



# 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.

Simulated cognitive topologies:  
automatically generating highly  
contextual maps for complex journeys

Lucas Godfrey



THE UNIVERSITY  
*of* EDINBURGH

Thesis submitted in fulfilment of the requirements for the degree of  
Doctor of Philosophy to The University of Edinburgh, 2019













# Declaration

I declare that this thesis has been composed solely by myself and that it has not been submitted, either in whole or in part, in any previous application for a degree. Except where otherwise acknowledged, the work presented is entirely my own.

Lucas Godfrey  
December 2019



# Abstract

As people traverse complex journeys, they engage in a number of information interactions across spatial scales and levels of abstraction. Journey complexity is characterised by factors including the number of actions required, and by variation in the contextual basis of reasoning such as a transition between different modes of transport. The high-level task of an A to B journey decomposes into a sequence of lower-level navigational sub-tasks, with the representation of geographic entities that support navigation during, between and across sub-tasks, varying relative to the nature of the task and the character of the geography. For example, transitioning from or to a particular mode of transport has a direct bearing on the natural level of representational abstraction that supports the task, as well as on the overall extent of the task's region of influence on the traveller's focus. Modern mobile technologies send data to a device that can in theory be context-specific in terms of explicitly reflecting a traveller's heterogeneous information requirements, however the extent to which context is explicitly reflected in the selection and display of navigational information remains limited in practice, with a rigid, pre-determined scale-based hierarchy of cartographic views remaining the underlying representational paradigm.

The core subject of the research is the context-dependent selection and display of navigational information, and while there are many and varied considerations in developing techniques to address selection and display, the central challenge can simply be articulated as how to determine the probability, given the traveller's current context, that a feature should be in the current map view. Clearly this central challenge extends to all features in the spatial extent, and so from a practical perspective, research questions centre around the initial selection of a subset of features, and around determining an overall probability distribution over the subset given the significance of features within the hierarchically ordered sequence of tasks.

In this thesis research is presented around the use of graph structures as a practical basis for modeling urban geography to support heterogenous

selections across viewing scales, and ultimately for displaying highly context-specific cartographic views. Through an iterative, empirical research methodology, a formalised approach based on routing networks is presented, which serves as the basis for modeling, selection and display.

Findings are presented from a series of 7 situated navigation studies that included research with an existing navigation application as well as experimental research stimuli. Hypotheses were validated and refined over the course of the studies, with a focus on journey-specific regions that form around the navigable network. Empirical data includes sketch maps, textual descriptions, video and device interactions over the course of complex navigation exercises. Study findings support the proposed graph architecture, including subgraph classes that approximate cognitive structures central to natural comprehension and reasoning. Empirical findings lead to the central argument of a model based on causal mechanisms, in which relations are formalised between task, selection and abstraction.

A causal framework for automatically determining map content for a given journey context is presented, with the approach involving a conceptual shift from treating geographic features as spatially indexed records, to treating them as variables with a finite number of possible states. Causal nets serve as the practical basis of reasoning, with geographic features being represented by variables in these causal structures. The central challenge of finding the probability that a variable in a causal net is in a particular state is addressed through a causal model in which journey context serves as the evidence that propagates over the net. In this way, complex heterogeneous selections for interactive multi-scale information spaces are expressed as probability distributions determined through message propagation.

The thesis concludes with a discussion around the implications of the approach for the presentation of navigational information, and it is shown how the framework can support context-specific selection and disambiguation of map content, demonstrated through the central use case of navigating complex urban journeys.

# Lay Summary

Urban navigation in large cities often involves interaction with navigational information in a variety of representational forms. While at a high-level a person's task is travelling from A to B, in practice journeys are structured as sequences of sub-tasks, for example finding the correct bus stop, finding the correct entrance etc. The nature of each of these tasks, and the context of a task in terms of both the broader journey and the particular character of the geography, all contribute to the information needs of the traveller with respect to scale, focus, level of detail and the overall level of abstraction. While modern mobile navigation technologies provide a technical foundation for surfacing information to a traveller that specifically relates to their current information needs, in practice the extent to which context is explicitly reflected in the selection and display of navigational information remains limited, with a rigid, pre-determined scale-based hierarchy of cartographic information persisting as the underlying basis of representation.

This thesis presents research into the design of data structures and various techniques that may serve as a practical basis for modeling urban geography, with a view to supporting the selection and display of navigational information in a way that formally expresses the information needs of the traveller with respect to the specific context of each sub-task. The research is based on an iterative, empirical research methodology and concludes with the presentation of a formalised approach to data modeling, selection and display.

Findings are presented from a series of 7 situated navigation studies that included research with an existing navigation application as well as experimental research stimuli. Hypotheses were validated and refined over the course of the studies, with a focus on the cognitive representation of a journey as a hierarchically structured sequence of regions. Empirical data includes sketch maps, textual descriptions, video and device interactions over the course of complex navigation exercises. Study findings support the proposed approach which is based on abstract formalisations that serve to approximate cognitive structures. These structures are shown to be central to

comprehension and reasoning within the situated context of navigation tasks, but more specifically, that they may be used as a formal basis for automatically selecting and displaying information. For example, the relationship between a primary action such as making a right turn, and the grouping and eventual disambiguation of geographic features that support the execution of the action. Empirical findings lead to the central argument of a model based on causal mechanisms.

An approach to structuring geographic data is presented, with the central challenge being expressed in terms of the probability that a feature should be present in a particular cartographic view and at a particular level of abstraction, given the current task and the overall journey context.

The thesis concludes with a discussion around the broader implications of the approach for the presentation of navigational information, and a discussion around how the model supports context-specific selection and disambiguation of map content, drawing on the central use case used throughout, namely navigating complex urban journeys. Complexity is characterised by factors including the number of actions required, and by variation in the contextual basis of reasoning such as transitions between transport modes.

# Acknowledgements

Thanks to my mother for her support on the long winding path to both the beginning and end of a PhD.

I would like to give particular thanks to my supervisor William Mackaness, to Chris Speed, and to Jeremy Morley.

From the University of Edinburgh I would like to thank Bruce Gittings, Owen Macdonald, Neil Stuart, James Coburn, Eduardo Serafin, Eric Laurier and Marc Metzger.

From the Ordnance Survey I would like to thank Nick Groome, Stefano Cavazzi, Shaun Murray-Brown and Tony Joyce, and I give thanks to both the Ordnance Survey and the Engineering and Physical Sciences Research Council for supporting this research.

For various discussions over the years I thank Rob Kitchin, Sabine Timpf, Phil Bartie, Aaron Quigley and Leonidas Doumas.

I also give thanks to Andrea Ballatore and all organisers of the Conference on Spatial Information Theory, past and future.

Finally, I would like to take a moment to remember Professor Jon Oberlander. Together with Chris Speed, Jon's vision led to my initial motivation for coming to Edinburgh. An inspiring person who I knew all too briefly.





# Contents

Abstract	iv
Lay Summary	vi
Acknowledgements	viii
List of Figures	xv
List of Tables	xxvii
Notation	xxix
<b>1 Introduction</b>	<b>1</b>
1.1 Thesis overview	2
<b>2 Background</b>	<b>5</b>
2.1 Continuous experience and situated representation	5
2.2 The structure of experience	7
2.3 The event horizon	9
2.4 The spatial context of experience	10
2.5 Concluding comments	12
<b>3 Research aim and objectives</b>	<b>15</b>
3.1 Event structure as a basis for the context-specific selection of geospatial data	15
3.2 Multi-cut selections as a practical basis for selection over a graph of features	21
3.3 Core research aim	25
3.4 Initial hypothesis	25
3.5 Research objectives	26
3.6 Scope of research	29
<b>4 Review</b>	<b>33</b>
4.1 Events as the cognitive basis for episodic experience and spatial goal acquisition	33

4.2	The schematisation of events and the formal representation of spatial knowledge _____	47
4.3	Information supporting goal acquisition across navigational episodes _____	55
4.4	Concluding comments _____	64
<b>5</b>	<b>Methodology _____</b>	<b>67</b>
5.1	The ontological basis of investigation _____	67
5.2	Situated versus laboratory studies _____	68
5.3	Observation, analysis and interpretation _____	70
5.4	Epistemological foundations of technology research _____	78
<b>6</b>	<b>Technical overview _____</b>	<b>81</b>
<b>7</b>	<b>Situated navigation studies _____</b>	<b>93</b>
7.1	Research design _____	93
7.2	Research stimuli _____	101
7.3	Concluding comments _____	122
<b>8</b>	<b>Summary of observations _____</b>	<b>125</b>
8.1	Observations on performance and design objectives _____	125
8.2	Empirical observations overview _____	128
8.3	Concluding comments _____	162
<b>9</b>	<b>Analysis and findings _____</b>	<b>165</b>
9.1	Part 1: The resource profile of cognitive extension _____	165
9.2	Part 2: Representing the structure of experience across navigation episodes _____	182
9.3	Part 3: Determining content through context-specific regionalisation _____	195
9.4	Concluding comments _____	201
<b>10</b>	<b>Model _____</b>	<b>203</b>
10.1	The causal basis of cognitive topology _____	204
10.2	Graph-based knowledge representation _____	207
10.3	A causal framework for determining the resource profile of cognitive extension _____	212
10.4	Graph data model _____	230
10.5	Concluding comments _____	239

---

<b>11</b>	<b>Discussion</b>	<b>241</b>
11.1	Data, modeling and use cases	241
11.2	Cartography and interaction	243
11.3	Concluding comments	246
<b>12</b>	<b>Conclusion</b>	<b>249</b>
12.1	Core research aim and objectives	249
12.2	Final comments	255
	<b>References</b>	<b>257</b>
	<b>Appendix 1: example study data</b>	<b>279</b>
1.1	Figure 50	280
1.2	Figure 52 (rotated)	282
1.3	Figure 52	286
1.4	Figure 53	287
1.5	Figure 54 (rotated)	288
1.6	Figure 55 (rotated)	292
1.7	Figure 56 (rotated)	294
1.8	Figure 58 (rotated)	297
1.9	Figure 59 (rotated)	299
1.10	Figure 60 (rotated)	303
1.11	Figure 61 (rotated)	304
1.12	Figure 62 (rotated)	306
1.13	Figure 67 (rotated)	308
1.14	Figure 68 (rotated)	310
1.15	Figure 68 (rotated)	312
1.16	Figure 74	314
	<b>Appendix 2: Navigation Errors</b>	<b>317</b>
	<b>Appendix 3: S3 Research Stimulus Views</b>	<b>323</b>
	<b>Appendix 4: IJC Pre-Print</b>	<b>327</b>



# List of Figures

1	Conceptual model illustrating two distinct overlapping episodes in memory (Eichenbaum et al., 1999).	<b>11</b>
2	Open Street Map graph data model, subgraph example using Neo4j (Neo4j, 2019).	<b>16</b>
3	In panel a) an event horizon of a hierarchical graph, and in b) multiple levels of abstraction in a single horizon leading to global context and local detail in the same selection (Quigley, 2001).	<b>17</b>
4	Quadtree layout (far left), then 5th, 4th and 3rd level horizon views of power system distribution network sparse matrix data, represented as a hierarchical graph (Quigley, 2001).	<b>18</b>
5	Example output for a journey across Edinburgh, with shifts in the form of representation driven by 'walk to bus' and 'bus to walk' transitions (Mackaness, Quigley and Tanasescu, 2011).	<b>19</b>
6	(left) 3D surface visualisation of an event horizon of an underlying graph based on two focal regions, and (right) an event horizon visualisation for multiple regions of local detail in the same selection (Quigley, 2001).	<b>20</b>
7	(top) a multi-cut selection from three source vertices of depth $k$ showing the underlying spatial index and the selected cells, and (bottom left) the same multi-cut selection showing the edges (blue vertical lines) from three source vertices extending 'down' from depth $k$ and edges from two vertices extending down from depth $k-1$ , and (bottom right) a 3D surface of the event horizon corresponding with the selection.	<b>22</b>

8	Schematic of the cognitive representation of a navigable environment as consisting of a set of place representations and the distances and directions between them (O’Keefe, 1998).	<b>39</b>
9	Illustrative example of the anchor-point view of spatial knowledge containing three types of action context, their spatial connectivity at multiple levels of abstraction and the generalised regions associated with each (Golledge, 1999).	<b>43</b>
10	(A) connections between distinct local charts, where places are linked across non-overlapping environments, (B) solving navigational problems by combining charts, and (C) efficient computation of a route (a-d) based on hierarchical knowledge of adjacent charts and the places contained within them (Poucet, 1993).	<b>44</b>
11	(a) A manifold of charts representing Southern California, and (b) its hierarchy. F is a reference point in a chart of {A, B, C}, G is a reference point in {E,D}, and H is a reference point of a global chart including F and G. Local charts may overlap. $X_i$ is the ratio between the number of branches at the $(i - 1)$ th and $i$ th level (Yoshino, 1991).	<b>45</b>
12	Hierarchical basis of navigation strategy: ‘fine-to-coarse’ route planning for two route types (‘B’ and ‘C’) where type refers to starting position relative to the global environmental structure (Wiener, Schnee and Mallot, 2004).	<b>46</b>
13	Schematic geometry’ of the ‘Enge’ station in Zurich, (Rüetschi and Timpf, 2004).	<b>50</b>
14	Hierarchical structure representing information about a decision point in the long-term memory of NAVIGATOR. The labels on the links between decision-points indicate orientation, number of steps, and street information (Gopal, Klatzky and Smith, 1989).	<b>54</b>
15	‘Interactive spider map’ including routing overview and focal topographic detail (Gonçalves Santos, 2019).	<b>61</b>

16	Disambiguation of route-based information including landmarks and decision-point (Zhao, Li and Li, 2011).	<b>62</b>
17	Schematisation of a route between two cities including journey-specific landmark, route and survey information (Schwering et al., 2019).	<b>62</b>
18	‘Focus + context’ Taipei metro map example (Wu, 2014).	<b>63</b>
19	The Naked City (Debord and Jorn, 1957).	<b>72</b>
20	(left) linear order relations and (right) sketch map analysis of ‘both-sided’ linear order relations and topological relation of landmark (‘library’) within a city block element (Wang and Schwering, 2015).	<b>74</b>
21	(a) landmarks and city-blocks in sketch and geo-referenced maps, (b) relations between city-blocks and landmarks using the RCC11, RCC5, and the coarse IA relations, (c) diagram representing possible alignments (1 and 2) of landmarks in city-blocks (CB1, CB2) with landmarks in city-blocks (CB1’, CB2’) in geo-references map (Jan et al., 2017).	<b>74</b>
22	Sketch maps and digitised schematics (Field, O’Brien and Beale, 2011).	<b>75</b>
23	Normalised image space in a mobile device context with $c$ denoting the image space centre.	<b>84</b>
24	a) key components of multi-scale view space, including the world space $W_i$ , a square subset of world space for the current view $A_j$ , the geographic location of the centre of the current view $v$ , the centre $c$ of the image space for the current view $I_i$ , and the effective ‘tile 0’ image space $I_0$ ; b) example selection where the current view corresponds with image space $I_0$ ; c) example selection after a ‘zoom’ interaction to a view corresponding with image space $I_i$ . Axis $x$ = easting, $y$ = northing, $\phi$ = map scale.	<b>85</b>
25	Locations of view event onsets before and after straight-line shortest distance processing. 1b shows locations after processing for a region including the journey origin of two study routes, with 1a showing the raw GPX traces for the	<b>86</b>

	same region across groups for study route 1 (red hues) and for study route 2 (purple hues). 2b shows view event onsets around the Edinburgh Haymarket train station for both study routes, with 2a showing raw GPX traces for participants across study route 2 only (purple hues).	
26	View event onset locations after shortest distance processing for study groups 2 and 3 (top – study route 1), and for study groups 5, 6 and 7 (bottom – study route 2), visualised over the underlying paths of traversal.	86
27	Example complex view curve $Q_i$ over a time interval, projected into Euclidean space given world space $W_i$ . Axis $x$ = easting, $y$ = northing, $\phi$ = map scale.	87
28	A road network that persists in view space from views at numerical scale A to views at scale C, with a disambiguation of features at scale B for illustration.	89
29	Differential scaling over an extent based on the disambiguation of focal regions and the deformation of metric relations. Vertical lines mark the Easting of the location of the primary nodes for the journey context. The result of zoom interactions are shown over three time intervals from $t_1$ to $t_4$ , with scale increasing for region $RA$ but not for region $RB$ , i.e. the distance between $c_{RA}$ and $d_{RA}$ increases, but the distance between $d_{RA}$ and $c_{RB}$ does not.	90
30	Example stills from navigation exercise video recordings.	94
31	The City of Edinburgh study route comparison between (top) study route 1 and (bottom) study route 2, visible through participant GPX traces.	98
32	Study route comparison - traversed paths over OS Highways centreline data for (top) study route 1; and (bottom) study route 2.	99
33	Journey overviews for both study routes.	100
34	Group 2 example video stills for the walking phases of study route 1, Participants G2_1 to G2_12.	100

35	SR1 overview, with the origin region denoted by a green bounding rectangle and the destination region denoted by a blue bounding rectangle.	<b>101</b>
36	SR2 overview, with the origin region denoted by a green bounding rectangle and the transfer region denoted by a blue bounding rectangle.	<b>101</b>
37	Google Maps views from G2 screen recordings.	<b>102</b>
38	Google Maps: (left) disambiguation of the entity ‘The Booking Office’; (middle) the same extent for the same route as the figure to the left; and (right) vector data rendering example (Google, 2018).	<b>103</b>
39	Lothian Buses application (Transport for Edinburgh, 2018).	<b>104</b>
40	40 S1 interactions for participant G3_10.	<b>107</b>
41	S1 interactions for participant G3_10 plus left-right region scaling markers.	<b>108</b>
42	Panel b) S1 metric distortions given an origin region focus (slider interaction to right of centre), showing local detail and global context in the same view; and panel a) the metric extent of regions illustrated over participant GPX traces for the study route.	<b>109</b>
43	Low detail views from study group 6 screen recordings.	<b>111</b>
44	A series of views from study group 6 screen recordings for the study route 2 origin region.	<b>112</b>
45	A series of views from study group 6 screen recordings for the study route 2 transfer region.	<b>112</b>
46	A series of views from study group 6 screen recordings for the study route 2 destination region.	<b>113</b>
47	The sequence of predetermined views in S3, as used for navigation by participants in study group 7 (see Appendix 3 for larger illustration).	<b>114</b>

48	Views in S3, as used for navigation by participants in study group 7, showing the decomposition of the first textual route direction into a sequence of 4 descriptions corresponding with a higher detail cartographic view.	<b>114</b>
49	Situated navigation studies equipment for data capture: mobile phone, wireless audio transmitter and receiver, video camera.	<b>115</b>
50	Examples of participant written directions for (left) G1_11, and (right) G2_8 (see Appendix 1.1 for larger illustration).	<b>117</b>
51	Example participant sketches: (top-left) G6_5, (top-right) G7_1, (bottom-left) G1_6, (bottom-right) G2_1 (see Appendix 1.2 for larger illustration).	<b>117</b>
52	Research study sketch maps: (top, left to right) Group 1, participants 1 to 12; Group 2, participants 1 to 12; Group 3, participants 1 to 12; and (bottom, left to right) Group 4, participants 1 to 8; Group 5, participants 1 to 8; Group 6, participants 1 (see Appendix 1.3 for larger illustration).	<b>119</b>
53	Research study raw data from directions exercise for a) study group 1; b) study group 2 ; c) study group 3; d) study group 4; e) study group 5; f) study group 6; and g) study group 7 (see Appendix 1.4 for larger illustration).	<b>120</b>
54	Example sketches showing the central role of linear representations in a) G1_11, b) G2_5, c) G3_2 and d) G5_5 (see Appendix 1.5 for larger illustration).	<b>130</b>
55	Addition of road network geometry after prompting shown in green, examples shown for a) G1_5, and b) G2_1 (see Appendix 1.6 for larger illustration). Bounding boxes (pink) simply used to highlight the road features in question.	<b>131</b>
56	SR1 study participant features and directions data (colour corresponds with sketch maps colour conventions). In panel a) an overview is shown, with panels b), and c) showing two example regions at a larger scale (see Appendix 1.7 for larger illustration) Relation between features, directions and structure of network shown through spatial distribution.	<b>132</b>

57	Region types implied by empirical data and derived from the underlying network (G2 data used for illustration). Circular regions derived from the location of network intersections, elongated regions formed around the navigable path connecting sequential intersections, and a broader region created by the overall extent of the journey. Labelled features are members of this third region type.	<b>133</b>
58	Outlier examples of landmarks not adjacent to route but explicitly conveyed in a) G2_1 (the 'smoke stack'), and b) G1_4, Edinburgh Castle (see Appendix 1.8 for larger illustration).	<b>134</b>
59	Regionalisation conveyed through sketch representations, marked by approximate bounding boxes (pink) - examples a) G2_10, b) G2_3, c) G4_1, and d) G5_4 (see Appendix 1.9 for larger illustration).	<b>136</b>
60	Survey-like sketch map representation, participant G1_10 (see Appendix 1.10 for larger illustration).	<b>137</b>
61	Metric distortion of spatial structure shown through sketch examples from participants a) G1_7, and b) G2_6 (see Appendix 1.11 for larger illustration).	<b>138</b>
62	Two example sketches, each based on a highly abstract regionalisation of the environment a) G6_2, and b) G3_5 (see Appendix 1.12 for larger illustration).	<b>139</b>
63	a) View curves for G2 (red) shown with the image space implied by the view centre (blue), and b) a view curve highlighted along the 'crossbar' of the 'A' form. Axis $x$ = easting, $y$ = northing, $\phi$ = map scale.	<b>145</b>
64	Clustering of study participant mobile view event onsets for G2. Regions A through F denote approximate groupings based on distinct causal relations between journey phase, physical environment and participant information needs. Hue of point objects and vertical markers denotes participant.	<b>146</b>
65	Image spaces for G2 participant map interactions where the point objects denote image space centres and the squares denote image space extents over the underlying geography.	<b>147</b>

66	Complex view curve example from participant G2_4. Participant location marked as 'current position' on the plane formed of easting and northing axes.	147
67	G3 example sketches based on exposure to prototype S1 (see Appendix 1.13 for larger illustration).	149
68	SR2 sketch map examples from participants a) G4_5, and b) G5_4 (see Appendix 1.14 for larger illustration).	156
69	Examples of the explicit conveyance of hierarchical representations in participant understanding, a) G7_2, and b) G7_5 (see Appendix 1.15 for larger illustration).	160
70	SR1, G2 view event onsets, where colour denotes the study participant.	166
71	SR2, G5 view event onsets, where colour denotes the study participant.	167
72	G2 views, shown by the implied boundaries and centres of the square subsets of world space implied by each map view, hues denote participant.	168
73	View curves in simplified view space for G2 participant map interaction, with the 3 images simply showing the set of views curves for G2 from different 'camera' angles.	168
74	Examples of participant representations of the indoor environment 'Waverley Train Station': a) G7_4, b) G5_4, c) G6_3, d) G4_1, e) G4_6, f) G4_8, g) G5_2; and h) participant G4_5's representation of the same indoor environment including the path of the route and the 'escalator' entity (see Appendix 1.16 for larger illustration).	178
75	Participant G6_2 travels down to incorrect platforms, back up to walkway, then 'back' down to main concourse.	179
76	Disambiguation of representations in example output from (Mackaness, Quigley and Tanasescu, 2011), including low detail selection across the extent (a), topographic features	183

	adjacent to the route (b, d), and a highly schematised view of the bus journey (c).	
77	Patterns of event segmentation inferred through G2 directions exercise textual data, overlaid on participant GPX traces and the OS Highways network (Ordnance Survey, 2019) 1a, turn; 1b, crossing, turn; 1c, crossing; 1d, crossing; 1e, crossing; 1f, crossing; 1g, transfer; 2a, transfer, crossing; 2b, crossing; 2c, turn; 2d, origin.	<b>188</b>
78	Participant textual directions example, G1_10.	<b>192</b>
79	TURN event schema example illustration.	<b>193</b>
80	1a-g, clustering of features based on the structure of the navigable network for G1-G3 participant data; 2, the region shown in ‘1’ within the context of the overall SR1 extent; and 3, the ‘residential area’ along the path towards the journeys final turning action. Axis $x$ = easting, $y$ = northing, $\phi$ = map scale.	<b>199</b>
81	Illustration characterising the spatial extent of region classes in simplified view space. Axis $x$ = easting, $y$ = northing, $\phi$ = map scale.	<b>200</b>
82	Features in multiscale view space as probability distribution determined through message propagation. Axis $x$ = easting, $y$ = northing, $\phi$ = map scale.	<b>205</b>
83	Views of two phase regions for two journey contexts with a shared origin.	<b>206</b>
84	As in Figure 83, views of two phase regions for two journey contexts with a shared origin, but with each showing a view at a smaller image space width, i.e. a larger scale.	<b>206</b>
85	A causal net consisting of three unobserved variables and the observed evidence $\xi_i$ .	<b>210</b>
86	Two graph patterns based on topological distance, where circular objects represent graph vertices.	<b>215</b>

87	Compound schema template 1: features that segment experience serve to further decompose the network representation leading to a new, more detailed network representation. In panel b) a feature that is understood to be a proxy for a potential action is used to derive the new $k$ -level network node $C$ .	<b>216</b>
88	Compound schema templates that provide the basis of key graph patterns, with a) template 2 as nested paths, b) template 3 as a region based on associating an action with a location, and c) template 4 as a region created by the sequential ordering of actions.	<b>217</b>
89	Features not adjacent to a route but lying within a spatial region formed around a navigable path.	<b>218</b>
90	SR2 journey overview showing the 5 phase regions.	<b>220</b>
91	Origin region in SR2 journey context, with ACTION_REGION feature clusters highlighted in red.	<b>220</b>
92	Propagating evidence $M_R^i$ for 5 scales in view space (a-e) starting at variable $x_{NK}^i$ , i.e. the variable representing the $i$ -th $k$ -level network node.	<b>222</b>
93	Illustration of the relationship between a graph vertex (circular object) that stores feature data (e.g. geometry data) and the representation of the feature as a variable (square object).	<b>223</b>
94	SR2 journey overview showing the 5 phase regions (the same illustration as Figure 70).	<b>227</b>
95	Result of variable activation for the SR2 journey context visualised in simplified view space with example feature data.	<b>227</b>
96	A view for a given journey context as a lateral slice across probabilistic view space (red) for a given view centre, i.e. a local chart in the manifold over the underlying set of features.	<b>228</b>
97	Data model.	<b>231</b>
98	Graph data model (white), plus temporary vertices and edges for the current session (orange).	<b>232</b>

99	Illustration of ACTION_REGION and PATH_REGION spatial extent.	<b>238</b>
100	‘Crossfade’ interaction concept at two zoom levels: (left) context-specific selection and (right) generic content fading into view.	<b>244</b>
101	Contextual zoom interaction concept example. Single interaction decomposes both graphical and textual information.	<b>245</b>
102	Sequential activation interaction concept, with views reflecting progression along a route shown ‘before and ‘after’ a user has moved a slider control.	<b>246</b>



# List of Tables

1	Study exercise types	<b>95</b>
2	Study groups summary information	<b>97</b>
3	Prototype research stimuli summary	<b>105</b>
4	Participant data	<b>116</b>
5	Navigation errors summary	<b>127</b>
6	Typical structure and content of textual directions	<b>140</b>
7	Example representations anchored by transfer actions	<b>141</b>
8	Example textual directions based on distinctive characteristics of the geography of the network	<b>141</b>
9	Example ‘chunked’ actions sequences conveyed through textual directions	<b>142</b>
10	Characteristics of information interaction during journey planning for study groups G2, G3, G5, G6 and G7	<b>143</b>
11	Example feedback for prototype research stimulus S1 during initial exposure by participant G3_1	<b>150</b>
12	Text directions – average count of the directions per study group plus the standard deviation for the group	<b>161</b>



# Notation

## Parametric view space

$\mathbb{R}^2$	Two-dimensional real coordinate space
$\mathbb{R}^3$	Three-dimensional real coordinate space
$W$	World space, denoting a subset of the plane
$A$	A square subset of world space
$I_0$	The unit square image space
$\rho$	The van Wijk and Nuij metric
$\mathbf{c}$	The centre of image space
$w$	Width of a given $A$
$\mathbf{x}$	Two-dimensional geometric vector

## View curves

$t$	Time instant
$\mathbf{v}$	Two-dimensional vector denoting the location of a view event onset
$\mathbf{c}$	Three-dimensional vector denoting the centre of a view
$V$	View event as a tuple containing $\mathbf{v}$ and $\mathbf{c}$ and $t$
$\mathcal{Q}$	View curve as a sequence of views over a time interval

## Property graph

$\mathcal{G}$	Labelled property graph
$\mathcal{V}$	Set of vertices
$\mathcal{E}$	Set of edges
$\lambda$	Function to assign labels
$\mu$	Function to assign properties
$L_v$	Set of vertex labels
$L_e$	Set of edge labels
$P_v$	Set of vertex properties

$P_e$	Set of edge properties
$\zeta$	Path in $\mathcal{G}$
$\Psi$	Traversal in $\mathcal{G}$
$\mathcal{L}_k$	Vertex depth

## Graph queries and results

$OD$	Origin-destination tuple
$\mathcal{R}$	Route sequence given $OD$ as an indexed family
$\delta$	Phase region
$\mathcal{S}$	Subgraph returned from $\mathcal{G}$ , given $\mathcal{R}$

## Causal net

$\Omega$	Causal net
$\alpha$	Function for creating an indexed family of variables
$\beta$	Function for creating an indexed family of arcs
$X$	Indexed family of variables
$Y$	Indexed family of arcs
$\mathbf{f}$	Finding vector
$\mathbf{M}$	Propagation matrix
$\Xi$	Evidence
$\theta$	Variable properties
$\psi$	Activation trace over $\Omega$
$T$	Set of functions for sending messages over an arc
$\Pi$	Set of functions for receiving a message over an arc
$\Gamma$	Set of functions for propagating evidence
$\omega$	The union of each $\psi$ for the current $\Xi$
$\sigma$	Function to map the resource profile $\omega$ to multi-scale view space





# 1 Introduction

*“...space is constantly [brought] into being as an incomplete solution to an on-going relational problem...”*

(Kitchin and Dodge 2011, p71)

In his seminal study on urban form, Alexander noted that the tendency to conceptualise phenomena through rigid hierarchical structures can be understood as driven by the convenience of the designer, as opposed to reflecting the experiential reality of perceived relations between entities across spatial scales. An entity within the overall fabric of the city may be a discrete, spatially bounded object, however it may also play very different roles in its interaction with people through the flows of daily life. Where the entity plays a role in cognitive processes supporting navigational decision-making, the context of the overall journey can further modulate the nature of the associative links between the entity and the wider geography, causing differences in the attentional processes that form the basis of spatial knowledge acquisition.

A ‘tree’ structure as a basis for knowledge representation can be understood as rigidly hierarchical, with a key property of this pervasive form being the separation of entities given their path of antecedence. In the prevailing design paradigm of mobile cartography, the vector tile approach can be understood as a direct example of such a structure. While we can ‘zoom’ from tile level 0 to tile level 21, we cannot view the tile at level 0 and a tile at level 21 simultaneously. Here the overall form of representation is fundamentally constrained by this underlying architecture. As Alexander noted however, in relation to the ‘semilattice’ structure that supports the representation of overlapping sets:

*“...a tree based on 20 elements can contain at most 19 further subsets of the 20, while a semilattice based on the same 20 elements can contain more than 1,000,000 different subsets.”*

(Alexander 1966, p5)

To view the selection and display of map content through this lens is to grasp the basic essence of the research presented in this thesis. The cartographer makes design decisions with respect to both the features that are visible at different scales and the form of their representation, however a rigid hierarchical approach means that the context of the traveller can only be explicitly reflected in navigational information in a very limited manner. The task of automatically configuring map content relative to a particular journey context can be understood as the central subject of the present research.

## **1.1 Thesis overview**

The overall approach to exploring the problem introduced in the opening paragraphs is to study interaction with navigational information in terms of situated cognitive processes, with a view to shaping automated methods such that they result in information that effectively addresses the context-specific information needs of the traveller. In the chapter that follows this introduction, background is given around a number of concepts that are central to the overall approach. In particular, the role of event perception in everyday experience, and the way that distinct experiential events can be used as way of understanding shared experiential patterns in navigation. The notion of perceptual events as part-whole structures is introduced, and initial linkages formed between cognitive theories and research into the modelling and selection of spatial data. In the third chapter the overarching aim of the research is presented, along with research objectives grouped by theme. This follows a detailed discussion around the initial research premise including analysis of the proposals found in (Mackaness, Quigley and Tanasescu, 2011). Here concepts from the preceding chapter are elaborated on, in particular the notion of drawing explicit relations between perceptual event structures and multi-scale data. A brief introduction to graph data models is provided, with further explanation in Chapter 10. The chapter concludes with some brief statements around the scope of the research. Chapter 4 then provides a review of key literature, with the review being divided into three sections. The first addresses the state of knowledge with respect to the cognitive representation of spatial experience and our natural navigational strategies. In the second section the review turns to a brief introduction of formalised theories of ‘cognitive mapping’, and finally the third section provides background in terms of navigational information itself, for example distinctions and overlap between knowledge acquisition through textual and graphical representations.

The methodology is presented in Chapter 5, including a summary of the ontological position of the research. A key decision was to conduct situated field studies as opposed to lab-based experiments, and so this chapter explains the reasoning behind

this decision, as well as a discussion around the associated methods. In Chapter 6 a technical overview is provided that aims to act as a reference for the reader if required. This overview covers the technical basis of interactive prototypes, and introduces conventions used in data visualisations in subsequent chapters.

Chapter 7 is a summary of the overall research design, the research activities, and the forms of data collected through participant observation across 7 studies. The research stimuli used in field studies are discussed, as well as additional detail around research objectives. The following chapter then provides a summary of observations made over the course of research. Chapter 9 provides detailed analysis divided into 3 parts. In the first part it is argued that the central task in automating map content selection should be understood as predicting the state of a coupled system. This finding contrasts with the initial hypothesis, with a key theme throughout the analysis being the importance of incremental knowledge acquisition that is not well supported by reducing the number of map views, even if selection within each view is context-specific based on our knowledge of the current journey. Building on the notion of predicting the state of a coupled system, the first part of the analysis also presents the key finding that modeling and selection of data is best understood in terms of the aim of simulating a resource profile of cognitive extension. This finding and the associated terminology are explained in detail, following relevant background information provided earlier in the thesis. In particular this analysis and argument builds on the extended cognition literature, with geospatial features cast as informational resources within extended cognitive processes.

Parts 2 and 3 of the analysis chapter then provide further specificity as to the concepts and classes found to be central to the approach. Image-schema theory plays an important role here, with cognitive topology being a consistent theme throughout. The analysis chapter lays the foundation for the presentation of the model in Chapter 10. A graph data model, together with causal methods of subgraph processing form the practical basis of formalisation, and so additional background is provided around these topics before the specification is presented. This takes the form of classes and methods, with the functionality of causal reasoning methods summarised in the explanation of a `VARIABLE_ACTIVATION` algorithm. The ultimate outcome of these techniques can be understood in terms of predicting a resource profile of cognitive extension, where informational resources are determined across multi-scale view space. Here the approach can simply be articulated in terms of determining, given a particular journey context, the probability that a feature should be in an output map view. Following the presentation of the model, a brief discussion is provided around outstanding questions, considerations for further research, and possible methods of utilising the approach in the display of information to users.



## 2 Background

A brief summary is provided of key research that has brought clear and practical explanation as to the cognitive basis of everyday experience, with a focus on the segmentation of activity into distinct experiential events that naturally form part-whole structures (are ‘partonomic’). A key insight is that events are partonomic in that they exist at various resolutions of attention such that an event can be an abstract representation of a number of events over shorter timescales that are experienced as being ‘within’ the higher-level event. Navigational episodes are discussed in terms of A to B journeys, with complexity arising out of the subdivision of activity into perceptually distinct units. The chapter aims to provide essential background and introduces the notion of event structure over an episode as being an *event horizon*. The proposal points toward the potential for formally representing events, with the ultimate aim of driving context-specific information selection to support spatial goal acquisition.

### 2.1 Continuous experience and situated representation

The segmentation of continuous experience is fundamental to attention and memory, which is to say, the ability to perceive a meaningful beginning and end to a segment of experience through time. It has been shown this basic ability shapes the development of our understanding of the world and the underlying mechanisms through which we act and plan for the future. A wide variety of cues can come to shape our perception of these ‘event boundaries’, with distinct events forming across spatial and temporal scales. Event overlap leads to the formation of connected sequences of experiential events, for example where an action or sequence of actions forms part of two distinct events, or where a perceptual experience driven by a distinctive cue occurs in more than one event. The natural ‘chunking’ of sequences of actions can be seen to lead to the formation of distinct events, with sequences of

events being bound together within broader motivational contexts, forming experiential episodes.

Central to the emergence of these phenomena are goals, which shape motivational hierarchies such that perception and action are contextualised at varying levels of abstraction over an episode. While an episode may emerge based on a series of distinctive events that were not in fact driven by internal goals, invariably the formation of episodes is driven by experience between an initial distinctive state and the acquisition of a goal. While there is often a spatial context to goal acquisition, clearly in the case of navigation, progression through navigable space is fundamental to the experience of transitioning between states.

In navigation the *cause* of an episode is a motivational state that leads to the formation of a goal, where achieving the goal requires a change of spatial position. High-level goals at the episodic level contextualise this primary goal and shape attention and cognitive control across subtasks resulting in the contextual basis for action. For example: through choice of transport mode, choice of route, speed of movement and so on.

The cognitive representation of navigable environments can therefore be seen as emerging out of situated action and interaction over series of distinct events, with information needs relative to the pursuit of goal acquisition similarly emerging out of situated experience over time.

### 2.1.1 PERCEPTION THROUGH SITUATED INTERACTION

Object perception arises from the perception of boundaries in space that lead to the identification and categorisation of perceptually cohesive structure. However, it has been argued the processes that support identification and categorisation of objects are more appropriately characterised as mental trajectories in continuous experience as opposed to the retrieval of static representations (Spivey and Dale, 2006). It would seem in fact, that the concept of memory refers to the thresholds that determine the paths of these mental trajectories through time. This view fits both with the interpretation of experimental observations of subjective experience, and with objective observations of shared patterns of neural activation across distinct experiential scenarios.

An example described by Rowlands is the experience of viewing a graphical depiction of a circle. The representational theory of mind would explain the identification and categorisation of the ‘circle’ entity in terms of matching perceptual experience to a static ‘circle’ representation in memory. The view

through the lens of situated cognition however would suggest that the crucial process of perceptual experience in this case is in fact *the pattern of change* in the movement of the eye, with the experience of a circle being more appropriately characterised as situated interaction between agent and environment as opposed to an environmental cue simply triggering the retrieval of an object representation from memory (Rowlands, 2005).

An example that is perhaps clearer is the case where there are two similar objects being perceived together such as the ‘cup’ and ‘bowl’ case discussed by Spivey. The visual similarity in these objects leads the viewer’s ‘thought trajectory’ to gradually arrive at a unitary response category, where categorisation is equivalent to the path of thought trajectory as crossing a threshold (Spivey, 2007). In this simplified example the basis of feature vector space is simply aspect ratio, i.e. width relative to height, which is a key distinguishing feature in the visual categorisation of cups versus bowls. If the task is to simply categorise an object as either a cup or a bowl, the argument is that this task-context discourages a person from attending to ‘gradations’ in the stimulus given the need for a discretised response. If however the task is to rate how ‘bowl-like’ each stimulus is, the underlying basis of attention shifts such that gradations in the stimulus are attended to differentially from the first task-context.

These object-based examples are really just a part of the picture however. In terms of the hierarchical view of motivation and control, attending to a discrete object happens at a low-level within the overall structure of perception and behaviour. Of greater concern to the study of navigation is how the segmentation of continuous experience occurs at the level of *events* rather than objects.

## 2.2 The structure of experience

As discussed in (Williams, Postans and Hodgetts, 2019), we parse continuous experience into meaningful chunks or segments, and groups of segments become bound together such that tighter bindings within a segment group lead to the coherent experience of events through time (Levine *et al.*, 2019). While spatial boundaries can serve to trigger the formation of event boundaries in memory (Bullens *et al.*, 2010), particularly in the case of compartmentalised environments, we do not solely rely on spatial boundaries, and our more general method of chunking experiential units is by way of action sequences, as discussed in the previous subsection.

So it can be seen that behavioural patterns are found in consequential experiences such as turns and decision-points (Brunec, Moscovitch and Barense, 2018), as well as structural regularities such as physical boundaries. Studies in both rodent and human navigation have shown that decision-points are subject to specialised neural coding, and that frequency of decision-points is a fundamental determinant of the neural and cognitive representation of an environment (Brunec, Moscovitch and Barense, 2018). Also significantly, it has been shown that cognitive agents segment routes during navigation (Alexander and Nitz, 2017), with decision-points acting as perceptual boundaries between route segments, and with route representations and the granularity of decision-making corresponding to a hierarchical structure of goals and sub-goals (Brunec, Moscovitch and Barense, 2018).

So an issue of central importance to cognitive agents during navigation is the segmentation of continuous experience, with segmentation into distinct events unfolding in a way that reflects our hierarchically structured representation of the spatial context of action sequences (Levine *et al.*, 2019). It should also be noted that the detection of event boundaries is understood to be a key trigger for attention (Radvansky and Zacks, 2017), and that our cognitive representations in a more general sense are fundamentally structured by these event boundaries (Speer, Zacks and Reynolds, 2007).

As discussed in (Zacks and Tversky, 2001), the perception of events can be interpreted as being analogous to object perception, but with the additional need to perceive boundaries in the temporal dimension as well as the spatial boundaries that support perception of discrete objects. Objects play a causal role in goal acquisition at the level of actions, with events playing a causal role across episodic experience, and activity being subdivided across the episode based on context-specific segmentation. As noted in (Zacks and Tversky, 2001), people actively modulate this process of subdividing the stream of activity, for example by increasing the subdivision into a larger number of distinct events when unpredictability in the contextual basis of action increases.

The generalisation of object perception to event perception also extends to the perceptual experience of events at varying levels of abstraction, and to partonomic relations between events across these levels of granularity. Goals drive the formation of task structure, with tasks then leading to a segmentation of experience that results in partonomically organised sets of events. The authors draw on the example in Newell and Simon, used to support the explanation of the computational model of planning expressed through their General Problem Solver:

*“...a high-level goal is recursively broken down into lower level goals until the outstanding lower level goals can be satisfied. For example, ‘getting to Boston’ can*

*be decomposed into ‘getting to the airport’ plus ‘getting on a plane’. ‘Getting to the airport’ can in turn be decomposed into ‘getting to the bus stop’ plus ‘getting on a bus’. This decomposition continues until the outstanding goals are represented as primitives...”*

(Zacks and Tversky, 2001, p14) with reference to (Newell and Simon, 1972)

From the above quote it is clear how complex navigational episodes may be studied in relation to goal decomposition and the natural segmentation of experience over time, with partonomic structure of events and goals ‘mirroring’ each other, from the episodic level, down to the level of ‘basic actions’.

### 2.3 The event horizon

In his 2012 paper, Radvansky discusses event models as a subtype of mental models (Johnson-Laird, 1983) that “...capture the contents and structure of the events that people experience...” (Radvansky, 2012). Work to formalise knowledge about event segmentation and the relationship between perceived event boundaries, comprehension and memory, led to the proposal of a framework called the Event Horizon Model. This framework supports the analysis and prediction of the relationships between context, behaviour, event segmentation and cognitive performance, for example in terms of distinguishing between event boundaries that improve or impair the commitment of experience to memory. This framework was discussed further in (Radvansky and Zacks, 2014), building on the description of event structure in (Zacks and Tversky, 2001), with the authors presenting the Event Horizon Model as ‘a theory for how experience is encoded in long-term memory’. More recent work to refine the theory can be found in (Radvansky and Zacks, 2017), with further explanation around the role of causal relations and the link between event boundaries and attention.

While the work of Radvansky and Zacks specifically concerns the set of events that are naturally perceived through continuous experience and that ultimately shape the content of episodic memory, in (Mackaness, Quigley and Tanasescu, 2011) it was noted that more generally an event horizon is:

*“...a boundary in spacetime, beyond which events cannot affect an outside observer...”*

(Mackaness, Quigley and Tanasescu, 2011, p213)

Research into event segmentation naturally deals with experiences through time and across spatial regions, for example experiments that have explored the

relationship between object memory formation and the traversal of compartmentalised environments (Pettijohn and Radvansky, 2016). In (Mackaness, Quigley and Tanasescu, 2011) however, the notion of an event horizon was considered in the domain of navigation specifically, with goal acquisition across the set of experiential events over the course of a navigational episode requiring external information relative to a sequence of distinct task-contexts. Here the proposal was to apply previous work by Quigley to the problem of modelling, selecting and displaying navigational information from multiple levels of representational abstraction given the specific context of the current journey. In other words explicitly representing information to support different navigational task-contexts in the same view, where the transitions between different forms of representational abstraction are ‘smoothed’ so as to lead to well-formed Gestalts. The underlying ambition expressed in the paper was summarised as:

*“...cognitively smooth transitions from the schematised to the topographic, varying in theme, at different levels of detail, and depending on task...”*

(Mackaness, Quigley and Tanasescu, 2011, p217)

The background to this proposal and further explanation as to how the conceptual approach could be applied to navigational information will be discussed further in Chapter 3.

## 2.4 The spatial context of experience

The picture of our natural representations of navigable space is one of local representations with metric qualities integrated into larger, looser representations that afford flexibility in the face of inconsistent spatial knowledge at varying resolutions of attention. Here space, time and memory are intertwined, with results from the study of the organisation of spatial and temporal knowledge pointing toward generalised mechanisms for the organisation of memory (Eichenbaum, 2017).

In reviewing work on human ‘spatial cognitive maps’ and exploring evidence for cognitive maps in animals, Poucet noted that a topological framework affords a “...loose, yet operational... representation that specifies the connectivity of space and its overall arrangement...”, while also allowing for the integration of metric information that enhances the ‘grain’ of the representation (Poucet, 1993, p174). Of central importance here is the observation that there is a fundamental distinction between the way that spatial experience is encoded within a ‘place’, and the way encoding occurs during traversals of the environment between distinctive places. Here ‘place’ refers to some locally bounded region of the environment that warrants

a transition to these higher resolution encoding mechanisms, often involving successive rotational movements to allow for the combination and association of a number of local ‘views’, and always involving an increase in attentional focus, particularly at decision-points (Janzen and Van Turenout, 2004).

In terms of a cognitive framework that might underpin these mechanisms of knowledge acquisition, (Eichenbaum *et al.*, 1999) argued that spatial representations are in fact constructed by way of ‘memory spaces’ that are not specific to spatial experience, but that given the importance of spatial context, result in the kind of




Figure 1 Conceptual model illustrating two distinct overlapping episodes in memory (Eichenbaum *et al.*, 1999).

spatial representations we observe both in terms of neural activation and in terms of overall behaviour. Here then the core neural and cognitive mechanisms of constructing episodic memory through time lead to the construction of coherent spatial representations incorporating various types of knowledge across various granularities. As opposed to being distinctly ‘spatial’ in the manner implied by the ‘spatial cognitive map’, this approach casts the underlying basis of spatial experience as forming around a shared experiential template in which distinct behaviourally significant events are bound into event sequences, and event sequences in turn are bound by overlap through shared events, echoing research discussed previously.

It is important to note that in the conceptualisation illustrated in Figure 1 the progression through an episode reflects the sequence of encoding through time as opposed to strictly relating to an encoding during movement through space, however clearly if a person is moving through space, these two aspects are both implicit in the sequencing of events. Nevertheless, it is important to note here that the ‘node’ represents the cell or cells that are activated during the memory of both episodes and serve to bind episodic memories into a broader experiential context through the representation of shared event types. In the context of navigation clearly such shared activations could relate to a ‘place’, however the ‘nodes’ are not spatial per se. In this conceptualisation then, spatial representations emerge out of “...a

constrained framework of associations...” (Eichenbaum *et al.*, 1999, p220) through embodied interaction with spatial environments. An important point here is that in addition to providing a framework that addresses the inconsistencies in the ‘spatial’ view of the cognitive map while still capturing its empirically proven elements, this account is also consistent with the broader understanding of the function of the hippocampus in terms of processing context (Nadel, 2008), and congruent with the finding that the hippocampus is “...selectively involved in behaviours that require working memory...” irrespective of whether the behaviour is specifically ‘spatial’ in nature (Olton, Becker and Handeimann, 1979, p313).

Work to understand how the cognitive basis of navigation might be specified (Leutgeb *et al.*, 2005) and conceptualised in the way discussed above is ongoing, for example through the novel experiment discussed in (Warren *et al.*, 2017) and elaborated on further in (Warren, 2019). In this study human subjects navigated both a Euclidean ‘control’ (virtual) environment and an environment that violated Euclidean geometry through the introduction of ‘wormholes’. These wormholes transported participants across the space in a ‘globally inconsistent’ manner. Observations of directional shortcuts supported established knowledge about the encoding of geometric spatial properties, but the wormholes biased the results to demonstrate ‘rips’, ‘folds’ and ‘ordinal reversals’ in spatial knowledge, even in this small virtual maze in which the entire environment could be traversed during each experiment. It is also of particular note that no participant recognised the global inconsistency in the test environment. Building on previous research such as (Muller, Stead and Pach, 1996), the authors proposed that this inconsistent cognitive representation is best characterised by a labelled graph in which local geometric knowledge is integrated into a flexible global structure through weighted edges between these nodal regions, a *cognitive graph*.

The natural acquisition of spatial knowledge and the contextual basis of reasoning about action in spatial environments both seem to be deeply entwined with subdivision of activity into distinct experiential events, with event paronomies suggesting a structure for episodic memory that supports the cognitive representation of navigable environments across multiple levels of granularity.

## 2.5 Concluding comments

The background presented in this chapter has sought to lay the foundation for the presentation of research that follows. Of central importance to the present research is the situated nature of cognitive processes, the notion of situated interaction with navigational information, and the guiding conceptualisation of experience over a

navigational episode as being formed of a partonomic event structure, where information needs vary relative to the contextual basis of goal acquisition within and across events.

The following chapter will now provide further explanation as to the core research aim and the overall research objectives, categorised by key theme.



## 3 Research aim and objectives

The premise underlying the present research is discussed in detail, before the central aim, key initial hypothesis and overall objectives are presented. The chapter builds on the explanation in Chapter 2 of the application of the event horizon concept to navigational information. Key research challenges are discussed briefly, with further explanation and more detailed hypotheses in Chapter 7.

### 3.1 Event structure as a basis for the context-specific selection of geospatial data

As discussed in the first chapter, goal acquisition within a given context leads to an event structure that arises out of the natural segmentation of continuous experience. The central aim of the present research then is to automatically determine the navigational information that supports spatial goal acquisition across these context-specific event structures for A-B journeys. The practical approach to addressing this challenge that was explored in (Mackaness, Quigley and Tanasescu, 2011), building on (Quigley, 2001), was to use graph structures as the basis for modelling and selection. A brief overview is given, followed by further explanation as to the application of graphs within the domain of navigational information.

#### 3.1.1 GRAPH DATA MODELS AND SPATIAL INFORMATION

Graphs are topological structures consisting of vertices, and edges, where edges represent relations between vertices. Property graphs use these edge and vertex constructs as the basis for modelling and storing data. For example, the graph model of features in Open Street Map shown in Figure 2, in which geographic point ‘nodes’ provide an example of graph vertices, and ‘FIRST\_NODE’ provides an example of an edge joining two vertices. The full graph implementation of the Open Street Map

data was developed to demonstrate the potential to return groups of Open Street Map features using graph patterns that would be highly inefficient to return through standard relational queries. This contrasts with a relational database approach in that data is not stored in tables, and native graph constructs can be used to define specific, arbitrary relations between objects.

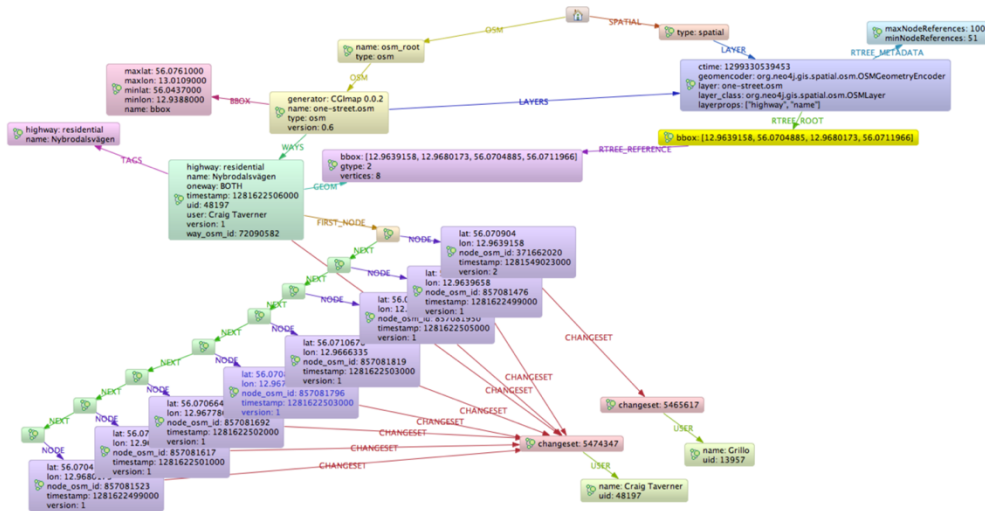


Figure 2 Open Street Map graph data model, subgraph example using Neo4j (Neo4j, 2019).

In real-world domains graphs often become very large and a key challenge is how to visualise these large structures in a meaningful way. Clearly in this Open Street Map example we are seeing a very small subset of the data. In (Mackaness et al., 2011) it was proposed that techniques for visualising large graphs might be adapted to provide a basis for data selection that could support the central aim as outlined later in this chapter. This proposal built on research in (Quigley, 2001) in which techniques were developed to select data such that both ‘local detail’ and ‘global context’ could be surfaced in the same view. The core research challenge explored by Quigley was how to visualise graphs with a very large number of vertices given the practical constraints of limited (screen) area. The basis of the approach was to vary the extent to which the viewer sees ‘into’ a hierarchical graph, with a ‘horizon’ summarising vertex depth across a selection.

In the context of navigation then, an event horizon can be thought of as the aspects of experience that influence goal acquisition as the traveller traverses from A to B. If I am travelling from London to Sydney, both the superordinate entity ‘Australia’ and the part of Australia that is ‘Sydney’ are key entities relevant to the present goal. One exists at a different spatial scale to the other, and so an event horizon that includes both is a manifold across both global and local regions, echoing the explanation in (Yoshino, 1991).

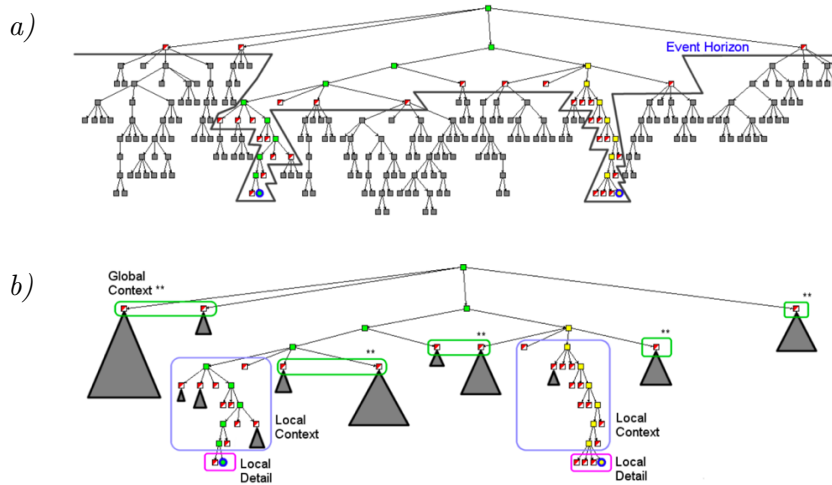


Figure 3 In panel a) an event horizon of a hierarchical graph, and in b) multiple levels of abstraction in a single horizon leading to global context and local detail in the same selection (Quigley, 2001).

If a set of data is modelled using a hierarchical graph with consistent vertex depths, a lateral selection across this structure results in a subset of data from the same level of abstraction. If the hierarchy were to be strictly based on spatial scale, then such a selection would lead to a subset of data at the associated scale for that depth, for example vertices storing features visible in satellite imagery, where vertex depth corresponds with spatial resolution. If however the selection were to ‘cut’ ‘down’ the graph from a source vertex at one level of abstraction to a target vertex at a lower level of abstraction, i.e. to a feature of higher detail, then data from multiple levels of abstraction could be returned in the same selection (see Figure 3).

As discussed by Quigley, to select data stored on a vertex, the horizon must be ‘deep’ enough to surface that vertex. If we consider the standard tile server approach to delivering cartographic information in which a zoom interaction leads to the retrieval and display of a new set of vector tiles at higher resolution, from the perspective of this horizon-based conceptualisation, the horizon is always ‘flat’ in that a subset of tiles can only be surfaced from the same level in the scale hierarchy, and so a zoom interaction serves to push this flat horizon ‘down’ from one uniform vertex depth to another. However, if instead of considering vector tiles as the data structures being stored on graph vertices, we consider the underlying cartographic features themselves – the building features, the text features and so on, we see that in this case the horizon is not flat. Determining the features that should be visible in any given view is one of the basic aims of cartography, so we can see that to apply the horizon conceptualisation where features are ‘elevated’ in the selection so as to be artificially visible at smaller scales, the horizon becomes a manifold, as discussed previously. So, a central task of the cartographer can then be viewed in terms of making design decisions that are equivalent to determining *the shape of*

*the event horizon*. In systems that automatically serve cartographic information to users, vector tiles serve as abstractions over this underlying manifold, and so the shape of the event horizon shifts as a user zooms relative to the underlying features, but in a predetermined manner irrespective of the specific task-contexts that require information over the course of the current journey.

So, we can see that the essential proposal in (Mackaness et al., 2011) was to model geospatial data and to create an overall system such that in essence the specific context of a journey modulates the shape of the event horizon over the underlying graph. If vertex depth were to correspond with the level of schematisation for example, then it follows that by modulating the shape of the event horizon, the level of schematisation could be non-uniform across the resulting view in a way that reflects the specific information needs of the current sequence of task-contexts.




Figure 4 Quadtree layout (far left), then 5th, 4th and 3rd level horizon views of power system distribution network sparse matrix data, represented as a hierarchical graph (Quigley, 2001).

Many fundamental questions arise from this proposal, partly stemming from the fact the original techniques proposed by Quigley were explicitly concerned with visualising non-spatial data. For example in the case of the power system distribution network data shown in Figure 4, the structure of this graph is purely derived from topological relations stored in a matrix, and so proximity of vertices visualised in the plane do not indicate spatial proximity of the underlying entities in geographic space. Clearly then, if vertices exist in the same cell of the quadtree that was used as a basis for clustering, a vertex that represents a cluster of vertices at a lower depth would not necessarily represent features from the same cell of a spatial index based on geographic location. So we see that if the techniques developed by Quigley were applied ‘as is’ to geospatial data in a hierarchical structure that could support this type of selection, the result would in fact be an abstract visualisation of the structure of the database, as opposed to a map view based on the geometry attributes stored on graph vertices.




Figure 5 Example output for a journey across Edinburgh, with shifts in the form of representation driven by 'walk to bus' and 'bus to walk' transitions (Mackness, Quigley and Tanasescu, 2011).

The example shown in Figure 5 does however point toward the potential for using graph structures as a basis for generating information for multiple task-contexts in the same view. In this example features were modelled in a graph database along with a representation of the route, with shifts in the form of representation being triggered by transitions between transport modes. Here we see that a transition to the bus leads to a reduction in topographic detail, with bus stop information being surfaced along traversed route segments. Also it should be noted that even along the sections of the route traversed on foot, the selection of topographic features seems to be aimed at supporting egocentric route knowledge as opposed to providing detail about the broader geographical context of the example journey, i.e. through selection appearing to be based on route-adjacency or some other measure of proximity to the current route


The aim of further investigating these ideas is to support the generation of mixed representations that reflect our natural navigation strategies such as those based on 'fine-to-coarse' representations in which abstraction increases beyond some local 'cusp' (Wiener & Mallot, 2003).

If we consider Figure 5, it is not clear however how the horizon conceptualisation would work in practice for this case. The bus network and the high detail topographic representations could be modelled on a continuum of thematic emphasis, however there is no explicit hierarchical relation between the two, so simply using selections that 'cut' 'down' the graph to return a subgraph would not in itself address this challenge. Graphs are however highly flexible, and so there is perhaps a different type of structure that would support this type of selection. As noted by Mackness et al:

*“...the question then becomes: are these the appropriate underlying data structures necessary to support the delivery of the different types of geographic information [previously described]...?”*

(Mackaness, Quigley and Tanasescu, 2011, p218)

A key aspect of the approach originally described by Quigley that suggests multi-scale geospatial information to support navigation might be modelled and selected in this way relates to returning local detail and global context in the same subgraph (Figure 3). If the selection of a subgraph is based on a ‘cut’ ‘down’ from a source vertex to a target vertex at a lower level of abstraction (a greater depth) then the data stored on the vertex at the ‘bottom’ of this subgraph represents local detail, clusters of vertices near the bottom or at the same depth as the target vertex represent local context, with vertices and vertex clusters near the ‘top’ or at the same depth as the source vertex representing global context. In the case of a ‘multi-cut’ in which global context and local detail may be returned for multiple focal regions across the extent based on multiple source vertices, we see a practical basis for returning the type of selections underlying the example illustrated in Figure 5.



*Figure 6 (left) 3D surface visualisation of an event horizon of an underlying graph based on two focal regions, and (right) an event horizon visualisation for multiple regions of local detail in the same selection (Quigley, 2001).*

If we view such selections in terms of the 3D surfaces that Quigley used to represent multi-cut selections (Figure 6), we can see how selection could relate to a geographic extent across multiple levels of abstraction, i.e. in the case where the index does in fact store features based on geographic location as opposed to the abstract representation of the underlying graph. The illustrations in Figure 6 do seem to point toward the potential application of the event horizon approach to selections across a geographic extent for multiple focal regions, as discussed previously in terms

of our natural navigation strategies, echoing the focal representations discussed in (Wiener, Schnee and Mallot, 2004) and in (Couclelis *et al.*, 1987).

## 3.2 Multi-cut selections as a practical basis for selection over a graph of features

In the previous section the potential to use graph structures as a basis for modeling and selecting information content relative to the context-specific needs of the journey was discussed. The approach proposed in (Mackaness, Quigley and Tanasescu, 2011) did not explicitly consider events in terms of the perceptually bounded segments of experience discussed in Chapter 2, however it is argued that the underlying proposal was in effect equivalent to matching information content from a database to the information needs of a traveller given goal decomposition and the resulting event structure for the current episode. Here the practicality of this proposal is explored in greater detail.

### 3.2.1 CONSIDERATIONS IN THE APPLICATION OF MULTI-CUT SELECTIONS TO DATA SUPPORTING NAVIGATIONAL INFORMATION

In Figure 7 a hypothetical multi-cut selection is visualised for spatial data across a sample extent as an aid to further considering the research agenda that follows from the underlying proposal. Here a multi-cut is comprised of three subgraphs from source vertices at depth  $\mathcal{L}_k$ . Two of these subgraphs extend first to depth  $\mathcal{L}_{k-1}$ , then to depth  $\mathcal{L}_{k-2}$ . The third subgraph, located between the other two, extends to depth  $\mathcal{L}_{k-1}$ , only.

If the vertices and relations in this example were modelled in a hierarchical compound graph it would imply that:

a) features at depth  $\mathcal{L}_k$  or any other depth other than leaf vertices are superordinate with features returned in a subgraph (referred to as an ‘inclusion

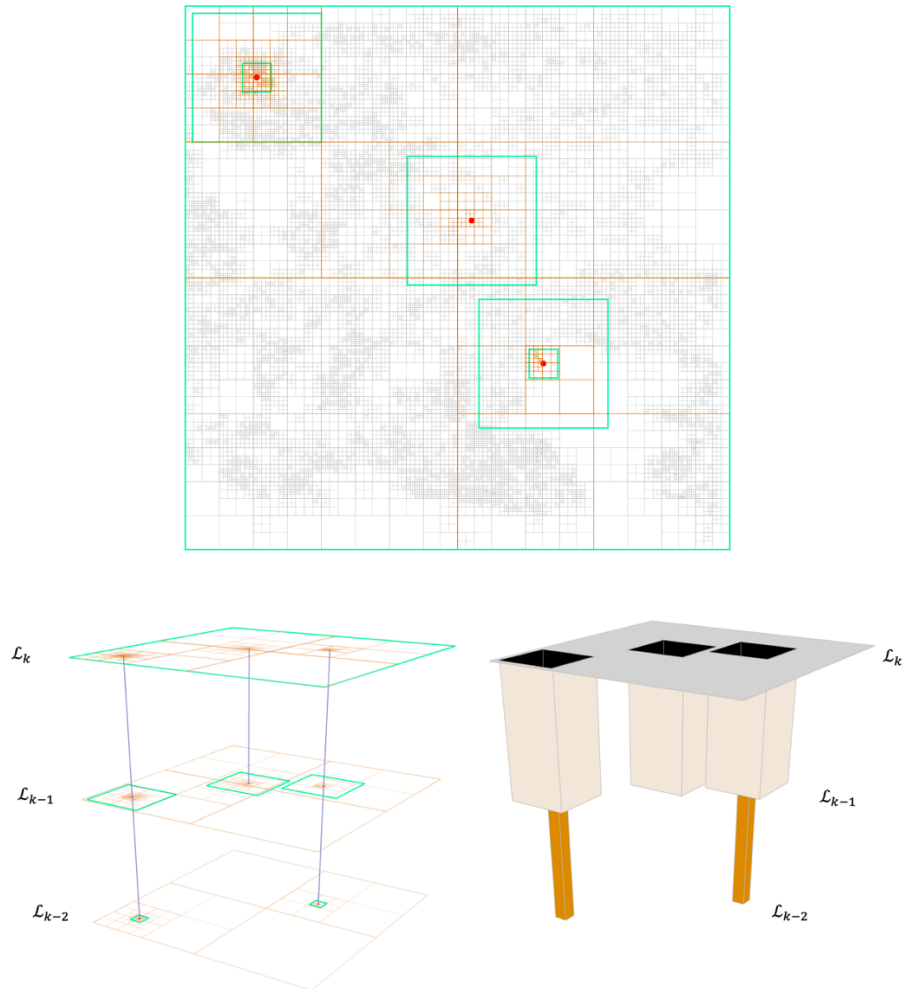


Figure 7 (top) a multi-cut selection from three source vertices of depth  $k$  showing the underlying spatial index and the selected cells, and (bottom left) the same multi-cut selection showing the edges (blue vertical lines) from three source vertices extending ‘down’ from depth  $k$  and edges from two vertices extending down from depth  $k-1$ , and (bottom right) a 3D surface of the event horizon corresponding with the selection.

tree’ in (Quigley, 2001)) being of explicit partonomic relation to source vertices and any other antecedent vertices in the ‘cut’;

b) if a feature is known to be required in the current selection, there is some mechanism for defining a traversal across vertex depths that yields a subgraph of which the feature is an element;

c) hierarchical ordering forms the basis for determining context-dependent content;

d) feature attributes stored on vertices that are returned by a multi-cut can be used to construct ‘cognitively smooth’ views that integrate data from multiple levels of abstraction.

To point ‘a’ above, it is not clear what these superordinate features would be. A city is often represented as a point feature at low resolution, and buildings can be represented at multiple levels of generalisation through processes of aggregation, but overall there are very few hierarchical orderings that are candidates for explicitly modelling the geography with a view to returning ‘cuts’ based on source vertices.

To point ‘b’, by definition a superordinate feature would decompose into a number of features at a lower level of abstraction. If a multi-cut were to begin at a source vertex at a high level of abstraction, in the context of geospatial data, it is not clear what the basis would be of a traversal over the graph leading to a specific feature at a lower level of abstraction. Following from this, there is also the issue that the same feature may play a different role in an alternate context. As discussed in (Bateman *et al.*, 2007), if we consider the case of a body of water meeting the land, this feature could be a boundary in one context, and in the case of a ferry network, a link. A link and a boundary are “...almost opposite: a boundary divides, a link connects...” (Bateman *et al.*, 2007, p101). And so we see that to effectively model context it may be that in fact the underlying graph and the methods of traversal must support multiple paths to the same vertex.

To point ‘c’, there is clearly a basis for considering context in terms of scale-based hierarchies, with focal regions in a hierarchical reference corresponding with tasks specific to the current journey. However, beyond spatial resolution it is not clear how hierarchical orderings would provide the architectural basis for selections that specifically support goal acquisition across experiential events. For example if a cluster close to the target of a multi-cut were to be returned to provide local context as outlined by Quigley, then how would thematic content of direct relevance to the current context be distinguished from the rest of the content in the local region? If we consider the example illustrated in Figure 5, we see a transition between the topographic content and the content specific to the bus network, but distinguishing between these representations is not inherently expressed through ordering relations alone.

Finally to point ‘d’ above, it is not clear how selections from multiple levels of abstraction will lead to well-formed Gestalts. Clearly if features from multiple representations are returned then there is the potential for superimposition, as discussed in (Dumont *et al.*, 2016), as well the potential for overlap and an unacceptable level of map complexity. More generally, there would seem to be the issue that the visual representation of features would have to be considered relative

to every possible multi-cut subgraph to which they could be an element, as opposed to the standard task in cartography of simply considering the Gestalt at each distinct scale. Returning to the discussion in (Dumont et al., 2017), this would seem to indicate the potential for intermediary representations not just across transitions over the scale hierarchy, but also between different forms of representation that may conceptually exist at the same scale. In terms of whether feature representations should be discrete or continuous in support of ‘cognitively smooth transitions’, this issue would also seem to lead to questions around continuous deformation and interpolation. A graph is by definition a discrete structure, and so it is not clear how continuous changes would be supported by discrete representations of features stored on graph vertices. For example would features need to have an attribute that could be used to parameterise the level of geometric distortion through continuous changes relative to a predetermined ‘bounds of distortion’ (Godfrey & Mackaness, 2017)? For further background to the concept of the bounds of distortion, refer to Appendix 4.

This point links with perhaps an even more significant question: to model geospatial features in such a way as to return the kind of mixed, context-specific representations as described, are our current cartographic features themselves in fact sufficient as a basis for achieving this? In other words, is this simply a challenge of how to form ordering relations between features that already exist in say Ordnance Survey and Open Street Map datasets? Or, is this challenge in fact only effectively resolved by also introducing new types of features not currently present in our databases? Intermediary representations would be one example. Another would be superordinate features that serve as abstract representations of features from multiple classes at lower levels of abstraction. In other words, to traverse an underlying graph to return a ‘cut’ subgraph, it may be necessary for edges to link superordinate features with features of different types at lower vertex depths, which is not naturally supported by current approaches.

This final point links with a proposed approach in (Mackaness et al., 2011) that may address a number of the challenges as described, in which the authors argued for “...entities to be grouped partonomically according to the function of the entity and its ‘fit’ with our conceptual views of the world...” (Mackaness et al., 2011, p219). The essential proposal then was to create hierarchical, partonomic relations based on functional categories such as: “...suburbs, industrial parks, recreation centres... road [networks], public transportation hubs...” (Mackaness et al., 2011, p219), and that if these entities across levels of abstraction could be ordered hierarchically, then queries equivalent to multi-cut selections could be applied to the structure as a basis for returning context-specific content. This proposal sits alongside a number of concepts and techniques that will be discussed further in Chapter 4, such as the GAP-edge forest (van Oosterom, 2005).

Finally, the examples given in (Mackaness et al., 2011) suggest an implicit aim to return a selection that integrates multiple representations into a single view. The ‘spider maps’ referenced in the paper, as well as the example output shown in Figure 5 that has been discussed extensively, both suggest a design objective of reconciling information into a single view. It is not however clear that this approach would be optimal, with ‘spider maps’ for example being fundamentally designed as large printed map products as opposed to small (mobile) interactive maps. In both cases it is implicit that zoom interactions are either reduced or removed entirely, but if this design aim does not turn out to be appropriate, then a key issue that naturally follows is how the overall approach described in this chapter could be realised in the case where there is in fact a requirement for the equivalent of multiple event horizons for a single journey context.

### 3.3 Core research aim

Summarising from the preceding discussion, the core aim of the present research can be understood in terms of:

**developing a framework for the automatic selection and display of a subgraph of features, where the resulting subgraph supports information needs over a set of navigational task-contexts for a given journey.**

Here the process of selection is equivalent to using context as a basis for automatically modulating the shape of an event horizon over an underlying graph of features. In terms of the goals that underly episodic experience, and the partonomic structure of events over navigational episodes, this core aim can be seen in terms of providing information to support goal acquisition across the distinct set of events that will come to comprise the current episode.

### 3.4 Initial hypothesis

Given the core research aim, the initial hypothesis underlying the proposed research agenda is based on the following logic:

**For any A-B route, if a person were to travel the route and successfully complete it, then upon reaching the destination their cognitive representation of the route would be appropriately modelled as a ‘cognitive graph’ (as discussed in Chapter 2).**

If however an artificial agent were to traverse a virtual environment that effectively represents the same underlying geography, and the agent develops spatial knowledge in a way that mimics the way the person would have developed their own cognitive representation over the course of journey completion, then the context-specific representation formed by the agent will be approximately isomorphic to the cognitive graph that would have been developed by the person.

If the selection of features from an underlying database is equivalent to the second case, then displaying this subset of features to the traveller *a priori* will be approximately equivalent to the person acquiring context-specific spatial knowledge through physical route traversal.

If we reflect on the description of an event horizon as ‘a boundary in spacetime beyond which events cannot affect an outside observer’, the hypothesis described above can be viewed in terms of predicting the experiential events that will affect the traveller’s progression towards goal acquisition, and matching content in the database with information needs over the course of those events.

In terms of the primacy of attention in the acquisition of spatial knowledge, the initial hypothesis described above implies that selection should be the result of shaping and constraining the path of graph traversals based on the likely processes of attention that would have led to the *a posteriori* cognitive representation of the journey, had the person in fact traversed the route physically. Irrespective of whether the system is ‘agent-based’ per se, this conceptualisation is equivalent to the underlying graph being a representation of a navigable environment, and a graph traversal for a given journey being equivalent to an agent acquiring spatial knowledge through situated action. For example, by being sensitive to the likely subdivision of activity that will come to shape the event structure for the current episode.

### 3.5 Research objectives

Following the core research aim and the hypothesis stated above, research objectives can be understood to be divided across two overall themes. The first set of objectives relate to the gaps in knowledge that require empirical investigation. The second set of objectives stem from the need to formalise knowledge such that map content can be automatically selected and displayed in a way that aligns with the empirical findings.

**1. Objectives: empirical research**

- a) *Which aspects of our knowledge of spatial cognition are of practical significance given the central task of automatically selecting map content to support multi-modal travel?*

It is necessary to specify aspects of the structure and content of cognitive representations developed through situated action in navigational task-contexts that are of relevance to the research aim. While a large body of knowledge concerning how spatial experience is encoded and operationalised already exists, what is not clear is how this knowledge should be used in the formalisation of information systems that automatically select and display map content based on context. Linked to this, a key question is to what extent is it in fact appropriate to explicitly drive information selection using this type of knowledge. For example, natural cognitive representations are known to be incomplete and distorted, so the extent to which these types of characteristic should be made explicit is an area for investigation. Another important overarching question linked to this issue is how can we predict the set of events and event sequences for a given journey by way of standard contextual data (location, destination, route)? Additional contextual data such as detailed profile information is deemed out of scope of the present research, and so the need to use data that is already available frames the practicality referenced in the high-level question that expresses this overall research objective.

In considering the initial hypothesis, is it in fact appropriate to conceptualise the problem in terms of an agent experiencing the environment from the egocentric perspective implied by mimicking the experience of a person physically traversing the route, or should the system not be constrained by this type of ‘agent’ conceptualisation?

- b) *Are any of the entities that are elements in distinct events or event sequences over navigational episodes additional to the standard data that is available to us?*

If we are to model, select and display map content such that it automatically reflects the natural information needs of the user, is it in fact sufficient to use entities that exist in our current databases, or are there additional classes of entity required for this approach to be viable?

In other words, is the central objective in the development of a system that reflects the core research aim simply to model relations and groupings between existing entities, or should additional entities be defined and formalised?

- c) *What should be the basis of clustering with respect to entities that play a role in multi-modal navigation?*

The central aim of automatically determining map content can be understood as one of clustering in the sense that a number of features must come together to form the basis of a map view. In other words, there must be some process for features stored in a database to be formally grouped given the aim of surfacing navigational information to a user. While in the prevailing paradigm a tree structure forms the underlying basis for this clustering as discussed previously, in terms of the core research aim, the basis of clustering is not known, and is a key research objective. It is hypothesised that addressing this issue requires exploring questions such as what form of abstraction is congruent with the traveller's information needs over action sequences associated with key experiential events, and given identified patterns of event sequences, what form of representational abstraction reflects information needs associated with traversal between distinct events?

Are we correct that if we can predict an approximation to a cognitive representation in advance and we can model information needs given context-specific event sequences then we will be able to match information content to a given journey context, i.e. will we be able to predict clusters in an effective way?

## **2. Objectives: data modelling, selection and display**

- a) *Are any additional entities required to support data modelling that are not explicitly present in the empirical data?*

To develop a data model that supports the core research aim, a key question is whether any classes of object are required that are not explicit in the empirical data but nevertheless support the overall functionality with respect to selecting and displaying a subgraph of features to support traveller information needs. This question extends beyond the

representation of physical entities to abstract entities such as those that might generalise over classes of features in a way that reflects situated experience.

- b) *What are the relations or orderings between entities that allow them to be surfaced in context-specific selections via graph queries?*

A key question is what could be the overall architecture of a graph that supports the storage of entities and relations in-line with the core research aim? For example, is this structure a hierarchical compound graph as described in (Mackaness et al., 2011)? Also, beyond this overarching structural question, what are the relations between individual entities that support modelling, selection and display, for example relations that capture the roles a subset of entities might play in the development of situated understanding.

- c) *How can interactive maps for multi-modal travel be automatically generated from data selected from the underlying database?*

While the focus is on data modelling and selection, ultimately the aim is to surface data in map views, and so a key objective is to determine how context-specific selections should support cartographic visualisations. Here there is an argument for working ‘backwards’ from target visualisations to the underlying data architecture that would support their manifestation by automated means. Of critical importance here is ensuring that automated methods lead to the generation of views that are well-formed Gestalts, and that align with the empirical findings with respect to traveller information needs.

### **3.6 Scope of research**

A model of the cognitive representation of complex urban journeys that is stable across individual preferences and behaviour types given an A-B navigation task is within scope. It should be highlighted however that it is only in scope so far as it relates to traveller information needs for which there is a clear method for determining selection from an underlying graph of features.

A method for modelling geospatial data is in scope as far as it supports the selection of data that addresses information needs for specific task-contents during the traversal of an A-B route, i.e. the data can be returned via graph queries.

The outcomes outlined above must use inferences about traveller context that take standard contextual data as input such as current location, destination and recommended route. In other words, the framework developed over the course of the present research is not based on any requirement for additional data that is not standard in current systems. For example, there is no need for a user to provide profile information. This approach means the methods developed will provide correct results even in a 'base case' but does not preclude the methods from being designed in such a way as to be scalable across further use cases.

Finally, an empirically validated framework for the display of context-specific selections on mobile devices is within scope. In other words, given methods for modeling and querying an underlying graph of features, it is within scope to also provide detail as to how the framework results in information that may be displayed to a user through a mobile device.

Activity that concerns routing is to be highlighted as out of scope. In other words, the study of cognitive tasks associated with route comparison and route selection. Here the scope of research is constrained to route planning, orientation and wayfinding given a recommended A-B route as a starting premise. Following from this, determining a multimodal route by automated means should also be highlighted as being out of scope. Here the starting premise is that a routing engine has already returned a recommended route based on a starting location and a destination





# 4 Review

Literature is reviewed across three high-level sections. The first addresses the state of knowledge with respect to the cognitive representation of spatial experience and our natural navigational strategies. In the second section, formalised theories of ‘cognitive mapping’ are briefly reviewed, including approaches that have integrated localised representations into looser global structures. Finally, the third section considers navigational information itself and provides further background as to the role of information in different forms and the integration of information such that various types of navigational tasks are supported through single views.

## 4.1 Events as the cognitive basis for episodic experience and spatial goal acquisition

Building on the background discussion in Chapter 2 this section reviews literature relating to the cognitive foundations of spatial experience, including the acquisition of spatial knowledge over time, the natural regionalisation of environments to support comprehension, reasoning and memory, and the way that naturally forming representations support flexible and efficient navigational strategies.

### 4.1.1 SITUATED REPRESENTATION AND URBAN ENVIRONMNETS

In the seminal ‘The Image of the City’ (Lynch, 1962), Lynch proposed that in terms of their form and structure, urban environments are naturally comprehended and experienced through a set of distinct foundational classes. The proposed classes were paths, edges, districts, nodes and landmarks. Echoing the origins of graph theory in terms of an effective understanding of urban form as being fundamentally topological in character, this representational framework expressed an essential insight that our natural cognitive representation of urban environments is one of distinct objects and collections being loosely bound together in a flexible global

representation, where an understanding of connectivity and spatial relations develops experientially over time. Here then the notion of the ‘mental image’ of the environment encompasses the ‘legible’ structure of the city, a representation that is partly, but not entirely visual, and is relative to each person:

*“Environmental images are the result of a two-way process between the observer and his environment. The environment suggests distinctions and relations, and the observer—with great adaptability and in the light of his own purposes—selects, organizes, and endows with meaning what he sees. The image so developed now limits and emphasizes what is seen, while the image itself is being tested against the filtered perceptual input in a constant interacting process. Thus the image of a given reality may vary significantly between different observers.”*

(Lynch 1962, p6)

Taken together with the proposed classes, the above quote succinctly captures a view on the natural form of our spatial knowledge that has been repeatedly shown in empirical studies across varying disciplines to be an appropriate characterisation. Indeed, across both cognitive science and neuroscience, as highly detailed knowledge of spatial representation has been uncovered, it has been shown that even at the biological level, a representational framework that is at least similar to Lynch’s proposal is congruent with the observed evidence.

#### 4.1.2 GOALS, MOTIVATIONS AND ATTENTION

Goal-directed behaviour arises out of commitment to a course of action (Shadlen and Kiani, 2013), or ‘commitment to a categorical proposition’ (Gold and Shadlen, 2007), and is intrinsically linked to processes that lead to expectations and the evaluation of circumstances relative to those expectations, with reward or ‘the value of consequences’ (Balleine, 2018) being shown to modulate these processes, for example by way of spatially encoded representations of goal proximity (Viard et al., 2011). In ‘ecologically valid’ scenarios however, the ‘cost’ of action cannot be completely known *a priori*, with costs and expected rewards requiring re-evaluation as the organism progresses through the environment. In other words, although the decision may have been made to traverse to goal A, many more decisions must be evaluated as a result of embodied action before A is actually reached (Lepora and Pezzulo, 2015).

A bias towards preferential consequences plays a central role within our broad perception and conception of causal relations (Schultz, Dayan and Montague, 1997), and in our embodied use of causal knowledge (Sloman and Lagnado, 2015). Attention has been studied extensively and has been shown to be fundamental to

cognition within this causal backdrop (Corbetta and Shulman, 2002), with the motivational impulses to either move ‘towards’ or ‘away from’ serving as basic motivational templates (Bradley, 2009). Attention has been found to be directly shaped by these motivational impulses, with a person’s goal, as shaped by motivations, leading to attentional bias as stimuli enter the flow of perceptual experience. Motivation leads to commitment to a goal, which is co-constituted with expectations, which in turn have been shown to modulate attention through ‘anticipatory bias’, for example directing focus to spatial regions in which a change in stimulus is predicted to be most likely (Small *et al.*, 2003). Attention therefore can be characterised in terms of both top-down ‘endogenous orientation’ directed toward expected stimuli of context-specific significance, and in terms of bottom-up ‘exogenous orientation’ (Raz and Buhle, 2006) through sensory information that captures some level of physical salience such as movement and visual contrast (Yantis and Egeth, 1999). Also it should be noted that attention extends to ‘temporal orientation’ (Coull *et al.*, 2000), with an example being anticipation and response to the duration of intervals, i.e. temporal spacing between distinct events (Miniussi *et al.*, 1999).

Here it is important to clarify that in the study of the cognitive basis of navigation, orientation has two meanings. On the one hand the notion of orientation in terms of wayfinding, describing the processes of determining the position of the body in space and position relative to environmental features and spatial goals. On the other, the more fundamental notion of the orientation of attention, as described above with respect to endogenous, exogenous and temporal orientation. Clearly however, these two notions of orientation are linked, with for example top-down endogenous attentional orientation based on long-term memory being “...essential for targeting behaviourally relevant objects or events embedded in complex environments...” (Summerfield *et al.*, 2006, p905). Here the authors also emphasise “...the importance of the regularities of contexts in guiding object and scene identification...” (Summerfield *et al.*, 2006, p905). So we see that memory influences the orientation of attention (Carr and Schissler, 1969), and that regular patterns of experience come to shape the way the environment is perceived, including the objects that are attended to (Corbetta, Patel and Shulman, 2008) and the manner in which an overall context is comprehended.

A picture emerges then in which cognitive processes are not simply passive, driven by external stimuli, but rather they shape the basis of the agent’s informational state through both epistemic action and on-going prediction (Clark, 2015). Clark’s articulation of the brain as an embodied ‘prediction machine’ can be seen within a broad, extensive and on-going research agenda to establish the basis of cognitive control, for example investigations into “adaptive fluctuations in the attentional focus on goal-relevant information...” (Muhle-Karbe, Jiang and Egner, 2018, p972),

and investigations into visual search with multiple targets (Ort *et al.*, 2019). It has also been shown that an ability to shift attention based on task-context is supported by the ability to modulate cognitive control itself through a continuum between ‘diffuse’ and ‘focused’ attention, with evidence suggesting that variation in the degree of attention diffusion serves a central role in task success (Amer, Campbell and Hasher, 2016).

In term of categorising behaviour overall, activity requiring goal-directed decision-making has been explored in contrast with habitual behaviours that do not require sensitivity to context (Miller *et al.*, 2018). So, goal-directed behaviour necessarily arises out of situations in which a need cannot be met simply by reverting to habit alone, and where the organism is sensitive to context and has actionable strategies for reducing context-specific uncertainty. This situation is complicated by the fact that in practice achieving goals requires the enaction and integration of both habitual and context-dependent directed behaviours, and it has been proposed that these two key behavioural modalities compete under conditions of uncertainty or ‘expected inaccuracy’ (Daw, Niv and Dayan, 2005). It has also been proposed that habitual behaviour itself forms the basis of goal selection and the development of reward mechanisms over time (Cushman and Morris, 2015).

Finally, it should also be noted that the influence of temporal context on goal-directed decision-making is pervasive. On the one hand the influence of costs in terms of evaluating decisions, such as ‘temporal discounting’ on general attributions of subjective value (Peters and Büchel, 2010). On the other, the cost of decision-making itself, or the ‘cost of accumulative evidence’ (Drugowitsch *et al.*, 2012). In other words, task complexity and the duration to a goal will lead to cumulative ‘cognitive costs’, for example costs associated with working memory demand. These costs will then influence a person’s attention and the manner in which a semantic segmentation emerges out of the underlying continuity of experience as time unfolds (Bor *et al.*, 2003).

#### 4.1.3 EVENTS, SEQUENCES AND EPISODES: A SPATIAL CONTEXT FOR EXPERIENCE

While an extensive understanding has been developed with respect to neural mechanisms that appear to somehow ‘map’ space, the evidence suggests that really we are observing processes of memory formation through time that are tightly bound to spatial experience, but not in fact some kind of systematic ‘Cartesian’ mapping system (Eichenbaum *et al.*, 1999). A systematic ‘spatial’ method of encoding would not lead to the kind of representational structures that seem to underly observed behaviour and would not lead to study participants being so easy to fool in

navigation experiments. Perhaps the most simple example here is the effect of introducing dynamically positioned extrinsic cues into experiments, such as movable boundaries (Moser, Kropff and Moser, 2008). As shown by (Skaggs, Moore and McNaughton, 1996), the movement of an environmental cue causes the pattern of activation in cells that ‘encode space’ to change, pointing to an underlying process of encoding which is in fact the “...independent representation of subsets of the spatial cues...” (Eichenbaum *et al.*, 1999, p212) as opposed to the representation of space itself. Evidence from empirical studies with humans has also helped to refute the notion that ‘space’ is mapped in some objective allocentric manner (Wilson and Wilson, 2018). More recently it has been suggested that entorhinal grid cells, rather than encoding a spatial ‘map’ per se, are perhaps more accurately described as encoding low-dimensional representations that are inputs into higher-order predictive processes (Stachenfeld, Botvinick and Gershman, 2017).

A key question then is what is the actual informational content generated by these neural mechanisms? It has been proposed that the encoding of geometry and distance observed in various studies may in fact be conveying topological information ‘downstream’, along with the metric knowledge that is assumed to originate in these processes (Dabaghian, Brandt and Frank, 2014). For example, given there is a natural temporal sequence to activation, it is likely that topological sequence is automatically encoded in some way, and given that through the normal course of movement through space, as place fields naturally overlap through shared activations, or are disjoint through clear boundaries between sets of activated neurons, key topological information such as boundaries, containment, sequences and intersections, are likely to be present (Dabaghian *et al.*, 2012). Empirical support for hippocampal place cells somehow passing topological information downstream was supported by the finding that as “...long as the sequence of spaces experienced remains the same... and the place fields still overlap, it would appear not to matter to the hippocampus if the place fields are stretched...” (Dabaghian, Brandt and Frank, 2014, p10). A proposal that speaks to the basic question of informational content then is that patterns of activation over time in the hippocampus may form the basis of topological ‘templates’ that serve as a basis of “...a spatial context for experience...” (Dabaghian, Brandt and Frank, 2014, p10). This notion is consistent with previous work that considered evidence the ‘cognitive map’ in navigation tasks seems to be not so much a topographic representation as a framework for constructing spatial memory (Eichenbaum *et al.*, 1999). See (Epstein *et al.*, 2017) for a more general review of the current state of knowledge regarding the ‘cognitive map’.

While an underlying framework for organising knowledge through the continuity of experience seems to result in some cases in the formation of ‘map-like’ representations, the more fundamental organising principle seems to be one of

supporting flexible behaviour across transitions in task-context and across shifts in the nature of available information. Here a picture emerges of our knowledge of navigable space as being comprised of local representations with metric qualities integrated into larger, looser representations that afford flexibility in the face of inconsistent spatial knowledge at varying resolutions of attention. Here space, time and memory are intertwined, with results from the study of the organisation of spatial and temporal knowledge pointing toward general mechanisms for the organisation of memory (Eichenbaum, 2017).

Following from Piaget's seminal work on cognitive development, Piaget and Inhelder built on key results to propose that the development of spatial cognition, along with a broader developmental progression from egocentric to allocentric reasoning, can be seen as beginning with the development of egocentric topological knowledge, through to reasoning in 'projective space', and finally Euclidean space (Piaget and Inhelder, 1956). As noted by Poucet, this led Piaget and Inhelder to adapt mathematical definitions of topology to an embodied view of topology, resulting in "...notions of proximity and neighbourhood, enclosure, continuity of lines, and order of succession points..." (Poucet, 1993, p168).

The study of topological spatial representations in cognition and its underlying neural basis can really be seen through two interrelated research agendas. On the one hand, investigations into the specific mechanisms that support the development of topological representations, for example the application of Hebbian learning theory (Hebb, 1949) to the study of pairwise synaptic relations between place cells in distinct firing fields (Muller, Stead and Pach, 1996). On the other hand, the study of comprehension and reasoning in spatial contexts based on the type of flexible representations observed by Piaget and Inhelder. To the first case, this work provides underlying experimental and theoretical support for cognitive representations of space with topological properties, however it is the second case that is really of direct relevance to the present research.

Work to understand how the cognitive basis of navigation might be specified (Leutgeb *et al.*, 2005) and conceptualised in the way discussed above is ongoing, for example through the novel experiment discussed in (Warren *et al.*, 2017) and elaborated on further in (Warren, 2019). In this study human subjects navigated both a Euclidean 'control' (virtual) environment and an environment that violated Euclidean geometry through the introduction of 'wormholes' that 'transported' participants across the space in a 'globally inconsistent' manner. Observations of directional shortcuts supported established knowledge about the encoding of geometric spatial properties, but the wormholes biased the results to demonstrate 'rips', 'folds' and 'ordinal reversals' in spatial knowledge, even in this small virtual maze in which the entire environment could be traversed during each experiment.

It is also of particular note that no participant recognised the global inconsistency in the test environment. Building on previous research such as (Muller, Stead and Pach, 1996), the authors proposed that this inconsistent cognitive representation is best characterised by a labelled graph in which local geometric knowledge is integrated into a flexible global structure through weighted edges between these nodal regions, a ‘cognitive graph’.

#### 4.1.4 GLOBAL REPRESENTATIONS AND SUBMAPS: STRUCTURING SPATIAL KNOWLEDGE FOR FLEXIBLE AND EFFICIENT BEHAVIOUR

While the notion of ‘cognitive maps’ remains a focus of research effort, the evidence clearly shows that a labelled graph is a far better characterisation than a Euclidean map based on extensive findings from various disciplines. This has been known for many years, for example in (O’Keefe, 1998), one of the founding figures in the theory of cognitive maps discusses experimental support for a structure very similar to the labelled graph discussed in (Warren *et al.*, 2017), with places as ‘patches’ of the environment being loosely bound by direction and distance information, and with paths through space being ordered sequences of places and the translation vectors

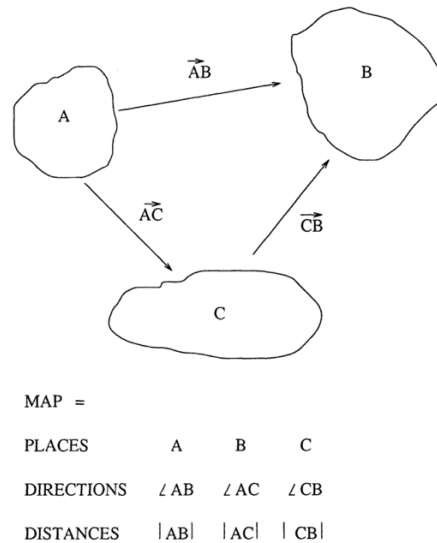


Figure 8 Schematic of the cognitive representation of a navigable environment as consisting of a set of place representations and the distances and directions between them (O’Keefe, 1998).

between these local ‘patches’, as illustrated in Figure 8. Clearly this view echoes the framework discussed in (Eichenbaum *et al.*, 1999) in which spatial representation and memory more generally are encoded as distinct events bound together in

sequences within overall episodic experience. Whether the argument is for a loose global structure bound together by estimated distance and direction values, or for overlapping memory chunks as the basic framework for binding the representation, there is in essence a consistent empirically supported view on the underlying cognitive framework for our knowledge of navigable environments.

The spatial context of key behavioural events and behaviourally distinctive places (Mou and McNamara, 2002) then leads to more detailed topographic representations within these ‘patches’ (Hafting *et al.*, 2005), with agents preferentially attending to local cues, and goal-directed orientation in the agent’s environment supporting selectivity of cognitive resource allocation. These processes have also been shown to extend beyond attention to objects within our immediate sensory experience (Corbetta, Patel and Shulman, 2008), and crucially, to ‘cue combination’ in which associative links between perceived stimuli lead to the ability to attend to contexts as opposed to discrete objects (Bar, Aminoff and Schacter, 2008). Contexts may arise out of simple associative links between objects, or more complex assemblages such as combinations of both idiothetic and allothetic cues (idiothetic cues are ‘internal’ cues triggered through self-motion, whereas allothetic cues arrive through an experience of the ‘external’ environment) (Chen *et al.*, 2017). It should be noted here that ‘object’ is simply used as a term to generalise over categories of relevant perceptual stimuli including surfaces (Knierim, 2002).

Beyond these processes shaping which distinct stimuli or associative cue networks are attended to in the first place, it has been shown that in fact, in the macaque at least, task goals actually “...modulate the neural representation of a stimulus...” itself (Wen *et al.*, 2019, p1) and (Mirabella *et al.*, 2007). Lab-based studies with humans have also shown that task directly affects representation, for example through the impact of both conceptual and perceptual tasks on the temporal dynamics of visual stimulus processing (Hebart *et al.*, 2018). There seems to be broad consensus that task-context directly affects not just the objects that are attended to but the representational form of perceptual content itself.

#### 4.1.5 REGION HIERARCHIES: SCALE AND COST IN REPRESENTATION AND REASONING ACROSS TASKS AND SUBTASKS

In navigation then, commitment to a course of action manifests in behaviours such as route selection, route planning, re-planning, path following, orientation, re-orientation, landmark recognition, map reading and following textual route descriptions. While arriving at a single comprehensive definition of navigation is problematic, what can simply be stated is that navigation concerns goal-directed

behaviours in which the goal has a location that is either spatially or metaphorically distinct from the current location, and that position, heading, physical displacement and proximity to the goal must be comprehended in some way.

How then do the various mechanisms that support people in shaping goals and acting on motivations in navigational scenarios come to be so coordinated? Embodied action unfolds such that processes occur in parallel and across various temporal and spatial scales, with the consensus view being that experimental results point toward an interplay between ‘bottom-up’ sensory data, ‘top-down’ cognitive control, and intermediary mechanisms that help shape how further sensory data should be collected and high-level goals adjusted given new states of knowledge (Clark, 2015).

An example explanation is that control and motivation are complementary aspects of active inference, where goals are contextualised at multiple levels of hierarchical abstraction, and updated belief then forms the basis of action at each level (Pezzulo, Rigoli and Friston, 2018). Here then, hierarchy relates to the ordering of processes of ‘motivated control’ within contextual constraints across an affordance level, a semantic level and an episodic level. So, for example when driving an ambulance - the recognition of a left turn at the affordance level, the recognition that the light is red at the semantic level, and the fact that the patient will die if the journey is held up at the episodic level. The active inference approach provides a computational basis for dealing with motivated control within and across these levels which may not necessarily reflect the actual neural architecture of goal-directed behaviour, but is however characteristic of the current state of knowledge with respect to the basis of behaviour across tasks and subtasks (Botvinick, Niv and Barto, 2009) and (Stoianov *et al.*, 2018). More generally, these findings and proposals can be seen within a wider research agenda that has sought to determine the structure of behaviour, with the notion of task hierarchies having both an empirical and mathematical foundation (Yee *et al.*, 2014).

It has been shown that hierarchy emerges naturally in spatial networks through the ‘global consequences’ of localised cost-benefit constraints (Louf, Jensen and Barthelemy, 2013), echoing broader findings around cost efficiencies gained through the emergence of hierarchical modularity in complex systems (Siyari, Dilkina and Dovrolis, 2019). In terms of navigation and wayfinding, while the hierarchical nature of spatial representation has been studied extensively, the notion of a hierarchical subdivision of goals and tasks has received far less attention and has often been treated as implicit and as essentially distinct from the spatial representation. In (A. Car and Frank, 1994) and (A Car and Frank, 1994) it was stated that in spatial reasoning hierarchical subdivisions can occur in terms of either the representation of task, as in (Timpf *et al.*, 1992) or the representation of space itself. Evidence

from neuroscience and cognitive science suggests such a distinction is problematic however, and introduces an artificial division between the representation of task and the representation of spatial granularity. Arguably in the case of the (Timpf *et al.*, 1992) paper, in which wayfinding is considered across a ‘planning level’, an ‘instructional level’ and a ‘driver level’, it is difficult to see how it is possible to separate the subdivision of task from the granularity of the spatial representation, and it would seem that the hierarchical basis of motivation and control in fact has an explicit link to the hierarchical representation of space, particularly when issues of cost are considered central to the evaluation of causal knowledge and the ultimate commitment to action. In (Raubal and Egenhofer, 1998) an approach to the analysis of wayfinding task complexity was proposed based on image-schema theory (Lakoff and Turner, 1989), and although this work did not explore the hierarchical structuring of wayfinding tasks, applying an analysis of task complexity to wayfinding contexts naturally lends itself to aligning the hierarchical representation of space with a hierarchy of tasks and subtasks.

This hierarchical view is also central to the anchor-point hypothesis that provided a further basis for understanding how spatial knowledge develops over time and how spatial representations are operationalised in practice (Couclelis *et al.*, 1987). Here the authors built on established observations of our strong natural tendency toward regionalisation, route segmentation and hierarchical ordering. The anchor-point hypothesis can therefore be seen as expressing spatial knowledge as having both metric and topological qualities, with key nodes serving to anchor the wider representation and exert regions of influence across spatial scales. Anchors such as salient landmarks can be seen as themselves being members of a hierarchical representation in which “...a process of ‘spreading activation’...” (Couclelis *et al.*, 1987, p103) leads to the establishment of associational links across regions and the development of intra-regional structure. As illustrated in Figure 9, this account leads to a view on spatial knowledge that is highly congruent with evidence discussed previously, with goals, attention and memory, by way of path following, leading to the representation of a loose, regionalised global structure that develops around localised events and event sequences. Research discussed previously from (Brunec, Moscovitch and Barense, 2018) and (Alexander and Nitz, 2017) represents just two examples of the on-going empirical support for the anchor-point hypothesis as being an appropriate framework for expressing the nature of our spatial knowledge acquired through everyday experience.




Figure 9 Illustrative example of the anchor-point view of spatial knowledge containing three types of action context, their spatial connectivity at multiple levels of abstraction and the generalised regions associated with each (Golledge, 1999).

Another important aspect of the anchor-point view is that it articulates not just a practical basis for the development of spatial knowledge, but also the basis for internal representations to be anchored to the experience of the external world, linking with broader work to understand how memory comes to be bound within spatial contexts, for example through frames of reference (Mou and McNamara, 2002).

It has been noted this view of representation is constructivist in nature, with disparate pieces of information, some of which may not be ‘map-like’ at all, being bound together in representational structures better characterised as ‘cognitive collages’ rather than cognitive maps (Tversky, 1993). While there may be anchors in a person’s representation that are causally linked to their individual experience, for example the location of some specific event in memory, there are a broad set of elements that come to shape the development of the understanding of regionalised structures that have been shown repeatedly to be stable across representations (across people). For example see (Tenbrink, 2012) for a review and a summary of key elements of route descriptions such as the starting point, intermediate decision-points and route segments.

Elaborating on the hierarchical nature of these representations, with reference to previous work by Golledge, Couclelis et al. note that ‘primary nodes’ anchor distinct

regions and the “...anchors and linkages between them provide a skeletal hierarchical structure for representing and organizing cognitive information about space...” (Couclelis *et al.*, 1987, p99). Again, this assertion has consistently been supported by experimental data, for example through studies on the role of landmarks and language cues as the basis for the formation of abstract regions (Schick *et al.*, 2019), and in navigation tasks such as perceived region connectivity in route planning (Wiener, Schnee and Mallot, 2004).




Figure 10 (A) connections between distinct local charts, where places are linked across non-overlapping environments, (B) solving navigational problems by combining charts, and (C) efficient computation of a route (a-d) based on hierarchical knowledge of adjacent charts and the places contained within them (Poucet, 1993).

This view of a hierarchical regionalisation of navigable space was echoed in work on animal navigation that sought to explain observations of flexible and efficient behaviour (Poucet, 1993), with Poucet hypothesising that animal navigation could be characterised in terms of processes of environmental learning that involve:

“a) the integration of local views into place representations; b) the formation of local charts of contiguous regions, in which vectorial information is initially conveyed by location-dependent reference frameworks; and c) the computation of an overall (location-independent) reference direction for each local chart...”

(Poucet, 1993, p174)(see Figure 10).

As Poucet noted, this approach would support the efficient use of cognitive resources in navigation, building on earlier work into the study of the efficiency of potential

hierarchical representations of geographic environments (Yoshino, 1991). In Yoshino's proposed framework, directed paths ('strings') through the environment

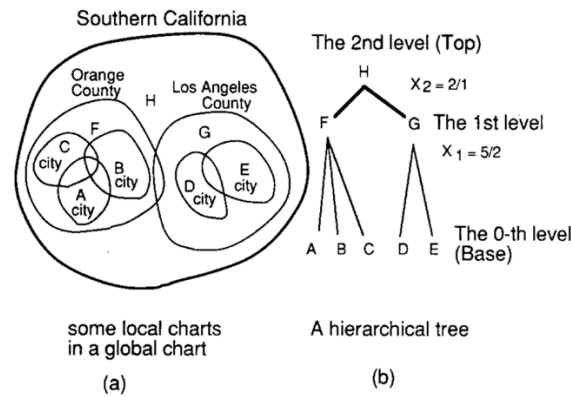


Figure 11 (a) A manifold of charts representing Southern California, and (b) its hierarchy. F is a reference point in a chart of {A, B, C}, G is a reference point in {E, D}, and H is a reference point of a global chart including F and G. Local charts may overlap.  $X_i$  is the ratio between the number of branches at the  $(i - 1)$ th and  $i$ th level (Yoshino, 1991).

serve as the basis for acquiring spatial knowledge over time, with the representation forming from 'separate pieces of information' to a 'partially organized set of local charts' (a 'manifold'), to a 'global map' binding together local charts. Reflecting the experiential reality of spatial knowledge this approach supports reasoning with incomplete knowledge and being able to integrate partial information. Over time the overall thrust of this proposal has continued to gain empirical support, and has strong parallels with more recent work that has sought to expand the neuroscientific study of spatial representation beyond contrived lab environments to more naturalistic environments that are 'compartmentalised', 'nested' and 'variable in time' (Derdikman and Moser, 2016). Building on previous work to study the fragmentation of grid cell representations in multicompartiment environments (Derdikman *et al.*, 2009), the authors note that experimental evidence does indeed show that neural representations based on patterns of hippocampal and entorhinal activity reflect this naturalistic complexity, with overall cell activity fragmenting "...into interconnected, rapidly changing and tightly coordinated submaps..." (Derdikman and Moser, 2016, p561). Here then, 'map fragments' can be seen as containing local 'constellations' of landmarks and as being 'anchored' to a reference frame, echoing results from studies such as the investigation of representations of connected environments using lab-based spatial learning and memory experiments (Han and Becker, 2014).

So it can be seen that cognitive strategies for navigation make use of hierarchical representations as a means of reducing ‘cognitive cost’, with work by (Wiener, Schnee and Mallot, 2004) providing one example of the broad empirical support for: a) regions being learned early on in spatial experience, b) region representations being used in reasoning about connectivity as opposed to reasoning using place

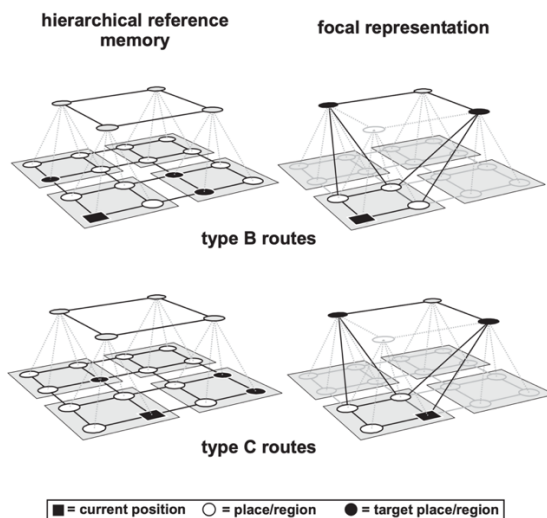


Figure 12 Hierarchical basis of navigation strategy: ‘fine-to-coarse’ route planning for two route types (‘B’ and ‘C’) where type refers to starting position relative to the global environmental structure (Wiener, Schnee and Mallot, 2004).

connectivity alone, for example in route planning (following (Poucet, 1993)), c) that region-level information can be used to overcome issues that could arise out of missing or imprecise information at the detailed level, and d) that region level representations serve to constrain the search space for specific locations by restricting search to a region known or hypothesised to contain the location. In Figure 12 these findings are expressed through the illustration of route planning in an experimental environment using a ‘focal representation’ that uses both fine-grained ‘place connectivity’ and coarse-grained ‘region-connectivity’. Here ‘focal representation’ refers to integrating a fine-grained representation of current and ‘close’ places, with a coarse-grained representation of distant places using a super-ordinate region. Moving from ‘hierarchical reference memory’ to a focal-representation is then a sequence of cognitive processes from constructing a coarse-grained representation of connectivity to the disambiguation of focal regions (‘coarse-to-fine’), then to a ‘fine-to-coarse’ representation, building on previous work in (Wiener and Mallot, 2003).

This model of region and place connectivity continues to gain empirical support, for example in a recent study into the integration of local and global ‘memory units’ (Strickrodt, Bühlhoff and Meilinger, 2019). Here it was observed that global representations are ‘consulted’ when information needs to be accessed beyond a local threshold or ‘regional border’. In terms of the structure of these local, regional and global contexts the approach echoes the proposal by Yoshino of traversed path sequences as being a key determinant of the natural segmentation of space. In (Strickrodt, Bühlhoff and Meilinger, 2019) experimental results were argued to show that study participants did not exclusively use the embedding of local reference frames as a basis for environmental representation, and neither did they use a single reference frame encompassing the whole environment. Rather, it was shown that memory unit formation occurs in a distinct manner across hierarchical levels of the representation, providing yet further evidence for the hierarchical nature of spatial representation as discussed in (Hirtle and Jonides, 1985). This work also provides further evidence for the flexibility of representation across spatial scales, for example the use of both ‘high-resolution egocentric snapshots’ as well as super-ordinate representations learned gradually over time (Ekstrom and Isham, 2017).

Finally, in terms of the relation between a discrete object and the wider representational structure, it has been suggested that experimental evidence points toward spatial memory as containing two ‘dissociable representations of the location of an object, “...a schematic map of where an object lies with respect to local geometry, and a representation of the identity and location of the environmental context...” (Marchette, Ryan and Epstein, 2017, p79), echoing the previous discussion on the relation between hierarchical reference frames and focal representations (Wiener, Schnee and Mallot, 2004).

## **4.2 The schematisation of events and the formal representation of spatial knowledge**

In the previous section, literature was discussed around how people naturally develop spatial knowledge, and around the navigational strategies that shape action across tasks and subtasks. Here, the review turns to the formalisation of cognitive theories of spatial knowledge, such that representations may be stored, queried and operationalised in practice. The first subsection addresses the underlying structure of spatial representations, with the second subsection then exploring the notion of formally representing situated experience in a spatial context, as opposed to simply representing static entities that are dissociated from the active roles they may play in cognitive processes.

#### 4.2.1 THE STRUCTURE OF SPATIAL REPRESENTATION

There is a deep relationship between the structure of a representation, its content, and the way it may be operationalised. While it has been shown that a fundamental prerequisite for the emergence of a stable complex system are processes that lead to the formation of hierarchical abstractions or inherent ‘hierarchical modularity’ (Siyari, Dilkina, and Dovrolis 2019), it has also been argued that strictly hierarchical structures rarely provide an effective underlying model of the world in terms of actual lived experience. An early and seminal argument made in urban studies and urban planning was that an idealized ‘tree’ structure leads to inaccurate characterisations of the basic structure of experience (Alexander 1966), with the central case study being our conceptualisations of the physical structure of cities.

A tree in this sense may be thought of as a connected acyclic graph in which any two vertices are connected by exactly one relation. For example a partonomic relation between an instance of ResidentialNeighbourhood and ResidentialBuilding, adapted from (Lüscher, Burghardt and Weibel, 2007). The graph is connected in that there exists a relation that connects every vertex to the rest of the structure and is acyclic in that no path through the graph forms a cycle in which a previously visited vertex may be visited again. In terms of describing a structure such as this in which relations are directed, the result is a set of hierarchical relations, and one that crucially leads to no overlapping subsets. Continuing with the example, each instance of ResidentialBuilding would have a mereological relation with one and only one instance of ResidentialNeighbourhood. In terms of a city then, an example would be to describe an instance of City as the root, with vertices representing neighbourhoods that are ‘within’ this city object as being descendent to this root vertex. The practical implication of the inability to form overlapping subsets means that neighbourhoods and their members will always remain distinct within this representation.

While there are scenarios in which comprehending and reasoning about a city through this type of strict hierarchy may be useful, there are also many in which this representation quickly comes into a severe tension with the underlying reality of how the environment is both understood and used (see Alexander for further explanation and an analysis around structural differences between ‘historic’ and ‘artificial’ cities). The practical implications of Alexander’s argument are that a structure that supports overlapping subsets is required if we are to better represent the lived reality of complex structures such as cities. Here Alexander proposed the semilattice as a candidate formal structure to address this fundamental issue. In contrast to a tree, a semilattice is a partially-ordered set with an upper or lower bound (a ‘least upper bound’ or ‘greatest lower bound’) for any non-empty finite subset. In straight-forward terms, it is a structure that supports ordering (‘the

neighbourhood is *in the city*'), but in which overlapping subsets may form, meaning there may be more than one path to reach the same vertex:

*“Both the tree and the semilattice are ways of thinking about how a large collection of many small systems goes to make up a large and complex system. More generally, they are both names for structures of sets.”*

(Alexander 1966, p1)

As succinctly expressed in the above quote, both these representations provide a structural substrate for organising groups of objects with hierarchical relations between them, however in the semilattice, the membership of groups may overlap, and so we are now able to formally express that ‘river A is an object that exists within both region Y and region Z’, or that the membership of ResidentialBuilding B is contentious, and may be understood as being contained within neighbourhood M or neighbourhood N. This insight has helped to provide a basis for exploring the structure of ‘cognitive spaces’ with respect to spatial knowledge (Hirtle 1995). Here then, with reference to (Hirtle, & Heidorn, 1993; Medyckyj-Scott, & Blades, 1992), Hirtle discusses the observation that there can be discord between the ‘internal’ cognitive representation of space and spatial objects, and the way the various underlying phenomena are in fact treated in systems such as those supporting navigation.

Returning to Alexander’s explanation, he conceptualised elements within subsets as ‘units’, with the overlapping interactions and experiences of people with these basic units giving rise to heterogeneous groupings. Rather than spatial objects with a relatively clear ontological position within a spatial representation, e.g. ResidentialBuilding, Alexander provided an illustrative example using objects on a smaller scale, some having a non-fixed physical location – a sidewalk, a newspaper rack, a newspaper, a traffic light. Through the simple, small scale example of a person being drawn to pause and read the headlines as they wait for the green man, we see in microcosm how objects may be grouped heterogeneously based on day-to-day lived experience, and how objects may be elements in many different *configurations*. While it is intuitively clear that such groupings of objects as newspapers and sidewalks form part of the fabric of everyday spatial experience, it is also clear the difficulty with which such objects and their grouping are able to be formally expressed in information systems.

While the formalisation of geographic knowledge has resulted in a variety of structural motifs, it is argued three high-level structures serve to characterise the overall landscape, on a continuum from rigid to flexible in terms of the way that relations between entities may form. A tree is the most rigid in that entities cannot belong to more than one set, as discussed. A semilattice is more flexible as it

structures relations based on a partial ordering which does afford such overlap. The most flexible then is the rhizome (Deleuze, Guattari and Guattari, 1987), or ‘creeping rootstalk’. In contrast with arborescent structures, this conceptualisation is based on the notion that any point on a rhizome is able to connect to any other point, and any part may be a point of ‘entry’ or ‘exit’ to or from the structure (Kamalipour and Peimani, 2015). While the rhizome concept is usually invoked in the theoretical discussion of ‘urban assemblages’ (Anderson and McFarlane, 2011), to treat it as a coherent structural form that could be used as a basis for modelling data is still feasible, albeit impractical.

Further research into the practical implementation of naturalistic knowledge structures has focused on the semilattice, with Timpf and Rüetschi arguing for it to be the strong candidate for supporting information systems that better reflect our natural understanding of geographic environments (Rüetschi and Timpf 2004). Here the authors discuss the ‘schematic’, ‘fragmentary’ and ‘partially-ordered’ nature of spatial knowledge, and propose ‘schematic geometry’ as a basis for modeling. This concept can be understood in terms of using the semilattice as a formal basis for representing *cognitive topology* (Lakoff, 1990). While the authors use a different term, the fact that the structure is a partial order of image-schemata would suggest it is effectively equivalent to a cognitive topology. In Figure 13, an example is




Figure 13 ‘Schematic geometry’ of the ‘Enge’ station in Zurich, (Rüetschi and Timpf, 2004).

provided for a public transport station in Zurich. In this example. The fact that both the ‘subway’ and the ‘main entrance’ can be understood as entities within more than one superordinate region clearly supports the overall argument for partial ordering as a critical aspect of any formal representation that reflects lived experience of the geography. Alexander discussed the ‘physically unchanging’ aspects of the city as being ‘spatial units’, and the example in Figure 13 also points toward the content of representation in terms of possible ‘units’, linking back to the

discussion in (Mackaness, Quigley and Tanasescu, 2011) around the grouping of entities according to function.

Reflecting on the present research, information needs vary over the event partonomy with respect to the appropriate level of abstraction, with an ambition being to surface objects from varying levels of abstraction in the same view. Using Figure 13, this could be understood in terms of surfacing an individual entity at the levels of ‘platform 2’ or ‘tramway stop’ in the same view as entities at the level of ‘station’, for example. Here research into ‘vario-scale’ data structures is highlighted, with methods of Generalized Area Partitioning pointing toward a possible way to address this issue in practice (van Oosterom, 2005). Reflecting on the event horizon proposal discussed in (Mackaness, Quigley and Tanasescu, 2011), here the author demonstrates how ‘map fragments’ that represent features from ordered levels of abstraction may be surface in single integrated views by taking a ‘non-horizontal slice’ across the overall structure. Following the explanation of Quigley’s approach in the Chapter 3, this can be understood in similar terms to a ‘multi-cut’ selection ‘down’ the graph.

#### 4.2.2 THE SCHEMATISATION OF EXPERIENCE

While maps support cognitive tasks by transforming a physically spatial problem into an abstract representation (Freksa *et al.*, 2017), of central concern in the present research is the representation of geography from the perspective of *experiential realism* (Mark and Frank, 1996). While physical entities clearly play a role in spatial experience, the spatial extent of a journey is such that only a small subset of entities within the extent may be surfaced in a view, even at large scales. Additionally, the role an entity plays in situated experience may also have a bearing on the appropriate level of abstraction or the way in which the entity is grouped with other entities at varying scales. Here then, the notion that spatial experience itself may be formalised through schematic representation offers a powerful way to bridge this divide between the objective nature of the geography and the entities that may be of specific contextual relevance within the current task hierarchy. The nature of this ‘naïve’ spatial knowledge (Egenhofer and Mark, 1995) has been shown to be fundamentally topological (Mark, 1999), with metric information being secondary to these topological structures. As discussed by Rumelhart and Ortony, the notion of such schematised representations or ‘schemata’ can be traced back to Kant, and can be understood as “...generalized concepts underlying objects, situations, events, sequences of events, actions, and sequences of actions...” (Rumelhart and Ortony 1977, p101). Here the authors describe four essential characteristics of schemata, that: 1) they have variables; 2) they can be embedded, one within the other; 3) they represent generic concepts which vary in their level of

abstraction; and 4) they represent knowledge as opposed to being definitions. These 4 characteristics together mean that schemata in the general sense are a very powerful way of expressing spatial knowledge. In fact our basic ability to form representations and to reason about phenomena in this way can itself be seen as arising from everyday spatial experience (Lakoff and Turner, 1989). While the cognitive linguistic theory of image-schemata is discussed throughout this thesis, in this subsection a brief review is provided as to the relevant research in terms of formalising these structures, as well as the related concept of ‘choremes’.

While primitive image-schemata can easily be identified in the language used to express an idea or an experience, of central concern is how these primitive objects combine to form representations of experience at a meaningful level of abstraction given the domain. For example the primitive VERTICALITY schema would most likely be expressed if a person was specifically asked to describe moving from a road to a raised pavement, however in the context of an overall description of a journey, it is very unlikely that the act of stepping onto a pavement after crossing a road would feature. So, it can be seen that formalising the schematic representation of experience is partly to do with determining the possible primitive elements, but as with all spatial information, arguably the key task is really in determining the appropriate level of abstraction. Following this, the question is how this type of representation might be formalised such that information may be stored, queried and displayed in a way that aligns with the lived experience of the domain.

It has been shown that image-schemata can be used as a basis for understanding spatial experience across sequences of subtasks (Raubal, Egenhofer and Pfoser, 1997), with schemata combining in ‘blocks’ that form repeated patterns of experience as observed through interviews. This links with research discussed in the previous subsection, with Timpf and Rüetschi proposing that image-schemata could be linked into partially ordered sets that represent naturalistic spatial knowledge of geography relative to wayfinding tasks. Here entities corresponding with image-schemata were linked in a graph representation formalising part-whole relations. In (Raubal and Worboys, 1999) however, a similar approach saw the image-schemata as a way to formalise spatial knowledge at each ‘view’ across an action graph. In other words, rather than a graph structure that represents part-whole relations it represents the transition of an agent between distinct places, where image-schemata then encode the salient aspects of spatial experience across this sequence of actions. Various approaches have been developed for formal reasoning based on image-schema theory, for example see (Frank and Raubal, 2000), (Amant *et al.*, 1999) and (Hedblom *et al.*, 2017), although it should be emphasised these are all approaches to qualitative spatial reasoning as opposed to being techniques that have been demonstrated in relation to map construction.

Returning to the discussion about event perception in Chapter 2, Tversky and Zacks noted the schematic representation of events encodes information about goals and plays a role in planning, and in the understanding past actions (Zacks and Tversky, 2001). While only the Raubal and Worboys paper from the examples discussed above explicitly includes a representation of movement between states, it is argued that all of these representations express action at least implicitly, and that by narrowing selection to small subsets of entities and relations that play specific roles in perceptual events, these subsets can be understood as expressing lived experience over time.

A more explicit example of the ‘schematisation of action’ in terms of navigation and wayfinding can be found in the adaptation of ‘choremes’ to the functional role of navigable networks. Brunet proposed choremes as ‘models’ that underpin repeated patterns in the ‘conceptual organisation of space’, where an emphasis was placed on patterns of change and interaction as opposed to discrete spatially bounded objects that exist wholly at a particular time instant; see (Brunet, 1980) and (Brunet, 1986). Following these conceptualisations, Klippel proposed ‘wayfinding choremes’ as a basis for disambiguating a ‘functional perspective’ from the static structure of the geography (Klippel, 2003). Here function refers to elements of journeys that are known to be key drivers of event segmentation such as the origin, the destination and decision points (Tenbrink, 2012). Again, it is argued that even if movement or a transition between states is explicitly represented, by disambiguating elements in the representation of a geography that play a specific role in navigation, something of the dynamic nature of the event sequence is still implicitly captured. From an ontological perspective then, even if actions are not explicit entities (Timpf, 2002), there is arguably a distinction between a feature that happens to be within a spatial extent and a feature that plays a direct role in cognitive processes.

There are however formalised representations that explicitly capture action, movement and uncertainty. Broadly speaking these approaches treat the role of a navigational agent explicitly, and to varying degrees can be understood as theories of cognitive mapping in that they aim to replicate the natural acquisition of spatial knowledge over time, or the natural basis of spatial reasoning given uncertainty in the choice of action. A central aspect of many such approaches is some form of representation of sequences of ‘places’ and or actions, where place sequences lead to the representation of paths, and groups of places lead to the representation of regions. A classic description of such an approach can be found in Kuiper’s ‘TOUR’ model (Kuipers, 1978). Examples of related approaches can be found in the ‘PLAN’ theory of ‘cognitive mapping’ (Chown, Kaplan and Kortenkamp, 1995), and the ‘NAVIGATOR’ model of environmental learning (Gopal, Klatzky and Smith, 1989), see Figure 14. Of particular note is the SSH (Spatial Semantic Hierarchy) model that further specified the essential proposals underpinning ‘TOUR’. These ‘theories

of cognitive mapping' and the related domain of cognitive robotics are of key relevance to the present research as they address aspects of modeling that are not present in the cartography literature that have the explicit aim of representing spatial knowledge in ways that reflects the situated experience and understanding of navigational agents.

In general, a navigational strategy such as 'exploration' or 'shortest path' provides the basis for progression between 'places' in these models, where a distinctive 'place' is experienced through 'views', and knowledge about distinctive locations is developed through the integration of views over time. Here there is a basic difference between knowledge developed during traversal between distinctive places and knowledge 'within' the perceptual boundary of a distinctive place. Metric information within a 'place' can then be integrated into flexible global structures, with new information integrated over time, and hierarchies developing both in terms of groupings of places, and in terms of the grouping and ordering of paths (Remolina *et al.*, 1999). Hierarchical graphs often form the basis of hierarchical representation (Zivkovic, Bakker and Kröse, 2006), with nested regions supporting efficient spatial reasoning of the kind discussed previously in relation to Wiener *et al* (Wiener, Schnee and Mallot, 2004).

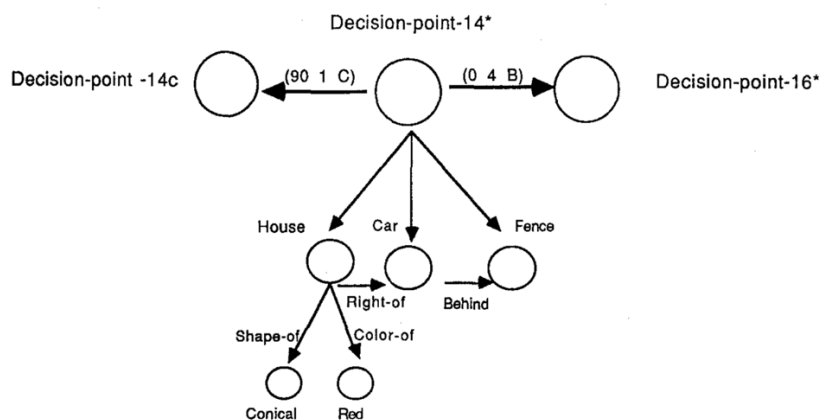


Figure 14 Hierarchical structure representing information about a decision point in the long-term memory of NAVIGATOR. The labels on the links between decision-points indicate orientation, number of steps, and street information (Gopal, Klatzky and Smith, 1989).

These approaches include two characteristics of critical relevance to the present research: both the formalised representation of situated navigation in terms of progression between distinctive states, and the flexible integration of metric and topological information such that navigation occurs seamlessly 'through' representations based on varying ontological foundations (Konolige, Marder-Eppstein and Marthi, 2011). While digital cartography has been traditionally bound

by fixed representations based on rigid scale hierarchies, the flexibility to express situated understanding that is inherent in many formal approaches to simulating the acquisition of spatial knowledge means these approaches may point the way to alternate methods of map content selection. In terms of using these types of models for explicitly generating map-like information, at present this can be understood to be essentially limited to the work on SLAM (Simultaneous Localisation and Mapping) in which algorithms support mobile robots in the task of developing semantic interpretations of sensor data such as room segmentation (Bormann *et al.*, 2016). It is argued in fact that the central task of the present research can in a sense be understood as *the opposite* of the task in SLAM. Across the SLAM work, the essential premise is that the robot does not have a pre-existing model of the environment, and develops a representation over time in a way that would be meaningful to a human. In the present research, an underlying database of features is the starting premise, with the central task being to understand the situated experience of the agent.

### 4.3 Information supporting goal acquisition across navigational episodes

Now a summary has been provided around our knowledge of the cognitive foundations of navigational behaviour and of the basic character of our ‘internal’ representations of navigable environments. Here the focus turns to external navigational information and how the form of this information both mirrors and augments our natural cognitive processes.

#### 4.3.1 EXTERNAL REPRESENTATIONS AS A BASIS FOR SPATIAL KNOWLEDGE ACQUISITION

External information serves to significantly reduce the resources required to acquire knowledge. Through the use of external information, for the same energy expenditure, our cognitive processes become ‘extended’ (Clark, 2013), with these informational ‘cognitive artefacts’ serving to ‘augment the mind’ by externalising memory and processing (Tversky, 2000). The ‘burden of memory’ (Sterelny, 2004) is externalised by ‘offloading’ data and information to artefacts that have some form of representational state (the wall of a cave, a piece of paper, and so on). Processing can be externalised by conducting calculations on an external device, for example by using a system that calculates quantitative relations and displays them

graphically, resulting in a complex processing problem being reduced to a simpler perceptual problem (Sterelny, 2004).

In terms of spatial environments, external information serves to expand our declarative knowledge, for example knowledge of the existence of a geographic feature, and serves to provide a short-cut to acquiring procedural knowledge such as landmark sequences, as well as a short-cut to configural knowledge such as the layout of the city (Golledge *et al.*, 1985). As Norman noted, these informational artefacts: “...serve a representational function that affects human cognitive performance...” (Norman, 1991, p7). Some key examples of spatial knowledge that may be acquired through interacting with representational cognitive artefacts include the knowledge of existential and universal quantifications of spatial entities, regionalisation at different scales, region connectivity, place connectivity, hierarchical ordering through containment and partonomic relations, sequential ordering, adjacency and proximity.

While information may be conveyed either graphically or textually, graphics allow for communication ‘using space and elements in space to express relations and meanings directly’ (Tversky, 2000). A key distinction then between textual and graphical information is flexibility. While graphics are able to convey specificity in spatial relations more efficiently (Stenning and Oberlander, 1995), textual information can convey both determinate and indeterminate information ‘with equal ease’ (Mani and Johnson-Laird, 1982). This has been expressed in terms of the difference between ‘mental model’ representations and propositional representations (Johnson-Laird, 1983), with pictorial information providing a particular immediacy to the instantiation and refinement of model-based representations that comprise a set of entities and their spatial relation (Jahn, Knauff and Johnson-Laird, 2007).

Focusing on maps specifically, following Charles Sanders Peirce’s classic description of a triadic relation between interpretant, sign and object (Peirce, 1969), maps can be seen as iconic representational cognitive artefacts (Heersmink, 2013) as they exhibit relatively high structural isomorphism between sign and object (the underlying geography). Maps support the acquisition of configural knowledge by reducing scale, but do so in a way that is systematically inconsistent, through the exaggeration of road and river widths for example. Crucially, maps also omit information in a systematic and targeted way, and so a map’s usefulness stems from a reduction of information as well as a reduction of scale (Tversky, 2000).

Maps can be seen as schematising and exaggerating information in a way that reflects our internal spatial representations. There is naturally a correspondence between the two in the way that information is omitted and regularised, through the mixing of perspectives, the reduction of dimensionality and by way of exaggeration (Tversky, 2000). Comprehending information in maps relies on

processes of matching, and so systematic inconsistency must reflect correspondences between internal and external representations:

*“The external geographic space is being visualized, measured, modelled and reported upon. But as soon as it is visualized on a screen in the form of a map... that visualization in itself is also a space.”*

(Davies, Li and Albrecht, 2010, p19)

These miniature, simplified spaces must then support comprehension based on correspondences between, for example Lynch’s foundational classes, and the depiction of entities using the basic graphical classes of points, lines and areas. Also, to successfully support extended cognitive processes (Clark, 2013), geographic representations require ‘portability of information’ in that representations should to some extent be constrained by conventions and regularities that mean information can be ‘ported’ to multiple distinct task-contexts. To draw on Sterelny’s example, the “...same map of a chunk of New Zealand’s fiordland will be used very differently by hikers, hunters and fishermen...” (Sterelny, 2004, 251).

Highly schematised representations serve to focus attention on topological relations, with information about connectivity and sequential ordering having primacy as opposed to metric spatial relations and topographic detail. Gestalt reasoning in the schematic network case is simplified by the focus on the intersection of lines as implying potential paths of movement, albeit within an overall layout that exhibits some configural correspondence with the spatial structure of the underlying geography. While octilinear layouts that constrain the angle of connecting network segments to 90 and 45 degrees have been shown to support the reduction in cognitive load of interpreting network maps, it has been found that the correspondence with the underlying geography is still a significant factor in legibility. Here it has been postulated that the octilinear layout of the London Underground Tube map may be successful in part because of a natural correspondence between this form of layout and the underlying physical network itself, i.e. ‘geographical conservation’ (Mackness and Reimer, 2014), as opposed to necessarily being a ‘gold standard’ network layout for urban route planning across all geographic configurations (Roberts *et al.*, 2013). As discussed in (Gallotti, Porter and Barthelemy, 2015), the ‘cognitive limits’ of planning in large cities can be quickly reached, particularly where a change of transportation mode is required, with highly schematised maps serving to make navigational decision-making in these environments an efficient and tractable process (Guo, 2011). While navigational information can be conveyed textually or through topographic representations, in the case of public transport the general convention of providing schematised network maps is founded in their essential, empirically supported utility (Bartram, 1980).

In the case of topographic maps, without the same emphasis on network topology, the plane may be regionalised (Bertin, 2011) in a far larger variety of ways, providing a basis for conveying more complex Gestalts. Topographic map design also lends itself to greater differences in representation across scales, with the level of schematisation being reduced as scale increases. While topographic representation still supports the procedural knowledge that is inherent in schematised network maps, the broader affordances of topographic maps lead to a greater capacity for supporting declarative and configural knowledge acquisition (Golledge, Dougherty and Bell, 1995). While highly schematised network views support legibility by nature, design constraints that support legibility in topographic maps reflect a more challenging set of issues, with the high-level considerations being visual perception (MacEachren, 1995), i.e. objects are large enough given visual acuity; and map complexity, i.e. the quantity of objects is not too dense for the available area (Harrie and Stigmar, 2010). Given the likely tasks of the map user at each map scale, a related consideration is the natural perceptual constraints that shape the number of distinct objects we are able to attend to, both in terms of visual perception and in terms of working memory (Miller, 1956).

While textual information clearly affords the ability to convey determinate, propositional knowledge, it should be noted that text can also support the development of mental models and hierarchical orderings. Even if a textual description has not been explicitly formulated to elicit the development of an internal representational model of entities and the relations between them, given a certain level of content in textual information we still naturally form ‘Gestalt-like representations’ such as spatial groupings (Dumitru and Joergensen, 2016). So, despite the physical dimensionality of language, the ‘linear sequence of symbols’ (Chomsky, 1959), configural spatial knowledge based on the construction of ‘snapshots’ of objects and the relations between them may still emerge from knowledge acquisition through text (Mani and Johnson-Laird, 1982). In fact, as shown in (Ferguson and Hegarty, 1994), these representations can form in a very similar way to the egocentric, hierarchical representations characterised by the anchor-point hypothesis, for example the way that certain landmarks serve as a reference for other features within a localised region.

The cognitive linguistic study of metaphor has shown that understanding, even of abstract concepts, can be interpreted as fundamentally grounded in everyday spatial experience, with ‘image-schemata’ (Lakoff and Turner, 1989) serving as ‘preconceptual’ structures derived from these experiences (Lakoff, 1987), as discussed previously. Here then we can see that there is effectively a transitive relation between aspects of the physical environment that are attended to, the ‘mental model’ of this environment that emerges through situated interaction, and the expression of the image-schematic understanding conveyed through language.

In other words, while there is empirical support for the construction of configural ‘environmental images’ based on textual ‘input’ to a human agent, the linguistic expression of understanding by the agent (i.e. ‘output’) then conveys the content and structure of this internal model to an ‘outside’ observer.

As discussed in (Baddeley, 2007), where textual information contains content that readily translates to an imagistic representation, it tends to be far easier for this content to be committed to memory than content for which an image, even if manifested associatively rather than implicitly, is not naturally forthcoming. It should be noted here that there is a distinction between a memory that is imagistic in the sense of visualising an object or scene, and the ‘image-schematic’ representations described by Lakoff and Turner. Image-schemata are proposed as ‘preconceptual’ in the sense that they serve to provide a spatial basis for understanding, however these structures are not images per se, but are better characterised as being compound sets of primitive topological relations. Image-schematic reasoning and the inference of subjects’ active image-schemata based on linguistic data will be discussed further in Chapter 5.

So, we see that there is in fact a deep relation between textual and graphical information. While each has relative strengths or ‘performance characteristics’, there can be overlap, correspondence and ‘equivalence of inferences’ between the cognitive representations that develop based on each (Stenning and Oberlander, 1995). Additionally it should be noted that while diagrams limit abstraction and ‘thereby aid processability’ through ‘specificity’ (Stenning and Oberlander, 1995), it is clearly also the case that in practice textual and graphical information are integrated and complement each other, with explicit integration offering the potential to reduce extraneous cognitive load, i.e. to reduce ‘cognitive activity that is engaged in because of the way the task is organized and presented’ that is not inherent to the underlying goal of knowledge acquisition (Sweller *et al.*, 1990).

#### 4.3.2 CONTEXT-DEPENDENT INFORMATION NEEDS

As discussed previously in this chapter, attention both shapes and is shaped by context. Different forms of information serve as external resources that are relatively more or less suited to supporting task-dependent cognitive processes. This was explored in the previous subsection with the short review of the practical implications for knowledge acquisition for presenting information in textual or graphical form, and in presenting information in more or less schematised graphical representations. In considering information for navigation however, a key issue is that in many cases a journey will include transitions between a number of task-contexts, with different forms of information being more or less suited to supporting

goal acquisition across different tasks. Consider the case of highly schematised network maps. While information in this form supports the focus of attentional resources on connectivity and sequential ordering by constraining the level of detail and relaxing absolute spatial relations, this approach is clearly only suitable for route planning and route following on the network itself, and a switch to a less schematised topographic representation would be required if, upon leaving the network, the traveller must navigate by foot through an unfamiliar area to the ultimate destination. Over the course of a journey then, particularly where more than one form of transport is required, sources of information may include schematised network maps, topographic maps, scheduling information, environmental cues, signage, textual route descriptions and so on.

For information to be coherently integrated into broader knowledge about the journey including transitions between contexts with distinct informational requirements, matching must occur such that some subset of the information is understood to relate to an event for which a subset of another form of information is also related. Due to the nature of the domain, events related to navigable environments always correspond with either a location or a spatial region. Overall context for a given journey is then the union of these task sequences, integrated through the overlap of behaviourally significant, or at least perceptually distinct events. For example in the case of navigating to a bus stop on foot, followed by a transition to the network because of boarding a bus, the location of this transition and the abstract representation of the bus stop can serve as a basis for matching a topographic map and a schematised network map, echoing the comment in (Bateman *et al.*, 2007) regarding spatial location as an ‘individual quality’ that can be used as a basis for combining representations that diverge in their ontological basis. A key affordance of a topographic map is that it can support a diverse range of graphical cues for matching between, say the map and the external environment, or between the map and a textual description.

#### 4.3.3 EXTERNAL INFORMATION TO SUPPORT SEQUENCES OF DISTINCT TASK-CONTEXTS

So it can be seen that information of a particular form can be more or less suited to supporting knowledge acquisition given the traveller’s context, and that over the course of a journey, it is often necessary to use information in different forms as a basis for gaining knowledge across event sequences with distinct attentional requirements.

In (Godfrey and Mackaness, 2017) a number of approaches were reviewed that aim to support informational needs across distinct task-contexts within integrated views.

In these examples the elements that serve the matching function as discussed previously in fact match distinct forms of representation in the same view, as opposed to serving a matching function across distinct views. These design approaches essentially seek to reduce extraneous cognitive load in that the information to support knowledge acquisition across distinct task-contexts is integrated, minimising the need for cognitive resources to support matching across distinct views (Sweller *et al.*, 1990). Following (Spiers, Maguire and Square, 2008) and the study of distinct ‘thought sequences’ as the agent traverses the environment, the target outcome of these approaches is that the integration of information to support the overall set of distinct thought sequences can be expressed in a single view, reflecting the lived experience of a journey as continuous, despite context transitions based on perceptual segmentation that are driven by behaviourally distinct events.



Figure 15 ‘Interactive spider map’ including routing overview and focal topographic detail (Beatriz Santos Gonçalves, 2019).

In the ‘hard boundary’ case, illustrated in Figure 15, distinct forms of representation are nested, with correspondences formed between the two by way of continuity between roads extending to the interior of the boundary and route representations beginning at the exterior of the boundary. Here task-contexts can be seen in terms of navigation from or to a bus stop within a central ‘hub’, and route planning from or to this central hub using the bus network. In the ‘focus map’ case, views also integrate a focal region within a broader geographical context, however rather than there being a hard boundary between a nested region and the broader representation, geometric distortion is used to integrate the ‘global’ and ‘local’ views by relaxing absolute spatial relations around the focal region. Techniques that explore the focus map case often have some similarity to the ‘fisheye’ views proposed in (Furnas, 1986).

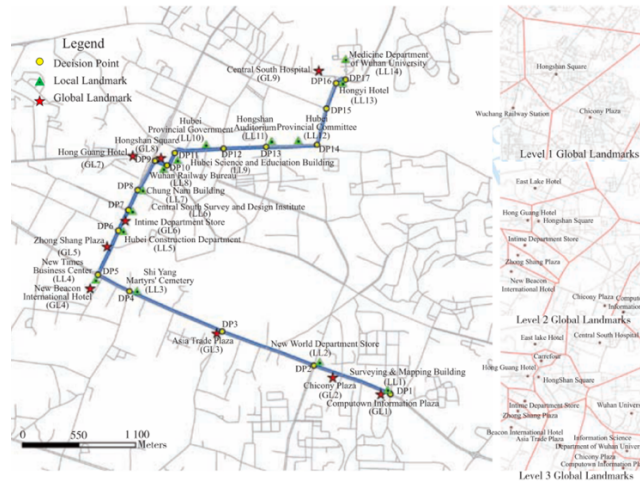


Figure 16 Disambiguation of route-based information including landmarks and decision-point (Zhao, Li and Li, 2011).

Another approach is ‘variable level of detail’, with content selected given the specific requirements of the current journey. For example in (Schmid, Richter and Peters, 2010) a route was extracted from the generic survey view, with topographic detail disambiguated at the start and end of the journey, reflecting the general observation that higher detail information is often required during event sequences that are close to the origin and the destination. In Figure 16, a similar approach can be seen from




Figure 17 Schematisation of a route between two cities including journey-specific landmark, route and survey information (Schwering et al., 2019).

Zhao et al in which route-based information is disambiguated including the sequence of anticipated decision-points. In the example in Figure 17 information is also disambiguated across the extent, however here a context-specific regionalisation of the geography leads to a schematised view that integrates explicit regions with route-based information (Schwering *et al.*, 2019).

An approach that extends ‘variable level of detail’ is ‘variable scale route’. Here, in addition to selection being context dependent and therefore non-uniform over the extent, geometric distortion is also applied so as to provide an explicit emphasis on certain journey segments using greater screen area. In (Agrawala and Stotle, 2001) the underlying observation was that car journeys can be characterised by periods of higher and lower frequency of decision-making, and so techniques were developed to not only select a variable level of detail, but to geometrically compress route segments through areas of lower detail, providing emphasis on the information that actually supports navigational decision-making. This approach can also be seen in the various implementations of ‘focus and context’, such the example in Figure 18.




Figure 18 ‘Focus + context’ Taipei metro map example (Wu, 2014)

Finally, linking back to the summary of research relating to the cognitive and neural basis of spatial knowledge, we can see that in all these cases, what is essentially happening is that information supporting distinct context-specific events sequences is being made explicit in selection and display. Taking the first case of the ‘hard boundary’ as an example, we can see that the central hub conveyed through the higher detail representation may be used by a traveller in orienting relative to any one of a number of nodal events within distinct journeys to, from or through this centre. Referring again to the ‘memory space’ view of the cognitive representation of navigable environments (Eichenbaum *et al.*, 1999), multi-modal journeys across the city can be seen to form episodes in memory, with transitions within the central

hub being distinct events that serve to link key sequences within episodes. To elaborate on this view of spatial knowledge in terms of task hierarchies, we can also see that detailed topographic views are often used to support action sequences that combine to form distinct events, whereas events themselves combine to form event sequences over continuous experience, with highly schematised views providing an overview of event sequences, and action sequences being loosely implied by network connectivity.

#### **4.4 Concluding comments**

Over the course of this chapter literature has been reviewed across our underlying knowledge of the cognitive representations that support navigation, the formal representation of this knowledge, and the way that external information expresses this knowledge in both textual and graphical forms. A particular focus has been the integration of different forms of representation that reflect natural approaches to navigation and wayfinding. While there are many examples of overlap across the literature, it is argued that a key gap is in using formalisations of the natural cognitive representation of a journey context as a means of explicitly driving map construction. While SLAM techniques are based around the formalisation of a number of observed characteristics of our natural spatial representations, these techniques seek to construct information based on sensor data as opposed to selecting information that is already stored in an underlying data model. Over the following chapters, empirical observations will be presented, and an overall argument developed as to how systems may express the situated context of a navigational agent in selection, and ultimately, how to generate context-specific views that support human agents in goal acquisition across complex navigational episodes.





# 5 Methodology

Methodologically the ultimate aim is to develop a framework for data modeling, selection and display that reflects the findings resulting from the empirical validation of hypotheses. Here the overall set of methods will be contextualised in terms of an overarching epistemological position, with particular focus on the analysis of textual and graphical data as a means to develop formalised representations that deliver on the core research aim and objectives.

## 5.1 The ontological basis of investigation

During the twentieth century cognitive theories took hold based on the notion that thought is somehow distinct from perceptual experience, that the flow of sensory information is only implicated in cognitive processes superficially, and that the key to thought is in fact the manipulation of symbols that are ‘stored’ in a manner distinct from the ‘recording devices’ (i.e. sense organs) that allow new data to ‘enter’ (Barsalou, 1999). Representation came to be viewed as ‘amodal’ in that the basis of cognitive representation was conceptualised as being de-coupled from the modal system (e.g. vision) due to symbolic representations bearing “...no correspondence to the perceptual states that produced them...” (Barsalou, 1999, p578). Barsalou argued that the major advances in logic and computation during the twentieth century helped to fuel these ideas, propagating the notion that our ‘internal’ representational system was somehow mirrored in these external constructs (such as semantic nets, lists etc).

Towards the end of the twentieth century however, the tension between this account and the growing empirical evidence became clear, and while the notion of cognition as symbolic computation still lingers, from the empirically validated perspective of neuroscience and cognitive science, distinctions between perception and the basis of representation have all but evaporated. For example, the finding that “...categorical knowledge is grounded in sensory motor regions [of the brain]...” (Barsalou, 1999,

p639) with reference to (Pulvermüller, 1999). Such findings can be seen within the broader context of cognition coming to be understood as fundamentally situated. While critiques of this view remain, for example see (Chatterjee, 2010), a basic ontological claim within the essential premise of the present research is that the situated or ‘embodied’ view of cognition is correct, and that states of knowledge should be understood as arising through embodied experience.

While the notion that an event as a segment of experience through time at a location or across a spatial region is quite clearly situated by definition, a key methodological consideration in the present research is the ontological status of information within this situated context. As discussed by Rupert, cognition can be seen as being ‘embodied’ in the sense that the physical or ‘nonneural’ body plays a central role in cognitive processes, and ‘embedded’ in the sense that cognition happens through interaction with the environment, which can be understood in terms of the use of *external resources* within cognitive tasks (Rupert, 2009). As discussed in (Clark, 2013), the mobile phone is clearly such a resource, with Clark and Chalmers’s extended mind theory (Clark and Chalmers, 1998) being drawn on in the analysis presented in Chapter 9. In terms of maps specifically however, while Sterelny considered maps in a discussion around epistemic artefacts and extended cognitive processes (Sterelny, 2004), in the extended cognition literature in general, maps are mentioned often but with no detailed analysis provided - they are mentioned in passing in support of a broader argument. Here then, when considering maps and other navigational information within the context of the situated empirical studies that were completed over the course of the present research, the analysis draws on the ‘processual’ view of maps presented in (Kitchin and Dodge, 2007). In summary, this view holds that cartographic information is best understood as emerging through practice, as opposed to being distinct from the situated action of the user.

## 5.2 Situated versus laboratory studies

Following the explanation above, the most significant decision in the research design was to focus on field as opposed to lab-based observation. If a particular variable must be isolated in research that relates to the study of cognitive processes, then studies must almost always take place in the lab due to the near impossibility of holding variables constant in field settings, for example the research presented in (Starrett and Ekstrom, 2018).

In terms of investigating the efficacy of navigational information, it is often reasonable for studies that focus on route selection and route-planning to be lab-based given that the lab provides a proxy of the lived context of these activities –

i.e. an indoor environment prior to the physical commencement of the journey. The work presented in (Gallotti, Porter and Barthelemy, 2016) is an example of research at the extreme opposite of naturalistic studies, given an outcome of city ranking by complexity of route planning despite the research involving no empirical observation at all (simply the mathematical analysis of network topology).

In (Roberts *et al.*, 2013) the investigation of metro map usability did involve eliciting data from participants, however study findings relating to the usability of maps when the participant does not actually leave the lab would seem to result in almost speculative findings. If a researcher presents navigational information to a participant in a lab and gives the participant a task, then provides a simplified version of the information that contains enough detail to complete the same lab-based task, then clearly the likely participant feedback will be that the information in the second case was easier to use. It would seem that this says little about the actual efficacy of the information, in contrast with the research presented in (Janarthanam *et al.*, 2013) in which participants were also asked to evaluate their experience of using geographical information, however in this case the users of an experimental system were providing feedback directly after physically completing each section ('leg') of the study route.

An experiment discussed in (Ceci, Rosenblum and DeBruyn, 1999) provides a vivid example to illustrate the central issue. Here 10-14 year olds were given tasks to complete in a 30-minute period, with one study group completing the tasks in a lab and the other in their own home. Here, beyond the intuitively obvious implications of conducting the same task in a different context, differences in behaviour for the group performing the tasks at home versus performing the tasks in a lab in fact had a measurable impact on task performance.

A comparison of observations made during the present research and observations made by Tversky in a lab setting provide support for the argument of field-based studies. Through the recall of geographical scenes during exercises in a lab-based context reported in (Tversky, 1998), some participants described major landmarks first, then provided a sequential 'tour' description. In the studies undertaken during the course of the present research however, *not a single participant* expressed their knowledge in this way, with *all participants* focusing on the representation of the path of traversal first, before adding representations of landmarks subsequently. Fundamentally there are different mechanisms of knowledge acquisition through the lived experience of navigation versus the recall of a representation that is dissociated from physical experience. The experience of the passing of time, the sheer volume of environmental cues, and many other factors can be seen to play a role. The need to physically reach the destination is argued to be of critical consequence to the

development of spatial knowledge through interaction with external information and the processes of attention that shape the perception of events over time.

### **5.3 Observation, analysis and interpretation**

The varied requirements for research data that follow from the objectives presented in Chapter 3 result in a need for data through which observations captured in one form can be contextualised and interpreted through data of other forms for the same participant. In the domain of the present research the potential forms of data can be understood as being either textual, graphical, or objective, and it is argued that no single form is sufficient for drawing conclusions relative to the overall research aim and objectives. Here the observation, analysis and interpretation of each will be discussed in turn.

#### **5.3.1 TEXTUAL DATA COLLECTION AND ANALYSIS**

Textual data includes both written and verbal communication and is often the only practical way of being able to collect data regarding meaning. While meaning can be inferred through posture and body position, for example, in terms of conceptual understanding, textual data is the primary resource. In fact, as discussed in (Barsalou, 2008), the embodied nature of cognition results in knowledge that is not dissociable from physical experience, and through language we can also come to understand the physical experience of the participant. The central aim of empirical observation in the present research is to determine the content and structure of participant knowledge developed through situated experience, and to validate hypotheses with respect to the role of external information throughout this process. Here it is argued that the key methods for achieving this aim with respect to textual data stem from the practical application of image-schema theory (Lakoff and Turner, 1989) and the invariance hypothesis (Lakoff, 1990).

The intuitive appeal, flexibility and empirical foundation of image-schema theory led to the application of the theory to spatial cognition soon after it spread through cognitive linguistics and the emergent ‘grounded’ cognition community. A key example of practical application that is of relevance to the methodological approach of the present research can be found in (Raubal, Egenhofer and Pfoser, 1997). Here the authors used image-schema theory in the study of wayfinding tasks and in the development of a formalisation of the structure of tasks as evident in the study data (following discussion in Chapter 4). In terms of the analysis of textual data gained through participant interviews, the authors term the process of determining the

underlying image-schematic structure of a participant’s understanding as ‘extraction’, i.e. the identification of cognitive topology through parsing natural language descriptions of wayfinding experiences. The overall approach was to conduct interviews including participant exposure to images of geographical scenes (airports in fact), to analyse the interviews and extract the image-schemata as described, and to develop formalised representations of the structure of wayfinding tasks from these ‘extracted’ schemata. For example use of the word “in...” can be interpreted as expressing the CONTAINER schema, and use of the word “on...” can be interpreted as expressing the SURFACE schema.

Compound schemata that reflect the kind of spatial knowledge expressed in natural language can therefore be formalised in terms of topological structures consisting of primitive schemata and relations. As discussed in (Hedblom and Kutz, 2015) a fundamental compound schema is SOURCE-PATH-TARGET (sometimes referred to as SOURCE-PATH-GOAL). This compound schema can be understood as being a directed LINK schema between two OBJECT schemata, for example. Cognitive topology is therefore the overall set of schemata and their relations that together form the relational basis for the cognitive representation of some aspect of experience at a meaningful level of abstraction. Clearly in our case the aspect of experience that is of central concern is the navigational episode.

While a SCALE image-schema has been proposed (Johnson, 1987), the compound image-schemata that are ‘extracted’ from textual data are topological structures that do not include any explicit representation of scale. It should be highlighted that it is argued that image-schemata can nevertheless be used to infer scale based on two forms of analysis. Firstly, that scale is implicit in cognitive topology through the density of image-schemata over spatial regions. In other words given that image-schemata relating to experiences during navigational episodes have a direct mapping to physical space, the spatial extent of the set of entities relating to a compound schema and the density of image-schemata over spatial regions have direct implications for the spatial granularity of the representation. This is most clearly apparent through the subdivision of the traversed path into a greater number of segments reflecting the increased density of decision-points or at least regions characterised by greater attention. The second form of analysis that is argued to provide information about scale is the situation in which a number of compound image-schemata are either implicitly or explicitly expressed in terms of a smaller number of compound image-schemata or even a single image schema. In other words where a compound schema is either explicitly or implicitly described in terms of the COLLECTION schema. Again taking the example of PATH, if in one instance (e.g. a sentence) two distinct locations are expressed in terms of a link formed by a single directed path, but in another instance (another sentence) the same two locations are expressed in terms of a sequence of more than one paths, then implicitly the

PATH compound schema described in the first case is a COLLECTION of the compound-schemata described in the second case, i.e. a smaller scale representation of the same sequence.

It should also be highlighted that it is argued the ‘extraction’ of image-schemata from textual data supports the identification and analysis of events. While the inference of a BOUNDARY image schema from textual data may indicate an event boundary, more generally compound-schemata describe the structure of spatial experience in terms of actions over time (Kuhn, 2007) and the role of entities within events, and so there is a natural correspondence between cognitive topology and event structure. This links with Klippel’s distinction between the structural and functional views of navigable environments discussed in Chapter 4 (Klippel, 2003), and with the related discussion around the schematisation of events.

The analysis of textual data extends beyond the application of image-schema theory of course, for example the identification of environmental variables (Raubal and Egenhofer, 1998), and Tversky’s observations of the way route and survey perspectives can be expressed through language and combined in ‘mixed descriptions’ (Tversky, 1998). Following the discussion in chapter 4, it can also be seen that anchors and hierarchy can be inferred through language in addition to more general propositional knowledge.

### 5.3.2 GRAPHICAL DATA COLLECTION AND ANALYSIS

The situationists utilised pictorial representations to convey ‘psychogeographic’ experience, for example the representation of Paris as a linked group of ‘segments’ experienced through wandering behaviour, each with its own ‘unity of atmosphere’ (translation found in (McDonough, 2016)), shown in Figure 19. Sketch maps and



Figure 19 *The Naked City* (Debord and Jorn, 1957).

the graphical collation of verbal descriptions have been and remain a key method of data collection in the investigation of spatial knowledge and have been used as a tool in research since at least the 1960's. In studies of the spatial knowledge of US cities including Boston and Los Angeles, Lynch noted how participant sketches seemed to 'grow in different ways', with the interpretation being that characterisations of the form of growth reflects something of the how the environments came to be understood. For example environmental images that were "...developed along, and then outward from, familiar lines of movement... [such] as branching out from a point of entrance..." (Lynch, 1962, p87). This observation of the implications of ordering were echoed in findings by Tversky in that ordering during the act of sketching "...reflects mental organization of the domain..." (Tversky 2002, p1), and fits within a broader set of observations that are known to be characteristic of spatial knowledge:

*"The image itself was not a precise, miniaturized model of reality, reduced in scale and consistently abstracted. As a purposive simplification, it was made by reducing, eliminating or even adding elements to reality, by fusion and distortion, by relating and structuring the parts. It was sufficient, perhaps better, for its purpose if rearranged, distorted, "illogical."*

*...However distorted, there was a strong element of topological invariance with respect to reality. It was if the map were drawn on an infinitely flexible rubber sheet; directions were twisted, distances stretched or compressed, large forms so changed from their accurate scale projection as to be at first unrecognizable. But the sequence was usually correct, the map was rarely torn and sewn back together in another order."*

(Lynch 1962, p87)

While Lynch's observations continue to be shown to reflect the essential character of spatial (particularly urban) knowledge, more formal methods of analysis have been developed and refined over time. Building on Lynch's observation of elements that are invariant between objective reality and cognitive representation, Wang and Schwering discuss invariant elements such as 'linear order' (Wang and Schwering, 2015). Here then, although sketches are two-dimensional, the route through the environment and implicitly the route 'through' the sketch leads to context-dependent interpretations of order, with element (e.g. landmark) sequences being preserved, and with the ordering of elements that extend through space relative to points or linear subspaces along traversed paths being disambiguated.




Figure 20 (left) linear order relations and (right) sketch map analysis of ‘both-sided’ linear order relations and topological relation of landmark (‘library’) within a city block element (Wang and Schwering, 2015).

In Figure 20 the traversed route  $R$  can be seen as comprising a sequence of oriented street segments  $r_1$ ,  $r_2$  and  $r_3$ . More than one element may occur at the same point along a route segment (the library and the bank) and one element may be adjacent to more than one route segment, for example landmarks on a corner, as shown here by the library. Here important details of the representation are made explicit, for example that the library is adjacent to  $r_1$  and  $r_2$ , and that  $s_1$  and  $s_2$  in fact share a position in a linear ordering relative to the traversed route. While in this example we see topological containment expressed in terms of, for example, the library as being within a disambiguated city block, in (Jan *et al.*, 2017) the analysis of these type of relations was developed further, with application of Region Connection

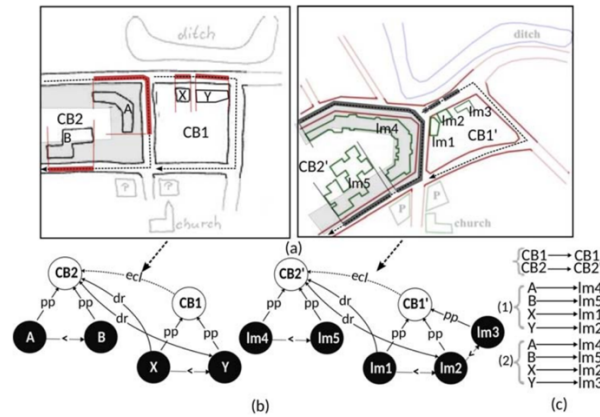


Figure 21 (a) landmarks and city-blocks in sketch and geo-referenced maps, (b) relations between city-blocks and landmarks using the RCC11, RCC5, and the coarse IA relations, (c) diagram representing possible alignments (1 and 2) of landmarks in city-blocks (CB1, CB2) with landmarks in city-blocks (CB1', CB2') in geo-references map (Jan *et al.*, 2017).

Calculus (see (Randell, Cui and Cohn, 1992) following Allen’s formalisation for hierarchical structures in time (Allen, 1981)) to the modelling of hierarchical clusters of elements using identified ‘city-blocks’, with a view to automatically aligning sketches with ‘geo-referenced’ maps.

A number of approaches have been developed for the formal analysis of sketch maps, such as techniques to analyse the geometric distortion of regions that contain distinct elements (Stryjewsja *et al.*, 2017), techniques for converting sketches of compartmentalised environments into labelled graphs (Mielle, Magnusson and Lilienthal, 2019), and even techniques to use sketches as a basis for communicating spatial knowledge to robotic systems using ‘qualitative landmark states’ (Chronis and Skubic, 2004). It has also been found that the order of recall can be used as a basis for inferring hierarchy in the cognitive representation (Taylor and Tversky, 1992) and (Tversky, 2002), and expanding on experimental insights related to hierarchy, the famous result that perceived region hierarchies can form the basis of the systematic distortions observed in our spatial knowledge (Stevens and Coupe, 1978).

Both in the illustration in Figure 20 discussed previously, and in Figure 22, it can be seen that the central task is to identify and disambiguate elements in the sketches. In terms of an overall characterisation of different approaches to sketch mapping however, while there is naturally a degree of subjectivity, the categorisation presented in (Appleyard, 1970) can be seen as a guiding framework, for example through structural patterns such as the ‘chain’ and the ‘mosaic’.

Finally, while image-schema theory was developed through research in cognitive linguistics, clearly the notion of ‘extracting’ image-schemata through analysis can

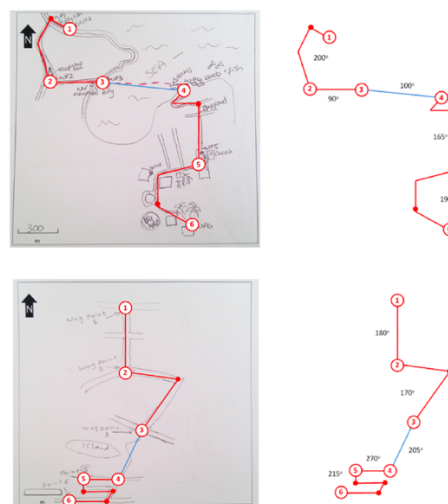


Figure 22 Sketch maps and digitised schematics (Field, O’Brien and Beale, 2011)

be extended to graphical data also, with the inherent flexibility of image-schemata meaning that cognitive topology expressed through textual and graphical data can be compared, and an overall picture of cognitive representation developed from multiple data sources.

### 5.3.3 OBJECTIVE DATA COLLECTION AND ANALYSIS

Objective data is data that is generated through the measurement of the change in a variable over time. In the study of urban navigation there are two sources of objective data – the participant themselves, and the device used for accessing information. Data that can be directly observed in terms of the participant relate to the position and location of the body, and the progression of the body through space. For example the location of the participant when they expressed something verbally, the direction in which they were or were not looking, their speed of movement etc. Physical actions can also be inferred through data from the device. For example by capturing the device's screen the timing, volume and type of interactions with the device can be known. Additionally the device is able to record the path of movement over the course of the journey, for example see (Kraak, 2003).

During research design a mobile eye-tracking device was also tested, however the trade-off in terms of the time required to conduct a meaningful analysis of the data was not determined to be appropriate given the core research aim.

### 5.3.4 SAMPLE METHODOLOGY

Referring back to the scope of the research as discussed in Chapter 3, the research aim is to be understood in terms of delivering highly contextual information but with only standard contextual data as input, i.e. current location, destination and recommended route. Implicit in this scope is that the system does not know anything about the user themselves, and therefore it is also implicit that the system must produce a correct result given the research aim irrespective of who a particular user is. Sample recruitment therefore sought as broad a demographic as possible, without a focus on use case or user group beyond *a person* navigating urban environments across multiple forms of transport network. An obvious distinction that could be made is between residents and tourists, however recruiting tourists was challenging given that studies took up to two hours and had to be arranged in advance. A compromise which is argued to be satisfactory given the study objectives was the inclusion of a number of participants who had only very recently moved to the city – two of whom had only arrived that week in fact. Overall then there was a broad

mix of participants not just in terms of demographic but also in terms of familiarity with the environment.

As for gender there are conflicting results in terms of the significance of this factor relative to the development of spatial knowledge and approach to navigational tasks. For example in both (Montello *et al.*, 1999) and (Castelli, Latini Corazzini and Geminiani, 2008) clear differences were reported, whereas in (Kitchin, 1996) it was found that any differences observed were better explained by socio-cultural factors than gender *per se*. Observations made over the course of the present research strongly support the latter view, with occupation being the only discernible factor that seemed to have an impact on the nature of participant spatial knowledge. That being said, an even mix of gender was a recruitment objective across the studies.

### 5.3.5 FROM OBSERVATIONS TO MODELS

The observation and analysis of textual, graphical and objective data has been discussed, with the ultimate aim of data collection and analysis being the development of a formalised framework for modeling, selection and display. Here the methods for moving from observations to the development of classes and relations is briefly summarised. Following (Mark, 1999) with reference to (ICA, 1997), ‘entity’ describes a ‘thing in the real world’, whereas ‘object’ can be understood as a ‘thing in the digital world’. In general, throughout this thesis where the subject of discussion is the cognitive representation of an underlying entity, simply the term ‘representation’ is used, although the reader should be mindful that ‘representation’ is used to describe digital ‘objects’ also.

#### **Entities**

The focus here is to determine the entities that play a role in representations that can be shown to be stable across preferences and behavioural types, following objective 1a. Clearly if the same entities appear to play consistent roles that are meaningful to goal acquisition across a significant number of participants then this would indicate candidates for the entities that should be explicitly represented in the underlying data model.

#### **Entity groups**

Following *Entities* above, consistent forms of grouping may inform the development of the data model – for example consistent groupings that include multiple entity types. Clearly also there is the question of scale and how groupings support abstraction, and so this is another aspect of the analysis that may directly support formalisation.

### **Regions**

As discussed previously, the natural regionalisation in the perception of the city relative to the current journey context is a key focus. It is anticipated that both textual and graphical data will support the identification of region types and region structure. Following the identification of entities and entity groups following it is also anticipated that regions will be naturally identifiable through this process, as opposed to the identification of ‘naked’ regions (Casati and Varzi, 1997) that exist independently of the entities that play a role in navigation.

### **Events**

Following *Entities*, *Entity groups* and *Regions* above, the underlying aim to is to formalise event types such that selection supports information needs across the event paratomy. Here it is anticipated that the formalisation of entities, entity groupings and context-specific regions will be developed in parallel with the identification and formalisation of event types. More generally, a requirement may become apparent to develop formal representations of other *perdurant enitites* (Bittner, Donnelly and Smith, 2004) that play a role in events but are not themselves of type *event*, i.e. other identifiable entities that do not wholly exist at a single point in time but are better characterised as processes over some time interval.

### **Relations**

Of central concern is the identification of relations across entities, events and regions such that data can be modelled and returned via graph queries. It is anticipated that candidate relations will be discovered by participants explicitly communicating their perception of a relation, or by analysis of the data such that candidate relations would logically return selections that can be shown to have an empirically validated basis given participant information needs over the episode.

## **5.4 Epistemological foundations of technology research**

Finally, given the overall aim of the present research, it should be highlighted that the methodology described above should be understood as being framed within an overarching epistemological position of pragmatism. Following Pierce’s pragmatic maxim (Peirce, 1878), in essence this means that the guiding principle of inquiry is the anticipated ‘practical consequence’ of the knowledge that results from the activity. This practical consequence can be understood in terms of ‘rational conduct’ that follows from the knowledge. An implication of this for the methodological

approach applied in the present research is that the structure of inquiry is *iterative*, meaning the state of knowledge is improved over time through experience. This is manifest through the division of research into 7 studies, with hypotheses being updated and refined over time. More specifically however, this epistemological position frames the evaluation of knowledge at each step of the inquiry in terms of the practical consequence of the knowledge given the central research aim.

Here it can be seen that empirical observation aims to guide the development of formalised representations along with the guiding principle of practical consequence as described. The guiding principle can be seen in terms of the aim to develop knowledge around the automated selection and display of navigational information as opposed to knowledge about cognition itself. While the aim is to develop a framework that supports representations which better align with the natural information needs of travellers, it is not to develop a cognitively plausible model per se. Some approaches seek biological plausibility in the formalisation of navigational behaviours (Milford and Schulz, 2014), and others explicitly seek cognitive plausibility, through for example explicit models of cue integration given sensory data (Madl *et al.*, 2018). The present research however should be understood as being constrained by the practical goal of generating output maps views, and so cognitive plausibility in the structure and content of formal representations is an objective only so far as it aligns with the appropriate levels of abstraction given the need to support spatial goal acquisition with external informational resources. Overall then, the aim of formalisation should be understood as being predictive rather than explanatory (Oren and Yilmaz, 2013), with simulation being the process of determining a prediction about the information required over the event partonomy given a database of features, as opposed to a cognitively plausible simulation based on sensory data that is analogous to the experience of a cognitive agent in a ‘real’ experiential context



## 6 Technical overview

In this short chapter, explanation is provided around the approach to the development of prototype research stimuli that may act as a reference when the reader is engaging with the presentation of subsequent content. It is not necessary to fully grasp the explanation of the development of prototype research stimuli, but it is included here for completeness. That being said, it should be emphasised that in subsequent visualisations of selections, view curves and persistence of feature representations across view space, the simplified geometry of these visualisations is not equal to the actual geometry of view space that underlies the implementation of interactive prototypes. To some readers this distinction may not be meaningful, but for others it could be a source of confusion, therefore explanation is provided.

The aim was to develop a highly flexible prototyping framework so there would be no constraints on the design and implementation of each stimulus presented to research participants. Because the scope of research initially included distortions of metric space, this flexibility had to extend beyond the representation of feature attributes to the underlying spatial substrate.

### 6.1.1 VECTOR TILE OVERVIEW

The tile server is the fundamental basis for sending cartographic data to a mobile device. The usual approach to cartographic representation is for the earth to be divided into square subsets with each square being a 256 x 256 px resolution. At zoom level 0 the entire surface of the earth is represented on a single tile, and at each proceeding zoom level, the length of the area represented doubles along each axis – i.e. the area on a single tile at zoom level 0 decomposes into a representation across 4 tiles at zoom level 1, and so on. If we take the equatorial radius to be 6,378.137 km (WGS-84), then the length of the equator is 40075.016686 km, and a single tile at zoom level 0 contains pixels that each represent a length of 156543.03m, i.e.

$$6378137.0 \cdot 2 \cdot \left\{ \frac{\pi}{256} \right\} \approx 156543.03.$$

Given this numerical model of spatial scale, in an intra-city context many of the lower resolution zoom levels are not relevant. In the case of journey across Edinburgh it is argued that given the geographical extent of the city, zoom level 14 is really the lowest resolution zoom level that is relevant, or approximately 1:35,000. If this defines our qualitative ‘lower bound’ on spatial scale, the upper bound is the highest resolution that is practical to render given the resulting size of cartographic objects on the screen or zoom level 20 (approximately 1:500).

Prior to 2013, Google Maps served raster tiles, but since then the standard for digital mapping has become the vector tile. In other words, although data is still divided into square subsets, the ‘tile’ is actually a set of vector geometries and attribute information. Highly contextual mapping that reflects the unique situation of the user would not be practical if a bespoke set of raster tiles had to be automatically generated, but given that the actual geometry and attribute data is sent to the client, this opens up the possibility to change or at least augment the data that is sent, and also it affords the possibility to render the data on the client in a way that is adaptive to the user’s particular context. If index-free adjacency is the foundational technology that makes this research possible in terms of data modeling and querying, it is the vector tile paradigm that leads to the research being possible on the front-end.

In practice, an implementation framework must include a unifying basis for dealing with a number of different coordinate systems. While the classical problem of cartography and coordinate reference systems is representing the three-dimensional surface of the earth in the two-dimensional plane, in practice the problem requires further steps because a) geographic coordinates must then somehow be expressed relative to a specific selection, and b) coordinates must ultimately be interpreted as locations in an array of pixel locations that correspond to the viewing window to be rendered to the screen.

In the case of Google Maps for example, there are four interrelated coordinate systems: geographic coordinates, world coordinates, pixel coordinates and tile coordinates. In the Google Maps approach world coordinates are based on a fractional  $x,y$  position relative to a single tile at zoom level 0; pixel coordinates express the pixel position at a zoom level by multiplying world coordinates by a scale factor, and tile coordinates provide a numerical basis for requesting tiles at a given zoom level (based on pixel coordinates and tile size).

### 6.1.2 RESEARCH STIMULUS PROTOTYPING FRAMEWORK

The following briefly summaries the technical basis of rendering geometries and attributes to a device, modelling a metric interactive view space, and representing a non-metric view space that is also interactive. No third-party mapping framework was used to allow for complete flexibility. As such, all feature classes were built from scratch using primitive classes implemented in JavaScript such as lines, points and vectors.

Feature geometries and attributes were either taken from features in OS MasterMap, OS Highways, OS Vector Map Local (Ordnance Survey, 2019), Open Street Map (OSM, 2019), created in Qgis (Qgis, 2019), or created using python scripts during geojson processing. Ground truth for point features relating to entities adjacent or near to the path of test routes was established and a database of these features was created manually.

The geographic extent of the city was treated as the effective ‘tile 0’, and pre-processing on the raw geojson geometries converted the Eastings and Northings to screen coordinates within the context of this effective ‘tile 0’ extent.

The JavaScript location API was used to get the latitude/ longitude data from the client’s browser. Code from Chris Veness’s geodesy library was adapted for various operations such as processing location data so that all vector geometries were reconciled to the same coordinate space. Specifically, code was adapted from the modules *dms*, *latlon-ellipsoidal* and *osgridref* (Chris Veness, [online]).

### 6.1.3 IMPLEMENTATION OF INTERACTIVE VIEW SPACE

The essential issue in implementation is that a model must be developed that accounts for the non-linear relationship between map scale and the average optical flow given a pan interaction. To maintain smooth, straight trajectories in the plane, requires that in fact, numerically speaking, paths through view space are curved (Van Wijk and Nuij, 2003), with the geometry of the curve and the ‘speed’ of movement being key factors. Conceptually, view space is the space that defines the constraints and overall behaviour of user pan and zoom interactions, i.e. it defines the overall set of possible Gestalts given the underlying data.




Figure 23 Normalised image space in a mobile device context with  $c$  denoting the image space centre.

The map as an interactive multiscale information space is represented by a parametric model, in which a relation is parameterised between the path of the centre of image space  $\mathbf{c}$ , and the width of image space over the arc length of the path. This is necessary for smooth multiscale interactions with consistent speed and straight trajectories in  $\mathbb{R}^2$ . The parametric representation is adapted from (Wijk and Nuij, 2004) with the fundamental difference being the use case of free user-driven interactions in parameterised view space, as opposed to multiscale animations with pre-determined paths. In other words the ‘destination’ of the path is always the current image space centre, as opposed to a centre at a future time instant.

A square subset  $A_j$  of world space  $W_i$  in  $\mathbb{R}^2$  is projected to a unit square image space  $I_0$ , where  $I_0$  is equivalent to a default zoom level (around zoom level 14, which qualitatively can be thought of as ‘city space’  $\sim 1: 36,000$ ). In our case in fact the extent of the overall journey region defines the size of this image with respect to the device’s screen. A ‘zoom’ interaction triggers a transition to a new image space, controlled by a change in the width  $w_i$ , as represented in normalised image space, of a Euclidean distance in  $W_i$ . The ratio of  $w_i$  to  $w_{i+1}$  is equivalent to the map scale.

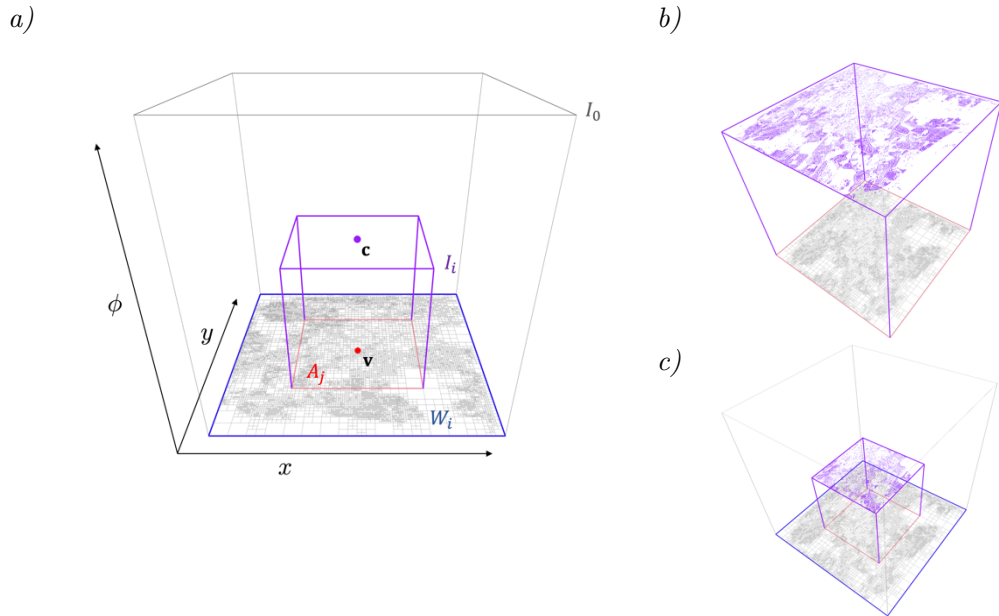


Figure 24 a) key components of multi-scale view space, including the world space  $W_i$ , a square subset of world space for the current view  $A_j$ , the geographic location of the centre of the current view  $v$ , the centre  $c$  of the image space for the current view  $I_i$ , and the effective ‘tile 0’ image space  $I_0$ ; b) example selection where the current view corresponds with image space  $I_0$ ; c) example selection after a ‘zoom’ interaction to a view corresponding with image space  $I_i$ . Axis  $x$  = easting,  $y$  = northing,  $\phi$  = map scale.

In practice this means that for each vector  $\mathbf{x}$  in a feature geometry located in  $W_i$ ,  $\mathbf{x}.x$  and  $\mathbf{x}.y$  are projected into image space based on the current view centre and current pan offset. They are transformed based on the new image space width, and then rendered to the screen using the inverse projection of this new vector.

The complete implementation is not central to the arguments presented in this thesis, but a final point for completeness is that the parameterisation uses a value of 1.565 for the metric  $\rho$ , and a value of 0.5 for the speed variable. These values were simply arrived at through experimentation given the extent of the study routes and the goal of a smooth visual result.

#### 6.1.4 VIEW EVENTS & VIEW CURVES

An important aspect of data collection was the observation of interactions with information through the mobile phone device. Research participants were instructed to record any interactions with the device using a screen capture application. These recordings were then used as the basis for analysing view event onsets by using the

timecode and matching it to the time on the GPX trace that was recorded over the duration of each study. In other words, on each occasion that a participant opened the device to check information (the beginning of a view event), the time at which a screen recording was initiated was matched to the location of the participant at that time, as captured through the device's location. An issue however is that GPX

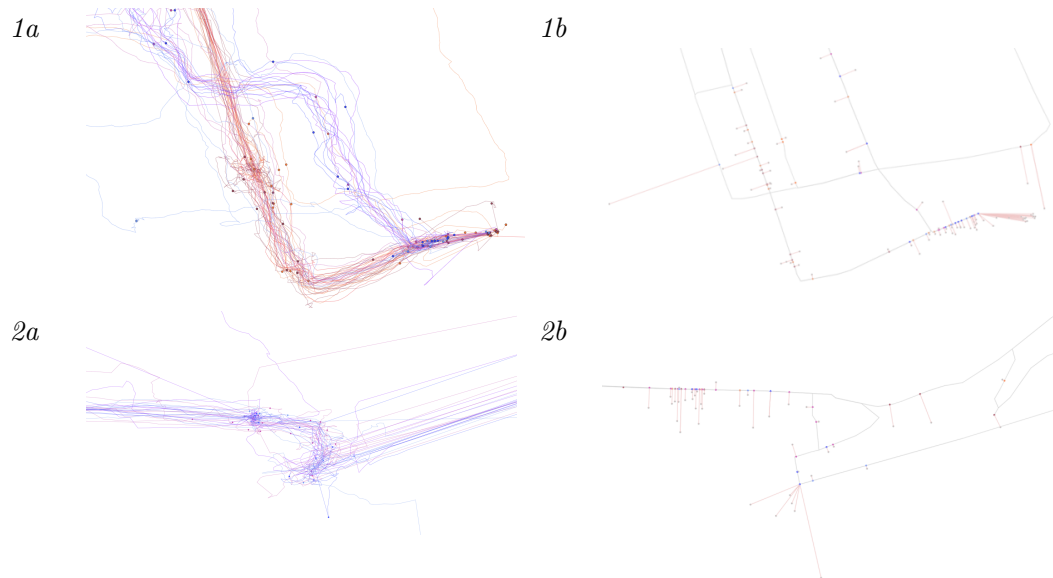


Figure 25 Locations of view event onsets before and after straight-line shortest distance processing. 1b shows locations after processing for a region including the journey origin of two study routes, with 1a showing the raw GPX traces for the same region across groups for study route 1 (red hues) and for study route 2 (purple hues). 2b shows view event onsets around the Edinburgh Haymarket train station for both study routes, with 2a showing raw GPX traces for participants across study route 2 only (purple hues).




Figure 26 View event onset locations after shortest distance processing for study groups 2 and 3 (top - study route 1), and for study groups 5, 6 and 7 (bottom - study route 2), visualised over the underlying paths of traversal.

traces are only approximate, and so to visualise the start of view events directly leads to visualisations that are hard to interpret. Where view event onsets are visualised then, the location of these point features corresponds with the intersection of the GPX location for the given time with the path of the routing network that was traversed during the study. i.e. shortest distance, as shown in Figure 25).

In addition to providing the data required to analyse view event onsets, the screen recordings also provided the data from which the scale of map views accessed over the course of view events could be derived, i.e. the equivalent width of image space.

By observing the centre of view space in screen recordings and the extent of the square subset of world space for each centre, the path of attention through multiscale space could be approximated. A view curve then is an approximation of the path of attention through interactive multiscale space that can be visualised in a Euclidean three-dimensional graphics context, as in Figure 27.




Figure 27 Example complex view curve  $Q_i$  over a time interval, projected into Euclidean space given world space  $W_i$ . Axis  $x$  = easting,  $y$  = northing,  $\phi$  = map scale.

The geometry of interactive view space as is not accurately modelled using a Euclidean space (van Wijk and Nuij suggest that the parameterisation of pan and

zoom interactions is in fact characterised by the hyperbolic geometry (Wijk and Nuij, 2004)). With this in mind, the Euclidean representation used in a number of figures in this thesis should be highlighted as a simplification for the purposes of visualisation as opposed to an accurate representation of the geometry of view curves.

**View event definition:**

A study participant interacts with information at one or more locations as they progress along the journey. The location of a view event onset, i.e. the person’s location as they begin an interaction with the mobile device, is modelled as a vector  $\mathbf{v}$ . At each  $\mathbf{v}$ , cartographic information is determined by a three-dimensional vector  $\mathbf{c}$  that denotes the centre of image space and the numerical scale, where a function maps features to the given  $\mathbf{c}$ . A view event  $V$  for a given time instant  $t$  can therefore be expressed as the tuple:

$$V = \langle \mathbf{v}, \mathbf{c}, t \rangle. \tag{6.1}$$

**View curve definition:**

A time interval groups view events based on the start and end of a screen interaction, and so the ordered sequence of view events over the interval models a view curve, expressed as the sequence  $Q$ , or an approximation to the path of attention through view space, contextualised by the viewer’s physical geographic location as represented by  $\mathbf{v}$ :

$$Q = (\mathbf{c}_i \dots \mathbf{c}_n \mid t_i \leq t \leq t_n). \tag{6.2}$$

The graphical depiction of the path of a curve assumes a straight trajectory in Euclidean space between each  $\mathbf{c}$ , which is the basis for the simplified representation used in view curve figures throughout.

6.1.5 PERSISTENCE VECTORS

In Chapter 10 a type of object named ‘persistence vectors’ will be discussed. These objects represent the persistence of a feature through multi-scale view space, i.e. the scale at which the feature materialises, and the range of scales over which it persists in the current context. It should be highlighted here that the visualisation of persistence vectors follows the same visualisation logic as discussed in 6.1.3. In other words the trajectory of a persistence vector through view space should not be

interpreted literally, but should be understood as a simplification for the purposes of visualising feature selections in Euclidean space. An illustration is provided in Figure 28 that shows the persistence of a road network between two scales in view space.

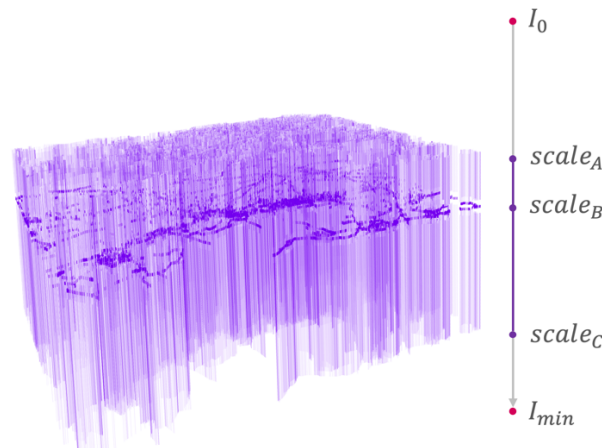


Figure 28 A road network that persists in view space from views at numerical scale A to views at scale C, with a disambiguation of features at scale B for illustration.

This way of graphically expressing the behaviour of objects that represent geospatial features is similar to the conceptualisation discussed by van Oosterom and Meijers of the generalisation of two-dimensional features as being like an ‘extrusion’ into the third-dimension (van Oosterom and Meijers, 2014).

### 6.1.6 FRAGMENTED VIEW SPACE

In Chapter 4, it was highlighted that the natural cognitive regionalisation of space leads to distortions formed by regions of higher detail that exhibit metric-like qualities, loosely bound in global structures that relax absolute distance and direction in support of flexible and efficient representations. Participants in Study Group 3 completed a navigation exercise using a research stimulus that incorporated these regionalised metric distortions into the presentation of cartographic information. Results from this study will be discussed subsequently, however a summary of the implementation is provided here for reference.

Event segmentation leads to the regionalisation of space based on task-context and the perception of event boundaries. In terms of a navigational episode that includes transitions between different forms of transport mode, these transitions serve as the basis of high-level event segmentation, conforming in fact to every known mechanism for the natural grouping of event sequences. The approach was to

approximate this segmentation such that cartographic features could be displayed in a way that was predicted to reflect an *a posteriori* representation of the journey, including metric distortions and regions of higher schematisation. The implementation determined this regionalisation using key nodes in the traversed network. Following the well documented tendency for route segmentation to be driven by the origin, destination, actions and intermediary decision points (Tenbrink, 2012), regionalisation was determined by deriving region centres using primary nodes, then deriving points of intersection between the path of the traversed route and context-specific region boundaries. This approach meant that regions

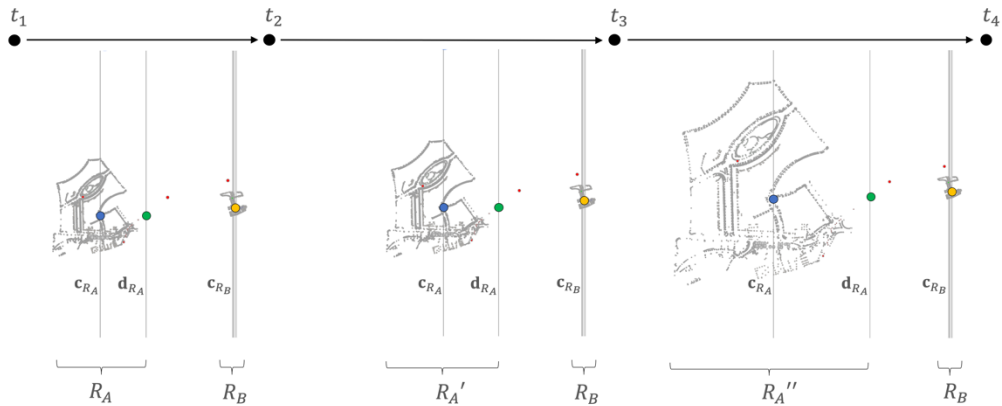


Figure 29 Differential scaling over an extent based on the disambiguation of focal regions and the deformation of metric relations. Vertical lines mark the Easting of the location of the primary nodes for the journey context. The result of zoom interactions are shown over three time intervals from  $t_1$  to  $t_4$ , with scale increasing for region  $R_A$  but not for region  $R_B$ , i.e. the distance between  $c_{R_A}$  and  $d_{R_A}$  increases, but the distance between  $d_{R_A}$  and  $c_{R_B}$  does not.

could be differentiated such that zoom interactions could affect scaling in a non-uniform manner over the extent, with the intersection of the traversed route with context-specific regions serving as control points in the deformation of metric relations. The approach also meant that content selection could be region-specific, with the method of transport dictating the level of schematisation and the level of detail for each region, following (Mackaness, Quigley and Tanasescu, 2011).

In Figure 29, zoom interactions cause the scale to increase for region  $R_A$  but not region  $R_B$ . This is achieved by deriving region centroids  $c_{R_A}$  (blue) and  $c_{R_B}$  (orange) using the location of the primary routing nodes that approximate the boundaries of the event sequence. To smooth the display of the regionalisation given the need to render metric feature geometries in focal regions, the intersection of the derived boundary of each region with the path of the traversed route is used as a control point (Godfrey and Mackaness, 2018). As region  $R_A$  is focal given the

interactions that occur through  $t_1$  to  $t_4$ , the relevant control point for the current deformation is point  $\mathbf{d}_{R_A}$  (green), as indicated in Figure 29.

Reflecting on the mathematical basis of view events discussed previously, the approach described above can be seen in terms of the fragmentation of view space at region boundaries. In this context, fragmentation means the trajectory of a view curve does not occur in a consistent vector space, with the abrupt change at an intersection point inspiring the term ‘fragmentation’ as opposed to ‘warping’ or some other term implying a more gradual change.



# 7 Situated navigation studies

A summary of the overall research design, activities, and the forms of data collected through participant observation are presented. Following the methodology, the essential logic of the approach was that if we are to learn how to effectively reflect a person's context in the selection and display of information, inquiry should be conducted in that same context, or in as close an approximation as is feasible. While lab-based studies are required when there is a need to strictly isolate variables, the ambition of the research called for a more naturalistic approach that would allow for the observation of people as they actively engaged in navigation tasks. There would be no way, for example, to effectively study the natural sequence of information interactions over the course of a journey without the participant actually completing the journey. In a lab setting you would rely on self-report which would be unlikely to offer more than a superficial level of insight given the need to observe the relationship between the participant, the environment and navigational information delivered through the mobile device.

## 7.1 Research design

### 7.1.1 NAVIGATION EXERCISES

Studies were divided across three high-level types of navigation exercise, the first with no participant exposure to information beyond environmental cues, the second using the established form of navigational information (Google Maps), the third using research stimuli developed as a means of investigating hypotheses in situated contexts. Research stimuli were developed as a means of validating and refining hypotheses around the selection and display of information, and around the overall development of a cognitive representation of the episode. These stimuli were presented either as interactive web-apps that supported pan and zoom, or as a 'click-through prototype' in which interactivity was limited to being able to progress through a predetermined series of views.

In all cases each study involved a single participant (with the exception of 1 baby). Participant's provided answers to a series of background questions via email before the exercise, and then a short discussion in person prior to each exercise meant these answers could be queried and elaborated on. Questions included general background such as travel habits, habits in terms of forms of information regularly accessed, level of confidence with urban navigation, knowledge of the city, work background and so on.

After providing answers to additional background questions in person, the participant was briefed on the exercise and the overall format of the study, with each study taking 1-1.5 hours for the first route, and 1.5-2 hours for the second route. In the case of participants that were being followed as opposed to following, time also had to be set aside for setting up a wireless microphone that was linked to the camera via a receiver. These participants were also briefed on using the screen capture application and had to demonstrate an understanding around how and when to record the screen of the testing device during the exercise.



Figure 30 Example stills from navigation exercise video recordings.

In the case of participant's that were able to use additional information (Google Maps, prototype stimulus, or other additional information), they were not briefed on the follow-up exercises in advance of the journey, but were simply told that there would be additional exercises toward the end of the session. In the case of participants who were to complete the test route by following the researcher, a risk was that they would not be adequately engaged with the journey given they had no decisions to make and so no inherent motivation to consciously attend to the experience. In these cases then, participants were briefed that while they should be

relaxed during the exercise, they also needed to bear in mind that they would need to absorb the experience sufficiently given that they would need to complete follow-up exercises including providing directions.

In the case of participants using Google Maps, they were provided with a destination verbally, and then had to talk through their thought process as they planned the journey. This verbal data was recorded using audio, with screen interactions recorded using a screen capture application on the device. Controlling the exact route was not possible as the Google app suggests different routes given the current options. Within this practical backdrop, participants were free to choose their route from the available options, however for the first study route it was a constraint that the journey had to involve a bus, and for the second route it was a constraint that the journey had to involve both a bus and a train. These constraints served to ensure study route journeys were fair representations of the level of complexity usually encountered while travelling across major urban areas. Once the participant was satisfied with the route, they were set-up with the wireless microphone. The researcher turned on a GPS tracker and followed the participant with the video camera as the participant talked through their thought process. The trade-off between practicality and the benefit of video recording on public transport meant that during phases of the journey in which the participant transferred to a bus or train, recording switched to audio only.

<b>Exercise type</b>	<b>Summary</b>
<i>Following</i>	No exposure to information on a mobile device – participant experiences the journey simply by following the researcher
<i>Google Maps</i>	Participant provided with destination verbally and uses Google Maps, plus any other information/ application they desire
<i>Research stimulus</i>	Participant provided with a research stimulus in which the journey destination and route are preselected (i.e. no alternatives)

*Table 1: Study exercise types*

In the case of participants who were provided with a research stimulus, the route to the final destination was predetermined, however clearly the form of information itself was unfamiliar, and so some time had to be spent ensuring the participant

reached a satisfactory level of confidence with their understanding of the journey. As with all route planning across both Google Maps participants and research stimulus participants, the participant had to verbalise their thought process. In the case of research stimulus participants however, extra time was naturally spent during initial exposure to the information as they formed an understanding as to the nature of the information and the form of its display.

In all cases, immediately after the destination had been reached, the researcher and participant began the follow-up exercises. A key factor in choosing test routes was the proximity of the final destination to a location where these follow-up exercises could be completed. This began with a short reflection on the experience followed by two exercises. The ordering of exercises was alternated between participants to address the issue of biasing the data in the second exercise given the experience of completing the first. These were a textual directions exercise and sketching exercise.

### 7.1.2 STUDY GROUPS AND PARTICIPANTS

64 participants took part across 7 studies.

An even number of participants were recruited to each study group with a view to a 50/50 gender split. This target was achieved in most cases, with the overall split across the 7 studies being 58% female and 42% male. The sample was skewed toward participants in their late 20s, however a broad mix of ages participated, from 19 to 75. Effort was also made to recruit a broad demographic including non-students. Students of geography, cartography and GIS were not accepted into the sample. Participants were provided with a £25 cash stipend for participation. Participants had to be over 18 years of age, fluent in English and 'smart phone' users.

Each study group will be referred to by the convention 'GX', i.e. G1 for group 1 and so on. Each time data relating to an individual participant is referred to, the participant will be referred to by group number and a participant number for that group, i.e. G1\_1 to G1\_12 for the first study group.

The research took place across two study routes which will be referred to as study route 1 – 'SR1', and study route 2 – 'SR2'. The three prototype research stimuli will be referred to as 'S1', 'S2' and 'S3'. This will be discussed further subsequently.

Group	Total participants	Sketch first participants	Directions first participants	Study route	Navigation type
G1	12	6	6	SR1	Following
G2	12	6	6	SR1	Google Maps
G3	12	6	6	SR1	S1
G4	8	4	4	SR2	Following
G5	8	4	4	SR2	Google Maps
G6	6	3	3	SR2	S2
G7	6	3	3	SR2	S3

Table 2: Study groups summary information

### 7.1.3 EMPIRICAL STUDIES OVERVIEW

While the routes taken during navigation studies are simply referred to in terms of study routes 1 and 2, in reality for groups 2 and 5, participants had freedom to choose their route within the overall constraint of taking a bus for SR1, and taking a train and a bus for SR2. Additionally, some participants in other groups who were given a predetermined route did not follow the given route exactly. Study route 1 or ‘SR1’ should therefore be understood in terms of a consistent origin and destination and a journey that involved walking to a bus stop, taking a bus, then walking to the final destination. Beyond these characteristics, the path of the routes vary. Similarly study route 2 or ‘SR2’ should be understood in terms of a consistent origin and destination, and journeys involving both a train and a bus, but again, there is some variation in the precise paths of routes taken.

In Figure 31 the two study routes are shown through the path of GPX traces recorded during navigation exercises. In Figure 32 the path of the routes in terms of the road centreline network is illustrated, using OS Highways centreline data (Ordnance Survey, 2019).

Study routes were chosen based on a number of criteria related to research objectives and practicality. Transitions between different forms of transport network was a prerequisite given that the aim was to study journeys that are characteristic of the type of routes taken in large cities as discussed.

It should be noted that in both study routes, while the direction of the arrows in Figures 33 and 34 are from left to right, in fact the journey origin of both routes is to the east of the destination. At the top of Figures 35 and 36 then, the direction

of arrows indicating the journey phase are shown from right to left. The arrows shown by the bounding rectangles in both Figures 35 and 36 simply indicate the overall direction of movement over the local region.

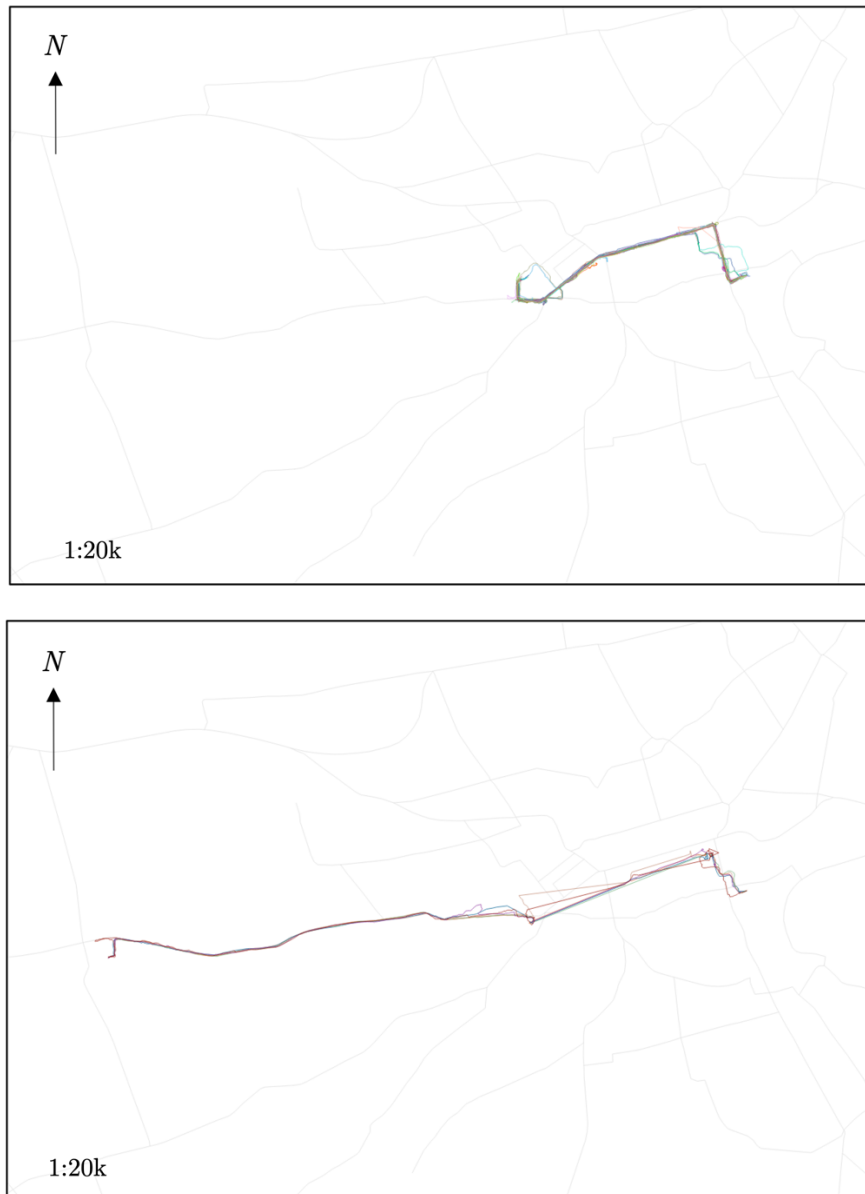


Figure 31 The City of Edinburgh study route comparison between (top) study route 1 and (bottom) study route 2, visible through participant GPX traces.

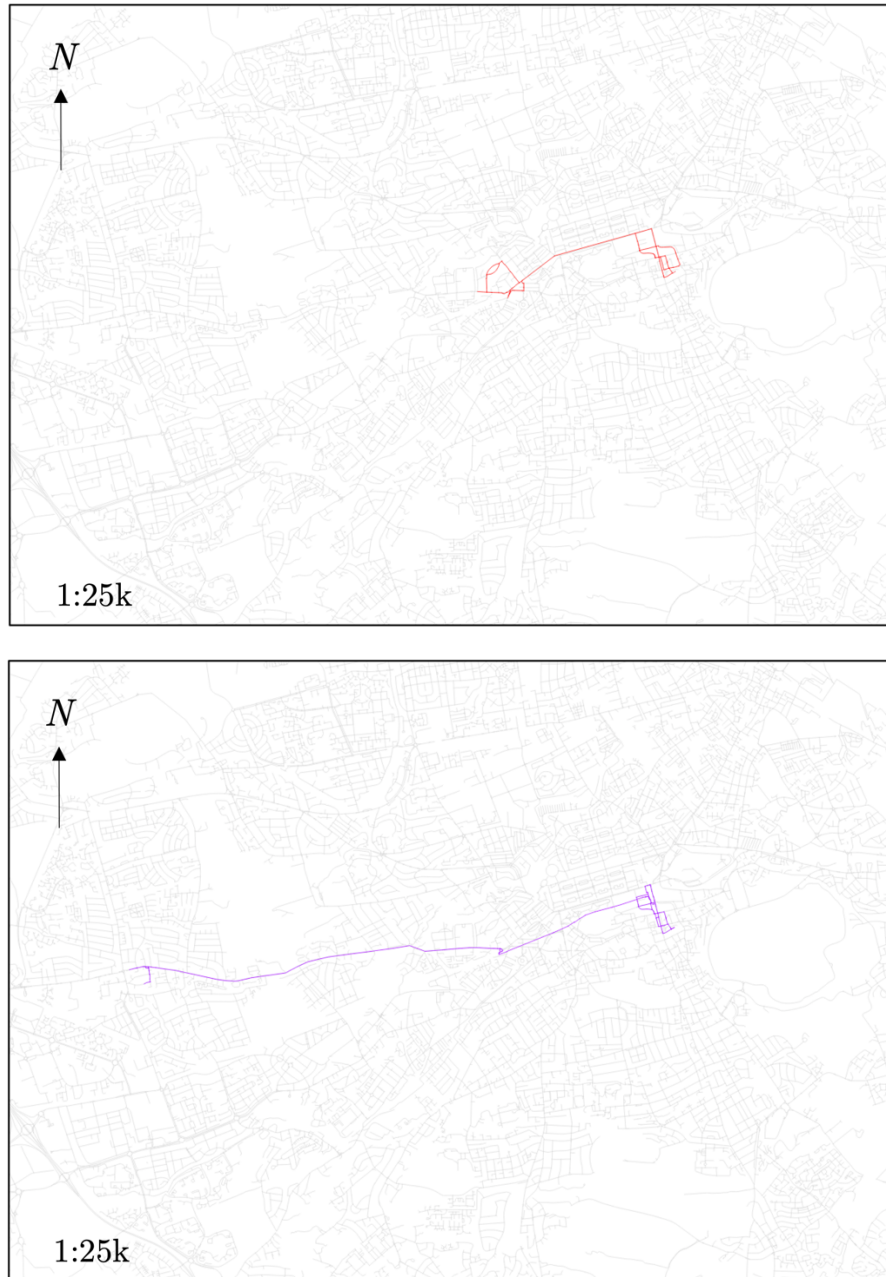


Figure 32 Study route comparison - traversed paths over OS Highways centreline data for (top) study route 1; and (bottom) study route 2.




Figure 33 Journey overviews for both study routes.

Group 2



Figure 34 Group 2 example video stills for the walking phases of study route 1, Participants G2\_1 to G2\_12.

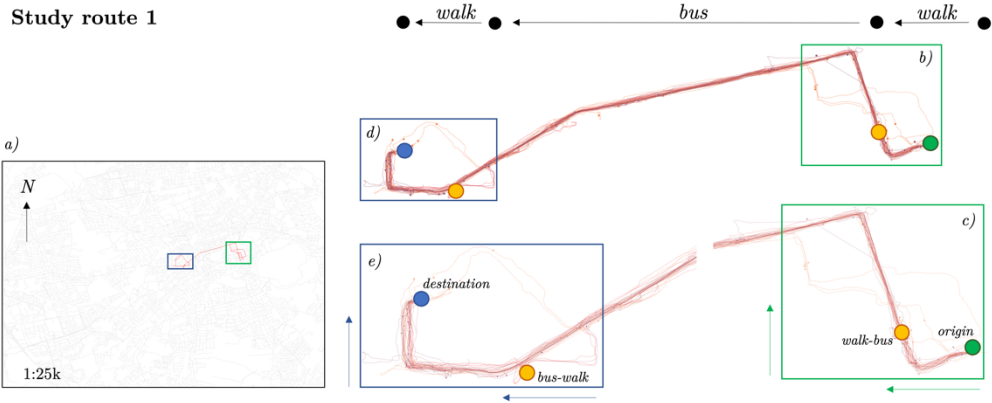


Figure 35 SR1 overview, with the origin region denoted by a green bounding rectangle and the destination region denoted by a blue bounding rectangle.

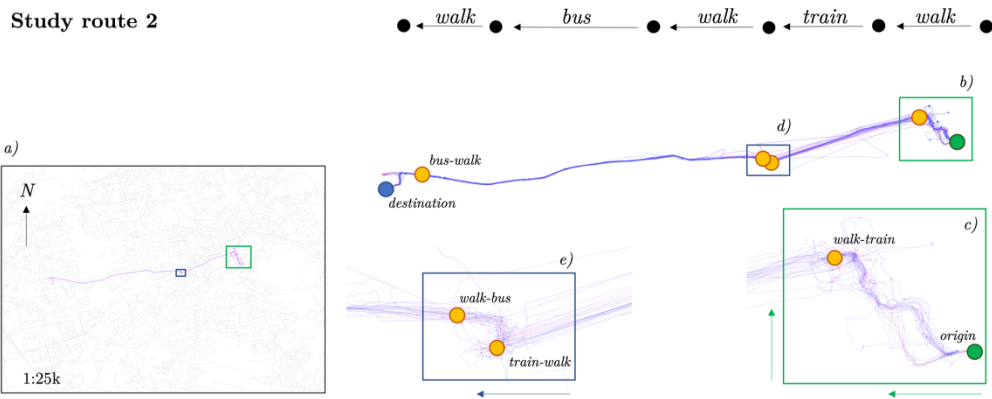


Figure 36 SR2 overview, with the origin region denoted by a green bounding rectangle and the transfer region denoted by a blue bounding rectangle.

## 7.2 Research stimuli

### 7.2.1 GOOGLE MAPS

In the studies using existing forms of navigational information the main information source was the Google Maps application. Participants were provided with the route destination verbally, and then spoke aloud as they planned the route within the

constraints discussed previously (a route including a bus for the first study route and a route including a bus and a train for the second).

As illustrated in Figure 37, the Google application is designed around a sequence of views of which the map itself only forms the basis of a subset. The application is designed such that as certain views are loaded they default to a predetermined image space (i.e. a view corresponding with a predetermined image space centre, following the explanation in the previous chapter). This issue was accounted for in the subsequent analysis of participant views.

Following previous comments around the difficulty of holding variables constant when studying ‘cognition in the wild’ (Brown, McGregor and Laurier, 2013), when studies involve a third party application this challenge can be compounded by the

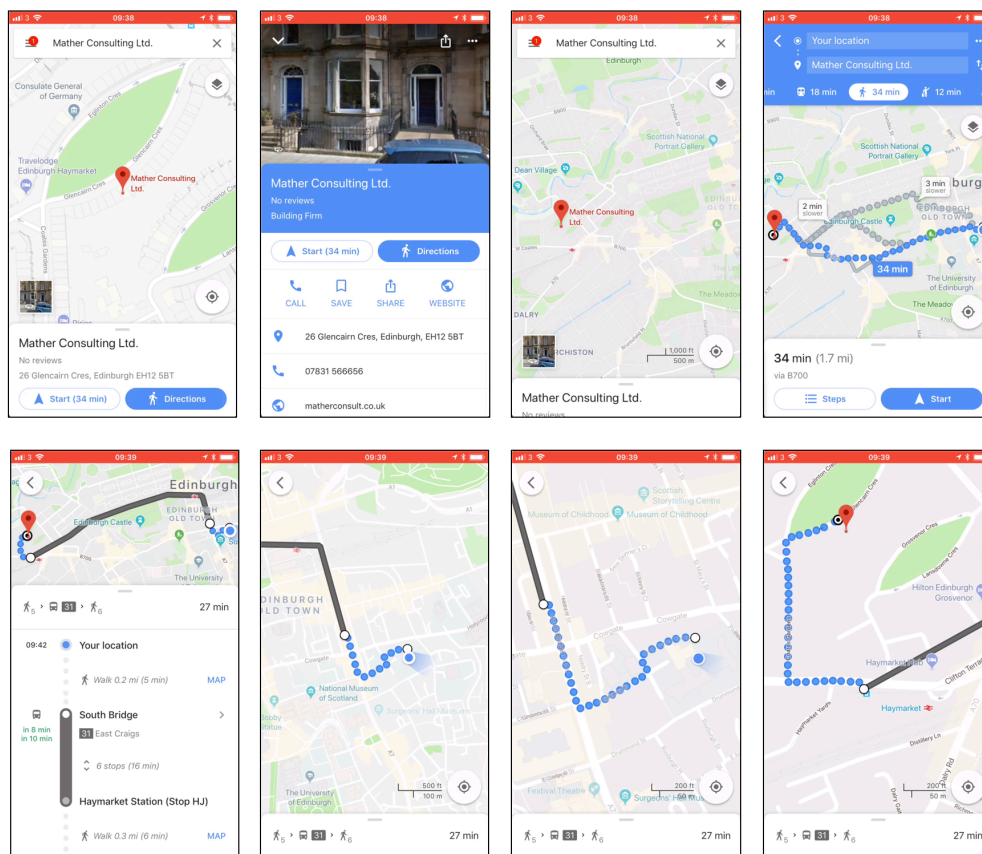


Figure 37 Google Maps views from G2 screen recordings.

on-going development and testing of the product by the third party. In Figure 38 panel 1a a view is shown that was captured by a participant during a study. Here the transition from walking to bus is located adjacent to the business ‘The Booking

Office'. Here a disambiguated polygon feature can be seen to represent this business, which is striking given that the ultimate aim of the present research is to select features that are predicted to play a direct role in navigation for the current episode. What was the reason this feature was disambiguated in this way however? Not only is the answer to this question not known, in panel 1b in the same figure, captured

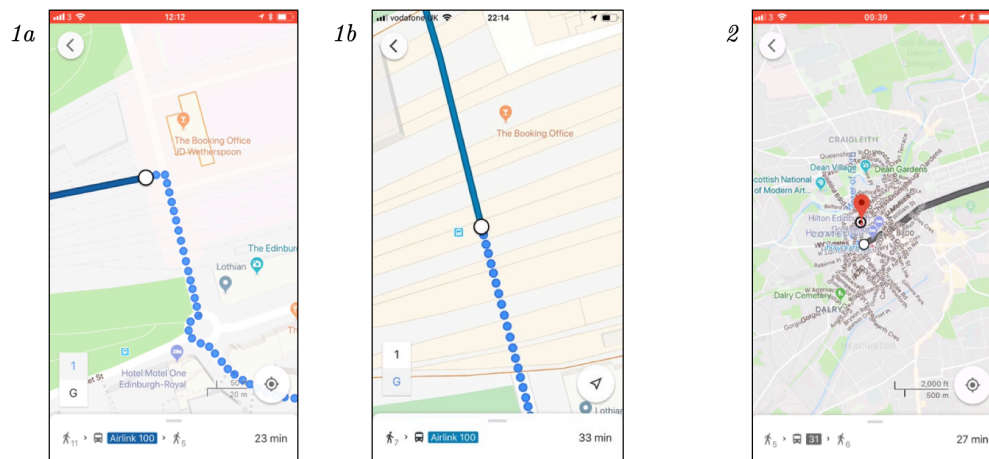


Figure 38 Google Maps: (left) disambiguation of the entity 'The Booking Office'; (middle) the same extent for the same route as the figure to the left; and (right) vector data rendering example (Google, 2018)

at a later date, the exact same image space (i.e. a map view corresponding with the same image space centre) for the same journey (taking 'Airlink 100') shows that not only has the path of the route been corrected to follow the road topography, also the train platform structures that are beneath the bridge on which the bus stop is located are emphasised. Most significantly however relative to the present research, the polygon feature is no longer disambiguated. These types of changes occurred regularly and the cartographic approach invariably changed then reverted back to an older style, implying some experiment had been completed. Whatever the reasons underlying the example shown over Figure 38 panels 1a and 1b, it should simply be highlighted that a flexible approach had to be taken with respect to both study design and data analysis.

Additionally, in Figure 38 panel 2 an example of vector data rendering is included as another example of the practical issues faced when investigating the use of data products outside of the lab. Google Maps was chosen as the main third party information source in the studies as it seemed to be the primary source of navigational information for the majority of people in Edinburgh. Aims of the studies involving third party applications were to support the overall objectives in terms of determining events and information needs, but also to act as a point of comparison with participants who completed navigation exercises using prototype

research stimuli. The aim then was not to study the Google application per se, and participants were free to access any information they needed. In almost all cases,

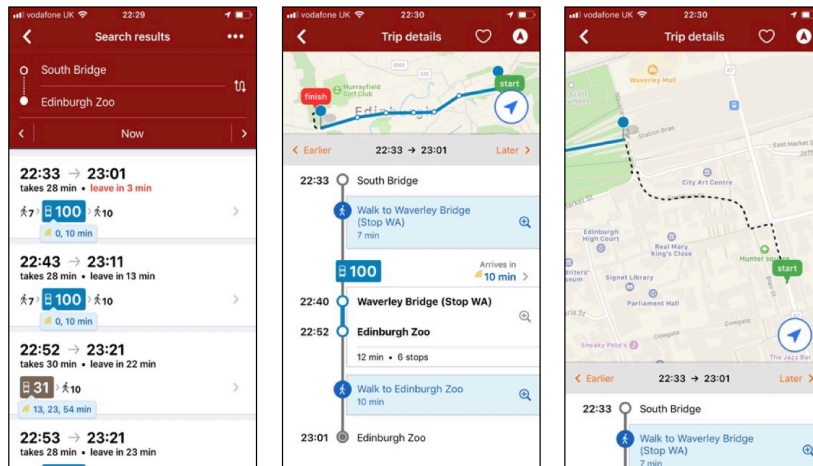


Figure 39 Lothian Buses application (Transport for Edinburgh, 2018).

this meant a switch to the Lothian Buses application which is the application provided by one of the city's main transport operators. The overall design of this application follows the approach of Google Maps, however participants expected that the Lothian Buses application would provide more accurate information about bus times.

## 7.2.2 PROTOTYPE RESEARCH STIMULI OVERVIEW

An overview of the high-level specification guiding the design of each prototype research stimulus is provided.

Prototype research stimulus	Study group	Summary
S1	G3	Structured around expected primary anchor nodes. Focal ‘origin’ and ‘destination’ regions connected by a schematised bus phase with deformation of the plane to smooth transitions between local and global views controlled by a slider interaction. Highly route-based cartographic information relative to the path of the current journey.
S2	G6	Highly route-based cartographic information but with a limited number of additional low-detail geographic features over the extent. Transitions between indoor and outdoor environments. Removal of the focal region deformation of the plane.
S3	G7	Click-through research stimulus including integration of nested textual direction information with the cartographic views.

Table 3: Prototype research stimuli summary

### Prototype stimulus S1 (G3, SR1)

The first prototype research stimulus will be referred to as “S1”, which was used by participants in G3.

In terms of research objective 1a, the study with stimulus S1 aimed to investigate the practical reality of explicitly reflecting natural cognitive representations of multi-modal journeys in the information provided. As discussed, while it is known that natural representations are loose, distorted and incomplete, what is not known is in what ways is it in fact appropriate for navigational information to explicitly reflect these characteristics, and in what ways.

So, the overall aim of data collection for G3 then can be understood in terms of validating the initial hypothesis presented in Chapter 3. To recap the hypothesis states that were an artificial agent were to traverse a virtual environment that effectively represents an underlying geography, and the agent develops spatial knowledge in a way that mimics the way a person would have developed their own cognitive representation over the course of journey completion, then the context-specific representation formed by the agent will be approximately isomorphic to the cognitive representation that would have been developed by the person. If the selection of features from an underlying database is equivalent to the second case, then displaying this subset of features to the traveller *a priori* will be approximately equivalent to the person acquiring context-specific spatial knowledge through physical route traversal. Conceptually this can be thought of in terms of a system generating a map view that reflects an *a posteriori* cognitive representation of the geography. Following the discussion in chapters 2 and 4 regarding the cognitive representations that naturally develop over a navigational episode, we can see that a map that expresses these characteristics will be determined through the following high-level specification:

- 1 a focus on route-based knowledge given the path of the route for the journey;
- 2 higher detail corresponding with increased attention at decision points and density of decision-making and the ‘grain’ of encoding being primary drivers of the level of schematisation;
- 3 an anchor-based representation with primary anchors serving as the basis of event boundaries in the upper levels of the event partonomy;
- 4 metric distortions across the extent, with higher detail local ‘charts’ corresponding with stable metric representations, and journey phases traversed on public transport leading to metric distortions between bounding regions.

Decision-points including intersections are understood to naturally subdivide action sequences and support event formation, however primary anchors are understood to be the locations of entities that serve a fundamental role in the episode (i.e. the origin and destination), as well as the locations of critical actions. In the case of SR1, at the top level these critical actions are the transfers between transport mode.

While a key aim was to investigate the integration of information supporting distinct event types in the same view following the discussion of the event horizon in chapters 2 and 3, the practical constraints of screen area meant that interactivity still had to be central to the implementation. This presented a challenge as the practical implication is that across a regionalisation of the plane, a single ‘zoom’ interaction should affect regions differently depending on the level of information abstraction required for goal acquisition across the associated event sequence. In the

traditional approach as shown in Figure 37, the ‘pinch and zoom’ interaction is the basis for moving ‘up’ or ‘down’ the scale hierarchy. One option here was to implement multiple sliders that could be used to individually control scaling of specified regions. The interaction concept that was ultimately implemented in the prototype stimulus was a single slider given that each slider itself takes up screen real estate. The benefit of a slider UI element as opposed to touch-based ‘pinch and zoom’ was the participant had a persistent visual reference, with the default slider position being centred and corresponding with a ‘zero’ zoom value (Figure 40).

In the case of SR1 there were two focal regions corresponding with the ORIGIN and DESTINATION primary anchors, and so a movement of the slider to the left caused the function that controlled destination region zoom to be active (the destination was in the east and the origin in the west). Similarly, a movement of the slider to the right caused the function that controlled the origin region zoom to be active. Each of these functions not only controlled the region-level zoom but also served to ‘pull’ a region centroid towards the centre of the screen. This implementation meant that pan interactions only needed to take place after a zoom interaction once a user was ‘inside’ the region of a local chart, and meant that a single form of interaction could be used to transition between high detail views in disparate local charts within the overall extent. In Figure 41 the result of this can be seen through the position of the slider and the size of the circles provided for illustration (blue for the

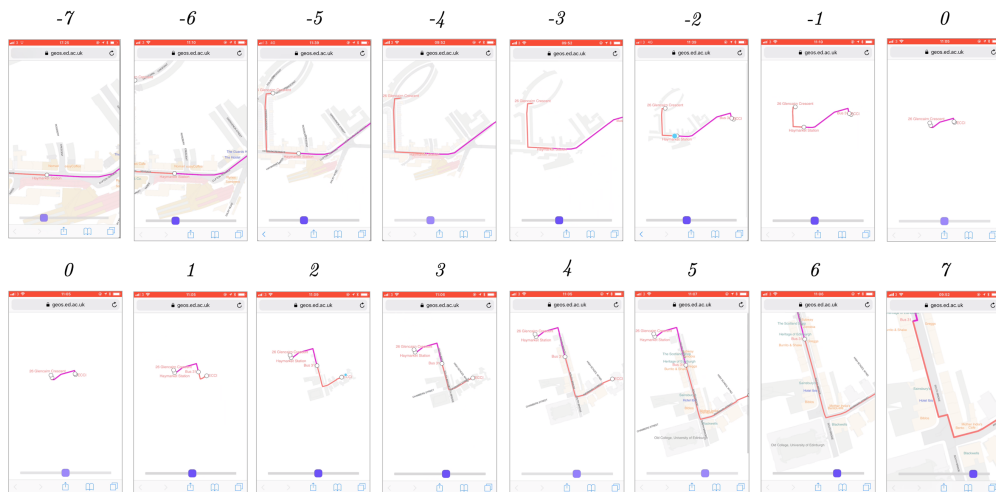


Figure 40 S1 interactions for participant G3\_10

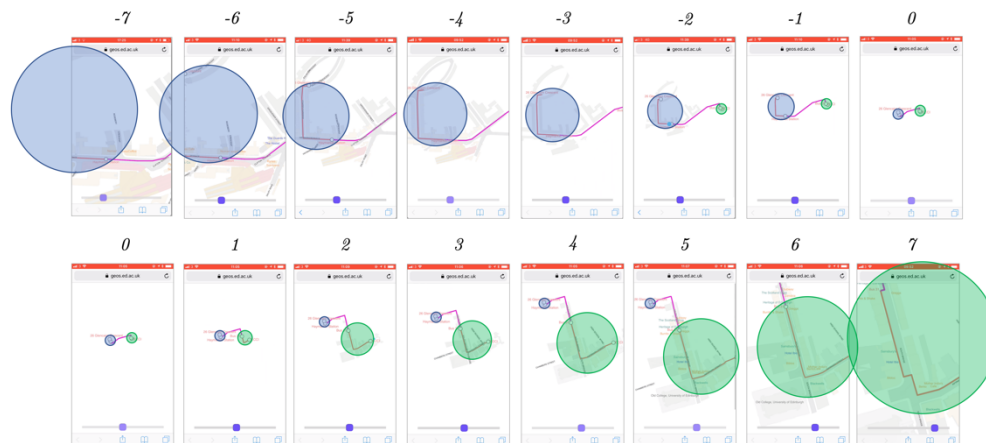


Figure 41 S1 interactions for participant G3\_10 plus left-right region scaling markers.

destination region and green for the origin region). It should be noted that in both Figure 40 and Figure 41 the numbers (-7 to 7) do not denote any numerical value in the application but are simply used to identify each view in the figure, with minus values denoting a sequence of views towards the destination and plus values denoting a sequence of views toward the origin.

As discussed in the technical overview the prototype stimulus was implemented as a JavaScript web application, accessible from any device with a web browser but optimised for mobile devices. In views ‘-2’ and ‘2’ a blue location marker is visible that is based on data from the JavaScript location API. Given the significant amount of vector geometry underpinning the dynamic regionalisation of the plane in this prototype stimulus, consistently locating this marker correctly in the view was challenging, with this issue being taken into consideration in the analysis of the data for this study.

In addition to the overall design objectives outlined above, research with S1 also sought to investigate the integration of local detail and global context in the same view. The practical realisation of this aim is illustrated in view ‘5’ in both Figure 40 and Figure 41, and is shown as a larger illustration in Figure 42 below.



Figure 42 Panel b)  $S1$  metric distortions given an origin region focus (slider interaction to right of centre), showing local detail and global context in the same view; and panel a) the metric extent of regions illustrated over participant GPX traces for the study route.

Here the tripartite regionalisation of the journey given two transport mode transfers is shown in panel 'a' overlaid on participant GPX traces for study route 1, and in panel 'b' after an origin region zoom interaction. Here the transfer point in the destination region as well as the destination region node are both visible in the same view as topographic detail in the origin region.

In terms of the features that are included in this prototype application, the high detail road and building geometries are taken from OS MasterMap (Ordnance Survey, 2019) as well as the road text and the surfaces and structures that related to the geography of the transport network, for example the bus stops. Additional point features were created manually based on ground truth.

The paths representing the route were created manually and were ordered such that low detail polyline geometries connected the origin to the first transfer node, the first transfer node to the second transfer node, and the second transfer node to the destination. As described in Chapter 6, control points served as the basis for smoothing the transition between regions after metric distortions. Zoom interactions beyond a certain value triggered a switch to higher detail polyline geometries within the focal regions.

**Prototype stimulus S2 (G6, SR2)**

The second prototype research stimulus will be referred to as “S2”, which was used by participants in G6.

The overall aim of data collection for G6 was to validate the interpretation of the findings from research across SR1 (i.e. findings from groups 1, 2 and 3). This included the findings from the research with S1 as well as the overall findings from across the three types of navigation exercise. Key differences then between S1 and S2 are the removal of the functionality that leads to metric distortions, higher fidelity topographic representations within focal regions, and additional low detail features across the geographic extent of the journey that are not adjacent to the route. In essence then, while geometric distortion was removed from the prototype research stimulus and additional features added to provide a sense of the wider geography, the stimulus can still be seen as essentially *reductive* in its design, with the selection of features being driven by the path of the current route and the sequence of decisions specific to the study route. Here ‘reductive’ is meant in the sense that detail was purposefully reduced to approximately the level expected to be exhibited in *a posteriori* participant cognitive representations, as opposed to a level which may in fact be required in terms of information needs over an episode.

Additionally, SR2 was a longer more complex journey involving 4 transfers between transport modes. This meant that S2 covered a much greater geographic extent and included 3 focal regions as opposed to 2, with the journey phase completed on foot between the end of the train journey and the beginning of the bus journey forming the additional high detail region (i.e. the additional ‘local chart’). The inclusion of a journey phase using a train as the transport mode also meant the study route included two indoor regions as the participant had to navigate inside the stations at each end of the train phase, and so S2 also includes transitions between indoor and outdoor environments.

In summary, the high-level design specification for S2 can be understood in terms of:

- 1 a focus on route-based knowledge given the path of the route for the journey (following S1);
- 2 higher detail corresponding with increased attention at decision points and density of decision-making and the ‘grain’ of encoding being primary drivers of the level of schematisation (following S1);
- 3 an anchor-based representation with primary anchors serving as the basis of event boundaries in the upper levels of the event partonomy (following S1);

- 4 feature selection based on a reductive principle but with additional features not adjacent to the path of the route providing a sense of the wider geography.

In Figure 43, panels a-d show example low detail overviews of the geography from study screen recordings. Here as few additional features were added as deemed possible given the aim of providing at least some survey-level content. These features were low detail polygons representing the city's major parks that were created for the prototype, and major roads, taken from OS Highways centreline data.



Figure 43 Low detail views from study group 6 screen recordings.

In Figure 44 panels a-d show example views from study screen recordings in which the participant increased the scale of the region bounded by the origin node and the first transfer node (inside 'Waverley Train Station'). Here selection can be seen at two levels of detail over the range of numerical scales, with OS MasterMap high detail topographic building polygons being superimposed on generalised building polygons from OS VML in the views shown in panels 'c' and 'd'.

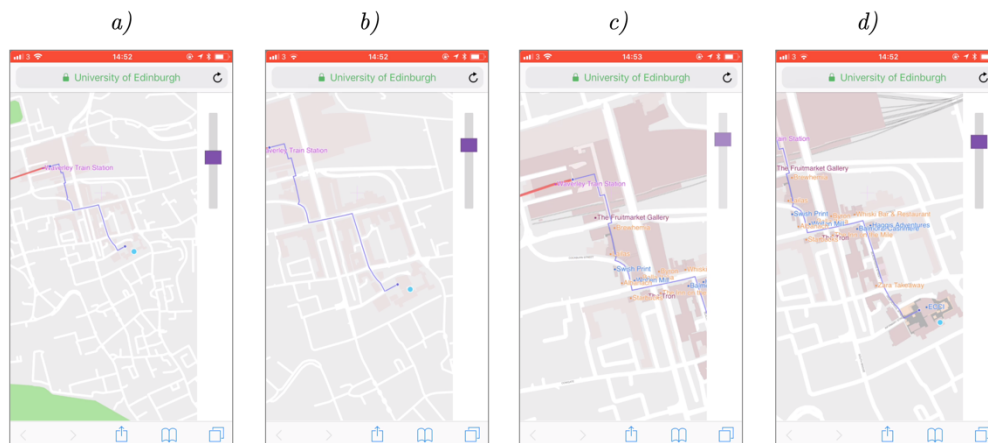


Figure 44 A series of views from study group 6 screen recordings for the study route 2 origin region.

Additionally, as illustrated in Figure 44, polygon features representing entities ‘inside’ the train station that were created for the prototype are visible in panel ‘c’ also.

In Figure 45, panels a-d show another sequence of views as a result of ‘zooming’ interactions, this time for the transfer region in which the traveller must exit the train, walk through the ‘Haymarket Train Station’, exit the train station, and locate

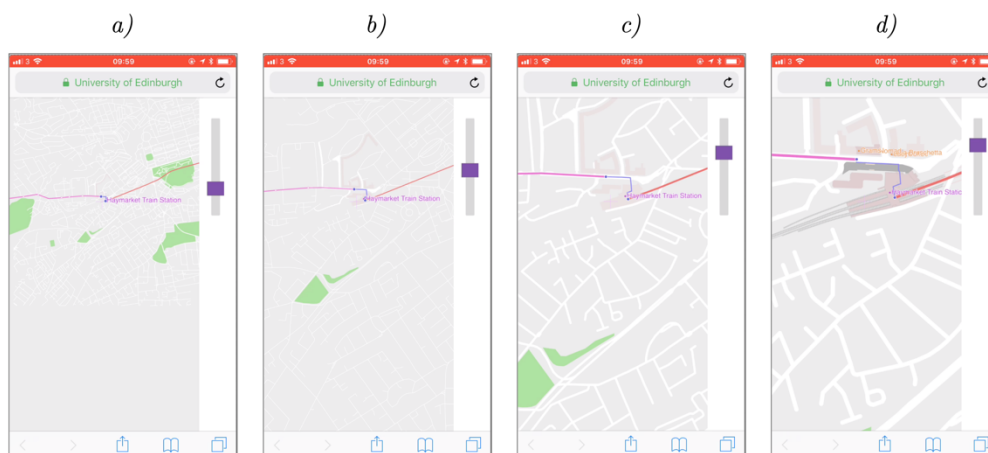


Figure 45 A series of views from study group 6 screen recordings for the study route 2 transfer region.

the correct bus stop. From the perspective of topographic detail this was an interesting study area as the bus stop was in fact located on a central ‘island’ between a main road and a tram line, and so the specific geometry of the topographic

representation was expected to play a role in locating the required bus stop, as opposed to being able to simply locate the bus stop relative to intersections that could be understood by way of a road network representation alone.

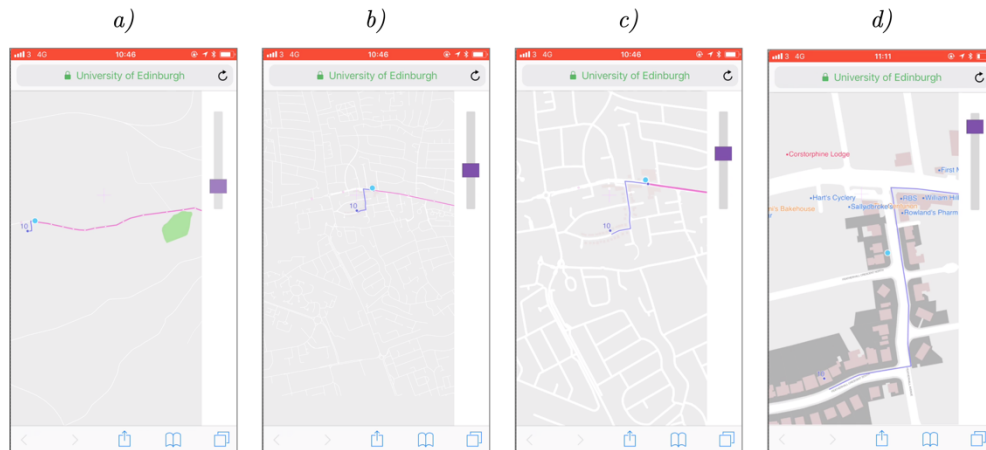


Figure 46 A series of views from study group 6 screen recordings for the study route 2 destination region.

In Figure 46, a sequence of views are shown for the third focal region in study route 2, with the region being bound by the bus stop and the final destination (a residential building).

### Prototype stimulus S3 (G7, SR2)

Study group 7 were presented with a prototype research stimulus that also included text directions integrated into the view, here referred to as ‘S3’. The purpose of this final study was to evaluate the experimental map content alongside text content. Comparing results from studies that were conducted with prototype research stimuli and the Google Maps information was problematic given that the Google Maps application is designed in such a way that the map content is just one aspect of the overall information provided. S3 allowed data collection to expand to travellers who were exposed to experimental map content but with the addition of further route information. This study also served as basis to conducted an initial piece of research into the concept of a ‘contextual zoom’ that, in parallel with the increase in scale of the map view, decomposes summarised text route directions into more detailed descriptions that directly correspond with the higher detail map content. S3 also provided opportunity for additional features to be included based on findings from study group 6.

Additionally, as S3 was implemented as a ‘click-through’ prototype with predetermined views, this forced participants to have exposure to all the experimental content. It was found in the G6 study with research stimulus S3 that



Figure 47 The sequence of predetermined views in S3, as used for navigation by participants in study group 7 (see Appendix 1 for larger illustration).

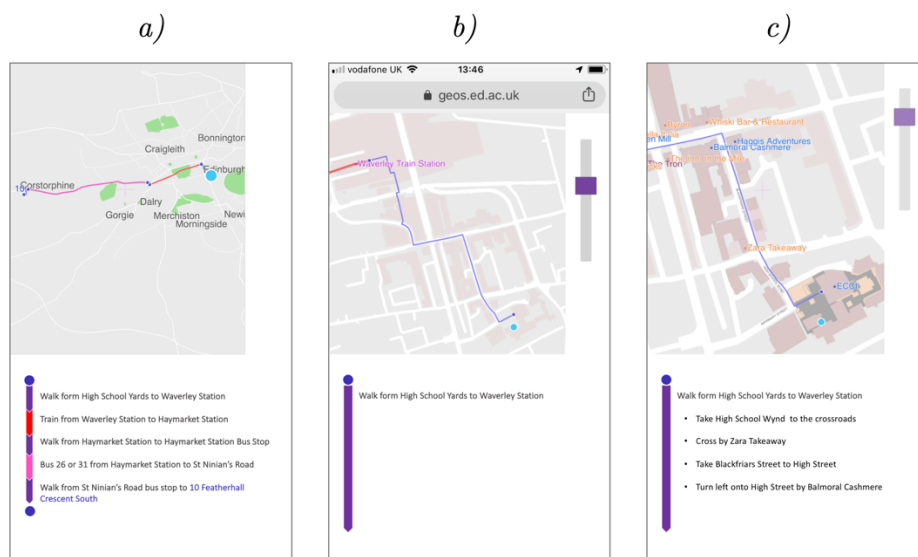


Figure 48 Views in S3, as used for navigation by participants in study group 7, showing the decomposition of the first textual route direction into a sequence of 4 descriptions corresponding with a higher detail cartographic view.

a number of participants did not increase the map scale sufficiently to have exposure to some of the high detail ‘local chart’ content, and so here was an opportunity to ensure that data could be collected based on participants having been exposed to all content for the SR2 journey context.

Note that Figure 48 simply shows enlarged versions of the first three views from Figure 47 so that example nested textual descriptions are legible.

### 7.2.3 DATA COLLECTION

GPX location data was captured using a mobile device for all participants, although for groups 1 and 4 this just served as a basis for data visualisation as opposed to analysis.

Video was captured using a high-definition camera which also served to record the output from the wireless microphone receiver. Participants wore a microphone connected to a transmitter with 60m range and so had total freedom of movement during the exercise.

During sections of journeys on buses, on trains and during the initial and final interviews, voice data capture switched to a standard mobile device as opposed to the wireless microphone receiver.



*Figure 49 Situated navigation studies equipment for data capture: mobile phone, wireless audio transmitter and receiver, video camera.*

A summary of the data collected over the course of the studies is provided in Table 4.

<b>Exercise stage</b>	<b>Participant data</b>
<i>Prior to exercise</i>	Written answers to background questions
<i>Briefing and planning</i>	Audio/ screen capture videos – during initial interview and route planning
<i>Origin to transfer/ transfer to transfer/ transfer to destination</i>	Audio/ video/ screen capture video
<i>During public transport travel</i>	Audio - during bus and train phases of the journey
Follow-up exercises	Audio/ written directions/ sketch

Table 4: Participant data

#### 7.2.4 DIRECTIONS EXERCISE

In the directions exercise participants were told to imagine that a person who they would not get to meet would need to complete the exact journey they had just completed. They needed to write directions that would be given to this hypothetical person, and while the exact format and content of these directions was up to them, they had to be numbered. This resulted in data that could be explicitly associated with the sketch using the direction's index number. While there was some variation in the style of these written directions there were clear trends and the reasons for outliers were established during the course of the interviews. The directions exercise served as a basis for eliciting the entities that played a role in navigation, but also crucially served as a basis for image schema extraction as discussed in the methodology, as well as the inference of event segmentation.

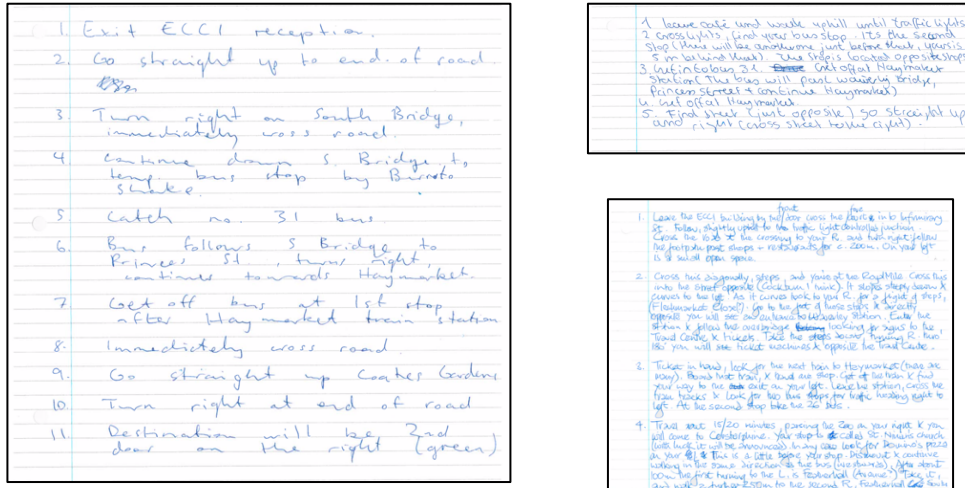


Figure 50 Examples of participant written directions for (left) G1\_11, and (right) G2\_8 (see Appendix 13.1 for larger illustration)

7.2.5 SKETCH EXERCISE

In explaining the requirements of the sketch exercise the words ‘map’ and ‘route’ were not used – participants were instructed to create a sketch that would help

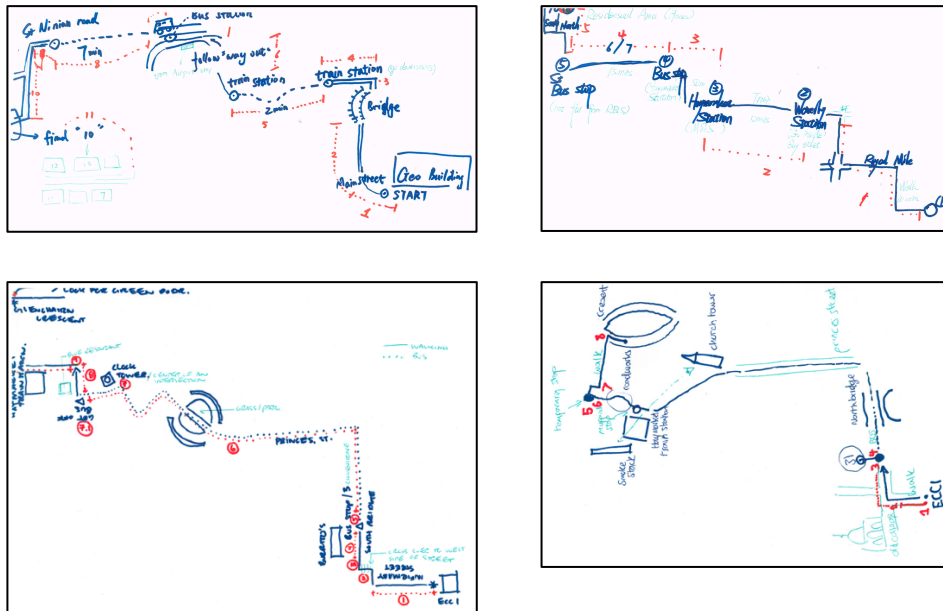


Figure 51 Example participant sketches: (top-left) G6\_5, (top-right) G7\_1, (bottom-left) G1\_6, (bottom-right) G2\_1 (see Appendix 13.2 for larger illustration).

someone to complete the journey, with 'journey' being the only description of the subject of the sketch used by the researcher. While it was implicit that the result would be something map-like, the participant in fact had complete freedom in both the format and content of the sketch.

The sketch exercise was split into three parts, with actions in each part identified through the colour of the pen used. The blue pen shows the sketch the participant created without any prompting beyond the initial instruction as described above. A duration for the exercise was not specified, however in all cases there was a natural point at which the participant clearly showed that they felt they had finished. There would be a slight relaxation in the body, the posture would change as they considered the sketch, they would put the blue pen on the table and sit back in their chair. This subconscious ritual that marked the end of the activity was effectively consistent across all participants.

Once this 'moment' had occurred, the researcher instructed them to look at the sketch and to think if, on reflection, there was anything they would add that would help the hypothetical third party to complete the journey that was not already included. It was emphasised that it was not necessary to add anything and if they felt happy with the sketch, no additions needed to be made. The participant was given a green pen for this stage, so any green markings on sketches can be understood as being conceived by the participant after this further prompt. Through this approach it was not only possible for the participant to convey their understanding of the journey, but the significance of entities and relations in the representation could be inferred from the order in which graphical objects were added.

Finally, after the participant was satisfied that nothing more should be added to the sketch, the instruction was given to use a red pen to mark the start and end of each of the directions conveyed in the numbered textual descriptions. Clearly if the sketch was the first of the two follow-up exercises for that participant, this final stage was completed after the directions exercise. This eliminated potential ambiguity in interpreting the location and spatial extent of the regions implied by each numbered direction and made the comparison of data simpler across participants and groups.



Figure 52 Research study sketch maps: (top, left to right) Group 1, participants 1 to 12; Group 2, participants 1 to 12; Group 3, participants 1 to 12; and (bottom, left to right) Group 4, participants 1 to 8; Group 5, participants 1 to 8; Group 6, participants 1 (see Appendix 13.3 for larger illustration).

This approach is argued to have been effective given the research aims. Methodological issues encountered with the approach will however be addressed in the discussion.

It could correctly be argued that the type of pens and the size of the paper given to participants had an impact on the sketches that were produced. This would however be true whichever format was chosen, so the approach taken here was to ensure consistency. Paper was size A4 and a number of participants required more than one sheet to finish their sketch. The approach taken here was not to mention the

possibility of using more than one sheet initially, but if it became clear they would run out of space, one and only one further sheet was offered. This in itself proved interesting as some people declined an extra sheet and split their sketch into parts that were not aligned, even though they had been given the option to extend the initial sketch across onto another page and maintain global alignment.




Figure 53 Research study raw data from directions exercise for a) study group 1; b) study group 2 ; c) study group 3; d) study group 4; e) study group 5; f) study group 6; and g) study group 7 (see Appendix 13.4 for larger illustration).

### 7.2.6 PRACTICAL ISSUES, CONSTRAINTS & STRATEGIES

While on the one hand a situated study provides for a potentially extremely rich set of field data, isolating variables is a clear challenge, and drawing conclusions as to the cause of some observation or set of observations, is often problematic. An example from the field studies was observing the timing and nature of interactions with information on the device during bus travel, when there is in fact variation in Edinburgh buses' on-board technology. Some buses have electronic displays with information about the upcoming sequence of bus stops, others provide no information at all. There was no way to predict prior to an individual study which type of bus it would be, and so it must be accepted that the available information between studies contain differences that will affect the experience of the participant, and therefore behaviour.

The main strategy for dealing with this issue and similar issues was to have a degree of consistency across studies such that, over the course of the 7 studies, general patterns of observed behaviour could be considered, in addition to the specific observations relating to the aims of each individual study. In other words, by building up a data set using 64 participants, an overall body of insight was developed from which to make reasonable inferences at the level of an individual study as to whether an observation was part of a trend or an outlier, or was most likely caused by some factor particular to that exercise that lay outside of the central scope of inquiry.

In terms of designing and developing experimental research stimuli, a fundamental constraint was that participants must be able to navigate using the information provided, to progress correctly and to ultimately reach the destination, having started 'at the beginning'. An example from the studies was the issue of public transport timing information. People are used to applications with live timing information, with the information often being central to navigational decisions. This type of issue posed a significant constraint on the extent to which research stimuli could take an experimental form, but for considerations such as live times, the approach taken was to compromise by supporting participants in progressing through verbal prompts, for example by offering guidance that buses would arrive regularly and to not be concerned with trying to catch a specific bus. This is provided as an example of the type of compromises that were necessary to elicit data in-situ, particularly when the participant was not only actively engaged in an exercise, but it was a pre-requisite of the study that the exercise be completed. While it was important to observe participant's own decision-making and navigational behaviours, it was also important that the exercise did not put participants in an undue state of stress, which would in itself taint the data.

In the case of critical errors such as waiting at an incorrect bus stop, the approach to verbal prompts was to wait until it appeared the final decision had been taken so it was clear the participant had misunderstood the navigational information. The researcher would then explicitly call out that a prompt was being provided before giving guidance.

### **7.3 Concluding comments**

An overview of the research undertaken has been presented including the forms of data collected and the overall approach in terms of structuring a series of studies to support the iterative development of an underlying framework based on empirical insight. In the following chapter further detail is provided as to the observations made over course of the 7 studies





## 8 Summary of observations

Following the overview of the navigation studies, a summary of key observations is now provided to give further background to the analysis and findings presented in Chapter 9. Firstly, a discussion around performance and design objectives considers the empirical evidence for being able to reasonably assess the efficacy of navigational information, concluding with 3 design criteria that emerged through the course of study. The main section within the chapter then goes on to provide an overview of key observations from the study groups across both SR1 and SR2, including observations from participant interactions prototype research stimuli and its effect on their acquisition of spatial knowledge.

### 8.1 Observations on performance and design objectives

Prior to the completion of the first few navigation exercises it was anticipated that a set of performance and design criteria would naturally emerge over the course of observations that could be refined and used in subsequent phases of research. Here ‘design criteria’ is meant in terms of criteria for the presentation of information that address observable aspects of traveller behaviour that could reasonably indicate the success of the information. It was hypothesised that these would include a) the number of device interactions, b) time spent engaging with information on the device, c) overall journey time, d) the extraneous cognitive load induced by the information, and e) the number of measurable errors in navigation.

In terms of a) the number of device interactions and b) the time spent engaging with the information, it was hypothesised that the lower the values for these metrics, the better performing the information. This follows from the logic that the more intuitively understandable the information, the less need there will be to re-check or to spend time engaging with the information in general. In

observing situated action and the development of knowledge through the participant's embodied experience of the journey, this logic turned out to be flawed however, and the hypotheses highly problematic. In summary, the only way of attributing performance to these metrics was to consider them within a broader context. A higher number of interactions could be a result of information not being readily understood and therefore the participant working to establish a clear picture. However, a higher number of interactions could also mean that *because* the information was clear, the participant was more engaged and had a higher level of trust, which led to more time spent absorbing and even exploring the information in the map. Equally, a lower number of interactions could be driven by clarity of information, with the participant therefore quickly understanding how to progress, but it could also mean that the information did not effectively address a gap in knowledge, with the participant therefore switching back to attending to environmental cues. So it was found that to use simple metrics in the analysis of these observations leads to highly misleading results. Given this issue, it was found to be too complex to quantify device interactions relative to the performance of the information in a meaningful way that would generalise over all participants.

In addition to the point discussed above, another finding was that interaction with information through pan and zoom in fact seemed to itself play a role in the development of spatial knowledge and the refinement of environmental representations over time. This will be elaborated on in Chapter 9.

With regard to c) journey time, there was in general no relationship between the competence of the participant, the degree to which they had understood the information, and the overall journey time. Study routes were explicitly designed for complexity and to reflect realistic distances in the context of large urban areas. Given this aspect of study design then, the amount of traffic or the timing of the arrival of a bus would almost always be the factors that drove overall journey time as opposed to participant understanding. Additionally, the natural pace of a person was a factor during phases traversed on foot, which again influenced journey time in a way that was either only loosely related to the level of understanding about the structure of the environment, or not at all.

While techniques have been developed for assessing extraneous cognitive load (d) in field-based research studies, these techniques affect the degree to which the study can be naturalistic – in essence the study must be controlled in such a way as it becomes about measuring extraneous cognitive only. Given the early stage nature of the present research, there was not sufficient grounds to focus on these type of cognitive load experiments at the expense of a broader more exploratory research

design. In addition to this issue, Google Maps includes the familiar GPS location marker and the current heading derived from the device accelerometer. Given that users who are not confident are able to literally ‘follow the blue dot’, it is not clear how an alternative approach to cartography could be shown to reduce extraneous cognitive load from this starting point in terms of orientation and wayfinding.

Errors in navigation (e) was clearly a factor that suffers from less ambiguity than the points discussed above and was considered in the analysis however comparison of this data across groups is argued to be problematic with the focus being the analysis of the cause of an error at the individual level. To elaborate on this point further, consider the count of the minor and major errors for G7 in Table 5. A minor error is defined as a physical action that was not a ‘critical’ mistake in the sense it jeopardised overall goal acquisition, but was ‘incorrect’ in that it demonstrated a misunderstanding of the information provided. This error type included being on the incorrect side of the street and having to double back to a crossing, and diverging from the route provided in a way that was not a ‘short-cut’ but still led to a sub-goal acquisition. Major errors were actions that required the intervention of the researcher so as to ensure the destination was reached. These errors included standing at an incorrect bus stop and walking past a turning such that the path the participant was taking no longer led to the destination.

While it is interesting that G7 had the lowest number of across the studies given that this was based on distinct stimulus type (including textual directions), when each participant’s familiarity with the city is taken into account, it becomes much less clear that this difference is meaningful, with participants who were more familiar with the city committing less errors, as one would expect. In G7 for example, overall the participants had an above average level of familiarity when compared with the other groups. Additionally, while G3 had less errors overall than G2, and G5 and G6 are comparable, this data is muddied by the fact that a number of participants using prototype research stimuli switched to Google Maps or the Lothian app at some point during the exercise. See Appendix 2 for the full breakdown of errors by each participant.

Group	Minor errors	Major errors
G2	6	1
G3	2	1
G5	4	1
G6	4	2
G7	0	1

*Table 5: Navigation errors summary*

Additionally, it was observed that even a design objective for information to support participants' acquisition of spatial knowledge is problematic. While some people are naturally inquisitive about the geography of a city, the essential goal of the traveller is simply to reach the destination as opposed to explicitly learning the layout of the environment. While navigational information may be used to develop propositional and configural knowledge in long-term memory, this is not the central role of navigational information for users, with short-term goal acquisition shaping information needs as opposed to the development of long-term memory.

Over the course of the field studies then, it was found that realistic design criteria must be looser than strict metrics and serve more as an overall guide given the need to support any user in goal acquisition for the current episode. These high-level design criteria are that:

- 1) everyone gets to the destination and in almost all cases they do not make any major errors such as boarding an incorrect bus or taking a street that moves them away as opposed to toward the destination;
- 2) cartographic objects and textual descriptions are correctly interpreted and in general correspond with user needs in terms of the level of abstraction for the given scale; and
- 3) that users feel a sense of control in terms of being able to determine their current location and in their ability to progress toward goals from that location.

To this third point, sense of control can simply be understood in terms of the city *feeling manageable*.

## 8.2 Empirical observations overview

Here an overview of observations is provided for the navigation exercises, sketch exercises, directions exercises and the prototype research stimuli. It should be emphasised that all sketch map figures show 'raw' data drawn by study participants. To reiterate the explanation provided on page 112, the colour convention used in the sketch maps is as follows:

- blue for the initial sketch, where the drawing is based simply on the direction to create a sketch to help a third party to complete the same journey
- green for additional elements that were added after prompting
- red to mark and number the extent of the textual directions on the sketch

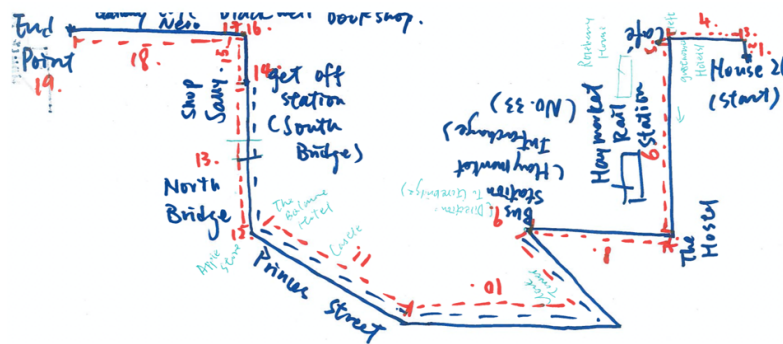
For all figures in this chapter, refer to page 94 for a view of the extent to scale.

### 8.2.1 OVERVIEW OF SR1 STUDIES (G1, G2, G3)

#### *Overview – SR1 sketches*

- The degree to which the linear depiction of the path of the route had primacy in the representation was greater than anticipated. Creating a sequence of ‘lines’ was the first action was begun immediately in almost every case, irrespective of the exercise type. Arguably this could be influenced by the fact that people are used to using applications that present journeys in terms of a sequence of linear paths, but irrespective of the reasons, the basic reality was that participants’ understanding was structured in this manner. In the majority of cases this took the form of a sequence of line objects, however occasionally a series of dots was used, for example to distinguish between mode of transport. Figure 54 provides 4 examples. It should be noted that the red dashes were added based on a prompt to clearly show the extent of the region described by the each numbered textual directions.

a)



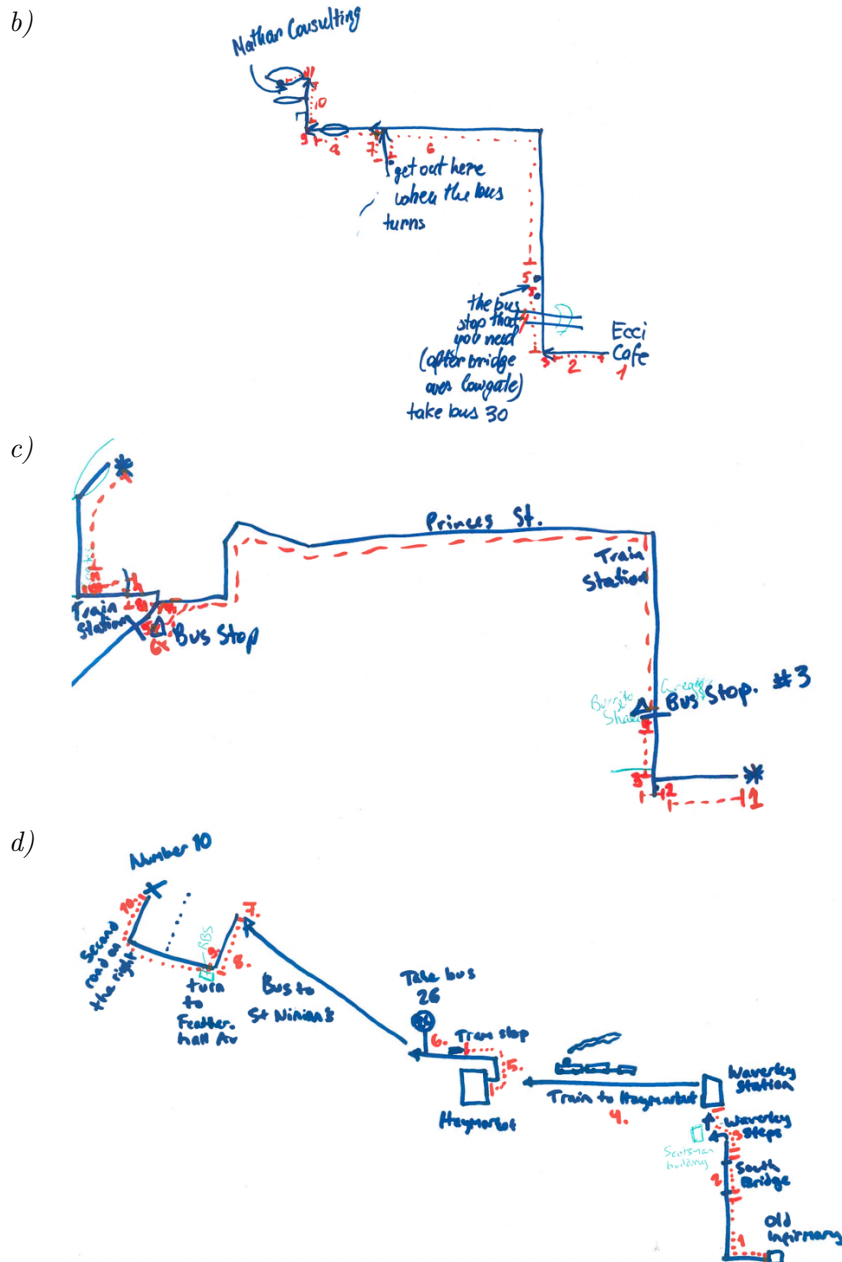


Figure 54 Example sketches showing the central role of linear representations in a) G1\_11, b) G2\_5, c) G3\_2 and d) G5\_5 (see Appendix 13.5 for larger illustration)

- Many participants did not see a need to specify the geometry of the road and felt the road network itself was implicit (see examples in Figure 54).
- Fewer intersections were explicitly conveyed as expected, however when given an opportunity to reflect on their sketch, adding intersection detail was a common action. Similarly, when prompted to consider if any

element should be added, approximations to the geometry of the road network was a class of the elements introduced into the sketch. See examples highlighted in figure 55.

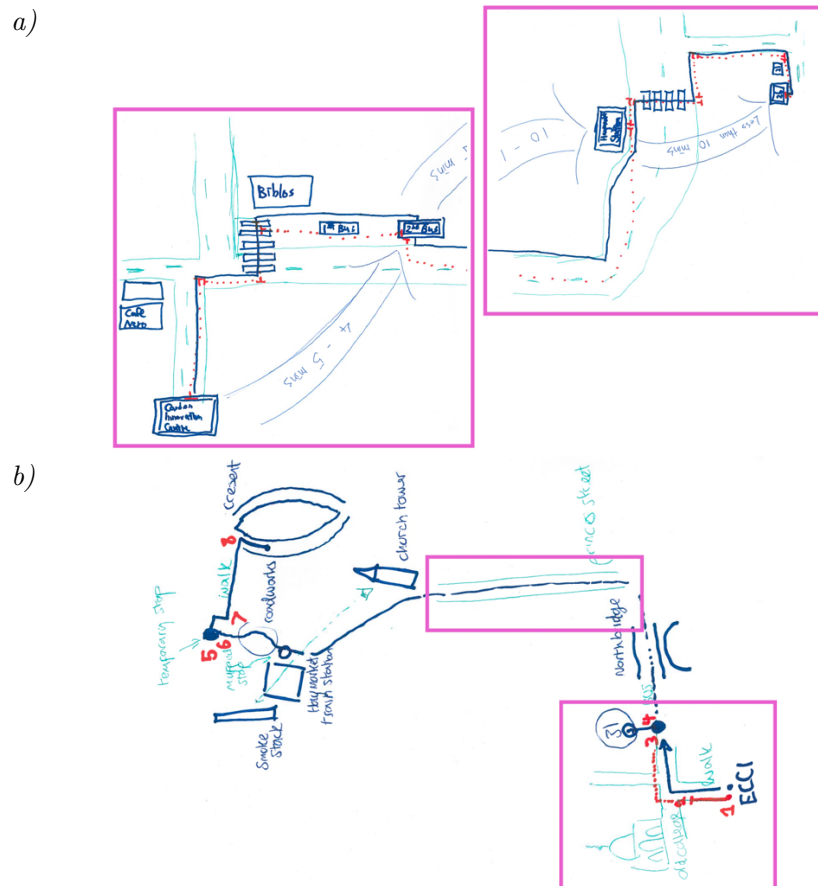


Figure 55 Addition of road network geometry after prompting shown in green, examples shown for a) G1\_5, and b) G2\_1 (see Appendix 13.6 for larger illustration). Bounding boxes (pink) simply used to highlight the road features in question.

- As expected, the explicit communication of features exhibited very high correlation with the location of primary actions and almost all other features were adjacent to the path of the route and used in such a way as to support orientation. This can be observed by visualising features and directions as point objects over the centreline network, as shown in Figure 56. This visualisation is derived from the sketch maps data, i.e. the features (blue and green) and directions (red) expressed by participants through drawings and text were digitised.



Figure 56 SR1 study participant features and directions data (colour corresponds with sketch maps colour conventions). In panel a) an overview is shown, with panels b), and c) showing two example regions at a larger scale (see Appendix 13.7 for larger illustration) Relation between features, directions and structure of network shown through spatial distribution.

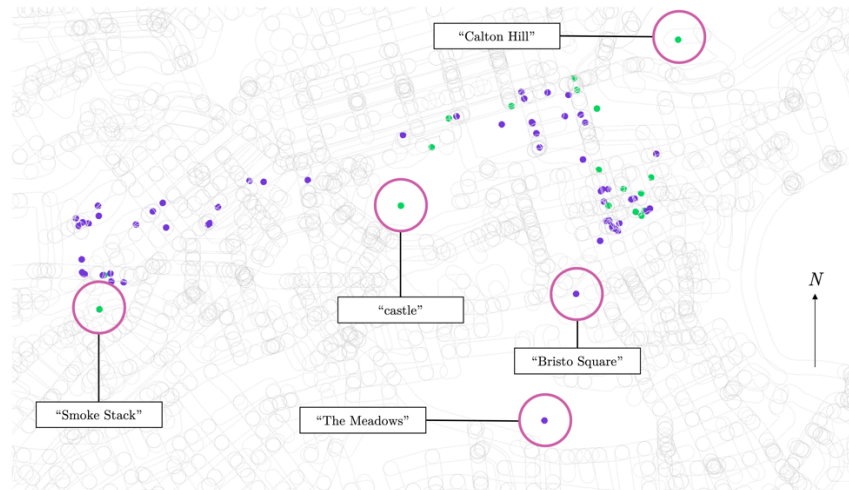


Figure 57 Region types implied by empirical data and derived from the underlying network (G2 data used for illustration). Circular regions derived from the location of network intersections, elongated regions formed around the navigable path connecting sequential intersections, and a broader region created by the overall extent of the journey. Labelled features are members of this third region type.

- Features conveyed by participants in empirical studies were observed to belong to one of three region types. In Figure 57 two of these region types are made explicit, and one is implied. In almost all cases, features that were conveyed by participants would be captured by creating circular regions around the structure of the network, or elongated regions that form around the navigable path connecting sequential intersections. Features that are not adjacent to the traversed path but were nevertheless conveyed were features that are both visually prominent such as the ‘Smoke Stack’ (a large chimney), the ‘castle’, and ‘Calton Hill’. These features can be understood as corresponding to a broader region type that only needs to include a very small subset of entities. In Figure 57 there are two features which in fact are not visible from the study route, ‘The Meadows’ and ‘Bristo Square’, however both of these were conveyed by a single participant who had passed through both of these features on the way to meet the researcher before commencement of the study. These outliers are argued to be in the dataset because they formed part of that specific participant’s cognitive representation given their experience immediately prior to the study, although they do support the broader point about including key features in the overall view of the journey’s extent as means of supporting reference memory.
- Features that were not strictly adjacent to the route corresponded with significant landmarks, however it was unexpected how seldom these types of landmarks were included (evident in Figure 67-a),, even after prompting whether about the opportunity to additional content. For example is was

expected that Edinburgh Castle would have been explicitly conveyed in the average case but this was not observed. In the few cases where this type of feature was conveyed, visual prominence was always a key driver. Examples in Figure 58 show visually prominent features that are not geographically adjacent to the path of the route but visible.

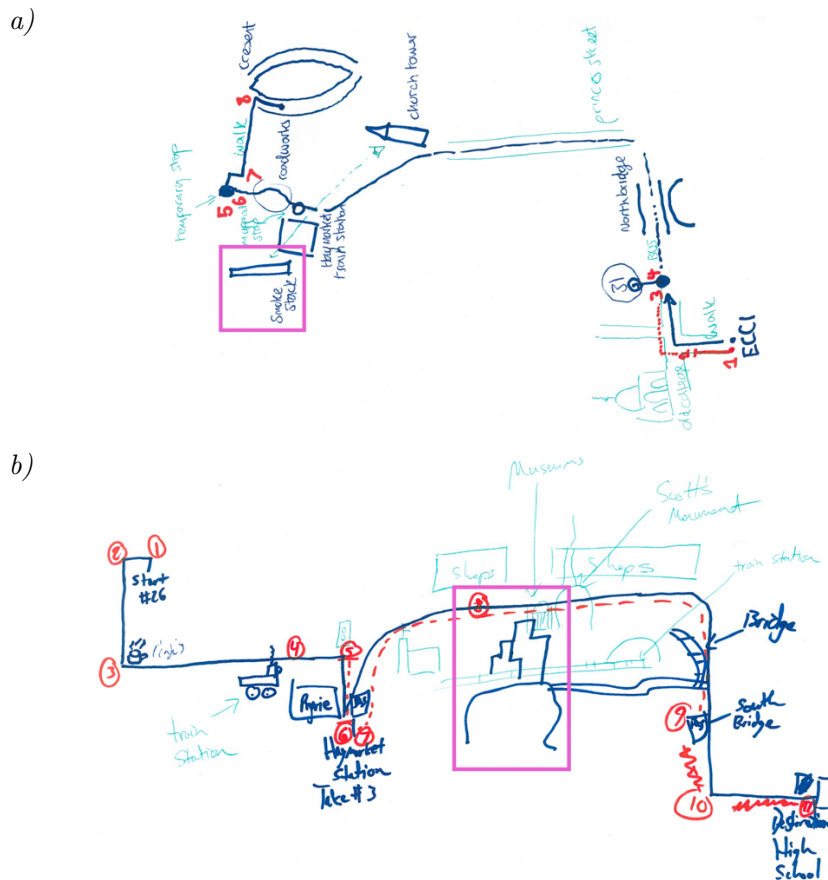
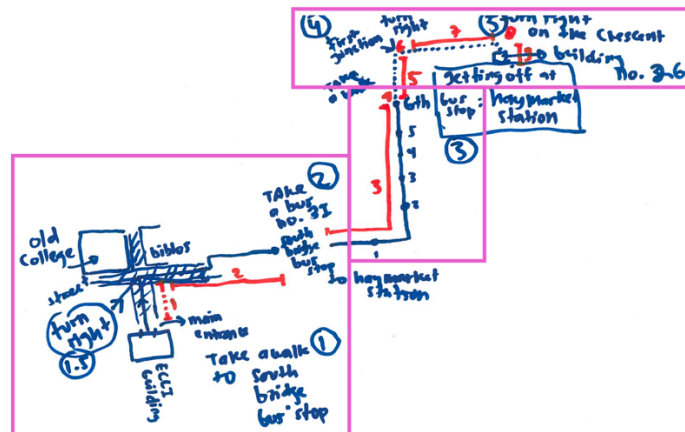


Figure 58 Outlier examples of landmarks not adjacent to route but explicitly conveyed in a) G2\_1 (the 'smoke stack'), and b) G1\_4, Edinburgh Castle (see Appendix 13.8 for larger illustration).


- Over the 3 studies for SR1, no feature was conveyed that is not visible from the route except for outliers such as those highlighted in Figure 57. While it was expected that the sketches would be strongly characterised in terms of route-based knowledge, the extent to which this was the case was striking. Particularly given the number of well-known and distinctive potential landmarks in the city.

- Very few objects were specified that represented entities other than buildings. These seemed to be in one of three categories which aligned with expectations based on the results of previous studies. Either a) because of some visually distinctive characteristic (e.g. the ‘crescent’), b) because of proximity to a critical action (e.g. the ‘clock tower’ near the bus stop), or c) because the feature played a role in the physical actions required to progress along the route (e.g. a crossing, a traffic light). The unexpected observation in the data in terms of features was the lack of features included because of some particular visual characteristic. The significance of landmarks was even more strongly driven by the functional relation to primary actions than expected, with features of particular visual salience seeming to play a far more diminished role than predicted.
- Overall, the structure of sketches in terms of the regionalisation of the representation followed the expected form, with phases of the journey completed on foot corresponding with higher detail and causing distortion in terms of direction and distance between the representation of these localised regions. (as shown in the examples in Figure 59).
- Only one participant created a survey view in approximately the correct proportions in terms of the different phases of the journey. This outlier seemed to be attributable to fact the participant was an architecture student and so had experience in sketching survey views (see Figure 60). Here the characteristic of a survey-like view is apparent through the relative proportions of the linear representations of the path of the route and through the overall regionalisation across the extent.


a)



b)



c)



d)

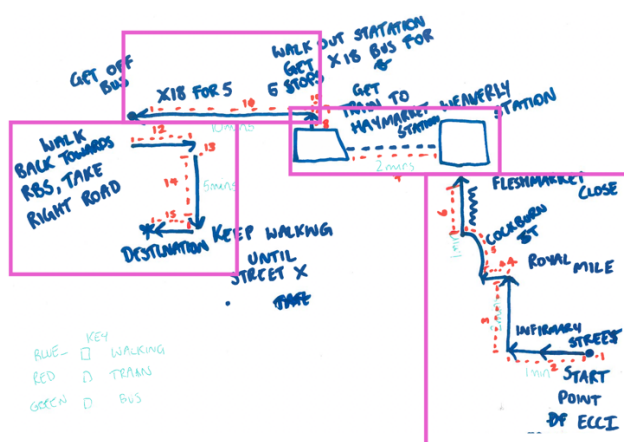


Figure 59 Regionalisation conveyed through sketch representations, marked by approximate bounding boxes (pink) - examples a) G2\_10, b) G2\_3, c) G4\_1, and d) G5\_4 (see Appendix 13.9 for larger illustration).

- An observation that is not easy to explain is that there was no clear difference overall in the sketches between those created by participants who were ‘following’ and those who were using Google Maps, however the participants in G3 who used the S1 prototype stimulus were in general clearly distinct in terms of the structure and content of the sketches they created. The S1 views were designed around a much-reduced set of content versus the Google Maps views, and this was reflected in the sketches, however all participants completed approximately the same journey and all had the opportunity to add further detail to the sketches to help the hypothetical third party. It seems a larger sample would be required to really be able to draw conclusions about this.




Figure 60 Survey-like sketch map representation, participant G1\_10 (see Appendix 13.10 for larger illustration).

- An important consideration in terms of the methodology was that objective analysis based on the overall layout of sketches proved very problematic. Prior to completion of the first few sketch exercises it had been anticipated that analysis would include objective measures of distortion, for example the differences in relative scale of distinct regions within the sketches, the level of angular distortion between regions, and so on. In practice it was observed that sketches did not effectively convey participant understanding in a way that could be consistently analysed using these measures. If the exercise had explicitly been to create an ‘accurate’ map, then such techniques could be applied. However, participants who created sketches that significantly disregarded the underlying spatial structure of the geography did not do so because this reflected their literal understanding of physical layout. See figure 61-a for an example sketch from a participant who did not see their layout as a literal representation of spatial structure).

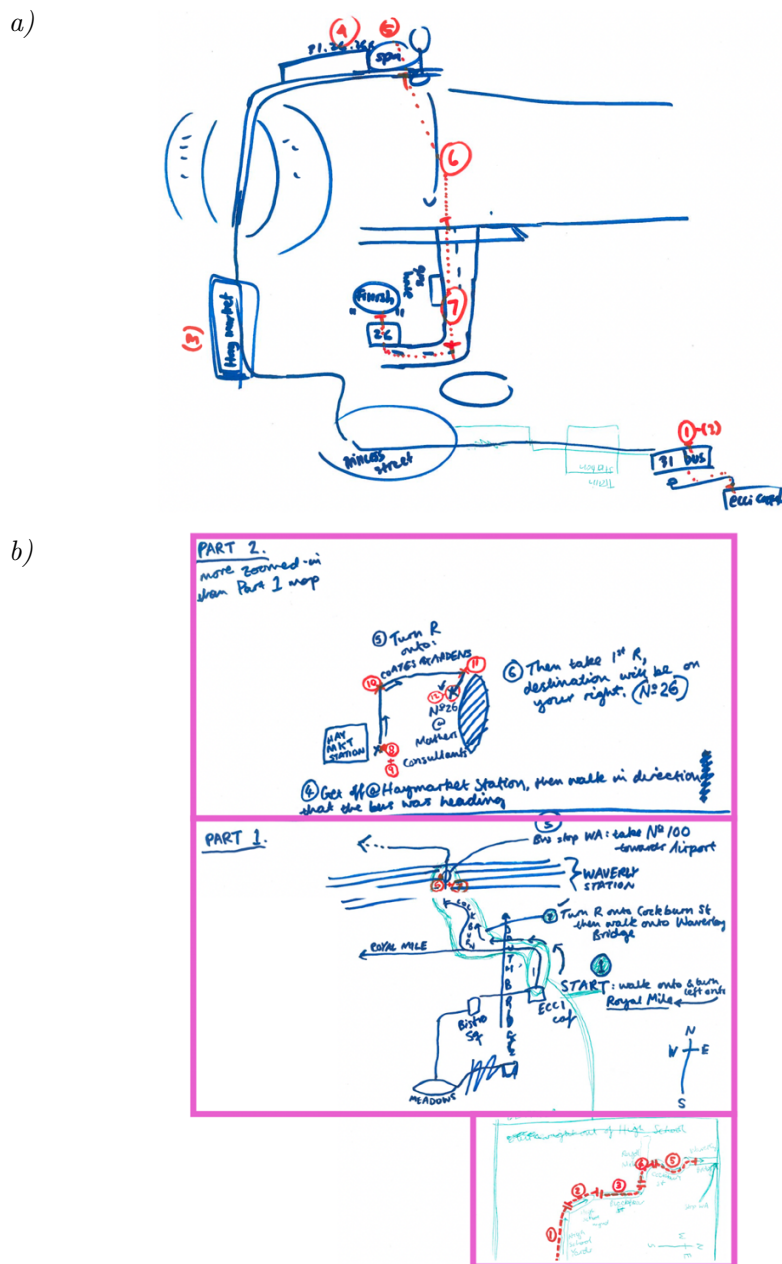


Figure 61 Metric distortion of spatial structure shown through sketch examples from participants a) G1\_7, and b) G2\_6 (see Appendix 13.11 for larger illustration).

- It seemed in fact that for some participants there was a sense that given the aim of the task, spatial accuracy was not an important objective. Here then it seemed there was an effect in terms of the interpretation of the task, as

opposed to significant spatial distortions conveying something particular about participants' underlying spatial knowledge.

- In Figure 61-b we see an example of a sketch that is explicitly divided into a series of submaps. There was a pervasive tendency for participants to convey regions traversed by public transport as more highly schematised than regions traversed on foot, and in cases such as the example in 61-b, even to completely abstract them away (shown through the 'jump' from "PART 1" to "PART 2"). However this did not signify literal understanding, and crucially did not reflect participant preferences in terms of the form of the information to be used for navigation itself. In other

a)



b)

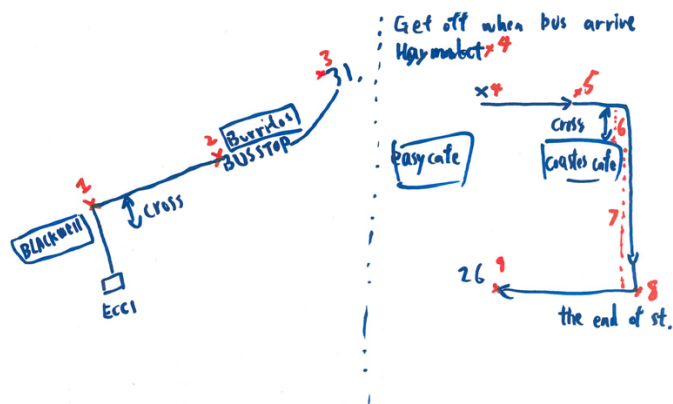


Figure 62 Two example sketches, each based on a highly abstract regionalisation of the environment a) G6\_2, and b) G3\_5 (see Appendix 13.2 for larger illustration).

words, after a journey was completed, journey phases completed on a bus, for example, were understood to have less significance in terms of the information required for navigation, but this understanding was only developed after the fact.

- To provide further examples with respect to the previous point, Figure 62 shows 2 further example sketches in which the representation is highly schematised and the overall structure of the geography depicted at a very high-level of abstraction. Here regions are shown based on their approximate topological relation but with no representation of the geography connecting the regions at all. In both cases the participants had issues with navigation including minor errors (see Appendix 1 for errors table), however when asked about their expectation of the efficacy of their sketches, both displayed confidence. This reflects a strong asymmetry in perception before/ during/ after the journey, however it does also provide support for a representation based on event segmentation, with the regionalisation in each of these sketches being driven by primary anchors such as transfers, and the structure of the navigable network then determining the more granular event structure within this overall partonomy.

### Overview – SR1 directions

- It is argued that the directions given by a participant reflect their perception of event structure and their overall understanding of the episode in terms of navigation. Within this overall context, the directions data strongly followed the expected pattern in terms of reflecting an event structure driven by actions and by the topological structure of the navigable network. For example in Table 6 textual directions convey actions such as ‘turning’ and ‘crossing’, together with descriptions of the navigable path itself (‘pedestrian crossing’, ‘cobbled street’), together with features that serve to anchor the representation such as the bus stop and the sushi restaurant.

Participant	Direction index	Text
G1_10	3	<i>At the first pedestrian crossing, cross over the road and continue north on south clerk street</i>
G2_1	7	<i>Walk up the cobbled street to the crescent</i>
G3_6	2	<i>Cross the street, walk a few mins to the right until you get to the 31 stop</i>
G3_9	4	<i>Turn right at Bento Sushi</i>

Table 6: Typical structure and content of textual directions

- Key actions included turns, road crossings, transfers between transport mode and path traversal. The only other clear driver of event segmentation

was the subdivision of the path of traversal due to either a salient aspect of the structure of the navigable network, or due to the presence of a visually distinct feature that formed a boundary cue.

Participant	Direction index	Text
G1_8	5	<i>You hop on the bus (route number 37 outside a souvenir shop)</i>
G2_1	6	<i>Once you alight cross over and take the road just passed</i>
G3_3	10	<i>When you are nearing Haymarket station press bell</i>
G3_9	11	<i>Exit bus at Haymarket train station</i>

Table 7: Example representations anchored by transfer actions

- In terms of distinctive aspects of the navigable network, this can be understood in terms of characteristics such as the length of road, geometry of road (e.g. straightness), the location of an intersection between two distinctive roads or at least the location of a change between one distinctive direction and another. Examples include the sharp right-hand turn as the bus turned left on Princess Street, and the turn at the end of Princes Street onto Shandwick Place (see examples in Table 8).

Participant	Direction index	Text
G1_2	7	<i>Then the bus will go to a long street that is like an avenue, it is called Princes Street</i>
G1_3	6	<i>...the bus will take you down South Bridge, then Princes Street and then will veer left towards Haymarket</i>
G2_1	8	<i>When you reach the crescent walk along the right side (near side) until you reach number 26 (your destination)</i>
G2_5	10	<i>Walk until the you reach the 2<sup>nd</sup> crescent (walk past 1 crescent)</i>

Table 8: Example textual directions based on distinctive characteristics of the geography of the network

- In terms of visually distinct boundary cues, North Bridge and the ‘crescent’ that was repeatedly conveyed serve as examples of this.
- In terms of the action sequences that were ‘chunked’ into directions, as expected these consisted almost entirely of progression along navigable paths given the underlying topological representation of the environment, and actions that corresponded with changes of direction.

Participant	Direction index	Text
G1_3	7	<i>Once you get off the bus you’ll need to cross the street and walk towards the station on the side of the street opposite the station</i>
G2_8	1	<i>Leave the café and walk uphill until the traffic lights</i>
G2_6	5	<i>Walk down Cockburn Street then cross roundabout to continue onto Waverley Bridge</i>
G3_1	2	<i>Walk all the way up to South Bridge and cross the road and go to the right</i>

Table 9: Example ‘chunked’ actions sequences conveyed through textual directions

### Overview – navigation styles and patterns of information interaction

- A common distinction in navigation styles for participants in G2 was planning to the destination versus planning to some intermediary point. It was observed that often participants would not surface a detailed view of the destination before leaving but would plan to either the start or end of the bus phase (i.e. the to the first or second transfer).
- In Table 10 styles of information interaction are shown for each participant group during the planning phase, where views of the geography are categorised as being either high, medium or low detail. ‘High’ refers to a series of interactions that led the participant to a detailed view of a small area, for example focusing on a small sequence of streets. ‘Medium’ refers to a lower level of detail in which a reduced amount of features are visible and polygon features are at a higher level of generalisation. ‘Low’ refers to a view in which only a very limited number of point features are visible, with the main map content simply being a generalised view of the road network. For each region, the letter representing the level of information abstraction refers to the highest level of detail the participant accessed. So if a region is

marked ‘*h*’ for example, this means that a high detail view was the lowest level of abstraction, but does not convey whether the participant also accessed an ‘*m*’ level view for the same region. This simplification serves to make the presentation of data in a summary table clearer.

- For SR1, the planning styles are characterised across the origin region O, the region based on the bus phase T, and the region around the destination D1. Only a single planning style is attributed to G3 because the S1 prototype used by this group essentially only supported a single planning style (i.e. the planning style characterised by G2-a also).

Route	Planning Style	O	T/T1	D1/T2	T3	D2
SR1	G2-a	<i>h</i>	<i>l</i>	<i>h</i>		
	G2-b	<i>h</i>	<i>l</i>	<i>m</i>		
	G2-c	<i>m</i>	<i>l</i>	<i>m</i>		
	G2-d	<i>m</i>	<i>m</i>	<i>h</i>		
	G2-e	<i>m</i>	<i>m</i>	<i>m</i>		
	G2-f	<i>h</i>	<i>m</i>	<i>m</i>		
	G3	<i>h</i>	<i>l</i>	<i>h</i>		
SR2	G5-a	<i>l</i>	<i>l</i>	<i>l</i>	<i>l</i>	<i>m</i>
	G5-b	<i>h</i>	<i>l</i>	<i>l</i>	<i>l</i>	<i>l</i>
	G5-c	<i>m</i>	<i>l</i>	<i>l</i>	<i>l</i>	<i>m</i>
	G5-d	<i>m</i>	<i>m</i>	<i>m</i>	<i>m</i>	<i>m</i>
	G5-e	<i>l</i>	<i>l</i>	<i>l</i>	<i>l</i>	<i>l</i>
	G6-a	<i>h</i>	<i>m</i>	<i>h</i>	<i>m</i>	<i>h</i>
	G6-b	<i>m</i>	<i>l</i>	<i>l</i>	<i>l</i>	<i>l</i>
	G6-c	<i>h</i>	<i>l</i>	<i>l</i>	<i>l</i>	<i>m</i>
	G7	<i>h</i>	<i>m</i>	<i>h</i>	<i>m</i>	<i>h</i>

Table 10: Characterisations of information interaction during journey planning for study groups G2, G3, G5, G6 and G7.

- It should be noted that ‘planning’ refers to the fact that these styles of information interaction were observed prior to the journey commencing – i.e. while the participant was still at the journey origin. The same analysis could be done at a number of points along the journey, or within a defined set of regions, however here it should simply be highlighted that travellers seek information at different levels of abstraction, and that a certain level

of heterogeneity in the patterns of travellers' information interaction occurred throughout the journey.

- In terms of the 'h-l-h' planning style for G3, apart from the connecting polyline between the two focal regions, there was no information available to participants outside of the focal regions (region 'O' and 'D1' in Table 10). This lack of information around broader geographical context drove participants to interact with the prototype such that a high level of detail in each region focal was accessed. As shown in the G2 planning styles however, it is not necessarily the preference for participants to access this level of information during the planning phase, with for example a tendency to access only a medium level of detail if the traveller has some familiarity with the geography.
- As journey length increases, the pressure on an agent to reduce exposure to information with low abstraction in the planning phase increases, as shown by the planning styles for G5.
- In terms of planning to the end of the bus phase, another distinction was whether the sequence of views consisted of a) a move to a lower detail, then a zoom back to a higher detail for a new focal region, or b) to pan along the path of the bus route at a medium level of detail.
- In Figure 64 view event onsets are shown for G2, in other words the locations at which a series of interactions with navigational information on the mobile device were initiated. Cluster groupings are based on the shared causal basis of traveller information needs.
- In grouping A view event onsets were driven by a motivation to check that the participant had caught the correct bus. For example checking the sharp right hand turn occurred as opposed to continuing straight or to the left.
- Grouping B is based on the observation of a motivation for beginning a device interaction just after a natural point of relaxation – there would be the initial anxiousness about being on the correct bus, followed by a moment of relaxation, followed by the idea to look ahead in anticipation of reaching the destination.
- In grouping C device interactions were also driven by a motivation to attend to later stages of the journey, however here the shared cause was the physical structure of the geography, with the shift in direction after a long straight prompting participants to shift focus.
- In the grouping marked D, participants knew the bus stop was in fairly close proximity and so attention shifted to more immediate visual cues.
- The groupings shown in E and F both relate to device interactions that began after the participant transferred from the bus to walking (i.e. for two different bus stops). This was a near pervasive pattern of interaction, with the strong preference on the bus being to focus on the task of identifying the correct stop, and the moment immediately after leaving the bus freeing

up attentional resources to switch to the next key task (walking to the destination).

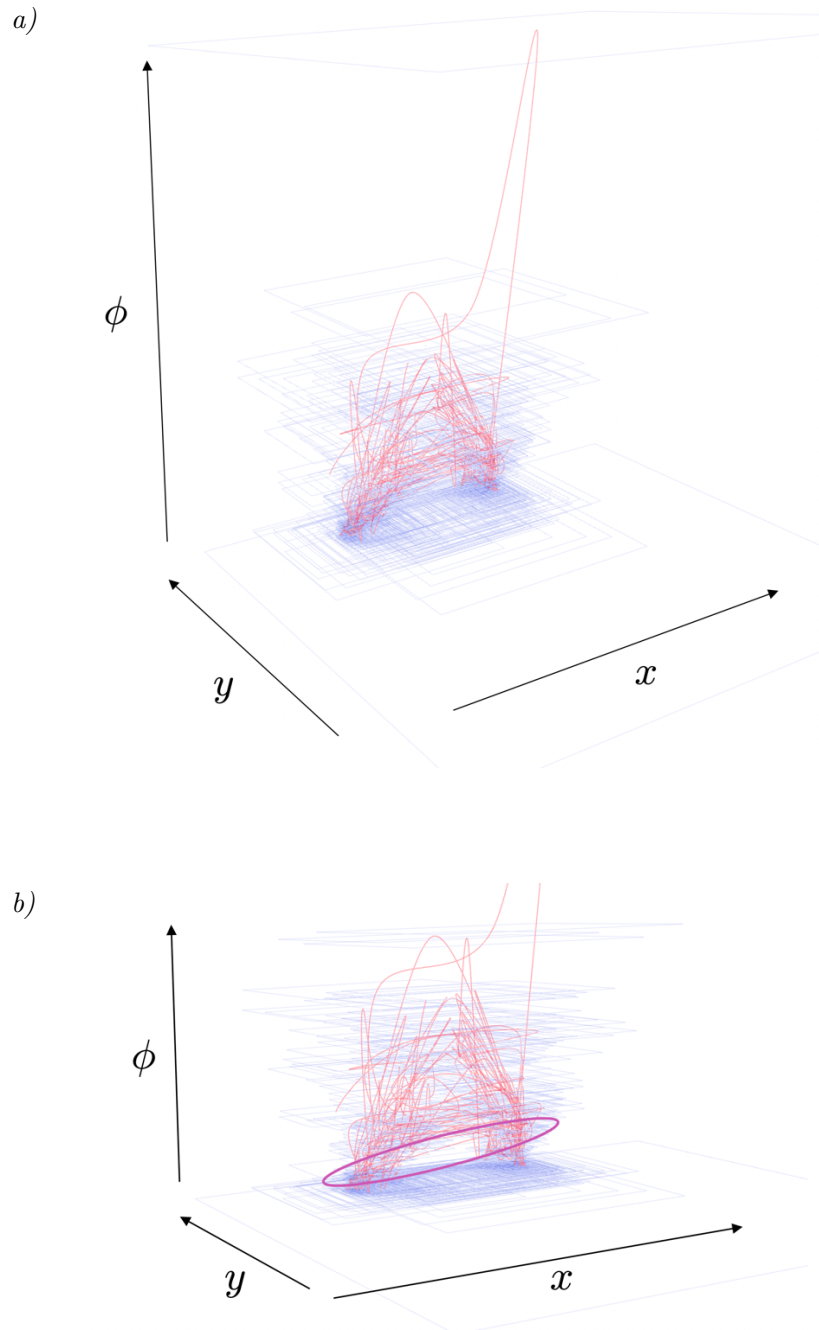


Figure 63 a) View curves for G2 (red) shown with the image space implied by the view centre (blue), and b) a view curve highlighted along the 'crossbar' of the 'A' form. Axis  $x$  = easting,  $y$  = northing,  $\phi$  = map scale.




Figure 64 Clustering of study participant mobile view event onsets for G2. Regions A through F denote approximate groupings based on distinct causal relations between journey phase, physical environment and participant information needs. Hue of point objects and vertical markers denotes participant.

- As expected, another style was to almost literally follow the location marker, with some participants even maintaining a view event for the duration of the bus journey – not attending to visual cues in the environment at all but simply waiting for the location marker to reach the graphical marker for the target bus stop.
- In terms of interacting with map views, the ability to control the position in view space seemed to play a very important role in knowledge acquisition. For example a regularly observed sequence of actions was for participants to repeatedly zoom in and out of the same region. This behaviour accounts for some of the density of image spaces shown in Figure 65. As opposed to reflecting lack of clarity of the information this type of interaction in fact seemed to support the development of hierarchical reference memory, with vague understanding being refined and reinforced over time.
- In terms of the bus route, there was a regular pattern of interaction, with the mobile device being accessed very soon after getting onto the bus and kept open for a short period after the bus began moving so as to be able to check the direction and route met expectation. Then there would be a period of relaxation, followed by a natural moment to check the device again, after which the device was referred to with increasingly short intervals between view events as the bus progressed toward the target bus stop.
- Most participants favoured checking the device after leaving the bus for planning the final phase to the destination. This overall pattern was very consistent.
- It was observed that characterising the content itself that would be accessed at particular locations along the route was problematic, and that in practice

it should be expected that a user will want easy access to any view that they feel is related to the current journey context, irrespective of the relation between the geographical extent within that view and the current location.

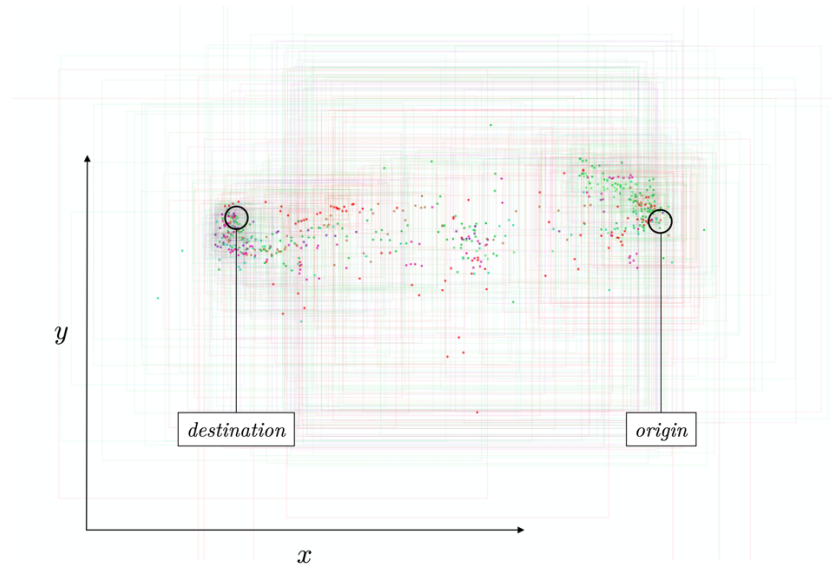


Figure 65 Image spaces for G2 participant map interactions where the point objects denote image space centres and the squares denote image space extents over the underlying geography.




Figure 66 Complex view curve example from participant G2\_4. Participant location marked as 'current position' on the plane formed of easting and northing axes.

- The previous point is illustrated in Figure 66, with a view curve showing the path of interaction through view space for a single view event executed at a location labelled 'current position'. This movement through view space reflects the importance of exploratory interactions as well as the importance revisiting views previously accessed. The strong sense from observing participants was that this type of interaction served to reinforce a cognitive representation of the geography relative to the current task context, with a loose representation being refined and embellished over time, as opposed to being learned 'all at once'.

### **Overview – S1 prototype stimulus findings**

- There was a clear influence of the prototype research stimulus on the sketches. In almost all cases no element was added that was not represented in the experimental views and the form of regionalisation strongly reflected the visualisation participants had been exposed to.
- While a tripartite regionalisation was anticipated in participant understanding of the journey, there was arguably more consistency in the structure and content of the representation conveyed in the G3 sketches. Two illustrative examples are shown in Figure 67 with the origin, destination and transfer primary anchors serving as the basis of regionalisation, and a de-emphasis on the journey phase completed by bus, shown by direction number '3' in 67-a and by direction number 4 in 67-b.
- Overall it can be said that the sketches clearly reflect the features present in S1 even if only a small number were added, as well as the overall regionalisation underlying the selection and visual display. In other words, a regionalisation based around the three main journey phases, in which the level of detail and level of schematisation follows a tripartite structure with higher detail 'local charts' for the phases on foot between the origin and the first transfer, and between the second transfer and the destination.
- The 'crescent' entity which has been discussed did not feature once in any of these sketches. This suggests that participants were drawing on their memory of the representation as opposed to drawing on knowledge acquired both through the representation and by other means (either existing knowledge, or through sensory experience during the exercise). Why this would be the case for the experimental views and not the Google Maps views is not clear however.
- Evaluating the first three points in the proposed specification was problematic as the geometric distortion proved to be a very ineffective way of presenting the information, with this issue really framing the overall observations from the study.




Figure 67 G3 example sketches based on exposure to prototype S1 (see Appendix 13.13 for larger illustration).

- Evaluating the first three points in the proposed S1 specification was difficult as the geometric distortion proved to be a very problematic way of presenting the information, with this issue really framing the overall observations from this study. Observations and empirical data included both positive and negative aspects of this way of presenting navigational information, with sentiment reflected in verbal responses in interviews, as well as in behaviour such as switching to alternate sources of information during the exercise.
- Sentiment expressed by participants can be understood to be 'positive' where either an explicitly positive statement was made or where the participant expressed understanding based on the information, and similarly

negative where either an explicitly negative statement was made, or the participant expressed a lack of satisfactory understanding given the task. An example is provided in Table 11 for participant G3\_1.

<i>Sentiment</i>	<i>Prompt</i>	<i>Quote</i>
<i>Neg</i>	<i>So talk me through what you can see</i>	<i>Is that us, where we are located?</i>
<i>Neg</i>	<i>What do you think this is showing you?</i>	<i>The route... oh wait... at first it just showed part of the bus route... it feels like a different route now.</i>
<i>Neg</i>	<i>What do you mean it feels like a different route?</i>	<i>Well at first it had a line... maybe it is the same... I don't know, it is a bit hard to read when it doesn't have all the extra... because it doesn't have everything around it that's not directly on the route.</i>
<i>Neg</i>	<i>What are you looking for?</i>	<i>Just more information, just of the whole town... so I just happen to know where Old College is but it's nice to have more information around it. So, to locate it – the route in the city. Oh yea, you can see it is South Bridge... but it's a bit out of perspective as well isn't it? Because if you see how long that is, compared to Haymarket... so that bit is shortened... so the whole route from South Bridge to Haymarket is about half of the whole of Infirmary Street here... which is a short street ...and then its bigger again, so it changes.. Is that because of how long it takes, because that is a walking bit? Because there you get on the bus... but still it wouldn't make sense but this should be longer.. it's quite difficult to navigate....</i>
<i>Neg</i>	<i>Talk me through what you think that</i>	<i>Well it's definitely part of a map... it's a bit like when people go hiking and they only take like little bits of the map that they plan to walk through... and it's like</i>

	<i>is showing you?</i>	<i>well, if you stay on course that might work but best to have some context as well... so I think with this you need quite a lot of information already to get the info out of it. I would personally like more context around the route. Also because it is kind of distorted distances it is kind of hard to position where it actually all is.</i>
<i>Neg</i>	<i>Why does that matter?</i>	<i>Just to give you an indication... at least when I think about a route.. like I have a map in my head... I have to go from that point to that point, so I've got to go to those streets. But here because you can't really judge it. Like if I wouldn't have known that between South Bridge and Haymarket, what that distance is, I would probably have thought... it's just around the corner, why would I need a bus for that, whereas... it's actually quite a while</i>
<i>Pos</i>		<i>I guess it helps you focus, like all the names of the different places... on the other hand it is very clear that... because it condenses it so much, like it's clear I just take... here from Infirmary Street, I take a right, and then get on a bus, and that's going to take one left turn and then basically I'm at Haymarket.</i>
<i>Pos</i>		<i>Oh wait it does say here....</i>
<i>Pos</i>	<i>What were you looking for?</i>	<i>On which side of the street I would have to take the bus... so it's just on the other side of the street... take a right then cross the street... and then at Burrrito and Shake, that's where the bus stop is...</i>
<i>Pos</i>	<i>I will just prompt you that you need to move the</i>	<i>Oh so that makes sense why that keeps moving... so it goes to the parts of the route then is zooms out on that... ah that makes sense... because I wondering when I had the slider here why it makes that</i>

	slider the other way	<i>bit so much bigger than that but it's a different part of the route</i>
<i>Pos</i>	<i>So talk me through what you are seeing</i>	<i>So that's actually pretty useful when you really zoom in on a different part... like the walking parts where you actually have to be more detailed... because the bus route... like its good just for visualising where the bus will go past, but you don't have to be that precise because the bus driver knows where he's going... that is actually quite useful.</i>
<i>Pos</i>	<i>So talk me through what we are going to do</i>	<i>So first just get back from Infirmary Street onto South Bridge, cross the road, get bus 31, on that until Haymarket Station, and there get off, cross the street again, then take a right up to Glencairn Crescent... seems simple enough.</i>

*Table 11: Example feedback for prototype research stimulus S1 during initial exposure by participant G3\_1*

- The data from participant G3\_1 is provided as an example as it serves to illustrate the characteristic response from participants in terms of the aspects of S1 that warranted further research and the aspects that proved to be too problematic. For example disambiguating regions that are predicted to be accessed at a higher level of detail was commented on as being a useful feature, but the distortion of the plane caused confusion and was not well received as a basis for presenting the map content. It should be noted that responses tended to become more positive as the participant came to understand the overall functionality of the prototype, but initially the slider interaction to move between focal regions tended not to be intuitively understood as it did not fit with existing mental models of how to interact with a map application. This point is illustrated in the positive and negative comments in Table 11 which are chronological and for the most part, consecutive.
- Comments about the usefulness of the approach expressed by G3\_1 can be compared with, for example participant G3\_2:
  - “It is kind of a neat idea to do the... zoom in at one end or the other depending on which way you toggle, but that is not intuitive at all, as you saw, for me at least.”*
- Clearly it could be argued that the novel slider interaction caused confusion as opposed to the geometric distortion itself, but even if further experiments

were done with the slider removed, the issue of interaction in fragmented view space remains.

- While the representation of journey phases completed on public transport were represented in sketch maps using higher levels of schematisation and even explicit compression/ distortion (see figure 67-a), the fact remains during journey planning and during the exercise itself, non-uniform representation of spatial relations caused issues in the understanding of distances and angles.
- A key finding from the S1 studies was that the selection of features was too extreme in terms of completely abstracting away geographical context, as illustrated by the feedback from participant G3\_1 in Table 11. While the focus on features adjacent to the route and within specific regions was well received, the intended reduction of cognitive effort did not result from this approach, with participants exhibiting difficulty in constructing a hierarchical reference. This observation was much clearer in the context of physically observing body language as opposed to the literal transcription of verbal feedback alone, with facial expressions, pauses and repeat interactions all pointing toward difficulty in comprehension where no features were visualised beyond those immediately adjacent to the path of the route.
- As shown in Appendix 2, G3 participants demonstrated both minor and major errors in navigation. The one major error occurred when participant G3\_12 walked past the correct turning and ended up 'off' the map. Together with the observations about the importance of broader geographical context, this points toward the need for content to be provided that is not specific to the current journey, even if attention is directed toward focal regions.

*RE specification point 1) a focus on route-based knowledge given the path of the route for the journey*

- Although it was difficult to evaluate point 1 above, it can be said that observations from this study do not support the complete disambiguation of features adjacent to the path of the route. While high detail topographic features were included, it seemed the Gestalt of the views still requires the conveyance of broader geographical context for the display of features to be meaningful.

*RE specification point 2) higher detail corresponding with increased attention at decision points and density of decision-making and the 'grain' of encoding being primary drivers of the level of schematisation*

- Again, it was difficult to evaluate point 2 given the issues with geometric distortion.

*RE specification point 3) an anchor-based representation with primary anchors serving as the basis of event boundaries in the upper levels of the event paronomy*

- Findings strongly supported point 3 across SR1 studies, with clear indication as to the overall event structure, not just at the level of primary anchors but down to large spatial scales (e.g. road crossings).
- This form of representation was strongly supported across studies. While it is explicit in the data that primary anchors serve to segment the route, it was also inferred from empirical observations that this segmentation is also the basis of determining regions of broader geographical context that serve an important role in comprehension and planning.
  
- *RE specification point 4) metric distortions across the extent, with higher detail local 'charts' corresponding with stable metric representations, and journey phases traversed on public transport leading to metric distortions between bounding regions*
- The slider functionality that meant the user could transition between the detailed view in one focal region to another through a single interaction had some positive responses, however using geometric distortion as a means of showing local detail and global context in the same view seemed to be incongruent with the way that spatial knowledge is naturally developed. The lack of free interaction in view space was one issue in terms of participants' sense of control. However, it seemed that more crucially, this way of presenting the information did not effectively support the development of hierarchical reference memory (Wiener, Schnee and Mallot, 2004), and consequently meant that it was difficult for participants to form a representation of the location entities with respect to the broader hierarchical reference. It would seem very likely that this way of presenting the information also interfered with the perception of a region topology for supporting knowledge around the spatial context of event sequences. These observations will be discussed further in Chapter 9.
- Clearly then a priority for research with an iteration of this prototype stimulus was to assess points 1-3 above without the interference of geometric distortion. Also as discussed, there was a clear signal that the selection was too extreme in terms of removing any features beyond those immediately adjacent to the path of route segments traversed on foot.

In summary, referring to the design criteria with respect to the evaluation of S1:

- 1) *everyone gets to the destination and in almost all cases they do not make any major errors such as boarding an incorrect bus or taking a street that moves them away as opposed to toward the destination*

A critical error was made by participants using S1 in which the researcher had to intervene, and a number of participants who did not make a critical error arguably would have had they not been allowed to switch to a different application at some point during the exercise.

- 2) *that cartographic objects and textual descriptions are correctly interpreted;*

Showing groups of features as ‘islands’ without any broader geographical context was not effective, with a need to support Gestalts that give a sense of the overall geography still being critical, even if a more contextual selection serves to disambiguate features that are of specific relevance to the current journey context.

- 3) *that users feel a sense of control in terms of understanding their spatial location and in their ability to progress toward goals from that location. To this final point, the sense of control can simply be understood in terms of the city feeling manageable;*

While the interaction concept involving the slider that smoothly transitioned between focal regions was well received, restricting interaction such that view space could not be explored beyond this predetermined path was not a design approach that was supported by the evidence. Being able to freely explore and to have control over the position in view space seemed to strongly contribute to the overall sense of control.

The central irony of the S1 prototype stimulus was that in terms of reflecting participants’ *a posteriori* cognitive representations of the journey it performed very well, as shown through the similarity between the prototype content and participant sketch maps, but in terms of the 3 high-level design criteria, it essentially failed. This meant that the initial hypothesis was shown to be incorrect, and provided important insight with respect to research objectives 1a, 1b and 1c. This will be discussed further in Chapter 9.

It is argued the findings over the course of the SR1 studies including the evaluation of prototype stimulus S1 meant that the design specifications detailed in the first 3 points of the S1 specification required further research, but that the specification described in the fourth point with respect to geometric distortion should be removed from S2.

8.2.2 OVERVIEW OF SR2 STUDIES (G4, G5, G6, G7)

In terms of G4 (following) and G5 (Google Maps), SR2 provided the opportunity to collect data for a more complex route observations made during the evaluation of prototype stimuli in G6 and G7. In addition to greater overall length and 3 versus 2 focal regions, SR2 also included two indoor environments and a large set of steps which provided further basis for observing information needs over the course of a journey.

Overview – SR2 sketches

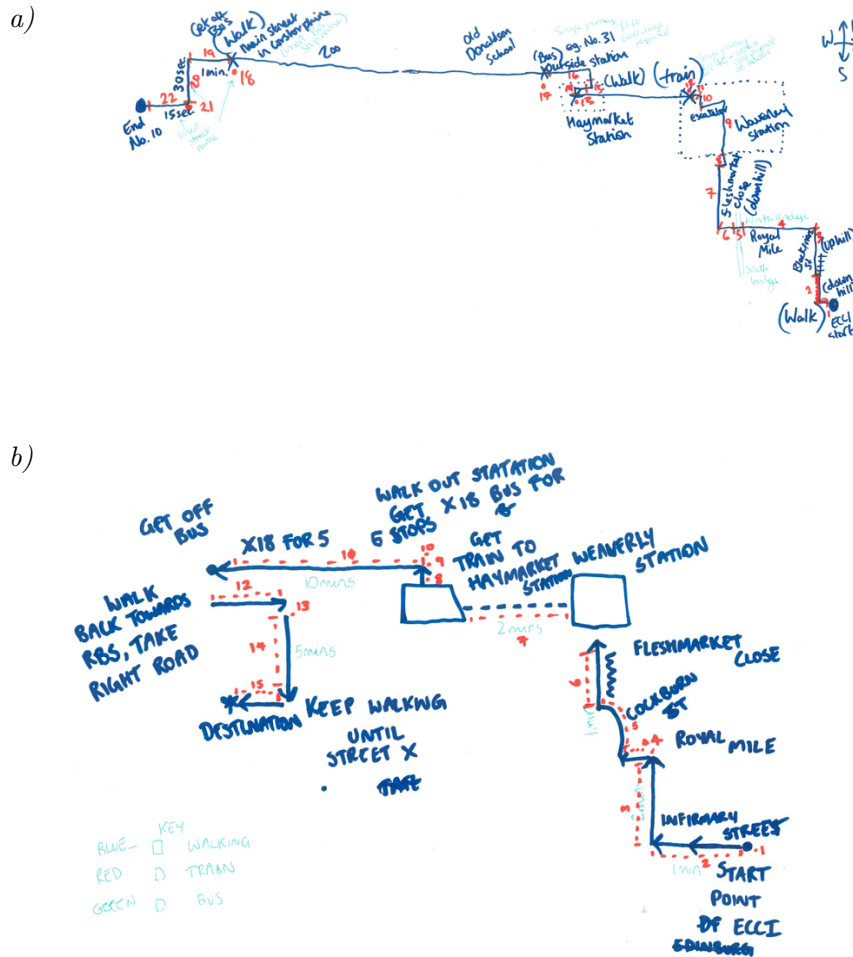


Figure 68 SR2 sketch map examples from participants a) G4\_5, and b) G5\_4 (see Appendix 13.14 for larger illustration).

- While the route itself covered a much larger geographic extent, participants were not provided with a larger size of paper. This naturally led to greater
- schematisation in the representation relative to the overall journey extent and so analysis and findings of SR2 sketch data in comparison with SR1 takes this issue into consideration. See figure 68 for illustrative examples.
- SR2 sketches across G4 and G5 were consistent with the observations from G1 and G2 in terms of the features expressed, the relation between features and the navigable network and the overall basis of regionalisation.
- It can be said that sketches were broadly in line with the content expressed in the S2 and S3 prototype research stimuli, however a key difference was the way participants expressed indoor environments versus the approach that was anticipated. This point will be discussed further in Chapter 9.
- It can be said that sketches were broadly in line with the content expressed in the S2 and S3 prototype research stimuli, however a key difference was the way participants expressed indoor environments versus the approach that was anticipated. This point will be discussed further in Chapter 9.

#### **Overview – SR2 directions**

- Naturally it was expected that the number of directions would increase given the nature of the SR2 route versus SR1. Additional comments about the number of directions are provided subsequently.

#### **Overview – S2 prototype stimulus findings**

- Following the observations from SR1 studies as described, the first three points from the experimental design specification were investigated in a prototype research stimulus without the geometric distortion. Also, additional features were added to support the conveyance of broader geographical context and the acquisition of participant hierarchical reference memory.

*RE specification point 1) a focus on route-based knowledge given the path of the route for the journey;*

- The selection of features in the stimulus strongly correlated with the cognitive representation across G4 and G5. The removal of the geometric distortion, and the addition of survey-level features and the free movement in view space all conspired to support a reduction in negative participant responses.

*RE specification point 2) higher detail corresponding with increased attention at decision points and density of decision-making and the 'grain' of encoding being primary drivers of the level of schematisation;*

- The display of feature selections based on the location of key actions was supported as a design approach, however it was found that the size and shape of the region corresponding with mode transfers meant that the effective 'anchor regions' needed to extend along the path of the route in both directions around these locations, as opposed to a more radial-type selection that characterised the features supporting pedestrian-level actions such as turns. While this was in fact hypothesised to be the case prior to the studies, the central insight here was that if features are to be prioritised in the selection for mode transfers, then features adjacent to 'walk to bus' actions are appropriately the ones prioritised, however features prior to the 'bus to walk' action should be prioritised, meaning a mechanism must be in place for distinguishing features at the extremes of these elongated action regions, with prioritisation in selection then following from the direction of travel relative to the region.

*RE specification point 3) an anchor-based representation with primary anchors serving as the basis of event boundaries in the upper levels of the event partonomy;*

- This approach was observed to be congruent with observed traveller information needs with additional findings around the relationship between information needs and events at different levels of abstraction over the event partonomy. While events formed around actions such as pressing the bell on the bus or purchasing a train ticket, it was not natural for these events to be explicitly expressed in graphical information, as reflected both in participant sketches, and in the extent to which high detail views were surfaced at all in field studies.

*RE specification point 4) feature selection based on a reductive principle but with additional features not adjacent to the path of the route providing a sense of the wider geography;*

- It was found that even with the addition of major parks and the overall city road network, further detail was still required in terms of the overall Gestalt of map views and the perception of regions such that hierarchical reference memory could be developed. In essence the finding was there should be broader detail representing the geography such that in any predicted view there would be a representation beyond simply the road network, even if no specific feature within this additional detail played a direct role in navigation over the current journey context. This is similar to the notion of 'desert fog' (Jul and Furnas, 1998) in the sense that overly reductive

selections that only provide broader geographical context at low levels of detail result in higher detail views with insufficient information to support orientation and wayfinding. . This was evident in comments about a lack of differentiation in the representation of the broader geography (“monochrome”, G5\_6), and in the tendency for participants to switch to the use of Google Maps during navigation exercises.

### **Overview – S3 prototype stimulus findings**

- In S3, the final prototype stimulus, aside from additional low detail text features (city borough names), the cartographic views were consistent from S2.
- As outlined in the previous chapter the key difference from S3 was the integration of textual directions into the views, with directions being nested such that for higher detail map views, the text direction also decomposed into more granular textual directions.
- The prototype was implemented as a series of static ‘click-through’ views however this approach meant there was no marker for the participant’s current location. This caused a significant challenge for one participant, although overall the group displayed the lowest level of errors when compared with other studies. It is argued the reason for this was partly due to the detailed provision of textual directions, however it should also be highlighted that the group had a higher level of familiarity with the city than other groups in the average case. There was one critical error for this group which was due to the lack of a GPS location marker.
- The implementation also restricted participants to only being able to see the predetermined views. Following the observations made during the G3 study with prototype research stimulus S1, lack of control over the map views was highly problematic. While in general this did not lead to errors during the studies, it did require fairly significant effort on the part of some participants in terms of spending time comprehending the sequence of actions required to complete the task.
- In terms of the contextual zoom interaction concept there was evidence from the follow-up directions exercise that this helped travellers to form a hierarchical representation, however to be able to meaningfully reflect an event paronomy in textual descriptions resulted in a large amount of text.

- a)
- ① Walk to Waverley  
 take right out of gate from High School Yards  
 continue along High School Yards  
 cross Cowgate & continue straight  
 turn left on High St  
 cross ~~road~~ N. Bridge  
 turn right down \_\_\_\_\_ road  
 take steps down to Station
  - ② Take train from Waverley to Haymarket
  - ③ Take 26 or 31 bus from bus stop at Haymarket to St Ninian
  - ④ Walk to 10 Featherstone Crescent South  
 walk along St. John's Road  
 Take first left and the second right on Featherstone Crescent South
- b)
- |       |  |
|-------|--|
| walk  | ① turn right after goes out of the Building.<br>pass the Zara Take away and keep walking<br>turn left until you <del>see</del> reached Royal Mile.<br>keep walking and cross the road (only use time)<br>turn right and walk down the road.<br>pass the Boots store and walk down the stairs.<br>Once get down the stairs you will see the Tram station. |
| train | ② take the train to Haymarket station. (only one stop)   |
| walk  | ③ gets out of the train station from the way out which face the street<br>wait for 26 or 31  |
| bus   | ④ after goes on the bus you will pass the M... S... and the zoo.<br>gets off at St N... stop   |
| walk  | ⑤ walk towards the direction that the bus goes<br>turn left when you see the Royal Bank of Scotland.<br>keep walking for 4-5 minutes. and turn right at ... road<br>the destination is on your right, follow the numbers on the house.   |

Figure 69 Examples of the explicit conveyance of hierarchical representations in participant understanding, a) G7\_2, and b) G7\_5 (see Appendix 13.15 for larger illustration).

- See Figure 69 for two examples from G7 where participants showed the formation of hierarchical understanding by chunking actions and event sequences based on a higher-level action or task (e.g. walking to the train station). While the data points towards there being potential for this ‘contextual zoom’ approach that simultaneously affects both the cartographic view and the sequence of textual descriptions, the idea requires for further research (see Chapter 11).

## 8.2.3 ADDITIONAL OBSERVATIONS ACROSS SR1 &amp; SR2

**Additional observations – directions exercise**

In Table 12 the average count of the numbered directions written by participants is shown for each group, and below that, the standard deviation of the count of directions to aid comparison between groups. It is argued that in the average case the numbered directions were an appropriate way of observing an approximation to the event structure developed by the participants over the episode. Following from this argument, the purpose of showing the data in Table 12 is to highlight that overall there was a reduction in the number of directions from the ‘following’ groups to the ‘Google Maps’ groups, and again from the ‘Google Maps’ groups to the groups that were navigating using the prototype research stimuli. Only the 15.36 value for G5 versus the 15 value for G4 is not consistent with this trend, however it seems to be such a small difference that overall the trend would seem to be noteworthy. Given the volume of studies it is not possible to draw clear conclusions from this observation, however it would seem to point toward the possibility that exposure to a larger number of features through perceptual experience would lead to a greater subdivision of action sequences, resulting in a larger number of perceived event for the episode. What an optimal number of events would be for a given journey in terms of successful navigation is not clear, however. It is known that increases in

G1 AV	11.92		G4 AV	15		G7 AV	13.76
G1 STDEV	3.9		G4 STDEV	6.87		G7 STDEV	4.18
G2 AV	9.42		G5 AV	15.36			
G2 STDEV	2.83		G5 STDEV	6.65			
G3 AV	8.67		G6 AV	12			
G3 STDEV	3.62		G6 STDEV	8.32			

*Table 12: Text directions – average count of the directions per study group plus the standard deviation for the group.*

unpredictability can lead to an increase in the subdivision of action sequences, which would imply that groups using the Google Maps application experienced lower unpredictability versus the ‘following’ groups (which would clearly be expected), and also that also groups using the prototype stimuli then experienced lower unpredictability versus the Google Maps groups. While this would be a very positive result from the perspective of design approaches to providing highly contextual navigational information over the episode, the issues encountered with S1 and S3

indicate that perhaps there is another factor at play rather than just a reduction in unpredictability alone.

### **Additional observations – sketch exercise**

As expected, there was a higher level of schematisation overall in the sketches produced for SR2 versus SR1. In terms the physical entities that were expressed, there was very little evidence for the need to include entities that are not already present in either OS, OSM or Google products. Examples that did appear related to features that played a role in physical path traversal such as the representation of the physical appearance of road crossings and distinctive road surface textures.

An aspect of the sketch exercises that was not anticipated however was an issue with the methodology in terms of participants being directed to talk aloud as they sketched. The exercise was explicitly framed as creating a sketch for someone they would not meet, however on reflection, in asking them to think aloud as they sketched, it was observed that this naturally put participants in the ‘frame of mind’ that the sketch was to be accompanied by explanation. This was apparent implicitly in that on a number of occasions the participant was providing explanation that was additional to the content of the sketch as opposed to simply explaining the thinking behind the elements that were included. In these instances it was necessary to remind the participant that all information conveyed to the hypothetical traveller needed to be within the sketch. This issue was explicitly apparent in that when participants were asked how, on reflection, they felt about a person actually using the sketch to navigate, a common comment was that they felt it needed to be provided in conjunction verbal explanation, despite the initial direction that the user was not known and they would not meet them. It should be noted that this issue was taken into consideration in the analysis.

## **8.3 Concluding comments**

The observations from navigation studies and the follow-up exercises support a framework for selection based around an approximation to an event partonomy driven by actions required in the current journey context and by event sequences across spatial scales.

In terms of the display of information, no evidence was found for explicitly reflecting the distortions that are regularly observed in ‘cognitive maps’ in the display of information itself. While it was hypothesised this may have been a way of

integrating global context with local detail on small screen devices, it is not in fact congruent with the natural acquisition of spatial knowledge, even if though these type of structures characterise spatial knowledge *a posteriori*.

Overall then it was observed that the display of context-specific selections should be driven by the locations of decisions and actions for the current journey context. While it is appropriate to alter the level schematisation given information needs over an event sequence, display and interactivity must still support Gestalts that include broader geographic detail and the flexibility to move freely in view space, even if the set of views accessed over the journey is anchored to the path of the route.



## 9 Analysis and findings

In the first part of the analysis, two fundamental findings are discussed. Firstly, that the central task in automating map content selection should be understood as predicting the state of a coupled system. Secondly, that the informational component can be determined by simulating a resource profile of cognitive extension for the current journey context. The theory and concepts that underpin these findings are explained, with analysis of the empirical data used throughout the explanation. This lays a foundation for the explanation in parts 2 and 3 which go into further detail around the implications of the empirical observations for the structure and content of the underlying graph of features. Part 2 focuses on the empirical basis of forming an explicit relation between journey context and the underlying data. Part 3 then goes on to present the basis of an approach to the context-specific regionalisation of the journey extent, where features that are predicted to support navigational tasks may be clustered across a reconfigurable content-hierarchy.

### **9.1 Part 1: The resource profile of cognitive extension**

In this first part of the analysis empirical data is considered that relates to the nature of participant interactions with information during navigation, and the manner in which these interactions play a role in the acquisition of spatial knowledge. Here the initial hypothesis is considered in light of these observations, and an overall theoretical foundation proposed to support the development of systems that automate the process of map content selection given a basic level of knowledge about a traveller's context.

### 9.1.1 THE DEVELOPMENT OF COGNITIVE REPRESENTATION OVER TIME

Following the explanation in Chapter 6 a ‘view event onset’ is simply the moment in space and time at which a participant begins an interaction with navigational information. These phenomena can be treated as objects of analysis by taking the time from the beginning of a device screen recording and then locating the participant at that time using the GPX data. In Figure 70 these locations are visualised for G2 participants, and in Figure 71, the locations are visualised for G5 participants. These visualisations do not of course show all the points along the route at which a participant was viewing information on the device – simply the locations where need was high enough given the immediate context to begin referring to external information. This is argued to be a sound approximation to the pattern of information needs over time, with vertical coloured lines serving to give a sense of the level of need by reconciling the view event onsets over the study group to a single dimension.

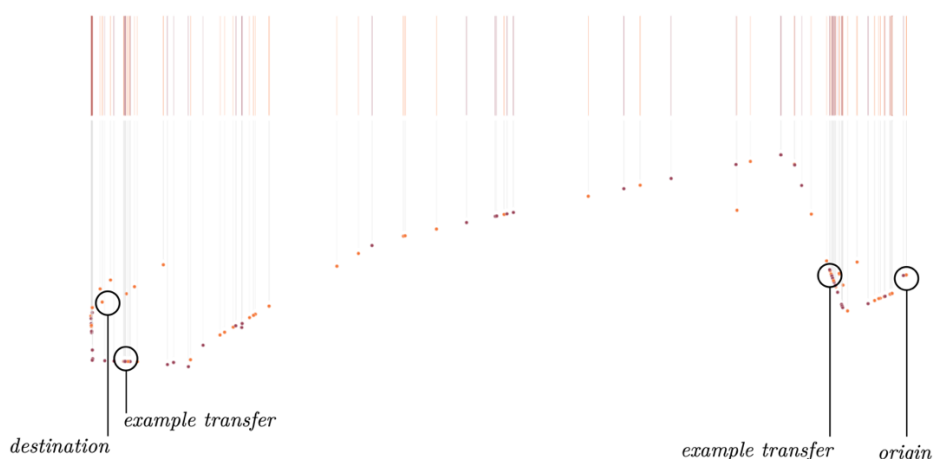


Figure 70 SR1, G2 view event onsets, where colour denotes the study participant.

In Figure 71 it appears there are relatively few view events at the location of the origin but this is simply because the method of visualisation resulted in these onsets coinciding at the exact same location. At the journey destination it shows only one view event began in the immediate vicinity as most participants were ‘in’ view events that were ongoing during physical arrival at the destination. Example transfers from walking to the bus and vice versa are annotated as they were the transfer locations for the most common route taken. Over the course of the bus journey it seems that the start of view events is fairly evenly spread but this impression is argued to be a result of the relatively short bus journey overall. In

other words the journey length serves to slightly obscure the general pattern of interaction discussed in the observations chapter. These view event onsets all relate to checking the bus was travelling as expected, followed by a short switch to environmental cues only, followed by a check on overall progress. A point at which a view event was initiated given a sense that the destination was approaching would then last up until the bus stop had been reached. An observation that this visualisation does not convey is the tendency of some participants to keep the current view event active for very long intervals, for example for the whole bus journey, and then simply checking the location marker versus the target transfer.




Figure 71 SR2, G5 view event onsets, where colour denotes the study participant

In Figure 71 essentially the same pattern of interaction is shown, however the key differences here are the 3 focal regions that result from the additional transfers, and the train journey, during which only one view event was initiated given that the current goal in that scenario is simply the next train station.

While these visualisations show the locations where sequences of views began, in Figure 72 the implied boundaries of image space for the view centres over each sequence of views is shown. This boundary is implied in the sense that clearly the form factor of a mobile device is not square, so these bounding boxes signify the selection that would have need to have been made given a user accessing information for each view centre, as opposed to indicating the actual area visible to the user (following the explanation in Chapter 6). There are 3 regions that are visible in this figure based on overall number of views, and while the visibility of the larger ‘middle’ region is slightly misleading as it includes views that Google Maps defaulted to as opposed to views that participants specifically ‘navigated’ to, it is argued this figure still gives a fair approximation as to the regions and scales at which information was required. Here the content viewed is implied by the scale as

determined by the width of image space (again, following the explanation in Chapter 6).




Figure 72 G2 views, shown by the implied boundaries and centres of the square subsets of world space implied by each map view, hues denote participant.

If we understand scale as being represented in a third dimension, as shown on three-dimensional axes in previous figures, then by using the width of image space as a guide for scale, the transition between views over time given a path between image space centres can be represented through view curves in simplified view space.

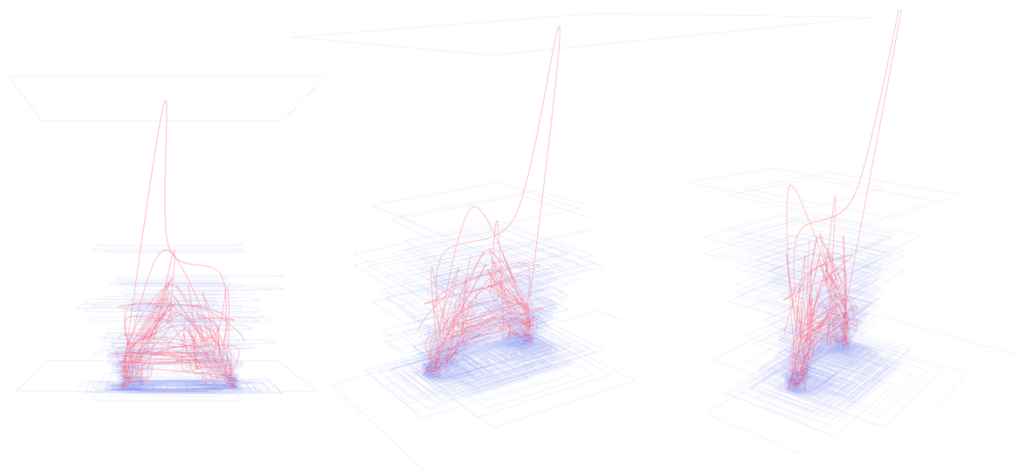


Figure 73 View curves in simplified view space for G2 participant map interaction, with the 3 images simply showing the set of views curves for G2 from different 'camera' angles.

In Figure 73, the path of attention through view space is represented by the red curves, with blue bounding boxes at the ‘bottom’ showing the square subsets of world space implied by the view centre, and the blue bounding boxes at the same ‘height’ as view centres representing the square subsets of view space that correspond with cartographic content. Again, the word ‘implied’ is used here because a) the actual view for a participant is not square even though the selection itself is based on a square area, and b) the cartographic content that corresponds with a given view centre is relative to the content that corresponds with that view centre in the current application. As noted in Chapter 8, even content that corresponds with a particular view centre in the same application may change over time, and so content corresponding to that view centre cannot be understood as necessarily being consistent between journeys.

An important point to highlight here is that while there are consistent patterns of behaviour in terms of the relationship between view curves and the location of the participant, it was found to be highly problematic to assume that only certain view curves would be associated with particular locations. Overall it was found that while it may be more or less likely for a participant to ‘visit’ certain information content from their current location, it should be assumed that any view curve that is likely for the current journey context may be required from any location over the course of the journey.

The view curves that are visualised in Figure 73 show a clear ‘A’ structure for the tripartite journey case, i.e. the case where a journey is characterised by two connected focal regions given two mode transfers (‘walk to bus’ and ‘bus to walk’). This structure conveys overviews of the journey extent at the apex of the ‘A’, the access of ‘mid-detail’ content along the ‘crossbar’ of the ‘A’, and high detail views in the focal regions at the ‘feet’ of the ‘A’. For journeys with a higher number of focal regions, this structure is then repeated, within a new overall journey extent ‘apex’. While it may be the case that not all travellers require all the information corresponding to the paths through view space that lead to this structure, for example not all travellers will access content from the ‘crossbar’ of the ‘A’, given the aim of the present research is to automatically determine content in the average case, then it is argued this structure provides an appropriate guide.

### 9.1.2 REFLECTING ON THE ONTOLOGICAL STATUS OF NAVIGATIONAL INFORMATION RESOURCES

Referring back to Figure 5 in Chapter 3, and the research premise discussed in (Mackaness, Quigley and Tanasescu, 2011), the proposal can be interpreted ontologically to be that the maps generated by the system should be information

resources that are wholly present at a particular point in time. The notion of integrating ‘global context’ and ‘local detail’ in the same view, and the related notion of integrating content from various levels of abstraction, both imply that extraneous cognitive load may be reduced by bringing all relevant information into the same experiential ‘moment’.

As noted by (Willis *et al.*, 2009), a comparison of spatial knowledge acquisition between paper and mobile maps showed that mobile map users suffered increased attentional requirements due to the ‘piecemeal’ presentation of the information, and that mobile device users acquired a more ‘fragmented’ cognitive representation. While the proposal outlined in the opening paragraph of this section can be understood in terms addressing these issues through context-specific selections based on the ‘event horizon’ conceptualisation discussed in Chapters 2 and 3, the fundamental issue noted by Willis *et al.* remains. The limits of screen size mean there is no way to meaningfully visualise the content necessary for the successful completion of a journey ‘all at once’. If we consider the content being accessed by participants at the ‘feet’ of the ‘A’ in Figure 73, i.e. high detail views of a localised region, were this content to be integrated into a broader view of the overall journey consisting of varying levels of schematisation to reduce the overall amount of information, at one point in time to be able to perceive and comprehend the content in one of these views, given visual acuity, still requires the user to pan and zoom such that the content ‘fills’ the screen at the expense of showing any broader geographical context, irrespective of the level of schematisation of other regions at that point in time. Referring to Figure 42 that shows an example view from the S1 prototype research stimulus, this illustration demonstrates the extent to which literally visualising content corresponding with a manifold over the underlying set of features leads to geometric distortion to the point where the meaning of the ‘global detail’ is quickly lost as the user transitions to view the focal region. There was no evidence from the empirical studies that reducing the number of distinct views supported knowledge acquisition, and in fact, it seemed to interfere with the natural way that participants developed their knowledge of the environment.

Central to these considerations is the role of interactivity. Returning to the proposal in (Mackness, Quigley and Tanasescu, 2011), in addition to implying that the information would exist wholly at a given point in time, the notion of integrating information from various levels of abstraction to reflect the overall journey in a single view also implies little to no interactivity. While an initial implementation was developed, again as shown in Figure 5, this output is assumed to be a static view, with no additional layers of information available to the user. In terms of the process of developing spatial knowledge illustrated by the figures in section 9.1.1, it was found that in fact interactivity itself seemed to play an important role, for example through participants repeatedly zooming in and out on the same area.

Rather than these interactions indicating a lack of clarity in the information as was hypothesised prior to the empirical studies, it was in fact as if these actions served to help construct and reinforce hierarchical reference memory over time (Wiener, Schnee and Mallot, 2004), as well as helping the participant to develop an understanding as to the location of discrete features within the global representational structure (Marchette, Ryan and Epstein, 2017). In observing participants, it seemed as though their understanding was being refined over time, with a loose, unstable form of knowledge gradually solidifying into a stable representation through which the location of discrete objects could be known within the overall global reference. Evidence was not found to support the presentation of a single static view that would be able to somehow result in the same knowledge as that developed through situated interaction over time across a larger number of views.

In summary, it has been found that while mobile devices lead to issues with the ‘continuity of attention’ (Oulasvirta *et al.*, 2005), the fundamental constraints of the physical form factor mean that the affordances of the device must be negotiated as opposed to being somehow circumvented. Referring to the figures in section 9.1.1, participants accessed views corresponding with many different representations of the underlying geography over the course of the journey, echoing the finding that mobile devices seem to discourage detailed planning through a small number of interactions in favour of developing knowledge gradually over the course of many incremental interactions (Willis *et al.*, 2009). It is argued the incremental nature of knowledge acquisition through small devices is a fundamental affordance that must be central to the design of methods that shape data selection and display. Arguably, the overall navigation strategy that follows from incremental knowledge acquisition does in fact support the reduction in cognitive load as information can be called upon in support of attaining a goal for a particular event sequence, as opposed to memorising information to support navigation across many event sequences, or even the whole journey. While the impact of the incremental approach has been found to be less effective in terms of acquiring configural knowledge and developing long-term memory of spatial environments (Ishikawa *et al.*, 2008), referring back to the 3 design objectives, these are not however central objectives for the user in the general case.

Returning to the present aim then in terms of automatically selecting and displaying information to support urban navigation, it was observed that mobile cartographic information can be understood in terms of the description of emergent cartography in (Kitchin and Dodge, 2007). As previously discussed, the notion that cartographic information emerges through practice as opposed to being distinct from the situated actions of the user. It should however be noted that while the description of maps as being ontogenetic in nature, i.e. as ‘emerging through practice’ (Kitchin and

Dodge, 2007), fits with the empirical observations made during the navigation studies, as we are concerned with specifying the entities and relations that may be formally represented in a system, it is not appropriate to reject navigational information as having an ontological status. Here then it is argued that navigational information ('the application') is productively understood as being a perdurant entity (Bittner, Donnelly and Smith, 2004), where features that form the content associated with this entity at a point in time play roles in situated, extended cognitive processes (Clark, 2013). That is to say the 'map' does not exist wholly at a single point in time, but unfolds in 'successive temporal parts' (Bittner, Donnelly and Smith, 2004).

An extended cognitive process is a process whereby the use of an external resource (a map, a compass and so on) has an impact on cognitive performance such that in terms of study, the resource can itself be understood as part of the cognitive process. According to this view, objects in the underlying geospatial database should be understood as being potential resources (De Landa, 2013), where their presence in the database is determined by the possibility that they may play a role in an extended cognitive process relative to navigation in general, and their selection is determined by the possibility they may play a role in an extended cognitive process for the current journey context specifically. Before discussing this point further it is necessary to elaborate on the analysis of how these information resources can come to be associated with a given journey.

### 9.1.3 PATTERNS OF EVENT STRUCTURE

In terms of the traveller's perceptual experience, journeys can be understood as being a sequence of behaviourally significant events, where events are comprised of tightly grouped action units (for example 'walking on cobbles', 'riding the escalator', 'sitting on the bus'), and overlapping events serve to bind distinct event sequences into a broader representation. Determining the informational resources that will be required by a traveller can then be understood in terms of matching information to needs associated with actions, events and event sequences for the current journey.

In terms of the cognitive representation of the traveller, to describe an event structure is to describe the structure of the episode given the segmentation of perceptual experience that has already occurred – i.e. after the fact, once the final event in the episode has passed. It is argued however, in terms of the scope of the present research, that while this representation is developed and refined over time, its formation can be understood as beginning during planning. A visually salient feature may serve as an event boundary cue given the actual situated experience of the traveller, however overall, particularly at the upper levels of the event

partonomy, primary actions serve as a predictable basis for the meaningful segmentation of experience. Here primary actions refer to the critical actions that support overall goal acquisition including transfers to and from public transport, and of course, the destination itself (refer to chapter X<Discussion> for discussion around the implication of alternative use cases for the segmentation of experience).

In terms of how this structure forms in practice, building on the previous discussion around the incremental nature of knowledge acquisition, it may be that during this planning stage, the ultimate destination is virtually ignored. Overall, based on empirical observation, there can be said to be 3 high level patterns of event structure that begin to form at this initial stage, even if the overall event structure coalesces into a representation that is largely stable across travellers *a posteriori*. Either the participant planned ‘to the end’, exploring the information such that an initial level of understanding was established for the whole route; or, they would plan to an intermediary decision point that corresponded with a primary action, in most cases the second transfer point for G2, for example; or, they would engage with the information in a very superficial way such that essentially no aspect of the journey was committed to memory, and then proceed to simply follow the location marker.

Two aspects of this should be emphasised. Firstly, that irrespective of how knowledge about the journey was initially committed to memory, upon completion a high-level segmentation based on primary actions was interpreted as being the stable basis for event structure in all cases. Secondly, that the selection and display of information should be of a form such that users’ information needs are supported irrespective which primary action served as the basis for the target event during planning (i.e. whether they plan to the ultimate destination or some intermediary location). Fundamentally, planning to an intermediary decision point can be understood as a key strategy for reducing cognitive load and making the environment *feel manageable*, and so navigational information that flexibly allows the user to focus on sub-tasks below the level of the overall A–B journey must be supported.

#### 9.1.4 TOWARDS AN EVENT HORIZON FOR INCREMENTAL KNOWLEDGE ACQUISITION

Returning to the initial hypothesis, the logic was that conceptually speaking, if an artificial agent were to traverse a virtual environment that effectively represents the underlying geography, and the agent develops spatial knowledge during this process in a way that mimics the way the person would have developed their own cognitive representation over the course of journey completion, then the context-specific representation formed by the agent would be approximately isomorphic to the

representation that would have been developed by the person. If the selection of features from an underlying database was equivalent to the second case, then displaying this subset of features to the traveller a priori would be approximately equivalent to the person acquiring context-specific spatial knowledge through physical route traversal.

As discussed in Chapter 8, the S1 prototype stimulus performed well in terms of approximating participant understanding of the journey in the way expressed through sketch maps and textual directions. This claim stems from the conveyance of route-based knowledge with clear emphasis on decision points, greater schematisation for route segments with lower density of decision-points, and the regionalisation of the overall extent such that focal regions of higher topographic detail were bound into a looser global structure with metric distortions between focal regions.

The thrust of the initial hypothesis was that by predicting this type of representation in advance, a natural form of spatial knowledge could be learned prior to the process of physically traversing the route – essentially a ‘short-cut’ to naturalistic spatial knowledge for the given journey. However, a key question following the field observations was that if it was understood after the fact that segments of a journey did not need to conform to absolute distance or direction in the sense of relaxing metric relations, then why would this matter a priori? It is argued that this type of integrated, distorted view that includes both global context and local detail expresses the situated understanding of a traveller who has completed the journey, but is not congruent with information needs over the course of navigation itself. While clearly it is possible for sketch maps to be used for navigation, the context in which this scenario would actually take place is when a person is explaining the journey to a traveller, using verbal and body language, with the sketch playing a role within this broader information context.

This finding and the general performance of the S1 prototype stimulus naturally prompted a need to revisit the initial hypothesis. If information needs cannot be automatically determined by approximating the *a posteriori* representation of the journey, even though it has been shown doing so can in fact be achieved via methods that could reasonably be automated, then what does the study data indicate is the solution? Reflecting on the notion of navigational information as being ‘processual’ and ‘emergent’ as discussed in Chapter 5, the evidence strongly points toward the selection and display of map content as needing to support the incremental acquisition of spatial knowledge over time, as opposed to integrating multiple forms of representation into a single view that wholly exists at a single moment in time.

Linked to this is the issue of interactivity. While a hypothesis prior to empirical study was that a reduction in the number of device interactions would be an indicator of more effective information, no evidence was found to support this, and in fact, participant behaviour during field exercises strongly conveyed that interactivity played an important role, both in participants' overall sense of control, and in the actual process of developing memory of a representation across spatial scales.

Both of the findings discussed above – the need to support incremental knowledge acquisition and the need to support interactivity, imply that the automatic selection and display of map content must be throughout interactive view space such that a user is able move through view space relative to current information needs (i.e. zoom, pan). In S1, the slider control that transitioned between focal regions was equivalent to a fixed view curve across fragmented view space. Referring to Figure 57 however, it is argued that free movement and smooth curves in consistent views space is optimal given the aim to support the incremental acquisition of hierarchical reference memory, and knowledge about discrete entities within the overall global reference.

So, we can see that the problem is in fact not one of determining a single integrated representation of the environment, but of determining a global structure comprised of a set of connected 'submaps' or 'local charts' that can be explored freely. Here free exploration does still occur within a range that is constrained by the current journey context, with the visualisation of the 'A' structure from the G2 data in Figure 73 giving a sense of such a range.

While on the one hand these findings simplify the approach to displaying information as the standard pan and zoom interactions may be utilised, this significantly complicates the process of determining a selection over an underlying graph of features as the 'event horizon' for the current journey context has to support any view the user is predicted to require throughout multiscale view space. While the standard 'tile server' method of providing hierarchically ordered representations in metric space does provide a basis for submaps bound into a global structure that may be freely explored, a fundamental question is how features may be selected such that, not only can they be associated with the current context, but also their 'locations' in multi scale view space can be known. In other words, beyond determining whether feature  $v$  is relevant to the process of navigation in the current journey context, over what range of scales should feature  $v$  be visible?

### 9.1.5 AGENT, APP AND ENVIRONMENT AS COUPLED SYSTEM

The need for an informational resource arises out of situated interactions between the agent (the traveller), information (the application), and the ‘external’ environment. Here a distinction is made between information that may feasibly be stored in the geospatial database and accessed through the device, and information that is embedded in the broader environment (for example signage, landmarks and so on). Here then, ‘environment’ is used to describe all information available to the agent in addition to the information available in the application, i.e. the information that is not under our control but nevertheless plays a role in spatial goal acquisition. Exposure to both information and environmental cues can cause a change in the state of knowledge of the agent, and a change of spatial position. This new position may then serve to prompt the user to seek more information, to validate expectations versus the new set of environmental cues, and so on.

This observation can be understood in terms of the relation between a human agent and an external resource as a ‘coupling’ (Heersmink, 2015). Here ‘resource’ simply means any entity that plays a causal role in processes that impact on cognitive performance. Resources can then be understood as being ‘recruited’ into cognitive processes (Bjørndahl *et al.*, 2019), with resources being utilised in situated epistemic actions (Kirsh and Maglio, 1994), i.e. actions that change the informational state of the agent. While ‘pragmatic action’ serves to change the physical state of the agent, for example ‘walking from the corner to the crossroads’, epistemic action serves to change informational state even if no change of spatial position has occurred, for example the ‘zooming in/ zooming out’ behaviour discussed previously. The agent and the ‘app’ can be understood as being coupled as the informational state of the agent causes a change in the state of the informational resource, which in turn causes a change in the informational state of the agent. This view echoes the classic study of plans as being derived from situated action (Suchman, 1985). Environmental cues play a key role here also, and so overall the coupled system can be understood as comprising of 3 high-level ‘components’, the agent, the entity through which informational resources are accessed (the ‘app’), and the overall environmental context. The cognitive representations that the human agent then develops by participating in this coupled system can be understood as being ‘local’ and ‘action-oriented’ opposed to ‘objective’ and ‘action-independent’, following (Clark, 1998).

So, through treating informational objects as resources in extended cognitive processes during incremental knowledge acquisition, it became apparent over the course of empirical observation that the central task in predicting the information needs of the user is not to approximate an a posteriori cognitive representation of the journey, following the initial hypothesis in chapter 3. It is argued in fact, that

given the evidence, the central task should be understood as predicting the state of a coupled system. While the focus of the present research is the navigational information itself, to be able to effectively automate the selection and display of map content, it is necessary to develop a solution that takes into consideration the ‘active causal role’ of the agent themselves (Clark and Chalmers, 1998), as well as the role of the environmental context.

To elaborate on this, and following the Resource Competition Framework proposed by (Oulasvirta *et al.*, 2005), a resource-based view of attention and extended cognitive processes in urban navigation can be understood in terms of the interacting components of a coupled system as described. Limitations in the agent’s attention and memory shape the nature and efficacy of these interactions, with the informational state of the agent changing over time. Here ‘task difference’ (Oulasvirta *et al.*, 2005) leads to variation in the demands of cognitive resources, which in turn shapes the resources required in terms of informational artefacts and environmental cues. Here attention can be understood as ‘distributed’ across device and environment, with a central issue for the present research being the determination of resource allocation relative to goals, the overall motivational hierarchy, and the situated constraints of cognitive processes (Navon and Gopher, 1979).

Here two illustrative examples from empirical studies are provided that ran counter to expectations in terms of a resource profile. The first concerning information that supports navigation in indoor environments, and the second being the development of region-level understanding using features at ‘medium’ scales.

To the first example: a hypothesis prior to SR2 empirical studies was that participants would convey their understanding of indoor environments in essentially the same way as the outdoor environment. In other words, it was anticipated that in depicting indoor environments the path of the route, cues adjacent to decision points, and some level of detail with respect to the structure of the environment adjacent to the path of the route would be included (in the SR2 prototype high detail polygon features were included in the representation of the indoor environments, including structures such as raised walkways, steps, escalators, ticket styles, ticket machines and stand-alone building structures).

As illustrated through examples in Figure 74 however, this was not the case, and almost every participant seemingly conveyed indoor environments based a single abstract representation of a CONTAINER schema.

This abstract representation serves to convey the BOUNDARY of the entity (‘the station’), supporting qualitative reasoning with respect to being either within (or

not within) the INTERIOR, and in terms of the ‘face’ of the BOUNDARY. For example in Figure 74 panels ‘d’ and ‘h’, ‘north’ is indicated on the sketch, and so it is conveyed that the transition from EXTERIOR to INTERIOR for the current journey occurs at the south face of the ‘Waverley Station’ entity. In the other sketches this same logic is implicit to varying degrees. In panel ‘h’ it can also be seen that the participant included the path of the route through the station, and

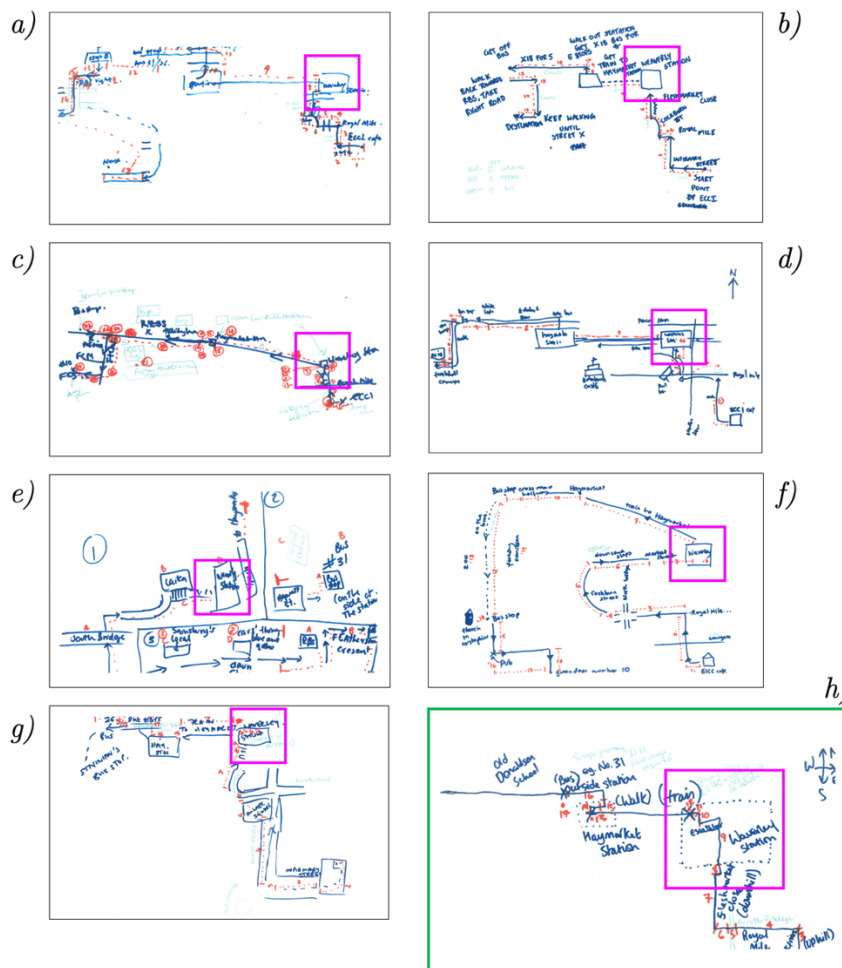


Figure 74 Examples of participant representations of the indoor environment ‘Waverley Train Station’: a) G7\_4, b) G5\_4, c) G6\_3, d) G4\_1, e) G4\_6, f) G4\_8, g) G5\_2; and h) participant G4\_5’s representation of the same indoor environment including the path of the route and the ‘escalator’ entity (see Appendix 13.6 for larger illustration).

also marked the entity ‘escalator’ (while on the one hand this example is an outlier within the overall data set, it can be seen that, as with other cases, the feature that did have primacy in the representation had a functional role in physical path

traversal as opposed to being visually salient per se). That being said, in text directions ‘the departure board’ was conveyed by some participants.

While an argument could be made that the size of the paper meant that the overall size of the sketch led to a space constraint for including high detail features, all participants were asked about this abstract form of representation in the follow-up interview, and the consistent explanation was that the nature of the environment was understood to be such that attention naturally switches to environmental cues, and so including detail in the sketch was not necessary. Here in the indoor environment, signage and overall architecture were understood to be the primary source of information for orientation and wayfinding, not the additional navigational aid (the device). This mental model was apparent in the study with participant G6\_2. The participant had only moved to the city two days previously and had never been inside the main train station before. Upon entering they immediately made a mistake that led to the participant being on a platform from which there was no access to the main area where tickets could be purchased. At no point did the participant reach for the mobile device despite the prototype stimulus showing the exact path of the route all the way to the ticket styles. When prompted to describe what he would normally do in such a situation, the response was to ask a person, not to return to the information on the device.




Figure 75 Participant G6\_2 travels down to incorrect platforms, back up to walkway, then ‘back’ down to main concourse

Clearly, a consideration here is that as mapping applications do not tend to perform well in indoor contexts, there is no established mental model for switching to a mobile device in this scenario – it is not a habitual behaviour (at the time of writing). That being said, based on discussions with participants during follow-up interviews and through observation of orientation and wayfinding in the indoor environments for SR2, the clear conclusion is that attentional resources habitually switch to environmental cues in the indoor case. Here then, it is highly problematic for a design objective of the ‘information component’ to be to reduce the cognitive load of navigation in indoor environments given that there would be strong *competition* between these resources and environmental cues, given the tendency to attend to the environment itself. This example has been highlighted as it provides a clear illustration of how a resource profile shifts over time given transitions

between task-contexts. Even in a case where high detail (i.e. small) entities may be used in support of spatial goal acquisition, the resource profile given the state of the coupled system may dictate that the appropriate resources expressed through the navigational aid may still be at a higher level of abstraction than these entities.

The second illustrative example of an unexpected finding with respect to features and level of abstraction concerned features from ‘medium’ levels of detail that were not adjacent to the path of the route. Here ‘medium’ is used to qualitatively describe scales in a range below the scales close to the effective tile 0 of the journey extent, and above scales close to the point at which a focal region ‘fills’ the screen. While observations made during the study with prototype stimulus S1 pointed toward a need for additional map content to support the development of hierarchical reference memory and knowledge of ‘global’ region connectivity, it was anticipated that given features to support route-based knowledge were included relative to primary actions, the participant would not require broader geographical context beyond the street network. Here the aim was to support the user in attending to only the objects that were predicted to specifically play a role in navigation given the current journey.

The additional low-detail features included the city’s major parks and the road network across the journey extent. Given this selection, even in conjunction with the set of context-specific features that were the primary focus of the study, users would still quite easily execute sequences of device interactions that led to the road network as being the only feature type in the view. This was a source of frustration, with the need to develop knowledge based on the perception of regions that extend beyond the immediate path of the route being hampered by lack of cartographic detail. Oulasvirta et al. described a ‘resource depletion penalty’ (Oulasvirta *et al.*, 2005) whereby insufficient cognitive resources lead to tasks being ‘slowed-down’, ‘postponed’, ‘put on hold’, or ‘terminated’ (Oulasvirta *et al.*, 2005). Here however, the issue was not in fact insufficient attentional resources but insufficient informational resources. In terms of the resource profile then, it was found that even if entities that are of specific relevance to the current journey are displayed, for example features adjacent to turning actions, there is still a requirement for broader map content such that a user does not arrive at a ‘naked’ view (Casati and Varzi, 1997), or a view in ‘desert fog’, following the explanation in (Jul and Furnas, 1998). It is argued this example serves to highlight the challenge of determining traveller information needs. The kinds of features that are implied by this broader region-level selection were not conveyed at all in sketch maps, and while this emphasises the primacy of the entities that were conveyed, it also serves to demonstrate that information needs may not be consciously understood and therefore difficult to ‘uncover’. It was only by removing detail from prototype stimuli that the effect of these oversimplified representations could be investigated.

### 9.1.6 CHARACTERISING RESOURCE PROFILES

Phenomena that play a role in extended cognitive processes can be understood as resources in that they are ‘recruited’ into these processes as a means of improving cognitive performance. The set of these resources for a given journey context will be described as a resource profile. This profile can be understood in terms of the resources required over time, given a sequence of states in the coupled system as described in the previous section. The present research, when viewed through this lens, can be understood in terms of determining, given a predicted resource profile, the appropriate selection of features from the underlying database, and the form of their display. Here the selection aims to support the incremental acquisition of spatial knowledge. While informational resources will be discussed subsequently, for clarity it should be highlighted here that resources in terms of the underlying database refer to objects representing cartographic features from varying levels of abstraction.

A resource profile should be understood in terms of a characterisation of high-level states, where likely shifts in attention are implicitly represented, even if the specific environmental cues that will come to play a role in the lived experience of the traveller are not known. Returning to the overall epistemological principle of the ‘rational consequence’ of knowledge, it should be emphasised that the notion of predicting the state of a coupled system is meant only in so far as it is practical in the sense of supporting the automatic selection and display of navigational information.

### 9.1.7 PART 1 SUMMARY

Part 1 has presented analysis that provides a basis for further specifying the content and functionality of a system given the central aim of the research. Findings have supported previous studies that concluded the acquisition of spatial knowledge on mobile devices is by nature incremental, with attention naturally being divided over a large number of interactions over time. While field observations supported the overall aim to provide context-specific selections, interactivity and free movement in multi scale view space were also found to be fundamental to the efficacy of delivering spatial information through these small devices.

In terms of the event horizon over the underlying graph of features, the implication of these findings is that, as opposed to the horizon being equivalent to a single integrated view of the geography, it must in fact support any view the traveller may require as knowledge is developed over time.

The argument from the situated cognition research that cognitive processes that productively utilise external resources should be understood in terms of cognitive extension provides a practical foundation for considering coupled systems as the subject of enquiry, with the aim of making explicit the nature of couplings between the agent and informational resources such as topographic representations, schematic representations and textual descriptions.

In Part 2, the analysis provides further explanation as to how an explicit relation may be formed between a given journey context and the underlying database.

## **9.2 Part 2: Representing the structure of experience across navigation episodes**

Part 2 begins with an analysis of the initial research proposal from (Mackaness, Quigley and Tanasescu, 2011) with a focus on the basis for forming relations between an urban journey and a set of informational resources that is evident from the initial implementation. A set of event types is presented based on research findings, and an approach developed in which predicted events may be located on the route network, with event-schemata associated with these events then serving to determine a resource profile, following the explanation in Part 1.

### **9.2.1 REFLECTIONS ON THE INITIAL RESEARCH PREMISE**

Returning to the proposals in (Mackaness, Quigley and Tanasescu, 2011), the integration of different forms of representation was conceptualised as being a three-dimensional space, defined by the axes: a) thematic emphasis, b) level of abstraction, and c) resolution. Thematic emphasis was conceptualised in terms of the degree to which, for example, connectivity is emphasised versus more complex Gestalts. Level of abstraction was conceptualised in terms of a continuum between topographic and schematic representation, and resolution was simply conceptualised as level of granularity. In the example shown in Figure 76, features were modelled in a graph database along with a representation of the route, with shifts in the form of representation being triggered by transitions between transport mode. Here we see that a transition to the bus leads to a reduction in topographic detail, with bus stop information being surfaced along traversed route segments. Also it should be noted that even along the sections of the route traversed on foot, the selection of topographic features seems to be aimed at supporting egocentric route knowledge as opposed to providing detail about the broader geographical context of the

example journey, i.e. through selection appearing to be based on route-adjacency or some other measure of proximity to the current route.




Figure 76 Disambiguation of representations in example output from (Mackness, Quigley and Tanasescu, 2011), including low detail selection across the extent (a), topographic features adjacent to the route (b, d), and a highly schematised view of the bus journey (c)

In Figure 76 we can see that while topographic features are in the selection, this selection seems to be determined by the path of the route as opposed to being uniform across the extent, and is therefore conceptually driven by a shift some way along the continuum between topographic abstraction and a more highly schematised form of representation. For this approach to be developed further a number of fundamental questions had to be explored.

The conceptualisation implies that a coherent view corresponds with every point in the space, and that to select a view that reflects the context-dependent information needs of the traveller, contextual data about the traveller must somehow form the basis of determining the point in this space. It also reflects the explicit ambition that if context changes, the vector from one point in the space to another would correspond with a smooth transition between views.

Questions that logically follow then, if this rough conceptualisation (thematic emphasis, level of abstraction, resolution) is to be formalised in practice, based on our knowledge of the context-dependent information needs of the user, are as follows:

1. Are these the correct axes in terms of the nature of abstraction and in terms of the number of dimensions? Is this in fact a multidimensional space beyond  $\mathbb{R}^3$ ?
2. Should this space be continuous, discontinuous, or should the axes express both continuous and discontinuous properties?
3. Does a unique view correspond with a point in this space, a region of the space, or does a point in this space in fact only appropriately correspond with the state of an individual feature or targeted subset of features? If the latter is the case, then by what means might the transition vector corresponding with the representation of a feature constrain the transition vectors of other features given the need for features to work together to form coherent views?
4. How might we specify mappings from data about the context of the user to a point or region in the space?
5. How can the results that are expressed by a point or region in the space be constrained such that they conform to the underlying rules of information visualisation (visual acuity, complexity, contrast, gestalt etc)?

In the example shown in Figure 76, it would seem in fact that only two of the dimensions are represented. Firstly, thematic emphasis as a continuum between the detailed structure of the geography and the information that is specific to the bus network. Secondly, the continuum between topographic and schematic representations. It is apparent then that both purpose and content explicit in these two dimensions are in fact deeply entangled, and so beyond a rough conceptualisation of a vector through a space as described, there is an outstanding question as to what specifically would be the basis of each continuum and what content would be tied to each. Here then it seems it could theoretically be the case that the same content is tied to more than one dimension.

The implementation that underpins the example shown in Figure 60 simply generates single static views, with the connection between context and content seeming to be based on the path of the route and transitions in transport mode. In both cases then it would seem that discrete representations form the basis for this implementation. In fact, it would seem there are only three discrete possibilities for content selection in this example: either a low detail selection as shown across the extent in the illustration (presumed to be point features in the output shown in Figure 76 box 'a'), the high detail topographic features surrounding route segments traversed on foot (Figure 76 boxes 'b' and 'd'), or the more highly schematised representation of the bus route (Figure 76 box 'c').

Following this analysis of the approach underpinning the example in Figures 60, we can see this implementation does not include 'intermediary representations'

(Dumont *et al.*, 2017) between content at different levels of abstraction. The view is static with no ability to transition to a different scale through a zooming action, and within this static view selection can be from one of the three distinct forms only. If ‘cognitively smooth transitions’ were to be supported by continuous processes such as deformation, aggregation and interpolation then clearly we would also need to specify the nature of intermediary representations. More generally however, even if transitions are to be between discrete representations of features, there are a number of considerations to explore such as whether features may ‘coexist’ or whether they may be ‘superimposed’ (Dumont, Touya and Duchene, 2016). In other words, whether multiple levels of abstraction for the same feature type may be present in the same view, or whether the same object may be represented simultaneously at different levels of abstraction in the same view.

Over the course of the empirical studies, through validating and refining techniques, it was found that an adaptation of the routing network is a powerful basis for modeling and selection. In (Mackaness, Quigley and Tanasescu, 2011) it seemed that the points of transition between transport mode served to trigger transitions in the form of representation. This approach has been developed such that an adapted routing network serves as the foundational for the overall framework presented in Chapter 10. The remainder of Part 2 presents the empirical basis for this argument in detail, drawing on the conceptual background presented in chapter 2.

### 9.2.2 GOAL HIERARCHY AND EVENT PARTONOMY IN URBAN TRAVEL

As discussed throughout this thesis, events can be understood as comprising the overall structure of experience, with the subdivision of activity into distinct events and event sequences being driven by goals at various spatial and temporal scales. In multi modal travel, the transition between modes of transport is a distinctive event, in which the agent has clearly changed state. Beyond this form of state transition however, how might we be able to draw explicit linkages between knowledge about the traveller and content in an underlying database that will support information needs?

While some features or regions may be particularly salient and as such naturally serve to segment experience through the formation of perceptual boundaries, in almost all cases in the field studies, it was the structure of the network that served as the basis for event segmentation and sequence formation.

In terms of the network itself, the Ordnance Survey Highways (Ordnance Survey, 2018) routing network was used as the initial representation. The adaptation of this data structure will be discussed further in Chapter 10.

### 9.2.3 DIRECTED PATHS AS THE FOUNDATIONAL BASIS OF COMPREHENSION AND REASONING

Across both sketches and textual route directions, directed paths were shown to form the fundamental basis of comprehension and reasoning. In terms of the ordering of actions in sketch exercises, the conveyance of the path of the route through a sequence of linear objects had clear primacy and served to anchor the broader representation. While certain landmarks served to anchor the perceptual regionalisation of the environment as expected, what was striking was the extent to which the structure of the network served as the primary basis for anchoring. While there was evidence that spatial characteristics of the road network played a role as discussed in Chapter 8, for example through driving event formation through distinctive changes in road geometry, overall the topological structure of the network was by far the most significant factor in event perception and sequence formation.

The directions exercise served to make explicit the participants' knowledge in terms of changes in spatial resolution and the implied schematisation of these directed paths. It is argued the subdivision of activity can be understood as being driven by primary, secondary and tertiary actions. Primary actions are critical to overall A-B route completion and serve to shape event formation over large spatial scales. So, in the context of multi modal travel, these include mode transfers and the destination. Secondary actions also serve to progress the agent toward a goal but drive event formation at much smaller spatial scales, for example a turning action at a network intersection. Tertiary actions occur between either two secondary actions or a secondary and a primary action, and serve to progress the agent toward a goal but do not serve to trigger a change to a new sub-task – i.e. traversing an intersection which is not the location of a turning action.

The level of schematisation was 'implied' in the empirical data in the sense that, although clearly all participant knowledge could be said to be schematised, variation in the level of schematisation was, as expected, driven by the density of decision points over the spatial extent.

Returning to the event horizon conceptualisation, the empirical observations are argued to show that a hierarchically ordered sequence of directed paths serves as the natural way to represent an event partonomy over the episode in the average

case. Here hierarchical orderings can be understood in terms of the primary, secondary and tertiary actions as discussed. This links with the articulation of regions as ‘path convex sets of places’ (Remolina *et al.*, 1999), where actions that serve as the basis for distinct behaviourally significant events correspond with ‘places’, and sets of actions of lower perceptual significance are grouped ‘within’ these path-based regions (i.e. regions formed through a spatially distinct ‘source’ and ‘target’).

#### 9.2.4 EVENT TYPES

The set of behaviourally significant events is fundamentally tied to use-case, with the focus in the present research being the general case of multi modal urban travel. Here then, a clear set of event types were confirmed, with each being appropriately represented within a route network data structure. What should be highlighted here is that these event types are not the only event types observed over the course of empirical studies, but rather, the event types that should be understood as directly corresponding with information needs given a resource profile.

##### ORIGIN

**Definition:** the journey origin

##### DESTINATION

**Definition:** the journey destination

##### TRANSFER

**Definition:** a change of transport mode

##### TURN

**Definition:** a pedestrian turning action

##### INTERSECTION

**Definition:** a traversal of an intersection that is explicit in the ordered sequence of routing nodes

##### CROSSING

**Definition:** a transition from one side of a street to the other

##### LEVEL-TRANSITION

**Definition:** a distinctive transition between levels (steps, escalator)

## ENTRANCE-EXIT

**Definition:** crossing a boundary between the INTERIOR and the EXTERIOR of an entity

## CHECK-POINT

**Definition:** an event comprised of a sequence of actions on the public transport network using information and environmental cues to determine either a) that the current route is correct, or b) the proximity to an upcoming transfer

In Figure 77 the locations of the start and end of textual descriptions of the journey for G2 participants are shown for two example regions. Here it can be seen that, if road crossing points are included as nodes in the underlying network representation,



Figure 77 Patterns of event segmentation inferred through G2 directions exercise textual data, overlaid on participant GPX traces and the OS Highways network (Ordnance Survey, 2019) 1a, turn; 1b, crossing, turn; 1c, crossing; 1d, crossing; 1e, crossing; 1f, crossing; 1g, transfer; 2a, transfer, crossing; 2b, crossing; 2c, turn; 2d, origin.

then event segmentation based on textual directions is effectively represented by the structure of the network. Groupings consisting of 3 locations, shown for example in 1d and 2a, show where the crossing was either communicated in terms of the road centreline or was communicated in terms of the explicit transition between one side of the street and the other. Overall, the significance of the CROSSING event type in the data points toward the need for the network representation to include both the centreline at small scales, and the navigable paths along both sides of the street at larger scales.

In terms of TRANSFER, TURN, INTERSECTION, ENTRANCE-EXIT and LEVEL-TRANSITION, these event types generalise across both indoor and outdoor environments.

Applying the event horizon conceptualisation to urban navigation was intended as a means to support the selection of content across a context-specific set of behaviourally distinct events, irrespective of the spatial scale at which those events occurred. The 9 event types defined above are presented as those that, not only may be explicitly determined for a given journey context, but also correspond with resource profiles given information interaction through small mobile devices. While it was anticipated that the research would lead to techniques for ‘mixing’ the representation of entities from both large and small spatial scales, the physical affordances of the mobile form factor mean that some of these smaller entities do not naturally fit within the resource profile.

The 9 event types presented above then form the basis of event sequences based on their ordering over the journey. Hierarchy can be formalised in terms of these sequential orderings. Clearly the highest-level event sequence is ORIGIN-DESTINATION, with this sequence decomposing into the sequences formed by the ORIGIN, the DESTINATION, and each TRANSFER for the current journey context. These sequences then decompose all the way ‘down’ to the level of CROSSING-CROSSING and INTERSECTION-INTERSECTION. This should be understood as the formal basis for representing event paronomies. The overall approach is highly flexible however, and any event type that could reasonably be represented as a node in a route network could be added given the use-case. This point will be addressed further in the discussion.

### 9.2.5 EVENT-SCHEMATA

The analysis of both sketch and directions data supported the development of event-schemata such that the role of entities in experiential events could be made explicit. Following the explanation in the methodology, this analysis used ‘extracted’ image-schemata, whereby the representation of geographic entities could be understood in terms of the event types as presented in the previous section. As discussed, image-schemata and compound image-schemata can be understood as the ‘preconceptual’ structures that characterise repeating patterns of experience. In terms of attention, ‘preconceptual’ structures can be understood as drivers of endogenous attentional orientation, as discussed in Chapter 4, where repeating patterns of experience are associated with successful goal acquisition across episodes. In terms of urban navigation then, this view provides the theoretical basis for interpreting empirical observations, where associative structures at various spatial scales come to shape

the set of entities that are of context-specific relevance to goal acquisition. Again, this links with previous work to model the schematic structure of events as opposed to features themselves, as discussed in Chapter 4.

A criticism of image-schema theory is that identifying and categorising both image-schemata and compound image-schemata is subjective, with no single agreed upon categorisation. In terms of the aims of the present research however, this criticism is argued not to be of practical significance, following the epistemological position outlined in Chapter 5.

Over the course of research then, the following primitive image-schemata were found to be of direct use in developing the framework presented in Chapter 10:

#### LINK

**Definition:** either a link between the agent and an object formed through visual attention, or the perception of a link between two distinct entities

#### OBJECT

**Definition:** any perceptually significant discrete entity at any scale – so, entities at the highest level of detail as well as abstract entities that generalise over a number of smaller entities

#### BOUNDARY

**Definition:** a boundary as a physical endurant as opposed to an event boundary per se

#### CONTAINER

**Definition:** an entity for which there may be the topological relation contains relative to the agent, given a traversed path through the entity

#### SURFACE

**Definition:** a distinctive surface over some spatial extent

#### CENTRE-PERIPHERY

**Definition:** an egocentric relation between the agent and 1 or more image-schemata, i.e. 1 in the case of the CONTAINER schema

Building on this set of primitive image-schemata, over the course of research the following compound image-schemata were found to be of direct use in developing the framework:

**SOURCE-PATH-TARGET****Definition:** a directed, navigable path between two distinctive locations**COLLECTION****Definition:** a perceptually significant grouping of entities

While additional image-schemata were identifiable in the analysis and inference of participant knowledge, both primitive and compound, they were found not to have practical significance given the aim of selecting and displaying map content. For example the VERTICALITY schema was expressed by a number of participants, however if, for example, this was in relation to the identification of steps or an escalator, in practice the SOURCE-PATH-TARGET compound schema is still able to capture this transition, where attributes show that that the target of a network segment traversed on foot is the source of a network segment representing a traversal of a ‘steps’ feature. At the abstract level that serves to structure the overall representation, the image-schemata listed above are argued to be sufficient. Again, this is specific to use-case, and so the image-schemata that serve as the basis for the representation of events could be altered or added to without affecting the overall approach.

An event schema is then an associative structure comprised of the image-schemata as described, for each event type. The power of image-schemata as a basis for analysis and modeling is that their flexibility means not only can they be studied in the context of both textual and graphical empirical data, but they are also able to represent perceptual structures at different spatial scales, and across scales.

In terms of the textual descriptions that study participants created, as expected the SOURCE-PATH-TARGET compound schema is pervasive. While the data indicates some variation in the perceptual grouping of event sequences, in general directions corresponded with the ‘chunking’ of 2 to 3 events – i.e. 1 to 2 event sequences. In the highlighted example in Figure 78, if we understand “...cross over the road...” to be an event, then implicitly there are 3 events and 2 event sequences that form the basis of this perceptual structure. This chunk begins at the intersection described in direction ‘2’, and ends at “...the first bus stop...” in direction ‘4’.

In terms of the image-schema ‘OBJECT’, while it could clearly be argued that a generalised feature such as a topographic representation of a ‘city block’ is a COLLECTION, for the purposes of the current explanation, given that the representation in the underlying database is a single feature, generalised features are to be understood as image-schemata of type OBJECT. COLLECTION then refers to a grouping of objects that are not represented by a single discrete feature

at higher abstraction but may nevertheless be bound together as a meaningful grouping given perceptual experience in a given context. There were few examples in the data of participants explicitly conveying generalised entities. Generalisation was implicit in the representation of the road network, and was explicit in the single instance of a park feature, as well as a number of cases where ‘some shops’, and ‘residential area’ were conveyed. The appropriate level of abstraction with respect to these entities is apparent in these cases, however more generally, following the discussion

- 
- 1) Walk out of the ECCI cante and to the top of Infirmary Street
  - 2) Turn ~~right~~ right on to South Clerk Street
  - 3) At the first pedestrian crossing, cross over the road and continue north on South Clerk Street
  - 4) At the first bus stop you come to, take the 31 bus signposted 'East Craig'
  - 5) Take the bus north on South Clerk Street / South Bridge, stay on as the bus turns left on to Princes Street
  - 6) Continue on the bus along the entire length of Princes Street, towards Haymarket
  - 7) After Haymarket Station, press the stop on the bus and get off at the next stop
  - 8) Immediately cross the road and head North / North-east on Coates gardens
  - 9) Follow the road round until you reach 'Crescent'
  - 10) Turn right on to the crescent until you reach no. 26
  - 11) You've arrived

Figure 78 Participant textual directions example, G1\_10

around the problematic effect of removing content that conveys broader geographic context, it is argued that generalised features should be included in selections to support the formation of Gestalts that do not lead to the ‘resource depletion penalty’ effect discussed previously (Oulasvirta *et al.*, 2005), i.e. sufficient resources for the task.

In terms of image-schemata associated with perdurant entities, arguably ‘bus’ entities and ‘train’ entities would be candidates for this. Live bus and train data were not included in the research prototypes however, but this should be considered in the further development of the approach.

In Figure 79, an example illustration shows a TURN event. While an event can be associated with the location of an action, here denoted by the red node, the event itself should be understood as extending both temporally and spatially, with the

spatial extent determined by the locations of features that are associatively bound to the action and therefore play a role in goal acquisition. Taking a common example of how a TURN event is perceived, event boundaries extend between 1 or 2 distinctive states forward and backward from the event over the network (i.e. a topological distance of 1 or 2 over the network intersections). Additionally, entities that are adjacent to both network intersections and traversed paths within this

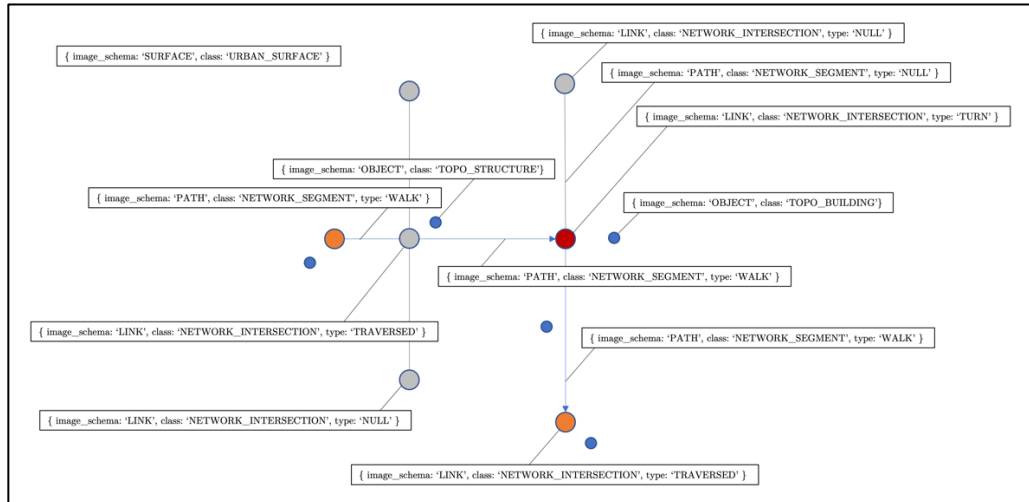


Figure 79 TURN event schema example illustration.

event also serve to extend the spatial region formed by the event out from the traversed path. In this illustration, ‘type’ attributes serve as the basis for representing context in that the same NETWORK\_INTERSECTION could be ‘NULL’, i.e. not traversed in the current journey; ‘TRAVERSED’, i.e. traversed in the current journey but not the location of a turning action; or ‘TURN’ – the location of a turning action.

A final point here, if we again consider the example textual description in Figure 78 (number ‘3’), is that irrespective of whether it is conveyed implicitly or explicitly, events overlap. Consider direction ‘2’: “Turn right on to South Clerk Street...”; direction ‘3’: “At the first pedestrian crossing, cross over the road and continue North on South Clerk Street...”; and ‘4’: “At the first bus stop you come to, take the 31 bus signposted ‘East Craigs’ ”. Traversing South Clerk Street occurs in all these descriptions, with the spatial extent of the description in ‘2’ overlapping the spatial extent in ‘3’, and similarly, the extent of the event conveyed in ‘3’ overlapping the event in ‘4’. This is the essence of the observation in (Eichenbaum *et al.*, 1999) discussed in chapter 2, that memory formation through events and event sequences during movement can be understood in terms of a ‘spatial context for experience’, as opposed to a ‘cognitive map’ per se. In other words, the

informational content in this form of representation is really derived through an encoding of the sequence of events, as opposed to the encoding of the geometry of the environment. While the physical shape of an environmental cue may serve to support orientation and the commitment of experience to memory, this is primarily due to the role the cue plays in the formation an event as opposed to the role it plays in the spatial encoding the local geometry, echoing the proposal in (Dabaghian, Brandt and Frank, 2014).

### 9.2.6 FROM EVENT-SCHEMATA TO COGNITIVE TOPOLOGY

In the context of the present research, cognitive topology (Lakoff, 1990) serves to describe the overall structure of experience over time. While actions and temporal boundaries may be implicit in compound image-schemata, for example the act of traversing a directed path and the duration of the traversal, clearly over an experiential episode, many compound image-schemata combine and interlink through continuous experience across tasks and subtasks. Reflecting on the ultimate research aim of selecting information that supports navigation across tasks and subtasks for an A – B journey, while an event partonomy represents the structure of experience at the level of overlapping event boundaries and the partonomic ordering of events, event-schemata serve to represent associative structures bound to actions within the spatial extent of these events. A cognitive topology in this case then is the overall set of image-schemata and compound-schemata that are active in the cognitive representation of all events over the episode. So, it is argued that incremental, context-specific spatial knowledge acquisition is supported by selections that reflect the features that match the predicted resource profile for the given journey, where a simulated cognitive topology can be understood as the basis for determining the resource profile.

### 9.2.7 PART 2 SUMMARY

Part 2 has shown that observations in the empirical data strongly support the case for the modeling and selection of information to be primarily structured around the navigable network. Context-specific hierarchical abstractions are supported by this approach which is argued to be fundamental to matching information resources to the predicted resource profile for a given journey. Building on the approach in (Mackanness, Quigley and Tanasescu, 2011) in which relations between context and content seemed to be based on the location of transport mode transitions, or TRANSFER events, an argument was developed for explicitly anchoring further event types to the route network such that event-schemata may be activated as a

result of a specific journey context. In Part 3 this argument is developed further with a focus on how features from various levels of abstraction should be ‘matched’ to the event-schemata that comprise the simulated cognitive topology.

### **9.3 Part 3: Determining content through context-specific regionalisation**

Spatial knowledge is known to be fundamentally topological in character, following the background review provided in Chapter 4. Localised representations with metric qualities are bound into looser global structures that support reasoning with incomplete knowledge and the flexible integration of new knowledge over time. In Part 2 of the analysis, this type of topological ‘template’ was discussed, with an empirically supported formalisation emerging through the study of event segmentation over navigational episodes. Event-schemata were presented as a basis for expressing the set of image-schemata corresponding with an associative network of features, where these features can be understood as being bound through attentional processes to the behaviourally significant action that serves as the ontogenetic basis of the event.

This account speaks to the underlying logic of selecting information resources given a simulated cognitive topology, however it lacks detail as to the specificity of the resources themselves. Here in Part 3, evidence is considered with respect to determining informational resources in terms of the appropriate level of representational abstraction, the clustering of features across submaps, and the practical basis of ensuring well-formed Gestalts given the overall aim of automatically generating selections over multi scale view space.

#### **9.3.1 ENTITY TYPES AND MULTIPLE REPRESENTATION**

In Part 2, OBJECT image-schemata were considered in relation to abstract ‘topological templates’. Here the entities corresponding with these OBJECT instances will be discussed given the overall aim of determining a resource profile of cognitive extension. During exercises over the course of empirical studies, participants conveyed their understanding of the entities they felt were significant in terms of the act of navigation. Entities were either conveyed explicitly through sketches and written directions, or the need for a representation of an entity as an informational resource was implied through feedback on prototype research stimuli.

Research participants had complete freedom in how they represented the geography to the hypothetical third party. Returning to a point highlighted in the methodology, participants were not directed to create a ‘map’, but simply ‘a sketch to help someone complete the exact same route’. While something ‘map-like’ was the implicit outcome of the exercise, how the city was represented is argued not have been constrained beyond the limitations and preconceptions of the participant themselves. Despite this however, the representation of entities did not diverge from standard cartographic features in any significant way. The exception to this was arguably the inclusion of additional textual information embedded in the sketch – i.e. explanatory notes as opposed to feature labels. These additions were noticeably more elaborate for the ‘sketch-first’ participants who had not already been through the process of writing textual directions, however it was common, particularly in the follow-up ‘green pen’ stage for textual pointers to be included in the sketch.

So, the conveyance of features at relatively small scales that generalise over a number of underlying entities included: the implicit generalisation of the path of the route and the geometry of the road network; residential areas where a change from a main road to a residential area was perceptually distinct; groups of commercial buildings corresponding with an adjacent road (i.e. grouped based on adjacency to a distinctive road as opposed to being grouped by ‘city block’); and parks and other ‘natural environment’ entities.

The selections underlying research prototypes included features from OS MasterMap corresponding with the SURFACE image schema, with polygon representations of surfaces being included in the selection where they were either intersected by the path of traversal or adjacent to the path. In practice, aside from road geometry features, this resulted in the selection of surface polygons in the localised region of the journey origin and the third transfer for SR1 (i.e. the region around the bus stop that was the location of the third TRANSFER).

A fundamental issue that was not anticipated prior to research was that it was very seldom that participants accessed map views at scales large enough for high detail geometries to be meaningful. Spatial knowledge was predominantly developed at scales ‘above’ the scale at which, say the geometry of an individual residential building was large enough to be of particular significance. The strong sense from observing users was that it was suboptimal to have to ‘zoom’ to large scales for viewing detailed topography. Clearly this type of interaction indicated that insufficient knowledge had been developed at smaller scales, with the overall conclusion being that the optimal informational resource profile on mobile devices emphasises knowledge acquisition through point and polyline objects, with polygon geometries providing a more general sense of geographical context. High detail topographic representations are likely to be accessed in only a very small number

of instances over the journey, and often, not at all. While the S1 research explored the potential to elevate high detail polygon features in the view through geometric distortion, the strong conclusion was that this interferes with the Gestalt of the view to such an extent that overall map effectiveness is reduced in comparison to standard approaches. High detail topographic polygon features can therefore be understood to be clustered around key decision points in the resource profile, with the overall emphasis over the profile being the effective use of point features, and generalised polygons that serve to support the overall Gestalt of a view, as well the perception of regions at scales beyond the spatial extent of egocentric experience for the journey.

In terms of point features, prototype research stimuli included, road names, business names from ground truth, text labels for train stations and bus stops, and borough text labels to provide some textual content at the ‘city-scale’. In terms of selections that correspond to active events, the point features that would be included in these selections would simply be dictated by whether they correspond to OBJECT instances. This point is addressed in the discussion. In terms of point features that support the acquisition of knowledge across broader allocentric regions, there was no evidence from empirical studies to suggest an approach that diverges from the standard approach – i.e. a prioritisation for labels associated with key landmarks.

As discussed, the specific approach presented in this thesis was developed through the study of the general case of urban travel. Clearly a use case such as ‘wheelchair user’ could have a significant impact on the set of features that would be of context-specific relevance to navigation. Here it seems likely the SURFACE image schema could play a more significant role in the patterns of experience across events for this case, for example. Adapting the approach to account for factors such as this is argued to be trivial given the necessary data. This will be addressed again in the discussion.

Perceptually distinct entities that are smaller than a building were also considered. For example a ‘ticket machine’ is clearly associated with the critical event ‘buying a ticket’, however within the overall context of the A-B journey, and given the constraints of screen area, no strong evidence was found to elevate the representation of this type of feature in the selection as they would immediately begin to compete with entities of greater significance to navigational decision-making at conceptual scales over larger spatial regions. In terms of including the geometry of this type of feature, for example bus shelters are in OS MasterMap and were included in the research stimuli, the practical reality is that the point feature representing the entity will effectively always be larger than the polygon representation. It was expected that the ticket machine in the train station would be conveyed in most cases in the sketch, however it was only included once.

Conveying very high detail entities such as this seemed to be more congruent with textual descriptions, however overall it is argued that cognitive load stemming from the act of searching for a very small environmental cue within the overall melee means this scale of OBJECT is argued, in the general case, to not be in the optimal information resource profile for supporting navigation.

### 9.3.2 SUBMAP CONTENT ACROSS MULTI-SCALE VIEW SPACE

In the previous subsection, the representation of entities that played a role in navigation was discussed. Here the focus shifts to the scales over which regionalisation occurs, such that features from varying levels of abstraction may be clustered into submaps that reflect the resource profile at the level of views. Reflecting on the finding that even if content could be arbitrarily combined to form integrated views across ‘global context’ and ‘local detail’, the most effective approach is still to support incremental knowledge acquisition given the natural division of attention over a relatively large number of views both before and during the journey. Here ‘submaps’ refer to clusters of features throughout multi scale view space given a resource profile. The term ‘submaps’ is used following studies into the neural representation of naturalistic spatial environments, where it has been argued the evidence suggests spatial knowledge is formed through the ‘interconnection of submaps’ as opposed to a consistent, homogenous form of representation (Derdikman and Moser, 2016).

Through an analysis of the empirical study data, it is argued that in the case of urban navigation, content should be understood as being structured across 5 region classes that each correspond with a distinct approach to selection and abstraction. These regions will be discussed in terms of the content associated with each, and in terms of their role in the overall regionalisation of the spatial extent given a journey context.

#### I-ZERO

**Definition:** an image space corresponding with the world space for the current journey, following the explanation in Chapter 6

#### JOURNEY EXTENT

**Definition:** an image space corresponding with a region bound by the locations derived from ORIGIN-DESTINATION

#### PHASE REGION

**Definition:** an image space corresponding with a region bound by two primary actions, e.g. ORIGIN-TRANSFER

## PATH REGION

**Definition:** a region derived from a PATH between the location of two events

## ACTION REGION

**Definition:** a localised region derived from the location a behaviourally significant action

Each of the region classes presented above can be understood as specifying a distinct basis for selection over the underlying graph. Reflecting on the proposal discussed in chapter 3 to use journey context as the basis for modulating the shape of an event horizon, each of these classes of region in effect serves as the basis for such a modulation.

Network structure serves as a proxy for the spatial basis of feature clusters in both the ACTION REGION and PATH REGION cases. Where an event schema is understood to be anchored to a location on the navigable network, an ACTION REGION then forms at this location based on a clustering of features that are associatively bound to the behaviourally significant action, i.e. an action which has a spatial location approximated by the location of the event schema. In the case of

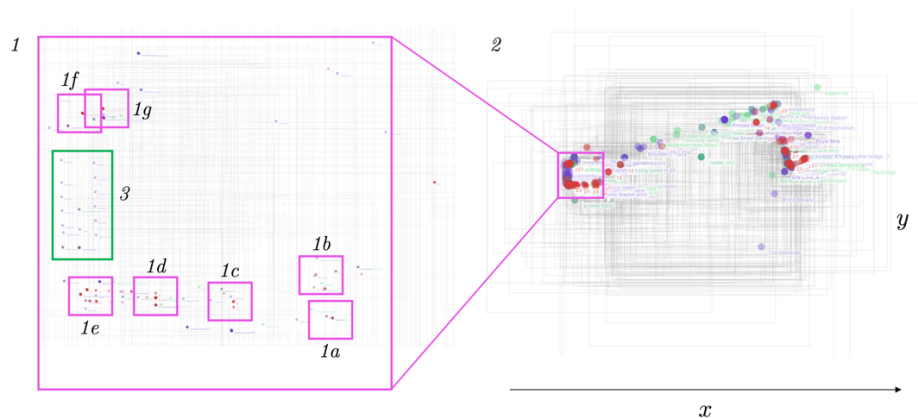


Figure 80 1a-g, clustering of features based on the structure of the navigable network for G1-G3 participant data; 2, the region shown in '1' within the context of the overall SR1 extent; and 3, the 'residential area' along the path towards the journey's final turning action. Axis  $x$  = easting,  $y$  = northing,  $\phi$  = map scale.

a PATH REGION, this region forms through the sequential ordering of events, where features that are adjacent to the PATH are understood to be of context-specific relevance to the current journey but play a less significant role than features corresponding with an instance of ACTION REGION. In both cases however, a crucial point is that features that are region members correspond with instances of OBJECT, which may be a feature from any level of abstraction. In this

way, event structure can serve as the basis for context-specific content hierarchies across view space. In other words, membership of a cluster associated with an instance of ACTION REGION can result in the prioritisation of the feature at smaller scales, irrespective of the feature's underlying level of abstraction. This same logic extends to path regions but with lower prioritisation.

In Figure 80 the entities that were explicitly conveyed by participants in study groups G1-G3 are visualised over the image space areas for the map views accessed by participants in G2. Here in panel '1' the clustering of features shows the correlation with the structure of the network, with 1a corresponding with the location of a TRANSFER, 1b-1e with the location of instances of CROSSING, 1e-1f with instances of TURN, and in 1g the DESTINATION. Panel '2' simply shows this same data in the context of the features conveyed over the whole SR1 journey. In the detail shown by the green box in '3', the 'residential area' is represented through features adjacent to the PATH.

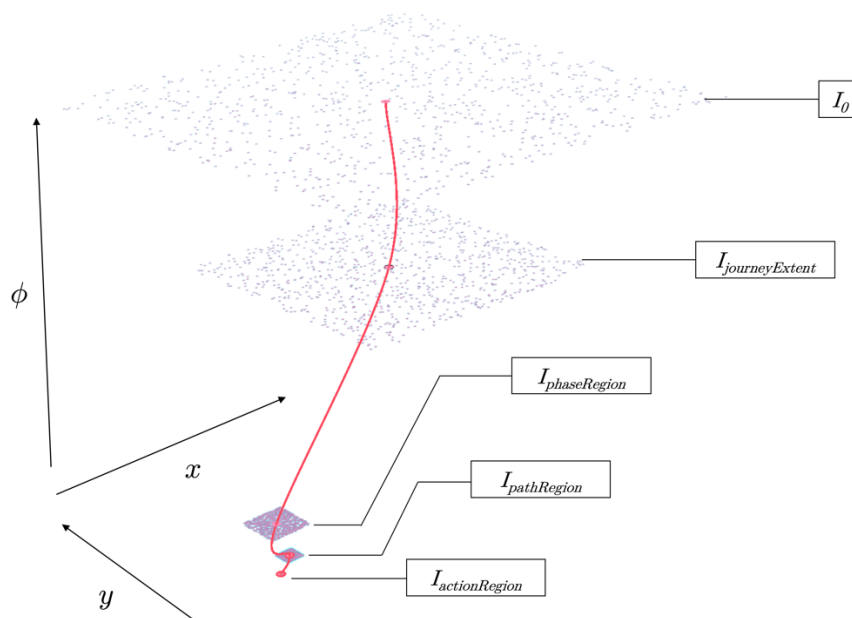


Figure 81 Illustration characterising the spatial extent of region classes in simplified view space. Axis  $x$  = easting,  $y$  = northing,  $\phi$  = map scale.

Returning to the proposed event-schemata, clearly the spatial regions that form based on the location of features associated with these structures are highly localised, and so at small scales it is not feasible to emphasise all features in the cluster. This will be explored further in the discussion.

In terms of region I-ZERO, this is simply the broader geographical context of the journey. A traveller may for example ‘zoom out’ to a view in which the JOURNEY EXTENT region can be contextualised relative to the overall layout of the city. While the extent to which content may be context-specific at these numerical scales is limited, the transition to the scale of a PHASE REGION serves as a basis to further prioritise low detail content for the current journey. See Figure 81 for an illustration of the spatial extent of regions in simplified view space.

### 9.3.3 PART 3 SUMMARY

In summary, the regionalisation of the journey extent can be understood in terms of map content that is reconfigurable based on context. Empirical studies showed the need to support incremental knowledge acquisition and free movement in view space, with the set of region classes presented in Part 3 serving as a basis for addressing these needs, while also clustering such that features may be matched to a resource profile for a specific journey context.

## 9.4 Concluding comments

The analysis has presented a number of arguments developed through the course of observation and hypothesis testing. The key finding that the central task in the effective automation of map content selection should be understood as predicting the state of a coupled system was explained in detail. The analysis went on to argue for the importance of the role of the navigable network in structuring the natural representation of a journey, and for the power of this form of representation in drawing explicit relations between context and content. Finally, map content itself was addressed, with issues of scale and regionalisation being considered throughout the presentation of a set of region classes. In the following chapter, the implementation of this approach will be presented, including a graph data model that supports the selection of features corresponding with event-schemata, and the processing of subgraphs such that content across interactive, multi-scale view space may be generated automatically.



# 10 Model

Building on the analysis and findings, the focus now shifts to the presentation of the formal approach developed over the course of navigation studies. Here ‘model’ is used as a term to generalise over both a graph data model and techniques for processing a subgraph for the purpose of automatically generating context-specific information. The chapter is divided into 4 parts. The first part provides background as to the central importance of causal reasoning within the overall approach. The second part elaborates on the brief explanation of graph technologies provided in Chapter 3, including the formal definition of a labelled property graph referenced throughout the explanation, as well as more specific background around the use of graphs in causal reasoning. The third part of the chapter then provides a detailed explanation of the framework that can be used as a basis for developing systems for the automatic generation of context-specific interactive multi-scale maps. This explanation covers the overarching approach, methods for constructing and querying a graph of features, subgraph processing as a basis for determining context-specific selections, and the link between the output of these techniques and the content that a user is able to view and interact with through a front-end system. Finally, the fourth section provides further detail as to the specification of the underlying data model, vertex classes and associated edge classes. The explanation also includes example graph transactions for demonstration.

In Chapter 9, the problem of how to automatically select and display map content to support navigation for an urban journey was argued to be best understood in terms of predicting the state of a coupled system. This coupled system consists of a navigational agent, an environment, and information delivered through a mobile device. Agents recruit resources into extended cognitive processes as a means of improving cognitive performance, with features in a geospatial database being potential resources within this overall picture. The approach presented in this chapter can then be understood in terms of determining the resource profile of cognitive extension for a given journey. In practice, this task is argued to be one of predicting the probability that a feature should be in a given view across multi scale view space.

## 10.1 The causal basis of cognitive topology

As a traveller plans their route, as they move away from their point of origin and as they evaluate, decide and act over the course of their journey, attention and the natural segmentation of experience shape the formation of a cognitive representation of the environment. While event-schemata have been proposed as a way to make explicit the associative patterns of experience across these episodes, embedded within the city is a latent context such that, location itself is not the only determinant of whether attentional processes lead to associative structures becoming *active* in the representation. Following the findings presented in chapter 9, to develop an approach in line with the overall project aim, it was found that map content to support navigation must be reconfigurable based on context, as opposed to their being a single configuration for a given spatial extent. Behaviourally significant actions can be understood as being *caused* by journey context, with the activation of event-schemata and a simulation of an overall cognitive topology also arising through causal processes. For example, if we understand a network intersection as having the *potential* to be the location of a turning action, the same network node could be traversed in various routes, but whether the event schema is activated is dependent on journey context. *There is a causal link between journey context and the activation of a specific event schema.* Given that features can be understood as matching instances of OBJECT in event-schemata, then it should be clear that if ‘Feature A’ matches an OBJECT in an active event schema for the current journey context, then there is a causal relation between the high-level journey context (the ORIGIN-DESTINATION), and *specific* features across levels of abstraction, for example the coarseness of generalisation.

The direct implication of this is that the same location in the network can serve to segment event sequences differently in different journey contexts, and so has a different causal role within the broader selection of features. Following the finding that selection must support the generation of content hierarchies to support pan and zoom through interactive multi-scale view space, these causal relations between journey context and features plays out across scales. Here then, a central challenge in the research was not simply to establish a basis for returning a subset of features across an event horizon to generate a single view, but to return a subset of features to support *any view* in multi scale view space given the journey context.

To clarify this point, if journey context results in a subset of features being returned from the underlying graph, how can it be known over which scales in view space (i.e. for which widths of image space) each feature should exist. In other words, there are two basic questions that must be answered for each feature in the journey extent: 1) should this feature be in the resource profile for the current journey context, and 2) if so, over what range of scales in view space should the feature

*persist?* To this second point, if we imagine a low-detail view in which both the origin and the destination are visible in the same portrait mobile view, there is very little physical area for map content. As the user zooms, content then materialises at some numerical scale. Persistence is then range of scales over which content remains in view space.




Figure 82 Features in multiscale view space as probability distribution determined through message propagation. Axis  $x$  = easting,  $y$  = northing,  $\phi$  = map scale.

Reflecting on the notion of map content as being *configurable*, the implication here is that as opposed to organising features in a predetermined hierarchical structure, features are instead organised such that a context-dependent hierarchy may be generated or each possible journey. Formal methods of causal reasoning were found to be the practical basis for achieving this, with causal reasoning being central to the ability for an agent to adapt to context (Waldmann and Hagmayer, 2013). Given that supporting incremental knowledge acquisition was found to be of central importance, and that causal reasoning fundamentally concerns the relationship between events over time, there was a natural fit between these techniques and the target functionality uncovered through navigation studies. There are also very natural linkages between formal methods of causal reasoning, the modeling and simulation of cognitive processes, and graph structures, and so research activity naturally converged on implementing and refining this type of approach. Exploring

these techniques also led to the succinct re-articulation of the central task as being to determine, given journey context, the probability that a feature should be in a given view.



Figure 83 Views of two phase regions for two journey contexts with a shared origin.

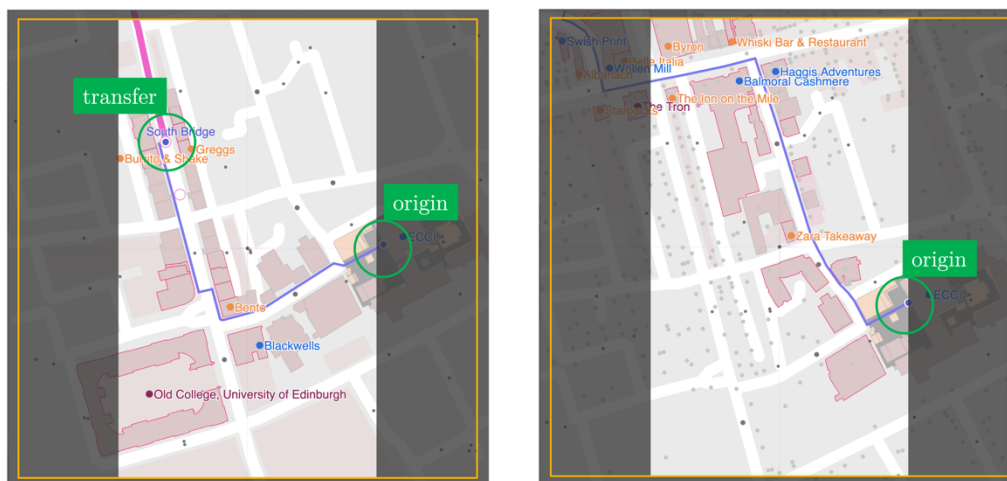


Figure 84 As in Figure 83, views of two phase regions for two journey contexts with a shared origin, but with each showing a view at a smaller image space width, i.e. a larger scale.

Referring to Figure 82, the content in this chapter presents a model that leads to the formal basis for this selection, whereby features are returned from an underlying graph and their membership across sets is determined such that an interactive multi-scale information space may be automatically generated. An annotated version of this figure is provided subsequently following further explanation, however this image is provided here to give the reader a sense of the outcome. In terms of providing clarity in the type of visualisation shown in Figure 82 it is necessary to

represent features simply as point objects. In Figures 83 and 84 then, example views for square subsets of view space are shown in the context of this square image space. These views can conceptually be understood as lateral ‘slices’ across probabilistic view space for given image space centres. Here two overlapping regions from distinct journey contexts are shown to highlight the manner in which content that is context-specific may be automatically determined using this approach.

## 10.2 Graph-based knowledge representation

It should be noted that the implementation technology used for demonstration describes graphs in terms of nodes and edges as opposed to vertices and edges. The term ‘vertices’ is applied throughout to distinguish between a graph vertex and a network node. Here then, any mention of ‘node’ refers to a network node, while in the implementation technology, any mention of ‘node’ refers to a graph vertex. In the description that follows it should now hopefully be clear that ‘network node’ refers to a type of data structure that is stored on a ‘graph vertex’.

### 10.2.1 PROPERTY GRAPH DEFINITION

The high-level graph data model that is used for storing and querying data is a *labelled property graph*. In other words a structure consisting of vertices and edges, where both vertices and edges are labelled and where both of these types of object can be associated with a set of key-value pairs. A vertex represents and stores information about an entity, an edge represents and stores information about the relationship between entities. Subgraphs comprised of different classes of vertices and edges can therefore be created to formally model phenomena of arbitrary complexity.

A property graph  $\mathcal{G}$  consists of a set of vertices  $\mathcal{V}$  and a set of edges  $\mathcal{E}$ . A subgraph is then a subset of vertices and edges in  $\mathcal{G}$ . There must exist a function  $\lambda_v$  for assigning labels to vertices, and a function  $\lambda_e$  for assigning labels to edges, i.e. there must exist a set of vertex labels  $L_v$  where:

$$\lambda_v: \mathcal{V} \longrightarrow L_v , \quad (10.1)$$

and a set of edge labels  $L_e$  where:

$$\lambda_e: \mathcal{E} \longrightarrow L_e . \quad (10.2)$$

There may also be a function  $\mu_v$  for assigning properties to vertices, and a function  $\mu_e$  for assigning properties to edges, i.e. a set of vertex properties  $P_v$  as key-value pairs where:

$$\mu_v: \mathcal{V} \longrightarrow P_v , \quad (10.3)$$

and a set of edge properties  $P_e$  as key-value pairs where:

$$\mu_e: \mathcal{E} \longrightarrow P_e . \quad (10.4)$$

For the purposes of the present research, a labelled property graph can therefore be defined as:

$$\mathcal{G} = \{ \mathcal{V}, \mathcal{E}, L_v, L_e, P_v, P_e \} . \quad (10.5)$$

Both edge and vertex id values must be unique, where uniqueness constraints are either imposed automatically by the database technology, or in the case of initially building the set of vertices from data, specifying uniqueness constraints explicitly.

Also it should be noted that given a graph  $\mathcal{G}$ , where the graph expresses a hierarchical ordering of vertices, topological distance can be understood in terms of vertex depth, following the explanation of the multi-cut example in Chapter 3 in which a  $k$ -level vertex can be understood as existing at depth  $k$ .

## 10.2.2 GRAPH DATA MODELS IN PRACTICE

A number of technologies have been developed for expressing data in graphs and for returning subgraphs via queries. The implementation of the framework presented in this thesis used the Neo4j graph database technology and its associated Cypher query language. In addition to the main database server, prototypes utilised the Neo4j Python Driver and the Neo4j Spatial server plugin. This approach meant that vertices could be constructed from csv data, and once both the set of vertices and the set of edges were stored in the graph, vertices with geometry properties could be added to a spatial index. In other words, prototypes supported both graph queries and spatial queries. Here the Neo4j technology makes use of the Java Topology Suite. While it is very useful for prototyping navigational applications to have a graph database technology that supports spatial indexing, it should be noted there is no aspect of the approach presented in this thesis that rests on any specific product or third-party library.

The key distinction between graph databases and relational databases is that in a ‘graph native’ technology, a relation between objects, or the ‘adjacency’ of vertices, is ‘index-free’. This offers great flexibility in making explicit the relations between instances of different types of entity in a way that is tractable from a query perspective, as one object is able to ‘point’ to another directly.

Querying a graph database can typically be understood in terms of two different approaches, either through imperative *traversals* or through declarative *pattern matching* (Rodríguez, 2015). Graph traversal means that a set of instructions defines a sequence of ‘steps’, moving across the graph from some initial starting ‘location’ (e.g. a starting vertex). The result of a traversal query is data from these ‘visited’ locations. In pattern matching, a template subgraph is expressed in the query itself, with the result of the query then being the data from any subgraph in the database that matches the pattern. In the approach presented in this thesis, the primary basis for returning data from a graph of features is the pattern matching approach.

### 10.2.3 CAUSAL NETS

A causal net is a graph structure that expresses causal relations between variables (Sloman and Lagnado, 2015). While a property graph can express the types of relationship between objects, a causal net expresses the change in state of child variables given changes in the state of parent variables that have a known causal connection to the child. Causal nets are graph structures, however the convention for the ontology of causal nets is to describe the constituent objects as variables and arcs as opposed to vertices and edges. In contrast with vanilla graph vertices, variables in a causal net in fact represent sample spaces, and ‘arcs’ is used by convention as the term implies directed links which clearly are a prerequisite for modeling causal relations. Given that the approach proposed in the present research requires that a causal net be instantiated from an underlying graph, this terminology also has the useful characteristic of serving to distinguish between a feature stored in the graph (a ‘vertex’), and a feature as represented in a causal net instantiated from a subgraph (a ‘variable’). Similarly, this serves to distinguish between relations stored as edges in the underlying graph, and the arcs that serve to propagate evidence in a causal net.

Following the definition in (Jensen, 2001), causal Bayes nets must consist of variables with a finite number of mutually exclusive states, and the patterns of causation as expressed by the set of arcs must be acyclic, i.e. there is no path from  $x_i$  to  $x_n$  such that  $x_i = x_n$ .

Following Bayesian methods of reasoning about causation, the power of the approach lies in the fact that the state of a variable can be entirely known from the state of its parents, as opposed to having to know anything about the states of variables that precede the parents in the flow of causation. This leads to reasoning being made tractable over large, complex patterns of causation, hence the applicability of the approach in relation to reasoning about real-world phenomena.

Key to the approach are the concepts of belief and *evidence*. Belief about states of the world can be expressed by the states of variables, but as new information or new evidence becomes available, belief can be updated based on the causal impact of this new evidence over the set of variables in the net.

In essence this can be expressed as:

$$P(x_{\theta_i} \mid \Xi), \tag{10.6}$$

or the probability that variable  $x$  is in state  $\theta_i$  given evidence  $\Xi$ . Here  $\theta$  expresses the state of the parents of  $x$ .




Figure 85 A causal net consisting of three unobserved variables and the observed evidence  $\xi_i$ .

Following conventions of hierarchical Bayesian modeling (Farrel and Lewandowsky, 2018), Figure 85 shows a causal net in which square objects denote variables with a discrete state space, white objects denote unobserved variables, and grey objects denote observed variables, or in this case, simply the evidence. Here then,  $P(x_{\theta_i}^k \mid \xi_i)$  is the conditional probability of the  $k$ -th variable in the net given the state of the  $i$ -th and  $j$ -th variables for the current evidence, or the  $k$ -th variable's *posterior* probability. A probability distribution therefore represents the state of all variables in the net, i.e. a distribution over  $X$ , where  $x \in X$ .

There are two forms of evidence that can form the basis of updating belief: *likelihood* evidence and *finding* evidence (Jensen, 2001). A finding is in essence a statement that given some evidence the probability of certain states of a given variable are impossible. Findings are specified as binary vectors, or equivalent data structures in terms of their information content. Likelihood evidence expresses belief over a range (0 to 1) and can therefore convey information that is not possible through finding evidence, i.e. a change in magnitude. In the approach presented in this chapter, finding evidence forms the basis of determining a probability distribution as we are only concerned with discrete structures. In other words, we need to know unequivocally whether a feature  $x_i$  should exist in a given view. The practical implications of incorporating likelihood evidence in further research will be considered in the Discussion chapter. Also it should be noted that while a variable can represent continuous state, the focus of the explanation provided in this chapter is variables with a finite number of discrete states. The implications of incorporating variables with continuous properties in future research will also be considered in the Discussion chapter.

If we understand an object representing a geospatial feature to be a variable  $x$  in in a given  $\Omega$ , where the probability the underlying object should be an informational resource is known based on the state of its parents, then the joint probability distribution over  $X$  represents *all* the features that are predicted to be informational resources for the evidence, i.e.

$$P(X) = \prod_{i=1}^n P(x_i \mid pa(x_i)) . \quad (10.7)$$

Here, following the notation in (Jensen, 2001),  $pa(x_i)$  simply refers to the parents of the variable.

#### 10.2.4 MESSAGE PROPAGATION

Updating belief in causal nets can be understood as a process of *message propagation*. A message is information that can be used as a basis for updating belief, i.e. a message conveys the state of parent variables. Message propagation is the process of passing evidence over the net, requiring that variables are able to *receive* messages via an incoming arc, and *send* messages over an out-going arc.

Various techniques have been developed for efficient message propagation, however for the purposes of demonstrating the approach presented here, it is

sufficient to simply define message propagation in terms of sending and receiving information given new evidence.

### **10.3 A causal framework for determining the resource profile of cognitive extension**

#### 10.3.1 SUMMARY OF APPROACH

Informational resources support goal acquisition over event structures that emerge through situated experience. However, these event structures, together with their implied resource profile, do not have a direct correspondence with a data structure that may be stored in memory. Instead the underlying graph must support the expression of event structures and implied resource profiles given *latent context*. In other words, the overall system must support selections that correspond with context-dependent resource profiles given updated beliefs, following the explanation in 10.2.3. It has been found that to achieve this requires an approach that can be thought of in terms of the dual processes of conscious and unconscious cognitive control as discussed in (Posner and Snyder, 2004). In this explanation of attention, interaction between ‘automatic activation processes’ and ‘conscious strategies’ determine cognitive performance, with a clear link to the interaction between long-term and working memory, echoing both (Gopal, Klatzky and Smith, 1989) and (Golledge *et al.*, 1985). This can also be understood in relation to the distinction between a ‘builder’ and a ‘searcher’ function, as discussed in (O’Neill, 1991).

The central insight is that a final selection must be determined through a process that is similar to conscious attention in which limited capacity is a fundamental constraint. The dual process can be seen then in terms of an initial selection of a *universe* of objects (i.e. the set of possible objects) that may or may not be attended to given the current context, followed by a process that is analogous to simulating conscious attention given both the context and the universe.

The initial subgraph returned from the database can be thought of as a universe of *candidate* objects, of which the objects and relationships that are in the resource profile for this event paratomy are a subset.

This subgraph is then used as a basis for instantiating a causal net, where variables in the net represent vertices in the graph, and arcs in the net represent edges between vertices. An important aspect of the approach then is that a query needs to not only return the data stored on graph vertices, but also needs to return the

structure of the data itself. In other words, it is necessary to return relationships as well as the data stored on vertices, such that casual relations may be constructed based on this structure. By using journey context as a basis for determining patterns of causation over the net, features can be selected such that an overall set of features and their persistence in view space can be understood as an approximation to the resource profile for a given journey.

It should be highlighted that the practical basis of the approach is the ‘ $k$ -level’ network. This term describes a data structure that includes nodes representing all event types on a centreline network. While a standard centreline network clearly includes the location of potential TURN actions, the ‘ $k$ -level’ network in the present research includes all locations representing event types, where integration into the data structure of such a location (e.g. a bus stop) simply uses shortest distance intersection.

### 10.3.2 JOURNEY CONTEXT & EVENT SEGMENTATION

As outlined in Chapter 9, the empirical research showed that for the purposes of selecting features that approximate the informational resource profile of urban navigation, an adaptation of a road centreline network can be used as the foundational basis of the data model and queries. It is assumed a routing engine has returned an ordered sequence of network nodes given an origin-destination query  $OD$ . This sequence from the  $k$ -level network is then used to derive a set of temporary vertices with properties including the associated  $k$ -level node id, and a property that stores a value representing the primary action that forms the basis of the associated event. This property can therefore be assigned one of the following values for the present use case: “TRAVERSED”, “TURN”, “TRANSER”, “LEVEL-TRANSITION”, “CROSSING”, “ENTRANCE-EXIT”, “CHECK-POINT”. Temporary graph vertices representing “ORIGIN” and “DESTINATION” must also be created for the session with a geometry property, but clearly they do not share an id with a network node graph vertex stored in memory. All network nodes in a sequence are to be traversed for the current route, with the “TRAVERSED” value simply serving to distinguish between intersections that are the locations of turning actions and those that are not.

The objects representing the traversed sequence of  $k$ -level network nodes  $r$  are stored in a route sequence array variable  $\mathcal{R}$  (i.e. a standard array in which an index set is created and each element in the index set if mapped to each object, in this case the objects that represent network nodes).

### 10.3.3 PROPAGATION MATRICES

The representation of context develops in stages. The initial basis of journey context is the set of locations expressed in the *OD* query. From there the sequence of network nodes and their updated properties are created and stored in an array variable as discussed in the previous section. The array variable  $\mathcal{R}$  is then passed to graph queries, and is also passed to a `GET_CONTEXT` function that assigns propagation matrices. Graph queries will be discussed in the following section.

The `GET_CONTEXT` function takes the objects in  $\mathcal{R}$  and based on the value of each primary action property, assigns *propagation matrices* that express patterns of activation over subgraphs of features in localised regions. Following the explanation in 10.2.3, these propagation matrices are comprised of *finding vectors* for each conceptual scale given the primary action property.

A finding on a variable  $x_i$  is a binary vector where each variable  $a$  in the vector is a categorical variable of the form:

$$a \in \{ 0, 1 \} . \quad (10.8)$$

The information content of a message is then a finding vector  $\mathbf{f}$ , where:

$$\mathbf{f} = (a_{i,j}, \dots, a_{i,n}) \quad (10.9)$$

A propagation matrix  $\mathbf{M}$  is a representation of findings for variables linked to an  $x_{NK}$  instance for a given  $\mathcal{S}$  in the current  $\Omega$ . If the resource profile of an event partonomy for a given domain can be modelled based on 5 scales with  $2^4$  possible patterns of activation for each finding, a propagation matrix based on an object in  $\mathcal{R}$  is therefore of the form:

$$\mathbf{M}_{\mathcal{R}} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \\ a_{51} & a_{52} & a_{53} & a_{54} \end{bmatrix} . \quad (10.10)$$

Here the 5 scales discussed in Chapter 9 are represented by the number of rows, and the 4 columns reflect the number of possible findings required given the structure of the causal nets instantiated from an underlying graph that conforms to the proposed data model (presented subsequently).

### 10.3.4 COGNITIVE TOPOLOGY AS GRAPH PATTERNS

Graph queries are based on a labelled sequence of  $k$ -level network nodes. This data is stored in an array and is passed into MATCH queries on the underlying graph. In other words, ‘NODE\_K’ vertices in the graph are matched based on id properties of objects in the traversed route sequence. These ‘NODE\_K’ graph vertices then serve as the source vertices in paths through the graph.

The data model has been designed such that a traversal of the graph can return the set of features that correspond with the resource profile of possible event-schemata for the current journey context. The practical basis for achieving this is by constructing the graph so that a repertoire of graph patterns match subgraphs that return the required features. Referring back to Chapter 4, in the context of the present research, cognitive topology can be understood as the overall structure of the cognitive representation of the geography for a navigational episode. The approach is to model the graph such that a *simulated* cognitive topology can be constructed out of the latent context expressed in the structure of the underlying graph. Subgraphs representing potential event-schemata can be returned via pattern matching, and so for the given journey context, features are returned that potentially correspond with the resource profile. Anything that could correspond





Figure 86 Two graph patterns based on topological distance, where circular objects represent graph vertices.

with an OBJECT schema in a cognitive topology for a navigational episode that may be included in a resource profile must be represented by a graph vertex. Event-schemata can be understood then as being expressed explicitly through graph patterns, or implicitly by the selections that result from queries.


In Figure 86 two graph patterns are shown. In panel ‘a’ the pattern corresponds with any subgraph which matches a topological distance of 1 from the source vertex and in panel ‘b’, the same thing is illustrated but for a topological distance of 2. Square brackets are used to denote relationship type, but the pattern match query implied by the ‘empty’ brackets means that the type of relationship is unimportant, i.e. any subgraph corresponding with this pattern will be returned.

It was found that five compound schema templates are required as the underlying basis of patterns which are able to express context given a graph of geospatial features. Here then, five distinct relations between features must be formalised for the successful automatic construction of multi-scale maps. The phrase ‘compound schema’ is borrowed from the cognitive linguistics literature discussed extensively throughout the thesis. It is used here to describe distinct groupings of entities that, taken together, reflect some fundamental repeated pattern of experience over a navigational episode. ‘Simulated cognitive topology’ can therefore be understood in terms of data structures based on a set of active schema templates for the current journey context.

a)



b)



*Figure 87 Compound schema template 1: features that segment experience serve to further decompose the network representation leading to a new, more detailed network representation. In panel b) a feature that is understood to be a proxy for a potential action is used to derive the new k-level network node C.*

In figure 87, the first compound schema template is illustrated. Here, features that segment experience serve to further decompose the network representation. So, while possible turning actions are naturally represented by the location of network nodes, additional actions such as transferring between different modes of transport are also represented within a single navigable network. In this example then, the feature could be a bus stop, and node C would be created by

finding the shortest distance intersection between the feature and an existing network segment. Clearly then the output of this operation is three new entities

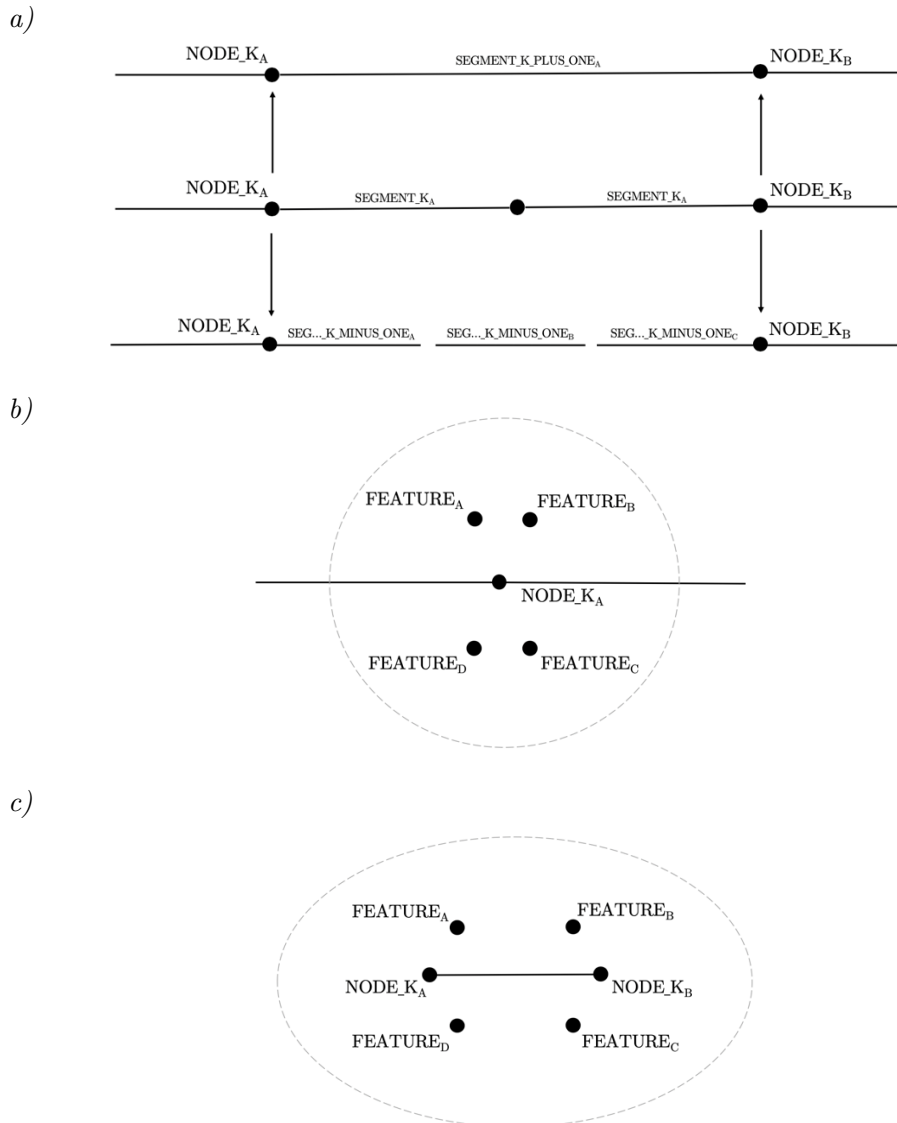


Figure 88 Compound schema templates that provide the basis of key graph patterns, with a) template 2 as nested paths, b) template 3 as a region based on associating an action with a location, and c) template 4 as a region created by the sequential ordering of actions.

in the database: node C, and two new network segments. While this approach is primarily aimed at representing the potential location of key actions, the same approach can be used to create a representation of a distinct feature that is also predicted to segment experience even if it is not tied to a behaviourally

significant action per se. In terms of a graph pattern then, returning a path through the navigable network is simply a case of matching on the ordered sequence of nodes returned by the routing engine.

The second compound schema template is the hierarchical ordering of network segments, such that an abstract representation of a path can be associated with more fine-grained representations, as shown in figure 88, panel a). Here the set of  $k$ -level nodes is used to derive both the lower detail ‘ $k$ -plus-one’ network segments, and higher detail ‘ $k$ -minus-one’ segments. In practice then, this template is one of nested sequences of SOURCE\_PATH\_TARGET schema representations, stored in a database as three resolutions of network segment. In the case of multi-modal travel, the set of  $k$ -level nodes that forms the basis of the overall network representation can be understood as supporting pedestrian navigation over areas that include a sequence of roads, the  $k$ -minus-one segments support pedestrian navigation by providing geometries at the pavement (sidewalk) level of an individual road, and the  $k$ -plus-one segments support navigation over larger distances traversed by public transport.

Figure 88 panel b) illustrates the third compound schema that represents a grouping based on the location of a behaviourally significant action. Features adjacent to the node that is associated with the action become associatively bound to the action,

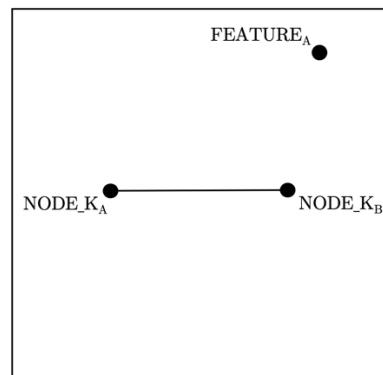


Figure 89 Features not adjacent to a route but lying within a spatial region formed around a navigable path.

and so form a region. An important point here is that features from multiple levels of abstraction may be bound to the same region (to the same node), which provides the practical basis for distinguishing context across multiple scales. In terms of the fourth template then, as illustrated in panel c), a sequence of actions represented by two  $k$ -level nodes forms an elongated region, with features adjacent

to the path connecting the two actions being associatively bound to this sequence of events. Again, this approach means that features from multiple levels of abstraction can be associated with the same region.

The fifth template illustrated in Figure 89 is not naturally expressed as a graph pattern as it is spatial, but it is derived from a graph pattern in the sense that the spatial region is defined by two bounding network nodes. This template reflects the need to include features not adjacent to the route, following the research finding about the need to ensure well-formed Gestalts that support the development of hierarchical reference memory.

### 10.3.5 FROM SIMULATED EXPERIENCE TO CONTEXT-DEPENDENT REGIONALISATION

As outlined in Chapter 9, event structures for navigational episodes that are of relevance to information needs across the task hierarchy can be approximated using three main classes of region: action regions, phase regions, and path regions that can themselves decompose into a set of smaller path regions at a lower level of abstraction (i.e. that chunk fewer action sequences). Each region type requires a distinct data structure such that a representation of the region and its associated features can be determined from the latent context discussed previously.

#### **Action regions**

As outlined in Chapter 9, key features of relevance to urban navigation can be adequately approximated by spatial clustering based on the location of actions that are of central importance to goal acquisition. The spatial distance from the location of the action can be qualitatively understood as the distance that leads to a cluster which includes features that are adjacent to the action, where qualitative adjacency is largely a function of visibility through rotational movements and topological distance – i.e. the first one or two distinctive features along a sequential ordering following the axis of each possible direction corresponding with a navigable path away from the location of the action. Action regions must be anchored to a location on the navigable network, with the spatial extent of the region being determined by the geometry of the features that are associatively bound to the location of the action.

#### **Path regions**

Similarly, spatial clusters in path regions derived from qualitative adjacency from the traversed path adequately approximate content that is of less significance but directly supports goal acquisition, with the underlying cognitive structure that is supported being an event sequence.

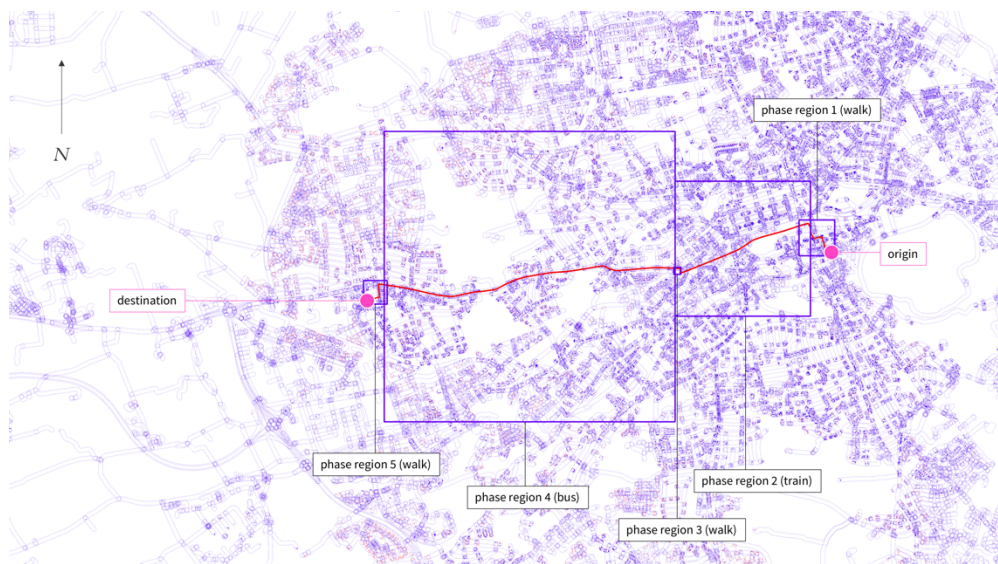


Figure 90 SR2 journey overview showing the 5 phase regions.

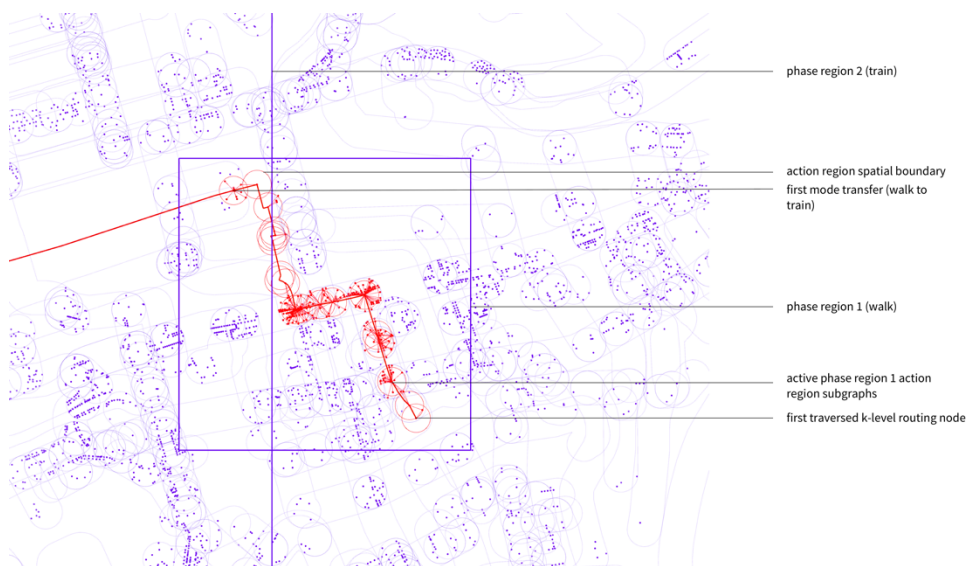


Figure 91 Origin region in SR2 journey context, with ACTION\_REGION feature clusters highlighted in red.

### Phase regions

In terms of regions that are segmented at the level of primary actions, no argument was found for these data structures to be anything more than a bounding box derived from the location of two primary actions on the network.

In Figure 90 the overall SR2 journey extent is shown, including the 5 phase regions given the 6 primary actions in the SR2 journey context. In Figure 91, the first phase region is shown at a larger scale. In both of these figures the boundary of path and action regions are shown, with feature centroids within these regions being shown also.

### 10.3.6 JOURNEY CONTEXT AS EVIDENCE

Following the explanation in 10.3.3, the function `GET_CONTEXT` takes  $\mathcal{R}$  as input, iterates through each object in  $\mathcal{R}$ , and creates a new object for each iteration comprising of the id property from the original object and a finding. Here the finding property is the associated propagation matrix assigned from a set of template matrices given the value of the  $r_{type}$  property.

Evidence is then the set of objects, which are themselves sets, comprising of the vertex id and a propagation matrix:

$$\xi_i = \{ v_{id}, \mathbf{M}_{\mathcal{R}}^i \}. \quad (10.11)$$

Evidence for a journey context  $\mathbf{M}_{\mathcal{R}}$  is the evidence that forms the basis of updating belief over the net, i.e.

$$\Xi = \text{GET\_CONTEXT}(\mathcal{R}) = ( \xi_1, \dots, \xi_n ). \quad (10.12)$$

The evidence  $\Xi$  expresses *the patterns of activation* over the net for the given journey context  $\mathcal{R}$ .

Although a different technique, this can be thought of in terms of graph cuts based on source and target vertices following Quigley (Quigley, 2001). Here however, a variable corresponding with a feature at the level of ‘local detail’ can be explicitly targeted at source.

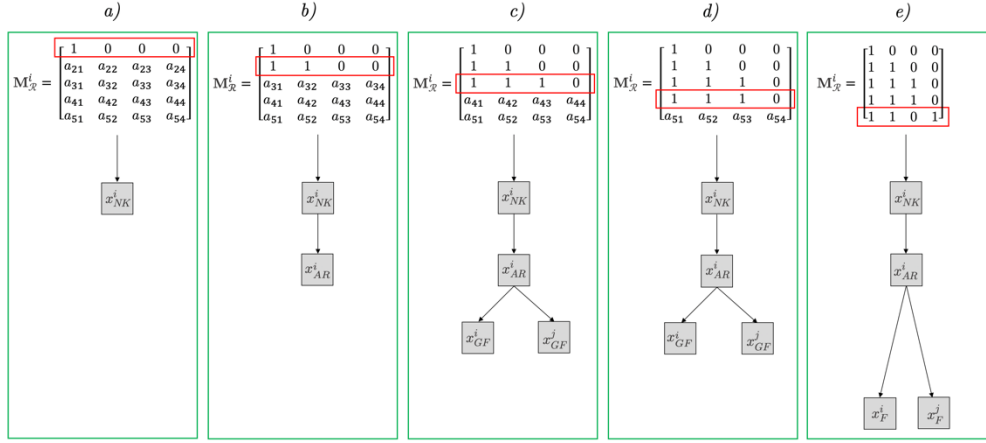


Figure 92 Propagating evidence  $M_x^i$  for 5 scales in view space (a-e) starting at variable  $x_{NK}^i$ , i.e. the variable representing the  $i$ -th  $k$ -level network node.

### 10.3.7 CONSTRUCTING A CAUSAL NET

Following the explanation in 10.3.2, the indexed family of variables  $X$  is the result of a CONSTRUCT\_VARIABLES function  $\alpha$  that takes the set of vertices in the subgraph  $\mathcal{S}_v$  as input i.e.:

$$\alpha: \mathcal{S}_v \longrightarrow X . \quad (10.13)$$

The indexed family of arcs  $Y$  is the result of a CONSTRUCT\_ARCS function  $\beta$  that takes the set of edges in the subgraph  $\mathcal{S}_e$  as input, i.e.:

$$\beta: \mathcal{S}_e \longrightarrow Y . \quad (10.15)$$

Note that the outputs of these two functions are not indexed families based on the index set of the underlying  $\mathcal{G}$ , but are families with index sets that are temporary, and specific to the current session.

As illustrated in Figure 93, vertex  $v$  in  $\mathcal{G}$ , that stores data such as a feature's geometry, is returned in a subgraph  $\mathcal{S}$ . The vertices in this subgraph are then represented by variables in the set  $X$ , i.e. for a feature of type  $NK$ , i.e.  $v_{NK} \in \mathcal{G}$ , and  $x_{NK} \in \Omega$ . It should be noted that although index sets for the graph vertices and the causal net variables are not the same here, 'i' is used as the index value for both  $v$  and  $x$  in this case to illustrate that the variable  $x$  is the representation of  $v$  in  $\Omega$ .

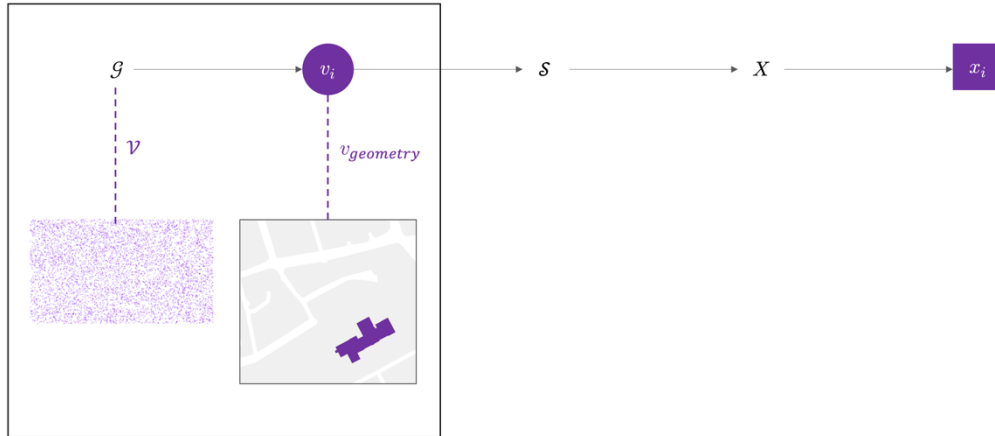


Figure 93 Illustration of the relationship between a graph vertex (circular object) that stores feature data (e.g. geometry data) and the representation of the feature as a variable (square object).

A causal net  $\Omega$  then, is comprised of an indexed family of variables  $X$  an indexed family of arcs  $Y$ , the evidence  $\Xi$  for the current journey, a set of functions for sending messages over an arc  $T$ , a set of functions for receiving messages over an arc  $\Pi$ , and a set of functions for propagating evidence  $\Gamma$ , i.e.:

$$\Omega = \{ X, Y, T, \Pi, \Gamma \} . \quad (10.16)$$

### 10.3.8 ACTIVATION VIA MESSAGE PROPAGATION

As discussed, variable activation is achieved through message propagation over the net given the evidence, where the result of message propagation is a change of state  $\theta$  for each variable  $x$  that matches the pattern of activation expressed through the finding vectors in the propagation matrix.

A function  $\tau$  for sending a message  $\mathbf{f}$  over an arc  $y$ , where  $\tau \in T$ , is defined as:

$$y_{i,\theta_i} = \tau(x_{i,\theta_i}) \quad (10.17)$$

i.e. the message passed over arc  $y_i$  is determined by the state  $\theta_i$  of the source variable  $x_i$ , where the evidence that forms the basis of message propagation is the vector  $\mathbf{f}_i$ .

A function  $\pi$  for receiving a message  $\mathbf{f}$  over an arc  $y$ , where  $\pi \in \Pi$ , is defined as:

$$x_{j,\theta_i} = \pi(y_{i,\theta_i}) \quad (10.18)$$

i.e. the state of variable  $x_j$  is determined by the message received over arc  $y_i$ .

Here the representation of source and target variable's in the  $y$  data structure is simply their index value in the current session. The assignment of send and receive functions for a given  $y$  is determined by the type of edge the arc represents.

### 10.3.9 ACTIVATION TRACES & PROBABILISITIC VIEW SPACE

The set  $\mathcal{U}$  is the sample space for the causal net  $\Omega$ , where:

$$\mathcal{U} = \{ X, Y \mid \mathcal{S} \} \quad (10.19)$$

Activation is then the process of determining the variables in  $X$  that are to be in the current resource profile, and over what range of scales they persist in view space. Note that variables representing  $k$ -level network nodes shall be denoted  $x_{NK}$ . The state of each  $x_{NK}$  variable is updated based on each variable's associated propagation matrix, i.e. updating the state of each  $x_{NK}$  leads to evidence being propagated over the net. An activation trace is then the subset of variables that have been activated for a given  $x_{NK}$ . Here the term 'trace' is used following the idea of a 'memory trace' (Vetere *et al.*, 2019).

The function `VARIABLE_ACTIVATION` runs propagation over the net, starting at each  $k$ -level routing node using the finding for the current scale. In other words, view space is divided into an integer number of scales, and `VARIABLE_ACTIVATION` runs this number of times. The finding is the pattern of activation for the parent NK variable for the activation trace at that scale.

Algorithm: VARIABLE_ACTIVATION
<i>Input:</i> A variable $x$ , the set of variables $X$ , the set of arcs $Y$ , an integer value $scale$ , a <i>finding</i> , a <i>selection</i> dictionary (or equivalent data type), and a <i>counter</i> value.
<i>Output:</i> The index of activated variables added to the scales at which the feature represented by that variable persists in view space.
<pre> 1.  <i>activatedVariables</i> <math>\leftarrow</math> [ ] 2.  <i>outArcs</i> <math>\leftarrow</math> <math>x.outArcs</math> 3.  <math>pMapping</math> <math>\leftarrow</math> MAP_FINDING_TO_ARCS(<i>finding</i>, <math>Y</math>, <i>outArcs</i>,       <i>counter</i>) 4.  <i>activatedVariables</i> <math>\leftarrow</math> PROPAGATE(<math>X</math>, <math>Y</math>, <math>pMapping</math>) 5.  <b>for each</b> <math>x</math> in <i>activatedVariables</i>: 6.      MERGE_INDEX(<math>scale</math>, <i>activatedVariables</i>, <i>selection</i>) 7.  <i>count</i> <math>\leftarrow</math> <i>counter</i> + 1 8.  <b>if</b> <i>activatedVariables</i> &gt; 0: 9.      <b>for each</b> <math>x</math> in <i>activatedVariables</i>: 10.         <b>if</b> <math>x.variableType</math> != 'nk': 11.             VARIABLE_ACTIVATION(<math>x</math>, <math>X</math>, <math>Y</math>, <math>scale</math>, <i>finding</i>,       <i>selection</i>, <i>count</i>) </pre>

Here further explanation is provided for each of the steps in the algorithm shown above:

1. An empty *activatedVariables* array is created
2. A property on the variable object that lists the out arcs is store in an *outArcs* variable. This is created when the variable is instantiated.
3. The function MAP\_FINDINGS\_TO\_ARCS is called. This function returns a list of out arcs for the current trace, using the counter as a way of expressing the reduction in the length of the finding given the topological distance of a variable from its associated NK variable. The result of this call is stored in *pMapping*.
4. The PROPAGATE function is called that returns the activated child variables on the filtered set of out arcs. The result is stored in the *activatedVariables* array.
5. Iterate through each variable object in *activatedVariables*
6. On each iteration, call the function MERGE\_INDEX. This function ensures that if a variable (a feature) is activated at a given scale, its index value is associated with the correct feature type at that scale, and that it is unique at that scale (i.e. it appears either 0 or 1 times at that scale).
7. The counter value is incremented so the depth or recursion can be known.
8. An 'if' clause to check if the *activatedVariables* array is not empty.
9. Iterate through each variable object in *activatedVariables*.

10. Ensure the variable is not of type NK.
11. Recursively call VARIABLE\_ACTIVATION for each of the activated variables using the incremented counter value.

As discussed, the result of VARIABLE\_ACTIVATION is referred to as an activation trace, which can now be defined as:

$$\psi = \text{VARIABLE\_ACTIVATION}(x, X, Y, \phi, \mathbf{f}, \textit{selection}, \textit{counter}) \quad , \quad (10.20)$$

where  $\phi$  represents the scale. This will be discussed further in the following subsection (10.3.10).

The context-specific selection of features, as represented by variables, can now be defined, where  $\omega$  represents the union of each activation trace  $\psi$ , given a causal net  $\Omega$ , and the evidence  $\Xi$ .

$$\omega = \bigcup_{i=0}^n \{ \psi_i \mid \Omega, \Xi \} \quad (10.21)$$

So,  $\omega$  can be understood as the resource profile for the given  $\mathcal{R}$ , or the probability distribution over  $X$  across multi-scale view space. As discussed in the findings, it is argued a selection of features equivalent to the features corresponding with an event horizon cannot be known from a single subgraph. In practice this can be thought of in terms of variable activation passing evidence over the net multiple times. Returning to the point about the dual processes of cognitive control, selecting the elements of the set  $X$  and the set  $Y$  can be understood as a process in which there is effectively no capacity constraint, graph patterns simply return data by matching on NK nodes and then returning subgraphs according to the defined patterns. The process of variable activation however can be understood in terms of conscious attention, with propagation over the net being roughly analogous to the formation of a cognitive representation of a localised region of the environment given the particular character of spatial goal acquisition for the current journey context.

In Figure 95, the output of variable activation is visualised for a view space based on 5 scales, using the phase region structure from SR2 and the feature types ‘BOROUGH\_TEXT’, ‘NK’, ‘OSM\_POINTS’, ‘VML\_POINTS’. Figure 94 is simply the same illustration as Figure 87 but is included as a visual reference for the phase regions.

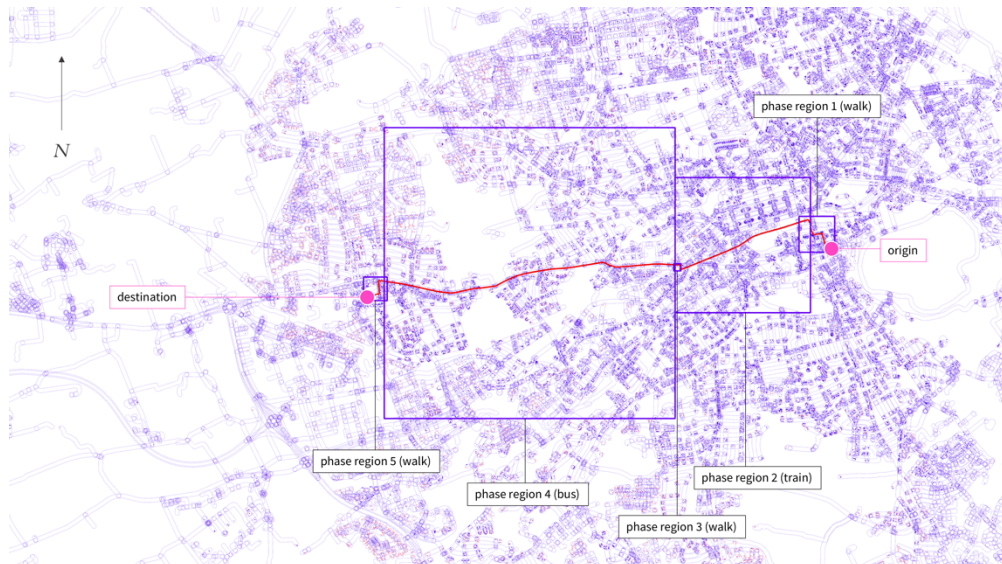


Figure 94 SR2 journey overview showing the 5 phase regions (the same illustration as Figure 70).

In Figure 95, example features are shown in terms of their persistence across view space, with an image space corresponding with ‘scale 1’ including the path of the route and the ‘BOUROUGH\_TEXT’ features, for example. A view corresponds with a square subset of view space, as illustrated in Figure 96.





Figure 95 Result of variable activation for the SR2 journey context visualised in simplified view space with example feature data.



*Figure 96 A view for a given journey context as a lateral slice across probabilistic view space (red) for a given view centre, i.e. a local chart in the manifold over the underlying set of features.*

In Figure 96 a view as a square subset of view space is shown, using polygon centroids and the network nodes bounding network segments as a basis for illustration. As these features are all of type POINT, the spatial index is simply a quadtree, as shown at the ‘base’ of view space in the figure. Again, this is just for illustration, with the spatial index of the features in the underlying database in fact being an Rtree. While the visualisation in Figure 95 can be understood in a similar way to the explanation of map generalisation as being like an ‘extrusion’ into the third dimension (van Oosterom and Meijers, 2014), a key conceptual distinction to grasp in interpreting this visualisation is that features from different levels of abstraction are projected ‘up’ into the space such that a lateral slice may include non-uniform abstraction of the same feature type, as opposed to the ‘non-horizontal’ slices discussed by van Oosterom et al. It is argued this is conceptually simpler to grasp given the ultimate output is a ‘flat’ 2D view.

### 10.3.10 AUTOMATICALLY GENERATING INTERACTIVE MULTISCALE MAPS

The result of message propagation cannot however be used directly to generate an interactive map. While the relevant subset of features from the overall subgraph returned from  $\mathcal{G}$  are now known, and their persistence is known for a range of scales, an interactive map that supports continuous zoom in fact requires a larger number of scale transition points. Over the course of model development it was found however that by running variable activation for a reduced number of scale transition points, a front-end application can then assign transition points to features based on type within each scale ‘band’. If we imagine scale transitions from zoom level 15 (1:15k), through to level 20 (1:500), then clearly there are 5 transition points through view space. In practice however, a broad range of content should not suddenly appear at each transition point, with screen area being used optimally by gradually introducing higher levels of detail across view space.

The approach presented in this chapter does not consider numerical zoom, but simply provides a subset of features and an overall ordering. While each scale is represented by an integer in the `VARIABLE_ACTIVATION` function, in fact these ‘scales’ should be understood as intervals, where the numerical scale at which a feature materialises lies between the upper and lower bound of the interval, as shown below for the case of 5 interval ranges.

$$\begin{aligned}\phi(0) &= \{ z \mid scale_0 \leq z < scale_1 \} \\ \phi(1) &= \{ z \mid scale_1 \leq z < scale_2 \} \\ \phi(2) &= \{ z \mid scale_2 \leq z < scale_3 \} \\ \phi(3) &= \{ z \mid scale_3 \leq z < scale_4 \} \\ \phi(4) &= \{ z \mid scale_4 \leq z \leq scale_5 \}\end{aligned}$$

The function  $\sigma$  maps the resource profile  $\omega$  determined through message propagation to numerical scales that directly determine the persistence of features across multiscale view space, and so:

$$S = \sigma_\omega, \quad (10.22)$$

where  $S$  is comprised of an index of features, and lists of feature indexes for each numerical scale.

A *persistence vector* is then implied by the ‘highest’ and ‘lowest’ points of the range over which a feature materialises in view space, following the illustration in Figure 28 (Chapter 6). It should be noted there is no literal ‘persistence vector’ data

structure. These objects are conceptualised to support explanation, however clearly these objects do have a mathematical basis and could be created from a given  $S$ . Similarly, there is no explicit ‘manifold’ data structure required in the approach, however, again the notion of a manifold is not entirely conceptual as a manifold object could be created based on a given  $S$ .

## 10.4 Graph data model

Developed and refined over the course of the situated navigation studies, the graph data model expresses the objects and relations that describe navigable urban environments such that a selection can be returned in-line with the likely resource profile of cognitive extension over a context-dependent event partonomy. The data model supports subgraphs that can be understood as expressing feature clusters associated with event-schemata, and provides the structural basis for instantiating a causal net as described in the previous sections.

### 10.4.1 DATA MODEL OVERVIEW

The data model is designed such that MATCH graph queries on ‘NODE\_K’ vertices serve as the basis for returning a subgraph. ‘NODE\_K’ vertices represent locations on a centreline network that may be understood as the approximate location of a behaviourally significant action. While a routing engine would not have to process this k-level network in the actual route calculation, it would have to return a sequence of nodes that logically determine the labelled sequence of ‘NODE\_K’ nodes for the journey context. The graph vertices in this high-level data model can be understood as representing the actual information resources from which to configure navigational information.

‘ACTION\_REGION’ is an explicit class of graph vertex in the model. While the relationship between ‘NODE\_K’ and ‘ACTION\_REGION’ is labelled ‘ACTIVATES’, really in the data model this can simply be understood as ‘LOCATION\_OF’. An ‘ACTION\_REGION’ can contain any type of feature from any level of abstraction.

An important point to highlight here is that while there is no PATH\_REGION label in the data model, ‘SEGMENT\_K’ can be understood as equivalent to PATH\_REGION, with features of any type and from any level of abstraction potentially being linked to a ‘SEGMENT\_K’ vertex via ‘HAS\_ADJACENT’, i.e. features adjacent to the traversed path of the route. ‘SEGMENT\_K’ features are

bound by ‘NODE\_K’ features, and so a movement from the location of one ‘NODE\_K’ feature to another is a movement along the path represented by a ‘SEGMENT\_K’ feature.

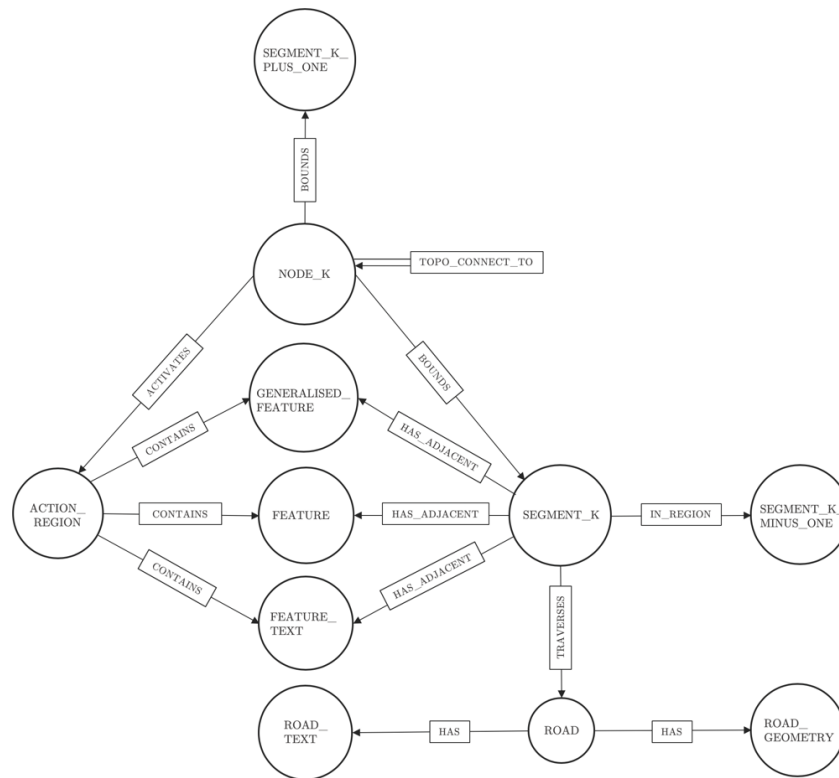


Figure 97 Data model.

The feature type ‘SEGMENT\_K\_MINUS\_ONE’ is simply a polyline feature within a high detail routing network so high detail paths can be visualised. In other words while ‘SEGMENT\_K’ refers to polyline segments from a centreline network, ‘SEGMENT\_K\_MINUS\_ONE’ refers to polyline segments that may implicitly represent a road crossing, a pavement or some other pedestrian pathway. Features of type ‘SEGMENT\_K\_MINUS\_ONE’ are not bound by ‘NODE\_K’ features, but can be logically understood at being within a PATH\_REGION that is bound by two ‘NODE\_K’ features.

Features of type ‘SEGMENT\_K\_PLUS\_ONE’ are features that comprise the representation public transport networks and so are at a higher level of schematisation than type ‘SEGMENT\_K’. These features are however still bound by ‘NODE\_K’ features.

The representation of roads simply follows the standard approach in which a number of distinct ‘ROAD\_GEOMETRY’ polygon objects can be associated with the same ‘ROAD’ object. Similarly, a number of ‘ROAD\_TEXT’ point objects can be associated with the same ‘ROAD’.

In terms of the overall representation of a journey, clearly the ORIGIN and DESTINATION are not represented by vertices in the underlying graph. The relationship between these temporary vertices and the underlying data model is shown in Figure X. PHASE\_REGION objects are also created based on the current journey context, with the overall importance of the k-level route network in the overall model should be clear from Figure 97.

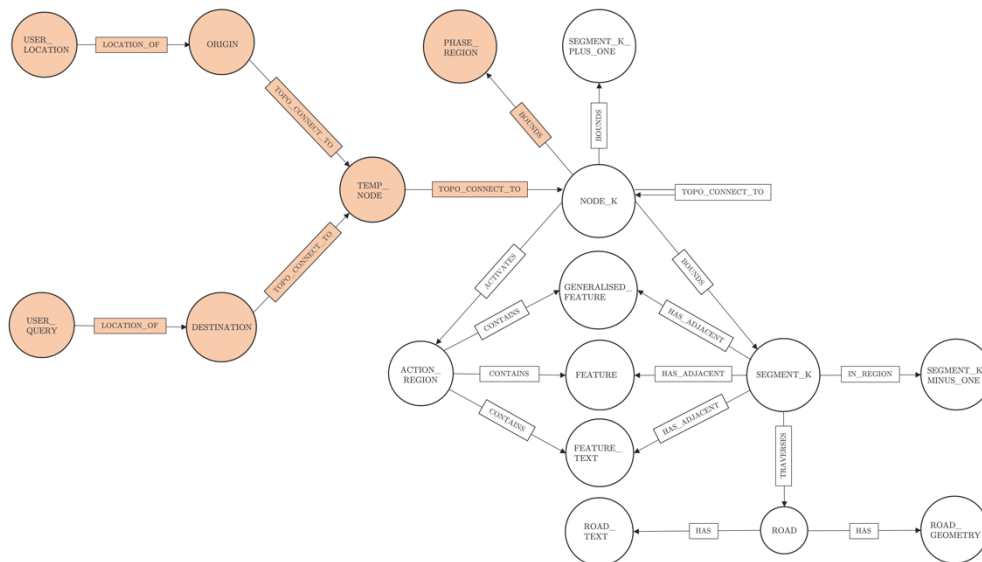


Figure 98 Graph data model (white), plus temporary vertices and edges for the current session (orange).

Returning to the event-schemata presented in the analysis, subgraphs based on graph patterns using this data model presented here can express any event schema.

#### 10.4.2 FEATURE VERTEX CLASSES

In the class definitions, data structures associated with vertex geometry properties  $v_{\text{geometry}}$  are not specified. This property should simply be interpreted in the context of the class – i.e. a point geometry for ‘TEXT\_FEATURE’, ‘ROAD\_TEXT’ and

‘NODE\_K’, a polygon geometry for ‘FEATURE’, ‘GENERALISED\_FEATURE’ and ‘ROAD\_GEOMETRY’, and a polyline geometry for network segments.

class: **GENERALISED\_FEATURE**

**Definition:** *Given the subset of vertices storing generalised polygon features  $GF \in \mathcal{V}$ ,  $v_{GF} \in GF$  is itself a set, where:*

$$v_{GF} = \{ GF_{label}, v_{properties} \}. \quad (10.23)$$

Label:  $GF_{label} = \text{“GENERALISED\_FEATURE”}$

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges: None

class: **FEATURE**

**Definition:** *Given the subset of vertices storing polygon features  $F \in \mathcal{V}$ ,  $v_F \in F$  is itself a set, where:*

$$v_F = \{ F_{label}, v_{properties} \}. \quad (10.24)$$

Label:  $F_{label} = \text{“FEATURE”}$

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges: None

class: **TEXT\_FEATURE**

**Definition:** *Given the subset of vertices storing text features  $TF \in \mathcal{V}$ ,  $v_{TF} \in TF$  is itself a set, where:*

$$v_{TF} = \{ TF_{label}, v_{properties} \}. \quad (10.25)$$

Label:  $TF_{label} = \text{“TEXT\_FEATURE”}$

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges: None

class: **ROAD**

**Definition:** *Given the subset of vertices storing road features  $R \in \mathcal{V}$ ,  $v_R \in R$  is itself a set, where:*

$$v_R = \{ R_{label}, v_{properties} \}. \quad (10.26)$$

Label:  $R_{label} = \text{"ROAD"}$

Properties:  $v_{properties} = \{ v_{id}, v_{textString} \}$

Out edges:

$(\text{ROAD})\text{--}[\text{HAS}]\text{--}(\text{ROAD\_GEOMETRY});$

$(\text{ROAD})\text{--}[\text{HAS}]\text{--}(\text{ROAD\_TEXT});$

class: **ROAD\_TEXT**

**Definition:** *Given the subset of vertices storing road text features  $RT \in \mathcal{V}$ ,  $v_{RT} \in RT$  is itself a set, where:*

$$v_{RT} = \{ RT_{label}, v_{properties} \}. \quad (10.27)$$

Label:  $RT_{label} = \text{"ROAD\_TEXT"}$

Properties:

$v_{properties} = \{ v_{id}, v_{textString}, v_{geometry} \}$

Out edges: None

class: **ROAD\_GEOMETRY**

**Definition:** *Given the subset of vertices storing road geometry features  $RG \in \mathcal{V}$ ,  $v_{RG} \in RG$  is itself a set, where:*

$$v_{RG} = \{ RG_{label}, v_{properties} \}. \quad (10.28)$$

Label:  $RG_{label} = \text{"ROAD\_GEOMETRY"}$

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges: None

## 10.4.3 NETWORK VERTEX CLASSES

As discussed the `NODE_K` and `SEGMENT_K` classes form the fundamental basis of the overall model, with addition segment classes provided to support the visualisation of both lower and higher detail paths.

class: **NODE\_K**

**Definition:** *Given the subset of vertices storing  $k$ -level network nodes  $NK \in \mathcal{V}$ ,  $v_{NK} \in NK$  is itself a set, where:*

$$v_{NK} = \{ NK_{label}, v_{properties} \}. \quad (10.29)$$

Label:  $NK_{label} = \text{"NODE\_K"}$

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges:

ACTIVATES (`NODE_K` →  
                  `ACTION_REGION`);

BOUNDS (`NODE_K` → `SEGMENT_K`);

BOUNDS (`NODE_K` →  
                  `SEGMENT_K_PLUS_ONE`);

class: **SEGMENT\_K**

**Definition:** *Given the subset of vertices storing  $k$ -level network segments  $SK \in \mathcal{V}$ ,  $v_{SK} \in SK$  is itself a set, where:*

$$v_{SK} = \{ SK_{label}, v_{properties} \}. \quad (10.30)$$

Label:  $SK_{label} = \text{"SEGMENT\_K"}$

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges:

(`SEGMENT_K`)—[`TRAVERSES`]→ (`ROAD`);

(`SEGMENT_K`)—[`IN_REGION`]→  
                  (`SEGMENT_K_MINUS_ONE`);

(`SEGMENT_K`)—[`HAS_ADJACENT`]→

(FEATURE);

(SEGMENT\_K)–[HAS\_ADJACENT]→  
(GENERALISED\_FEATURE);

(SEGMENT\_K)–[HAS\_ADJACENT]→  
(TEXT\_FEATURE);

class: **SEGMENT\_K\_MINUS\_ONE**

**Definition:** *Given the subset of vertices storing k-1-level network segments  $SKM1 \in \mathcal{V}$ ,  $v_{SKM1} \in SKM1$  is itself a set, where:*

$$v_{SKM1} = \{ SKM1_{label}, v_{properties} \}. \quad (10.31)$$

Label:  $SKM1_{label} =$

“SEGMENT\_K\_MINUS\_ONE”

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges: None

class: **SEGMENT\_K\_PLUS\_ONE**

**Definition:** *Given the subset of vertices storing k+1-level network segments  $SKP1 \in \mathcal{V}$ ,  $v_{SKP1} \in SKP1$  is itself a set, where:*

$$v_{SKP1} = \{ SKP1_{label}, v_{properties} \}. \quad (10.32)$$

Label:  $SKP1_{label} =$

“SEGMENT\_K\_PLUS\_ONE”

Properties:  $v_{properties} = \{ v_{id}, v_{geometry} \}$

Out edges: None

#### 10.4.4 ACTION REGION VERTEX CLASS

While the SEGMENT\_K class forms the practical basis of selecting feature clusters corresponding with a PATH\_REGION, ACTION\_REGION is provided as a distinct class, along with associated out edge types.

class: **ACTION\_REGION**

**Definition:** *Given the subset of vertices storing action regions  $AR \in \mathcal{V}$ ,  $v_{AR} \in AR$  is itself a set, where:*

$$v_{AR} = \{ AR_{label}, v_{properties} \}. \quad (10.33)$$

Label:  $AR_{label} = \text{"ACTION\_REGION"}$

Properties:  $v_{properties} = \{ v_{id} \}$

Out edges:

(ACTION\_REGION)–[CONTAINS]→  
(GENERALISED\_FEATURE);

(ACTION\_REGION)–[CONTAINS]→  
(FEATURE);

(ACTION\_REGION)–[CONTAINS]→  
(FEATURE\_TEXT);

#### 10.4.5 EXAMPLE GRAPH TRANSACTIONS

Example queries are provided for illustration.

Constructing action region subgraphs was implemented based on the spatial distance of features from the action region node location (see Figure 91).

Query example, cypher run from function ‘FEATURE\_SEARCH’:

```
(“CALL spatial.closest('layer', {lat: $lat, lon: $lon}, $distance) YIELD
node WHERE NOT (node:NK)
RETURN COLLECT({ar_id: $ar_id, feature: node})", lat=lat, lon=lon,
ar_id=ar_id, $distance=distance)
```

The query above selects features within the distance specified by the ‘distance’ variable from each ACTION\_REGION and returns the id of each ACTION\_REGION along with the vertex objects representing the selected features.




Figure 99 Illustration of ACTION\_REGION and PATH\_REGION spatial extent.

Query example, cypher run from function ‘CONSTRUCT\_ACTION\_REGIONS:

```
(“MATCH (ar:AR { id: $source_id }), (feature { id: $target_id })
MERGE (ar)-[edge:CONTAINS_FEATURE]->(feature)
RETURN ar, type(edge), feature”, source_id=source_id,
target_id=target_id)
```

The query above then takes the source and target id’s from the action regions and their associated feature id’s, as returned from the ‘FEATURE\_SEARCH’ function, and creates new edge objects in the graph. In other words, this query constructs action region subgraphs given the target features returned from ‘FEATURE\_SEARCH’ for each action region vertex.

#### 10.4.6 EXAMPLE GRAPH TRANSACTIONS FOR A GIVEN JOURNEY CONTEXT

Returning features in a phase region was simply implemented based on a bounding box determined by the location of primary actions in the journey context.

Query example, cypher run from function ‘SPATIAL\_QUERY:

```
(“CALL spatial.bbox('layer',{lon: $ll_lon, lat: $ll_lat }, {lon:
$str_lon, lat: $tr_lat})
YIELD node WHERE (node:VML_POINTS)
RETURN COLLECT(node)", ll_lon=lower_left_lon, ll_lat=lower_left_lat,
tr_lon=top_right_lon, tr_lat=top_right_lat)
```

This query takes the coordinates of the bounding box of a phase region and returns features within the region as defined. Here the example returns VML\_POINTS, i.e. Ordnance Survey Vector Map Local point features (Ordnance Survey, 2019).

Below are example queries to return features based on graph patterns, given a journey context. The ‘NODE\_K’ sequence is stored in an array variable, with topological distance from ‘NODE\_K’ graph vertices determining the results

```

pattern_one_nk_vertices = tx.run("WITH $route AS arr
MATCH pattern_one=(i)-[]->(j)
WHERE i.id IN arr
RETURN COLLECT(DISTINCT(i))", route=route)

pattern_one_child_vertices = tx.run("WITH $route AS arr
MATCH pattern_one=(i)-[]->(j)
WHERE i.id IN arr
RETURN COLLECT(DISTINCT(j))", route=route)

pattern_one_paths = tx.run("WITH $route AS arr
MATCH pattern_one=(i)-[]->(j)
WHERE i.id IN arr
RETURN relationships(pattern_one), i, j", route=route)

pattern_two_children_of_children_vertices = tx.run("WITH $route AS arr
MATCH pattern_two=(i)-[]->(j)-[]->(k)
WHERE i.id IN arr
RETURN COLLECT(DISTINCT(k))", route=route)

pattern_two_paths = tx.run("WITH $route AS arr
MATCH pattern_two=(i)-[]->(j)-[]->(k)
WHERE i.id IN arr
RETURN relationships(pattern_two), i, j, k", route=route)

```

## 10.5 Concluding comments

This chapter has presented a formal basis for automatically selecting context-specific map content. The graph data model and overall causal framework are argued to produce selections across multi-scale view space that reflect the information needs of users as observed throughout the empirical studies. The final output after a function has been called to map features to the numerical zoom can be understood as the resource profile of cognitive extension following the analysis presented in the previous chapter. Both the analysis and the techniques presented in this chapter will now be considered in a brief discussion.



# 11 Discussion

Reflecting on the research findings and the formalised approach presented in the previous chapter, here additional discussion is provided around a number of key issues across the underlying data model, the integration of additional forms of data, and a number of considerations with respect to cartographic outputs.

## 11.1 Data, modeling and use cases

Here a brief discussion is provided with respect to the underlying data, to modeling data across a range of use cases and also to the possibility for further research to explore the use of likelihood evidence as the basis of automated reasoning.

### 11.1.1 DETERMINING MEMBERSHIP IN ACTION REGION AND PATH REGION SUBGRPAHS

As discussed in the previous chapter, the approach taken was simply to create ACTION\_REGION and PATH\_REGION subgraphs by using spatial queries based on the location of NODE\_K objects, and the position of SEGMENT\_K objects respectively. In the context of the present research this simple ‘buffer’ based approach is argued to yield adequate results. This should however be a focus for further research. Questions clearly include how to determine the appropriate distance of the buffers given a range of geographical forms. It is argued however that a more complex issue is how to deal with the location of point features. While point features play such a central role in urban navigation with mobile devices, the location of text is distributed such that point features are not tightly clustered. While the motivation for this is clear in terms of standard approaches, it may mean that a label associated with a building adjacent to an intersection is not included in a selection, for example (i.e. where the label is centred in the overall polygon but is

relatively far away from the path of the network). Overall it would seem that the appropriate approach would be to determine membership of subgraphs stored in memory using topological methods, however this required further consideration.

Linked to this issue is the more general problem of scale. To specifically select features that are adjacent to a significant action naturally means that tightly grouped clusters are selected which cannot all be visualised at small scales irrespective of whether they are known to be of particular relevance. Two areas that require further research on this point are a) how to determine which features in the cluster should have primacy in the scale hierarchy when the features are in fact all of the same type, and b) how to better visualise labels of tightly clustered features at small scales. To the latter point, a key aspect of the approach results in the ability to limit the amount of generic information over the broader geographical extent. This area could be made better use of by allowing route-specific labels to ‘float’ over this wider area. The extent to which this would be appropriate in practice however is not known and requires further exploration.

#### 11.1.2 DATA INTEGRATION AND ALTERNATE USE CASES

On reflection the general use case of urban travel does not perhaps demonstrate the true power of the approach as by definition there is not sufficient basis for tailoring the content to a more specific user group. It is argued the specification in the model chapter is highly flexible, and the real potential of the approach would be more clearly demonstrated with more specific use cases. The only prerequisite for a system to be developed based on the approach is for there to be data describing the approximate location of significant actions, or at the least the location of likely event boundaries. From there the  $k$ -level routing network can be updated and the rest of the framework implemented largely ‘as is’. To give an example, Prandi et al. presented work to enrich geo-data with attributes of specific relevance for disabled travellers (Prandi *et al.*, 2014). If such data is known, then creating a more tailored  $k$ -level network representation is straight-forward. While an additional task would then be to specify the precedence of objects in context-specific scale hierarchies. For example, in the case of a wheelchair user perhaps an object corresponding with the SURFACE schema might have greater significance than a TURN action. Such refinements based on use-case are argued however to be a low barrier to scaling the approach.

### 11.1.3 IMPLICATIONS OF INCORPORATING LIKELIHOOD EVIDENCE

Evidence as encoded in propagation matrices in the specification was in the form of ‘finding’ evidence, i.e. binary vectors that determine patterns of activation over the net. There is no reason however that the approach could not be adapted to incorporate ‘likelihood’ evidence that not only encodes pattern of activation but also expresses magnitudes. A graphical example would be to encode gradations in hue such that, for example, features with lesser roles in event-schemata were displayed with less visual prominence. From a research perceptive however, a less superficial example would be to include some explicit basis of ‘cue competition’ during propagation (Chen *et al.*, 2017). Research is on-going into processes with Bayesian characteristics that shape spatial cognition (Madl, 2016), and there may perhaps be a way to express ‘weighting’ in the selection of cartographic features by building on the approach proposed in this thesis.

## 11.2 Cartography and interaction

### 11.2.1 GENERALISATION AND ABSTRACTION

A key question over the course of research was whether ‘mixing’ entities from different levels of abstraction might lead to new approaches to cartographic generalisation. In practice it was found that on mobile devices, detailed polygon features play a significantly diminished role in representation, with the focus being on polyline and point features, set within a rough geographical context. While there is perhaps an argument for grouping features based on the extent and overlap of ACTION\_REGION and PATH\_REGION clusters, it is difficult to assess to the practical impact of this approach on actual task performance. Overall the approach in Google Maps of generalising based on commercial and residential areas was strongly supported however. In terms generalisation techniques that would be of greater significance relative to task performance, approaches for determining hierarchical precedence of closely clustered point features would seem to be an area for further research.

### 11.2.2 INTERACTION CONCEPTS

Three interaction concepts are presented as starting point for further research. While none of these concepts have been tested with users, they are all based on the observations and analysis made during the course of the present research.

Overall these concepts demonstrate a number of practical possibilities following the approach presented in the preceding chapters, including: a) styling on individual features, b) construction of context-specific scale hierarchies, c) the disambiguation of feature subsets and ‘crossfading’ between disambiguated and generic selections, d) non-uniform abstraction for both polyline and polygon features within single views, e) creating explicit relations between cartographic content and additional content, and f) the sequential representation of feature subsets.

### Interaction Concept 1: CROSSFADE

While research has shown that context specific selections can support navigation, the optimal approach is to provide user control over content filtering.



Figure 100 ‘Crossfade’ interaction concept at two zoom levels: (left) context-specific selection and (right) generic content fading into view.

In the ‘crossfade’ concept illustrated in Figure 100 the user is shown the predicted selection for their journey context but has the ability to ‘fade in’ detail from across the broader geography. This interaction concept is based around the sound mixing ‘crossfader’ control that moves between two input sources. Sense of control was a key issue when providing highly specific sections, and clearly if a traveller does not for some reason take the exact recommended route, some way of returning to a more generic view is required.

### Interaction Concept 2: CONTEXTUAL ZOOM

In addition to supporting context specific selections at multiple scales, the approach also supports the integration of map content with additional hierarchically ordered content. While the amount of text is clearly an issue it is argued there was sufficient evidence for this concept to be explored further. Here then, nested content is accessible by a single zoom interaction operating on both cartographic content and textual descriptions, as shown in Figure 101.

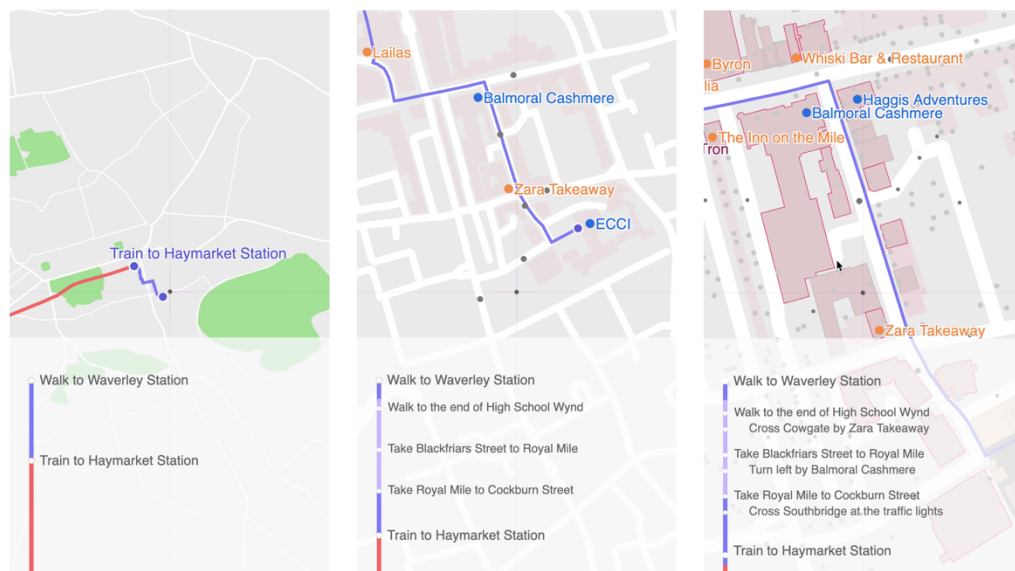


Figure 101 Contextual zoom interaction concept example. Single interaction decomposes both graphical and textual information.

### Interaction Concept 3: SEQUENTIAL ACTIVATION

Map users naturally understand journeys as sequences of directed paths that define distinct regions of features. In the sequential activation concept, the user is able to cycle through this ordered sequence of collections and see semantic disambiguation of activated features for the current path region, as illustrated in Figure 102.







# 12 Conclusion

The core research aim and objectives are now revisited with summary comments. Each will be addressed in turn, followed by some final concluding remarks.

## 12.1 Core research aim and objectives

### 12.1.1 CORE RESEARCH AIM

**To develop a framework for the automatic selection and display of a subgraph of features, where the resulting subgraph supports information needs over a set of navigational task-contexts for a given journey.**

The overall approach that is argued to deliver on the core research aim consists of an underlying graph architecture that supports the expression of *latent context*. It was found that while subgraphs can be stored in memory that correspond to feature clusters relevant to goal acquisition across *potential* experiential events, the overall regionalisation of the journey extent and the structure of the partonomic representation of the episode cannot be ‘stored’ and must be constructed given a particular journey context.

It was also found that the correct approach to content selection recognises the fundamentally incremental nature of knowledge acquisition on small devices, with the practical implication being the selection of content not just for a single integrated view, but for a multi-scale information space that supports any view a user may require throughout this space.

The framework developed over the course of the research includes a flexible high-level graph data model, and a number of algorithms for processing subgraphs in a way that loosely approximates attentional processes given a task-hierarchy. This was found to be best understood in terms of a resource profile of cognitive extension,

where the selection of ‘informational resources’ is achieved by representing the underlying geospatial features as variables in a causal net.

### 12.1.2 INITIAL HYPOTHESIS

The initial hypothesis was found to be flawed. While the prototype formalisation of natural spatial knowledge (S1) performed well in terms of approximating an *a posteriori* cognitive representation of the geography (i.e. by predicting the structure and content of sketch maps), it was found that the notion of a ‘short-cut’ to spatial knowledge acquired ‘after the fact’ did not effectively capture the information needs of the traveller over the course of the journey itself. This hypothesis was then rearticulated in terms of predicting the state of a coupled system across the event structure, as opposed to predicting a representation at a single ‘snapshot’ in time.

### 12.1.3 OBJECTIVES THEME 1: EMPIRICAL RESEARCH

1a)

**Which aspects of our knowledge of spatial cognition are of practical significance given the central task of automatically selecting map content to support multi-modal travel?**

The sketches and textual directions demonstrated that in the general case, participants shared a cognitive representation of the journey that was highly stable across preferences and behaviour types. This representation broadly followed the expected event structure, with primary actions serving to segment events near the ‘top’ of the partonomy, and distinct events almost always forming around the structure of the navigable network.

Event segmentation and the natural subdivision of experience into action sequences is argued to be a powerful approach in terms of analysing behaviour with a view to linking map content with journey context. Of central concern was the selection of data across scales and levels of generalisation. The study of event segmentation and the analysis of action sequences in the case of multi-modal travel was key in forming a highly flexible theoretical foundation. Similarly, image-schema theory was found to be a powerful basis for developing formalisations that make explicit the observed concepts, entities and relations that support spatial goal acquisition. The flexible nature of image-schemata meant they could be used as a way to express structures that are nested, that exist at different scales, and that support integrated

representations from varying levels of abstraction. Linked to existing knowledge of event boundaries and cognitive topology, knowledge around anchor-points and hierarchical reference memory also proved to be significant.

While it is argued the hypothesis was validated that it is possible to predict the *a posteriori* cognitive representation of the journey in advance in the general case, the usefulness of this was found to be limited to the selection of features and feature clusters as a means of focusing attention. Explicitly reflecting the natural tendency for understanding to lead to distortions of direction and distance between local regions was not found to be an effective means of integrating ‘local detail’ and ‘global context’ in the same view. Also, it was found that limiting selections to features that strictly match egocentric route knowledge was not effective, with the final approach lying somewhere between these purposefully extreme experimental scenarios and the prevailing design paradigm that emphasises views at a ‘medium’ level of detail.

Overall the ‘agent’ conceptualisation that formed the basis of the initial hypothesis was found to be a useful device in the development of the approach, however the model presented in this thesis is not ‘agent-based’ per se. That being said, the causal nature of the relationship between task and attention was key, with a probabilistic view of attention, comprehension and memory serving as the fundamental basis for determining traveller information needs. Overall, it is argued there is a strong connection between the approach proposed here for determining map content and previous work in cognitive robotics to automatically determine semantic understanding of spatial experience via sensor data.

1b)

**Are any of the entities that are elements in distinct events or event sequences over navigational episodes additional to the standard data that is available to us?**

No strong evidence was found for non-standard graphical representations, and in fact the central role of simple point features was underestimated during the initial stages of the research. It was found that high detail polygon features play a significantly diminished role on mobile devices, with users often not even ‘zooming’ far enough to reach scale transition points at which these features become visible. Beyond the expected observation of a need for relatively higher schematisation during journey phases traversed on public transport, a key observation in terms of representation was the significance of point features across the scale hierarchy.

In terms of the entities themselves, while evidence was found for the importance of small-scale features that relate to the environment's physical affordances, for example road surfaces and road crossings, overall there was little evidence of a need to represent entities that would be considered non-standard in terms of current data sets. That being said, the user preference for accessing information at low or medium levels of detail does require that generalisation of the features in the broader geographical context is a key focus. For example, the importance of distinguishing regions based on perceptual distinctions such as being 'commercial' or 'residential' areas, versus conveying accuracy in the geometry of region boundaries.

In addition to discrete entities, another consideration over the course of studies was the explicit representation of continuous processes. This was however not observed to be of relevance beyond the representation of the path of the route. Providing location information on approaching buses was not something included in prototype research stimuli but was clearly something that would be of value, however even in this case it is arguable as to whether the current location of a bus is really the representation of a continuous process. Overall, the strong conclusion was that the focus should remain on discrete entities that may be specified in terms of set membership.

1c)

### **What should be the basis of clustering with respect to entities that play a role in multi-modal navigation?**

It was found that events in the context of multi-modal urban travel can clearly be categorised into distinct classes, and that the nature of attentional processes can be anticipated for each type to an appropriate extent – for example confirmation as to the 'shape' of the region of influence of a transition to public transport versus a transition from public transport back to pedestrian path traversal.

The application of image schema theory is argued to have been successful with regard to modeling the features that are likely to play roles within significant events. Within the context of selecting features for display on mobile devices however, the issue of scale in terms of these tightly clustered groups of features meant that in practice it was difficult to assess an approach to determining feature-event membership beyond a simple spatial buffer.

The informational resource profile for urban navigation was found to include features that provide allocentric spatial context, even in the case where a highly context-dependent selection of features that are visible from the path of the route are disambiguated. From the perspective of clustering then, it was found that there

are features that play a behaviourally significant role but are not strictly located ‘within’ the localised regions that form around experiential events, which leads to a distinction between clusters of features formed at the level of actions (and directed paths between distinctive consecutive actions), and the features that can be considered to be within a broader ‘phase’ region but are still of significance in the development of hierarchical reference memory.

In terms of the issue of scale, it was found that supporting incremental knowledge acquisition over time is a primary aim, and so automated methods of clustering features based on context need to support multiple views across view space, as opposed to leading to single integrated views that capture the information ‘all at once’.

#### 12.1.4 OBJECTIVES THEME 2: DATA MODELLING, SELECTION AND DISPLAY

2a)

**Are any additional entities required to support data modelling that are not explicitly present in the empirical data?**

Overall there was very little evidence found for the need to formally represent entities that are not in standard datasets. In terms of the graph model, beyond the explicit specification of an ACTION\_REGION vertex class, graph vertices essentially represent standard cartographic objects. Really the key data structure that was required in the approach that is non-standard is the  $k$ -level network that includes the locations of potential actions. This data structure has more in common with techniques used in the cognitive robotics research discussed Chapter 4 than with cartographic techniques, and so this could be viewed as the ‘novel’ component of the approach in terms of the underlying data.

It was found that by including the locations of potentially significant actions on an adapted centreline network, if a routing engine returned a result that included an ordered sequence of labelled nodes, then this could form the basis of predicting reasonable approximations to an event paronomy for a given journey. Arguably the  $k$ -level routing network structure is in fact ‘visible’ in the sketch and textual data, and so the only entities that are required but are not present in the sketch data are generalised features in broader PHASE\_REGION extents that support the development of hierarchical reference memory.

2b)

**What are the relations or orderings between entities that allow them to be surfaced in context-specific selections via graph queries?**

Approximate locations of behaviourally significant actions are argued to be the correct basis for integrating distinct forms of information, echoing the initial approach proposed in (Mackaness, Quigley and Tanasescu, 2011). The intuitive connection between graph-based approaches to knowledge representation and the formalisation of spatial knowledge was found to be a significant benefit in the development of the approach. In terms of the use of a graph database, an alternative approach was not experimented with during the course of research, so it is difficult to provide comparative comments. That being said, the approach to query processing is very much graph-centric, and given that within the framework the entire purpose of returning a query from the database is to instantiate a causal net, it would seem that graph structures are the natural choice for the underlying data architecture.

In terms of the relations and orderings that allow features to be surfaced in context-specific groupings, the role of the  $k$ -level network was critical. Additionally, the finding that feature clusters can be determined across ACTION\_REGIONS, PATH\_REGIONS and PHASE\_REGIONS can be understood as the fundamental basis for returning candidate features using a combination of both graph and spatial queries.

Following the discussion in Chapter 4 around the tree, the semilattice and the rhizome, it was found that it seems necessary in systems that support the expression of naturalistic spatial context for there to be more than one path to certain types of graph vertices, as illustrated in the data model (i.e., *not* a hierarchical compound graph). Overall then, the research presented here supports the long-standing argument that the underlying structure that is best suited to expressing representations which reflect situated experience is the semilattice.

2c)

**How can interactive maps for multi-modal travel be automatically generated from data selected from the underlying database?**

It was found that to effectively reflect context, data selection should occur through two distinct stages which are argued to reflect the distinction between the dual stages of unconscious and conscious cognitive control. This ‘dual process’ includes

an initial stage that is not limited by capacity (or at least, less limited) in which a relatively large number of objects are selected as candidates, followed by a processing stage that is analogous to conscious attention. In practice this involves selecting a subgraph based on graph patterns, using the subgraph to instantiate a causal net, and then processing this data structure using causal reasoning given the specific character of the current journey.

In terms of the presentation of information to the user, it was found that free movement in interactive multi-scale view space is optimal on a mobile device, and that pan and zoom interactions play an important role in the situated acquisition of spatial knowledge. The approach presented in this thesis addresses this by showing how to determine content throughout view space, with views being conceptually understood as lateral slices across this 3D structure.

## **12.2 Final comments**

It has been shown that by formalising a representation of urban geography based on the potential role features may play in extended cognitive processes, the underlying data can be configurable in a way that supports context-specific selection across multi-scale information space. Reflecting on Alexander's description of a semilattice as supporting 1,000,000 potential subsets from a structure consisting of 20 elements, the overall conceptual approach and associated methods presented in this thesis show how this type of flexibility may be achieved in providing navigational information, based only on very superficial knowledge regarding the user and their current journey. Following the observation that rigid knowledge structures are born of convenience to the designer as opposed to the traveller, it is hoped the work presented here demonstrates how we might better support peoples' natural cognitive processes as they traverse complex urban environments



# References

- Agrawala, M. and Stotle, C. (2001). Rendering effective route maps: Improving usability through generalization. In Fiume, E. (eds), *SIGGRAPH '01 Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 241–249.
- Alexander, A.S. and Nitz, D.A. (2017). Spatially Periodic Activation Patterns of Retrosplenial Cortex Encode Route Sub-spaces and Distance Traveled. In *Current Biology*, 27(11), 1551-1560.
- Alexander, C. (1966). A city is not a tree. In *Architectural Forum*, 122(1), 58–62.
- Allen, J. F. (1981). An Interval-Based Representation of Temporal Knowledge. In Hayes, P.J. (eds), *Proceedings of the International Joint Conference on Artificial Intelligence*, Vancouver B.C., 221–226.
- Amant, R. S., Morrison, C.T., Chang, Y.H., Cohen, P.R. and Beal, C.R. (2006). An image schema language. In Fum, D. (eds), *Proceedings of the International Conference on Cognitive Modeling*, Trieste, 634–640.
- Amer, T., Campbell, K.L. and Hasher, L. (2016). Cognitive Control As a Double-Edged Sword. In *Trends in Cognitive Sciences*, 20(12), 905–915.
- Anderson, B. and McFarlane, C. (2011). Assemblage and geography. In *Area*, 43(2), 124–127.
- Appleyard, D. (1970). Styles and Methods of Structuring a City. In *Environment and Behaviour*, 2, 100–117.
- Baddeley, A. (2007). *Working Memory, Thought and Action*. Oxford University Press.

- Balleine, B.W. (2018). The Motivation of Action and the Origins of Reward. In Morris, R., Bornstein, A., and Shenhav, A. (eds), *Goal-Directed Decision Making: Computations and Neural Circuits*, 429-455. Academic Press.
- Bar, M., Aminoff, E. and Schacter, D.L. (2008). Scenes unseen: The parahippocampal cortex intrinsically subserves contextual associations, not scenes or places per se. In *Journal of Neuroscience*, 28(34), 8539–8544.
- Barsalou, L.W. (1999). Perceptual symbol systems. In *The Behavioral and Brain Sciences*, 22(4), 577–660.
- Barsalou, L.W. (2008). Grounded Cognition. In *Annual Review of Psychology*, 59(1), 617-645.
- Bartram, D.J. (1980). Comprehending spatial information: The relative efficiency of different methods of presenting information about bus routes. In *Journal of Applied Psychology*, 65(1), 103–110.
- Bateman, J., Borgo, S., Luettich, K., Masolo, C. and Massakowski, T. (2007). Ontological Modularity and Spatial Diversity. In *Spatial Cognition and Computation*, 7(1), 97–128.
- Bertin, J. (2011). *Semiology of Graphics*. Esri Press (re-print).
- Bittner, T., Donnelly, M. and Smith, B. (2004). Endurants and perdurants in directly depicting ontologies. In *AI Communications*, 17(4), 247–258.
- Bjørndahl, J.S., Fusaroli, R., Østergaard, S. and Tylén, K. (2019). Thinking together with material representations: Joint epistemic actions in creative problem solving. In *Cognitive Semiotics*, 7(1), 103–123.
- Bor, D., Duncan, J., Wiseman, R.J. and Owen, A.M. (2003). Encoding strategies dissociate prefrontal activity from working memory demand. In *Neuron*, 37(2), 361–367.
- Bormann, R., Jordan, F., Li, W., Hampp, J. and Hägele, M. (2016). Room segmentation: Survey, implementation, and analysis. In Shademan, A. et al. (eds), *Proceedings of the IEEE International Conference on Robotics and Automation*, 1019–1026.

- Botvinick, M.M., Niv, Y. and Barto, A.C. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement perspective. In *Cognition*, 113(3), 262–280.
- Bradley, M. M. (2009). Natural selective attention: Orienting and emotion. In *Psychophysiology*, 46(1), 1–11.
- Brown, B., McGregor, M. and Laurier, E. (2013). iPhone in vivo: video analysis of mobile device use. In Mackay, W.E. (eds), *Proceedings of CHI 2013*, 1031–1040.
- Brunec, I.K., Moscovitch, M. and Barense, M.D. (2018). Boundaries Shape Cognitive Representations of Spaces and Events. In *Trends in Cognitive Sciences*, 22(7), 637–650.
- Brunet, R. (1980). La composition des modèles dans l'analyse spatiale. In *L'Espace Géographique*, 4, 253–265.
- Brunet, R. (1986). La carte-modèle et les choremes. In *Mappemonde*, 4, 1–5.
- Bullens, J., Nardini, M., Doeller, C.F., Braddick, O., Postma, A. and Burgess, N. (2010). The role of landmarks and boundaries in the development of spatial memory. In *Developmental Science*, 13(1), 170–180.
- Cablitz, G.H. (2001). Geographical categories: An ontological investigation. In *International Journal of Geographical Information Science*, 15(7), 591–612.
- Car, A. and Frank, A. (1994). General principles of hierarchical spatial reasoning - the case of wayfinding. In Waugh, T.C. (eds), *Advances in GIS research, Proceedings of the 6th symposium*, Edinburgh, 2, 646–664.
- Car, A and Frank, A. (1994). Modelling a hierarchy of space applied to large road networks. In Nievergelt, J. (eds), *IGIS'94: Geographic Information Systems, Proceedings of International Workshop on Advanced Research in GIS*, Monte Verità, 15–24.
- Carr, S. and Schissler, D. (1969). The City as a Trip: Perceptual Selection and Memory in the View from the Road. In *Environment and Behavior*, 1(1), 7–35.
- Casati, R. and Varzi, A. C. (1997). Spatial Entities. In Stock, O. (eds), *Spatial and Temporal Reasoning*, 73–96. Springer.

- Castelli, L., Latini Corazzini, L. and Geminiani, G.C. (2008). Spatial navigation in large-scale virtual environments: Gender differences in survey tasks. In *Computers in Human Behavior*, 24(4), 1643–1667.
- Ceci, S., Rosenblum, T. and DeBruyn, E. (1999). Laboratory versus field approaches to cognition. In Sternberg, R.J. (eds), *The Nature of Cognition*, 385–408. MIT Press.
- Chatterjee, A. (2010). Disembodying cognition. In *Language and Cognition*, 2(1), 79–116.
- Chen, X., McNamara, T.P., Kelly, J.W. and Wolbers, T. (2017). Cue combination in human spatial navigation. In *Cognitive Psychology*, 95, 105–144.
- Chomsky, N. (1959). On certain formal properties of grammars. In *Information and Control*, 2(2), 137–167.
- Chown, E., Kaplan, S. and Kortenkamp, D. (1995). Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. In *Cognitive Science*, 19(1), 1–51.
- Chronis, G. and Skubic, M. (2004). Robot navigation using qualitative landmark states from sketched route maps. In *Proceedings of the IEEE International Conference on Robotics and Automation 2004*, New Orleans, 1530–1535.
- Clark, A. (1998). *Being there: putting mind, body and brain together again*. MIT Press.
- Clark, A. (2013). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Clark, A. (2015). Radical predictive processing. In *Southern Journal of Philosophy*, 53(S1), 3–27.
- Clark, A. and Chalmers, D. (1998). The Extended Mind. In *Analysis*, 58(1), 7–19.
- Corbetta, M., Patel, G. and Shulman, G.L. (2008). The Reorienting System of the Human Brain: From Environment to Theory of Mind. In *Neuron*, 58(3), 306–324.

- Corbetta, M. and Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. In *Nature Reviews Neuroscience*, 3(3), 201–215.
- Couclelis, H., Golledge, R.G., Gale, N. and Tobler, W. (1987). Exploring the anchor-point hypothesis of spatial cognition. In *Journal of Environmental Psychology*, 7, 99–122.
- Coull, J.T., Frith, C.D., Büchel, C. and Nobre, A.C. (2000). Orienting attention in time: Behavioural and neuroanatomical distinction between exogenous and endogenous shifts. In *Neuropsychologia*, 38(6), 808–819.
- Dabaghian, Y., Mémoli, F., Frank, L., Carlsson, G. (2012). A Topological Paradigm for Hippocampal Spatial Map Formation Using Persistent Homology. In *PLoS Computational Biology*, 8(8): e1002581.
- Dabaghian, Y., Brandt, V.L. and Frank, L.M. (2014). Reconceiving the hippocampal map as a topological template. In *eLife*, 3: e03476.
- Davies, C., Li, C. and Albrecht, J. (2010). Human Understanding of Space. In Haklay, M. (eds), *Interacting with Geospatial Technologies*, 19–35. Wiley-Blackwell.
- Debord, G. and Jorn, A. (1957). The naked city: illustration de l'hypothèse des plaques tournantes en psychogéographie. Rijksbureau voor Kunsthistorische Documentatie.
- Deleuze, G. and Guattari, F. (1987) A thousand plateaus: Capitalism and schizophrenia. Athlone.
- Derdikman, D., Whitlock, J.R., Tsao, A., Fyhn, M., Hafting, T., Moser, M. and Moser, E.I. (2009). Fragmentation of grid cell maps in a multicompartiment environment. In *Nature Neuroscience*, 12, 1325–1332.
- Derdikman, D. and Moser, E.I. (2016). A Manifold of Spatial Maps in the Brain. In *Trends in Cognitive Sciences*, 14(June), 561–569.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. In *Journal of Neuroscience*, 32(11), 3612–3628.
- Dumitru, M.L. and Joergensen, G.H. (2016). Gestalt Reasoning with Conjunctions and Disjunctions. In *PLoS One*, 11:e0151774.

- Dumont, M., Touya, G. and Duchêne, C. (2018). Alternative transitions between existing representations in multi-scale maps. In Brewer, C.A. (eds), *Proceedings of the International Cartographic Conference 2017, ICA*, Washington DC.
- Dumont, M., Touya, G. and Duchêne, C. (2016). A comparative study of existing multi-scale maps: what content at which scale? In *International Conference on GIScience Short Paper Proceedings*, 1(1), 88-91.
- Egenhofer, M.J. and Mark, D.M. (1995). Naive Geography. In Cohn, A.G. and Mark, D.M. (eds), *Proceedings of the Conference on Spatial Information Theory, COSIT 95*, Semmering, 1- 15.
- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. and Tanila, H. (1999). The hippocampus, memory, and place cells: is it spatial memory or a memory space? In *Neuron*, 23(2), 209–226.
- Eichenbaum, H. (2017). On the Integration of Space, Time, and Memory. In *Neuron*, 95(5), 1007–1018.
- Ekstrom, A.D. and Isham, E.A. (2017). Human spatial navigation: representations across dimensions and scales. In *Current Opinion in Behavioral Sciences*, 17, 84–89.
- Epstein, R.A., Patai, E.Z., Julian, J.B. and Spiers, H.J. (2017). The cognitive map in humans: Spatial navigation and beyond. In *Nature Neuroscience*, 20(11), 1504–1513.
- Farrel, S. and Lewandowsky, S. (2018). *Computational Modeling of Cognition and Behaviour*. Cambridge University Press.
- Ferguson, E. L. and Hegarty, M. (1994). Properties of cognitive maps constructed from texts. In *Memory & Cognition*, 22(4), 455–473.
- Field, K., O'Brien, J. and Beale, L. (2011). Paper maps or GPS? Exploring differences in wayfinding behaviour and spatial knowledge acquisition. In Ruas, A. (eds), *Proceedings of the International Cartographic Conference, ICA, 2011*, Paris, 1–8.

- Frank, A.U. and Raubal, M. (2000). Formal Specifications of Image Schemata for Interoperability in Geographic Information Systems. In Nualláin, S. (eds), *Spatial Cognition: Foundations and Applications*, 213–232. John Benjamins Publishing.
- Freksa, C., Barkowsky, T., Dylla, F., Falomir, Z., Olteteanu, A. and van de Ven, J. (2017). Spatial problem solving and cognition. In Zacks, J.M. and Taylor, H.A. (eds), *Representations in Mind and World: Essays Inspired by Barbara Tversky*, 156–183. Routledge.
- Furnas, G.W. (1986). Generalized Fisheye Views. In Mantei, M. (eds), *CHI'86 Proceedings*, 16-23.
- Gallotti, R., Porter, M.A. and Barthelemy, M. (2015). Information measures and cognitive limits in multilayer navigation. arXiv:1506.01978.
- Gallotti, R., Porter, M.A. and Barthelemy, M. (2016). Lost in transportation: Information measures and cognitive limits in multilayer navigation. In *Science Advances*, 2(2):e1500445.
- Godfrey, L. and Mackaness, W. (2017). The bounds of distortion: truth, meaning and efficacy in digital geographic representation. In *International Journal of Cartography*, 3(1), 31–44.
- Godfrey, L. and Mackaness, W. (2018). Artificial Cognitive Maps: Selecting Heterogeneous Sets of Geographic Objects and Relations to Drive Highly Contextual Task-Oriented Map Views. In Fogliaroni et al. (eds), *Proceedings of Workshops and Posters at the 13<sup>th</sup> International Conference on Spatial Information Theory, COSIT 2017, L'Aquila*, 57-62.
- Golledge, R.C., Dougherty, V. and Bell, S. (1995). Acquiring Spatial Knowledge: Survey Versus Route-Based Knowledge in Unfamiliar Environments. In *Annals of the Association of American Geographers*, 85, 134–158.
- Golledge, R.G., Smith, T.R., Pellegrino, J.W., Doherty, S. and Marshall, S.P. (1985). A conceptual model and empirical analysis of children's acquisition of spatial knowledge. In *Journal of Environmental Psychology*, 5(2), 125–152.
- Golledge, R.G. (1999). *Wayfinding Behavior*. John Hopkins University Press.
- Gonçalves Santos, S.B. (2019). Interactive Spider Maps for providing public transports information. Thesis. Universidade Do Porto.

- Gopal, S., Klatzky, R.L. and Smith, T.R. (1989). Navigator: A psychologically based model of environmental learning through navigation. In *Journal of Environmental Psychology*, 9(4), pp.309–331.
- Guo, Z. (2011). Mind the map! The Impact of transit maps on path choice in public transit. In *Transportation Research Part A: Policy and Practice*, 45(7), 625–639.
- Hafting, T., Fyhn, M., Molden, S., Moser, M. and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. In *Nature*, 436(7052), 801–806.
- Han, X. and Becker, S. (2014). One spatial map or many? Spatial coding of connected environments. In *Journal of Experimental Psychology: Learning Memory and Cognition*, 40(2), 511–531.
- Harrie, L. and Stigmar, H. (2010). An evaluation of measures for quantifying map information. In *ISPRS Journal of Photogrammetry and Remote Sensing*, 65(3), 266–274.
- Hebart, M.N., Bankson, B.B., Harel, A., Baker, C.I. and Cichy, R.M. (2018). The representational dynamics of task and object processing in humans. In *eLife*, 7:e32816.
- Hebb, D.O. (1949). *The Organization of Behaviour: A Neuropsychological Theory*. John Wiley & Sons.
- Hedblom, M.M., Kutz, O., Mossakowski, T., and Neuhaus, F. (2017). Between CONTACT and SUPPORT: Introducing a Logic for Image Schemas and Directed Movement. In Esposito, F. et al. (eds), *Proceedings of the 16th International Conference of the Italian Association for Artificial Intelligence (AI\*IA 2017)*, Bari, 256– 268.
- Hedblom, M.M. and Kutz, O. (2015). Choosing the Right Path: Image Schema Theory as a Foundation for Concept Invention. In *Journal of Artificial General Intelligence*, 6(1), 22–54.
- Heersmink, R. (2013). A Taxonomy of Cognitive Artifacts: Function, Information, and Categories. In *Review of Philosophy and Psychology*, 4(3), 465–481.

- Heersmink, R. (2015). Dimensions of integration in embedded and extended cognitive systems. In *Phenomenology and the Cognitive Sciences*, 14(3), 577–598.
- Hirtle, S.C. and Jonides, J. (1985). Evidence of hierarchies in cognitive maps. In *Memory & Cognition*, 13(3), 208–217.
- ICA (1997) Spatial Database Transfer Standards 2: Characteristics for Assessing Standards and Full Descriptions of the National and International Standards in the World. Pergamon.
- Ishikawa, T., Fujiwara, H., Imai, O. and Okab, A. (2008). Wayfinding with a GPS-based mobile navigation system: A comparison with maps and direct experience. In *Journal of Environmental Psychology*, 28(1), 74–82.
- Jahn, G., Knauff, M. and Johnson-Laird, P.N. (2007). Preferred mental models in reasoning about spatial relations. In *Memory & Cognition*, 35(8), 2075–2087.
- Jan, S., Schwering, A., Schultz, C. and Chipofya, M. (2017). Cognitively plausible representations for the alignment of sketch and geo-referenced maps. In *Journal of Spatial Information Science*, 14(2017), 31–59.
- Janarthanam, S., Lemon, O., Bartie, P., Dalmás, T., Dickinson, A., Liu, X., Mackaness, W., and Webber, B. (2013). Evaluating a City Exploration Dialogue System Combining Question-Answering and Pedestrian Navigation. In Scheutze, H. et al. (eds), *Proceedings of the 51st Annual Meeting of the Association of Computational Linguistics*, Sofia, 1660–1668.
- Janzen, G. and Van Turenout, M. (2004). Selective neural representation of objects relevant for navigation. In *Nature Neuroscience*, 7(6), 673–677.
- Jensen, F.V. (2001). Bayesian networks and decision graphs. Springer.
- Johnson-Laird, P.N. (1983). Mental models: Towards a cognitive science of language, inference, and consciousness. Cambridge University Press.
- Johnson, M. (1987). *The Body in the Mind: the Bodily Basis of Meaning, Imagination and Reason*. University of Chicago Press.
- Jul, S. and Furnas, G.W. (1998). Critical zones in desert fog: Aids to multiscale navigation. In *UIST (User Interface Software and Technology): Proceedings of the ACM Symposium*, 97–106.

- Kamalipour, H. and Peimani, N. (2015). Assemblage Thinking and the City: Implications for Urban Studies. In *Current Urban Studies*, 3(4), 402–408.
- Kirsh, D. and Maglio, P. (1994). On distinguishing epistemic from pragmatic action. In *Cognitive Science*, 18(4), 513–549.
- Kitchin, R. and Dodge, M. (2007). Rethinking Maps. In *Progress in Human Geography*, 31(3), 331–344.
- Kitchin, R. and Dodge, M. (2011). *Code/ Space*. MIT Press.
- Kitchin, R. (1996). Are there sex differences in geographic knowledge and understanding? In *The Geographical Journal*, 162(3), 273–286.
- Klippel, A. (2003). *Wayfinding Choremes*. Thesis. University of Bremen.
- Knierim, J.J. (2002). Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. In *Journal of Neuroscience*, 22(14), 6254–6264.
- Konolige, K., Marder-Eppstein, E. and Marthi, B. (2011). Navigation in hybrid metric-topological maps. In Salas, J. (eds), *Proceedings - IEEE International Conference on Robotics and Automation*, 3041–3047.
- Kraak, M. (2003). The Space-Time Cube Revisited from a Geovisualization Perspective. In *Proceedings of the 21st International Cartographic Conference (ICC)*, 1988-1995.
- Kuhn, W. (2007). An Image-Schematic Account of Spatial Categories. In Winter, S. et al. (eds), *Proceedings of the Conference on Spatial Information Theory, COSIT 2007*, Melbourne, 152–168.
- Kuipers, B. (1978). Modeling spatial knowledge. In *Cognitive Science*, 2(2), 129–153.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press.
- Lakoff, G. (1990). The Invariance Hypothesis - is abstract reason based on image-schemas? In *Cognitive Linguistics*, 1(1), 39-74.

- Lakoff, G. and Turner, M. (1989). *More than cool reason: a field guide to poetic metaphor*. University of Chicago Press.
- De Landa, M. (2013). *Intensive Science and Virtual Philosophy*. Bloomsbury Publishing.
- Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., Moser, M. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. In *Science*, 309(5734), 619–623.
- Levine, D., Buchsbaum, D., Hirsh-Pasek, K. and Golinkoff, R.M. (2019). Finding events in a continuous world: A developmental account. In *Developmental Psychobiology*, 61(3), 376–389.
- Louf, R., Jensen, P. and Barthelemy, M. (2013). Emergence of hierarchy in cost-driven growth of spatial networks. In *Proceedings of the National Academy of Sciences of the United States of America*, 110(22), 8824–8829.
- Lüscher, P., Burghardt, D. and Weibel, R. (2007). Ontology-driven Enrichment of Spatial Databases. In *Proceedings of the 10th ICA Workshop on Generalisation and Multiple Representation*, Moscow, 1–13.
- Lynch, K. (1962). *The Image of the City*. MIT Press.
- MacEachren, A.M. (1995). *How Maps Work: Representation, Visualization, and Design*. The Guilford Press.
- Mackaness, W., Quigley, A. and Tanasescu, V. (2011). Hierarchical Structures in Support of Dynamic Presentation of Multi Resolution Geographic Information for Navigation in Urban Environments. In Jones, C.A. et al. (eds), *Proceedings of the 19th GIS Research UK Annual Conference*, University of Portsmouth, 216– 224.
- Mackaness, W. and Reimer, A. (2014). Generalisation in the Context of Schematised Maps. In Burghardt, D., Duchene, C. and Mackaness, W. (eds), *Abstracting Geographic Information in a Data Rich World: Methodologies and Applications of Map Generalisation*, 299-328. Springer International.
- Madl, T. (2016). *Bayesian Mechanisms in Spatial Cognition - Towards Real-World Capable Computational Cognitive Models of Spatial Memory*. Thesis. The University of Manchester.

- Madl, T., Franklin, S., Chen, K. and Trapp, R. (2018). A computational cognitive framework of spatial memory in brains and robots. In *Cognitive Systems Research*, 47, 147–172.
- Mani, K. and Johnson-Laird, P.N. (1982). The mental representation of spatial descriptions. In *Memory & Cognition*, 10(2), 181–187.
- Marchette, S.A., Ryan, J. and Epstein, R.A. (2017). Schematic representations of local environmental space guide goal-directed navigation. In *Cognition*, 158, 68–80.
- Mark, D.M. (1999). Spatial representation: a cognitive view. In *Geographical information systems: principles and applications*, 1, 81–89.
- Mark, D.M. and Frank, A.U. (1996). Experiential and formal models of geographic space. In *Environment and Planning B: Planning and Design*, 23(1), 3–24.
- McDonough, T. (1994). Situationist Space. In *October (Journal)*, 67, 58–77.
- Mielle, M., Magnusson, M. and Lilienthal, A. (2019). URSIM: Unique Regions for Sketch Map Interpretation and Matching. In *Robotics*, 8(2), 43.
- Milford, M. and Schulz, R. (2014). Principles of goal-directed spatial robot navigation in biomimetic models. In *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369: 20130484.
- Miller, G. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. In *Psychological review*, 101(2), 343–352.
- Miniussi, C., Wilding, E.L., Coull, J.T. and Nobre, A.C. (1999). Orienting attention in time: Modulation of brain potentials. In *Brain*, 122(8), 1507–1518.
- Mirabella, G., Bertini, G., Samengo, I., Kilavik, B.E., Frilli, D., Libera, C.D. and Chelazzi, L. (2007). Neurons in Area V4 of the Macaque Translate Attended Visual Features into Behaviorally Relevant Categories. In *Neuron*, 54(2), 303–318.

- Montello, D.R., Lovelace, K.L., Golledge, R.G. and Self, C.M. (1999). Sex-related differences and similarities in geographic and environmental spatial abilities. In *Annals of the Association of American Geographers*, 89(3), 515–534.
- Moser, E.I., Kropff, E. and Moser, M. (2008). Place Cells, Grid Cells, and the Brain's Spatial Representation System. In *Annual Review of Neuroscience*, 31(1), 69–89.
- Mou, W. and McNamara, T.P. (2002). Intrinsic Frames of Reference in Spatial Memory. In *Journal of Experimental Psychology: Learning Memory and Cognition*, 28(1), 162–170.
- Muhle-Karbe, P.S., Jiang, J. and Egner, T. (2018). Causal evidence for learning-dependent frontal lobe contributions to cognitive control. In *Journal of Neuroscience*, 38(4), 962–973.
- Muller, R.U., Stead, M. and Pach, J. (1996). The hippocampus as a cognitive graph. In *Journal of General Physiology*, 107(6), 663–694.
- Nadel, L. (2008). The Hippocampus and Context Revisited. In Mizumori, S. (eds), *Hippocampal Place Fields: Relevance to Learning and Memory*. 3-15. Oxford University Press.
- Navon, D. and Gopher, D. (1979). On the economy of the human-processing system. In *Psychological Review*, 86(3), 214–255.
- Newell, A. and Simon, H. (1972). Human problem solving. Prentice Hall.
- Norman, D. (1991). Cognitive Artifacts. In Carroll, J.M. (eds), *Designing Interaction*. Cambridge University Press.
- O'Keefe, J. (1998). The Spatial Prepositions in English, Vector Grammar and the Cognitive Map Theory. In Bloom et al. (eds), *Language and Space*, 277-316. MIT Press.
- O'Neill, M. (1991). A biologically based model of spatial cognition and wayfinding. In *Journal of Environmental Psychology*, 11(4), 299–320.
- Olton, D.S., Becker, J.T. and Handeimann, G.E. (1979). Hippocampus, space, and memory. In *The Behavioral and Brain Sciences*, 2(1979), 313–365.

- van Oosterom, P. (2005). Variable-scale topological data structures suitable for progressive data transfer: The GAP-face tree and GAP-edge forest. In *Cartography and Geographic Information Science*, 32(4), 331–346.
- van Oosterom, P. and Meijers, M. (2014). Vario-scale data structures supporting smooth zoom and progressive transfer of 2D and 3D data. In *International Journal of Geographical Information Science*, 28(3), 455–478.
- Oren, T. and Yilmaz, L. (2013). Philosophical Aspects of Modeling and Simulation. In Tolk, A. (eds), *Ontology, Epistemology, and Teleology for Modeling and Simulation*, 157-172. Springer.
- Ort, E., Fahrenfort, J.J., Reeder, R., Pollmann, S. and Olivers, C. (2019). Frontal cortex differentiates between free and imposed target selection in multiple-target search. In *NeuroImage*, 202(2019):116133.
- Oulasvirta, A., Tamminen, S. and Roto, V. (2005). Interaction in 4-second bursts: The fragmented nature of attentional resources in mobile HCI. In Kellogg, W. (eds), *CHI 2005 - Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 919–928.
- Peirce, C. (1969). *Collected Papers of Charles Sanders Peirce, Volumes I and II – Principles of Philosophy and Elements of Logic*. Harvard University Press.
- Peirce, C. S. (1878). How to make our ideas clear. In *Popular Science Monthly*, 12(January), 286–302.
- Peters, J. and Büchel, C. (2010). Neural representations of subjective reward value. In *Behavioural Brain Research*, 213(2), 135–141.
- Pettijohn, K.A. and Radvansky, G.A. (2016). Walking through doorways causes forgetting: Environmental effects. In *Journal of Cognitive Psychology*, 28(3), 329–340.
- Pezzulo, G., Rigoli, F. and Friston, K.J. (2018). Hierarchical Active Inference: A Theory of Motivated Control. In *Trends in Cognitive Sciences*, 22(4), 294–306.
- Piaget, J. and Inhelder, B. (1956). *The child’s conception of space*. Routledge & Kegan Paul.

- Posner, M. and Snyder, C. (2004). Attention and Cognitive Control. In Balota, D.A. and Marsh, E.J. (eds), *Cognitive psychology: Key Readings*, 205–223. Psychology Press.
- Poucet, B. (1993). Spatial Cognitive Maps in Animals: New Hypotheses on Their Structure and Neural Mechanisms. In *Psychological Review*, 100(2), 163–182.
- Prandi, F., Soave, M., Devigili, F., De Amicis, R. and Astyakopoulos, A. (2014). Collaboratively collected geodata to support routing service for disabled people. In Gartner, G. and Huang, H. (eds), *Proceedings of the 11th International Symposium on Location-Based Services*, (November), 67–85.
- Pulvermüller, F. (1999). Words in the brain’s language. In *Behavioral and Brain Sciences*, 22(2), 253–279.
- Quigley, A.J. (2001). Large Scale Relational Information Visualization, Clustering, and Abstraction. Thesis. Newcastle University.
- Radvansky, G.A. (2012). Across the Event Horizon. In *Current Directions in Psychological Science*, 21(4), 269–272.
- Radvansky, G.A. and Zacks, J.M. (2014). Event Cognition. Oxford University Press.
- Radvansky, G.A. and Zacks, J.M. (2017). Event boundaries in memory and cognition. In *Current Opinion in Behavioral Sciences*, 17, 133–140.
- Randell, D., Cui, Z. and Cohn, A. (1992). A Spatial Logic based on Regions and Connection. In Nebel, B. (eds), *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, 165–176.
- Raubal, M. and Egenhofer, M. (1998). Comparing the complexity of wayfinding tasks in built environments. In *Environment and Planning B*, 25, 895–913.
- Raubal, M., Egenhofer, M.J. and Pfoser, D. (1997). Structuring Space with Image Schemata: Wayfinding in Airports as a Case Study. In Hirtle, S. C., and Frank, A. U. (eds), *Proceedings of the Conference of Spatial Information Theory, COSIT 97*, Laurel Highlands, 85–102.

- Raubal, M. and Worboys, M. (1999). A Formal Model of the Process of Wayfinding in Built Environments. In Freksa, C., and Mark, D. M. (eds), *Proceedings of the Conference of Spatial Information Theory, COSIT 99*, Stade, 381-399.
- Raz, A. and Buhle, J. (2006). Typologies of attentional networks. In *Nature Reviews Neuroscience*, 7(5), 367–379.
- Remolina, E., Fernandez, J.A., Kuipers, B. and Gonzalez, J. (1999). Formalizing regions in the spatial semantic hierarchy: An ah-graphs Implementation approach. In Freksa, C., and Mark, D. M. (eds), *Proceedings of the Conference of Spatial Information Theory, COSIT 99*, Stade, 109–124.
- Roberts, M.J., Newton, E.J., Lagattolla, F.D., Hughes, S. and Haslera, M.C. (2013). Objective versus subjective measures of Paris Metro map usability: Investigating traditional octolinear versus all-curves schematics. In *International Journal of Human Computer Studies*, 71(3), 363–386.
- Rodriguez, M.A. (2015). The gremlin graph traversal machine and language. In Cheney, J. (eds), *DBPL 2015 - Proceedings of the 15th Symposium on Database Programming Languages*, Pittsburgh, 1–10.
- Rowlands, M. (2005) Situated Representation. Thesis. Maastricht University.
- Rumelhart, D.E. and Ortony, A. (1977). The Representation of Knowledge in Memory. In Anderson, R.C. et al. (eds), *Schooling and the Acquisition of Knowledge*, 99–135. Routledge.
- Rupert, R. D. (2009). Cognitive Systems and the Extended Mind. Oxford University Press.
- Schick, W., Halfmann, M., Hardiess, G., Hamm, F. and Mallot, H.A. (2019). Language cues in the formation of hierarchical representations of space. In *Spatial Cognition and Computation*, 9(3), 252–281.
- Schmid, F., Richter, K.F. and Peters, D. (2010). Route aware maps: Multigranular wayfinding assistance. In *Spatial Cognition and Computation*, 10(2–3), 184–206.
- Schultz, W., Dayan, P. and Montague, P.R. (1997). A neural substrate of prediction and reward. In *Science*, 275(5306), 1593–1599.

- Schwering, A., Krukar, J., Galvao, M., Schick, W. and Löwen, H. (2019). Wayfinding Through Orientation: Schematizing Landmark, Route and Survey Information in a Single Map. In Gartner, G. et al. (eds), *Proceedings of the 2nd Schematic Mapping Workshop*, Vienna, 1–5.
- Siyari, P., Dilkina, B. and Dovrolis, C. (2019). Emergence and Evolution of Hierarchical Structure in Complex Systems. In Ghanbarnejad, F. et al. (eds), *Dynamics On and Of Complex Networks III*, 23–62. Springer.
- Skaggs, E., Moore, K.M. and McNaughton, L. (1996). Binding of Hippocampal CA1 Neural Activity to Multiple Reference Frames in a Landmark-Based Navigation Task. In *Cell*, 16(2), 823–835.
- Sloman, S.A. and Lagnado, D. (2015). Causality in Thought. In *Annual Review of Psychology*, 66(1), 223–247.
- Small, D.M., Gitelman, D.R., Gregory, M.D., Nobre, A.C., Parrish, T.B. and Mesulam, M. (2003). The posterior cingulate and medial prefrontal cortex mediate the anticipatory allocation of spatial attention. In *NeuroImage*, 18(3), 633–641.
- Speer, N.K., Zacks, J.M. and Reynolds, J.R. (2007). Human brain activity time-locked to narrative event boundaries: Research article. In *Psychological Science*, 18(5), 449–455.
- Spiers, H.J., Maguire, E.A. and Square, Q. (2008). The dynamic nature of cognition during wayfinding. In *Journal of Environmental Psychology*, 28(3), 232–249.
- Spivey, M. (2007). *The Continuity of Mind*. Oxford University Press.
- Spivey, M.J. and Dale, R. (2006). Continuous dynamics in real-time cognition. In *Current Directions in Psychological Science*, 15(5), 207–211.
- Stachenfeld, K.L., Botvinick, M.M. and Gershman, S.J. (2017). The hippocampus as a predictive map. In *Nature Neuroscience*, 20(11), 1643–1653.
- Starrett, M.J. and Ekstrom, A.D. (2018). Perspective: Assessing the flexible acquisition, integration, and deployment of human spatial representations and information. In *Frontiers in Human Neuroscience*, 12(July), 1–9.

- Stenning, K. and Oberlander, J. (1995). A cognitive theory of graphical and linguistic reasoning: Logic and implementation. In *Cognitive Science*, 19(1), 97–140.
- Sterelny, K. (2004). Externalism, epistemic artefacts and the extended mind. In Schantz, R. (eds), *The Externalist Challenge*, 239–254. De Gruyter.
- Stevens, A. and Coupe, P. (1978). Distortions in judged spatial relations. In *Cognitive Psychology*, 10(4), 422–437.
- Stoianov, I.P., Pennartz, C.M., Lansink, C.S. and Pezzulo, G. (2018). Model-based spatial navigation in the hippocampus-ventral striatum circuit: A computational analysis. In *PLoS Computational Biology*, 14(9).
- Strickrodt, M., Bühlhoff, H.H. and Meilinger, T. (2019). Memory for navigable space is flexible and not restricted to exclusive local or global memory units. In *Journal of Experimental Psychology: Learning Memory and Cognition*, 45(6), 993–1013.
- Stryjewsja, D., Kwoka, K., Szymanowska, P. and Janda-Debek, B. (2017). CogMap Analyst – a quantitative analysis of the structure and content characteristics of sketch drawings of cognitive maps of urbanized spaces. In *Polish Journal of Applied Psychology*, 14(4), 9–34.
- Suchman, L.A. (1985). *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge University Press.
- Summerfield, J.J., Lopsien, J., Gitelman, D.R., Mesulam, M. and Nobre, A.C. (2006). Orienting attention based on long-term memory experience. In *Neuron*, 49(6), 905–916.
- Sweller, J., Chandler, P., Tierney, P. and Cooper, M. (1990). Cognitive Load as a Factor in the Structuring of Technical Material. In *Journal of Experimental Psychology*, 119(2), 176–192.
- Taylor, H.A. and Tversky, B. (1992). Descriptions and depictions of environments. In *Memory & Cognition*, 20(5), 483–496.
- Tenbrink, T. (2012). Relevance in spatial navigation and communication. In Stachniss, C. et al. (eds). *Spatial Cognition VIII*, Kloster Seeon, 358–377.

- Timpf, S., Volta, G.S., Pollock, D.W. and Egenhofer, M.J. (1992). A conceptual model of wayfinding using multiple levels of abstraction. In Frank, A.U. et al. (eds), *Proceedings – International Conference GIS - Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, Pisa, 348–367.
- Timpf, S. (2002). Ontologies of wayfinding. In *Networks and Spatial Economics*, 2(1), 9–33.
- Tversky, B. (1993). Cognitive maps, cognitive collages, and spatial mental models. In Frank, A. U., and Campari, I. (eds), *Proceedings of the Conference of Spatial Information Theory, COSIT 93*, 14–24.
- Tversky, B. (1998). Spatial Perspective in Descriptions. In Bloom, P. et al. (eds), *Language and Space*, 463-492. MIT Press.
- Tversky, B. (2000). Some Ways that Maps and Diagrams Communicate. In Freksa, C. et al. (eds), *Spatial Cognition II*, 72–79. Springer.
- Tversky, B. (2002). What do sketches say about thinking? In Davis, R. et al. (eds), *AAAI Spring Symposium, Sketch Understanding Workshop, Stanford University, AAAI Technical Report*, Menlo Park, 148–151.
- Vetere, G., Tran, L.M., Moberg, S., Steadman, P.E., Restivo, L., Morrison, F.G., Ressler, K.J., Josselyn, S.A. and Frankland, P.W. (2019). Memory formation in the absence of experience. In *Nature Neuroscience*, 22(6), 933–940.
- Waldmann, M.R. and Hagmayer, Y. (2013). Causal Reasoning. In Reisberg, D. (eds), *Oxford handbook of cognitive psychology*, 733-752, Oxford University Press.
- Wang, J. and Schwering, A. (2015). Invariant spatial information in sketch maps - a study of survey sketch maps of urban areas. In *Journal of Spatial Information Science*, 11(2015), 31–52.
- Warren, W.H., Rothman, D.B., Schnapp, B.H. and Ericson, J.D. (2017). Wormholes in virtual space: From cognitive maps to cognitive graphs. In *Cognition*, 166, 152–163.
- Warren, W.H. (2019). Non-Euclidean navigation. In *Journal of Experimental Biology*, 222:jeb187971.

- Wen, Z., Yu, T., Yang, X. and Li, Y. (2019). Goal-directed processing of naturalistic stimuli modulates large-scale functional connectivity. In *Frontiers in Neuroscience*, 12(January), 1–12.
- Wiener, J.M., Schnee, A. and Mallot, H.A. (2004). Use and interaction of navigation strategies in regionalized environments. In *Journal of Environmental Psychology*, 24(4), 475–493.
- Wiener, J. and Mallot, H. (2003). “Fine-to-Coarse” Route Planning and Navigation in Regionalised Environments. In *Spatial Cognition and Computation*, 3(4), 331–358.
- Van Wijk, J.J. and Nuij, W. A. (2003). Smooth and efficient zooming and panning. In Munzner, T. (eds), *Proceedings - IEEE Symposium on Information Visualization*, InfoVis 03, 15–22.
- Van Wijk, J.J. Van and Nuij, W.A. (2004). A Model for Smooth Viewing and Navigation of Large 2D Information Spaces. In *IEEE Transactions on Visualization and Computer Graphics*, 10(4), 447–458.
- Williams, A.N., Postans, M. and Hodgetts, C.J. (2019). How the human brain segments continuous experience. In *Journal of Neuroscience*, 39(17), 3172–3174.
- Willis, K.S., Hölscher, C., Wilbertz, G. and Li, C. (2009). A comparison of spatial knowledge acquisition with maps and mobile maps. In *Computers, Environment and Urban Systems*, 33(2), 100–110.
- Wilson, S.P. and Wilson, P.N. (2018). Failure to demonstrate short-cutting in a replication and extension of Tolman et al.’s spatial learning experiment with humans. In *PLoS ONE*, 13(12): e0208794.
- Wu, H. (2016). Focus+Context Metro Map Layout and Annotation. In Spagnuolo, M. (eds), *SCCG '16: Proceedings of the 32nd Spring Conference on Computer Graphics*, Slomenice, 41–47.
- Yantis, S. and Egeth, H.E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. In *Journal of Experimental Psychology: Human Perception and Performance*, 25(3), 661–676.

- Yee, D., Solway, A., Diuk, C., Córdova, N., Barto, A.G., Niv, Y., Botvinick, M.M. (2014). Optimal Behavioral Hierarchy. In *PLoS Computational Biology* 10(8), 1-10.
- Yoshino, R. (1991). A note on cognitive maps: An optimal spatial knowledge representation. In *Journal of Mathematical Psychology*, 35(3), 371–393.
- Zacks, J.M. and Tversky, B. (2001). Event Structure in Perception and Conception. In *Psychological Bulletin*, 127(1), 3–21.
- Zhao, W., Li, Q. and Li, B. (2011). Spatial cognition driven context-adaptive route directions. In *Journal of Remote Sensing*, 15(6), 1171-1188.
- Zivkovic, Z., Bakker, B. and Kröse, B. (2006). Hierarchical map building and planning based on graph partitioning. In *Proceedings - IEEE International Conference on Robotics and Automation*, Orlando, 803–809.



# Appendix 1: Example Study Data


## (1.1) Figure 50

1. Exit ECC1 reception.
2. Go straight up to end. of road
3. Turn right on South Bridge, immediately cross road.
4. Continue down S. Bridge to temp. bus stop by Burroto skate.
5. Catch no. 31 bus.
6. Bus follows S Bridge to Prince's St., turn right, continue towards Haymarket.
7. Get off bus at 1st stop after Haymarket train station.
8. Immediately cross road.
9. Go straight up Coates Gardens.
10. Turn right at end of road.
11. Destination will be 2nd door on the right (green).


1. Leave café and walk uphill until traffic lights
2. Cross lights, find your bus stop. It's the second stop (there will be another one just before that, yours is 5m behind that). The stop is located opposite shops
3. Get in bus 31. ~~Get~~ Get off at Haymarket Station (The bus will pass Waverly bridge, Prince's street & continue Haymarket)
4. Get off at Haymarket.
5. Find street (just opposite) go straight up and right (cross street to the right).

1. Leave the ECCI building by the <sup>front</sup> door, cross the <sup>road</sup> ~~to~~ in to Infirmary St. Follow, slightly uphill to the traffic light controlled junction. Cross the road at the crossing to your R. and turn right; follow the footpath past shops + restaurants for c. 200m. On your left is a small open space.
2. Cross this diagonally, steps and you're at the Royal Mile. Cross this into the street opposite (Cockburn I think). It slopes steeply down & curves to the left. As it curves look to your R. for a flight of steps, (Fleshmarket Close?). Go to the foot of these steps & directly opposite you will see an entrance to Waverley Station. Enter the station & follow the overbridge ~~looking~~ looking for signs to the Travel Centre & tickets. Take the steps down, turning R. thro' 180° you will see ticket machines & opposite the Travel Centre.
3. Ticket in hand, look for the next train to Haymarket (there are many). Board that train, & board the stop. Get off the train & find your way to the ~~exit~~ exit on your left. Leave the station, cross the train tracks & look for two bus stops for traffic heading right to left. At the second stop take the 26 bus.
4. Travel about 15/20 minutes, passing the Zoo on your right & you will come to Corstorphine. Your stop is ~~at~~ called St. Ninian's church (with luck it will be announced). In any case look for Domino's pizza on your ~~left~~. This is a little before your stop. Dismount & continue walking in the same direction ~~as~~ the bus (westwards). After about 100m the first turning to the L. is Featherhall (Avenue?) Take it, and walk ~~2~~ further 250m to the second R, Featherhall ~~at~~ South No 10 is on your right about 100m down.

(1.2) Figure 52 (rotated)

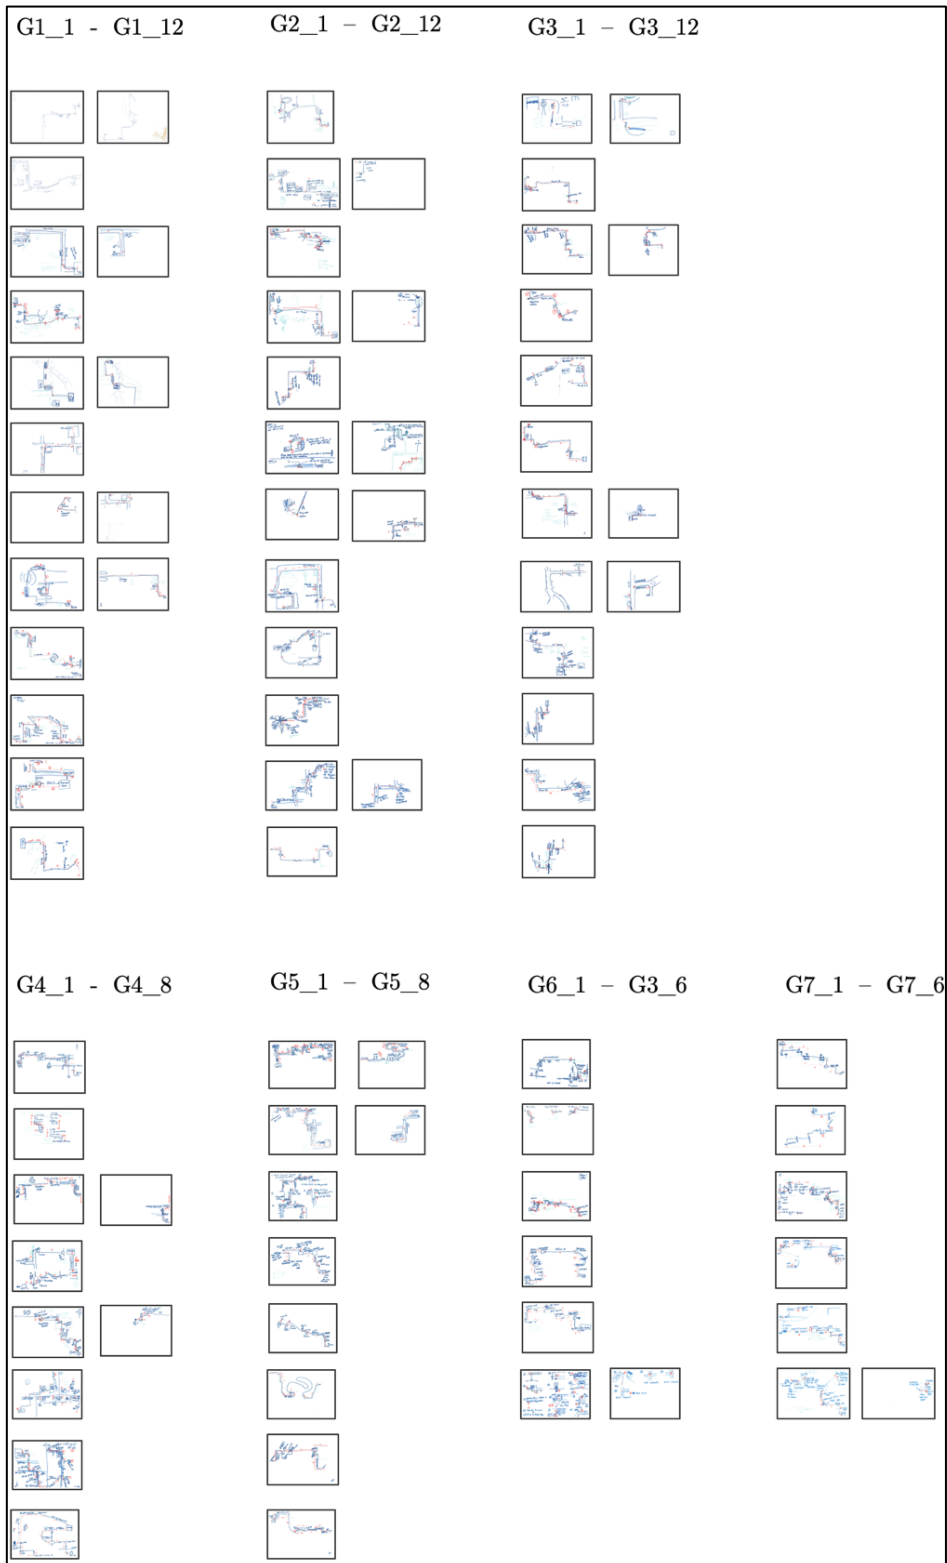







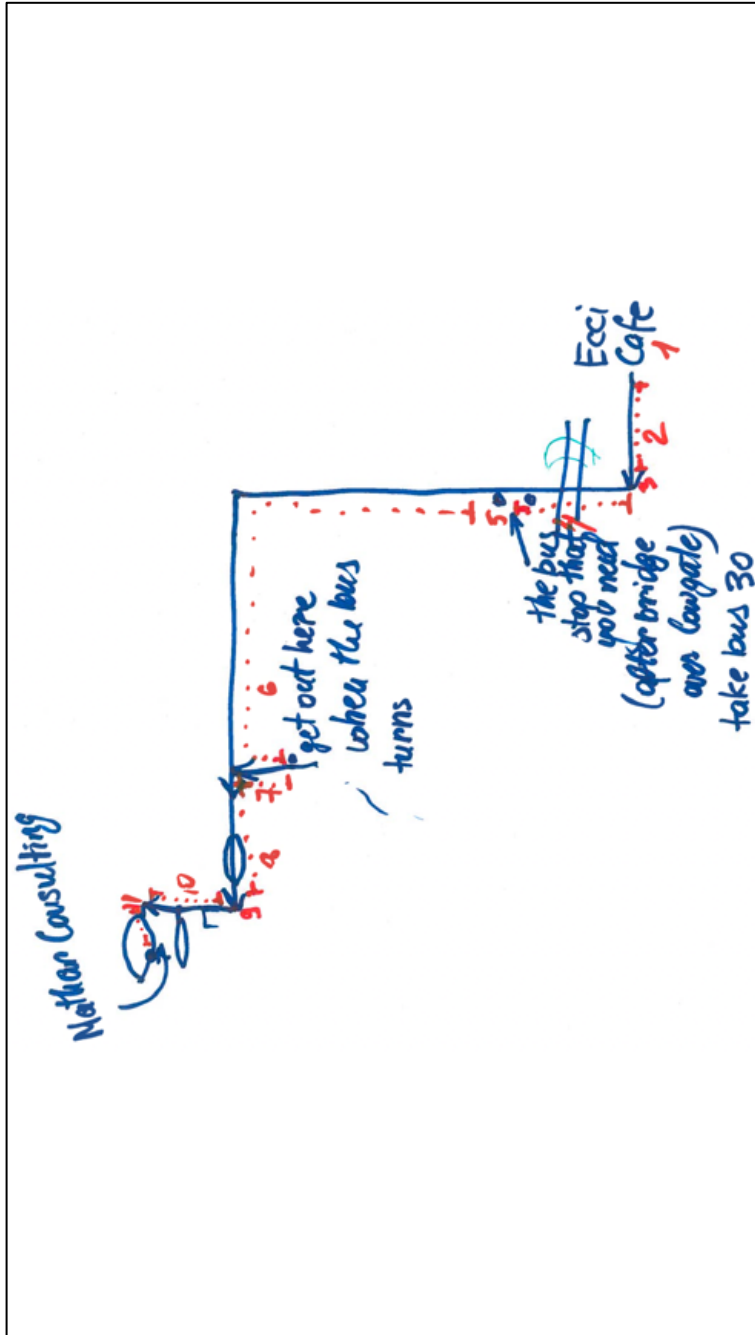



(1.3) Figure 52




**(1.4) Figure 53**












(1.6) Figure 55 (rotated)



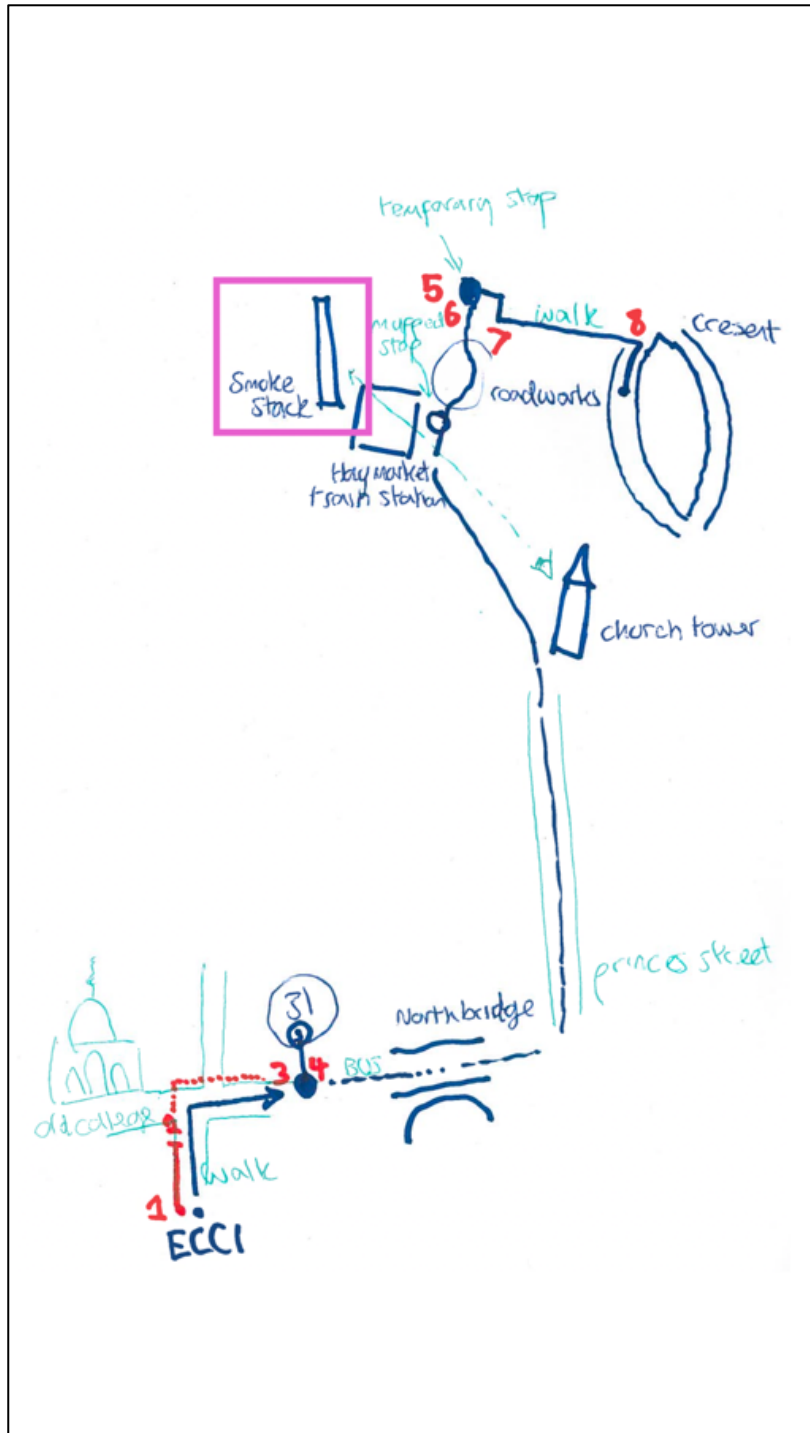


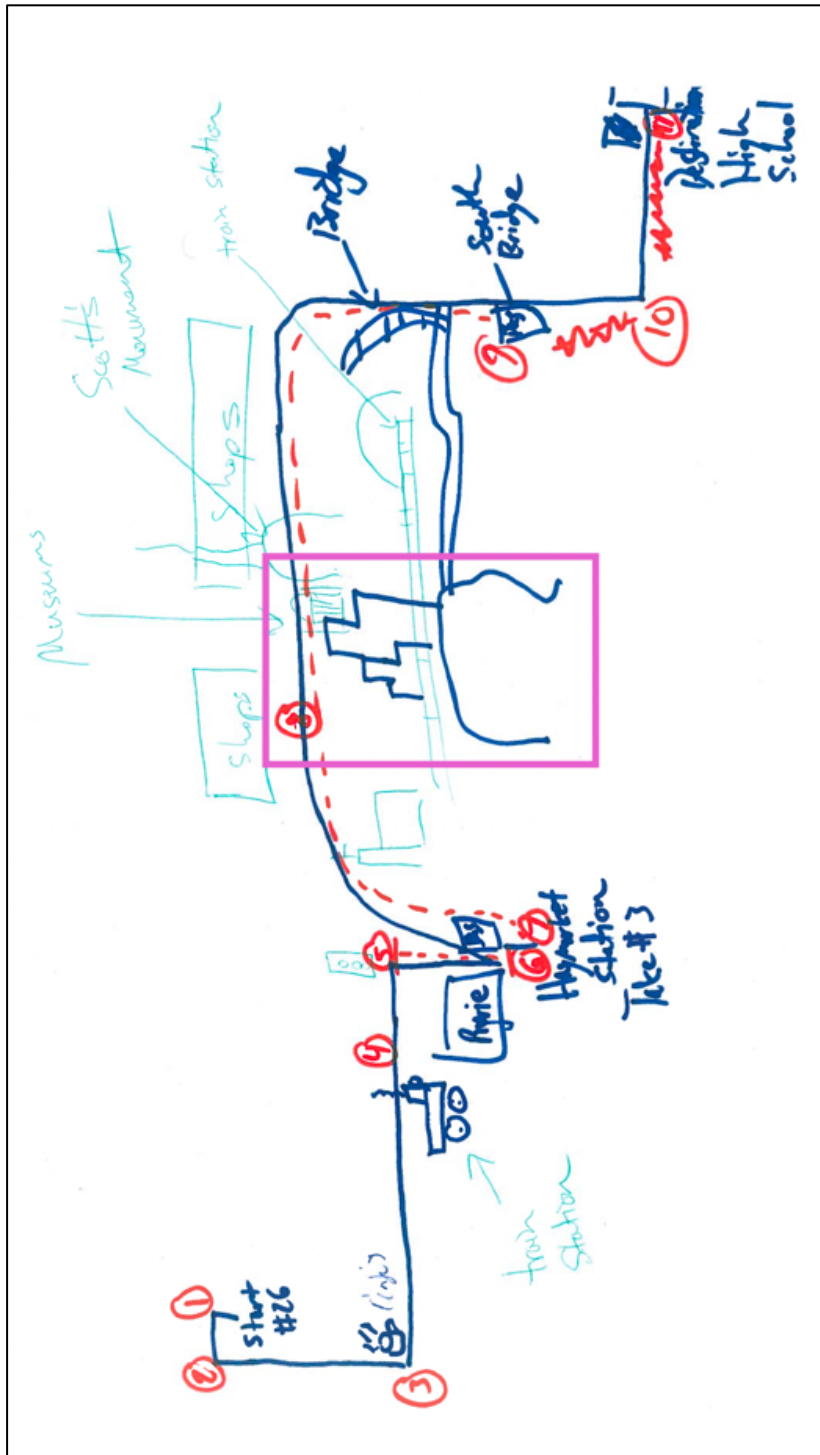




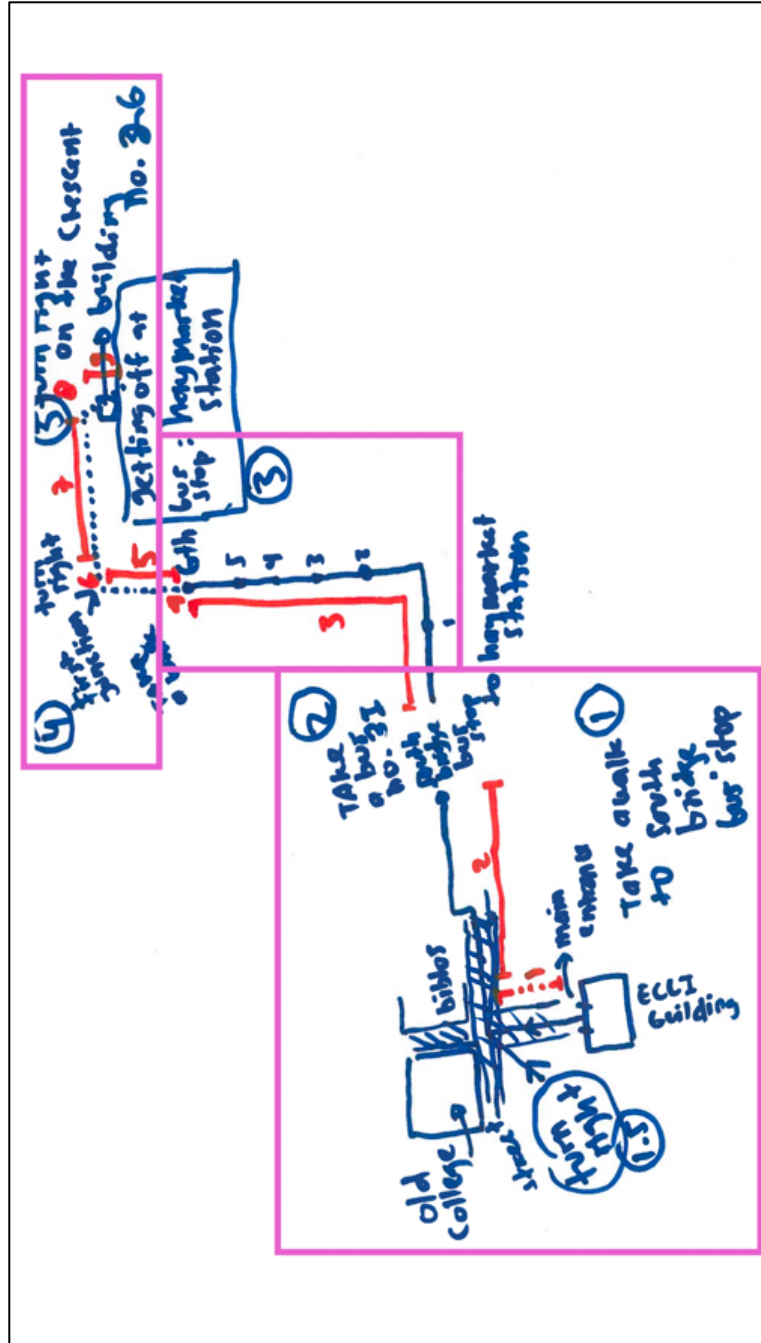



(1.8) Figure 58 (rotated)

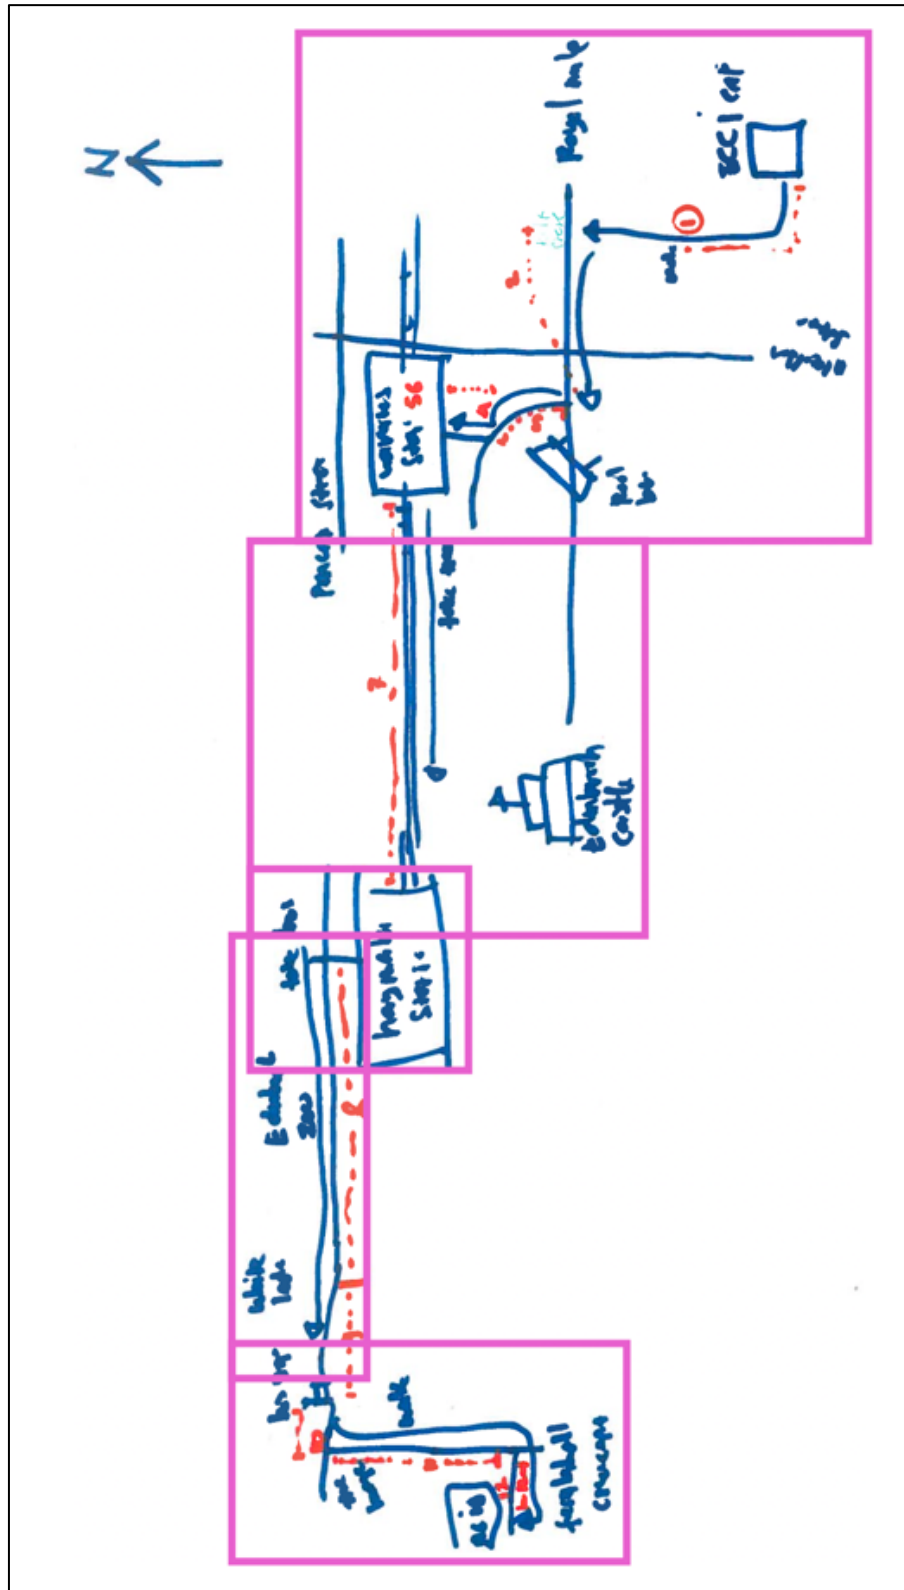


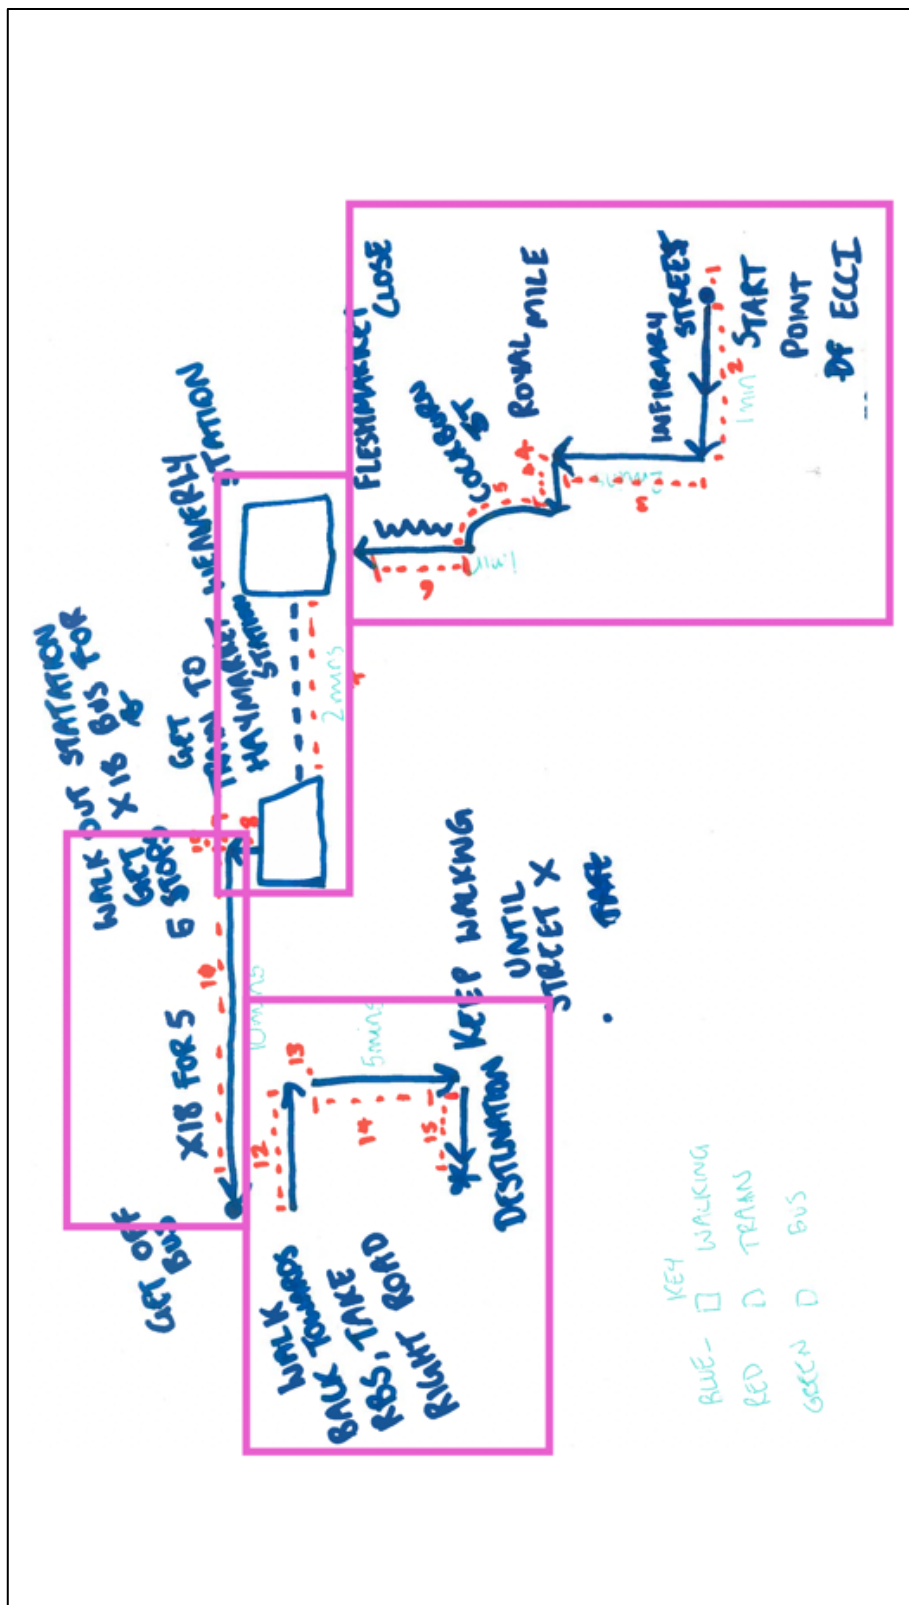


(1.9) Figure 59 (rotated)











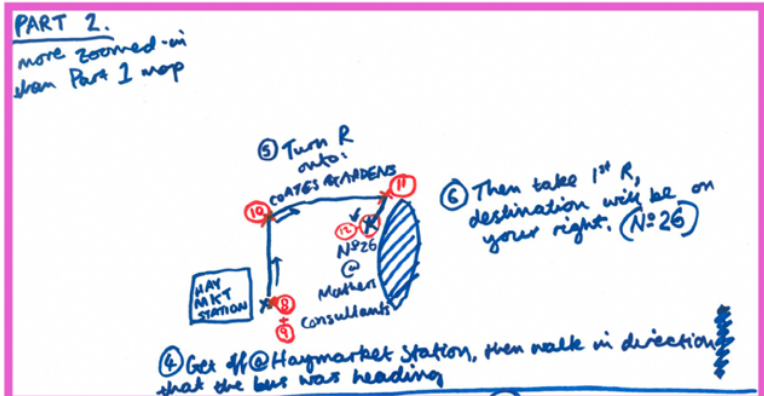


(1.10) Figure 60 (rotated)

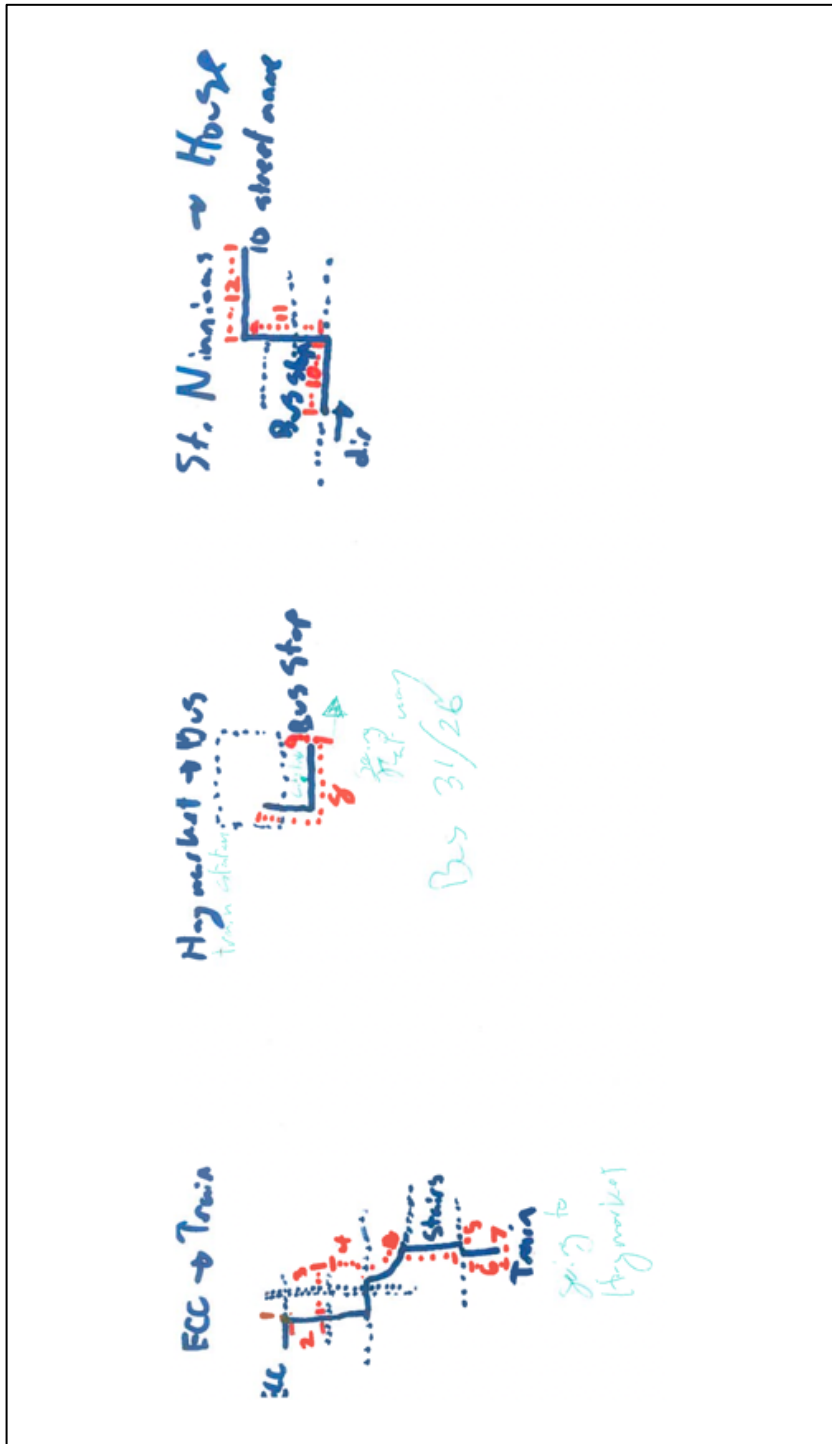



(1.11) Figure 61 (rotated)





(1.12) Figure 62 (rotated)







Get off when bus arrive  
Hymanbet 4


the end of st.

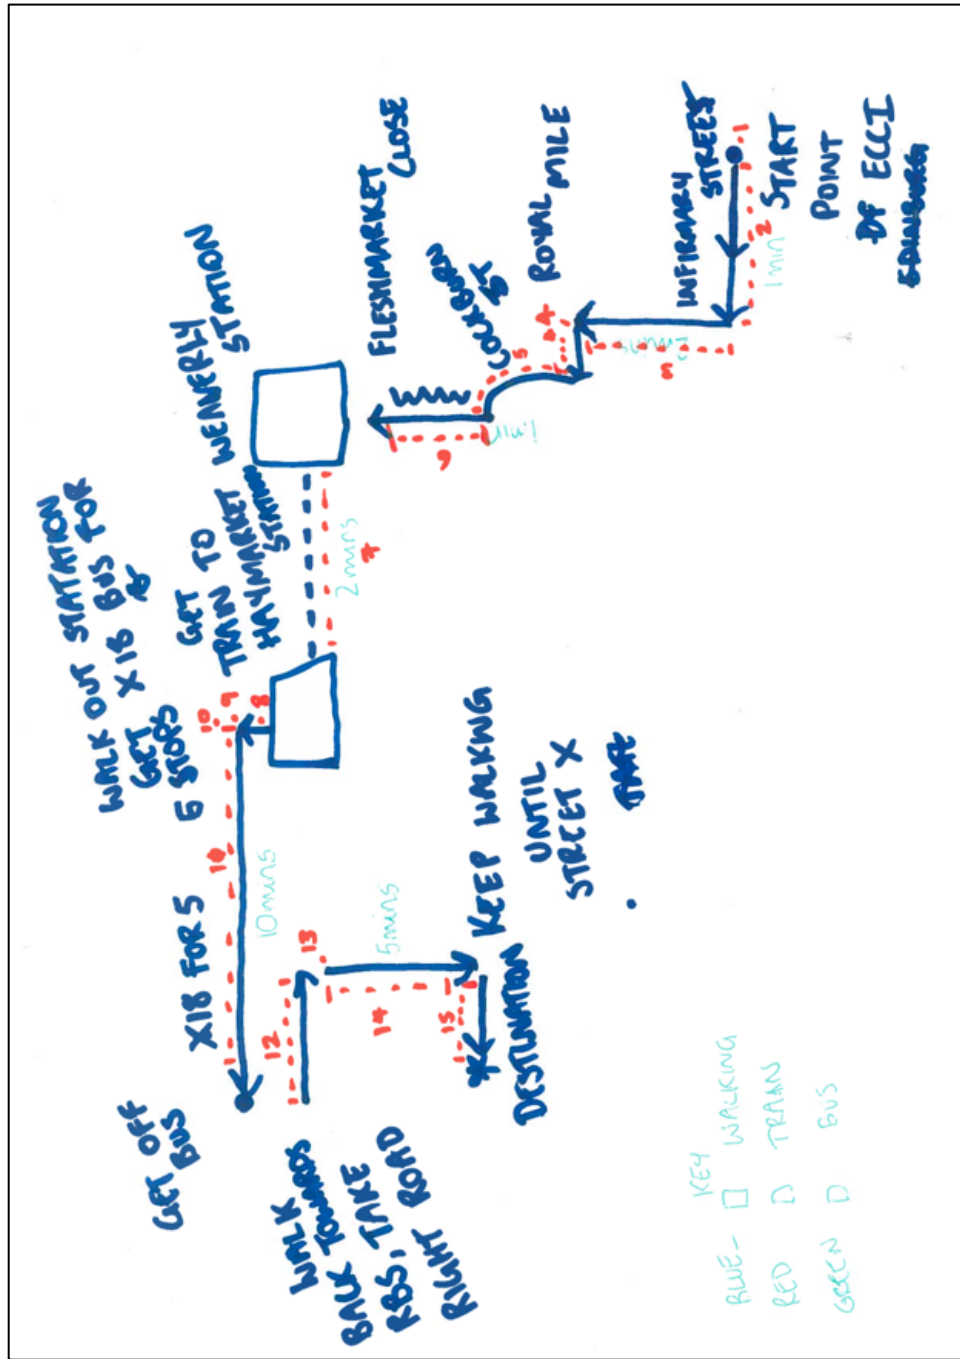
(1.13) Figure 67 (rotated)



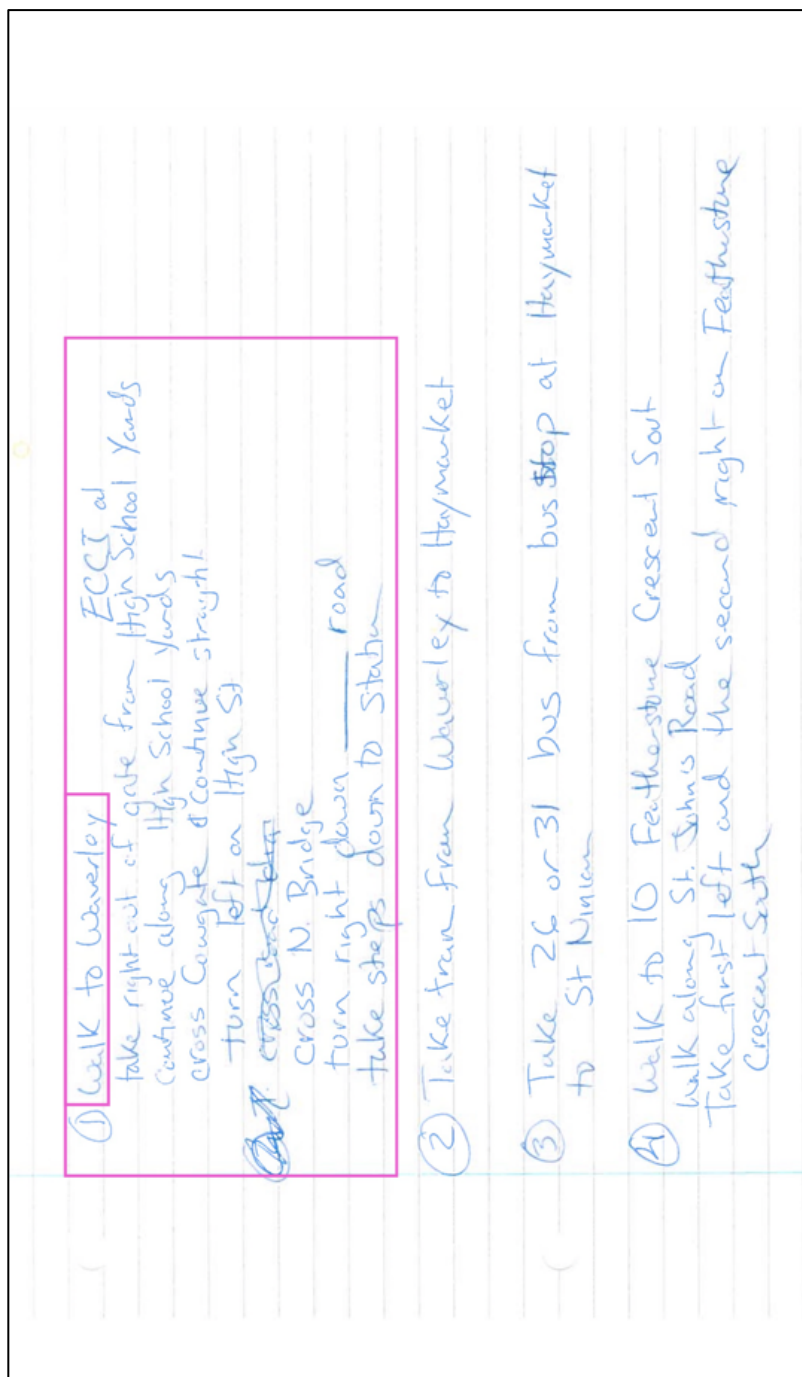


(1.14) Figure 68 (rotated)



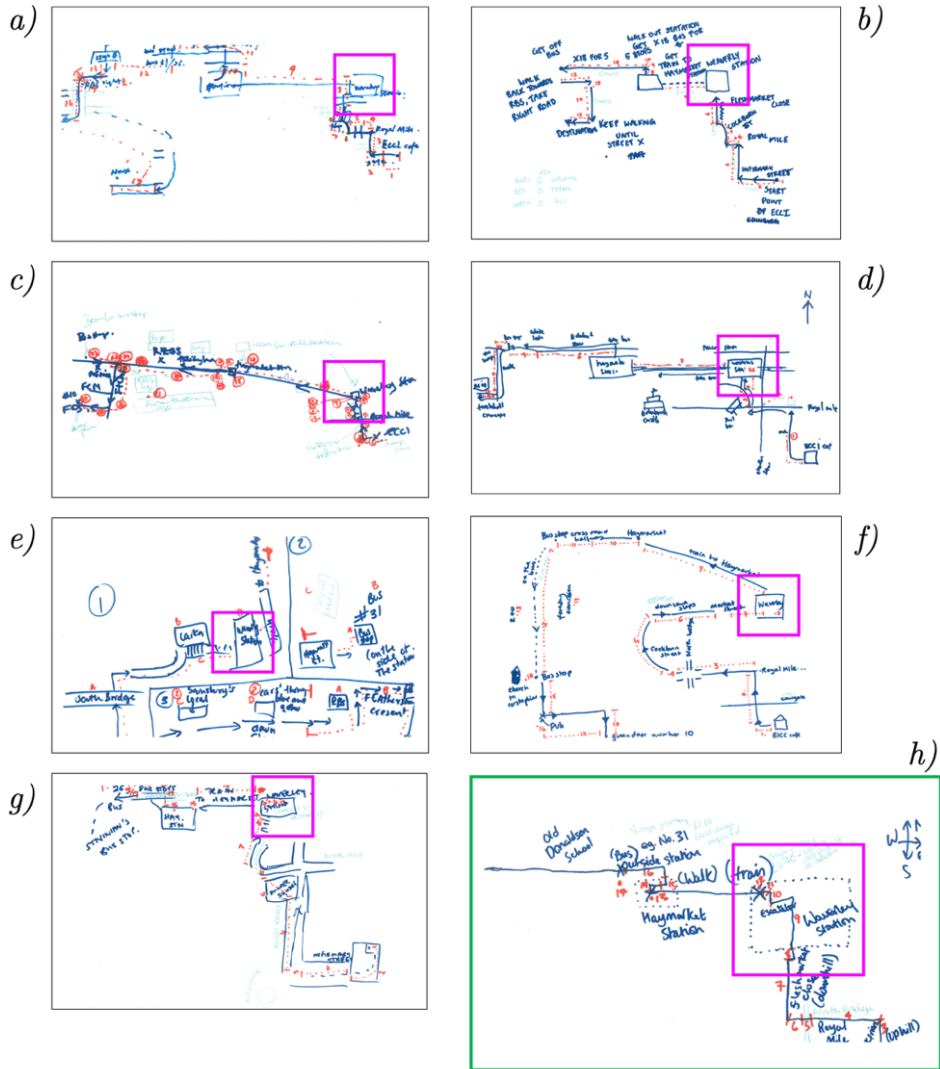


(1.15) Figure 68 (rotated)



walk	<p>① turn right after goes out of the Building.          Pass the Zara. Take away and keep walking.          turn left until you <del>see</del> reached Poyol mile.          keep walking and cross the road (only use time)          turn right and walk down the road.          pass the Boots store and walk down the stairs.          Once get down the stairs you will see the Train station.</p>
train	<p>② take the train to Haymarket station. (only one stop)</p>
walk	<p>③ gets out of the train station from the way out which face the street          wait for 26 or 31</p>
bus	<p>④ after goes on the bus you will pass the N... S... and the zoo.          gets off at st N... stop</p>
walk	<p>⑤ walk towards the direction that the bus goes          turn left when you see the Royal Bank of Scotland.          keep walking for 4-5 minutes and turn right at ... road          the destination is on your right, follow the numbers on the house.</p>

(1.16) Figure 74







## Appendix 2: Navigation Errors

<b>Group</b>	<b>Participant</b>	<b>Minor errors</b>	<b>Major errors</b>	<b>Comments</b>
G2	G2_1	0	0	No errors
	G2_2	0	1	Couldn't find the bus stop using Google Maps
	G2_3	0	0	No errors
	G2_4	1	0	Did not initially understand which side of the road the bus stop was on - had to cross the road and double back
	G2_5	1	0	Got off at a bus stop prior to the correct stop given the journey destination
	G2_6	0	0	No errors
	G2_7	0	0	No errors
	G2_8	1	0	Initially went to incorrect bus stop but corrected error
	G2_9	1	0	Initially went to incorrect bus stop but corrected error - it should be noted however that correct bus stop was closed and the journey required the use of a temporary stop which was note reflected in the information provided
	G2_10	1	0	Got off at a bus stop prior to the correct stop given the journey destination
	G2_11	0	0	No errors
	G2_12	1	0	Incorrectly crossed street and had to cross back
G3	G3_1	0	0	No errors and did not check any other information
	G3_2	0	0	No error but used Google Maps to confirm position on street network and Lothian app to look up bus times and route
	G3_3	0	0	Used Lothian to clarify bus route and times and Google Maps to confirm street names

	G3_4	0	0	Used Google Maps to confirm street network after getting off bus at different stop to the one shown in the prototype (intended stop was closed) : "I see the easyCoffee and the Nomad Cafe in the corner bit I don't see that... I see the salon and the Coates Café... "
	G3_5	1	0	Did not initially understand which side of the road the bus stop was on - had to cross the road and double back
	G3_6	0	0	No errors and did not check any other information
	G3_7	0	0	No errors and did not check any other information
	G3_8	1		Incorrect side of street
	G3_9	0	0	No errors and did not check any other information
	G3_10	0	0	No error but used the Lothian and Google apps to work out where to get off the bus
	G3_11	0	0	No errors and did not check any other information
	G3_12	0	1	Walked past the turning to the destination and ended up 'off the map'... had to switch to Google Maps
G5	G5_1	1	0	Got off bus after correct stop and had to double back
	G5_2	0	0	No errors
	G5_3	0	0	No errors
	G5_4	1	0	Got off bus after correct stop and had to double back
	G5_5	2	0	Incorrect exit at train station; got off bus after correct stop and had to double back
	G5_6	0	0	No errors
	G5_7	0	0	No errors
	G5_8	0	1	Incorrectly waited at tram stop on side of road as opposed to bus stop on the central 'island'
G6	G6_1	0	0	No errors and did not check any other information
	G6_2	1	0	Incorrect path inside train station

	G6_3	1	0	Got off bus after correct stop and had to double back. Used Google Maps to find final destination
	G6_4	0	1	Took incorrect route during first pedestrian phase - had to be prompted
	G6_5	2	1	Took incorrect route but transfer still reached; took <i>incorrect</i> turning in final <i>pedestrian</i> phase; waited at tram stop by side of road instead of bus stop on central island
	G6_6	1	0	Took incorrect street and doubled back
G7	G7_1	0	1	Became lost without location marker and had to switch to Google Maps
	G7_2	0	0	No errors and did not check any other information
	G7_3	0	0	No errors and did not check any other information
	G7_4	0	0	No errors and did not check any other information
	G7_5	0	0	Switched to Google Maps to use location services and find bus stop (final transfer)
	G7_6	0	0	No errors and did not check any other information





# Appendix 3: S3 Research Stimulus Views







Appendix 4: International  
Journal of Cartography Pre-  
Print

The following is a pre-print version of a paper publishing in the International Journal of Cartography, Volume 3 (2017), Issue 1.

### **Abstract**

Even maps that strive for a precise representation of spatial relationships use techniques of distortion to embed a geographic extent within a two-dimensional plane – be it a page or a screen. The graphical design of geospatial information does however conform to a consensus around the conceptual limits of that distortion, with an overall design framework that constrains the design to ensure comprehension and the effective recognition of geographic entities and relations. Although constraints are necessary, we argue that the advent of digital technology, particularly in mobile mapping, warrants re-examination of the parameters of these distortions. Here we introduce the concept of ‘the bounds of distortion’ as a device for considering the conceptual boundaries of map design, and as a foundation for further work to investigate how these bounds may be redefined to better support map users with more effective graphical information. The focus here is navigational information, and in particular urban navigation and interaction with the graphical representation of urban geography and public transport networks.

# The bounds of distortion: truth, meaning and efficacy in digital geographic representation

Lucas Godfrey<sup>1</sup> and William Mackaness<sup>1</sup>

1. The School of GeoSciences, University of Edinburgh,  
Edinburgh, United Kingdom

**Abstract:** Even maps that strive for a precise representation of spatial relationships use techniques of distortion to embed a geographic extent within a two-dimensional plane – be it a page or a screen. The graphical design of geospatial information does however conform to a consensus around the conceptual limits of that distortion, with an overall design framework that constrains the design to ensure comprehension and the effective recognition of geographic entities and relations. Although constraints are necessary, we argue that the advent of digital technology, particularly in mobile mapping, warrants re-examination of the parameters of these distortions. Here we introduce the concept of ‘the bounds of distortion’ as a device for considering the conceptual boundaries of map design, and as a foundation for further work to investigate how these bounds may be redefined to better support map users with more effective graphical information. The focus here is navigational information, and in particular urban navigation and interaction with the graphical representation of urban geography and public transport networks.

**Keywords:** Spatial cognition, Wayfinding, Navigation, Mobile cartography, Map effectiveness, Cartosemiotics, Multiscale mapping, Focus maps

## 1. Distortion and the graphical communication of geographic space

While in many domains of knowledge and practice, distortion is seen as something counter-productive, as something to be avoided, distortion is in fact central to the effective visual communication of information. As Barbara Tversky observed: “We change the truth to tell a bigger truth...” (Tversky, 2015). Faced with either a two or three-dimensional space, with limited area or volume, and a fundamental constraint of human-scale, we must make decisions as to how, through a process of abstraction and simplification, we can effectively convey the essential qualities of underlying geographic phenomena. Diagrams allow us to convey topological, temporal and spatial relationships with an immediacy that is difficult and often impossible with words alone. More specifically, cartography has formalised the use of two-dimensional space for the purpose of communicating caricature and geographical relationships, with Robinson reflecting on the ambition to reduce the spatial characteristics of a large area so as to “bring things into view...” (Robinson, 1995). Despite its negative connotations, distortion is a necessary part of the process of bringing ‘things into view’, reflected in the choice of projection and map generalisation methodologies in order to support communication of spatial relationships in ‘large-scale’ or ‘transperceptual’ space (Downs et al., 1977; Freundschuh et al., 1997).

The conceptual limits to these boundaries of distortion vary with map type and anticipated audience (for example varying according to whether the map is schematic, thematic, topological or topographic). In this paper we introduce the concept of the bounds of distortion, and consider five key approaches that have been explored in terms of seeking to redefine these bounds, before proposing opportunities for future work. The central aim is to consider how the graphical display

of geospatial relationships may benefit from a reconsideration of the assumptions that underlie cartographic design. This work is particularly concerned with the efficacy of navigational information in the context of the small screen devices that have become the dominant platform of interaction for most map users.

It should be highlighted that our concern is not maps that 'look' distorted, but rather, we are concerned with cartographic practice that makes use of techniques which elevate the needs and capacities of the map user above the goal of displaying absolute spatial relations. Here we argue for an imperative to develop cartographic approaches which draw less sharp a distinction between metric and non-metric spaces, and linked to this, draws less of a distinction between the structural and functional representations of geography. This line of investigation is closely linked with the transition from ink to data and the practical implications of using 'smart' mobile devices as our primary means of accessing spatial information.

## **2. Foundations of digital geographic representation**

We feel it is necessary to begin with a brief consideration of some key technical and conceptual foundations which underpin cartographic practice and are important in the framing of our arguments.

### ***2.1 Space, maps and mappings***

The most primitive spatial structure is the topological relation. Topology can describe structures in  $n$  dimensional space as the fundamental properties of topological relations can be seen as being independent of an 'embedding space'. Euclidean space provides an embedding space for either two or three dimensional structures, with positional information being lost in the process of transforming from a Euclidean space to one that is purely topological. This issue of data loss is common to all representational processes in which the phenomenon is represented in an alternate, often lower dimensional space, moving from the large-scale to the small-scale. Throughout the cartographic process, scale and level of detail act to govern the structure of the conceptual space in which geographic entities and their relations are displayed. While there have been many approaches proposed to representing relative spatial relations, the prevailing map design paradigm is significantly biased toward a formalisation of the physical world that describes an absolute space – a container within which the entities lie.

### ***2.2 Truth and meaning in cartography***

Distortion is both real and imagined. It is 'imagined' in the sense that people do not 'see' maps completely objectively but are influenced by a number of cognitive quirks that lead to a distorted understanding even if the distortion is not present in the map itself. A simple example is our tendency to see vertical lines as being shorter than horizontal lines of the same length (Held, 1971). Another, more far reaching example is the issue that perception does not happen in isolation but is a process that occurs in parallel with [mental] projection, which has a 'looser coupling' to stimuli and allows us to 'see' what could be there, rather than what is actually presented to us (Kirsh, 2009).

Beyond these issues of veracity however, is the communication and perception of meaning, something that is often of far greater significance to spatial decision making. From the perspective of semiotic analysis, the meaning of a map is a function of the 'triadic structure', in which the underlying geography is the 'object', the map is the 'sign', and the map user, the 'interpreter'. Within this construct the map is usually considered to be an 'icon', or a sign that is considered to display the quality of 'likeness' to the object (Schlichtmann, 2009). Linking with this, the meaning conveyed by our graphical depiction of geography has also been investigated

extensively from the perspective of Gestalt theory, with Gestalt principles being called upon to ensure our language of points, lines and polygons effectively conveys a coherent sense of form, for example in defining figure-ground relationships through perceptual segregation (Goldstein, 2010), and supporting the communication of cohesive forms through visual continuity (Carvalho & Moura, 2009).

The key overall issue here is *recognition*. Whatever form a map takes, the interpreter must be able to recognise which underlying object a graphical feature represents. If two buildings are generalised into a single form at lower levels of detail, this is acceptable if the map user recognises that they are now looking at a generalised urban area rather than a single large building. These transition points at which we find ourselves looking at a fundamentally different phenomenon remain a challenge for automated map generalisation. Thus we might argue that this is a ‘semantic distortion’. Similarly, there are occasions where a distortion of the angle or length of a connecting edge is necessary in order to convey the essential character of the geography being visualised. For example, we seek to preserve the gently flexing path of the railway line because this is a defining property that differentiates it from other phenomena. We suffer other distortions in order to preserve the conveyance of this distinguishing property, and so, we see trade-offs among different distortions in order to convey various characteristic properties and ultimately, to communicate meaning to the user.

### 3. A consensus on the conceptual limits of distortion in map design

While approaches to generalisation specify the exact process and logic by which spatial data is displayed, our concern here is the more fundamental basis of our approach to distortion, and in particular, the logic that defines an ‘upper bound’ to the acceptable level. The bounds of distortion then, refers to the region between a minimum and maximum level of distortion, where distortion is defined as a progressive transformation, such that absolute position is de-emphasised to support communication of some aspect of the geography.

The minimum level of distortion is the point at which it is physically possible to visualise the various data relating to the geography in either two or three dimensions. In three dimensions, this can be viewed as a distortion due to the fact the data will not completely describe the underlying phenomena (e.g. due to sample frequency). This issue is of course also true in two dimensions, but with the added distortion required to project the three-dimensional nature of the phenomena onto a two-dimensional plane. Distortion is then increased by various means, for example through map generalisation techniques we select a sub-set of the overall set of objects based on scale-dependent rules, and we alter the relative accuracy to emphasise functional aspect of the geography, for example exaggerating road widths so the structure of the road network remains visible at smaller scales.

Our concern is the upper extreme of this region - the point beyond which, in any given design context, further distortion of absolute accuracy leads to a reduction in the effectiveness of the geographical information in terms of how well it supports decision-making. In other words, the point after which further distortion leads to a decrease in comprehension on the part of the map user. We propose that, at present, this upper bound was established prior to the advent of digital mapping technologies, and that given the capabilities of contemporary technologies, for example the potential for mobile-based solutions to automatically take account of contextual factors in the information that is ultimately presented to the user, there is a case for seeking to redefine this upper bound.

In our analysis of this issue, we will distinguish between the representation of metric and non-metric spaces, and the associated approach to distortion in each. In a geographical context, metric space describes a space in which the Euclidean distance between all geographic objects is defined

by a distance function on the overall set. This provides a consistent mapping from reality to the conceptual model, and on to the target structure. A topographic map would be one such example. Non-metric space in a geographical context includes maps that relax this approach, usually with the purpose of highlighting the *functional*, as opposed to *structural* character of the geography. In other words, maps that do not adhere to metric spatial relations make this departure with a view to increasing the saliency of key elements of the overall spatial structure that help to 'specify action' (Klippel and Richter, 2004). When considering the bounds of distortion, it is necessary to distinguish between these two high level cases, and also to consider how the two are integrated.

### ***3.1 The bounds of distortion in metric geographic representation***

The core design convention in this case is that a metric space defines the absolute location of all elements of the overall set. A key property of this approach is that changes of scale do not affect relative distances. Limits to visual acuity require cartographic distortions via the process of generalisation (e.g. exaggerated road widths and simplification of geometric form). Ultimately scale (or level of detail) governs the limit to these graphical treatments; map real estate constraints mean that as we zoom out, each phenomenon will reach a 'conceptual cusp' (Muller 1991, Mackaness et al., 2005) or 'generalisation point' (Ratajski, 1967) – a point at which it is no longer feasible or meaningful to visualise that phenomenon. This links with the notion of 'map capacity' (Ratajski, 1967), whereby map content is optimised for the given area relative to scale and thematic focus. This reflects the hierarchical view of geography that is formalised in the data models which support geographic representation within metric spaces.

### ***3.2 The bounds of distortion in non-metric geographic representation***

In non-metric space, the design convention is focused around schematisation, often with topological relations maintained. Within this conceptual framework, cardinal bearing is maintained where possible, but is not strict. From a semiotic perspective, there is a consistent graphical treatment of iconic elements with emphasis on network structure and the broader overarching caricature of the phenomena being represented. In transportation themed maps, this approach supports a functional view of the geography, focused on supporting the user in a route orientation strategy (Skagerlund et al., 2012). The challenge becomes one of supporting both 'turn-by-turn' approaches to conceptualising space, together with a survey strategy, in which an overview of the geography is comprehended.

Integrating the two approaches – metric and non-metric, is often achieved by means of nesting one within the other, with a hard boundary between the two. In the case of integrating the two (i.e. on the same 'page'), the upper bound of distortion is essentially defined as the point at which the area contained within the 'box' corresponds with the area 'underneath' the box, to the extent that coherence is maintained between the content 'inside' and 'outside'. A variety of graphical techniques are used to both warn the user of these two different spaces and of the connections between them, for example using continuity principles to show how different graphical elements in each representation are recognised as being the same road.

## ***4. Five design approaches that explore the bounds of distortion***

As part of our investigation into how distortion in cartographic practice may be reframed to focus more on the contextual needs of the user, we reflect on five approaches.

### ***4.1 Hard Boundary***

The simplest approach is that of a 'Hard Boundary' - an approach that supports the display of

two or more spatial constructs within a single ‘page’. Usually this approach is used to display a metric and non-metric space in such a way as a user may be able to quickly switch between the different views and orientate themselves in either, relative to the other. The classic example of this Hard Boundary between two views is a hub and spoke or ‘spider’ map, in which a higher level of detail is shown near the centre of the extent, with highly schematised routes that cover a larger area being displayed around the topographically correct region. Figure 1, a Lothian Buses map of the bus network in the City of Edinburgh, is an example of this. In the example in Figure 1, the main central street and the area immediately surrounding it are shown in a higher level of detail, with the city’s bus routes shown in terms of their entry and exit points in the context of this central hub area. Even though schematisation of the larger scale map has been designed specifically to

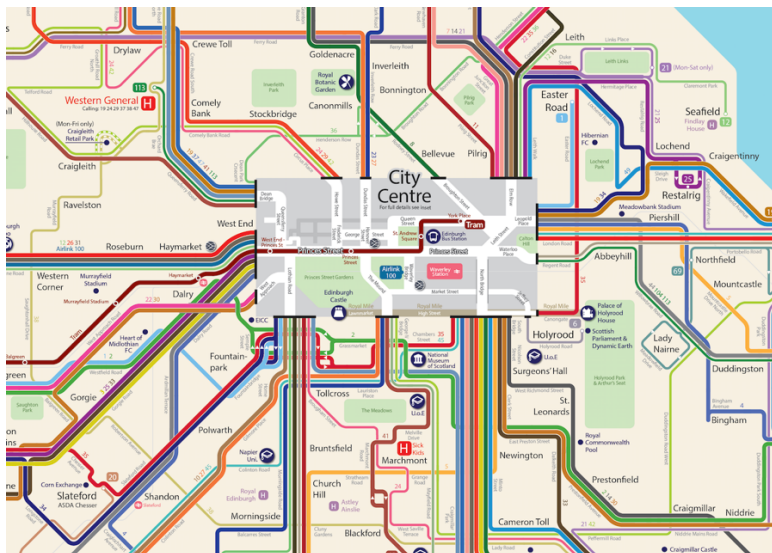


Fig 1. Lothian Buses Network Map (Finlay, 2016)


provide consistency with the roads in the smaller scale view, this design still fits within the Hard Boundary approach, with a clear boundary separating the two different representations of space.

#### 4.2 Global Transformation

Metric space in the context of cartography and GIS is synonymous with a transformation that preserves certain spatial aspects of the underlying geography – for example minimising the distortion of angles between locations (conformal). This is however a specific and narrow interpretation of metric space, with a metric space more generally being any space where a distance function is applied to all members of the set. Global Transformation is an approach that can be used to reinterpret the representation of the entire extent to communicate a derivative property of the underlying spatial structure. It refers then to metric and metric-like approaches which do not conform to traditional representational objectives. While there is no upper limit on the variety of applications of this approach, it is most often used as a way of representing spatial relationships in terms of time geography (Hägerstrand, 1985). A map in this case, is used to communicate time taken to travel as a derivative of the lived experience of the geography, as opposed to the absolute length of the distances within the given scale. Simple applications of this approach include the emphasis of denser urban areas which typically take longer to traverse than more suburban areas, despite the absolute distances being shorter.

Another example of Global Transformation is Benedikt Gross’s ‘Metrography – London Tube Map to large scale collective mental map’ (Figure 2). It uses the structure of a non-metric space to transform the metric representation of the city. This is not a map we could use for day-to-

day wayfinding, but we argue it is a mapping of people's sense of place insofar that millions of Londoners organise their thinking around the distorted space of the London Underground schematic.



**Fig 2.** Metrography - London tube map to large scale collective mental map (Gross, 2012)

Such a representation reminds us of Anchor Point Hypothesis posited by Couclelis et al. (1987). The Anchor-Point Hypothesis describes the phenomena in which people's cognitive conception of space is distorted by their variable level of familiarity and association with certain parts of a geographical environment, and that this leads to a unique, personal understanding of the structure of that environment. Technology now enables the modelling of an individual's familiarity and associations thus enabling cognitive and behavioural aspects to take a more central role in the design of personalised maps.


This idea has been further explored through Gross's 'MapMap Vauxhall' project, in which Gross created spatial transformations of a region of London based on research participant's mental conception of the area expressed through the recording of individual sketch maps (Figure 3). While this process was manual rather than automated, it is nevertheless an interesting example of a Global Transformation methodology that redefines the bounds of distortion in order to fit more closely with an individual's natural perception of space.



**Fig 3.** MapMap Vauxhall - Mashup Mental Maps and OpenStreetMap (Gross, [online]).

### 4.3 Focus Maps

The third approach is through the use of ‘focus maps’, with early investigations into the efficacy of maps on mobile devices using this approach (Zipf et al., 2002). Focus maps use computational and graphical techniques to draw the user’s attention to certain aspects of the geography.



**Fig 4.** Two mapping functions applied to the vertices of a regular grid (Hauert et al., 2011) To that end they seek to bridge the functional and structural approaches to map content by emphasising some aspect of the environment to support decision making (Figure 5).

Harrie et al.’s variable-scale mapping for small scale cartography is an early example of integrating multiple scales into single map views with mobile devices in mind, however a high level of visible distortion is apparent in this approach.

‘Focus, Glue and Context’, is a term coined by Yamamoto et al. (2009). Instead of the distortion being spread across the central region (Figure 5), it refers to the idea of containing the distortion in a narrow band around the focused part of the map. This ring is the ‘glue’ between the focus area and the surrounding context, which can be seen as a variation on the ‘fisheye’ approach that aims to “...provide a balance of local detail and global context...” (Furnas, 1986).

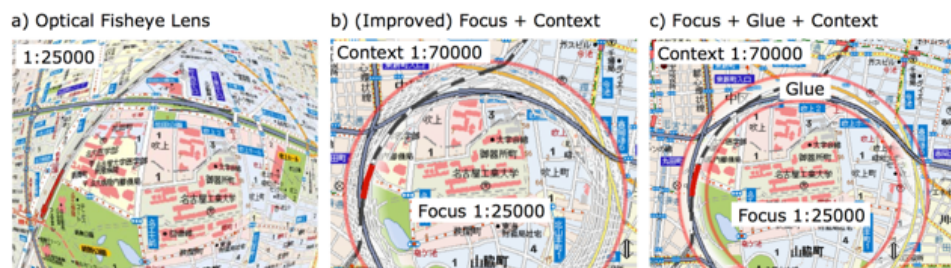


Fig 5. From optical fish-eye to ‘focus, glue and context’ (Yamamoto et al., 2009)

A user may find fisheye distortion disconcerting. By containing the distortion in the ring (the focus, glue, context approach), it is possible to make the distortion less apparent – conforming more to expected norms.

A mixing of these ideas is reflected in Haunert et al.’s use of focus in the context of representing road networks. Here we see a more subtle use of Focus, Glue and Context that supports map legibility through the use of ‘locally valid scale factors’ (Haunert et al., 2011) (Figure 7).



Fig 6. User-selected focus region and zoom factor (Haunert et al., 2011)

#### 4.4 Variable Scale Route

The fourth approach is one pioneered by Agrawala and Stotle (2001). The Variable Scale Route approach makes use of highly schematised representations to emphasise journey segments that are best displayed with varying levels of detail. The LineDrive project adopted this idea for long car journeys. The issue was long journey segments (often with the least need for guidance), consumed the largest portion of space on the device’s screen. The solution was to reduce the space devoted to long journey segments in order to provide more space for detailed segments. Here the distortion varies as a function of the complexity of decision making.

As with other approaches to the automation of schematic maps for mobile devices (e.g. Anand et al., 2004, Anand et al., 2006), maps that use the Variable Scale Route approach conceal the nature of this distortion by suppressing structural and contextualising information that would otherwise expose the distortion - the focus being on the conveyance of functional information almost exclusively (Klippel et al., 2005).



Fig 7. LineDrive route map (Agrawala et al., 2001)

#### 4.5 Variable Level of Detail

The prevailing design paradigm for digital maps makes use of variable levels of detail in so far as the level of detail is changed after each generalisation point is crossed (Ratajski, 1967), i.e. a transition across ‘conceptual cusps’ based on either zooming in or zooming out. What is outside of the current bounds of distortion however is a variable level of detail within each scale, based on the user task. Thus Variable Level of Detail in the context of redefining the bounds of distortion can be seen as another form of focus map, with functional information given precedence but without completely suppressing the broader structural content of the map, as with the Variable Scale Route. In the Chorematic Focus Map method (Klippel et al., 2004) and other approaches within this category, the objective is to tease out the salient information but to leave an intermediary level of detail that supports navigation and goes beyond a simple A-B route representation.

In Mackaness et al. (2011), multimodal journeys that cross transport networks are shown with varying levels of detail dependent on sub task. For example a higher level of detail would be shown where the task is on foot as opposed to via public transport (see Figure 9.). In a similar approach, Schmid et al. (2010) proposed emphasising detail at the start and end of the journey (as is often required in prototypical journeys), however this was proposed within the context of ‘route aware’ maps that embed emphasised route information within the broader structural map content, which can from there be further simplified to provide map views that convey ‘disambiguated’ start and end environments (as in illustration ‘c’ in Figure 10).



Fig 8. Rule-based output of varying level of detail for a multimodal route (Mackanness et al., 2011)




Fig 9. Survey map, extracted route and route with disambiguated start environment and extended destination environment (Schmid et al., 2010)

## 5. Key opportunities for future work

By considering the strengths and weaknesses of these five design approaches we propose a research agenda that seeks to improve map efficacy through a reworking of these bounds of distortion in the context of digital and personalised mapping. Ultimately, the aim is to formalise a design framework that supports an increase in comprehension on the part of the map user, however the challenge is finding where this revised ‘upper bound’ to the level of distortion sits for each context.

### 5.1 Developing and validating solutions

In terms of developing and testing prototypes to determine which approaches result in a shift ‘up’ the comprehension curve, and which result in a decrease in the user’s comprehension, it is very much contextual. The optimal level of distortion is relative to the type of geography and

the specific task the map user is engaged in, for example finding a residential address on foot versus finding a landmark on public transport. This issue of context is however central to the argument presented in this paper.

### ***5.2 Integrating structural and functional information: variable forms of spatial representation in single map views***

A key area that is yet to be effectively designed and implemented is the presentation of multiple forms of space in single map views. The aim here would be to develop Global Transformation operations such that we can derive and visualise the optimal form of visualisation for each stage of the journey – i.e. for each ‘sub-task’ (Mackness et al., 2011). While Schmid et al.’s route-aware maps go some way to achieving this, route-aware maps only address Variable Level of Detail, and do not allow for the integration of varying underlying spatial representations. While Variable Scale Route maps can integrate different forms of space, they do this at the expense of geographical context, so the ambition here is to find an approach to automation that effectively satisfies both a route and a survey orientation strategy.

## **6. Concluding comments**

We have introduced the bounds of distortion as a conceptual device that supports the analysis of established map design, with a particular aim to lay the foundation for further work to redefine these bounds in the context of the efficacy of maps for navigation. While truth is a key issue in cartographic representation, we have argued that the recognition of meaning is of greater influence on the effectiveness of navigational information, and that making better use of the limited space on the screens of our ubiquitous mobile devices means that there is considerable potential value in extending the work already done to redefine the way distortion is applied in cartographic design. We reviewed five high level approaches, and have highlighted their strengths and weaknesses in the context of users’ core orientation strategies.

To date generalisation techniques have focused on algorithms that respond to legibility constraints within narrow changes of scale, with a focus on topographic mapping (with emphasis on locational accuracy). We would propose a broader and more ambitious remit that considers wider scale transitions across a broader range of thematic and schematic maps. And even beyond this, a remit that explores a significant loosening of the boundaries of distortion, one that facilitates the optimal embedding of multi scale spaces, and ultimately one that chimes more closely our conceptualisation of, and familiarity with place.

## **Acknowledgements**

We would like to thank the Ordnance Survey and the United Kingdom Engineering and Physical Sciences Research Council for their support.

## **References**

Anand, S., Ware, J.M. and Taylor, G.E. (2004). Map Generalization for OSMasterMap Data in Location Based Services & Mobile CIS Applications, Proceedings of 12th International Conference on Geoinformatics - GeoSpatial Information Research : Bridging the Pacific and Atlantic, pp 54-60

Anand, S., Taylor, G., Thomas, N. and Ware, J.M. (2006). Automatic Production of Schematic Maps for Mobile Applications. Transactions in GIS. Blackwell Publishing, Jan.

Anand, S, Ware J.M. Taylor G.E., (2006). Generalization and schematization of large scale digital geographic datasets for MobileGIS applications. In: Dynamic & mobile GIS: investigating change in space and time. Taylor & Francis Ltd, London.

Agrawala, M. and Stotle, C. (2001). Rendering Effective Route Maps: Improving Usability Through Generalization. SIGGRAPH'01 Proceedings of the 28th annual conference on Computer graphics and interactive techniques, pp241-249

Carden, T. (2011). Travel Time Tube Map [online]  
Available at:  
<https://randometc.github.io/travel-time-tube-d3/>

Couclelis, H., Gale, N., Golledge, R., & Tobler, W. (1987). Exploring the anchor-point hypothesis in spatial cognition. Journal of Environmental Psychology 7, 99-122.

Downs, R. and Stea, D. (1977). Maps in Minds: Reflections on Cognitive Mapping. New York, Harper and Row.

Edwardes, A., Burghardt, D. and Weibel, R. (2005). Portrayal and Generalisation of Point Maps for Mobile Information Services. Map-Based Mobile Services. Meng, L., Zipf, A., and Reichenbacher, T. (Eds.). Springer.

Finlay, K. (2016). Lothian Buses Network Map. Transport for Edinburgh

Freksa, C., Barkowsky, T., Klippel, A. and Richter, K. (2005). The Cognitive Reality of Schematic Maps. Map-Based Mobile Services. Springer.

Freundschuh, S. M. and Egenhofer, M. J. (1997). Human conceptions of spaces: Implications for geographic information systems. Transactions in GIS, 2, 361–375

Furnas, G.W. (1986). Generalized Fisheye Views. Human Factors in Computing Systems CHI '86 Conference Proceedings, 16-23. 1986.

Gross, B. (2012). Metrography – London Tube Map to large scale collective mental map [online]  
Available at:  
<http://scale-collectibenedikt-gross.de/log/2012/02/metrography-london-tube-map-to-large-scale-collective-mental-map>

Gross, B. [online]. MapMap Vauxhall - Mashup Mental Maps and OpenStreetMap  
Available at:  
<http://benedikt-gross.de/log/2011/10/mapmap-vauxhall-mashup-mental-maps-and-openstreetmap/>

Hägerstrand, T. (1985). Time-geography: Focus on the Corporeality of Man, Society, and Environment. In: The Science and Praxis of Complexity. Tokyo: The United Nations University. pp. 193-216.

Hauert, J.H. and Sering.L. (2011) Drawing Road Networks with Focus Regions. IEEE Transactions on Visualization and Computer Graphics (Proc. Information Visualization 2011), 17(12):2555-2562

Harrie, L., Sarjakoski, L. T. and Lehto, L. (2002). 'A variable-scale map for small-display cartography', Proceedings of the Joint International Symposium on 'GeoSpatial Theory, Processing and Applications' (ISPRS / Commission IV, SDH2002), Ottawa, Canada.

Held, R. (Ed). (1971). Image, Object and Illusion. W.H. Freeman and Company

Klippel, A. and Richter, K.F. (2004). Chorematic Focus Maps. In Georg Gartner (Ed.), Location Based Services & Telectography, pp. 39–44, Geowissenschaftliche Mitteilungen. Technische Universität Wien, Wien.

Hauert, J.H. and Sering.L. (2011) Drawing Road Networks with Focus Regions. IEEE Transactions on Visualization and Computer Graphics (Proc. Information Visualization 2011), 17(12):2555-2562

- Kirsh, D. (2009). Projection, Problem Space and Anchoring. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2310-2315). Austin, TX: Cognitive Science Society.
- Klippel, A., Knuf, L., Hommel, B., & Freska, C. (2004). Perceptually Induced Distortion in Cognitive Maps. *Spatial Cognition IV: Reasoning, Action, Interaction*, 204-213
- Klippel, A. and Richter, K.F. (2004). Chorematic Focus Maps. In Georg Gartner (Ed.), *Location Based Services & Telecartography*, pp. 39-44, *Geowissenschaftliche Mitteilungen*. Technische Universität Wien, Wien.
- Mackness, W. and Chaudhry, O. (2005). Exploring representational issues in the visualisation of the geographical phenomenon over large changes in scale. GISRUK, Glasgow
- Mackness, W., Burghardt, D. and Duchene, C. (2014). *Map Generalisation: Fundamental to the Modelling. Abstracting Geographic Information in a Data Rich World*, Burghardt, D., Duchene, C. and Mackness, W. (Eds.). *Lecture Notes in Geoinformation and Cartography*. Springer
- Mackness, W.A., Quigley, A. and Tanasescu, V. (2011) *Hierarchical Structures in Support of Dynamic Presentation of Multi Resolution Geographic Information for Navigation in Urban Environments*. GISRUK.
- Mackness, W. and Reimer, A. (2014). *Generalisation in the Context of Schematised Maps. Abstracting Geographic Information in a Data Rich World*, Burghardt, D., Duchene, C. and Mackness, W. (Eds.). *Lecture Notes in Geoinformation and Cartography*. Springer
- Muller, J. C. (1991). *Generalisation of Spatial Databases*, in Maguire, D. J., Goodchild, M., and Rhind, D. (Eds.), *Geographical Information Systems: London*, Longman Scientific, p. 457-475.
- Ratajski, L. (1967): "Phenomene des points de generalisation. In *International Yearbook of Cartography*, K. Kirschbaum and K. H. Meine, Eds. Bonn-Bad: Godesberg, Vol 7, pp. 143-152.
- Robinson. A.H., Morrison, J.L., Muehrcke, P.C., Kimerling, A.J. and Guptil, S.C. (1995). *Elements of Cartography*. John Wiley & Sons, 6th Ed.
- Richardus, P. and Adler, R.K. (1972). *Map Projections*. North-Holland Publishing Company.
- Richter, K.F., Peters, D., Kuhnmünch, G., and Schmid, F. (2008). What Do Focus Maps Focus On? *Spatial Cognition VI – Learning, Reasoning, and Talking about Space*, Freksa, C., Newcombe, Gärdenfors, N.P., and Wöflfl, S. (Eds.). 154-70. Berlin: Springer.
- Robinson. A.H., Morrison, J.L., Muehrcke, P.C., Kimerling, A.J. and Guptil, S.C. (1995). *Elements of Cartography*. John Wiley & Sons, 6th Ed.
- Schlichtmann, H. (2009). Overview of the semiotics of maps. *Proceedings of the 24th International Cartographic Conference, Santiago de Chile, Chile, 15-21 November* [online]  
Available at:  
[http://icaci.org/files/documents/ICC\\_proceedings/ICC2009/](http://icaci.org/files/documents/ICC_proceedings/ICC2009/)
- Schmid, F. (2008). Knowledge based wayfinding maps for small display cartography. *Journal of Location Based Services*, 2:1, 57-83, DOI: 10.1080/17489720802279544
- Shmid, F., Richter, K.F. and Peters, D. (2010). Route Aware Maps: Multigranular Wayfinding Assistance. *Spatial Cognition & Computation*, 10:2-3, 184-206, DOI: 10.1080/13875861003592748
- Skagerlund, K., Kirsh, D. and Dahlbäck, N. (2012). Maps in the Head and Maps in the Hand. *Proceedings of the 34th Annual Cognitive Science Society*.
- Tversky, B. (2015). How Diagrams Communicate. ISVIS [online]  
Available at:  
[https://www.youtube.com/watch?v=B\\_X0PGgJjTs](https://www.youtube.com/watch?v=B_X0PGgJjTs)
- Yamamoto, D., Ozeki, S. and Takahashi, N. (2009). Focus+Glue+Context: An Improved Fisheye Approach for Web Map Services. *ACM GIS '09*, November 4-6, 2009. Seattle, WA
- Zipf, A. and Richter, K.F. (2002). Using Focus Maps to Ease Map Reading: Developing Smart Applications for Mobile Devices. *KI Special Issue Spatial Cognition 02(4)* 35

