoforms and AC-dependent cascades

In this study we have investigated the regulation of AC5 through interaction with the G α_{olf} and G α_i subunits. All membrane-bound AC isoforms are known to be stimulated by G α_s , a close homologue of G α_{olf} , while only ACs 1, 5 and 6 are inhibited by G α_i [6]. With this in mind, it is interesting to ask whether our findings concerning AC5 regulation, particularly the presence of an inactive ternary complex in the signalling network, could also be valid for cascades containing ACs 1 and 6. The sequences of rat G α_s and G α_{olf} are highly similar with an identity of 76.0% and similarity of 90.0% (S13 Fig). Restricting the comparison to the amino acid residues within 6 Å of AC5 in the modelled *apo* AC5 · G α_{olf} complex, the identity rises to 95.8%, with only a single position differing. For this reason, it is reasonable to assume that our findings regarding the formation of G α_{olf} -containing complexes are also applicable to G α_s . Indeed, our modelling of the AC5 · G α_{olf} complexes assumes that we can take crystal structures of AC · G α_s complexes as template structures. It should be noted that G α_s is known to be deactivated more slowly than G α_{olf} [47], which could reduce the ability of G α_s -activated networks to detect subsecond coincident signals (see [34]).

Previously, we have performed electrostatic analyses of mouse AC isoforms, to identify regions of electrostatic similarity within similarly regulated isoforms [48]. In that work, we showed that the molecular electrostatic potentials surrounding ACs 1, 5 and 6 in aqueous solution were very similar in the G α_i binding region of AC5, suggesting that the location of binding, and the bound orientation, of G α_i on ACs 1 and 6, could be similar. This electrostatic similarity was due to a more negative character, compared to other AC isoforms, which is complementary to the largely positive potential of the face of G α_i that contains the switch II region that is thought to interact with AC.

The sequence identities of AC isoforms with respect to AC5, in the C1 domain to which G α_i binds, show that ACs 1 and 6 are most similar, at 73.5% and 94.7%, respectively (Fig 5B). Considering only those residues predicted to be involved in the interaction with G α_i , the similarities are 64.7% and 91.2%. From the very high electrostatic and sequence similarity, it seems reasonable to assume that our findings should hold for AC6.

While AC1 is inhibited by $G\alpha_i$, the level of inhibition is much lower, with higher levels only seen for its forskolin or calmodulin-activated states [49], therefore directly applying our findings to AC1 cascades is more difficult. In the AC1 C1 domain sequence, there are three charge-altering substitutions in the region formed by the C-terminal end of α 3-helix and the loop connecting this helix to the β 4-strand (Fig 5C). These substitutions give this region a more negative character, and therefore it is reasonable to assume that they should not destabilize the binding mode we find for $G\alpha_i$ on AC5 to a large extent. In a previous mutagenesis study [21], including one of these substitutions (N559D by Uniprot sequence, N480D by sequence in [21]), which is also present in the non- $G\alpha_i$ inhibited ACs 2, 4, 7 and 8 (S14 Fig), as a mutation in AC5 was found to produce only a small reduction in its inhibition by $G\alpha_i$ (less than 2-fold increase in IC50). A more significant effect on the binding of $G\alpha_i$ to AC1 may occur due to a difference in the C-terminal residue of the α 3-helix of the C1 domain. In AC1, there is an alanine in this position, while in ACs 5 and 6, it is a valine. The wild-type AC5 construct used in [21] differed from the canonical sequence by having a methionine in this position (476 by their sequence numbering, 555 in the Uniprot sequence). They showed that mutating this residue to match the canonical sequence reduced the IC50 of $G\alpha_i$ by a third, while mutation to alanine gave a greater than 30-fold increase in IC50. Due to this apparent reduction in the affinity of AC1 for $G\alpha_i$, further MD simulations may be required to confirm the stability of a $G\alpha_{olf} \cdot AC1 \cdot G\alpha_i$ ternary complex. The lower sequence similarity between the C1 domains of ACs 1 and 5 also suggests that the allosteric effects on both the $G\alpha_{olf}/G\alpha_s$ binding groove and the active site could be different in a putative AC1 ternary complex. Again further MD simulations would be required to investigate this, as well as to further unravel the different computational properties of the AC isoforms found in different synapses.

Assumptions and limitations

As for any simulation study aimed at investigating *in vivo* subcellular processes, there are limitations that should be discussed. First, adenylyl cyclases, together with other components of the signal transduction networks they participate in, are organized as parts of multiprotein signaling complexes and/or are localized in structured microdomains in the cell which serve to compartmentalize the effects of the produced cAMP and efficiently activate downstream components of the networks and, ultimately, effectors [35–37, 39, 50]. The kinetic model here, on the other hand, assumes mass-action kinetics for the species included, i.e. it disregards any organization into multiprotein signaling complexes and instead describes a well-mixed solution of molecular species. This means that it reproduces the experimental measurements on cAMP in a partly phenomenological way. For example, AC5 needs to be presented with appropriate proportions of $G\alpha_{olf}$ and $G\alpha_i$ to reproduce the levels of cAMP, which may not necessarily be the same as the amounts of these proteins in real synapses. Second, a recent study has demonstrated that AC5 and the heterotrimeric G_{olf} protein are pre-assembled into a signaling complex and suggested that upon activation by the GPCR, the $G\alpha_{olf}$ subunit rearranges rather than physically dissociates from the $G\beta\gamma$ subunit [37]. This would imply an increase in the forward rate constant for AC5 and $G\alpha_{olf}$ binding predicted with the BD simulations. In light of this study, the effect of the hindered $G\alpha_{olf}$ binding to $AC5 \cdot G\alpha_i$ is not clear, since its main advantage of improving coincidence detection occurs precisely for values of k_{f1} around the BD estimates. It remains to be seen how much of an increase in the forward rate constant pre-assembly confers. A similar situation likely occurs for $G\alpha_i$, as well - the GPCRs and AC5 form oligomeric complexes that could include the G proteins [39], which would mean that the predicted value for the forward rate constant for AC5 and $G\alpha_i$ binding, k_{f2} , is also probably an underestimate. A larger value for k_{f2} is beneficial for coincidence detection both with the

simultaneous and the hindered simultaneous binding scheme (Fig 4G–4I). Finally, we should underline that the enzyme is additionally regulated by protein kinase A (PKA), calcium ions, nitric oxide, and the G $\beta\gamma$ subunit, and regulation via the transmembrane domains has been proposed (but not demonstrated so far) [4, 35, 51]. PKA is activated by cAMP and is the most common kinase to elicit the various downstream responses of the signal transduction network. PKA also inhibits AC5 via phosphorylation, and this is probably feedback that serves for signal termination. Calcium also inhibits AC5, an interaction which, due to excitatory synaptic input, might also help terminate its activity. The G $\beta\gamma$ enhances the effect of G α_s on AC5, but has no effect alone, and nitric oxide has an inhibitory effect whose purpose is also unknown. The measures for coincidence detection used here do not include these additional regulatory interactions, and this would provide a more complete picture of the enzyme's regulation in studies of various signaling scenarios.

On the same topic of molecular organization, while it is known that the D₂ and A_{2a} receptors are colocalized and form oligomeric complexes in D₂ MSNs, the same has not been demonstrated for the D₁ and M₄ receptors in D₁ MSNs. However, the D₁ and M₄ receptors have individually been reported in the postsynaptic density in striatal tissue (among other compartments) by electron microscopy [52,53], and activation of the M₄ receptors counteracts activation of the D₁ receptors both in striatal slices, as described above, and in striatal membrane preparations [17]. This makes it likely that they indeed colocalize enough to enable both G proteins to interact with AC5.

Furthermore, concerning the structural simulations carried out in the present study, it is appropriate to highlight some aspects. First, while the *apo* and the *holo* ternary complexes are relatively stable over about 2.1 and 1.1 μ s of MD simulation time, respectively, we cannot exclude that on a longer time scale such complexes could show a dissociation of the different protein units. In this regard, it is worth mentioning that both simulations of the *apo* and *holo* forms of the ternary complex have been repeated starting with new velocities for about 0.6 and 1.0 μ s, respectively, leading to the same overall conclusions described above. Second, in the systems simulated here, only the catalytic domains of AC5 are considered, while the transmembrane domains are not included. Although the transmembrane domains are important for the proper dimerization of the catalytic domains [54], the functionality of AC5 was experimentally found to be maintained upon removal of the transmembrane domains [21, 55]. It should also be noted that the structure of the AC5 catalytic domain modelled in this study does not include the C1b domain (whose structure has not yet been determined). The absence of this domain has been shown to reduce the sensitivity of AC5 to G α_i [21]. Nevertheless, the MD simulations performed here and by Van Keulen and R othlisberger [24] show that G α_i is able to inhibit the activity of AC5 in the absence of the C1b domain. This suggests that the lower inhibition is due to the reduced affinity of G α_i in the absence of C1b, rather than to a change in the internal mechanism of inhibition. Were it possible to include C1b in the BD simulations, this might lead to an increase in the rate of G α_i association, which would lead to increased synergy in both the simultaneous and hindered binding schemes.

As mentioned above, the membrane-anchoring is also neglected in the BD simulations in which freely diffusing solutes are assumed. This assumption has two main effects on the predicted association rate constants, which alter the predictions in opposite directions, leading to some degree of cancellation of errors. The anchoring of AC5 and G α_{off} in the membrane would add additional diffusional constraints that would potentially increase rates *in vivo*, by reducing the search space, while the slower diffusion of the lipid anchor in the membrane would slow down association.

It should also be noted that the simulated systems in this work were built by homology modelling using the then available X-ray crystal structures of AC as described in the Methods

section and also reported in previous studies [24, 56]. While this paper was in review, structures of AC9 including its two six-helix transmembrane domains (but not the C1b domain) were reported [57]. In order to check possible overlaps between atoms of the $G\alpha_i$ protein in our modelled ternary complex and atoms of the new AC9 structure (and of the membrane bilayer), we fitted the $C\alpha$ atoms of the C1a and C2a domains on the respective atoms of the AC9 protein (see S15 Fig). The figure clearly shows that the position of the $G\alpha_i$ protein in the ternary complex does not overlap with the membrane. Moreover, the conformation, orientation and binding mode of the $G\alpha_i$ to AC5 would also be perfectly suitable for AC9 complexation, providing additional support for our model. Finally, although not all parts of the AC9 enzyme could be resolved, this new cryo-EM structure could provide a basis for future modelling work to investigate the effects of membrane tethering on AC5 regulation.

Conclusions

In this work, we have investigated the stability and activity of a $G\alpha_{olf} \cdot AC5 \cdot G\alpha_i$ ternary complex by MD simulation and found that the complex appears to be stable on the microsecond time scale, but is unable to catalyze ATP conversion. Using BD simulations, we have made predictions of the association rate constants for the formation of both binary $AC5 \cdot G\alpha_i$ complexes, and the subsequent association of the second $G\alpha$ subunit to form the ternary complex.

Kinetic modelling of the AC5 signal transduction network showed that the predictions of the structure-based simulations maximize the ability of the network to recognize coincident neuromodulatory signals, with coincidence detection enhanced by both the presence of the ternary complex, and its reduced activity. Taken together, these results provide evidence that AC5 has evolved to perform coincidence detection of transient changes in the amounts of $G\alpha_{olf}$ and $G\alpha_i$ proteins, such that a brief deactivation of the $G\alpha_i$ signaling branch is needed to gate the $G\alpha_{olf}$ signal through. For the corticostriatal synapse on D_1 MSNs, this implies that both the transient rise in dopamine and the decrease in acetylcholine levels are necessary to trigger synaptic plasticity.

Methods

Modelling of binary AC5 complexes

The crystal structure (PDB ID 1AZS) of the ATP-free $AC \cdot G\alpha_s$ complex [20] was used as a template for the catalytic region of AC5 and $G\alpha_{olf}$ in the *apo* ternary complex. The template used for the initial complex conformation included 1AZS's C1 and C2 domains (more specifically, canine AC5 for C1 and rat AC2 for C2) for modelling the AC5 structure (UniprotKB Q04400) from *Rattus norvegicus* as well as the $G\alpha_s$ ' structure for the initial $G\alpha_{olf}$ (UniprotKB P38406) conformation from *Rattus norvegicus* [20, 58, 59]. The modelled structure of the myristoylated *Rattus norvegicus* $G\alpha_i$ subunit (UniprotKB P10824) interacting with guanosine-5'-triphosphate (GTP) and Mg^{2+} was taken from [56]. $G\alpha_i$ also has several isoforms, and the one referred to here is $G\alpha_{i1}$. Myristoylation is a post-translational modification of the N-terminus of $G\alpha_i$ that results in the covalent attachment of a 14-carbon saturated fatty acid to the N-terminal glycine residue of $G\alpha_i$ via an amide bond. The modelled AC5 and $G\alpha_i$ structures were used for docking $G\alpha_i$ on AC5's C1 domain to finalise the initial conformation of the ternary complex. This *apo* ternary complex setup was also used to simulate the *apo* forms of $AC5 \cdot G\alpha_i$ and $AC5 \cdot G\alpha_{olf}$ by removing the subunit not to be considered.

The crystal structure (PDB ID 1CJK) of the catalytic AC domains with a bound ATP analogue (Adenosine-5'- γ -thio-triphosphate), ATP γ S, and a $G\alpha_s$ interacting with the AC protein, was used as a template for the *holo* ternary complex. The template used for modelling

the active AC5 (UniprotKB Q04400) conformation in the ternary complex included the C1 and C2 domains from the PDB file 1CJK [22, 58, 59]. The $G\alpha_i$ subunit present in 1CJK was used as a template for the initial $G\alpha_{\text{olif}}$ (P38406) structure from *Rattus norvegicus* [22, 58, 59]. The modelled structure of the myristoylated *Rattus norvegicus* $G\alpha_i$ subunit (UniprotKB P10824) interacting with GTP and Mg^{2+} was taken from [56]. The active myristoylated *Rattus norvegicus* $G\alpha_i$ is referred to simply as $G\alpha_i$ because only a myristoylated form of $G\alpha_i$ was used in all simulations.

Modelling of ternary AC5 complexes

Membrane-bound ACs consist of two membrane-bound regions and two cytosolic domains. The latter form the active site of the enzyme and its structure has been determined by crystallography. The catalytic domains, C1 and C2, are located close to the membrane due to AC5's transmembrane domains, but remain entirely solvated. The crystal structure templates, used for modelling the complexes, were employed to determine the orientation of $G\alpha_{\text{olif}}$ on the C2 domain. The HADDOCK web server [60] was used for docking ten conformations of the active $G\alpha_i$ subunit to AC5's catalytic domains in the *apo* and *holo* forms as described in [24]. The active region of $G\alpha_i$ was defined in HADDOCK as a large part of the alpha helical domain (112-167), the switch I region (175-189) and the switch II region (200-220), allowing for a large unbiased area on the $G\alpha_i$ protein surface to be taken into account during docking. The active region of AC5's C1 domain was defined as the $\alpha 1$ helix (479-490) and the C-terminal region of the $\alpha 3$ helix (554-561) because experimentally it has been found that $G\alpha_i$ is unable to interact with C2 and its main interactions with AC5 are with the C1 domain [21]. Ten snapshots of $G\alpha_i$ were used for docking the $G\alpha$ subunit to the catalytic domain of AC5. These snapshots were extracted at time intervals of 0.5 ns from the end of the classical MD trajectory of $G\alpha_i$ (around 1.9 μ s) described in [56]. The same three criteria for complex selection as in [24] were applied: (1) the absence of overlap between the C2 domain and $G\alpha_i$, (2) no overlap with the GTP binding region of $G\alpha_i$ and the C1 domain of AC5, and (3) presence of similar complexes in the top-ten docking results of the docking calculations performed for all ten used $G\alpha_i$ conformations.

Classical molecular dynamics simulations

The Gromacs 5.1.2 software [61, 62] was used to perform the simulations. The *apo* and *holo* $G\alpha_{\text{olif}}$ · AC5 · $G\alpha_i$ systems, which were simulated for 2.1 μ s and 1.2 μ s respectively, each include two GTP molecules. In addition, the *holo* complex incorporates four Mg^{2+} ions and one ATP molecule, while in the *apo* form only three Mg^{2+} ions are present. Both *apo* and the *holo* systems were solvated in about 162,000 water molecules and 150 mM KCl. They were simulated at a temperature of 310 K and a pressure of 1 bar, maintained using the Nosé-Hoover thermostat and an isotropic Parrinello-Rahman barostat, respectively. The force fields used for the protein and the water molecules were AMBER99SB [63] and TIP3P [64]. For GTP and ATP, the force field generated by Meagher et al. was used [65]. The adjusted force field parameters for Cl^- and K^+ were taken from [66]. The Mg^{2+} parameters originated from [67] and the parameter set for the myristoyl group was taken from [56]. Electrostatic interactions were calculated with the Ewald particle mesh method with a real space cutoff of 12Å. The van der Waals interactions also had a cutoff of 12 Å. Bonds involving hydrogen atoms were constrained using the LINCS algorithm [68]. The integration time step was set to 2 fs.

The *apo* binary complexes, AC5 · $G\alpha_i$ and AC5 · $G\alpha_{\text{olif}}$, used for comparison to the ternary complexes, were built via the same procedure and were each simulated for 3 μ s.

The first step in the equilibration procedure of the protein systems included the energy minimisation of the protein complex together with the ligands (Mg^{2+} , GTP, ATP), referred to as the complete complex. Position restraints of $1000 \text{ kJ mol}^{-1} \text{ nm}^{-2}$ were applied to the structure of the complete complex during energy minimisation. The next step that was performed was the simulation of the complete complex under canonical NVT (constant number of atoms (N), constant volume (V) and constant temperature (T)) conditions, starting from the energy minimised structure, with position restraints of $1000 \text{ kJ mol}^{-1} \text{ nm}^{-2}$ on the complete complex. The length of the NVT run was 2 ns. The third step included the performance of a 4 ns isothermal-isobaric NPT run (constant number of atoms (N), constant pressure (P) and constant temperature (T)) with position restraints of $1000 \text{ kJ mol}^{-1} \text{ nm}^{-2}$ on the complete complex. The fourth step was an NPT simulation of 4 ns, with position restraints of $1000 \text{ kJ mol}^{-1} \text{ nm}^{-2}$ on the complete complex except for the hydrogens of the proteins. The fifth step contained an NPT run of 4 ns in which the backbone of the proteins were restrained by $1000 \text{ kJ mol}^{-1} \text{ nm}^{-2}$ as well as the ligands. The sixth step was an unrestrained NPT simulation of at least 10 ns, which was prolonged depending on the RMSD convergence of the proteins in the equilibrated system.

Forward rate constant estimation via Brownian dynamics simulations

Brownian dynamics (BD) simulations were performed to estimate the forward association rate constants in the two schemes in Fig 3. The simulations were performed using the SDA 7 software package [69], with the associating species represented in atomic resolution as rigid bodies. The inter-species interactions were modeled using the effective charge model [70], with the electrostatic desolvation term described by Elcock et al. [71], and following the parameterization of Gabdouliline and Wade [72].

The atomic structures of the reactant species in the forward reactions shown in Fig 3B and 3C were taken from clustered snapshots of the MD simulations described above, except for the structures of *apo* AC5 and $G\alpha_i$ used to calculate k_{f} , which were obtained from simulations performed by Van Keulen and R othlisberger [24]. The electrostatic potential of each snapshot of the reactant species was calculated via solution of the linearized Poisson-Boltzmann equation (PBE) using the APBS PBE solver [73], such that at the grid boundaries, the electrostatic potentials matched those of atom-centered Debye-H uckel spheres. The atomic charges of the protein residues were taken from the Amber force field, with the charges of the myristoyl moiety as described in [56] and the GTP charges as described by Meagher et al. [65]. The low-dielectric cavity ($\epsilon_r = 4$) was described using Bondi atomic radii [74] and the smoothed molecular surface definition of Bruccoleri et al. [75], while the solvent was modelled using a dielectric constant of 78, and a 150 mM concentration of salt with monovalent ions of radius 1.5  . For the single species reactants, solution of the linearized PBE generated cubic potential grids with 129 grid points per dimension, spaced at 1  , while larger grids with 161 points per dimension were generated for the binary reactants. Effective charges were calculated using the *ECM* module of SDA 7 [69, 70], with charge sites placed on the heteroatoms of ionized amino acid side chains and termini, the phosphate oxygen and phosphorus atoms of ATP and GTP, and Mg^{2+} ions. The effective charges of each solute were fitted such that, in a homogeneous dielectric environment, they reproduced the solute's electrostatic potential computed by solving the PBE in a skin bound by the surfaces described by rolling probe spheres of radii 4 and 7   along the molecular surface of the solute. Electrostatic desolvation potentials were calculated using the *make_edhalj_grid* tool in SDA 7 [69].

For each reaction, four BD simulations of 50 000 trajectories were performed using each MD snapshot, and rate constants were calculated using the Northrup, Allison and

McCammion formalism [76]. The mean and standard deviations across these four simulations was then determined. The average value for the corresponding rate constant was computed across all MD snapshots. The infinite dilution diffusion coefficients of each reaction species were calculated using HYDROPRO [77, 78] with the exception of those of the AC5 · Gα complex reactants in the ternary complex forming reactions, for which the diffusion coefficients of AC5 were used. In each simulation trajectory, the position of the center of AC5, or the reactant complex, was fixed at the center of the simulated volume, while the initial position of the center of the reactant Gα subunit was placed on the surface of a sphere of radius *b*, centered on the other reactant, with *b* taken as equal to the sum of the maximal extent of the distance of an atom of either reactant from the reactant center, plus the maximal extent of any interaction grid point to the solute center plus 30 Å. The simulations continued until the reactants diffused to a separation *c*, where *c* = 3*b*. Trajectories were assumed to have formed reactive encounter complexes when two independent native contacts between the two reactants reach a separation of 6.5 Å or less. Native contacts were defined as a pair of hydrogen bonding heteroatoms, separated by less than 5 Å in the bound complex. Two native contacts were assumed to be independent if the heteroatoms on the same solute that form the contacts were separated by more than 6 Å. This definition of an encounter complex has been shown to result in calculated protein-protein association rate constants that correlate well with experimental values [72].

Kinetic model of the signal transduction network

The kinetic model of the signal transduction network is a system of coupled ordinary differential equations with mass-action kinetics modeling the network's biochemical reactions. For example, for the reaction $A + B \xrightarrow{k_f} C$, the rate at which it occurs is given with:

$$v = \frac{d[C]}{dt} = -\frac{d[A]}{dt} = -\frac{d[B]}{dt} = k_f[A][B] - k_r[C]$$

In order to reduce the number of rate constants that would need estimation, our aim was to use a minimal model with which we could still study the coincidence detection ability of the enzyme and capitalize on the predictions of the molecular simulations. We have used two versions of the model, one with the allosteric exclusion and the other with the simultaneous binding scheme for AC5 regulation in Fig 3. The full reaction networks are given in S8 Fig. The two versions of the model have 8 and 16 rate constants, respectively. In Fig 3, we have additionally used versions of the model which included cAMP production to illustrate the correspondence between *k_c* and the levels of cAMP and thus rationalize the use of *k_c* as a proxy for the cAMP levels. The reactions and parameters for cAMP production and degradation have been taken from [15].

There are no receptors included in the model, and the Da↑ and ACh↓ inputs are modeled as changes in the rate constants for the conversion of inactive to active G proteins. Pools of inactive G_{oif} and G_i are activated at rates of *k_{fG_{oif}}* = 5s⁻¹ and *k_{fG_i}* = 5s⁻¹ when Da or ACh is present. This eliminates the parameters that would have corresponded to the reactions of ligand and receptor binding, and G protein and receptor binding. Additionally, we have omitted the heterotrimeric nature of the G proteins, i.e. we have not included the Gβγ subunit in the model. The G proteins are modeled as simply switching between an active and inactive form. The value for the rate of Gα_i activation is chosen so that there is a high resting-level inhibition of AC5 by the tonic level of ACh, which is a plausible assumption based on recent experimental results [14]. The value for *k_{fG_{oif}}* is, similarly, chosen to achieve amounts of active G_{oif} high enough to drive the binding reactions with AC5 and AC5 · Gα_i forward. The total

amounts of AC5 and the two G proteins used for this model are $n_{AC5} = 1500$ nM, $n_{Golf} = 1500$ nM and $n_{Gi} = 6000$ nM.

The activated $G\alpha_{olf}$ and $G\alpha_i$ interact with AC5 and can form each of the binary complexes and the ternary complex. Their deactivation is done by the intrinsic GTPase activity of the G proteins themselves, but is increased by AC5 for the case of $G\alpha_s$ (a homologue to $G\alpha_{olf}$) at least fivefold and the regulator of G protein signaling 9-2 (RGS9-2) for $G\alpha_i$ 20 to 40 times [41, 79], for which reason we have used values of $k_{rG_{olf}} = 0.5s^{-1}$ and $k_{rG_i} = 5s^{-1}$. If deactivated, the G proteins unbind from AC5. For the reverse (unbinding) rates of the G proteins from the binary AC5 complexes, we use values 100 times greater than the corresponding forward rate constants, which is the order of magnitude fitted in [23]. The reverse rates of G protein unbinding from the ternary complex, are increased by an order of magnitude compared to the reverse rates for unbinding from the respective binary complexes to qualitatively incorporate possible reduced stability of the ternary complex compared to the binary complexes on longer time scales (see Results section).

The remaining parameters, the forward rate constants of the G proteins' binding to AC5 and the binary complexes $AC5 \cdot G\alpha_{olf}$ and $AC5 \cdot G\alpha_i$, were estimated with the BD simulations (see Table 1). We have also varied these to explore their effects on the network's ability to perform coincidence detection.

Measures of coincidence detection

As was defined in the introduction, for the signal transduction network that we consider, coincidence detection means to respond with significant amounts of cAMP only when the two incoming signals $Da \uparrow$ and $ACh \downarrow$ coincide in time and in the spatial vicinity of the receptors. Note that there are two aspects of coincidence detection in the definition:

- distinguishing between the inputs $Da \uparrow + ACh \downarrow$, and a $Da \uparrow$ or $ACh \downarrow$ alone, and
- responding strongly enough with amounts of cAMP that are physiologically relevant.

To quantify how well the network distinguishes between the different inputs, we use the synergy quantity, and to quantify the strength of the response, we use the average catalytic rate, both defined below.

Average catalytic rate. The average catalytic rate is an average of the catalytic rates of the unbound form of AC5 and each of the complexes with the $G\alpha$ subunits. For the allosteric exclusion scheme in Fig 3B, where the ternary complex does not form, it is:

$$k_c = p_1 k_{c,AC5} + p_2 k_{c,AC5 \cdot G\alpha_{olf}} + p_3 k_{c,AC5 \cdot G\alpha_i}$$

with the weights

$$p_1 = \frac{[AC5]}{[AC5] + [AC5 \cdot G\alpha_{olf}] + [AC5 \cdot G\alpha_i]},$$

$$p_2 = \frac{[AC5 \cdot G\alpha_{olf}]}{[AC5] + [AC5 \cdot G\alpha_{olf}] + [AC5 \cdot G\alpha_i]},$$

$$p_3 = \frac{[AC5 \cdot G\alpha_i]}{[AC5] + [AC5 \cdot G\alpha_{olf}] + [AC5 \cdot G\alpha_i]}.$$

being the amounts of each enzyme species in the allosteric exclusion scheme as a fraction of

the total concentration of AC5 in the system. For the simultaneous binding scheme where the ternary complex does form, the average catalytic rate is

$$k_c = q_1 k_{c,AC5} + q_2 k_{c,AC5 \cdot G\alpha_{off}} + q_3 k_{c,AC5 \cdot G\alpha_i} + q_4 k_{c,G\alpha_{off} \cdot AC5 \cdot G\alpha_i}$$

with

$$q_1 = \frac{[AC5]}{[AC5] + [AC5 \cdot G\alpha_{off}] + [AC5 \cdot G\alpha_i] + [G\alpha_{off} \cdot AC5 \cdot G\alpha_i]},$$

$$q_2 = \frac{[AC5 \cdot G\alpha_{off}]}{[AC5] + [AC5 \cdot G\alpha_{off}] + [AC5 \cdot G\alpha_i] + [G\alpha_{off} \cdot AC5 \cdot G\alpha_i]},$$

$$q_3 = \frac{[AC5 \cdot G\alpha_i]}{[AC5] + [AC5 \cdot G\alpha_{off}] + [AC5 \cdot G\alpha_i] + [G\alpha_{off} \cdot AC5 \cdot G\alpha_i]},$$

$$q_4 = \frac{[G\alpha_{off} \cdot AC5 \cdot G\alpha_i]}{[AC5] + [AC5 \cdot G\alpha_{off}] + [AC5 \cdot G\alpha_i] + [G\alpha_{off} \cdot AC5 \cdot G\alpha_i]}.$$

We assume that the catalytic rate of the unbound AC5 is $k_{c,AC5} = 0.1s^{-1}$, and this rate is scaled by factors of $\alpha_{G\alpha_{off}}$ and $\alpha_{G\alpha_i}$ when the respective regulator G protein subunit binds:

$$k_{c,AC5 \cdot G\alpha_{off}} = \alpha_{G\alpha_{off}} k_{c,AC5}$$

$$k_{c,AC5 \cdot G\alpha_i} = \alpha_{G\alpha_i} k_{c,AC5}$$

The factors of stimulation and inhibition of AC5 are set to be $\alpha_{G\alpha_{off}} = 200$ ([23]) and $\alpha_{G\alpha_i} = 0.01s^{-1}$. For the catalytic rate of the ternary complex, we use the result of the MD simulations that the ternary complex is inactive, $\alpha_{G\alpha_{off},G\alpha_i} = \alpha_{G\alpha_i}$, i.e.:

$$k_{c,G\alpha_{off} \cdot AC5 \cdot G\alpha_i} = \alpha_{G\alpha_i} k_{c,AC5},$$

except in Fig 3D, 3E and 3F, where the catalytic rate of the ternary complex is varied to investigate its effect on coincidence detection.

Synergy. The synergy of two signals s_1 and s_2 is defined as

$$S(t) = \frac{r(s_1, s_2, t)}{r(s_1, t) + r(s_2, t) - r_{ss}},$$

where r_{ss} is the response at steady state. This quantity measures how strong the response r of the signal transduction network is for two coincident signals compared to the responses for single signals. The synergistic effect of the input signals can be examined in light of any quantity of interest in the network that is affected by the inputs, such as the level of activated PKA, for example [15]. Not having included PKA or cAMP in the kinetic model, we use the average catalytic rate k_c as a proxy for the amount of cAMP produced, since the latter directly depends on k_c . That is, the synergy of a simultaneous $Da \uparrow$ and $ACh \downarrow$ is:

$$S(t) = \frac{k_c(Da \uparrow, ACh \downarrow, t)}{k_c(Da \uparrow, t) + k_c(ACh \downarrow, t) - k_{c,ss}}.$$

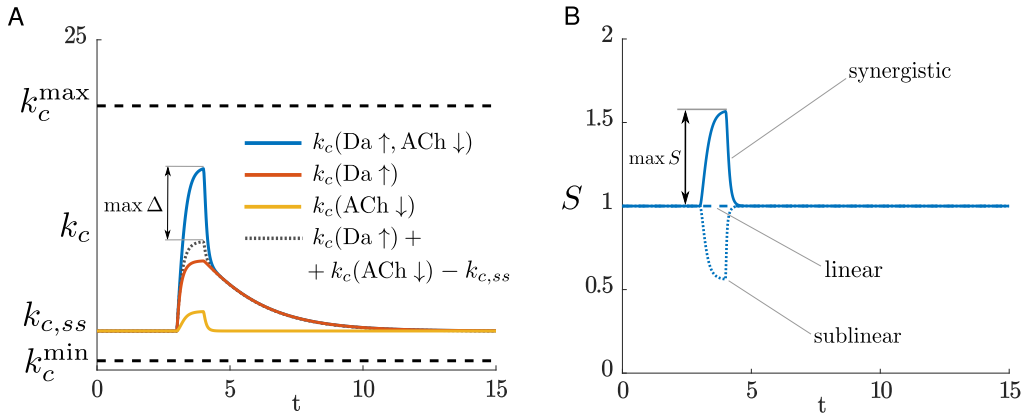


Fig 6. A A cartoon of the average catalytic rate during each of the input situations $Da \uparrow + ACh \downarrow$, $Da \uparrow$, and $ACh \downarrow$. B The synergy corresponding to panel (A). In addition, the dashed and dotted lines illustrate how the synergy for a linear and a sublinear response would look like, respectively.

<https://doi.org/10.1371/journal.pcbi.1007382.g006>

$S(t) > 1$ indicates a nonlinearly greater response in the presence of the two coincident signals, $S(t) = 1$ indicates a linear response to the coincident signals, and $S(t) < 1$ indicates a sublinear response. Hence, the case $S(t) > 1$ indicates that the signal transduction network can perform coincidence detection, whereas $S(t) \leq 1$ indicates an inability to do so.

Example traces for k_c and the corresponding synergy are given in Fig 6. Using

$$\Delta = k_c(Da \uparrow, ACh \downarrow, t) - (k_c(Da \uparrow, t) + k_c(ACh \downarrow, t) - k_{c,ss})$$

to express the difference between the response of the network for two coincident signals and the response for single signals, the expression for the synergy can also be rewritten as

$$S(t) = \frac{k_c(Da \uparrow, ACh \downarrow, t) - k_{c,ss}}{k_c(Da \uparrow, t) + k_c(ACh \downarrow, t) - k_{c,ss}} = 1 + \frac{\Delta}{k_c(Da \uparrow, t) + k_c(ACh \downarrow, t) - k_{c,ss}}$$

The difference in the responses, Δ , determines how big the synergistic effect of the input signals is (Fig 6).

Fig 6 is an example depicting how k_c relates to the synergy. There are minimum and maximum bounds on k_c : it would attain the minimum value k_c^{\min} if all of AC5 were inhibited by $G\alpha_i$, that is, only the catalytically inactive complex $AC5 \cdot G\alpha_i$ exists, where k_c^{\min} is the catalytic rate of $AC5 \cdot G\alpha_i$ (see above). Analogously, it would attain the maximum value k_c^{\max} if all of AC5 were bound in the catalytically active complex $AC5 \cdot G\alpha_{off}$, where k_c^{\max} is its catalytic rate (see above). In the models we use in this study, AC5 is never fully occupied by either of the $G\alpha$ subunits, and hence k_c is always between the minimum and maximum bounds. To maximize Δ , one would need to have k_c grow towards the maximum as much as possible during $Da \uparrow + ACh \downarrow$, and to have it as low as possible at steady state, and when each of the signals $Da \uparrow$ and $ACh \downarrow$ comes alone.

The metric C combines S and k_c . A combination of the average catalytic rate and the synergy gives a metric which can be helpful to evaluate whether the network both distinguishes between the different input situations and responds strongly. We use a product of k_c and S for this purpose, $C = Sk_c$.

Supporting information

S1 Fig. Root-mean-square deviations of three apo complexes: $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_{i1}$, $AC5 \cdot G\alpha_{i1}$ and $AC5 \cdot G\alpha_{\text{olf}}$. (Top panel) Individual RMSD values for C1, C2, $G\alpha_{i1}$ and $G\alpha_{\text{olf}}$ in the ternary complex only including the backbone of the domains. In the case of $G\alpha_{i1}$, residues 33 to 345 were taken into account during the RMSD calculation as the C and N termini are not structured. (Middle panel) Individual RMSD values for C1, C2 and $G\alpha_{i1}$ in the binary $AC5 \cdot G\alpha_{i1}$ complex only including the backbone of the domains. In the case of $G\alpha_{i1}$, residues 33 to 345 were taken into account during the RMSD calculation as the C and N termini are not structured. (Bottom panel) Individual RMSD values for C1, C2 and $G\alpha_{\text{olf}}$ in the binary $AC5 \cdot G\alpha_{\text{olf}}$ complex only including the backbone of the domains.

(EPS)

S2 Fig. Time evolution of the secondary structures for AC5 (top), $G\alpha_i$ (middle), and $G\alpha_{\text{olf}}$ (bottom) along the trajectory of the *apo* ternary complex.

(EPS)

S3 Fig. Time evolution of the secondary structures for AC5 (top), and $G\alpha_i$ (bottom), along the trajectory of the *apo* binary complex $AC5 \cdot G\alpha_i$.

(EPS)

S4 Fig. Time evolution of the secondary structures for AC5 (top), and $G\alpha_{\text{olf}}$ (bottom), along the trajectory of the *apo* binary complex $AC5 \cdot G\alpha_{\text{olf}}$.

(EPS)

S5 Fig. Time evolution of the number of hydrogen bonds present in the three simulated *apo* complexes along the respective MD trajectories.

(EPS)

S6 Fig. Root-mean-square fluctuations per residue calculated on the protein backbone of the different subunits (from top to bottom, AC5:C1, AC5:C2, $G\alpha_i$, and $G\alpha_{\text{olf}}$) of the three simulated *apo* complexes.

(EPS)

S7 Fig. Radius of gyration calculated along the MD trajectories of the three simulated apo complexes, $G\alpha_{\text{olf}} \cdot AC5 \cdot G\alpha_i$, $AC5 \cdot G\alpha_i$, and $AC5 \cdot G\alpha_{\text{olf}}$. The dashed lines indicate the values of the radius of gyration in the initial structures.

(EPS)

S8 Fig. The full kinetic models of the signal transduction networks used in this study.

(EPS)

S9 Fig. The effect of the interaction motif between AC5 and the regulatory $G\alpha$ subunits on coincidence detection. For the allosteric exclusion and simultaneous binding schemes, respectively, the amounts of each enzyme species as a percentage of the total amount of AC5 are shown for the cases of $Da \uparrow + ACh \downarrow$ (A, F), $Da \uparrow$ (B, G), and $ACh \downarrow$ (C, H). (D, I) Average catalytic rate for each scheme. (E, J) cAMP levels for each scheme.

(EPS)

S10 Fig. Time window for coincidence detection. (A) The detection window for the allosteric exclusion scheme and simultaneous binding scheme. Arrows are the time differences between $ACh \downarrow$ and $Da \uparrow$ chosen for the traces below. (B) The percentage of each AC5 species as a fraction of the total amount of AC5, (C) average catalytic rate, (D) synergy for the allosteric exclusion scheme. (E), (F), and (G) are the same quantities for the simultaneous binding scheme.

Note the shared axes.
(EPS)

S11 Fig. (A) The maximum of the synergy, (B) the maximum of k_c , (C) the maximum of the metric C as dependent on the rate of $G\alpha_i$ deactivation, k_{rGi} .
(EPS)

S12 Fig. Behavior of the two signal transduction schemes for different values of the association rate constants. From left to right: the percentage of enzyme species for $Da \uparrow + ACh \downarrow$, the percentage of enzyme species for $Da \uparrow$, the average catalytic rate, and the synergy, for the simultaneous binding scheme for (A) $k_{f1} = k_{f4} = 0.002 \text{ (nMs)}^{-1}$, $k_{f2} = k_{f3} = 2 \text{ (nMs)}^{-1}$, and (B) $k_{f1} = k_{f4} = 2 \text{ (nMs)}^{-1}$, $k_{f2} = k_{f3} = 2 \text{ (nMs)}^{-1}$ and the hindered simultaneous binding scheme for (C) $k_{f1} = k_{f4} = 0.002 \text{ (nMs)}^{-1}$, $k_{f2} = k_{f3} = 2 \text{ (nMs)}^{-1}$ and (D) $k_{f1} = k_{f4} = 2 \text{ (nMs)}^{-1}$, $k_{f2} = k_{f3} = 2 \text{ (nMs)}^{-1}$. Note the shared axes.
(EPS)

S13 Fig. Structure of $G\alpha_{olf}$ in the modelled $AC5 \cdot G\alpha_{olf}$ complex and sequence alignment of $G\alpha_{olf}$ and $G\alpha_s$. The structure of $G\alpha_{olf}$ in the modelled $AC5 \cdot G\alpha_{olf}$ apo complex (A). The highlighted regions show the switch II helix residues that interact the C2 binding groove on AC5 (magenta) and other amino acid residues that are within 6 of AC5 in the modelled structure. The sequence alignment of rat $G\alpha_{olf}$ (GNAL) and $G\alpha_s$ (GNAS2) (B). The magenta and green regions show the residues highlighted in (A). The yellow region shows the N-terminal residues not included in the structure used in this work.
(EPS)

S14 Fig. Multiple sequence alignment for all mouse AC isoforms with the colors matching those the structure in Fig 5A. The sequences were taken from Uniprot, and aligned using Clustal Omega within Uniprot. The red and blue boxes show positions where AC1 has substitutions compared to AC5, as described in Fig 5.
(EPS)

S15 Fig. Initial modelled configuration of the $G\alpha_{olf} \cdot AC5 \cdot G\alpha_i$ ternary complex used in our classical MD simulations, after fitting of $C\alpha$ atoms of C1a and C2a domains on the respective $C\alpha$ atoms of the AC9 protein (PDB ID: 6R3Q). AC9 (pale yellow) consisting of the transmembrane domain (TM), the helical domain (HD), and the two pseudo-symmetric C1a and C2a domains, is fully shown in cartoon representation. The two domains C1a (blue) and C2a (red) of AC5 in complex with $G\alpha_i$ (cyan) and $G\alpha_{olf}$ (gray) proteins are also represented as a cartoon.
(EPS)

S16 Fig. The allosteric exclusion scheme does not support coincidence detection for saturating concentrations of $G\alpha_{olf}$.
(EPS)

S1 Table. Bimolecular association rate constants $(\text{nMs})^{-1}$ for the forward reactions computed via BD simulations. Each rate constant was calculated using a number of snapshots from MD simulations. The reported numbers for each snapshot are the mean values estimated from 4 BD simulations of 50 000 trajectories (standard deviation in parentheses). Rate constants were calculated for complexes including apo and holo AC5 (superscripts a and h respectively).
(XLSX)

S1 Text. The allosteric exclusion scheme inherently lacks the ability for coincidence detection.
(PDF)

Acknowledgments

The authors thank Pietro Vidossich for scientific discussions during this project.

Author Contributions

Conceptualization: Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, Jeanette Hellgren Kotaleski.

Funding acquisition: Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, Jeanette Hellgren Kotaleski.

Investigation: Neil J. Bruce, Daniele Narzi, Daniel Trpevski, Siri C. van Keulen, Anu G. Nair, Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, Jeanette Hellgren Kotaleski.

Methodology: Neil J. Bruce, Daniele Narzi, Daniel Trpevski, Siri C. van Keulen, Anu G. Nair, Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, Jeanette Hellgren Kotaleski.

Project administration: Jeanette Hellgren Kotaleski.

Supervision: Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, Jeanette Hellgren Kotaleski.

Writing – original draft: Neil J. Bruce, Daniele Narzi, Daniel Trpevski, Siri C. van Keulen, Anu G. Nair.

Writing – review & editing: Neil J. Bruce, Daniele Narzi, Daniel Trpevski, Siri C. van Keulen, Anu G. Nair, Ursula Röthlisberger, Rebecca C. Wade, Paolo Carloni, Jeanette Hellgren Kotaleski.

References

1. Glatt CE, & Snyder SH. (1993). Cloning and expression of an adenylyl cyclase localized to the corpus striatum. *Nature*, 361, 536–538. <https://doi.org/10.1038/361536a0> PMID: 8429907
2. Mons M, & Cooper DMF. (1994). Selective expression of one Ca²⁺-inhibitible adenylyl cyclase in dopaminergically innervated rat brain regions. *Molecular Brain Research*, 22, 236–244. [https://doi.org/10.1016/0169-328x\(94\)90052-3](https://doi.org/10.1016/0169-328x(94)90052-3) PMID: 8015383
3. Changeux J-P, & Edelstein SJ. (2005). Allosteric mechanisms of Signal Transduction. *Science*, 308, 1424–1428. <https://doi.org/10.1126/science.1108595> PMID: 15933191
4. Linder JU. (2006). Class III adenylyl cyclases: Molecular mechanisms of catalysis and regulation. *Cellular and Molecular Life Sciences*, 63(15), 1736–1751. <https://doi.org/10.1007/s00018-006-6072-0> PMID: 16786220
5. Linder JU, & Schultz JE. (2003). The class III adenylyl cyclases: Multi-purpose signalling modules. *Cellular Signalling*, 15(12), 1081–1089. [https://doi.org/10.1016/s0898-6568\(03\)00130-x](https://doi.org/10.1016/s0898-6568(03)00130-x) PMID: 14575863
6. Sadana R, & Dessauer CW. (2009). Physiological roles for G protein-regulated adenylyl cyclase isoforms: insights from knockout and overexpression studies. *Neurosignals*, 17, 5–22. <https://doi.org/10.1159/000166277> PMID: 18948702
7. Anholt RRH. (1994). Signal integration in the nervous system: adenylyl cyclases as molecular coincidence detectors. *Trends in Neurosciences*, 17(1), 37–41. [https://doi.org/10.1016/0166-2236\(94\)90033-7](https://doi.org/10.1016/0166-2236(94)90033-7) PMID: 7511849
8. Bourne H. R., & Nicoll R. (1993). Molecular machines integrate coincident synaptic signals. *Cell*, 72(1001), 65–75. [https://doi.org/10.1016/S0092-8674\(05\)80029-7](https://doi.org/10.1016/S0092-8674(05)80029-7)

9. Delmeire D, Flamez D, Hinke SA, Cali JJ, Pipeleers D, & Schuit F. (2003). Type VIII adenylyl cyclase in rat beta cells: Coincidence signal detector/generator for glucose and GLP-1. *Diabetologia*, 46(10), 1383–1393. <https://doi.org/10.1007/s00125-003-1203-8> PMID: 13680124
10. Impey S, Wayman G, Wu Z, & Storm DR. (1994). Type I Adenylyl Cyclase Functions as a Coincidence Detector for Control of Cyclic AMP Response Element-Mediated Transcription: Synergistic Regulation of Transcription by Ca²⁺ and Isoproterenol. *Molecular and Cellular Biology*, 14(12), 8272–8281. <https://doi.org/10.1128/mcb.14.12.8272> PMID: 7969163
11. Lustig KD, Conklin BR, Herzmark P, Taussig R, & Bourne HR. (1993). Type II adenylylcyclase integrates coincident signals from G_s, G_i, and G_q. *Journal of Biological Chemistry*, 268(19), 13900–13905. PMID: 8390980
12. McVey M, Hill J, Howlett A, & Klein C. (1999). Adenylyl cyclase, a coincidence detector for nitric oxide. *Journal of Biological Chemistry*, 274(27), 18887–18892. <https://doi.org/10.1074/jbc.274.27.18887> PMID: 10383385
13. Mons N, Guillou JL, & Jaffard R. (1999). The role of Ca²⁺/calmodulin-stimulable adenylyl cyclases as molecular coincidence detectors in memory formation. *Cellular and Molecular Life Sciences*, 55(4), 525–533. <https://doi.org/10.1007/s000180050311> PMID: 10357223
14. Nair AG, Castro LRV, Khoury M, El, & Vincent P. (2019). The high efficacy of muscarinic M4 receptor in D1 medium spiny neurons reverses striatal hyperdopaminergia. *Neuropharmacology*, 146, 74–83. <https://doi.org/10.1016/j.neuropharm.2018.11.029> PMID: 30468798
15. Nair AG, Gutierrez-Arenas O, Eriksson O, Vincent P, & Hellgren Kotaleski J. (2015). Sensing Positive versus Negative Reward Signals through Adenylyl Cyclase-Coupled GPCRs in Direct and Indirect Pathway Striatal Medium Spiny Neurons. *Journal of Neuroscience*, 35(41), 14017–14030. <https://doi.org/10.1523/JNEUROSCI.0730-15.2015> PMID: 26468202
16. Onali P, Olanas MC. (2002) Muscarinic M4 receptor inhibition of dopamine D1-like receptor signalling in rat nucleus accumbens. *European Journal of Pharmacology* 448(2–3): 105–111, [https://doi.org/10.1016/s0014-2999\(02\)01910-6](https://doi.org/10.1016/s0014-2999(02)01910-6) PMID: 12144929
17. Jeon J, Dencker D, Wörtwein G, Woldbye DPD, Cui Y, Davis AA, Levey AI, Schütz G, Sager TN, Mørk A, Li C, Deng C-X, Fink-Jensen A and Wess J (2010). A Subpopulation of Neuronal M4 Muscarinic Acetylcholine Receptors Plays a Critical Role in Modulating Dopamine-Dependent Behaviors. *Journal of Neuroscience* 30 (6) 2396–2405, <https://doi.org/10.1523/JNEUROSCI.3843-09.2010> PMID: 20147565
18. Ferre S, Quiroz C, Woods AS, Cunha R, Popoli P, Ciruela F, Lluís C, Franco R, Azdad K, and Schiffmann SN. (2008) An Update on Adenosine A2A-Dopamine D2 Receptor Interactions: Implications for the Function of G Protein-Coupled Receptors. *Current Pharmaceutical Design*, 14, (15) 1468–1474(7) <https://doi.org/10.2174/138161208784480108> PMID: 18537670
19. Yapo C, Nair AG, Clement L, Castro LR, Hellgren Kotaleski J, & Vincent P. (2017). Detection of phasic dopamine by D1 and D2 striatal medium spiny neurons. *Journal of Physiology*, 595(24), 7451–7475. <https://doi.org/10.1113/JP274475> PMID: 28782235
20. Tesmer JJ, Sunahara RK, Gilman AG, & Sprang SR. (1997) Crystal structure of the catalytic domains of adenylyl cyclase in a complex with G α . *GTPgammaS Science* 278:1907–1916. <https://doi.org/10.1126/science.278.5345.1907> PMID: 9417641
21. Dessauer CW, Tesmer JGG, Sprang SR, & Gilman AG. (1998). Identification of a Gi Binding Site on Type V Adenylyl Cyclase. *Journal of Biological Chemistry*, 273(40), 25831–25839. <https://doi.org/10.1074/jbc.273.40.25831> PMID: 9748257
22. Tesmer JGG, Sunahara RK, Johnson RA, Gosselin G, Gilman AG, & Sprang SR. (1999). Two-metal-Ion catalysis in adenylyl cyclase. *Science*, 285(5428), 756–760. <https://doi.org/10.1126/science.285.5428.756> PMID: 10427002
23. Chen-Goodspeed M, Lukan AN, & Dessauer CW. (2005). Modeling of Gas and Gai regulation of human type V and VI adenylyl cyclase. *Journal of Biological Chemistry*, 280(3), 1808–1816. <https://doi.org/10.1074/jbc.M409172200> PMID: 15545274
24. Van Keulen SC, & Rothlisberger U. (2017b). Exploring the Inhibition Mechanism of Adenylyl Cyclase Type 5 by N-terminal Myristoylated Gai1:GTP. *PLoS Computational Biology*, 13(9), e1005673. <https://doi.org/10.1371/journal.pcbi.1005673> PMID: 28892485
25. Frezza E, Martin J, & Lavery R. (2018) A molecular dynamics study of adenylyl cyclase: The impact of ATP and G-protein binding. *PLoS ONE* 13(4): e0196207. <https://doi.org/10.1371/journal.pone.0196207>
26. Hollingsworth S, and Dror RO. (2018). Molecular Dynamics Simulation for All. *Neuron* 99(6): 1129–1143. <https://doi.org/10.1016/j.neuron.2018.08.011> PMID: 30236283
27. Stein M, Gabdoulline RR, Wade RC. (2007). Bridging from molecular simulation to biochemical networks. *Current Opinion in Structural Biology*, 17(2), 166–72. <https://doi.org/10.1016/j.sbi.2007.03.014> PMID: 17395455

28. Boras BW, Hirakis SP, Votapka LW, Malmstrom RD, Amaro RE, & McCulloch AD. (2015). Bridging scales through multiscale modeling: A case study on protein kinase A. *Frontiers in Physiology*, 6(SEP), 250. <https://doi.org/10.3389/fphys.2015.00250> PMID: 26441670
29. Kitano H. (2002). Computational systems biology. *Nature*, 420(November), 206–210. <https://doi.org/10.1038/nature01254> PMID: 12432404
30. Tesmer JGG and Sprang SR. (1998). The structure, catalytic mechanism and regulation of adenylyl cyclase. *Current opinion in structural biology*, 8(6), 713–719. [https://doi.org/10.1016/S0959-440X\(98\)80090-0](https://doi.org/10.1016/S0959-440X(98)80090-0) PMID: 9914249
31. Steitz TA. (1993). DNA- and RNA-dependent DNA polymerases. *Current Opinion in Structural Biology*, 3(1), 31–38. [https://doi.org/10.1016/0959-440X\(93\)90198-T](https://doi.org/10.1016/0959-440X(93)90198-T)
32. Nakamura T, Zhao Y, Yamagata Y, Hua YJ, and Yang W. (2012). Watching DNA polymerase η make a phosphodiester bond. *Nature*, 487(7406), 196–201. <https://doi.org/10.1038/nature11181> PMID: 22785315
33. Zhang YF, and Cragg SJ. (2017) Pauses in Striatal Cholinergic Interneurons: What is Revealed by Their Common Themes and Variations? *Frontiers in Systems Neuroscience* 11:80. <https://doi.org/10.3389/fnsys.2017.00080> PMID: 29163075
34. Nair AG. (2018). Modeling Biochemical Network Involved in Striatal Dopamine Signaling (TRITA-EECS-AVL; 2018:7). *KTH Royal Institute of Technology*.
35. Cooper DF. (2003) Regulation and organization of adenylyl cyclases and cAMP. *Biochem. J.* 375, 517–529. <https://doi.org/10.1042/BJ20031061> PMID: 12940771
36. Cooper DF, and Crossthwaite AJ. (2006) Higher-order organization and regulation of adenylyl cyclases. *Trends in Pharmacological Sciences*, 27(8): 426–431. <https://doi.org/10.1016/j.tips.2006.06.002> PMID: 16820220
37. Xie K, Masuho I, Shih C-C, Cao Y, Sasaki K, Lai CWJ, Han P-L, Ueda H, Dessauer CW, Ehrlich ME, Xu B, Willardson BM, Martemyanov KA. (2015) Stable G protein-effector complexes in striatal neurons: mechanism of assembly and role in neurotransmitter signaling. *eLife* 4: e10451. <https://doi.org/10.7554/eLife.10451> PMID: 26613416
38. Hahn DK, Tusell JR, Sprang SR and Chu X. (2015) Catalytic Mechanism of Mammalian Adenylyl Cyclase: A Computational Investigation. *Biochemistry*, 54(40), 6252–6262. <https://doi.org/10.1021/acs.biochem.5b00655> PMID: 26393535
39. Navarro G, Cordermi A, Casadó-Anguera, Moreno E, Cai N-S, Cortés A, Canela EI, Dessauer CW, Casadó V, Pardo L, Lluís C, and Ferré S. Evidence for functional pre-coupled complexes of receptor heteromers and adenylyl cyclase. *Nature Communications* 9, Article number: 1242 (2018), <https://doi.org/10.1038/s41467-018-03522-3>
40. Iwamoto T, Iwatsubo K, Okumura S, Hashimoto Y, Tsunematsu T, Toya Y, Hervé D, Umemura S and Ishikawa Y (2004). Disruption of type 5 adenylyl cyclase negates the developmental increase in G α expression in the striatum. *FEBS Letters*, 564, [https://doi.org/10.1016/S0014-5793\(04\)00333-3](https://doi.org/10.1016/S0014-5793(04)00333-3)
41. Xie K, Masuho I, Brand C, Dessauer CW, Martemyanov KA (2012) The complex of G protein regulator RGS9-2/G β 5 controls sensitization and signaling kinetics of type 5 adenylyl cyclase. *Sci Signal* 5:ra63, <https://doi.org/10.1126/scisignal.2002922> PMID: 22932702.
42. Taussig R, Tang W-J, Hepler JR, and Gilman AG. (1994) Distinct Patterns of Bidirectional Regulation of Mammalian Adenylyl Cyclases. *The Journal of Biological Chemistry*, Vol. 269, No. 8, pp. 6093–6100 PMID: 8119955
43. Fino E, Glowinski J, & Venance L. (2005) Bidirectional Activity-Dependent Plasticity at Corticostriatal Synapses. *The Journal of Neuroscience* 25(49): 11279–11287 <https://doi.org/10.1523/JNEUROSCI.4476-05.2005> PMID: 16339023
44. Pawlak V and Kerr JND. (2008) Dopamine Receptor Activation is Required for Corticostriatal Spike-Timing-Dependent Plasticity. *The Journal of Neuroscience* 28(10): 2435–2446 <https://doi.org/10.1523/JNEUROSCI.4402-07.2008> PMID: 18322089
45. Shen W, Flajolet M, Greengard P, and Surmeier J. (2008) Dichotomous Dopaminergic control of striatal synaptic plasticity. *Science* 321(5890): 848–851 <https://doi.org/10.1126/science.1160575> PMID: 18687967
46. Fisher SD, Robertson PB, Black MJ, Redgrave P, Sagar MA, Abraham WC, and Reynolds JNJ (2017), Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo. *Nature Communications*, 8:334, <https://doi.org/10.1038/s41467-017-00394-x> PMID: 28839128
47. Lui Y-H, Wenzel-Seifert K, & Seifert R. (2001). The olfactory G protein G α olf possesses a lower GDP-affinity and deactivates more rapidly than G α short: Consequences for receptor-coupling and adenylyl cyclase activation. *Journal of Neurochemistry*, 78(2), 325–338. <https://doi.org/10.1046/j.1471-4159.2001.00422.x> PMID: 11461968

48. Tong R, Wade RC, & Bruce NJ. (2016). Comparative electrostatic analysis of adenylyl cyclase for isoform dependent regulation properties. *Proteins: Structure, Function and Bioinformatics*, 84(13), 1844–1858. <https://doi.org/10.1002/prot.25167> PMID: 27667304
49. Taussig R, Iniguez-Lluhi JA, & Gilman AG. (1992). Inhibition of Adenylyl Cyclase by Giα. *Science*, 261, 218–221.
50. Dessauer C. (2009) Adenylyl Cyclase–A-kinase Anchoring Protein Complexes: The Next Dimension in cAMP Signaling. *Molecular Pharmacology* 76 (5): 935–941 <https://doi.org/10.1124/mol.109.059345> PMID: 19684092
51. Brand C. S., Sadana R., Malik M., Smrcka A. V., & Dessauer C. W. (2015). Adenylyl Cyclase 5 Regulation by Gβγ Involves Isoform-Specific Use of Multiple Interaction Sites. *Molecular Pharmacology*, 88, 758–767. <https://doi.org/10.1124/mol.115.099556> PMID: 26206488
52. Hersch SM, Gutekunst CA, Rees HD, Heilman CJ and Levey AI. Distribution of m1-m4 muscarinic receptor proteins in the rat striatum: light and electron microscopic immunocytochemistry using subtype-specific antibodies. *Journal of Neuroscience* 1994; 14 (5) 3351–3363; <https://doi.org/10.1523/JNEUROSCI.14-05-03351.1994>
53. Hersch SM, Ciliax BJ, Gutekunst CA, Rees HD, Heilman CJ, Yung KK, Bolam JP, Ince E, Yi H and Levey AI, Electron microscopic analysis of D1 and D2 dopamine receptor proteins in the dorsal striatum and their synaptic relationships with motor corticostriatal afferents. *Journal of Neuroscience* 1995; 15 (7) 5222–5237; <https://doi.org/10.1523/JNEUROSCI.15-07-05222.1995>
54. Seebacher T, Linder JU, & Schultz JE. (2001). An isoform-specific interaction of the membrane anchors affects mammalian adenylyl cyclase type V activity. *European Journal of Biochemistry*, 268, 105–110. <https://doi.org/10.1046/j.1432-1327.2001.01850.x> PMID: 11121109
55. Whisnant RE, Gilman AG, & Dessauer CW. (1996) Interaction of the two cytosolic domains of mammalian adenylyl cyclase. *Proceedings of the National Academy of Sciences of the United States of America* 93(13): p. 6621–6625. <https://doi.org/10.1073/pnas.93.13.6621> PMID: 8692867
56. Van Keulen SC, & Rothlisberger U. (2017a). Effect of N-Terminal Myristoylation on the Active Conformation of Gαi1-GTP. *Biochemistry*, 56(1), 271–280. <https://doi.org/10.1021/acs.biochem.6b00388> PMID: 27936598
57. Qi C, Sorrentino S, Medalia O, and Korkhov VM. (2019) The structure of a membrane adenylyl cyclase bound to an activated stimulatory G protein. *Science* 364(6438): 389–394. <https://doi.org/10.1126/science.aav0778> PMID: 31023924
58. Eswar N, Webb B, Marti-Renom MA, Madhusudhan MS, Eramian D, Shen M, Pieper U, & Sali A. (2006). Comparative protein structure modeling using modeller. *Current Protocols in Bioinformatics*, 5–6.
59. Marti-Renom MA, Stuart AC, Fiser A, Sánchez R, Melo F & Sali A. (2000). Comparative protein structure modeling of genes and genomes. *Annual Review of Biophysics and Biomolecular Structure*, 29 (1):291–325.
60. De Vries SJ, van Dijk M, & Bonvin AM. J. J. (2010). The haddock web server for data-driven biomolecular docking. *Nature Protocols*, 5(5):883–897, 4. <https://doi.org/10.1038/nprot.2010.32> PMID: 20431534
61. Bekker H, Berendsen HJC, Dijkstra EJ, Achterop S, Van Drunen R, Van der Spoel D, Sijbers A, Keestra H, and Renardus MKR. (1993). Gromacs: A parallel computer for molecular dynamics simulations. *Physics Computing*, 92:252–256.
62. Abraham MJ, Murtola T, Schulz R, Páll S, Smith JC, Hess B, Lindahl E. (2015). GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* 1, 19–25.
63. Hornak V, Abel R, Okur A, Strockbine B, Roitberg A & Simmerling C. (2006). Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins: Structure, Function, and Bioinformatics*. 65(3), 712–725.
64. Jorgensen WL, Chandrasekhar J, Madura JD, Impey RW & Klein ML. (1983). Comparison of simple potential functions for simulating liquid water. *The Journal of Chemical Physics*. 79(2), 926–935.
65. Meagher KL, Redman LT, & Carlson HA. (2003). Development of polyphosphate parameters for use with the AMBER force field. *Journal of Computational Chemistry*, 24(9), 1016–1025. <https://doi.org/10.1002/jcc.10262> PMID: 12759902
66. Joung IS & Cheatham TE III. (2008). Determination of alkali and halide monovalent ion parameters for use in explicitly solvated biomolecular simulations. *The Journal of Physical Chemistry B*. 112(30), 9020–9041. <https://doi.org/10.1021/jp8001614> PMID: 18593145
67. Allnér O, Nilsson L, & Villa A. (2012). Magnesium ion–water coordination and exchange in biomolecular simulations. *Journal of Chemical Theory and Computation*., 8(4), 1493–1502. <https://doi.org/10.1021/ct3000734> PMID: 26596759

68. Hess BP. (2008). LINCS: A parallel linear constraint solver for molecular simulation. *Journal of Chemical Theory and Computation*, 4(1), 116–122. [7 https://doi.org/10.1021/ct700200b](https://doi.org/10.1021/ct700200b) PMID: 26619985
69. Martinez M, Bruce NJ, Romanowska J, Kokh DB, Ozboyaci M, Yu X, Öztürk MA, Richter S, & Wade RC. (2015). SDA 7: A modular and parallel implementation of the simulation of diffusional association software. *Journal of Computational Chemistry*, 36(21), 1631–1645. <https://doi.org/10.1002/jcc.23971> PMID: 26123630
70. Gabbouline RR & Wade RC. (1996). Effective Charges for Macromolecules in Solvent. *Journal of Physical Chemistry*, 100(9), 3868–3878. <https://doi.org/10.1021/jp953109f>
71. Elcock AH, Gabbouline RR, Wade RC, & McCammon JA. (1999). Computer simulation of protein-protein association kinetics: acetylcholinesterase-fasciculin. *Journal of Molecular Biology*, 291(1), 149–162. <https://doi.org/10.1006/jmbi.1999.2919> PMID: 10438612
72. Gabbouline RR & Wade RC. (2001). Protein-protein association: investigation of factors influencing association rates by brownian dynamics simulations. *Journal of Molecular Biology*, 306(5), 1139–1155. <https://doi.org/10.1006/jmbi.2000.4404> PMID: 11237623
73. Baker NA, Sept D, Joseph S, Holst MJ, & McCammon JA. (2001). Electrostatics of nanosystems: application to microtubules and the ribosome. *Proceedings of the National Academy of Sciences U. S. A.*, 98(18), 10037–10041. <https://doi.org/10.1073/pnas.181342398> PMID: 11517324
74. Bondi A. (1964). van der Waals Volumes and Radii. *The Journal of Physical Chemistry*, 68(3), 441–451. <https://doi.org/10.1021/j100785a001>
75. Bruccoleri R, Novotny J, Davis M, & Sharp K. (1997). Finite difference Poisson-Boltzmann electrostatic calculations: Increased accuracy achieved by harmonic dielectric smoothing and charge antialiasing. *Journal of Computational Chemistry*, 18(2), 268–276. [https://doi.org/10.1002/\(SICI\)1096-987X\(19970130\)18:2<268::AID-JCC11>3.0.CO;2-E](https://doi.org/10.1002/(SICI)1096-987X(19970130)18:2<268::AID-JCC11>3.0.CO;2-E)
76. Northrup SH, Allison SA, & McCammon JA. (1984). Brownian dynamics simulation of diffusion-influenced bimolecular reactions. *The Journal of Chemical Physics*, 80(4), 1517. <https://doi.org/10.1063/1.446900>
77. García De La Torre J, Huertas ML, & Carrasco B. (2000). Calculation of hydrodynamic properties of globular proteins from their atomic-level structure. *Biophysical Journal*, 78(2), 719–730. [https://doi.org/10.1016/S0006-3495\(00\)76630-6](https://doi.org/10.1016/S0006-3495(00)76630-6) PMID: 10653785
78. Ortega A, Amorós D, & García de la Torre J. (2011). Prediction of hydrodynamic and other solution properties of rigid proteins from atomic- and residue-level models. *Biophysical Journal*, 101(4), 892–898. <https://doi.org/10.1016/j.bpj.2011.06.046> PMID: 21843480
79. Scholich K, Mullenix JB, Wittpoth C, Poppleton HM, Pierre SC, Lindorfer MA, Garrison JC, Patel TB. (1999) Facilitation of signal onset and termination by adenylyl cyclase. *Science* 283:1328–1331, <https://doi.org/10.1126/science.283.5406.1328> PMID: 10037603

Article B



Local, calcium- and reward-based synaptic learning rule that enhances dendritic nonlinearities can solve the nonlinear feature binding problem
(2024) eLife13:RP97274

Local, calcium- and reward-based synaptic learning rule that enhances dendritic nonlinearities can solve the nonlinear feature binding problem

Zahra Khodadadi^{1*}, Daniel Trpevski¹, Robert Lindroos^{1†}, Jeanette Hellgren Kotaleski^{1,2†}

¹ Department of Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden

² Department of Neuroscience, Karolinska Institutet, Stockholm, Sweden

***For correspondence:**

(ZK)

[†] These authors contributed equally to this work

Abstract

This study explores the computational potential of single striatal projection neurons (SPN), emphasizing dendritic nonlinearities and their crucial role in solving complex integration problems. Utilizing a biophysically detailed multicompartmental model of an SPN, we introduce a calcium-based, local synaptic learning rule that leverages dendritic plateau potentials. According to what is known about excitatory corticostriatal synapses, the learning rule is governed by local calcium dynamics from NMDA and L-type calcium channels and dopaminergic reward signals. In addition, we incorporated metaplasticity in order to devise a self-adjusting learning rule which ensures stability for individual synaptic weights. We demonstrate that this rule allows single neurons to solve the nonlinear feature binding problem (NFBP), a task traditionally attributed to neuronal networks. We also detail an inhibitory plasticity mechanism, critical for dendritic compartmentalization, further enhancing computational efficiency in dendrites. This *in silico* study underscores the computational capacity of individual neurons, extending our understanding of neuronal processing and the brain's ability to perform complex computations.

Introduction

Classically, single neurons in the nervous system have been thought to operate as simple linear integrators where the nonlinearity of dendrites can be neglected (McCulloch and Pitts, 1943). Based on this simplification, powerful artificial neural systems have been created, outperforming humans on multiple tasks (Silver *et al.*, 2018). However, in recent decades it has been shown that active dendritic properties participate in shaping neuronal output. Not only the axon spikes, but also the dendrites can display nonlinear integration of input signals (Antic *et al.*, 2010). Dendritic nonlinearities endow a neuron with the ability to perform sophisticated dendritic computations, expanding its computational power beyond what is available with the somatic voltage threshold and making it similar to a multilayer artificial neural network (Poirazi, Brannon and Mel, 2003).

A dendritic nonlinearity common among projection neurons in several brain areas is the NMDA-dependent plateau potential (Oikonomou *et al.*, 2014). Plateau potentials are regenerative, all-or-none, supralinear voltage elevations triggered by spatio-temporally clustered glutamatergic input (Schiller *et al.*, 2000; Polsky, Mel and Schiller, 2004; Losonczy and Magee, 2006; Major *et al.*, 2008; Larkum *et al.*, 2009; Lavzin *et al.*, 2012; Xu *et al.*, 2012). Such plateaus require that nearby spines are coactivated, but the requirement is perhaps somewhat loose as even single branches have been proposed to act as computational units (Losonczy and Magee, 2006; Branco and Häusser, 2010). With that said, multiple dendritic regions, preferentially responsive to different input values or features, are known to form with close dendritic proximity (Jia *et al.*, 2010; Chen *et al.*, 2011; Varga *et al.*, 2011). Such functional synaptic clusters are present in multiple species, developmental stages and brain regions (Kleindienst *et al.*, 2011; Takahashi *et al.*, 2012; Winnubst *et al.*, 2015; Wilson *et al.*, 2016; Iacuruso, Gasler and Hofer, 2017; Scholl, Wilson and Fitzpatrick, 2017; Niculescu *et al.*, 2018; Kerlin *et al.*, 2019; Ju *et al.*, 2020). Hence, multiple features are commonly clustered in a single dendritic branch, indicating that this could be the neural substrate where combinations of simple features into more complex items occur.

Combinations of features in dendritic branches further provide single neurons with the possibility to solve linearly non-separable tasks, such as the nonlinear feature binding problem (NFBP) (Tran-Van-Minh *et al.*, 2015; Gidon *et al.*, 2020). In the most elementary form, the NFBP consists of discriminating between two groups of feature combinations. The problem is nonlinear since the neuron should learn to respond only to specific feature combinations, even though all features are represented by the same amount of synaptic input. Common example features used are two different shapes combined with two different colors, giving in total four combinations of which the neuron should respond to only two feature combinations (exemplified in Fig. 1A and 1B).

As a task, the NFBP is relevant to brain regions which perform integration of multimodal input signals, or signals representing different features of the same modality (Roskies, 1999). It is usually illustrated with examples from the visual system, as in Fig. 1A (Roskies, 1999; von der Malsburg, 1999; Tran-Van-Minh *et al.*, 2015). A region that integrates multimodal inputs, such as sensory information and motor-related signals, is the input nucleus of the basal ganglia, the striatum (Reig and Silberberg, 2014; Johansson and Silberberg, 2020), and this system will be used in the present modeling study. We will here, however, continue to illustrate the NFBP with the more intuitive features borrowed from the visual field, although for dorsal striatum these features would rather map onto different sensory- and motor-related features. Plateau potentials and some clustering of input have been demonstrated in SPNs (Plotkin, Day and Surmeier, 2011; Oikonomou *et al.*, 2014; Du *et al.*, 2017; Hwang *et al.*, 2022; Day *et al.*, 2024; Sanabria *et al.*, 2024).

In addition to integrating convergent input from the cortex and the thalamus, the striatum is densely innervated by midbrain dopaminergic neurons which carry information about rewarding stimuli (Schultz, 2007; Matsuda *et al.*, 2009; Surmeier *et al.*, 2010). As such, the striatum is thought to be an important site of reward learning, associating actions with outcomes based on neuromodulatory cues. In this classical framework, peaks in dopamine (Da) signify rewarding outcomes and pauses in dopamine represent omission of expected rewards (Schultz, Dayan and Montague, 1997). Dopamine signals further control the synaptic plasticity of corticostriatal synapses on the SPNs (Fig. 1C). In direct pathway SPNs (dSPN) expressing the D₁ receptor, a dopamine peak stimulates the D₁ receptor, which together with significant calcium influx through NMDA receptors triggers synaptic strengthening (long-term potentiation - LTP). Conversely, when little/no dopamine is bound to the D₁ receptors, as during a dopamine pause, and there is significant calcium influx through L-type calcium channels, synaptic weakening occurs (long-term depression - LTD) (Shen *et al.*, 2008; Fino *et al.*, 2010; Plotkin *et al.*, 2013) (see Fig. 1D).

If dopamine peaks are associated with the relevant feature combinations in the NFBP and dopamine pauses with the irrelevant ones, it can trigger LTP in synapses representing the relevant feature combinations and LTD in those representing irrelevant combinations. If, after learning, the relevant feature combinations have strong enough synapses so they can evoke plateau potentials while the irrelevant feature combinations have weak enough synapses so they don't evoke plateaus, the outcome of this learning process should be a synaptic arrangement that could solve the NFBP (Fig. 1D) (Tran-Van-Minh *et al.*, 2015). In line with this, it has been demonstrated that the NFBP can be solved in abstract neuron models where the soma and dendrites are represented by single electrical compartments and where neuronal firing and plateau potentials are phenomenologically represented by

instantaneous firing rate functions (Legenstein and Maass, 2011; Schiess, Urbanczik and Senn, 2016). Good performance on the NFBP has also been demonstrated with biologically detailed models (Bicknell and Häusser, 2021). This solution used a multicompartmental model of a single pyramidal neuron, including both excitatory and inhibitory synapses and supralinear NMDA depolarizations. Synapses representing different features were randomly dispersed throughout the dendrites and a phenomenological learning rule - dependent on somatic spike timing and high local dendritic voltage - were used to optimize the strength of the synapses. The solution did, however, depend on a form of supervised learning as somatic current injections were used to raise the spiking probability of the relevant feature combinations.

In this article we develop a local, calcium-dependent synaptic learning rule based on the known synaptic machinery of the corticostriatal synapse onto dSPNs (Shen *et al.*, 2008), which, guided by the reward signals from dopaminergic neurons, can solve the NFBP in a multicompartment model of a dSPN. The learning rule is initially tested assuming pre-existing clustered synapses for each individual feature, potentiating synaptic clusters to the point where they can evoke robust plateau potentials for the relevant feature combinations of the NFBP, and weakening synapses representing irrelevant features. We also apply the learning rule on more randomly distributed synapses, and results suggest that branch-specific plasticity might be important for the single neuron to solve nonlinear problems. Although brain systems integrating multimodal inputs, such as the striatum, somehow solve nonlinear problems at the network/systems level, it is not known whether any individual neuron in the brain actually solves the NFBP on a regular basis. Our investigation suggests, however, that single SPNs have the computational capacity to solve linearly non-separable tasks by utilizing information of the organism's success (here represented by a dopamine peak) and failures (here a dopamine pause). This might also generalize to other projection neurons that can display dendritic plateaus, such as pyramidal neurons. Depending on neuron and synapse type, however, the specific feedback mechanisms (which in this case is dopamine) will have to be mapped to another neuromodulatory signal.

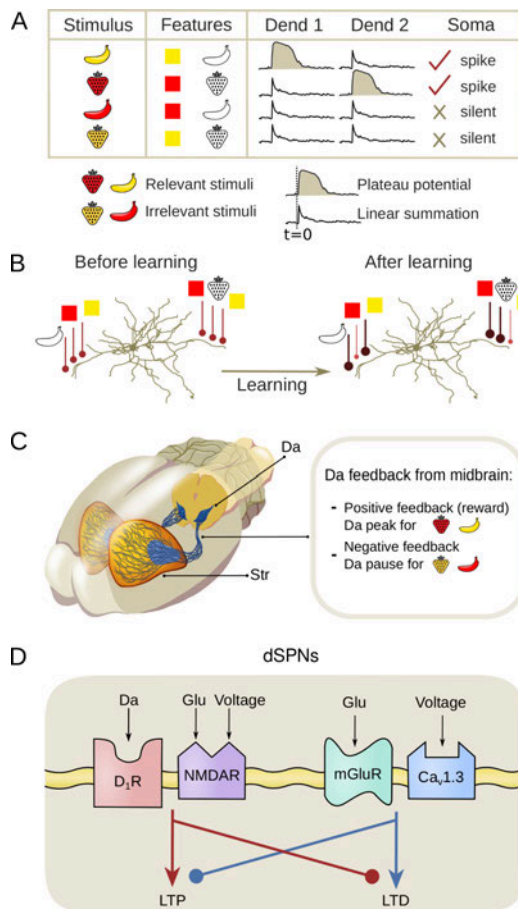


Figure 1: Learning Mechanisms in direct pathway Striatal Projection Neurons (dSPNs) for the Nonlinear Feature Binding Problem (NFBP)

A: Inputs and assumed supralinearity that could solve the NFBP: The NFBP is represented with an example from visual feature binding. In the simplest form of the NFBP, a stimulus has two features, here shape and form, each with two possible values, strawberry and banana, and red and yellow, respectively. The NFBP consists of responding with neuronal spiking to two of the feature combinations, corresponding to the relevant stimuli (red strawberry and yellow banana), and remaining silent for the other two feature combinations which represent the irrelevant stimuli (yellow strawberry and red banana). Assuming that each feature is represented with locally clustered synapses, a solution of the NFBP can be achieved when the co-active clusters on a single dendrite, representing the features of a relevant stimulus, evoke a plateau potential, thus supralinearly exciting the soma. Conversely, co-activation of synaptic clusters for the irrelevant combinations should not evoke plateau potentials.

B: Dendritic Learning: Illustration of how synaptic plasticity in SPNs may contribute to solving the NFBP for a pre-existing arrangement of synaptic clusters on two dendrites. A plasticity rule which strengthens only synaptic clusters representing relevant feature combinations, so that they produce robust supralinear responses, while weakening synapses activated by irrelevant feature combinations, could solve the NFBP.

C: Dopamine (Da) Feedback in Learning: dopaminergic feedback from the midbrain to the striatum (Str) guides the learning process, differentiating between positive feedback for relevant stimuli and negative feedback for irrelevant stimuli. Positive feedback represented by dopamine peaks is necessary for LTP, and negative feedback represented by a dopaminergic pause is necessary for LTD.

D: Signaling pathways underlying synaptic plasticity in dSPNs: Illustrations of molecular components at the corticostriatal synapse that modify synaptic strength (redrawn from Shen et al., 2009). NMDA calcium influx, followed by stimulation of D₁ dopamine receptors (D₁Rs), triggers LTP (while inhibiting the LTD cascade). L-type

calcium influx and activation of metabotropic glutamate receptors (mGluRs) when D₁Rs are free of D₁ triggers LTD (while counteracting the LTP cascade).

Results

Characterization of the dendritic NMDA-dependent nonlinearities in the model

The nonlinear, sigmoidal voltage sensitivity of the NMDA receptors is a crucial element for formation of dendritic plateau potentials. We use a model for generating plateau potentials first presented in Gao *et al.*, (2021) and adjusted to SPNs (Trpevski *et al.*, 2023). To produce robust, all-or-none plateau potentials, glutamate spillover from the synaptic cleft that activates extrasynaptic NMDA receptors is included in the model. Glutamate spillover occurs when the total synaptic weight (normalized value) of the activated synapses reaches a threshold value (see Methods). The threshold value here is set to be equivalent to the total weight of 10 clustered synapses with weights of 0.2 each (weight 0.2 corresponds to 0.5 nS). Figure 2A shows the somatic membrane potential following synaptic activation of a cluster of synapses of increasing size and the corresponding local spine membrane potential (averaged over all spines in the cluster) as well as the NMDA and L-type calcium accumulated in a single spine (also averaged over all spines in the cluster). A plateau potential is generated when a critical level of total NMDA conductance in a dendritic segment is reached (accomplished here by the addition of more synapses in a cluster and by glutamate spillover). Reaching the spillover threshold produces a sudden and robust increase in NMDA conductance, caused by the activation of extrasynaptic NMDA receptors (where the clearance of glutamate is assumed to be slower), thus triggering an all-or-none plateau potential as investigated in Trpevski *et al.*, (2023).

The synaptic input to the neuron is provided through the activation of a cluster of synapses at the indicated location in Fig. 2B, and gives the voltage and calcium responses in Fig. 2A. Figure 2C shows the maximal amplitudes of the somatic voltage and of the NMDA and L-type calcium signals in the dendritic spines, averaged over 10 trials and over 10 dendrites, and shown for three different synaptic weights in the clusters. The “baseline” results in Fig. 2C are within the range of the initial synaptic weights of excitatory synapses in all remaining figures in the article, and thus illustrate a possible initial situation before learning. Stronger and weaker synapses require a smaller and a larger cluster to trigger a plateau potential, respectively. In the simulation using “strengthened” synapses the synaptic weights in the cluster are 40% greater than in the “baseline” case, and hence need fewer synapses to trigger plateau potentials. Conversely, with weaker synapses where weights are 25% smaller than the “baseline” case, more synapses are needed to evoke a plateau.

To summarize, the dSPN model exhibits the dendritic nonlinearities required for solving the NFBP. If a cluster of strengthened synapses can reliably generate robust plateau potentials, this increases the likelihood for somatic spiking compared to when a more gradual NMDA dependent nonlinearity occurs in the dendrite. This also means that the neuron can reliably spike following the activation of a set of synapses strengthened (see illustration in Fig. 1). Conversely, a cluster of weakened synapses will most likely not generate plateau potentials,

and thus the neuron will spike with much lower probability following activation of such a cluster.

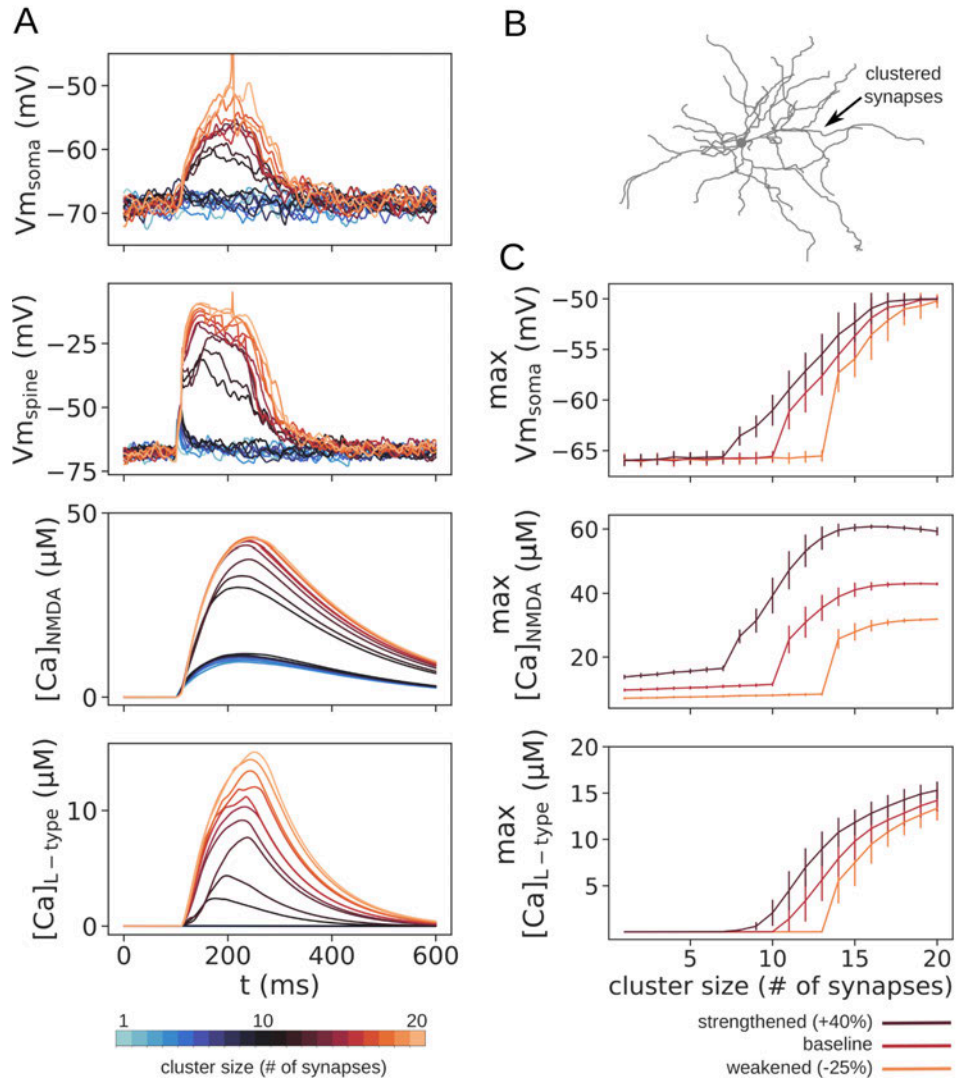


Figure 2: Characterization of dendritic plateau behavior in the model. A: Somatic voltage, spine voltage, NMDA calcium ($[Ca]_{NMDA}$), and L-type calcium ($[Ca]_{L-type}$) evoked by a cluster varying in size from 1 to 20 synapses. A plateau potential is evoked after a threshold level of NMDA conductance is exceeded, here set at 10 synapses with a weight of 0.2 each (corresponding to the "baseline" weights in C). The traces for spine voltage, $[Ca]_{NMDA}$, and $[Ca]_{L-type}$ are averages over all activated spines in the cluster.

B: Schematic of the neuron morphology with an arrow indicating the stimulated dendritic branch in A.

C: Maximal amplitude of the measures shown in A averaged over 10 dendritic locations. The curves represent clusters with different synaptic weights: baseline (0.2), strengthened (0.28) and weakened weights (0.15). Somatic voltages higher than the action potential threshold were set to -50 mV. A synaptic background noise is used in all simulations to elevate the membrane potential to ranges seen in vivo (Reig and Silberberg, 2014).

Characterization of the plasticity rule

To characterize the learning rule we start with a simple setup where three features (each illustrating either a color or a shape) are distributed onto two dendritic branches, so that one relevant and one irrelevant feature combination can be represented per dendrite (see Fig. 3A). Each feature is represented with 5 synapses, and we start with assuming that those synapses are already organized in pre-existing clusters. In addition to background synaptic noise inputs (compare Fig. 1), we also added 108 distributed glutamatergic synapses that were activated together with all four stimuli, i.e. they were feature-unspecific. SPNs have a very hyperpolarized resting potential, and the additional feature-unspecific synapses allow the neurons to spike often enough in the beginning of the learning process so that a dopamine feedback signal would be elicited and trigger learning in the activated synapses. The neuron model was then activated with a sequence of feature combinations (960 here) including equal amounts of relevant (i.e. 'red strawberry' and 'yellow banana') and irrelevant feature combinations (i.e. 'yellow strawberry' and 'red banana'). Thus, since the neuron always spikes initially due to the feature unspecific synapses, dopamine peaks and dopamine pauses arrive equally much. Fig. 3A shows an illustration of the setup. When the neuron spikes for the relevant feature combinations, dopamine rewards are delivered, triggering LTP in those active synapses with NMDA Ca levels within the LTP kernel, while spiking for the irrelevant feature combinations elicits a dopamine pause as feedback, instead triggering LTD as a function of L-type Ca. As learning progresses, the cell learns to respond only to the relevant feature combinations (see example in Fig. 3B). Initially all four stimuli, 'yellow banana' and 'red banana' in dendrite 1, and 'red strawberry' and 'yellow strawberry' in dendrite 2, elicited robust supralinear responses (as they together reached the threshold for glutamate spillover in our model). After learning the neuron could differentiate between the two sets of stimuli. The relevant feature combinations associated with a reward continued to provoke a plateau potential, eliciting somatic spiking. In contrast, the neuron's response to irrelevant feature combinations was notably decreased following LTD in the synapses on the dendrite where this feature is irrelevant.

Figure 3C shows the evolution of synaptic conductances during the learning process for dendrite 1 (dendrite 2 is not shown, but has the same behavior). The synapses representing the relevant feature combination in this dendrite ('yellow' and 'banana') are typically strengthened, eventually encoding this stimulus robustly. Conversely, the synapses for the feature 'red', activated during the irrelevant feature combination ('red banana') are all weakened, making the dendrite only weakly responsive to this stimulus following learning. Note that LTD also occurs in some of the synapses representing the features yellow as well as 'banana' as these features are also parts of the irrelevant stimuli (yellow strawberry and 'red banana', respectively). Nevertheless, the reward process still causes the remaining yellow synapses to strengthen, albeit to a lower level than the synapses for the feature 'banana' due to the LTD that has occurred (when combined with strawberry). This is because of the learning rule that successively slides the LTP NMDA Ca dependent plasticity kernel over training (see Methods). This means that our learning rule tends to stabilize the number of synapses that are needed to perform the task, but not necessarily all the synapses carrying the relevant features. Depending on the initial local Ca response in the synapse to relevant and irrelevant stimuli, individual synapses might either be preferentially recruited into the LTP or LTD process.

The NMDA calcium levels shown in Fig. 3D, left panel, shows examples on how the different synapses stabilize at their particular conductances/weights (for three example synapses marked with arrows in 3C). With repeated dopamine rewards, the metaplasticity kernel moves towards lower calcium levels, while a dopamine pause moves the kernel to higher calcium levels (i.e. it becomes easier for a synapse to strengthen for a slightly lower local calcium level if a reward has been seen often and also if the synapse is not being activated too often during reward omissions). Because the rate of metaplasticity kernel adaptation towards higher calcium levels is faster when the neuron makes errors (omitted rewards) on the NFBP, the kernels for all features move higher. This results in LTD in the 'red' synapses on the dendrite, while LTP develops in most of the 'yellow' and 'banana' synapses in dendrite 1. The upward shift in the kernel serves to prevent excessive strengthening in the weakened 'red' synapses on dendrite 1 when rewards for 'red strawberry' arrive (due to learning in dendrite 2), as they no longer fall within the optimal calcium range for LTP. In contrast to this, most synapses for 'yellow' and 'banana' are strengthened initially, but as their LTP kernels move downwards due to rewards and a sufficiently high NMDA calcium locally, no further LTP occurs. This means that in our learning rule, synapses tend to stop changing their weights when the neuron can perform the NFBP.

The feature-unspecific synapses exhibit different types of behaviors, which are primarily determined by the initial positioning of their calcium levels with respect to the optimal learning zone of the LTP plasticity kernel (Fig. 3C and 3D, right panels). The common trend of gradual weakening of these synapses, triggered by co-activation with irrelevant stimuli and subsequent dopaminergic pauses, is due to the metaplasticity kernel shifting away from the calcium level necessary for LTP. This shift increases the likelihood of LTD (initially during the training this occurs half of the time) This general weakening is exemplified by the purple trace in the right panel of Fig. 3D.

However, not all feature-unspecific synapses follow this trend. Some synapses, like the one depicted in blue, are situated close enough to the input clusters to experience increased NMDA Ca levels, enabling them to strengthen their synaptic connections. Finally, there are also a few feature-unspecific synapses that randomly happen to trigger calcium levels within the optimal learning region of the LTP plasticity kernel, causing enhanced LTP at the beginning (Fig. 3D, right panel, green trace). However, as the LTD amplitude is proportional to the L-type calcium level, the LTD-drive eventually becomes too high, pushing the calcium level below the LTP-kernel from which a period of LTD-only follows until synapses stabilize. To summarize, initially all synapses, clustered and feature-unspecific, may experience some LTP as dopamine rewards are delivered initially half of the time. However, as omitted rewards occurring for irrelevant stimuli cause the LTP metaplasticity kernel to slide higher, only clustered synapses signaling relevant feature combinations (and synapses close to them), where supralinear voltage effects cause high NMDA calcium, remain in the optimal learning region of the LTP kernel and a few undergo LTP before stabilizing. Note that learning is always "on", i.e. weight updates happen for each training example, meaning that there are no separate training and testing phases. The reason why synapse weights stop changing and stabilize is because the neuron performance improves so fewer and fewer omitted rewards occur (and this halts the LTD process) while at the same time the LTP kernel slides downwards when the neuron responds correctly most of the time and thus receives a reward (and this halts the LTP process).

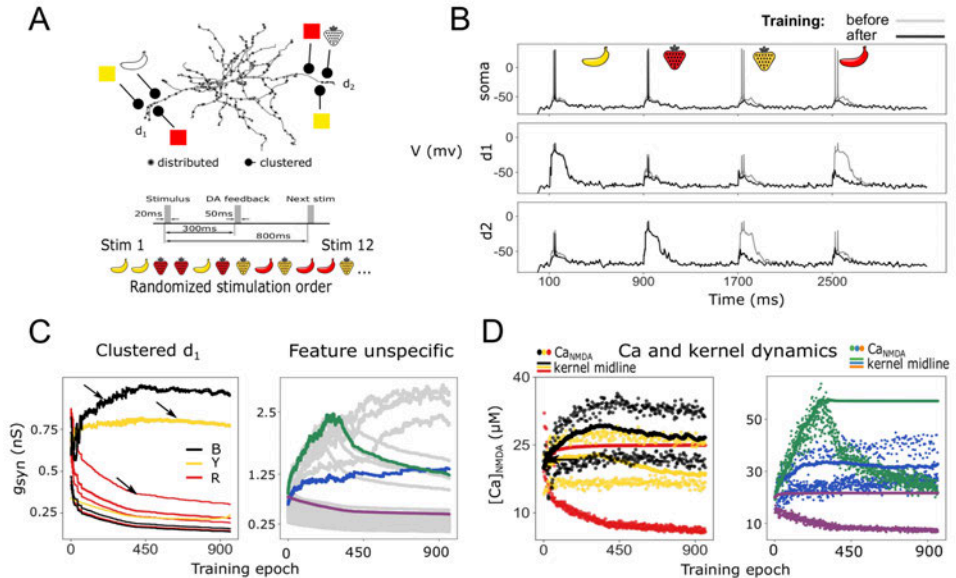


Figure 3: Example of setup and learning-induced synaptic plasticity A: Illustration of input configuration. Upper panel shows the arrangement of the features in two dendrites. Each dendrite has synaptic clusters for three features allowing the representation of each of the four feature combinations if seen from the whole neuron's perspective. Also, the feature combinations allow for only one relevant feature combination per dendrite. The middle panel illustrates the stimulation protocol used in the simulation where stimuli presentation is followed by a dopaminergic feedback signal only if the neuron spikes. The two features representing a stimulus are active within 20 ms, and dopamine feedback, lasting for 50 ms, is delivered 300 ms after the beginning of the stimulus. Two stimuli are spaced 800 ms apart, to allow for the calcium dynamics to reach baseline levels. The bottom panel illustrates the stimulus sequence over the full learning task: all stimuli are equally present in a sequence of 12 stimuli.

B: Example voltage in the soma and the two dendrites before and after learning. Each dendrite stops responding to the irrelevant feature represented by its synaptic clusters.

C: Evolution of synaptic conductances throughout learning. The left panel shows the conductances of the clustered synapses in one of the dendrites (dendrite 1, d_1). B, Y and R stand for 'banana', 'yellow' and 'red', respectively. The right panel shows the distributed feature-unspecific synapses. The initial synaptic conductances are set to 0.25 ± 0.05 (around 0.625 nS). The purple trace exemplifies a synapse that is weakened, while the green and blue traces exemplify synapses that are close to the clusters (blue) or by chance (green) have a sufficiently high local NMDA calcium level for LTP to dominate.

D: Peak calcium (dots) and plasticity kernel dynamics (solid lines) during learning. The left panel shows the NMDA calcium of a single 'banana' and a single 'yellow' synapse that undergo LTP as well as for a single red feature synapse undergoing LTD in dendrite 1 (marked with arrows in the left panel in C). The right panel shows the NMDA calcium for the feature-unspecific synapses identified in the right panel of C.

Characterization of input combinations that can be learned

After showing that the SPN can learn to separate relevant from irrelevant stimuli in the single example in Fig. 3, we next generalized the setup by giving different innervations of the four features to the two dendrites and recorded the performance of the SPN on the NFBP as learning progressed. Out of all possible feature innervations to two dendrites, we used only the innervations containing both relevant feature combinations arranged at least one per dendrite, as these are the cases with enough feature innervation to possibly solve the NFBP (for an illustration see Fig. 4A). We then estimated how the performance developed over

training, i.e. whether the SPN spiked for the right feature combinations and was silent for the irrelevant ones. Performance of 100% indicates that the neuron spikes only for the relevant stimuli and is silent for the irrelevant stimuli. Performance of 50% can indicate two situations: i) either the neuron spikes for all four stimuli, or ii) is silent for all four stimuli. Performance of 75% indicates spiking behavior between the following two cases: i) the SPN spikes for only one of the relevant feature combinations, remaining silent for the other three, or ii) it is silent for one of the irrelevant feature combinations while spiking for the other three. In this study we consider the NFBP solved when the performance is greater than 87.5%, exemplified by the situation where the SPN always spikes for one relevant stimulus, and at least half of the time for the other relevant stimulus, while remaining silent for the irrelevant stimuli. In the cases with two or three features innervating one dendrite (where one relevant feature combination is present per dendrite) the mean performance is above 90%, indicating that both stimuli are learned (Fig. 4B, purple traces). In the cases when all four features can innervate at least one of the dendrites, typically only one relevant feature combination (i.e. 'red strawberry' or 'yellow banana') could be correctly learned (Fig. 4B, light blue traces). In this case the same feature combination is usually encoded in both dendrites. In Fig. 4C the performance over the last 160 training examples of the simulation is shown, for each of the four stimuli separately. Fig. 4C₁ shows the performance for the purple traces in Fig. 4B, while Fig. 4C₂ shows the same for the light blue traces in Fig. 4B. This illustrates that for the purple traces in Fig. 4B, the NFBP is solved by the neuron, and the mistakes are usually made by firing for some of the irrelevant feature combinations. For the light blue traces in Fig. 4B, the NFBP is usually not solved, with somatic spiking for the relevant stimuli occurring only around 75% of the time. Additionally, spiking for the irrelevant stimuli is increased.

Optimal learning is achieved at intermediate distances from soma through excitatory plasticity

We also investigated the impact of synaptic positioning on learning (Fig. 4 D) when using the same settings as in Fig. 3, but varying dendritic locations. Our results predict that the best performance on the NFBP is obtained with synaptic clusters positioned at intermediate somatic distances from the soma (Fig. 4E). From Fig. 4E one can infer that after learning the proximal synapses have actually decreased when trained on the NFBP as the performance is around 50%. The neuron stays silent for the irrelevant feature combinations (blue dots), which is the correct response, but also the neuron can't respond to the relevant feature combinations (red dots). For successively more distal synapses the performance increases and then slightly decreases for the most distal clusters that sometimes fail to evoke somatic spiking for the correct feature combinations (red dots). This result can be conceptually explained in the following way. The electrotonic properties of dendrites dictate that synapses near the soma, in the most proximal regions, are less capable of inducing supralinear potentiation underlying plateau potentials (Du *et al.*, 2017). This is due to the soma acting as a current sink, resulting in smaller localized voltage changes (and hence a lower input resistance in accordance with Ohm's law). Consequently, these synapses cannot easily evoke dendritic nonlinearities necessary for solving the NFBP, and hence the performance with proximal clusters is low for NFBP. Note that in our simulations we allow glutamatergic synapses on spines quite close to the soma, although SPN dendritic spines are relatively rare at more proximal distances than 40-50 μm from the soma (Wilson *et al.*, 1983).

In contrast, the most distal dendritic regions are electrically more isolated and have a higher local input resistance, enabling larger voltage changes locally and thus also higher local calcium concentrations when synapses are activated in our model. This allows even a small number of active synapses to generate local supralinear NMDA-dependent responses. Such ease of elevating the local calcium, seems advantageous, but in fact results in decreased performance on the NFBP for the following reasons. In the distal synaptic clusters related to irrelevant stimuli, local calcium levels might stay in the optimal learning region of the LTP kernel despite omitted rewards occurring regularly and thus the weakening of the synapses takes longer. This causes increased somatic stimulation for irrelevant stimuli, sometimes leading to spiking over a longer period during the training. In addition, distally evoked dendritic signals, including plateau potentials from strengthened synaptic clusters, naturally attenuate more before they reach the soma, and sometimes fail to elicit somatic spiking. As a consequence, compared to more optimally placed clusters, the proportion of negative feedback signals is increased and that of positive feedback is decreased. The altered feedback results in a decreased capacity to selectively reinforce the active synapses to only relevant stimuli and weaken synapses activated with irrelevant stimuli.

The ideal learning zone for NFBP in our simulations thus lies at an intermediate somatic distance, roughly 100-150 μm from the soma, where synapses can effectively contribute to learning the NFBP (Fig. 4 E). In this zone synaptic changes are more likely to impact the neuron's firing probability as the dendritic plateau potential at this location causes a larger elevation of the somatic potential, and thus synapses at this distance benefit more from the dopamine feedback loop. Note that the prediction that proximal and very distal synapses are less likely to contribute to the solving of the NFBP doesn't imply they are not important for more 'linear' learning contexts. For instance, if we had trained the neuron to only respond to one single stimulus, such as 'red strawberry', both proximal and very distal synapses representing that stimulus would of course be able to both strengthen or weaken as well as contribute to spiking of the neuron following learning.

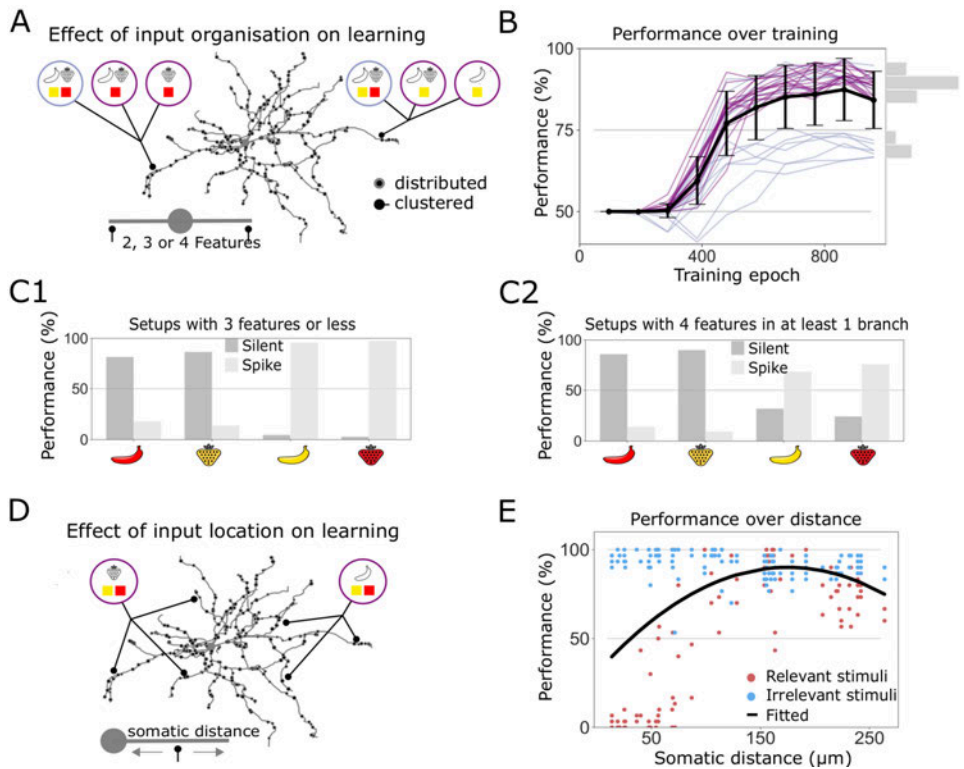


Figure 4: Impact of feature combinations and synaptic cluster locations on NFBP learning performance

A-C: NFBP configurational analysis. Configurations are categorized based on the number of features per dendrite, those with two or three features (purple traces) and those with all four features on at least one dendrite (light blue traces).

A: Illustration of the setup of the task where two, three or four features are given in two dendritic locations.

B: Performance trajectories for the combinatorial task illustrated in A. Traces of the performance of the model over time for 31 unique feature configurations. Light blue lines show combinations with four features in at least one dendrite, while purple lines show input with maximally 3 input in the local dendritic branches. Right-side histograms display the distribution of the end performance.

C: Outcome results split on stimuli for configurations where maximally three features in one dendritic location (C₁) or at least one dendrite has all four features present in a single dendrite (C₂).

D and E: End performance in a three-feature configuration as a function of cluster location. D illustrates the location of the inputs, and E shows end performance in the three-pattern configuration as a function of somatic distance of the synapse clusters. Initial synaptic weight is 0.25 ± 0.05 in all the simulation experiments.

The possible role of inhibitory plasticity in learning

In our initial simulations we assumed that only the excitatory synapses could undergo plasticity during learning, and we identified two critical observations that highlight possible areas for improvement. The first is a vulnerability to noise, which resulted in a performance around 90 percent, as illustrated in Fig. 4B. The second is decrease in performance observed across very distal synapses, as detailed in Fig. 4E. These findings prompted us to explore the potential role of inhibitory synapses in addressing these challenges.

We therefore developed a phenomenological inhibitory plasticity rule to enhance dendritic nonlinearities (detailed in the Methods section). The rule is designed to compartmentalize a dendrite so that it mostly responds to excitatory inputs that cause the strongest dendritic activation. It achieves this by weakening inhibitory connections at highly active excitatory synapses based on the local increase in voltage-dependent calcium, reinforcing their dominance, and strengthening them where excitatory activity is lower (and thus as well local calcium concentration, see Methods for details).

To demonstrate our inhibitory plasticity rule, we use the same excitatory synapse setup as in Fig. 4 to which we add four inhibitory synapses near each cluster in the middle of the dendritic branch, representing each of the four features (Fig. 5A). Thus, a single feature activates both the excitatory and inhibitory synapses. To achieve a level of depolarization and spike probability comparable to that in our excitatory-only setup, we increased the distributed synaptic input from 108 to 144. Alongside this, we began with low inhibitory synaptic weights. This was key to maintaining higher baseline activity in our model as starting with strong inhibitory weights could excessively suppress excitatory activity as e.g. inhibitory inputs close to clustered synapses effectively can counteract the NMDA-dependent nonlinearities (Doron *et al.*, 2017; Du *et al.*, 2017; Dorman, Jędrzejewska-Szmek and Blackwell, 2018). This setup ensured that we could clearly observe how learning modified synaptic connections without initial excessive inhibition preventing the dendritic nonlinearities initially.

As in the example in Fig. 3, the conductances of the excitatory inputs representing 'yellow' and 'banana' increase to cluster the 'yellow banana' pairing, while the weights of synapses representing the feature 'red' decrease (Fig. 5D, left panel). Conversely, the inhibitory synapses associated with the 'yellow' and 'banana' features, linked to excitatory synapses, are instead further weakened, while those linked to 'red' and 'strawberry' features in the same compartment are strengthened (Fig. 5D, right panel). This outcome of the inhibitory plasticity rule effectively prevents the 'red banana' and 'yellow strawberry' stimuli from triggering spikes due to the excitatory inputs to this particular dendrite, thereby compartmentalizing dendrite 1 to be responsive to mainly 'yellow banana'.

We also show the dynamics of peak calcium levels associated with both excitatory and inhibitory synapses for each task (Fig. 5E). For excitatory synapses, for which strengthening NMDA calcium influx is a key factor, the patterns in peak calcium levels behave as in the example in Fig. 3D (Fig. 5E, right panel). On the inhibitory side, the peak calcium levels, influenced by voltage-gated calcium channels as shown in the right panel of Fig. 5E, displayed different behavior, resulting from the inhibitory plasticity rule. The threshold levels of calcium for inhibitory synapses followed the highest excitatory synaptic activity, which in our task corresponded to the 'yellow banana' input. Conversely, the minimum threshold level was able to surpass all other synaptic activities, effectively designating it for 'yellow banana' only.

We also compared the performance for different feature configurations with added inhibitory synapses to the results for the excitatory-only setup from Fig. 4B. The results show not only that learning with inhibitory plasticity is faster, i.e. requires less training examples, but also achieves high performance, nearing 100%. This demonstrates the potential impact of

inhibitory synapses in enhancing learning capabilities, as well as in performing the NFBP computation (Fig. 5B).

Inhibitory plasticity additionally produced a marked improvement in performance when the location of the synaptic clusters was varied, especially prominent at very distal locations (Fig. 5C). Incorporating inhibitory plasticity allows for more rapid, robust, and accurate learning. Unlike excitatory plasticity, which relies on dopaminergic feedback signals, our inhibitory plasticity model follows a passive rule that just aligns with the local excitatory activity. It reinforces the most active excitatory synapses within a dendritic branch by decreasing the inhibitory synapses corresponding to the same features there, and conversely, it strengthens inhibitory synapses for features for which excitatory activity is less prominent. This 'winner-takes-all' strategy not only accelerated the learning process by requiring fewer training examples but also enhanced performance consistency across different dendritic locations, including the distal regions where excitatory-only models faltered to some extent when challenged with the NFBP. The inhibitory rule's ability to function without dopaminergic feedback streamlines the learning process, leading to higher accuracy and a more stable and robust set of synaptic modifications, regardless of the synaptic cluster's position along the dendrite. This demonstrates the influence inhibitory synapses could potentially exert in fine-tuning dendritic responsiveness and refining the neuronal circuitry critical for learning.

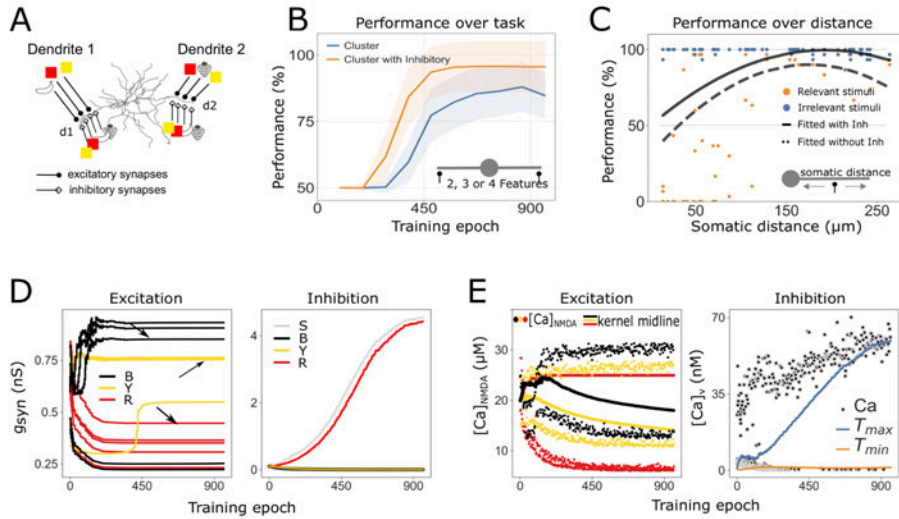


Figure 5: Effects of inhibitory inputs on performance

A: Dendritic input configuration with inhibitory synapses added. Illustrations depict two dendritic branches, each with synaptic connections from three excitatory and four inhibitory features, but the general setup is the same as in Fig. 4.

B: Displays average performance for configurations with varying pattern configurations as a comparison between the setup with (orange) and without inhibitory plasticity (blue).

C: Shows task-specific performances for dendritic locations of the clustered synapses, with individual dots and curves representing fitted performance curves with (solid) and without plastic inhibition (dashed).

D: Synaptic conductance changes during learning. Left panel shows excitatory synaptic conductances in dendrite 1 during learning. The right panel shows the inhibitory synaptic conductances. B, Y, R and S stands for 'banana', 'yellow', 'red' and 'strawberry'.

E: Peak calcium (dots) and plasticity thresholds dynamics (lines) over the learning. The left panel shows examples of the postsynaptic calcium amplitude of the clustered spines in dendrite 1 (marked with arrows on the left panel in D). The right panel tracks calcium levels at inhibitory synapses. The upper threshold (T_{max}) captures peak calcium levels while the lower threshold (T_{min}) identifies the next highest. Stability achieved at both thresholds enhances contrast between the activity levels and makes the dendrites respond preferentially to one relevant feature combination at that dendritic site.

Dendritic Branch Plasticity and Synaptic Learning – Investigating the Role of Randomly Distributing Excitatory and Inhibitory Synapses on the Dendrites

We finally challenged our plasticity rule by relaxing the assumption that single features are represented by pre-clustered synapses on specific dendritic branches. The synaptic plasticity rule was therefore investigated by randomly distributing 200 excitatory synapses signaling the different features across 30 dendrites. Each feature was represented by 40 excitatory synapses and an additional 40 feature-unspecific excitatory synapses were used. Figure 6A illustrates the setup and exemplifies the pre- and post-learning synaptic weights for both excitatory and inhibitory synapses. Our objective was to examine the learning dynamics in the absence of assumed synaptic clustering and to determine the capability of the single neuron to learn the NFBP. We initiated the experiment with synaptic conductance set to 0.3 ± 0.1 (around 0.75 nS), focusing first on purely excitatory synapses. This higher initial weight and increased variance, compared to the initial weights of 0.25 ± 0.05 (around 0.75 nS) with clustered synapses, were chosen to compensate for the reduced efficacy of

non-clustered synaptic inputs in producing sufficient depolarization and calcium influx, essential for effective learning and synaptic plasticity.

Figure 6B (right panel) shows limited learning capacity on the NFBP, which corresponded with our hypothesis: that without clustering-induced supralinearities due to spillover in the dendrites, distributed synapses would be too weak to trigger high enough differences in local calcium levels for separating the correct and incorrect feature combinations. To confirm that the lack of dendritic nonlinearities is the cause for not learning the NFBP, we extended this concept by introducing a hypothetical branch-specific spillover mechanism assuming that spillover occurs if synapses are on the same dendritic branch. This hypothetical construct assumes that concurrent activation of a critical mass of co-activated synapses in a single dendritic branch would trigger both plateau potentials and a significant rise in calcium levels despite that the synapses are not spatially pre-clustered, endowing the dendrites with supralinear responses even without closely clustered synapses. We also tested a more gradual spillover model proposing that each synapse could contribute incrementally to a cumulative spillover effect, instead of using a sharp glutamate threshold (see Trpevski et al, 2023). This modification provided a reduced level of nonlinearity compared to the branch-specific spillover model but still led to an increase in learning effectiveness. As shown in Fig. 6B, adding the hypothetical nonlinearities to the model increases the performance towards solving part of the NFBP, i.e. learning to respond to one relevant feature combination only. The performance increases with the amount of nonlinearity.

We next extended our investigation by including 60 inhibitory synapses, 15 for each feature, dispersed randomly over the 30 dendrites. Performance improved with the introduction of these additional nonlinearities, especially in combination with the branch-specific spillover mechanisms. This finding emphasizes the importance of dendritic branch plasticity and calcium nonlinearity in coordinating synaptic modifications for both excitatory and inhibitory inputs. In Fig. 6C, we report that a subset of neurons (5 out of 31) successfully solved the NFBP (reaching a performance of at least 87.5%).

For a more granular analysis of the successful cases, we selected an example that learned the NFBP. In Fig. 6D, we show the somatic and dendritic voltages for four dendritic branches where we found successful encodings of the relevant stimuli. The voltage traces show the somatic and dendritic responses to all four stimuli after learning. For each dendrite we also estimated the cumulative synaptic conductances for both excitatory and inhibitory synapses at the dendritic branch midpoint. Notably, in dendrites numbered 1 and 4, we observed an enhancement of excitatory inputs for the 'yellow banana' pattern and inhibitory inputs for the 'red strawberry' pattern. Conversely, dendrites 2 and 3 displayed the opposite arrangement.

These results show that for an adequate random innervation of distributed synapses, where the necessary features for a relevant stimulus innervate the same dendritic branch with enough synapses, that stimulus can be stored on that dendrite. In this way, the NFBP can also be learned if the two relevant stimuli are encoded on different dendrites and each of them can trigger a supralinear dendritic response. Since on average the two features representing a relevant stimulus do not innervate a single dendrite with enough synapses, the stimulus is not stored in a single dendrite, but is distributed across the dendritic tree. And since distributed synapses summate more linearly at the soma, only one of the relevant stimuli can be typically encoded by the dendritic tree.

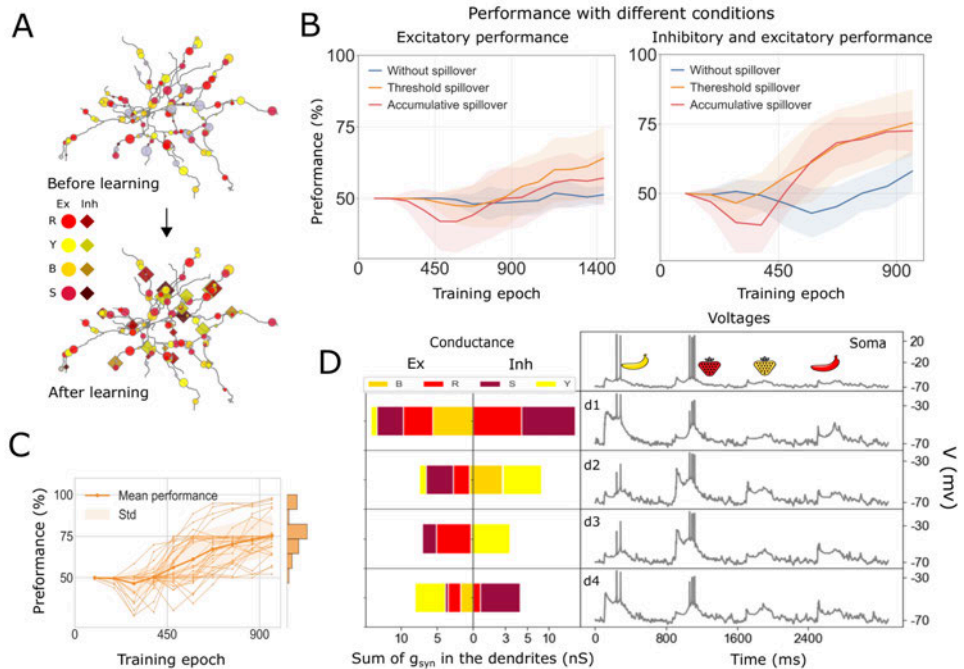


Figure 6: Performance analysis of learning using distributed synaptic inputs

A: Example illustration of synaptic distribution before (top) and after learning (bottom) of the 200 excitatory and 60 inhibitory inputs.

B: Learning performance with two spillover models, branch-specific thresholded and accumulative, without (left) and with plasticity of inhibitory synapses (right).

C: Performance trends of 31 distributions with branch-specific thresholded spillover and inhibitory plasticity.

D: Example of summed synaptic conductances (left) and voltage (right) in the soma, and four example dendrites (d1-d4) of one model following successful learning of the NFBP. The sums of both excitatory (Ex) and inhibitory (Inh) inputs are shown.

Discussion

In this article we studied whether single neurons can solve linearly non-separable computational tasks, represented by the NFBP, by using a biophysically detailed multicompartment dSPN model. Based on the synaptic machinery of corticostriatal synapses onto dSPNs, we propose a learning rule that uses local synaptic calcium concentration and dopamine feedback signals: rewards for relevant stimuli and omitted rewards for irrelevant stimuli. Assuming first that single features in the NFBP are represented by clustered synapses, we show that the learning rule can solve the NFBP by strengthening (or stabilizing) synaptic clusters for relevant stimuli and weakening clusters for the irrelevant stimuli. The feature combinations for the relevant stimuli, stored in strengthened synaptic clusters, trigger supralinear dendritic responses in the form of plateau potentials, which is an important ingredient for the solution of the NFBP, as plateaus significantly increase the likelihood of neuronal spiking in SPNs in a robust way (Du *et al.*, 2017).

The location of the synaptic clusters along the dendrites influenced the performance on the NFBP. In our model the region for optimal performance was 100-150 μm away from the soma, at about the same distance as where the somatic depolarization and induced spike probability, following activation of clustered input, is largest in an older version of the model (Lindroos and Hellgren Kotaleski, 2021). Clusters placed further away produce smaller somatic depolarizations, due to dendritic filtering (Major *et al.*, 2008), and as a consequence does not control the likelihood of somatic spiking as decisively. As the supralinear response is necessary for discriminating the relevant from the irrelevant stimuli, the performance with distally placed clusters decreases somewhat.

We further verified that supralinear dendritic responses were necessary to solve the NFBP by using randomly distributed synapses instead of clustered for each feature. In this scenario, only one relevant stimulus is sometimes learned in a dendritic branch by the randomly distributed synapses. The random setup was tested with two glutamate spillover models: The first model assumed that thresholded glutamate spillover is branch-specific, building on the notion that the single branch act as a single computational unit (Losonczy and Magee, 2006; Branco and Häusser, 2010). The second assumed that spillover is accumulative (Trpevski *et al.*, 2023). Both spillover versions increase the performance of our neuron model, further certifying the necessity of supralinear responses for solving the NFBP. In the real dendritic branches, however, also diffusion of signaling molecules within the branch likely contributes significantly both to the plasticity in already existing synapses as well as to local structural plasticity, both of which could increase branch specific supralinearities.

By using a phenomenological inhibitory plasticity rule based on the BCM formalism (Bienenstock, Cooper and Munro, 1982), we also show that inhibitory synapses can significantly improve performance on the NFBP. This is in line with earlier theoretical studies where negative synaptic weights were required to solve the NFBP (Schiess, Urbanczik and Senn, 2016). In our setup with pre-existing synaptic clusters, inhibitory synapses made learning faster and increased performance by inhibiting supralinear NMDA responses for the irrelevant stimuli. This was specifically true in distal dendrites where the input impedance is higher (Branco, Clark and Häusser, 2010). The threshold for plateau initiation is also lower in distal dendrites compared to proximal (Losonczy and Magee, 2006) which likely will further extend the influence of inhibition in this region (Doron *et al.*, 2017; Du *et al.*, 2017). Similarly, in the scenario with distributed synapses, inhibition enables one of the relevant stimuli to be reliably encoded in the dendritic branch by strengthening the inhibitory synapses for features different from those of the encoded stimulus. Together, it therefore seems like inhibition not only has a role in learning (Chen *et al.*, 2015; Cichon and Gan, 2015), but also improves the ability of the neuron to discriminate between stimuli with shared features.

Although our learning rule only occasionally solves the NFBP when used with randomly distributed synapses, it can always learn to perform a linearly separable task, such as learning to respond to only one relevant stimulus (such as red strawberry). Moreover, the learning rule is general enough so that in addition to the feature-specific inputs related to the task, it can handle feature-unspecific inputs that might or might not be related to the NFBP. Finally, the learning rule is always “on”, continuously updating synapses with each stimulus presentation, which is a more realistic mechanism compared to using separate training and

testing phases as in the field of machine learning. The synaptic weights automatically stabilize in the model when the performance improves. That is, rewards are then seen very regularly as the neuron has learned to spike for the relevant stimuli, while omitted rewards seldomly occur as the neuron stays silent when the irrelevant stimuli are provided.

When formulating the learning rule in this article, our goal was to base it on what is known regarding the synaptic machinery in corticostriatal synapses. This implied that the learning rule is based on the local calcium activity and on dopamine signals. Feedback from the dopamine system can also be viewed as an innate, evolutionarily encoded “supervisor”, which instructs neurons which feature combinations are beneficial and which ones should be avoided. However, in our case we do not use additional excitation to promote somatic spiking for only the relevant feature combinations, and in that sense the learning rule does not require a supervised learning paradigm. Since the SPNs rest at very hyperpolarized membrane potential, our setup includes distributed excitatory inputs which are feature-unspecific in order to make sure that the neuron spikes for all stimuli, especially at the beginning of training. These additional inputs are on average weakened as learning progresses (as they are activated for all stimuli and thus often receive negative feedback). That general or noisy inputs are reduced during learning is in line with the observed reduction of execution-variability during motor-learning as a novice becomes an expert (Kawai *et al.*, 2015). Also, the underlying neuronal representation of corticostriatal synapses undergo a similar change during learning (Santos *et al.*, 2015).

Since each synapse has its own calcium response, it is important for the learning rule to be able to follow individual synaptic activities. The LTP plasticity kernel in our model is for this reason itself plastic, meaning that it changes its calcium dependence as a function of the history of reward and punishment. This setup helps the model separate clustered synaptic input from the feature-unspecific input at the beginning of training, as only clustered synapses will see enough calcium to fall within the plastic range. We further use an asymmetric metaplasticity rule, where negative feedback causes a larger shift of the LTP kernel than a positive. This was necessary in order to stop LTP in synapses participating in LTD. Similarly to the classical loss-aversion tendency described in economic decision theory (Kahneman and Tversky, 1979), the model hence predicts that negative feedback will have a bigger impact in changing a well learned behavior on the single cell level than a positive feedback. Dopamine signaling has also been linked as a neural substrate to the decision making theory mentioned above (Stauffer *et al.*, 2016).

It is not known whether single neurons solve the NFBP or other linearly non-separable tasks. However, many brain nuclei receive convergent inputs from numerous other brain nuclei, acting as integratory hubs (van den Heuvel and Sporns, 2013). Since feature binding evidently occurs in the brain, and functional clusters for single features such as visual stimulus orientation, receptive fields, color, or sound intensity exist on single neurons (Chen *et al.*, 2011; Wilson *et al.*, 2016; Iacaruso, Gasler and Hofer, 2017; Scholl, Wilson and Fitzpatrick, 2017; Ju *et al.*, 2020), it is possible that the NFBP is a relevant task for neurons to solve. How brain regions with different synaptic machinery than the striatal dSPN might solve the NFBP remains a question, and reliance on other neuromodulatory signals may be part of the answer. For example, in the striatum, the indirect pathway SPNs (iSPN) have analogous synaptic machinery to the one in dSPNs, requiring calcium influx from the same sources for LTP and LTD, but are differently responsive to dopamine (Shen *et al.*, 2008). In

iSPNs a dopaminergic pause, together with a peak in adenosine, is required to trigger LTP, whereas a dopamine peak without peaks in adenosine rather promotes LTD (Shen *et al.*, 2008; Nair *et al.*, 2015). Therefore, we expect that an analogously formulated learning rule will also solve the NFBP in iSPNs, activating them for irrelevant feature combinations to e.g. suppress movement, and suppressing their activity for relevant feature combinations to facilitate movement. In addition, LTP in dSPNs might require co-activation of other neuromodulatory systems, such as a coincident acetylcholine pause with the dopamine peak, which we have not explicitly included in the model (Nair *et al.*, 2015; Bruce *et al.*, 2019; Reynolds *et al.*, 2022).

Method

In this paper we introduce a local, calcium- and reward-based synaptic learning rule, constrained by experimental findings, to investigate learning in SPNs. The learning rule operates based on the local calcium concentration in conjunction with plateau potentials and other dendritic nonlinearities and shows that such events enable learning of the nonlinear feature binding problem (NFBP). The learning rule is embedded in a biophysically detailed model of a dSPN built and simulated in the NEURON software, v8.2 (Carnevale *et al.*, 2006). Here we will focus on the setup of the learning rule and only give a short summary of the neuron and synapse models, focusing on changes compared to previously published versions. For a detailed description of the neuron model setup, see Lindroos *et al.*, (2018), Lindroos & Hellgren Kotaleski, (2021), and Trpevski *et al.*, (2023).

Neuron model

In short, the dSPN model used here was sourced from a collection of biophysically detailed models, including a reconstructed morphology and all of the most influential ion channels, including six calcium channels, each with its own voltage dependence and dendritic distribution (Lindroos and Hellgren Kotaleski, 2021). In accordance with Trpevski *et al.*, (2023), the model was further extended with synaptic spines on selected dendrites. Each spine was modeled as two additional compartments consisting of a neck and a head region, and contains voltage-gated calcium channels of types R ($\text{Ca}_v2.3$), T ($\text{Ca}_v3.2$ and $\text{Ca}_v3.3$), and L ($\text{Ca}_v1.2$ and $\text{Ca}_v1.3$); the addition of explicit spines did not change the basic behavior of the model, such as the response to current injections, etc.

Calcium sources used in learning

The intracellular calcium concentration is separated into distinct pools that are used during the learning process. For learning in glutamatergic synapses, one pool for NMDA-evoked calcium concentration ($[\text{Ca}]_{\text{NMDA}}$) is used, and another for L-type calcium concentration ($[\text{Ca}]_{\text{L-type}}$), to reflect the different synaptic plasticity responses of the SPN's biochemical machinery to these two calcium sources in the corticostriatal synapse (Shen *et al.*, 2008; Fino *et al.*, 2010; Plotkin *et al.*, 2013). Both pools are based on the calcium influx from the corresponding source (NMDA and the L-type channels $\text{Ca}_v1.2$ and $\text{Ca}_v1.3$, respectively). Similarly a third pool, used for inhibitory plasticity, collects the current from all voltage dependent Ca channels (T-, R-, L- and N-type, $[\text{Ca}]_v$). All pools include extrusion mechanisms in the form of a calcium pump as well as a one-dimensional time-decay. The calcium pump follows the implementation in Wolf *et al.*, (2005). The parameters for the

[Ca]_{NMDA} model were manually tuned to match the quantities reported in Dorman et al., (2018). The voltage gated calcium channel conductances in the spines were also manually tuned to match the relative calcium proportions in Carter and Sabatini, (2004) and Higley and Sabatini, (2010), as well as the calcium amplitudes due to stimulation with backpropagating action potentials (Shindou, Ochi-Shindou and Wickens, 2011). Spatial diffusion was not included in the present model.

Glutamatergic synaptic input

In our study, we used a synaptic model that includes both AMPA and NMDA conductances activated on spines. Additionally, we considered extrasynaptic NMDA conductances on dendritic shafts directly under the spines to account for glutamate spillover if the synapses were sufficiently stimulated and spatially clustered. Spillover has been shown to be important for generating robust all-or-none dendritic plateaus (Trpevski *et al.*, 2023).

The conductance of extrasynaptic NMDARs during spillover matches their synaptic counterparts. Glutamate spillover in a set of clustered synapses occurs in the model if a glutamate threshold is met, and in that case a presynaptic spike also activates the corresponding extrasynaptic NMDA conductance (Trpevski *et al.*, 2023). The threshold mechanism involves summing the synaptic weights for each NMDA synapse activated in the cluster, and checking whether their sum is greater than a threshold level. The synaptic weight is represented by a weight parameter (w) which scales the maximal conductance of the synapse (the change of the weights are described by Eq. 1-3). We exemplify this with the equation for calculating the NMDA current I_{NMDA} :

$$I_{NMDA}(t) = g_{syn}(t)(V(t) - E_{NMDA}); g_{syn}(t) = wg_{max}g(t)$$

where $V(t)$ is the membrane voltage, E_{NMDA} is the reversal potential of the NMDA synapse, g_{syn} is the NMDA conductance which has a time varying part, ranging from 0 to 1, scaled by a fixed part comprised of the maximal NMDA conductance and the synaptic weight. In our simulations a cluster needs to reach a summed weight of 2 to reach the glutamate threshold for spillover, at which point the extrasynaptic NMDARs are activated. This is modeled in NEURON by summing the NMDA synaptic weights upon each spike arrival, with the sum normalized so that e.g. ten synapses of 0.2 weight reach the threshold. An integrate-and-fire cell in NEURON generates a spike to extrasynaptic NMDARs when this threshold is met. This model allows for fewer synapses to reach the threshold as synaptic weight increases, a phenomenon linked to LTP and enhanced glutamate spillover due to astroglial process withdrawal (Henneberger *et al.*, 2020).

Extrasynaptic NMDA synapses were not included for feature-unspecific (non-clustered) synapses although we explored the role of a potential branch-specific spillover mechanism, where non-clustered synapses localized on the same branch still could interact through spillover mechanisms, either assuming thresholded or accumulative spillover as detailed in Trpevski *et al.*, (2023). The AMPA and NMDA synapse models were taken from Gao *et al.*, (2021) and are a variation of the saturating synapse models in Destexhe *et al.*, (1994). See Trpevski *et al.*, (2023) for a detailed description of how the synaptic input was set for the SPN model.

Learning rule

Excitatory synaptic plasticity

The learning rule is based on experimental findings showing that striatal LTP depends on NMDA-channel activation and the presence of dopamine, while LTD is dependent on activation of the L-type Ca channel $Ca_v1.3$ and the mGluR5 receptor in the absence of dopamine (low dopaminergic tone) (Shen *et al.*, 2008; Fino *et al.*, 2010; Plotkin *et al.*, 2013; Yagishita *et al.*, 2014; Fisher *et al.*, 2017; Shindou *et al.*, 2019). Under basal dopamine levels, no significant synaptic plasticity is assumed to occur. The rule describes a reward-based learning scheme in which a dopamine peak increases the synaptic weight based on one function (plasticity kernel) while a dopamine pause decreases the weights based on another function. Hence, dopamine acts as a switch selecting which of the LTP or LTD pathways should be activated (as illustrated in Fig. 1D). The plasticity kernels for LTP and LTD depend on NMDA calcium and L-type calcium, respectively, and are described below.

LTP

The LTP process is triggered by increased dopamine levels and is based on calcium influx through the NMDA channel. Here the synaptic strength is updated based on a sliding bell-shaped kernel in such a way that the maximal increase is obtained for peak calcium levels close to the kernel midpoint, while higher or lower peak calcium levels result in a smaller increase. This sets an upper limit on the synaptic strength and guarantees that the weights will not grow beyond a certain limit (Zenke and Gerstner, 2017; Zenke, Gerstner and Ganguli, 2017). Since each synapse has its own calcium level, the sliding of the bell-shaped kernel enables the precise tuning of each synaptic weight separately (see Fig. 7A for an illustration).

Mathematically, the bell-shaped kernel for the LTP process is represented by the derivative of the sigmoid function σ (Eqs. 1 and 2) with a slope that can be adjusted through the β parameter. The LTP learning rule (Eq. 3) describes the increase in synaptic weight (w) of the inputs, as a function of peak calcium level ($[Ca]_{NMDA}$), the learning rate (η_1), and the midpoint (θ_{ltp}) and slope of the sigmoid curve (β_1).

$$\sigma(Ca, \theta, \beta) = \frac{1}{1 + e^{-\beta(Ca - \theta)}} \quad \text{Eq. 1}$$

$$\sigma'(Ca, \theta, \beta) = \beta \cdot \sigma(Ca, \theta, \beta) (1 - \sigma(Ca, \theta, \beta)) \quad \text{Eq. 2}$$

$$dw_{ltp} = \eta_1 \cdot \sigma'(Ca_{NMDA}, \theta_{ltp}, \beta_1) \quad \text{Eq. 3}$$

LTD

The LTD process is triggered by a dopamine pause and is dependent on L-type calcium. The LTD plasticity rule (Eq. 4), describes a threshold level of the calcium level necessary for LTD

to occur (Shindou, Ochi-Shindou and Wickens, 2011), after which the decrease in synaptic weight is linearly proportional to the amplitude of the peak calcium level, and is scaled by the learning rate (η_2 , see left panel of Fig. 7B for an illustration). The peak calcium threshold is implemented with a sigmoid function, whose slope parameter, β_2 , was set to a high-value to make the curve resemble a step function. The midpoint, θ_{ltp} , of the LTP kernel was also increased during LTD, as described in the next section (Fig. 7B right panel).

$$dw_{ltd} = \eta_2 \cdot Ca_{v1.3} \cdot \sigma(Ca_{v1.3}, \theta_{ltd}, \beta_2) \cdot w_{ltd} \quad \text{Eq. 4}$$

Metaplasticity

Metaplasticity is a form of regulatory mechanism changing the state of the synapses in such a way as to influence subsequent learning. In our model we implemented this as a reward-history-dependent change of the location of the calcium range over which the LTP kernel was operating - a pause in dopamine triggered a shift towards higher calcium concentrations while a peak in dopamine pushed the LTP kernel in the opposite direction. The metaplasticity was also implemented using a kernel with the same midpoint as the LTP kernel, but with a wider calcium dependence. Together the setup with the two dynamically moving kernels allowed for a wide range of calcium levels to induce plasticity. This setup typically gave rise to a few characteristic situations:

1. For peaks in dopamine (see Fig. 7C for an illustration),
 - a. In synapses with low calcium levels, the kernel will be shifted closer to the observed level and thereby eventually enable LTP in synapses that are regularly activated during rewards, despite that they initially don't generate big changes in calcium.
 - b. In synapses with already high calcium levels, the kernel will be shifted away from the observed level, and thereby protect the neuron from excessive LTP and instead stabilize the weight (as illustrated with the red circle moving to the blue one in Fig. 7C).
2. For dopaminergic pauses, the kernel will be shifted away from the observed level and thereby reducing the likelihood of LTP in synapses that are undergoing LTD (see Fig. 7D for an illustration).

The LTP kernel update was further asymmetric with regards to dopamine peaks and pauses in such a way that a pause caused a larger shift of the kernel than a peak. This further reduced the likelihood of inducing LTP in synapses often participating in LTD or in synapses randomly activated with regard to the dopamine feedback signal. The metaplasticity kernel thus adjusts the LTP kernel for these processes based on ongoing neural activity. These changes are described by Eq. 5, where the LTP kernel midpoint is updated based on the learning rate (η_s), the peak NMDA calcium level ($[Ca]_{NMDA}$), and the slope of the sigmoid curve (β_3). The learning rate, η_s , captures the described asymmetry in response to dopamine feedback (see Table 1). This setup can give rise to the following situations:

1. For randomly alternating dopamine peaks and pauses,
 - a. In a few synapses that initially happen to generate very high calcium levels the synapses can become even stronger as their postsynaptic calcium response is still within the shifted LTP kernel. This explains e.g. the increase in strengths of the synapses illustrated with the blue and green traces in the

right panel of Fig. 3C. (However, as the computational performance of the neuron successively improves these synapses stabilize when the neuron starts to receive less dopamine pauses).

- b. In synapses that generate somewhat less postsynaptic calcium responses the general trend is that they decrease as their calcium level becomes lower than the LTP kernel. This is illustrated in Fig. 7D with the red circle shifting to the blue one; also e.g. compare the purple trace in the right panel of Fig. 3C.

In summary, the slower shifting metaplasticity kernel, compared to the weight update, allows the individual synapses to stabilize when the neuron has learned a specific task and thus doesn't receive negative reward feedback that often. Also, the rule allows initially weaker synapses to be recruited to undergo LTP if they are more consistently co-active with a reward than with an omitted reward (an illustration of this can be seen in Fig. 5D, where one of the yellow synapse traces jumps upwards during the middle of the training session) .

$$d\theta_{ltp} = \eta_s \cdot \sigma'(Ca_{NMDA}, \theta_{ltp}, \beta_3) \quad \text{Eq. 5}$$

Inhibitory synaptic plasticity

In contrast to the well studied mechanistic underpinnings of glutamatergic plasticity in e.g. SPNs, much less is known about how inhibitory synapses might be updated during learning. The inhibitory plasticity rule developed here is therefore more phenomenological and exploratory in nature and was developed to enhance nonlinearities in the local dendrite. The rule is based on the Bienenstock-Cooper-Munro (BCM) formalism (Bienenstock, Cooper and Munro, 1982). In the BCM rule there is a threshold level of synaptic activity below which LTD is triggered and above which LTP is triggered (Fig. 7E). Inspired by (Gandolfi *et al.*, 2020; Ravasenga *et al.*, 2022) and, we use calcium from all voltage-gated calcium channels ($[Ca]_v$) as the indicator of excitatory synaptic activity near the dendritic shaft where an inhibitory synapse is located. The inhibitory rule is designed to passively observe and respond to the surrounding excitatory synaptic activity, governed by local calcium influx, but without reliance on dopamine or explicit feedback. It operates at a slower pace to ensure alignment with the excitatory activity levels, thus enhancing the contrast in activity by amplifying local differences in excitatory synaptic efficacy. The weight change is further dynamic and dependent on the weight itself, in such a way that large and small weights get small updates. This stabilizes the weight at the end of learning in the inhibitory synapses, together with the dynamical threshold level meeting the local calcium level, and prevents the weight from taking on large or negative values. Large weights are then close to w_{inh}^{max} and small weights close to 0 (see Fig. 5D).

The rule is described by Eq. 6-9 where Eq. 6 describes the basic plasticity curve used in Eq. 7 to control the change in inhibitory weights (dw_{inh}). In these equations, the parameters a and b set the activity values for the two terms in Eq. 6 while β and T control the steepness and intersection with the x-axis.

The parameter c recalibrates the baseline for T_{min} , ensuring T_{min} and T_{max} intersect at zero, as depicted in Fig. 7E (right panel, indicated by a green asterisk *).

The upper threshold T_{max} gradually shifts toward the highest calcium level observed in each synapse while the lower threshold T_{min} shifts towards a level below the maximum calcium

level (Eqs. 8 and 9; see Fig. 7E right panel for an illustration). Once the thresholds stabilize, the synapse will also reach stability. The goal of this is to enhance the difference in depolarization/calcium between the most active level, and the next one.

Further, depending on whether the synapse was active or inactive during the period of the calcium influx, the learning rate of Eq. 7 (η_{act}) takes on negative or positive values. Because of this, the active synapses are depressed when the calcium concentration is higher than the upper threshold T_{max} and potentiated if the concentration is between the two thresholds, while inactive synapses follow the inverse relationship. No change happens below the lower threshold (see Fig. 7E left panel for an illustration). This framework emphasizes the inhibitory synapse capacity to follow and dynamically amplify the contrast between local excitatory synaptic activity levels.

$$\Omega(Ca, T_{min}, T_{max}, a, b) = \frac{a}{1 + e^{-\beta_4(Ca - 0.004 \cdot T_{min})}} + \frac{b}{1 + e^{-\beta_4(Ca - 0.004 \cdot T_{max})}} \quad \text{Eq. 6}$$

$$dw_{inh} = \eta_{act} \cdot \Omega(Ca_v, T_{min}, T_{max}, a_1, b_1) \cdot w_{inh} \cdot (w_{inh}^{max} - w_{inh}) \quad \text{Eq. 7}$$

$$dT_{max} = \eta_7 \cdot \Omega(Ca_v, T_{min}, T_{max}, a_2, b_2) \quad \text{Eq. 8}$$

$$dT_{min} = \eta_8 \cdot \Omega(Ca_v + c, T_{min}, T_{max}, a_3, b_3) \quad \text{Eq. 9}$$

Training Procedure

We present the neuron with a sequence of 960 stimuli, which, as described in Fig. 1, can come in four varieties (four possible feature combinations). The learning rule is always “on”, i.e. synapses can always be updated, meaning that there are no separate training and testing phases in which synapses are plastic and frozen, respectively. We used two distinct synaptic arrangements for when representing the features: the clustered setup and the distributed setup.

Clustered setup: This setup is based on the assumption of pre-existing clustered synapses for each feature. The features of a stimulus are represented here by the shape and color of bananas and strawberries (Fig. 1A). Features were allocated to two dendritic branches. Each branch had each feature represented with five synapses clustered closely on a single dendrite. Depending on the feature combination, two, three, or four features were represented in clusters on one or both dendritic branches (see Fig. 3A, 4A and 4D for example). Additionally, 108 feature-unspecific synapses were distributed throughout the dendrites, activated concurrently with all stimuli to enhance the probability of spiking (as plateau potentials together with the general background synaptic noise used do not often lead to spikes in SPNs, cf Fig 2A). To demonstrate the inhibitory plasticity rule, four inhibitory synapses near each cluster (in the middle of the dendritic branch), representing each of the

four features, were added (Fig. 4A). With this setup a single feature activates both excitatory and inhibitory synapses. To match the level of depolarization seen in our excitatory-only setup, the number of distributed feature-unspecific synaptic inputs was increased to 144. The initial conductance of inhibitory synapses is set at 0.1 ± 0.01 nS.

Distributed setup: In contrast, this setup examines learning dynamics in neurons without pre-existing synaptic clustering for individual features. A total of 200 excitatory synapses were randomly distributed over 30 dendrites. Each feature was represented by 40 excitatory synapses, and an additional 40 feature unspecific excitatory synapses were used (Fig. 6A). We initiated this experiment with excitatory synaptic weights at 0.3 ± 0.1 (around 0.75 nS). This higher initial weight was chosen to compensate for the reduced efficacy of non-clustered synaptic inputs in producing sufficient depolarization and calcium influx. Extending our investigation, we added 60 inhibitory synapses, 15 for each feature, dispersed randomly over the 30 dendrites. We initiated this extended experiment with excitatory synaptic weights at 0.45 ± 0.1 (around 1.125 nS), aiming to maintain the same baseline voltage activity as in the excitatory-only case. Figure 6A illustrates this setup, highlighting the pre- and post-learning synaptic weights for both excitatory and inhibitory synapses. The four stimuli were presented in random order three times each within a block of 12 stimuli, followed by another reshuffled block of 12, and so on, for 960 trials. Each trial consisted of a stimulus presentation, lasting 20 ms, followed by a reward cue arriving 300 ms after stimulus onset, and lasting for 50 ms (Fig. 3A). During the stimulus presentation, the stimuli-related synapses receive one randomly-timed spike per synapse. We opted not to quantitatively model the dopamine receptor response to dopamine concentrations. Instead, dopamine levels from the reward-based learning system are represented with +1 (dopamine peak) for relevant stimuli if the neuron spikes, and -1 (dopamine pause) for irrelevant stimuli if the neuron spikes, and 0 for baseline levels when the neuron was silent, at which time the calcium amplitudes were measured, and the synaptic weights were updated during the presence of the reward cue.

The time between two stimuli is 800 ms, long enough to allow for the voltage, calcium (and all other state variables in the model) to return to their baseline values.

Table 1: Excitatory Plasticity Parameters

Parameter	Description	Value
η_1	Learning rate for $[\text{Ca}]_{\text{NMDA}}$ -dependent plasticity	$1.5 \times 10^{-5} [\mu\text{S} \cdot \text{mM} \cdot \text{ms}^{-1}]$
η_2	Learning rate for $[\text{Ca}]_{\text{L-type}}$ -dependent plasticity	$3 \times 10^{-3} [\text{ms}^{-1} \cdot \text{mM}^{-1}]$
$\eta_{\text{s, LTP}}$	Learning rate for θ_{LTP} during Da peaks	$1 \times 10^{-7} [\text{mM}^2 \cdot \text{ms}^{-1}]$
$\eta_{\text{s, LTD}}$	Learning rate for θ_{LTP} during Da pauses	$4 \times 10^{-7} [\text{mM}^2 \cdot \text{ms}^{-1}]$
θ_{LTP}	Midpoint of LTP kernel (initial value)	0.02 [mM]

θ_{LTD}	Constant LTD threshold parameter	$7 \times 10^{-5} [mM]$
β_1	Parameter controlling width of LTP kernel	$10^3 [mM^{-1}]$
β_2	Parameter controlling steepness of LTD threshold	$10^5 [mM^{-1}]$
β_3	Parameter controlling width of metaplasticity kernel	$334 [mM^{-1}]$

Table 2: Inhibitory Plasticity Parameters

Parameter	Description	Value
η_{act}	Learning rate for the modification of inhibitory weight for active and inactive synapses	<i>active</i> : -0.055 <i>inactive</i> : 0.055 $[\mu S^{-1} \cdot ms^{-1}]$
w_{inh}^{max}	Maximum synaptic strength	$5 [nS]$
η_7	Learning rate for modifying T_{max}	$9 \times 10^{-4} [mM \cdot ms^{-1}]$
η_8	Learning rate for modifying T_{min}	$-5 \times 10^{-5} [mM \cdot ms^{-1}]$
a_1	Parameter controlling inhibitory weight	-1
b_1	Parameter controlling inhibitory weight	3
a_2	Parameter controlling T_{max}	-1
b_2	Parameter controlling T_{max}	3
a_3	Parameter controlling T_{min}	-2
b_3	Parameter controlling T_{min}	3
c	Calcium concentration offset	$6 \times 10^{-4} [mM]$

β_4	Parameter for controlling the steepness of the curves	$2.5 \times 10^3 [mM^{-1}]$
-----------	---	-----------------------------

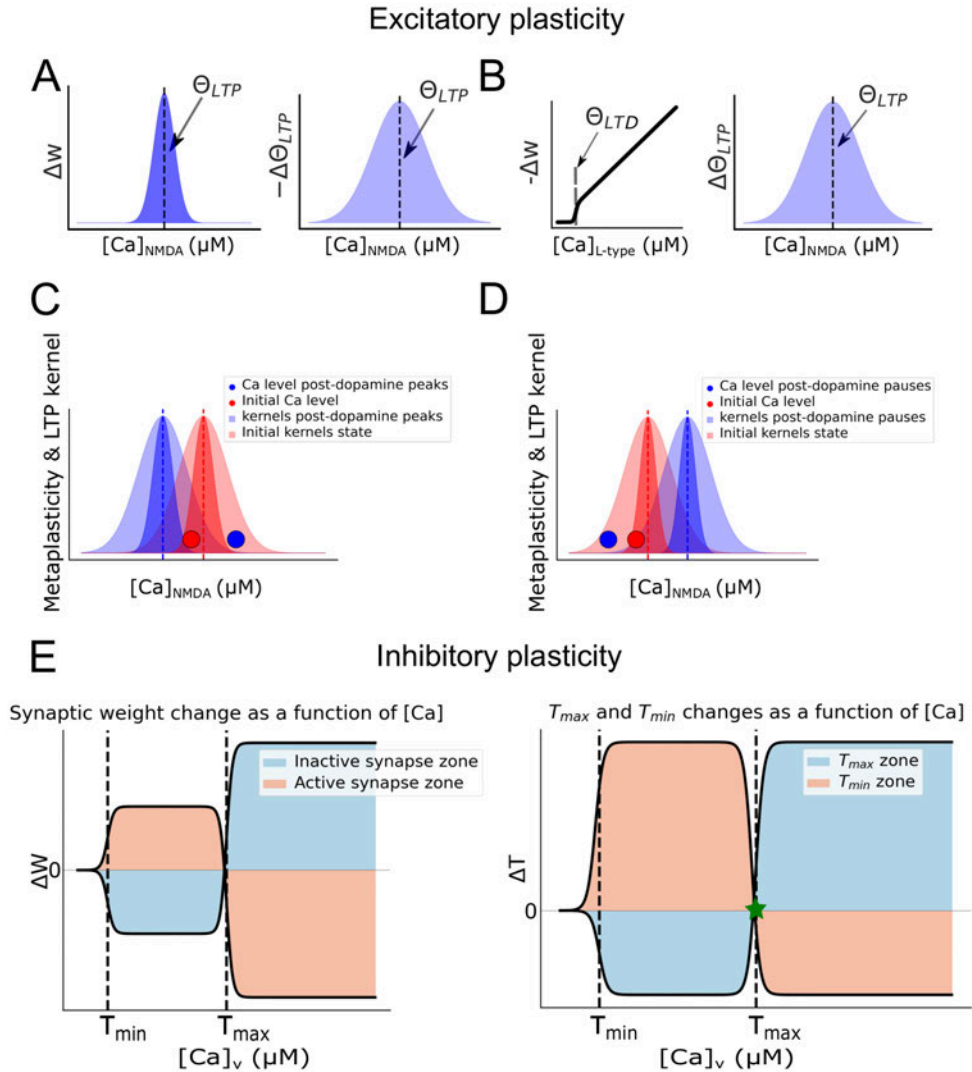


Figure 7: Synaptic Plasticity Rules: Calcium and Dopamine Interactions in Synaptic Weight Modification

A: Synaptic weight updates during a dopamine peak. (Left) LTP kernel is a bell-shaped curve which determines an optimal $[Ca]_{NMDA}$ region in which synaptic weight is increased. (Right) A wider bell-shaped kernel, i.e. the metaplasticity kernel, determines how the LTP kernel (the optimal region for plasticity) slides along the calcium level ($[Ca]_{NMDA}$) axis following a Da peak.

B: Synaptic weight updates during LTD. (Left) The LTD plasticity kernel. This kernel is constant. The LTD threshold is constant and set at 70 nM (Right). Metaplasticity describes how the LTP kernel slides along the calcium axis following a Da pause.

C: A schematic of how the LTP kernel window is updated following a dopamine peak. As NMDA calcium levels increase following activation of the strengthened synapse due to a dopamine peak (illustrated with the red circle jumping to the blue circle), the LTP kernel slides down and in this example the strengthened synapse (blue circle) stabilizes, and doesn't increase in strength even though additional rewards arrive.

D: A schematic showing the update of the LTP window following a dopamine pause leading to that synaptic weight and calcium response decrease as a function of L-type Ca (illustrated with the red circle jumping to the blue one). Here, the LTP kernel slides up towards higher NMDA calcium levels, and thus the weakened synapse (blue circle) is more unlikely to be recruited into the LTP window (unless it in the future is activated with Da peaks more regularly than with Da dips).

E: Depicts the inhibitory plasticity rule. (Left) Changes in synaptic weight for active (orange) and inactive (blue) synapses based on voltage dependent calcium levels in the dendritic shaft at the location of the inhibitory synapse, and minimum (T_{\min}) and maximum (T_{\max}) threshold. (Right) Functions for sliding the minimum and maximum thresholds with voltage dependent calcium level. The asterisk denotes the semi-stable zero level where the curves for active and inactive synapses meet.

Acknowledgement

We would like to thank members of the Hellgren Kotaleski laboratory for helpful discussion of various aspects of the manuscript. DT would like to thank Ana Kalajdjieva for illustrating the mouse brain. We acknowledge the use of Fenix Infrastructure resources, which are partially funded from the European Union's Horizon 2020 research and innovation programme through the ICEI project under the grant agreement No. 800858.

Simulations were also performed on resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at PDC KTH partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

The study was supported by the Swedish Research Council (VR-M-2020-01652), Swedish e-Science Research Centre (SeRC), Science for Life Lab, EU/Horizon 2020 no. 945539 (HBP SGA3) and No. 101147319 (EBRAINS 2.0 Project), European Union's Research and Innovation Program Horizon Europe under grant agreement No 101137289 (the Virtual Brain Twin Project), and KTH Digital Futures.

Author contributions

ZK and DT implemented the model, ran the simulations and analyzed the data with support from RL.

JHK and RL supervised all aspects of the study.

All authors contributed to the writing of the manuscript.

References

- Antic, S.D. *et al.* (2010) 'The decade of the dendritic NMDA spike', *Journal of neuroscience research*, 88(14), pp. 2991–3001.
- Bicknell, B.A. and Häusser, M. (2021) 'A synaptic learning rule for exploiting nonlinear dendritic computation', *Neuron*, 109(24), pp. 4001–4017. e10.
- Bienenstock, E.L., Cooper, L.N. and Munro, P.W. (1982) 'Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 2(1), pp. 32–48.
- Branco, T., Clark, B.A. and Häusser, M. (2010) 'Dendritic discrimination of temporal input sequences in cortical neurons', *Science*, 329(5999), pp. 1671–1675.
- Branco, T. and Häusser, M. (2010) 'The single dendritic branch as a fundamental functional unit in the nervous system', *Current opinion in neurobiology*, 20(4), pp. 494–502.
- Bruce, N.J. *et al.* (2019) 'Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals', *PLoS computational biology*, 15(10), p. e1007382.
- Carter, A.G. and Sabatini, B.L. (2004) 'State-dependent calcium signaling in dendritic spines of striatal medium spiny neurons', *Neuron*, 44(3), pp. 483–493.
- Chen, S.X. *et al.* (2015) 'Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning', *Nature neuroscience*, 18(8), pp. 1109–1115.
- Chen, X. *et al.* (2011) 'Functional mapping of single spines in cortical neurons in vivo', *Nature*, 475(7357), pp. 501–505.
- Cichon, J. and Gan, W.-B. (2015) 'Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity', *Nature*, 520(7546), pp. 180–185.
- Day, M. *et al.* (2024) 'GABAergic regulation of striatal spiny projection neurons depends upon their activity state', *PLoS biology*, 22(1), p. e3002483.
- Destexhe, A., Mainen, Z.F. and Sejnowski, T.J. (1994) 'An efficient method for computing synaptic conductances based on a kinetic model of receptor binding', *Neural computation*, 6(1), pp. 14–18.
- Dorman, D.B., Jędrzejewska-Szmek, J. and Blackwell, K.T. (2018) 'Inhibition enhances spatially-specific calcium encoding of synaptic input patterns in a biologically constrained model', *eLife*, 7. Available at: <https://doi.org/10.7554/eLife.38588>.
- Doron, M. *et al.* (2017) 'Timed Synaptic Inhibition Shapes NMDA Spikes, Influencing Local Dendritic Processing and Global I/O Properties of Cortical Neurons', *Cell reports*, 21(6), pp. 1550–1561.
- Du, K. *et al.* (2017) 'Cell-type-specific inhibition of the dendritic plateau potential in striatal spiny projection neurons', *Proceedings of the National Academy of Sciences of the United States of America*, 114(36), pp. E7612–E7621.
- Fino, E. *et al.* (2010) 'Distinct coincidence detectors govern the corticostriatal spike timing-dependent plasticity', *The Journal of physiology*, 588(Pt 16), pp. 3045–3062.

Fisher, S.D. *et al.* (2017) 'Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo', *Nature communications*, 8(1), p. 334.

Gandolfi, D. *et al.* (2020) 'Inhibitory Plasticity: From Molecules to Computation and Beyond', *International journal of molecular sciences*, 21(5). Available at: <https://doi.org/10.3390/ijms21051805>.

Gao, P.P. *et al.* (2021) 'Local glutamate-mediated dendritic plateau potentials change the state of the cortical pyramidal neuron', *Journal of neurophysiology*, 125(1), pp. 23–42.

Gidon, A. *et al.* (2020) 'Dendritic action potentials and computation in human layer 2/3 cortical neurons', *Science*, 367(6473), pp. 83–87.

Henneberger, C. *et al.* (2020) 'LTP Induction Boosts Glutamate Spillover by Driving Withdrawal of Perisynaptic Astroglia', *Neuron*, 108(5), pp. 919–936.e11.

van den Heuvel, M.P. and Sporns, O. (2013) 'Network hubs in the human brain', *Trends in cognitive sciences*, 17(12), pp. 683–696.

Higley, M.J. and Sabatini, B.L. (2010) 'Competitive regulation of synaptic Ca²⁺ influx by D2 dopamine and A2A adenosine receptors', *Nature neuroscience*, 13(8), pp. 958–966.

Hwang, F.-J. *et al.* (2022) 'Motor learning selectively strengthens cortical and striatal synapses of motor engram neurons', *Neuron*, 110(17), pp. 2790–2801.e5.

Iacaruso, M.F., Gasler, I.T. and Hofer, S.B. (2017) 'Synaptic organization of visual space in primary visual cortex', *Nature*, 547(7664), pp. 449–452.

Jia, H. *et al.* (2010) 'Dendritic organization of sensory input to cortical neurons in vivo', *Nature*, 464(7293), pp. 1307–1312.

Johansson, Y. and Silberberg, G. (2020) 'The Functional Organization of Cortical and Thalamic Inputs onto Five Types of Striatal Neurons Is Determined by Source and Target Cell Identities', *Cell reports*, 30(4), pp. 1178–1194.e3.

Ju, N. *et al.* (2020) 'Spatiotemporal functional organization of excitatory synaptic inputs onto macaque V1 neurons', *Nature communications*, 11(1), p. 697.

Kahneman, D. and Tversky, A. (1979) 'Prospect theory: An analysis of decision under risk', *Econometrica: journal of the Econometric Society*, 47(2), p. 263.

Kawai, R. *et al.* (2015) 'Motor cortex is required for learning but not for executing a motor skill', *Neuron*, 86(3), pp. 800–812.

Kerlin, A. *et al.* (2019) 'Functional clustering of dendritic activity during decision-making', *eLife*, 8. Available at: <https://doi.org/10.7554/eLife.46966>.

Kleindienst, T. *et al.* (2011) 'Activity-dependent clustering of functional synaptic inputs on developing hippocampal dendrites', *Neuron*, 72(6), pp. 1012–1024.

Larkum, M.E. *et al.* (2009) 'Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle', *Science*, 325(5941), pp. 756–760.

Lavzin, M. *et al.* (2012) 'Nonlinear dendritic processing determines angular tuning of barrel cortex neurons in vivo', *Nature*, 490(7420), pp. 397–401.

Legenstein, R. and Maass, W. (2011) 'Branch-specific plasticity enables self-organization of

nonlinear computation in single neurons', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 31(30), pp. 10787–10802.

Lindroos, R. *et al.* (2018) 'Basal Ganglia Neuromodulation Over Multiple Temporal and Structural Scales-Simulations of Direct Pathway MSNs Investigate the Fast Onset of Dopaminergic Effects and Predict the Role of Kv4.2', *Frontiers in neural circuits*, 12, p. 3.

Lindroos, R. and Hellgren Kotaleski, J. (2021) 'Predicting complex spikes in striatal projection neurons of the direct pathway following neuromodulation by acetylcholine and dopamine', *The European journal of neuroscience*, 53(7), pp. 2117–2134.

Losonczy, A. and Magee, J.C. (2006) 'Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons', *Neuron*, 50(2), pp. 291–307.

Major, G. *et al.* (2008) 'Spatiotemporally graded NMDA spike/plateau potentials in basal dendrites of neocortical pyramidal neurons', *Journal of neurophysiology*, 99(5), pp. 2584–2601.

von der Malsburg, C. (1999) 'The what and why of binding: the modeler's perspective', *Neuron*, 24(1), pp. 95–104, 111–25.

Matsuda, W. *et al.* (2009) 'Single nigrostriatal dopaminergic neurons form widely spread and highly dense axonal arborizations in the neostriatum', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 29(2), pp. 444–453.

McCulloch, W.S. and Pitts, W. (1943) 'A logical calculus of the ideas immanent in nervous activity', *The Bulletin of mathematical biophysics*, 5(4), pp. 115–133.

Nair, A.G. *et al.* (2015) 'Sensing Positive versus Negative Reward Signals through Adenylyl Cyclase-Coupled GPCRs in Direct and Indirect Pathway Striatal Medium Spiny Neurons', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 35(41), pp. 14017–14030.

Niculescu, D. *et al.* (2018) 'A BDNF-Mediated Push-Pull Plasticity Mechanism for Synaptic Clustering', *Cell reports*, 24(8), pp. 2063–2074.

Oikonomou, K.D. *et al.* (2014) 'Spiny neurons of amygdala, striatum, and cortex use dendritic plateau potentials to detect network UP states', *Frontiers in cellular neuroscience*, 8, p. 292.

Plotkin, J.L. *et al.* (2013) 'Regulation of dendritic calcium release in striatal spiny projection neurons', *Journal of neurophysiology*, 110(10), pp. 2325–2336.

Plotkin, J.L., Day, M. and Surmeier, D.J. (2011) 'Synaptically driven state transitions in distal dendrites of striatal spiny neurons', *Nature neuroscience*, 14(7), pp. 881–888.

Poirazi, P., Brannon, T. and Mel, B.W. (2003) 'Pyramidal neuron as two-layer neural network', *Neuron*, 37(6), pp. 989–999.

Polsky, A., Mel, B.W. and Schiller, J. (2004) 'Computational subunits in thin dendrites of pyramidal cells', *Nature neuroscience*, 7(6), pp. 621–627.

Ravasenga, T. *et al.* (2022) 'Spatial regulation of coordinated excitatory and inhibitory synaptic plasticity at dendritic synapses', *Cell reports*, 38(6), p. 110347.

Reig, R. and Silberberg, G. (2014) 'Multisensory integration in the mouse striatum', *Neuron*, 83(5), pp. 1200–1212.

Reynolds, J.N.J. *et al.* (2022) 'Coincidence of cholinergic pauses, dopaminergic activation and depolarisation of spiny projection neurons drives synaptic plasticity in the striatum', *Nature communications*, 13(1), p. 1296.

Roskies, A.L. (1999) 'The binding problem', *Neuron*, 24(1), pp. 7–9, 111–25.

Sanabria, B.D. *et al.* (2024) 'Cell-Type Specific Connectivity of Whisker-Related Sensory and Motor Cortical Input to Dorsal Striatum', *eNeuro*, 11(1). Available at: <https://doi.org/10.1523/ENEURO.0503-23.2023>.

Santos, F.J. *et al.* (2015) 'Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning', *eLife*, 4, p. e09423.

Schiess, M., Urbanczik, R. and Senn, W. (2016) 'Somato-dendritic Synaptic Plasticity and Error-backpropagation in Active Dendrites', *PLoS computational biology*, 12(2), p. e1004638.

Schiller, J. *et al.* (2000) 'NMDA spikes in basal dendrites of cortical pyramidal neurons', *Nature*, 404(6775), pp. 285–289.

Scholl, B., Wilson, D.E. and Fitzpatrick, D. (2017) 'Local Order within Global Disorder: Synaptic Architecture of Visual Space', *Neuron*, 96(5), pp. 1127–1138.e4.

Schultz, W. (2007) 'Multiple dopamine functions at different time courses', *Annual review of neuroscience*, 30, pp. 259–288.

Schultz, W., Dayan, P. and Montague, P.R. (1997) 'A neural substrate of prediction and reward', *Science*, 275(5306), pp. 1593–1599.

Shen, W. *et al.* (2008) 'Dichotomous dopaminergic control of striatal synaptic plasticity', *Science*, 321(5890), pp. 848–851.

Shindou, T. *et al.* (2019) 'A silent eligibility trace enables dopamine-dependent synaptic plasticity for reinforcement learning in the mouse striatum', *The European journal of neuroscience*, 49(5), pp. 726–736.

Shindou, T., Ochi-Shindou, M. and Wickens, J.R. (2011) 'A Ca(2+) threshold for induction of spike-timing-dependent depression in the mouse striatum', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 31(36), pp. 13015–13022.

Silver, D. *et al.* (2018) 'A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play', *Science*, 362(6419), pp. 1140–1144.

Stauffer, W.R. *et al.* (2016) 'Components and characteristics of the dopamine reward utility signal', *The Journal of comparative neurology*, 524(8), pp. 1699–1711.

Surmeier, D.J. *et al.* (2010) 'The role of dopamine in modulating the structure and function of striatal circuits', *Progress in brain research*, 183, pp. 149–167.

Takahashi, N. *et al.* (2012) 'Locally synchronized synaptic inputs', *Science*, 335(6066), pp. 353–356.

Tran-Van-Minh, A. *et al.* (2015) 'Contribution of sublinear and supralinear dendritic integration to neuronal computations', *Frontiers in cellular neuroscience*, 9, p. 67.

Trpevski, D. *et al.* (2023) 'Glutamate spillover drives robust all-or-none dendritic plateau potentials—an investigation using models of striatal projection neurons', *Frontiers in cellular neuroscience*, 17, p. 1196182.

Varga, Z. *et al.* (2011) 'Dendritic coding of multiple sensory inputs in single cortical neurons in vivo', *Proceedings of the National Academy of Sciences of the United States of America*, 108(37), pp. 15420–15425.

Wilson, C.J. *et al.* (1983) 'Three-dimensional structure of dendritic spines in the rat neostriatum', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 3(2), pp. 383–388.

Wilson, D.E. *et al.* (2016) 'Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex', *Nature neuroscience*, 19(8), pp. 1003–1009.

Winnubst, J. *et al.* (2015) 'Spontaneous Activity Drives Local Synaptic Plasticity In Vivo', *Neuron*, 87(2), pp. 399–410.

Wolf, J.A. *et al.* (2005) 'NMDA/AMPA ratio impacts state transitions and entrainment to oscillations in a computational model of the nucleus accumbens medium spiny projection neuron', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 25(40), pp. 9080–9095.

Xu, N.-L. *et al.* (2012) 'Nonlinear dendritic integration of sensory and motor input during an active sensing task', *Nature*, 492(7428), pp. 247–251.

Yagishita, S. *et al.* (2014) 'A critical time window for dopamine actions on the structural plasticity of dendritic spines', *Science*, 345(6204), pp. 1616–1620.

Zenke, F. and Gerstner, W. (2017) 'Hebbian plasticity requires compensatory processes on multiple timescales', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 372(1715). Available at: <https://doi.org/10.1098/rstb.2016.0259>.

Zenke, F., Gerstner, W. and Ganguli, S. (2017) 'The temporal paradox of Hebbian learning and homeostatic plasticity', *Current opinion in neurobiology*, 43, pp. 166–176.

Article C



Glutamate spillover drives robust all-or-none dendritic plateau potentials—an in silico investigation using models of striatal projection neurons
(2023) *Frontiers in Cellular Neuroscience* 17:1196182



OPEN ACCESS

EDITED BY

Srđjan D. Antic,
University of Connecticut Health Center,
United States

REVIEWED BY

William W. Lytton,
Downstate Health Sciences University,
United States
Pierre Apostolides,
University of Michigan, United States

*CORRESPONDENCE

Daniel Trpevski



Jeanette Hellgren Kotaleski



RECEIVED 29 March 2023

ACCEPTED 31 May 2023

PUBLISHED 29 June 2023

CITATION

Trpevski D, Khodadadi Z, Carannante I and
Hellgren Kotaleski J (2023) Glutamate spillover
drives robust all-or-none dendritic plateau
potentials—an *in silico* investigation using
models of striatal projection neurons.
Front. Cell. Neurosci. 17:1196182.
doi: 10.3389/fncel.2023.1196182

COPYRIGHT

© 2023 Trpevski, Khodadadi, Carannante and
Hellgren Kotaleski. This is an open-access
article distributed under the terms of the
[Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/).
The use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in this
journal is cited, in accordance with accepted
academic practice. No use, distribution or
reproduction is permitted which does not
comply with these terms.

Glutamate spillover drives robust all-or-none dendritic plateau potentials—an *in silico* investigation using models of striatal projection neurons

Daniel Trpevski^{1*}, Zahra Khodadadi¹, Ilaria Carannante¹ and
Jeanette Hellgren Kotaleski^{1,2*}

¹Science for Life Laboratory, Department of Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden, ²Department of Neuroscience, Karolinska Institutet, Stockholm, Sweden

Plateau potentials are a critical feature of neuronal excitability, but their all-or-none behavior is not easily captured in modeling. In this study, we investigated models of plateau potentials in multi-compartment neuron models and found that including glutamate spillover provides robust all-or-none behavior. This result arises due to the prolonged duration of extrasynaptic glutamate. When glutamate spillover is not included, the all-or-none behavior is very sensitive to the steepness of the Mg^{2+} block. These results suggest a potentially significant role of glutamate spillover in plateau potential generation, providing a mechanism for robust all-or-none behavior across a wide range of slopes of the Mg^{2+} block curve. We also illustrate the importance of the all-or-none plateau potential behavior for nonlinear computation with regard to the nonlinear feature binding problem.

KEYWORDS

glutamate spillover, plateau potentials, NMDA spikes, gating function, computational modeling, nonlinear dendritic computation, clustered synapses, magnesium block of NMDA receptors

1. Introduction

Plateau potentials are all-or-none voltage elevations that occur in neuronal dendrites in various brain regions, including striatal projection neurons (SPNs), which last tens to hundreds of milliseconds (Plotkin et al., 2011; Oikonomou et al., 2014; Du et al., 2017). They are involved in processing sensory information, as well as triggering synaptic plasticity (Lavzin et al., 2012; Xu et al., 2012; Gambino et al., 2014; Kumar et al., 2018). Clustered synapses, which facilitate the induction of dendritic NMDA spikes and plateau potentials, have been shown to form during development and learning (see e.g., review by Kastellakis and Poirazi, 2019). Plateau potentials are generated on stretches of dendrites around and above 20 μm of length, where the N-methyl-D-aspartate (NMDA) receptors that mediate them must have a cumulative conductance above a threshold level (Antic et al., 2010). Experimentally, plateau potentials have been evoked by activating clusters of synapses by glutamate uncaging or by stimulating stretches of dendrites by glutamate iontophoresis or repeated synaptic stimulation (Milojkovic et al., 2004; Major et al., 2008; Plotkin et al., 2011; Oikonomou et al., 2012; Du et al., 2017; Kumar et al., 2018; Gao et al., 2021). In the cases of progressively increasing glutamate stimulus in equal increments, a supralinear, all-or-none response in the amplitude of the somatic voltage is observed (Major et al., 2008; Oikonomou et al., 2012; Gao et al., 2021).

In this study, we deal with the question of how to model the all-or-none behavior of plateau potentials in multicompartment neuron models and discuss the implications of the model for the precise mechanism of generation of plateau potentials experimentally and *in vivo*.

The NMDA receptors are ionotropic receptors activated by glutamate and glycine as a co-factor that are additionally blocked by Mg^{2+} ions at resting membrane potentials. The amount of blockage depends on the local voltage, and it has been determined that this dependence is a sigmoidal nonlinearity (Figure 1C), where higher voltage relieves the block (Jahr and Stevens, 1990b). The generation mechanism is a positive feedback loop that starts by partially alleviating the Mg^{2+} block on NMDA receptors (NMDARs) when the dendritic voltage reaches a threshold value, usually around -50 mV. This causes the flow of an inward current through the NMDA receptors, raising the voltage in a cycle that progressively alleviates more of the Mg^{2+} block and allows for more inward current. Thus, a sustained voltage elevation is produced, which lasts as long as there is glutamate to bind to the NMDARs. The sigmoidal shape of the voltage nonlinearity is a crucial ingredient for the all-or-none behavior of the plateau potentials since, for linear increases in voltage, it provides a “sudden” relief of Mg^{2+} ions from the NMDARs. The binding of neurotransmitters to the NMDAR initiates a transition from the closed receptor state to the open state. The behavior of an NMDAR in the open state is described with the diagram in Figure 1A, also called the three-state model, since it includes three receptor states: open, blocked, and closed. An NMDAR in the open state can transition to two different non-conducting states, closed and blocked. The transitions between the open and blocked states depend on Mg^{2+} concentration (denoted as $[Mg^{2+}]$) and voltage, while the transitions from the open and blocked states to the closed state are voltage- and $[Mg^{2+}]$ -independent. Voltage crossing the -50 mV threshold unblocks some of the receptors, initiating the self-sustaining positive feedback loop. Neurotransmitter removal closes the receptor pore (the receptor transitions to the closed state), thus stopping the inward current and lowering the voltage. Lowering the voltage causes remaining open receptors (with bound neurotransmitter) to transition to the blocked state, which because of the sigmoidal nonlinearity rapidly lowers the voltage back to the resting state.

1.1. NMDAR state diagrams and the gating function

The three-state model has been shown to be suitable for high Mg^{2+} concentrations ($[Mg^{2+}] > 0.2$ mM), but it can be easily expanded to take into account lower concentrations by introducing an additional, $[Mg^{2+}]$ -independent, blocked state (Figure 1B, the four-state model; Jahr and Stevens, 1990a). All the transition rates of the four-state model have been derived from single-channel measurements and reported in Table 1 (reproduced from Jahr and Stevens, 1990a). The rate of leaving the open state in the three-state model ($O \rightarrow B$) is a , and the corresponding rate for the four-state model ($B_1 \leftarrow O \rightarrow B_2$) is $a_1 + a_2$. The transition rate from blocked to open is b in the three-state model, whereas the corresponding

rate in the four-state model is $b_1 a_1 / (a_1 + a_2) + b_2 a_2 / (a_1 + a_2)$ because the rate of leaving each blocked state is weighted by the probability that particular state had been entered. Similarly, the closing rate from the blocked state is B for the three-state model and $B_1 a_1 / (a_1 + a_2) + B_2 a_2 / (a_1 + a_2)$. The Mg^{2+} block is expressed with a gating function, $g(V)$, which gives the fraction of open NMDARs for a given voltage value. As shown in Jahr and Stevens (1990b) it is possible to express $g(V)$ in terms of the transition rates as:

$$g(V) = \frac{1}{1 + \frac{(a_1 + a_2)(a_1 B_1 + a_2 B_2)}{A a_1 (b_1 + B_1) + A a_2 (b_2 + B_2)}}. \quad (1)$$

The single-channel analysis showed that b_1 and b_2 are much larger than B_1 and B_2 (also visible in Table 1); hence, the following is a good approximation:

$$g(V) \approx \frac{1}{1 + \frac{(a_1 + a_2)(a_1 B_1 + a_2 B_2)}{A a_1 b_1 + A a_2 b_2}}. \quad (2)$$

Moreover, when $[Mg^{2+}]$ is greater than a few hundred micromolar, a_2 becomes much larger than a_1 since the former increases linearly with $[Mg^{2+}]$ while the latter is independent of it, and hence $g(V)$ for physiological $[Mg^{2+}]$ can be further approximated as:

$$g(V) \approx \frac{1}{1 + \frac{B_2 a_2}{A b_2}} \approx \frac{1}{1 + \frac{B a}{A b}}. \quad (3)$$

Finally, taking into account the transition rates in Table 1, the gating function can be written in its most commonly used form:

$$g(V) = \frac{1}{1 + \eta [Mg^{2+}] e^{-\alpha V}}, \quad (4)$$

where α is the steepness of the curve, η is related to the voltage value $V_{1/2}$ where half of the receptors are free of Mg^{2+} according to $V_{1/2} = \frac{\ln(\eta[Mg^{2+}])}{\alpha}$, $[Mg^{2+}]$ is the concentration of extracellular Mg^{2+} in mM and V is the membrane potential in mV (see Supplementary material for details and Supplementary Figure 1). As described in the Supplementary material, the values of α and η are determined by the transition rates between different states in the 4-state model, which in turn are determined by the structural properties of the NMDARs.

The gating function adequately describes the behavior not only of single channels but also of macroscopic NMDAR currents, which allows for the determination of the parameters α and η from macroscopic current measurements. Such macroscopic current measurements from NMDARs, as well as gating functions fitted to such measurements, have been reported numerous times since the work of Jahr and Stevens (1990b), and are summarized in Table 2. For completeness, in Table 2 we have added parameter values which were not estimated in the original articles. The fitting procedure, as well as the results of the fittings, are described in the Methods and Results sections, respectively. Table 2 shows a natural variability in the parameters of the Mg^{2+} block gating function which depends on the NMDAR subunit composition, as well as on the concentration of other intra- and extra-cellular ions (Qian et al., 2005; McMenimen et al., 2006; Qian and Johnson, 2006; Retchless et al., 2012). Since the sigmoidal Mg^{2+} block is a crucial ingredient for the all-or-none behavior of plateau potentials, we investigate the effect of the parameter variability on the behavior of the plateau potentials.




FIGURE 1

(A) The 3-state model for the functioning of open NMDARs, reproduced from Jahr and Stevens (1990b). It has three states: open (O), blocked (B), and closed (C). The transition rates between the open and blocked states, a and b , are voltage-dependent, while the transitions to the closed state, A and B , are not. The rate of blocking the open state, a , is additionally linearly dependent on $[Mg^{2+}]$. Transmitter removal causes the transitions to the closed state. (B) The four-state model for the functioning of NMDARs used in Jahr and Stevens (1990b) to derive the gating function for the Mg^{2+} block. It contains an additional blocked state, B_1 , the transition to which (a_1) is $[Mg^{2+}]$ -independent. The other transition rates are analogous to the three-state model, and are given in Table 1. (C) The gating function is expressed as a sigmoidal curve and shows the proportion of open NMDARs as a function of voltage. The steepness of the curve and its position along the x-axis, described with the parameters α and η in Equation (4) are usually fitted from macroscopic NMDA current measurements. The black line is the curve with the parameter values determined in Jahr and Stevens (1990b), and the gray dashed line is the curve with the corrected parameter value for η as in Ecker et al. (2020), which accounts for the junction potential.

TABLE 1 Transition rates for the four-state model reproduced from Jahr and Stevens (1990a,b).

Transition rate	Value
a_1	$\exp(-0.016 \cdot V - 2.91)$ [ms^{-1}]
a_2	$C \cdot \exp(-0.045 \cdot V - 6.97)$ [$\mu M^{-1} ms^{-1}$]
b_1	$\exp(0.009 \cdot V + 1.22)$ [ms^{-1}]
b_2	$\exp(0.017 \cdot V + 0.96)$ [ms^{-1}]
A	$\exp(-2.847)$ [ms^{-1}]
B_1	$\exp(-0.693)$ [ms^{-1}]
B_2	$\exp(-3.101)$ [ms^{-1}]

The transition rates were derived by single-channel analysis. C represents the magnesium concentration in μM and V the voltage in mV.

1.2. Computational models of dendritic plateau potentials

Studies that model plateau potentials and/or NMDA spikes fall into two groups. In the first group are studies that simply include the voltage-dependent Mg^{2+} block of the NMDARs in an NMDA synapse model and activate a cluster of such synapses on a dendrite to generate the plateau potential. The second group comprises one recent study that includes the effect of glutamate spillover when modeling plateau potentials (Gao et al., 2021). Glutamate spillover is thought to occur when a cluster of synapses is stimulated more strongly or repeatedly, and the amounts of exocytosed glutamate surpass the ability of astrocytes to take it up, thus causing it to spill over from the synaptic cleft into the space surrounding the spine necks and dendritic shafts. Here it stimulates extrasynaptic NMDARs (eNMDARs), which can aid the generation of plateau potentials (Chalifoux and Carter, 2011). The eNMDARs have also been found to have a substantial contribution to neuronal upstates in striatal SPNs, which in turn have been hypothesized to be driven by dendritic plateau potentials (Oikonomou et al., 2014; Garcia-Munoz et al., 2015).

None of the studies have explicitly focused on modeling the all-or-none behavior of the plateau potentials. However, all studies have changed the originally reported Mg^{2+} block parameters by

Jahr and Stevens (1990b) without explicitly stating a reason for the change. The parameter values for α , η , and $[Mg^{2+}]$ that have been used in the different studies are exemplified in Table 3. In addition, all of these studies except one use steeper sigmoid curves compared to the one reported by Jahr and Stevens (1990b).

We analyze the two different types of models of generating plateau potentials and find that the all-or-none behavior is very sensitive to the steepness of the Mg^{2+} block when no glutamate spillover is included in the model and, conversely, very robust with respect to the steepness of the Mg^{2+} block when glutamate spillover is included. Without glutamate spillover, in order to obtain the all-or-none behavior of the plateau potentials, one needs to increase the steepness of the Mg^{2+} block since the originally reported curve is usually too shallow and provides graded rather than all-or-none voltage elevations. In light of the variability in the steepness of the gating function (Table 2), these results indicate that glutamate spillover might be the reason why plateau potentials exhibit a robust all-or-none quality in experimental recordings.

A robust all-or-none plateau potential behavior can significantly improve nonlinear computation, as we illustrate for the nonlinear feature binding problem. We conclude by discussing the implications of the glutamate spillover model in light of the many NMDAR isoforms, the experimentally observed all-or-none behavior of plateau potentials, as well as the predicted role for dendritic computations.

2. Methods

The code implementation is done in Python 3, and the simulations in NEURON (Carnevale and Hines, 2005) are used via the Python interface.

2.1. Neuron models

The SPN model is taken from the collection of models published in Lindroos and Hellgren Kotaleski (2021). Spines in the neurons are added as additional compartments consisting of

TABLE 2 Gating function parameter values from experimental data.

	NMDAR type	[Mg] (mM)	α (/mV)	η (/mM)	Cell type	Ionic concentration (mM)
Nowak et al. (1984)	?	0.5	0.04	1.33	Mouse mesencephalic or striatal neurons	140 [Na] _o
Jahr and Stevens (1990b)	?	1	0.062	0.28	Rat hippocampal neurons	165 [Na] _o / 150 [Cs] _i
Chen and Huang (1992)	?	0.03	0.05	0.49	Trigeminal subnucleus caudalis neurons	140 [Na] _o / 125 [Cs] _i
Sharma and Stevens (1996)	N1/2A	3	0.06	0.28	HEK293 cells	145 [Na] _o / 160 [Cs] _i
McMenimen et al. (2006)	N1/2B	2	0.06	0.42	Xenopus laevis oocytes	96 [Na] _o
	N1/2B	0.2	0.05	3.3		96 [Na] _o
Chiu and Carter (2022)	?	1	0.074	0.11	CD1 mice Layer 2/3 pyramidal neurons	155 [Na] _o , 2 [Ca] _o , 2.5 [K] _o / 130 [Cs] _i or 128 [K] _i
	?	0.7	0.074	0.104		155 [Na] _o , 1.2 [Ca] _o , 4.2 [K] _o / 130 [Cs] _i or 128 [K] _i
	?	0.8	0.071	0.119		155 [Na] _o , 1.3 [Ca] _o , 3.8 [K] _o / 130 [Cs] _i or 128 [K] _i

Parameter values are determined from fits to examples of experimental data. Fits for all articles except Jahr and Stevens (1990b) were done as described in Section 2.

a neck and head with lengths and diameters $l_{\text{neck}} = 0.5 \mu\text{m}$, $l_{\text{head}} = 0.5 \mu\text{m}$, and $d_{\text{neck}} = 0.125 \mu\text{m}$, $d_{\text{head}} = 0.5 \mu\text{m}$, respectively. Axial resistance in all compartments is $150 \text{ Ohm} \cdot \text{cm}$, except for the spine neck, where it is $1,130 \text{ Ohm} \cdot \text{cm}$ (Dorman et al., 2018). Spines contain the inwardly rectifying potassium channel, with a conductance equal to that of the parent dendritic shaft segment. Voltage-gated calcium channels of types R ($\text{Ca}_v2.3$), T ($\text{Ca}_v3.2$ and $\text{Ca}_v3.3$), and L ($\text{Ca}_v1.2$ and $\text{Ca}_v1.3$) are added to the spines as well, and their conductances have been manually tuned to match the relative proportions determined in Carter and Sabatini (2004), and Higley and Sabatini (2010), as well as calcium concentration amplitudes arising from backpropagating action potential (bAP) stimulation as in Figure 2 of Shindou et al. (2011).

2.2. Background synaptic noise

Background synaptic noise is implemented by adding synapses across the dendrites according to the excitatory and inhibitory synapse densities reported in Cheng et al. (1997). These synapses are not on spines. We only explicitly model spines where we have clustered synaptic inputs. To increase the speed of the simulations, in compartments where the synaptic density specifies adding more than one synapse, only one synapse was added with scaled-up input frequency to account for the actual number of synapses that would need to arrive in that compartment.

2.3. Synaptic inputs

Synaptic inputs consist of AMPA and NMDA synapses located on spines within clusters. Glutamate spillover is modeled by including extrasynaptic NMDA conductances on the dendritic shafts situated directly beneath the spines. For glutamate spillover, we put as many extrasynaptic NMDA synapses as the NMDA

TABLE 3 Examples of studies that have modeled plateau potentials and the parameters used for the Mg^{2+} block.

	[Mg] (mM)	α (/mV)	η (/mM)
Rhodes (2006)	1 and 2	0.08	0.28
Major et al. (2008)	1.8	0.08	0.11
Farinella et al. (2014)	1	0.08	0.3
Poleg-Polsky (2015)	1	0.08	0.25
Doron et al. (2017)	1	0.08	0.28 and 1.45
Du et al. (2017)	1	0.07	0.33
Dorman et al. (2018)	1.4	0.099	0.055
Kumar et al. (2018)	1	0.08	0.25
Ecker et al. (2020)	1	0.062	0.38
Gao et al. (2021)	1	0.08	0.25

Ecker et al. (2020) have corrected η for the omitted junction potential in Jahr and Stevens (1990b), and use the same value for α . Note that Major et al. (2008) have used a different form of the equation, and we have converted their parameter values to match the form in Equation (4).

synapses on the spines. The AMPA and NMDA synapse models are taken from Gao et al. (2021), which are a variation of the saturating synapse models in Destexhe et al. (1994) implemented in the NEURON simulator (Carnevale and Hines, 2005). These are kinetic models which operate according to two different kinetic schemes depending on the presence of neurotransmitters. When neurotransmitters arrive at the postsynaptic site, receptor

dynamics are described by a kinetic scheme that models switching between closed (C) and open (O) states with rates of α and β :



where T is the transmitter concentration. After neurotransmitter has been cleared from the synaptic cleft, receptor dynamics evolves according to a kinetic scheme that describes just the closing of the receptor with a rate β :



When a presynaptic spike arrives, neurotransmitter levels are assumed to always reach a fixed saturating concentration, T_{\max} , in the synaptic cleft, i.e., are represented by a pulse with amplitude T_{\max} and duration of T_{dur} . A presynaptic spike that arrives while the neurotransmitter pulse is still on lengthens the pulse duration by T_{dur} .

Extrasynaptic receptors are modeled by the same kinetic scheme, with two differences: (i) T_{dur} is much longer, modeling the spillover effect of glutamate in the extrasynaptic space, and (ii) the synaptic spike there arrives with a delay of 1.5 ms after the glutamate threshold has been reached (taken from Szapiro and Barbour, 2007, as we describe in the next section). The channel properties of NMDARs and eNMDARs have been reported to be similar (Clark et al., 1997).

We use the model given in Gao et al. (2021) but with modified parameter values, given in Table 4. For example, in the simulation code for the model in Gao et al. (2021) different maximal transmitter concentration, T_{\max} , for the synaptic AMPA and NMDA synapses is used. Although difficult to justify physiologically, this difference would not affect the results in Gao et al. (2021), since the AMPA synapses are only used to activate the NMDARs and eNMDARs, and since the effect of the different concentrations can be offset with suitable values for the synaptic conductances. We have nevertheless used the same T_{\max} for both AMPARs and synaptic NMDARs. In addition, the model in Gao et al. (2021) uses the same T_{\max} for NMDARs and eNMDARs. We have changed this, since, even though the extrasynaptic glutamate concentration is not precisely known, extrasynaptic spillover-induced currents are very sensitive to low-affinity AMPA antagonists, suggesting lower extrasynaptic glutamate concentrations (Szapiro and Barbour, 2007). Low glutamate concentrations during spillover have also been reported in Okubo et al. (2010). As for the maximal synaptic glutamate concentration, T_{\max} , it has been found to be very close to 1 mM at cultured hippocampal synapses (Clements et al., 1992).

We have used a short T_{dur} for NMDARs (the same as for AMPARs), and long T_{dur} for eNMDARs. Lastly, we have modified the α and β parameters for AMPARs to more realistic values. Also note that we have denoted transmitter duration and transmitter concentration with T_{dur} and T_{\max} , respectively, which are given with C_{dur} and C_{\max} in the simulation code in Gao et al. (2021).

eNMDARs have been found to have both diffuse and punctual distribution along dendritic shafts, and electrophysiological

measurements have reported their conductance to be 20–60% of the total NMDAR conductance (Harris and Pettit, 2007; Petralia, 2012; Papouin and Oliet, 2014). Even though synapses are represented as punctual inputs in NEURON, the distribution of eNMDAR synapses along 20–30 μm of a dendrite resulted in us effectively modeling a diffuse distribution of eNMDARs. The maximal conductance of an extrasynaptic NMDA synapse in the model is the same as the maximal conductance for a synaptic NMDA synapse. The total synaptic surface area in the clustered spines is around 10% of the surface area of the dendritic shaft where a cluster is situated, making the density of eNMDARs 10 times smaller than the density of NMDARs on the spine.

2.4. Evoking plateau potentials

Plateau potentials are evoked by clustered synaptic input. Clusters ranged from 1 to 40 synapses for the model without spillover and from 1 to 20 synapses for the models with spillover; in SPNs, which are spherically symmetric, a cluster is placed on one of the dendrites, ~ 120 – $140 \mu\text{m}$ from the soma. The difference in the size of the clusters for the cases without and with glutamate spillover is in order to compare the same total NMDAR conductance in both cases. Since both cases have equal spine NMDAR conductance and the case without spillover has no eNMDARs, this requires that it possesses either a twice larger cluster size or a twice larger spine NMDAR conductance to match the total NMDAR conductance in the spillover scenarios; we opted for the former. Clustered synaptic input consists of one presynaptic spike per synapse, all arriving randomly within a 30 ms interval. When no glutamate spillover is modeled, synaptic inputs arrive only to the spines. For glutamate spillover, the same presynaptic spike is delivered to the extrasynaptic NMDA conductance placed in the dendritic shaft under that spine. The study by Gao et al. (2021) models clusters ranging from 10 synapses upwards, which leaves the question of how to treat glutamate spillover for clusters with fewer than 10 synapses open. We have explored two scenarios, which we call thresholded spillover and accumulative spillover. In thresholded spillover we assume glutamate spills over suddenly and simultaneously to all eNMDARs once a glutamate threshold is reached, with a delay of 1.5 ms. The delay is due to experimental data indicating spillover currents arise after 1–2 ms of synaptic currents (Szapiro and Barbour, 2007). In accumulative glutamate spillover there is no threshold for glutamate, and a presynaptic spike delivered to a synaptic NMDA synapse is also delivered to the corresponding extrasynaptic NMDA conductance after a delay of 1.5 ms. This means that every synaptic activation causes spillover and activates the corresponding eNMDARs. An additional difference between the model in Gao et al. (2021) and our study is that in Gao et al. (2021), larger synaptic clusters also have synapses with bigger weights (parameter w) and longer transmitter pulses (parameter T_{dur}), whereas in our study we have kept these parameters fixed when increasing cluster size.

The glutamate threshold in thresholded spillover is implemented in the following way. Each NMDA synapse has

TABLE 4 Parameters for the synapse models with glutamate spillover.

	g_{\max} (nS)	w	T_{dur} (ms)	T_{\max} (mM)	α [(/ms mM)]	β (/ms)	K_d (μ M) (nS)
AMPA	1.5	0.4	1	1	12.5	0.25	20
NMDA	3.5	0.4	1	1	4	0.01	2.5
eNMDA	3.5	0.4	$50 + 200w$	0.2	4	0.01	2.5

In the cases without spillover the NMDA conductance in the spines is the same as in the cases with spillover, i.e., NMDA $g_{\max} = 3.5$ nS, and eNMDA $g_{\max} = 0$ nS. The synaptic weight parameter is w (described in the text).

a weight parameter (w), which is a dimensionless number between 0 and 1 that scales the maximal conductance of the NMDAR. We have set up the model so that a cluster of at least ten synapses with a weight >0.4 needs to be activated to reach the threshold level of glutamate at which spillover occurs, and a spike is delivered to the eNMDARs. We implemented this in NEURON by adding the NMDA synaptic weight to a variable that conceptually models the concentration of glutamate each time a spike arrives at a spine. Added weights are normalized so that ten activated synapses with a weight >0.4 reach the glutamate threshold. To implement this variable, we have in fact used a NEURON integrate-and-fire cell that generates a spike to the eNMDARs in the cluster once it reaches a threshold value of 1.

The two glutamate spillover scenarios we consider are idealized scenarios, and biological glutamate spillover is most likely somewhere in between, varying among brain regions.

2.5. Calculation of phase plots for a single membrane compartment

To investigate the role of local membrane conductances in generating plateau potentials, we performed phase plot analysis for a single electrical compartment which we aimed to make similar to a dendritic region where plateau potentials are evoked. For this purpose, in the single compartment, we inserted the ionic channels Na_f , K_{af} , K_{as} , K_{ir} , K_{dr} with maximal conductances taken from a dendritic segment in the model SPN at ~ 120 μ m from the soma, where clustered inputs are located in most of the results showing voltage traces. (In NEURON nomenclature, the dendritic segment is located at position 0.6 of dendrite 3.) Calcium channels were omitted to simplify the analysis since they have been found to have a minor effect on the voltage, and SK and BK channels have also been omitted since they depend on intracellular calcium concentration, which we do not model in the single compartment. The equation describing the current in the single membrane compartment is:

$$\sum I_{ion} + I_{NMDA} + I_{leak} + I_c = 0,$$

or

$$C_m \dot{V} = - \sum I_{ion} - I_{NMDA} - I_{leak},$$

where

$$\begin{aligned} I_{ion} &= g_{ion} G_{\infty}^{ion} (V - E_{ion}) \\ I_{NMDA} &= g_{\max}^{NMDA} g(V) (V - E_{NMDA}), \\ I_{leak} &= g_{leak} (V - E_{leak}). \end{aligned}$$

In the equations above $I_c = C_m \dot{V}$ is the capacitive current, E_{ion} , E_{NMDA} , and E_{leak} are the reversal potentials and g_{ion} , g_{\max}^{NMDA} , and g_{leak} conductances of the ionic, NMDAR, and leak currents, respectively, and $g(V)$ is the gating function describing the Mg^{2+} block. The construction of the phase plots consists of calculating the stationary membrane current as dependent on the voltage (as if a voltage value has been "on" for a long time), which means that the gating variables in the ion channels have their asymptotic, time-independent values. These are denoted as G_{∞}^{ion} . For example, for the Na_f channel, this variable is $G_{\infty}^{Na_f} = m_{\infty, Na_f}^3 \cdot h_{\infty, Na_f}$. Since, in addition to the NMDA current, bistable membrane behavior requires a shunting current of longer duration, such as K_{ir} or $GABA_B$ currents (comparable to the duration of NMDA currents), and since the available K currents in this dendritic location have small conductances relative to the NMDA conductances, we have added a leak conductance in the same order of magnitude as the NMDA current to be able to obtain phase curves in the bistable regime (Lazarewicz et al., 2006; Shoemaker, 2011; Sanders et al., 2013). As described in the Results section, the y-axes of the phase plot show $C_m \dot{V}$, i.e., they show the current:

$$I = - \sum I_{ion} - I_{NMDA} - I_{leak}.$$

2.6. Procedure for fitting the gating function parameters from experimental data

We follow the procedure presented in Jahr and Stevens (1990b) to fit the parameters α and η we list in Table 2. We first digitize the data using WebPlotDigitizer (Rohatgi, 2022) for current-voltage relationships for the available Mg^{2+} concentration from the figures in the articles, from which we then derive the conductance-voltage relationships. When converting the current amplitudes to conductance, the different voltage reversal potential across different experimental conditions was taken into account. Specifically, following Ohm's Law, $g_{NMDA} = \frac{I_{NMDA, peak}}{V_{holding} - V_{rev}}$. Conductance-voltage relationships were normalized with the maximal conductance of

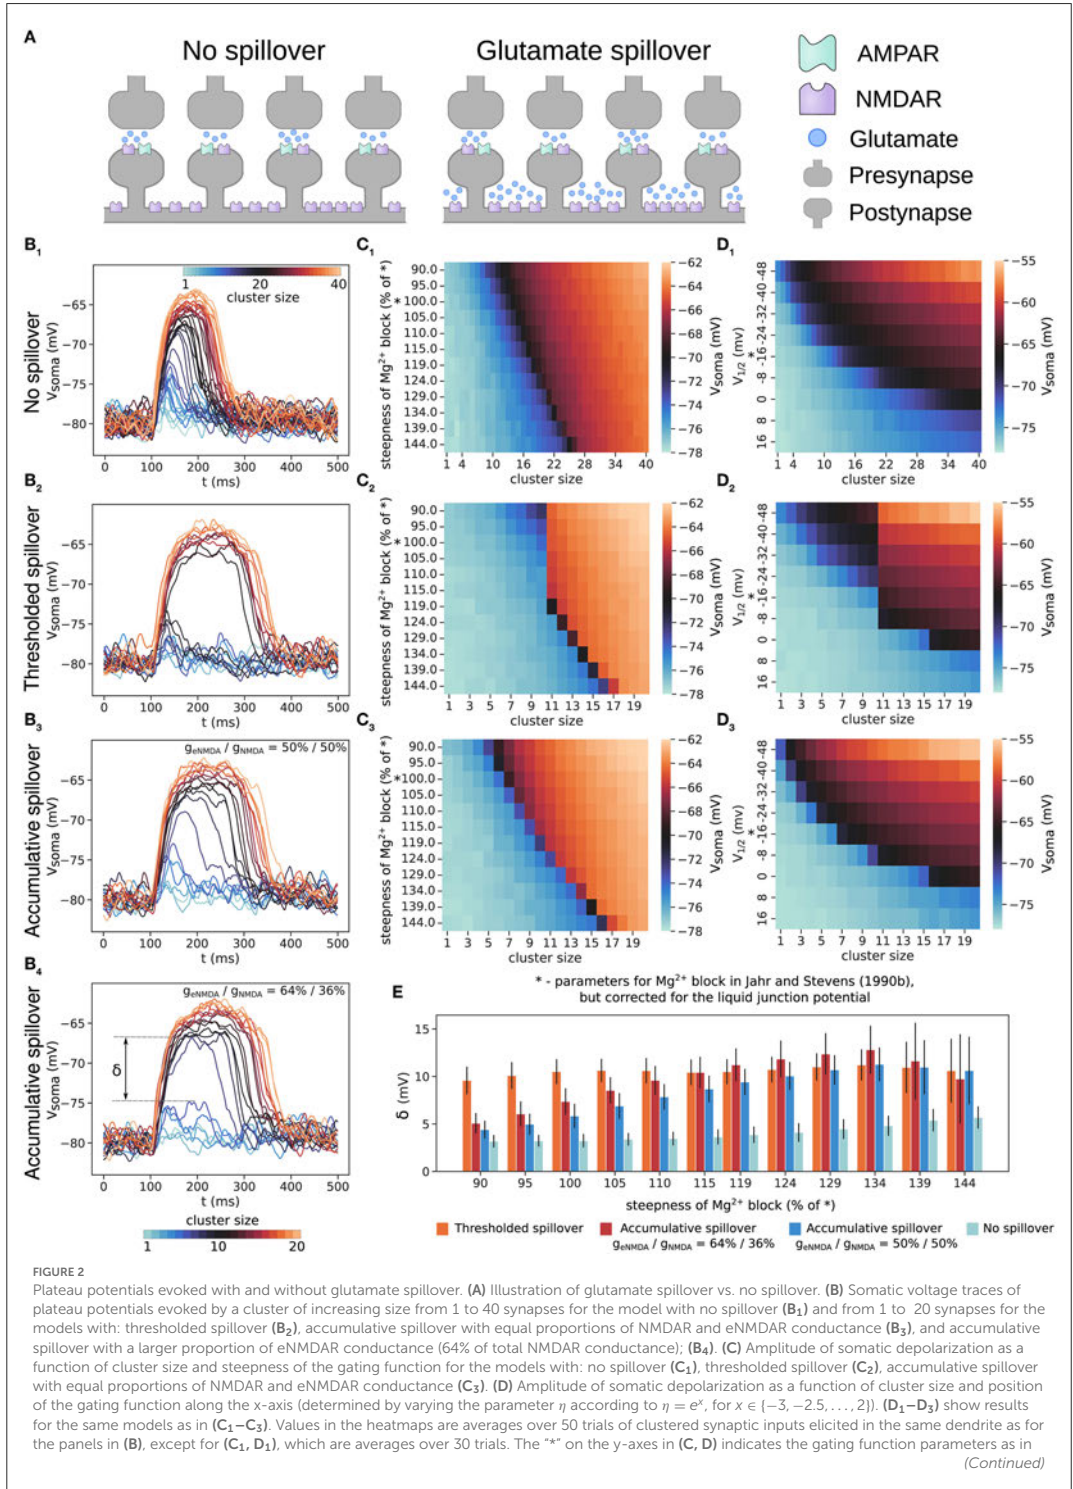


FIGURE 2

Plateau potentials evoked with and without glutamate spillover. (A) Illustration of glutamate spillover vs. no spillover. (B) Somatic voltage traces of plateau potentials evoked by a cluster of increasing size from 1 to 40 synapses for the model with no spillover (B₁) and from 1 to 20 synapses for the models with: thresholded spillover (B₂), accumulative spillover with equal proportions of NMDAR and eNMDAR conductance (B₃), and accumulative spillover with a larger proportion of eNMDAR conductance (64% of total NMDAR conductance); (B₄). (C) Amplitude of somatic depolarization as a function of cluster size and steepness of the gating function for the models with: no spillover (C₁), thresholded spillover (C₂), accumulative spillover with equal proportions of NMDAR and eNMDAR conductance (C₃). (D) Amplitude of somatic depolarization as a function of cluster size and position of the gating function along the x-axis (determined by varying the parameter η according to $\eta = e^x$, for $x \in \{-3, -2.5, \dots, 2\}$). (D₁–D₃) show results for the same models as in (C₁–C₃). Values in the heatmaps are averages over 50 trials of clustered synaptic inputs elicited in the same dendrite as for the panels in (B), except for (C₁, D₁), which are averages over 30 trials. The "*" on the y-axes in (C, D) indicates the gating function parameters as in (Continued)

FIGURE 2 (Continued)

Jahr and Stevens (1990b), but corrected for the liquid junction potential. (E) The maximal difference between two consecutive voltage traces, δ , also called the size of the voltage jump, for the four models in (B). Large jumps indicate all-or-none plateau potentials, and small jumps are consistent with graded NMDA spikes/plateau potentials. δ for no spillover was calculated taking every second voltage trace (steps of 7 nS in NMDA conductance), so that it is comparable to the δ calculated for (B₃, B₄). Results are averages over 50 trials and 11 different dendrites, with plateaus elicited at approximately the same distance from the soma, except for the model with no spillover, where they are averages over 30 trials in 11 different dendrites. Error bars represent standard deviation.

each curve, which is the point obtained for the highest applied voltage (usually +50 mV or above) when all of the NMDARs are free from Mg^{2+} . Finally, the normalized conductance-voltage relationships are fitted to Equation (4) using the Python function `scipy.optimize.curve_fit`.

3. Results

3.1. Thresholded glutamate spillover causes robust all-or-none behavior of plateau potentials

In Figure 2, we compare the plateau potentials generated by a cluster of increasing size with the two spillover models and without glutamate spillover and with the originally reported parameters for the Mg^{2+} block but corrected for the value of the junction potential (i.e., the parameters found in Jahr and Stevens, 1990b but corrected as in Ecker et al., 2020). The somatic voltage traces in Figures 2B₁, B₂ show that plateau potentials generated without glutamate spillover exhibit a graded increase in amplitude, whereas those generated with thresholded glutamate spillover exhibit an all-or-none response. As described in the Methods, in order to compare the same total NMDAR conductance between the two scenarios, the model with no spillover has a twice larger cluster size than the models with glutamate spillover.

We also varied the steepness of the gating function by varying the parameter α in Equation (4), and the amplitude of the plateau potentials thus generated is shown as heatmaps in Figures 2C₁, C₂. The all-or-none behavior can be seen as a sharp jump in the colors of the heatmaps for a small increase in cluster size (i.e., stimulus strength for evoking a plateau). Without glutamate spillover the all-or-none behavior is present only for steep sigmoidal curves, whereas with thresholded glutamate spillover, it is always present. (Steeper gating functions require more excitation to overcome the Mg^{2+} block in both cases). As described above, we also estimated the gating function parameters for available experimental data (Table 2, sigmoid curves plotted in Supplementary Figure 2). Most of these gating functions are more shallow than the originally reported one by Jahr and Stevens (1990b), and result in graded NMDA potentials in the model without spillover, but still provide all-or-none plateau potentials in the model with glutamate spillover. When varying the parameter η , which shifts the gating function along the x-axis, the same results regarding the all-or-none behavior are observed for the models with thresholded spillover and without glutamate spillover as when varying the parameter α (Figures 2D₁, D₂). Shifting the gating functions to very depolarized values disables the generation of plateau potentials, as seen from the heatmaps in Figures 2D₁,

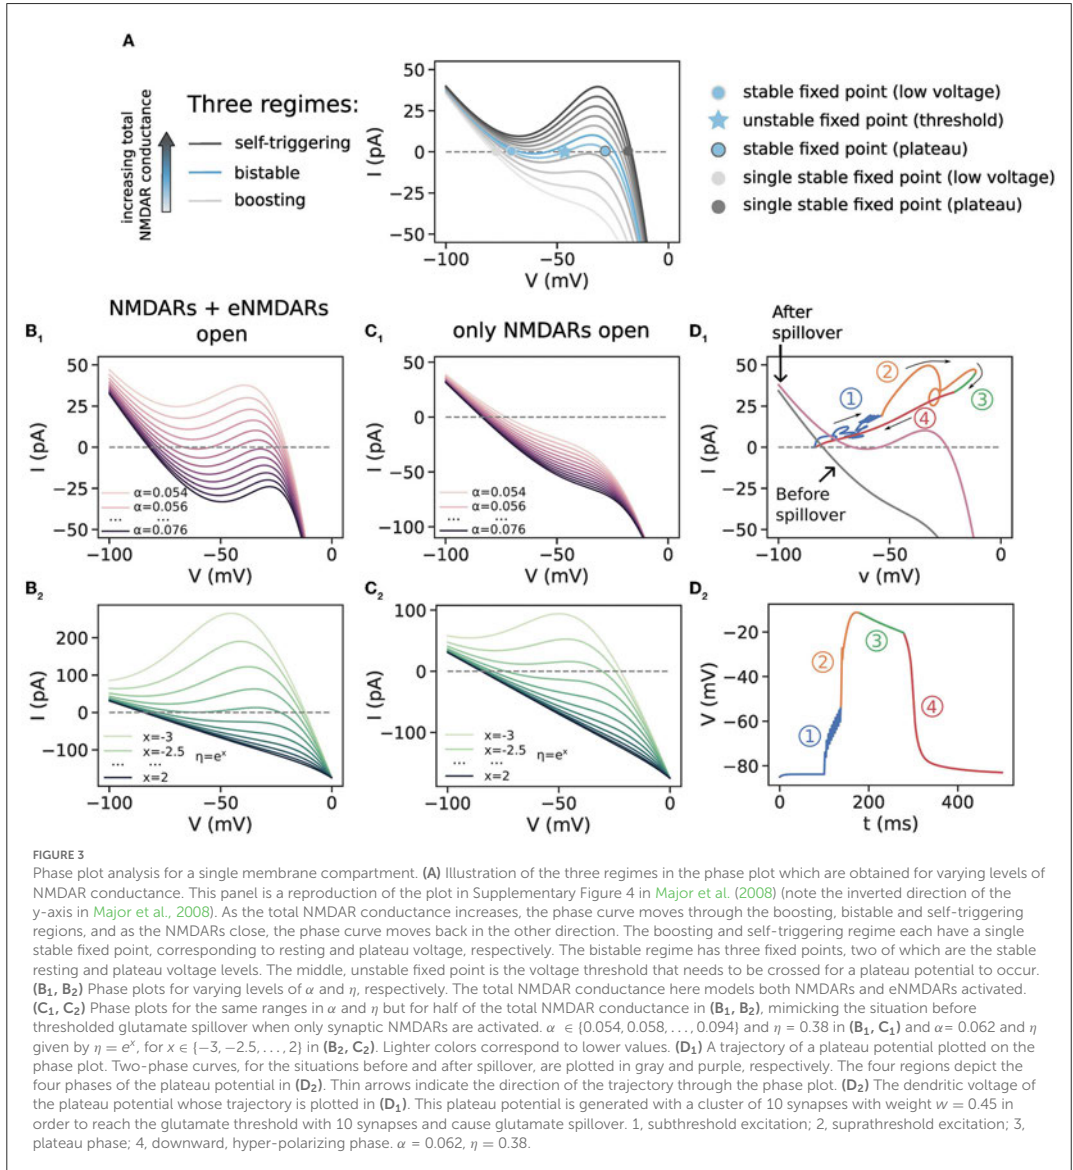
D₂. In Figure 2E, we compare the maximal difference between the somatic voltage amplitudes of two consecutive voltage traces, termed the “size of the jump,” δ (indicated in Figure 2B₄). In our view, the size of the jump, δ , is essentially proportional to the system’s nonlinearity. For example, a large sudden jump (≈ 10 mV) indicates all-or-none behavior, whereas small jumps (of a few mV) would indicate graded potentials. In our model, thresholded glutamate spillover produced voltage jumps of 8–12 mV, while no spillover produced jumps of ≈ 3 mV, typically (Figure 2B₁).

In addition, the somatic amplitude of plateau potentials generated with thresholded spillover and without glutamate spillover decays with increasing the cluster distance from the soma as shown in Supplementary Figure 3 (Major et al., 2008).

3.2. Importance of extrasynaptic NMDAR activation for all-or-none plateau potentials: accumulative vs thresholded spillover

Physiologically relevant parallel fiber stimulation protocols cause gradual accumulation of glutamate with each stimulation pulse in cerebellar extrasynaptic space (Okubo et al., 2010). To explore the effects of such more gradual spillover, we performed simulations where each activation of spine NMDARs also activates the corresponding eNMDARs in the dendritic shaft under the spine after a delay, without the need to first reach a glutamate threshold. (the accumulative spillover described in Section 2).

The results of these simulations show that accumulative glutamate spillover exhibits plateau potential behavior which is between the behavior without spillover and with thresholded glutamate spillover (Figures 2B₃, C₃, D₃). Shallow gating functions provide graded plateau potentials, but all-or-none behavior arises much sooner for increasing steepness than the in case without spillover. With thresholded spillover, clusters of up to 10 synapses only access the spine NMDAR conductance, and clusters with more than 10 synapses activate both NMDAR and eNMDAR conductances, thus producing the large voltage jump before and after spillover (Figure 2B₂). With accumulative spillover, eNMDARs are activated with a short delay after each NMDAR activation, giving rise to the more graded increase in plateau potential amplitude, as well as duration (Figure 2B₃). Moreover, for a larger proportion of eNMDAR conductance in the total conductance, the behavior is more all-or-none (Figures 2B₄, E). This suggests that there may be particular conditions under which all-or-none plateau potentials appear and that plateau potential behavior might vary among brain regions.



3.3. Spillover induces all-or-none plateau potentials by facilitating the switch to bistable or self-triggering regime

To explain the results in the previous section, we use phase plot analysis of the membrane voltage calculated for a single membrane compartment (Figure 3). Phase plot analysis reveals the role of different membrane conductances in generating nonlinear behavior such as plateau potentials. In particular, it has been found that apart from the nonlinear, voltage-dependent excitatory NMDA

current, an opposing, shunting current is also required acting at least on the longer timescale of NMDA currents. This role can be played by K_{ir} channels, GABA_B currents, a passive leak current, and/or an axial current diverting charge to the rest of the neuron (Lazarewicz et al., 2006; Shoemaker, 2011; Sanders et al., 2013).

We begin with a short description of how to read the phase plots (Figure 3A). A voltage phase plot normally contains voltage on the x-axis and the rate of voltage change (the time derivative of the voltage, \dot{V}) on the y-axis. We have instead plotted the membrane current on the y-axis, which is simply a linear rescaling

of \dot{V} with the membrane capacitance, $I = C_m \dot{V}$. The intersections of the phase curve with the x-axis correspond to fixed points in the voltage dynamics ($\dot{V} = 0$). Since our single compartment is a one-dimensional system (voltage is the only independent state variable), a point in the phase plot can move only along the phase curve. Initial voltage points that correspond to positive values of the phase curve ($\dot{V} > 0$) will get increased in time, whereas those corresponding to negative values get decreased. Voltage change stops when a point on the phase plot reaches a stable fixed point (where $\dot{V} = 0$). Points near unstable fixed points are repelled away from them. The trajectory from the multicompartment SPN model shown in Figure 3D₁, on the other hand, can take on any values in the phase plot. This is because in the recorded dendritic compartment, the gating variables of the ion channels are also time-varying, i.e., they are additional state variables in the model. We note that the phase plot is different from the commonly used and similarly named I–V plot, which can be found, for example, in the aforementioned articles measuring macroscopic NMDA currents.

Depending on the total NMDAR conductance in a dendritic region, NMDARs can exhibit three types of nonlinear dynamics, as shown in Figure 3A: boosting, bistable, and self-triggering, of which the latter two correspond to plateau potentials (Schiller and Schiller, 2001). Boosting behavior occurs for small or moderate amounts of total NMDAR conductance which is not enough to trigger plateau potentials. In this case other excitatory inputs are boosted in amplitude by activating the longer-lasting NMDA potentials. The phase curve has one stable fixed point to the left of the diagram at a low voltage. Bistable behavior is exhibited when the phase curve has an inverted-N shape with three interceptions with the x-axis (fixed points), which happens for larger amounts of total NMDAR conductance. In this case a plateau potential is generated for an excitation that exceeds the voltage threshold determined by the middle (unstable) fixed point. Such an excitation causes a fast excursion in the phase plot to the right (higher) stable fixed point, which causes the all-or-none behavior (region 2 in the phase plot in Figure 3D₁), which corresponds to the upward swing in the corresponding dendritic voltage in Figure 3D₂, and animation in Supplementary data). Even larger total NMDAR conductance gives rise to the so-called self-triggering behavior. In this case the phase curve, now situated above the bistable region, is inverted-N-shaped and has one stable fixed point to the right of the diagram which corresponds to high voltage. In this case, any excitation produces a plateau potential by carrying the voltage to the single fixed point. Thus, as the total NMDAR conductance increases with successive activations of synapses in a cluster, the phase curve sweeps potentially across all three regions in the phase plot. On the other hand, closing of the NMDARs moves the phase curve back in the other direction, collapsing the inverted-N shape back toward a straight line and terminating the plateau potential with the voltage dropping to the low-voltage, single stable fixed point, which corresponds to the resting voltage.

To elicit plateau potentials, the total NMDAR conductance must be enough to situate the phase curve in the bistable or self-triggering region (the threshold level of NMDAR conductance being at the border between the boosting and the bistable region). An example of a plateau potential generated with thresholded glutamate spillover, with the corresponding phase plot trajectory,

is given in Figures 3D₁, D₂. The gray phase curve in Figure 3D₁ is the one obtained when only synaptic NMDAR conductance is activated, whereas the purple curve represents the total NMDAR conductance with spillover, when eNMDARs are activated. After the subthreshold excitation by nine synaptic inputs (region 1), the 10th input causes spillover and the dendritic voltage crosses the threshold determined by the unstable fixed point. This initiates the self-sustaining suprathreshold excitation (region 2), which is followed by the plateau region (region 3). The end of the extrasynaptic transmitter pulse starts the downward, hyperpolarizing phase (region 4).

To examine the role of the gating function parameters α and η , we constructed phase plots with varying α and η . Figures 3B₁, B₂ are a model of the total NMDAR conductance when both NMDARs and eNMDARs are activated during spillover, as well as the case when all of this total conductance is situated in the spines of the equivalent model without spillover (Figure 2B). Varying the steepness of the Mg²⁺ block shifts the phase curve vertically, and all three phase plot regimes may be obtained (Figure 3B₁). Shallow sigmoidal curves have only one fixed point and are in the self-triggering region, whereas steep sigmoidal curves have three fixed points and are in the bistable region. Very steep sigmoidal curves are in the boosting region. In Figures 3C₁, C₂, on the other hand, the total NMDAR conductance is halved, to model only the synaptic NMDAR conductance accessible before thresholded spillover. In this case, all phase curves for the single membrane compartment are in the boosting region, regardless of the gating function steepness.

These results from the phase plots can explain the plateau potential behavior with and without spillover. Total NMDAR conductance increases with increasing cluster size—larger clusters have greater total NMDA conductance (Figure 2B, cluster size). When there is no glutamate spillover, all-or-none behavior is observed only for steep sigmoidal curves, since the corresponding phase curves are in the bistable regime, while for the shallow sigmoidal curves, all synaptic clusters of increasing size produce graded excitations as in Figure 2B₁, because the phase curve is in the self-triggering region with a single, high-voltage fixed point (Figure 3B₁). Graded excitations trigger NMDA spikes first and plateau potentials later, as the total NMDA conductance increases. On the other hand, in the model with glutamate spillover, activating only the synaptic NMDA conductance (extrasynaptic NMDARs are inactive) is insufficient to elicit plateau potentials (the phase curve is most likely in the boosting region as in Figure 3C₁). When spillover occurs, the activation of the eNMDARs provides enough NMDA conductance to elicit a plateau potential. This shifts the phase curve from the boosting to the self-triggering region for shallow gating functions and to the bistable region for steep gating functions. The duration of the phase curve in both of these regions is significantly prolonged due to glutamate spillover. This shift (and inversion) of the phase curve occurs suddenly because, as described in the Methods section, in thresholded glutamate spillover all the eNMDAR conductance is simultaneously activated. For shallow sigmoidal curves, it is this sudden inversion that causes the all-or-none plateau potential behavior. For steep sigmoidal curves, the voltage excitation also needs to be higher than the voltage threshold for a plateau to occur. In this way, glutamate spillover

causes all-or-none plateau potentials for a wide range of gating function steepness. Accumulative spillover, on the other hand, is very similar to no spillover. Each successive synaptic activation gradually increases the total NMDAR conductance, slowly moving the phase curve across the phase plot.

Varying the parameter η shifts the gating function along the x -axis. Shifting the function to higher potentials flattens (linearizes) the phase curve. While all three regimes can be obtained with various values of η , shifting the gating function to higher potentials results in only boosting behavior and disappearance of plateau potentials both with and without glutamate spillover (Figures 2D₁–D₃, 3C₁, C₂).

3.4. Long glutamate duration drives all-or-none plateau potentials

We verify that it is indeed the length of glutamate spillover that causes all-or-none plateau potential behavior by manipulating the duration of the transmitter pulse (the parameter T_{dur} , Figure 4). As explained in the Methods section, T_{dur} for the synapses onto spines is much shorter than T_{dur} for eNMDARs. Increasing T_{dur} in the case without spillover results in all-or-none plateau potential behavior. Again, increasing NMDAR conductance by increasing cluster size moves the stable fixed point further to the right in the phase plot in Figure 3A. When entering the bistable regime, the all-or-none jump appears in Figure 4A₁, and lasts because the increased T_{dur} keeps the phase curve in the bistable regime for a longer time. On the other hand, decreasing the eNMDAR T_{dur} in the case with spillover produces rather graded plateau potentials with much smaller voltage jumps (Figure 4B₁). The all-or-none behavior in Figure 4A₁ and graded behavior in Figure 4B₁ are also evident in the sharp and gradual color transitions in the respective heatmaps in Figures 4A₂, B₂, as well as in the comparison of the voltage jump, δ in Figure 4C.

3.5. Low extrasynaptic glutamate concentration can produce all-or-none plateau potentials

The extrasynaptic glutamate concentration (the parameter T_{max} for eNMDARs) used in the simulations above is 200 μM , an order of magnitude higher than the measurements reported in Okubo et al. (2010). NMDARs have high-affinity for glutamate, in the low-micromolar range, so they should be able to respond to low glutamate concentrations. We have tested the thresholded spillover model with lower values for extracellular glutamate concentration (the parameter T_{max}) and the results are shown in Figure 5. Concentrations of 10 and 20 μM , significantly higher than the NMDAR dissociation constant, $K_d = 2.5 \mu\text{M}$, produce all-or-none plateau potentials (Figures 5A, B). For concentrations of 5 μM , the size of the voltage jump is reduced, although the plateau potentials retain their all-or-none quality. Concentrations comparable to the NMDAR K_d do not activate the eNMDARs fully and produce graded plateau potential amplitude (Figures 5D, E).

3.6. All-or-none behavior of plateau potentials enables nonlinear computation, such as solving the nonlinear feature binding problem

Linearly non-separable tasks, an example of which is the nonlinear feature binding problem (NFBP) cannot be solved by the perceptron, a computational unit with a single nonlinearity, and require a network of such artificial neurons. As such, these tasks are traditionally used as a benchmark for the computational capabilities of a computational unit. Possessing dendritic nonlinearities in addition to the somatic nonlinearity, single biological neurons have been shown to be equivalent to at least a two-layer artificial neural network (Poirazi et al., 2003) and should, in principle, be able to solve such tasks (Tran-Van-Minh et al., 2015). Although it is not known whether the NFBP is solved by single neurons *in vivo*, and whether it is a relevant task for any brain region, human cortical layer 2/3 pyramidal neurons have been found to possess calcium-mediated dendritic spikes whose voltage dependency enables the dendrite exhibiting them to solve the exclusive OR (XOR) problem in simulations (Gidon et al., 2020), and the XOR problem is related to the NFBP (Cazé et al., 2013). These dendritic nonlinearities occur in the apical dendrites, while the SPN dendrites correspond to basal dendrites in pyramidal neurons, where NMDA nonlinearities occur. We, therefore, use the NFBP as a benchmark to demonstrate how all-or-none plateau potentials enable a single neuron to solve this task.

We illustrate the NFBP with an example from visual feature binding (Figure 6A). A stimulus has two features (shape and color), each of which can have two values (strawberry or banana for the shape, and red or yellow for the color), for a total of four possible combinations. Each stimulus excites the neuron with the same amount of excitation, on average. The task consists of responding to two of the feature combinations (the relevant stimuli) with a somatic spike and remaining silent for the other two (irrelevant stimuli). This means that the same amount of total excitation should be processed differently, and the plateau potentials offer a natural solution to the task: the synapses for the relevant stimuli should be clustered in separate dendrites (Figure 6B), and should elicit a plateau potential upon arrival of a relevant stimulus which will drive somatic spiking. Irrelevant stimuli would activate only half of the cluster and should not elicit a plateau potential (Tran-Van-Minh et al., 2015). This setup with thresholded glutamate spillover provides perfect performance on the NFBP (Figures 6C, D). Without spillover, the graded amplitudes of the plateaus cause somatic spiking for the irrelevant stimuli as well, and the neuron cannot perform this task (Figures 6C, E, F).

4. Discussion

In this computational study we investigated how to generate robust all-or-none dendritic plateau potentials using striatal projection neuron models as a test case. We also illustrated how such plateau potentials can enhance computation. Dendritic nonlinearities in the form of NMDA spikes and plateau potentials offer an enhanced ability for dendritic computation compared to

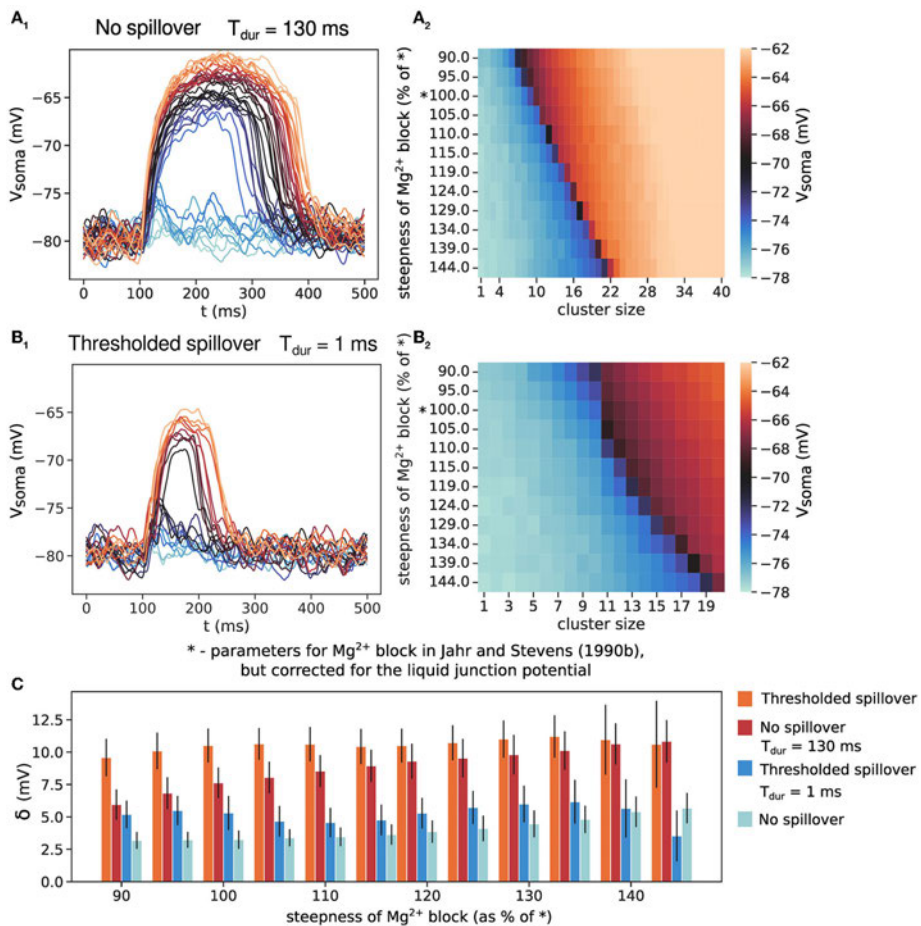
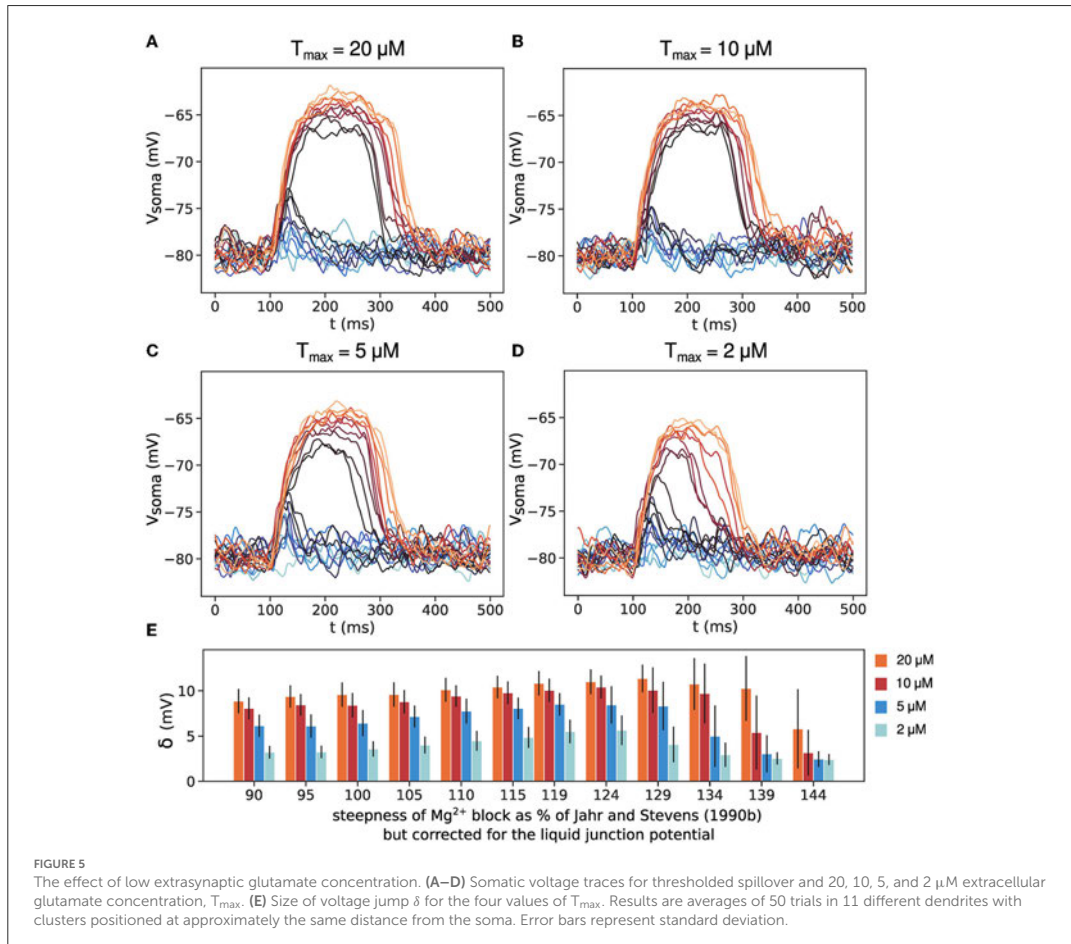


FIGURE 4

The effect of transmitter pulse duration, T_{dur} , on the all-or-none behavior of plateau potentials. (A₁) No glutamate spillover but long transmitter pulse duration produces all-or-none plateau potentials. Cluster size up to 40 synapses. (A₂) Somatic amplitude of plateau potential when varying the steepness for the model in (A₁). (B₁) Thresholded glutamate spillover and short transmitter pulse in the extrasynaptic space produces NMDA potentials with a small all-or-none jump. Cluster size up to 20 synapses. (B₂) Somatic amplitude of plateau potential when varying the steepness for the model in (B₁). The "*" on the y-axes in (A₂, B₂) indicates the gating function parameters as in Jahr and Stevens (1990b), but corrected for the liquid junction potential. (C) The size of the voltage jump, δ , for the scenarios in (A, B). Results for thresholded spillover and no spillover from Figure 2E are added for ease of comparison. Results are averages of 30 trials and 11 different dendrites for (A₁) and 50 trials and 11 different dendrites for (B₁), with clusters positioned at approximately the same distance from the soma. Error bars represent standard deviation.

the linear summation of excitatory inputs (Poirazi et al., 2003; Lavzin et al., 2012; Xu et al., 2012; Tran-Van-Minh et al., 2015; Kumar et al., 2018). By summing co-activated clustered inputs supralinearly, they can allow for the preferential detection of certain input patterns or combinations of inputs. For example, in the mouse barrel cortex, layer 4 spiny stellate neurons generate local and global multi-branch NMDA spikes that contribute substantially to the angular tuning of these neurons (Lavzin et al., 2012). In addition, distal dendrites of layer 5 pyramidal neurons in the barrel cortex might integrate correlated sensory and motor information via dendritic nonlinearities in order to produce a signal

related to object localization during active sensing tasks (Xu et al., 2012). Dendritic plateau potentials are also essential processes for triggering synaptic plasticity both *in vitro* and *in vivo* (Gambino et al., 2014; Cichon and Gan, 2015; Brandalise et al., 2016). For example, in the somatosensory cortex, rhythmic sensory whisker stimulation can induce synaptic long-term potentiation (LTP) in layer 2/3 pyramidal cells, which is triggered by plateau potentials generated through the cooperative activity of the intracolumnar lemniscal and thalamocortical paralemniscal synaptic circuitry in the absence of somatic spiking (Gambino et al., 2014). Additionally, different motor learning tasks induce dendritic calcium spikes on

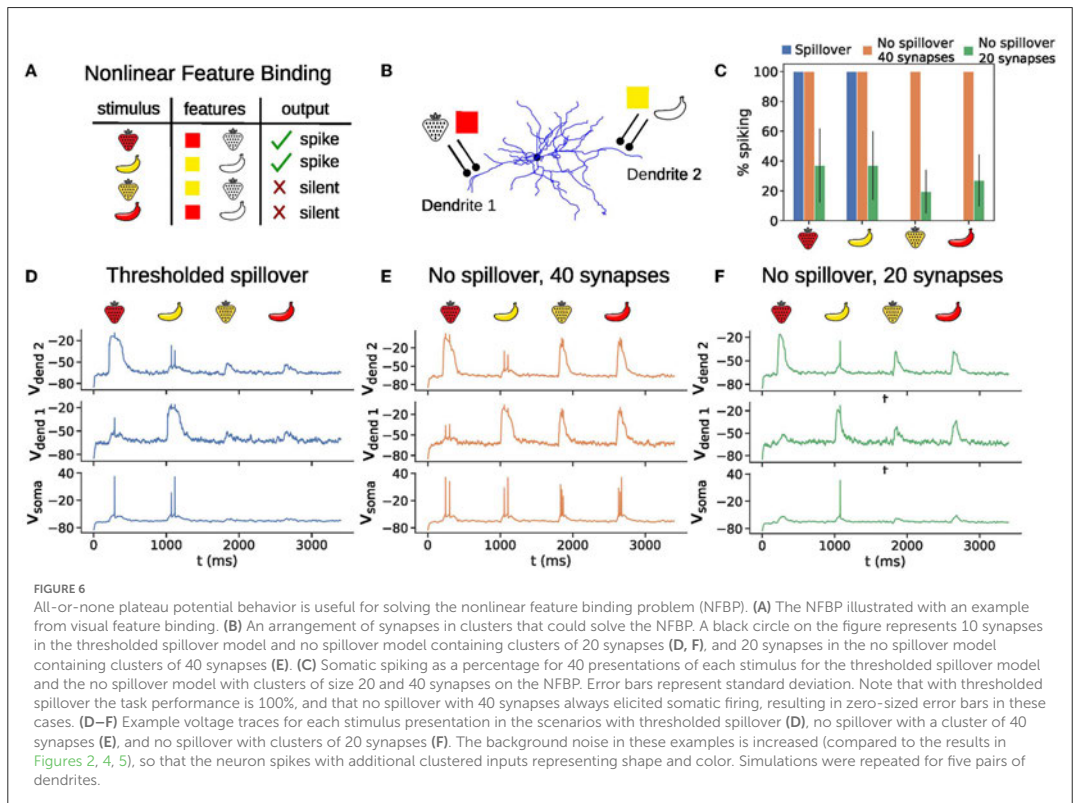


different apical tuft branches of layer 5 pyramidal neurons in the mouse motor cortex (Cichon and Gan, 2015). These task-related, branch-specific calcium spikes cause long-lasting potentiation of postsynaptic dendritic spines active at the time of spike generation, suggesting a role for dendritic nonlinearities in storing new information without disrupting previously acquired memories (Cichon and Gan, 2015). Lastly, dendritic nonlinearities have been suggested to enable a neuron to solve linearly non-separable tasks such as the NFBP (Tran-Van-Minh et al., 2015).

Experimentally evoked plateau potentials *in vitro* display all-or-none behavior, and this feature is not easily nor robustly captured when modeling. A critical ingredient for the all-or-none property is the sigmoidal shape of the Mg^{2+} block voltage dependence, described by the gating function in Equation (4). The many NMDAR isoforms show different sensitivity to Mg^{2+} , which results in a range of possible shapes of the sigmoid gating function in terms of its steepness and the voltage required to overcome the Mg^{2+} block. We have studied how the variability of these gating function

properties affects the all-or-none plateau potential behavior. We found that the all-or-none property is very sensitive to the steepness of the gating function. Shallow gating functions are in the self-triggering region of the phase plot and result in graded plateau potential amplitudes for increasing cluster size, whereas steep gating functions are in the bistable region and result in all-or-none plateau potentials (Figures 2B₁, B₂, 3B₁). Importantly, including glutamate spillover in the models provides a robust all-or-none quality of the plateau potentials for the whole tested range of Mg^{2+} block steepness.

Almost all the gating functions we fitted to available experimental data are more shallow than the originally reported sigmoidal curve by Jahr and Stevens (1990b), which already results in graded plateau potential amplitude in the SPN model (Figure 2B; Table 2). Conversely, all of the studies that model plateau potentials, except for Ecker et al. (2020), that correct for the omitted junction potential in Jahr and Stevens (1990b), have increased the steepness of the originally reported gating



function without explicitly stating a reason for it. As shown in Figure 3B₁, increasing the steepness of the gating function can move the phase curve in the bistable region, allowing to robustly generate all-or-none plateau potentials, thus offering an explanation for the steeper gating functions in modeling studies. Moreover, in light of the natural variability in the Mg^{2+} block gating function, the question remains as to how SPNs generate the all-or-none plateau behavior robustly in experiments with increasing stimulation. According to the results above, glutamate spillover is predicted to be a part of the plateau potential generation mechanism, providing robust all-or-none behavior across a wide range of slopes of the Mg^{2+} block gating function.

Whether, *in vivo* plateau potentials are graded or all-or-none is not known. However, in Gambino et al. (2014) for example, robust, seemingly all-or-none plateau potentials are evoked with co-activation of the intracolumnar lemniscal and thalamocortical paralemniscal pathways, while lemniscal pathway activation alone significantly reduces the probability of evoking plateau potentials. In this study we have illustrated the importance of such all-or-none plateau potentials in a commonly used benchmark task, the NFBP, showing that graded plateau potentials cannot perform this task compared to the perfect performance with

all-or-none plateau potentials. It is not known whether the NFBP is solved by single neurons in any brain region on a regular basis, but plateau potentials, whether used for the NFBP or in the detection/representation of single features, have been suggested to prepare the neuron to spike in a robust fashion as a response to the particular plateau-evoking input (Antic et al., 2018). *In vitro*, distally evoked plateau potentials in the basal dendrites of pyramidal neurons typically do not evoke somatic spiking by themselves (Milojkovic et al., 2004; Major et al., 2008; Gao et al., 2021), and neither do plateau potentials in SPNs (Plotkin et al., 2011; Du et al., 2017). This corresponds to the situation we modeled in Figures 2, 4, 5. However, in the presence of additional inputs or background noise, plateau potentials drive a robust somatic response to the particular input pattern that evoked them, as we exemplify in Figure 6. According to the predictions from this study, the prolonged activation of extrasynaptic NMDARs is critical for robust plateau potential generation.

Despite important advances such as the confirmation of glutamate spillover *in vivo*, the confirmation of its role in neuronal signaling in the cerebellum, and reports of extracellular glutamate concentrations in experimentally evoked spillover in brain slices in the low micromolar ranges, many important

details await further studies (Szapiro and Barbour, 2007; Okubo et al., 2010; Papouin and Oliet, 2014; Rusakov and Stewart, 2021). Some prominent missing details are: (i) the lifetime of glutamate in the extrasynaptic and perisynaptic spaces; (ii) the conditions under which possible astrocytic glutamate release and/or reversal of glutamate uptake occurs; as well as (iii) control of eNMDAR activation by the co-agonist glycine or d-serine (Malarkey and Parpura, 2008; Papouin and Oliet, 2014). Assuming high NMDAR affinity for glutamate, we have shown that 10–20 μM of extracellular glutamate can produce all-or-none plateau potentials if glutamate persists in the extracellular space for long durations, indicating that physiological concentrations may play a large role in signal integration. Recent findings that LTP boosts glutamate spillover by initiating withdrawal of perisynaptic astroglial processes suggest dynamic regulation of the synaptic and perisynaptic environment, including eNMDAR activation, for the purposes of signal processing and learning (Henneberger et al., 2020; Rusakov and Stewart, 2021). Facilitation of the prolonged eNMDAR activation would, in turn, increase the probability that the already strengthened synapses can evoke a robust dendritic plateau potential.

One direction for future work would be to investigate how the distinct gating functions of various NMDAR isoforms affect dendritic computations in different neuron types in the brain, as well as to further quantify the contribution of extrasynaptic NMDA receptors for these computations (Zhou et al., 2015). Another direction for future work would be to develop more detailed diffusion models of glutamate concentration activating extrasynaptic NMDARs for a quantitatively more precise understanding of the phenomenon in particular synapse types. However, since NMDARs have high affinity for glutamate, even low extrasynaptic glutamate concentrations also generate robust plateau potentials for a prolonged extrasynaptic glutamate signal (as shown in Figure 5). The synaptic model that we used assumes fixed transmitter concentration for the duration of the transmitter pulse both in the synaptic cleft and in the extrasynaptic space. Glutamate in the synaptic cleft rapidly reaches its peak concentration and is also quickly cleared due to diffusion and reuptake (Clements et al., 1992). In contrast to that, extrasynaptic glutamate dynamics is likely more variable. Under conditions of temporary saturation of astrocytic glutamate uptake and possibly Ca^{2+} -dependent astrocytic glutamate exocytosis and/or reverse operation of astrocytic glutamate transporters (Malarkey and Parpura, 2008), extrasynaptic glutamate concentration could remain elevated long enough to produce all-or-none plateau potentials. Further experimental studies of extrasynaptic glutamate are needed to clarify these issues in different brain regions.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found at: <https://github.com/danieltrpevski/Spillover/tree/master> and on ModelDB (McDougal et al., 2017).

Author contributions

DT designed the study, performed the simulations, and wrote the first draft of the manuscript. ZK contributed to the study's design and was involved in the simulations. IC performed the parameter estimation from experimental articles and wrote sections of the manuscript. JHK supervised the project. All authors contributed to the article and approved the submitted version.

Funding

This study was supported by the Swedish Research Council (VR-M-2020-01652), Swedish e-Science Research Centre (SeRC), EU/Horizon 2020 no. 945539 (HBP SGA3), and KTH Digital Futures.

Acknowledgments

The simulations were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) at PDC KTH partially funded by the Swedish Research Council through grant agreement no. 2018-05973. We acknowledge the use of Fenix Infrastructure resources, which are partially funded from the European Union's Horizon 2020 research and innovation programme through the ICEI project under the grant agreement No. 800858. We thank Dr. Robert Lindroos for helpful discussion and useful feedback.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fncel.2023.1196182/full#supplementary-material>

VIDEO S1

Video animation of the plots in Figures 3D₁, 3D₂. The trajectory traversing the phase plot in the upper panel corresponds to the voltage timecourse in the lower panel. For details see figure legend of Figures 3D₁, 3D₂.

References

- Antic, S. D., Hines, M., and Lytton, W. W. (2018). Embedded ensemble encoding hypothesis: the role of the "prepared" cell. *J. Neurosci. Res.* 96, 1543–1559. doi: 10.1002/jnr.24240
- Antic, S. D., Zhou, W.-L., Moore, A. R., Short, S. M., and Ikonomu, K. D. (2010). The decade of the dendritic NMDA spike. *J. Neurosci. Res.* 88, 2991–3001. doi: 10.1002/jnr.22444
- Brandalise, F., Carta, S., Helmhets, F., Lisman, J., and Gerber, U. (2016). Dendritic NMDA spikes are necessary for timing-dependent associative LTP in CA3 pyramidal cells. *Nat. Commun.* 7:13480. doi: 10.1038/ncomms13480
- Carnevale, N. T., and Hines, M. L. *The NEURON Book*. New York, NY: Cambridge University Press (2005).
- Carter, A. G., and Sabatini, B. L. (2004). State-dependent calcium signaling in dendritic spines of striatal medium spiny neurons. *Neuron* 44, 483–493. doi: 10.1016/j.neuron.2004.10.013
- Cazé, R. D., Humphries, M., and Gutkin, B. (2013). Passive dendrites enable single neurons to compute linearly non-separable functions. *PLoS Comput. Biol.* 9:e1002867. doi: 10.1371/journal.pcbi.1002867
- Chalifoux, J. R., and Carter, A. G. (2011). Glutamate spillover promotes the generation of NMDA spikes. *J. Neurosci.* 31, 16435–16446. doi: 10.1523/JNEUROSCI.2777-11.2011
- Chen, L., and Huang, L.-Y. M. (1992). Protein kinase C reduces Mg²⁺ block of NMDA-receptor channels as a mechanism of modulation. *Nature* 356, 521–523. doi: 10.1038/356521a0
- Cheng, H.-W., Rafols, J. A., Goshgarian, H. G., Anavi, Y., Tong, J., and McNeill, T. H. (1997). Differential spine loss and regrowth of striatal neurons following multiple forms of deafferentation: a Golgi study. *Exp. Neurol.* 147, 287–298. doi: 10.1006/exnr.1997.6618
- Chiu, D. N., and Carter, B. C. (2022). Synaptic NMDA receptor activity at resting membrane potentials. *Front. Cell. Neurosci.* 16:916626. doi: 10.3389/fncel.2022.916626
- Cichon, J., and Gan, W.-B. (2015). Branch-specific dendritic Ca²⁺ spikes cause persistent synaptic plasticity. *Nature* 520, 180–185. doi: 10.1038/nature14251
- Clark, B. A., Farrant, M., and Cull-Candy, S. G. (1997). A direct comparison of the single-channel properties of synaptic and extrasynaptic NMDA receptors. *J. Neurosci.* 17, 107–116. doi: 10.1523/JNEUROSCI.17-01-00107.1997
- Clements, J. D., Lester, R. A. J., Tong, G., Jahr, C. E., and Westbrook, G. L. (1992). The time course of glutamate in the synaptic cleft. *Science* 258, 1498–1501. doi: 10.1126/science.1359647
- Destexhe, A., Mainen, Z. F., and Sejnowski, T. J. (1994). An efficient method for computing synaptic conductances based on a kinetic model of receptor binding. *Neural Comput.* 6, 14–18. doi: 10.1162/neco.1994.6.1.14
- Dorman, D. B., Jkedorzevska-Szmek, J., and Blackwell, K. T. (2018). Inhibition enhances spatially-specific calcium encoding of synaptic input patterns in a biologically constrained model. *eLife* 7. doi: 10.7554/eLife.38588.022
- Doron, M., Chindemi, G., Müller, E., Markram, H., and Segev, I. (2017). Timed synaptic inhibition shapes NMDA spikes influencing local dendritic processing and global I/O properties of cortical neurons. *Cell Rep.* 21, 1550–1561. doi: 10.1016/j.celrep.2017.10.035
- Du, K., Wu, Y.-W., Lindroos, R., Liu, Y., Rózsa, B., Katona, G., et al. (2017). Cell-type-specific inhibition of the dendritic plateau potential in striatal spiny projection neurons. *Proc. Natl. Acad. Sci. U.S.A.* 114, E7612–E7621. doi: 10.1073/pnas.1704893114
- Ecker, A., Romani, A., Sárny, S., Káli, S., Migliore, M., Falck, J., et al. (2020). Data-driven integration of hippocampal cal synaptic physiology *in silico*. *Hippocampus* 30, 1129–1145. doi: 10.1002/hipo.23220
- Farinella, M., Ruedt, D. T., Gleeson, P., Lanore, F., and Silver, R. A. (2014). Glutamate-bound NMDARs arising from *in vivo*-like network activity extend spatio-temporal integration in a L5 cortical pyramidal cell model. *PLoS Comput. Biol.* 10:e1003590. doi: 10.1371/journal.pcbi.1003590
- Gambino, F., Pagès, S., Kehayas, V., Baptista, D., Tatti, R., Carleton, A., et al. (2014). Sensory-evoked LTP driven by dendritic plateau potentials *in vivo*. *Nature* 515, 116–119. doi: 10.1038/nature13664
- Gao, P. P., Graham, J. W., Zhou, W.-L., Jang, J., Angulo, S., Dura-Bernal, S., et al. (2021). Local glutamate-mediated dendritic plateau potentials change the state of the cortical pyramidal neuron. *J. Neurophysiol.* 125, 23–42. doi: 10.1152/jn.00734.2019
- García-Munoz, M., Lopez-Huerta, V. G., Carrillo-Reid, L., and Arbutnot, G. W. (2015). Extrasynaptic glutamate NMDA receptors: key players in striatal function. *Neuropharmacology* 89, 54–63. doi: 10.1016/j.neuropharm.2014.09.013
- Gidon, A., Zolnik, T. A., Fidzinski, P., Bolduan, F., Papoussi, A., Poirazi, P., et al. (2020). Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science* 367, 83–87. doi: 10.1126/science.aax6239
- Harris, A. Z., and Pettit, D. L. (2007). Extrasynaptic and synaptic NMDA receptors form stable and uniform pools in rat hippocampal slices. *J. Physiol.* 584, 509–519. doi: 10.1113/jphysiol.2007.137679
- Henneberger, C., Bard, L., Panatier, A., Reynolds, J. P., Kopach, O., Medvedev, N. L., et al. (2020). LTP induction boosts glutamate spillover by driving withdrawal of perisynaptic astroglia. *Neuron* 108, 919–936.e11. doi: 10.1016/j.neuron.2020.08.030
- Higley, M. J., and Sabatini, B. L. (2010). Competitive regulation of synaptic Ca²⁺ influx by D2 dopamine and A2A adenosine receptors. *Nat. Neurosci.* 13, 958–966. doi: 10.1038/nn.2592
- Jahr, C., and Stevens, C. (1990a). A quantitative description of NMDA receptor-channel kinetic behavior. *J. Neurosci.* 10, 1830–1837. doi: 10.1523/JNEUROSCI.10-06-01830.1990
- Jahr, C., and Stevens, C. (1990b). Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *J. Neurosci.* 10, 3178–3182. doi: 10.1523/JNEUROSCI.10-09-03178.1990
- Kastellakis, G., and Poirazi, P. (2019). Synaptic clustering and memory formation. *Front. Mol. Neurosci.* 12:300. doi: 10.3389/fnmol.2019.00300
- Kumar, A., Schiff, O., Barkai, E., Mel, B. W., Poleg-Polsky, A., and Schiller, J. (2018). NMDA spikes mediate amplification of inputs in the rat piriform cortex. *eLife* 7:e38446. doi: 10.7554/eLife.38446.018
- Lavzin, M., Rapoport, S., Polsky, A., Garion, L., and Schiller, J. (2012). Nonlinear dendritic processing determines angular tuning of barrel cortex neurons *in vivo*. *Nature* 490, 397–401. doi: 10.1038/nature11451
- Lazarewicz, M. T., Ang, C.-W., Carlson, G. C., Coulter, D. A., and Finkel, L. H. (2006). Analysis of nmda-dependent voltage bistability in thin dendritic compartments. *Neurocomputing* 69, 1025–1029. doi: 10.1016/j.neucom.2005.12.038
- Lindroos, R., and Hellgren Kotaleski, J. (2021). Predicting complex spikes in striatal projection neurons of the direct pathway following neuromodulation by acetylcholine and dopamine. *Eur. J. Neurosci.* 53, 2117–2134. doi: 10.1111/ejn.14891
- Major, G., Polsky, A., Denk, W., Schiller, J., and Tank, D. (2008). Spatiotemporally graded NMDA spike/plateau potentials in basal dendrites of neocortical pyramidal neurons. *J. Neurophysiol.* 99, 2584–601. doi: 10.1152/jn.00011.2008
- Malarkey, E. B., and Pappas, V. (2008). Mechanisms of glutamate release from astrocytes. *Neurochem. Int.* 52, 142–154. doi: 10.1016/j.neuint.2007.06.005
- McDougal, R. A., Morse, T. M., Carnevale, T., Marengo, L., Wang, R., Migliore, M., et al. (2017). Twenty years of modeldb and beyond: building essential modeling tools for the future of neuroscience. *J. Comput. Neurosci.* 42, 1–10. doi: 10.1007/s10827-016-0623-7
- McMenimen, K. A., Petersson, E. J., Lester, H. A., and Dougherty, D. A. (2006). Probing the mg2+ blockade site of an n-methyl-d-aspartate (NMDA) receptor with unnatural amino acid mutagenesis. *ACS Chem. Biol.* 1, 227–234. doi: 10.1021/cb6000944
- Milojkovic, B. A., Radojicic, M. S., Goldman-Rakic, P. S., and Antic, S. D. (2004). Burst generation in rat pyramidal neurons by regenerative potentials elicited in a restricted part of the basilar dendritic tree. *J. Physiol.* 558, 193–211. doi: 10.1113/jphysiol.2004.061416
- Nowak, L., Bregestovski, P., Ascher, P., Herbet, A., and Prochiantz, A. (1984). Magnesium gates glutamate-activated channels in mouse central neurones. *Nature* 307, 462–465. doi: 10.1038/307462a0
- Oikonomou, K. D., Short, S. M., Rich, M. T., and Antic, S. D. (2012). Extrasynaptic glutamate receptor activation as cellular bases for dynamic range compression in pyramidal neurons. *Front. Physiol.* 3:334. doi: 10.3389/fphys.2012.00334
- Oikonomou, K. D., Singh, M. B., Sterjanaj, E. V., and Antic, S. D. (2014). Spiny neurons of amygdala striatum, and cortex use dendritic plateau potentials to detect network UP states. *Front. Cell. Neurosci.* 8:292. doi: 10.3389/fncel.2014.00292
- Okubo, Y., Sekiya, H., Namiki, S., Sakamoto, H., Inuma, S., Yamasaki, M., et al. (2010). Imaging extrasynaptic glutamate dynamics in the brain. *Proc. Natl. Acad. Sci. U.S.A.* 107, 6526–6531. doi: 10.1073/pnas.0913154107
- Papouin, T., and Oliet, S. H. R. (2014). Organization, control and function of extrasynaptic nmda receptors. *Philos. Trans. R. Soc. B Biol. Sci.* 369:20130601. doi: 10.1098/rstb.2013.0601
- Petralia, R. S. (2012). Distribution of extrasynaptic NMDA receptors on neurons. *Sci. World J.* 2012:267120. doi: 10.1100/2012/267120
- Plotkin, J. L., Day, M., and Surmeier, D. J. (2011). Synaptically driven state transitions in distal dendrites of striatal spiny neurons. *Nat. Neurosci.* 14, 881–888. doi: 10.1038/nn.2848
- Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal neuron as two-layer neural network. *Neuron* 37, 989–999. doi: 10.1016/S0896-6273(03)00149-1

- Poleg-Polsky, A. (2015). Effects of neural morphology and input distribution on synaptic processing by global and focal NMDA-spikes. *PLoS ONE* 10:e0140254. doi: 10.1371/journal.pone.0140254
- Qian, A., Buller, A. L., and Johnson, J. W. (2005). Nr2 subunit-dependence of NMDA receptor channel block by external Mg^{2+} . *J. Physiol.* 562, 319–331. doi: 10.11113/jphysiol.2004.076737
- Qian, A., and Johnson, J. W. (2006). Permeant ion effects on external Mg^{2+} block of nr1/2d NMDA receptors. *J. Neurosci.* 26, 10899–10910. doi: 10.1523/JNEUROSCI.3453-06.2006
- Retchless, B. S., Gao, W., and Johnson, J. W. (2012). A single GluN2 subunit residue controls NMDA receptor channel properties via intersubunit interaction. *Nat. Neurosci.* 15, 406–413. doi: 10.1038/nn.3025
- Rhodes, P. (2006). The properties and implications of NMDA spikes in neocortical pyramidal cells. *J. Neurosci.* 26, 6704–6715. doi: 10.1523/JNEUROSCI.3791-05.2006
- Rohatgi, A. (2022). *Webplotdigitizer: Version 4.6*.
- Rusakov, D. A., and Stewart, M. G. (2021). Synaptic environment and extrasynaptic glutamate signals: the quest continues. *Neuropharmacology* 195:108688. doi: 10.1016/j.neuropharm.2021.108688
- Sanders, H., Berends, M., Major, G., Goldman, M. S., and Lisman, J. E. (2013). NMDA and gabab (kir) conductances: the “perfect couple” for bistability. *J. Neurosci.* 33, 424–429. doi: 10.1523/JNEUROSCI.1854-12.2013
- Schiller, J., and Schiller, Y. (2001). NMDA receptor-mediated dendritic spikes and coincident signal amplification. *Curr. Opin. Neurobiol.* 11, 343–348. doi: 10.1016/S0959-4388(00)00217-8
- Sharma, G., and Stevens, C. F. (1996). A mutation that alters magnesium block of N-methyl-D-aspartate receptor channels. *Proc. Natl. Acad. Sci. U.S.A.* 93, 9259–9263. doi: 10.1073/pnas.93.17.9259
- Shindou, T., Ochi-Shindou, M., and Wickens, J. (2011). A Ca^{2+} threshold for induction of spike-timing-dependent depression in the mouse striatum. *J. Neurosci.* 31, 13015–13022. doi: 10.1523/JNEUROSCI.3206-11.2011
- Shoemaker, P. A. (2011). Neural bistability and amplification mediated by NMDA receptors: analysis of stationary equations. *Neurocomputing* 74, 3058–3071. doi: 10.1016/j.neucom.2011.04.018
- Szapiro, G., and Barbour, B. (2007). Multiple climbing fibers signal to molecular layer interneurons exclusively via glutamate spillover. *Nat. Neurosci.* 10, 735–742. doi: 10.1038/nn1907
- Tran-Van-Minh, A., Cazé, R. D., Abrahamsson, T., Cathala, L., Gutkin, B. S., and DiGregorio, D. A. (2015). Contribution of sublinear and supralinear dendritic integration to neuronal computations. *Front. Cell. Neurosci.* 9:67. doi: 10.3389/fncel.2015.00067
- Xu, N., Harnett, M. T., Williams, S. R., Huber, D., O'Connor, D. H., Svoboda, K., et al. (2012). Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature* 492, 247–254. doi: 10.1038/nature11601
- Zhou, X., Chen, Z., Yun, W., Ren, J., Li, C., and Wang, H. (2015). Revisiting the function of extrasynaptic nmda receptor in excitotoxicity running title: NMDA receptor and excitotoxicity. *Neuroscientist* 21:337. doi: 10.1177/1073858414548724

