

Theoretical Computational Models for the Cognitive Map

Dissertation

der Mathematisch-Naturwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen
zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

vorgelegt von
Tristan Baumann
aus Konstanz

Tübingen
2023

Gedruckt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Eberhard Karls Universität Tübingen.

Tag der mündlichen Qualifikation:

21.11.2023

Dekan:

Prof. Dr. Thilo Stehle

1. Berichterstatter:

Prof. Dr. Hanspeter A. Mallot

2. Berichterstatter:

Prof. Dr. Martin V. Butz

Contents

Abbreviations and symbols	7
Abbreviations	7
Mathematical symbols	8
Summary	11
Zusammenfassung	13
List of publications	15
Authorship	17
Bibliography	19
1 Introduction	21
1.1 Navigation: from routes to maps	22
1.1.1 Stimulus-response association	22
1.1.2 The cognitive map	24
1.2 The neural basis of the cognitive map	26
1.3 Wayfinding and regions	29
1.4 Structure and content of the cognitive map	30
1.5 Aim of this thesis	33
Bibliography	35
2 Results	41
2.1 Mallot et al. (2020)	41
2.2 Brucklacher et al. (2021)	42
2.3 Baumann and Mallot (2023a)	43
2.4 Baumann and Mallot (2023b)	44
Bibliography	44
3 Discussion	47
3.1 Population coding of routes and wayfinding	48
3.1.1 The role of preplays and graph search	50
3.2 Regions and remapping	52
3.2.1 Neurological evidence for region coding	55
3.3 The form of spatial knowledge	56
3.3.1 Non-spatial domains	59

3.4	Open questions and future research	60
	Bibliography	61
Appendix: Individual publications		65
A1 Mallot et al. (2020)		67
A1.1	Introduction	67
	A1.1.1 Parsimonious Representations of Space	67
	A1.1.2 Dual Population Coding	69
A1.2	Navigation Algorithm	70
	A1.2.1 Feature Detection	70
	A1.2.2 Feature Matching	72
	A1.2.3 Graph Edge Formation	73
	A1.2.4 Edge Labeling and Reference Direction	75
	A1.2.5 Pathfinding and Voting	77
A1.3	Evaluation	80
	Bibliography	83
A2 Brucklacher et al. (2021)		85
A2.1	Introduction	85
A2.2	Path-planning in neural state-action networks	86
A2.3	Results	88
A2.4	Discussion	91
A2.5	Conclusion	92
	Bibliography	93
A3 Baumann and Mallot (2023a)		95
A3.1	Introduction	96
A3.2	Materials and Methods	101
	A3.2.1 Model overview	101
	A3.2.2 Grid Cells	103
	A3.2.3 Place Cells	106
	A3.2.4 Combined attractor and inverse place cell - grid cell activation	106
	A3.2.5 Remapping and gateway database	107
	A3.2.6 Simulation and exploration procedures	109
	A3.2.7 Virtual environments	110
	A3.2.8 Data analysis	111
A3.3	Results	112
	A3.3.1 Parallel rooms (Fig. A3.5)	112
	A3.3.2 Hairpin maze (Fig A3.6)	114
	A3.3.3 L-shaped rooms (Fig. A3.7)	116
A3.4	Discussion	118
	A3.4.1 Implications for biology	119
	A3.4.2 Remapping and place learning	121
	A3.4.3 Remapping without the MEC	123

Bibliography	124
A4 Baumann and Mallot (2023b)	131
A4.1 Introduction	132
A4.1.1 The cognitive map	132
A4.1.2 Distorted maps and non-Euclidean environments	135
A4.1.3 Evidence from wormhole experiments	136
A4.2 Materials and methods	137
A4.2.1 Data acquisition	137
A4.2.2 Graph and map setup	139
A4.2.3 Model comparison and data analysis	141
A4.3 Results	141
A4.3.1 Embeddings	141
A4.3.2 Dataset 1: Route-finding and shortcuts	143
A4.3.3 Dataset 2: Rips and folds	143
A4.4 Discussion	145
Bibliography	147

Abbreviations and symbols

Abbreviations

In general, terms in this thesis are spelled out in full and not abbreviated, except for the cases below. The list includes mostly terms where the abbreviation is more common than the full name, such as CA1 or fMRI.

CA1-4	<i>Cornu Ammonis</i> , subfields of the hippocampus.
CI	Confidence interval
fMRI	Functional magnetic resonance imaging
HC	Hippocampus
HRL	Hierarchical reinforcement learning
MDS	Multidimensional scaling
MEC	Medial entorhinal cortex
mPFC	Medial prefrontal cortex
MWU	Mann-Whitney-U-test
PC	Place cell
RSC	Retrosplenial cortex
SLSQP	Sequential least squares programming
STDP	Spike-time dependent plasticity
SURF	<i>Speeded-up Robust Features</i> , image feature detection algorithm.
SVF	<i>Synaptic vector field</i> , assigns movement vectors to positions.
U-SURF	Upright SURF
VR	Virtual reality

Mathematical symbols

Mathematical symbols for each included publication in alphabetical order. Latin and Greek letters are listed separately. Bold symbols refer to vectors or tensors. The codes [A1-4] refer to the full texts in the appendix and are used throughout the thesis.

[A1] Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations

c	Graph update counter
\mathbf{d}_i	SURF feature descriptor, 64-dimensional vector
E	Set of all edges α_{ij}
F	Set of all SURF features f
F_g	Set of features visible at the goal
F_t	Set of features visible at time step t
f	Individual SURF feature
i, j, k, l	Feature enumerators
J_t	Indices of outgoing edges α_{ij} at time point t
l	Learning step
N_i	Neighborhood, set of simultaneously visible features to f_i
P	Set (Bundle) of paths p
p	Dijkstra path
t	Time step
α_{ij}	Direction from f_i to f_j
β_i	Bearing of f_i
$\hat{\beta}_i$	Stored bearing of f_i
γ	Generic angle
$\boldsymbol{\eta}$	Heading vector
θ_N	Neighborhood threshold
θ_S	Descriptor similarity threshold
κ	Stiffness constant
ν	Allocentric reference direction

[A2] Hierarchical Planning in Multilayered State-Action Networks

A^-	Weight update strength constant
j, k	Position-encoding neurons
s_j	Time since the last spike of neuron j
$w_{k,j}$	Connection from neuron k to j
w^{P1}	Connections to directly neighboring neurons
w^{P2}	Connections to neurons one position away
w^{PR}	Connections between region and position neurons
w^{RR}	Connections between region neurons
τ_0, τ_{STDP}	time constants

[A3] Gateway identity and spatial remapping in a combined grid and place cell attractor

D	Gateway database
d	Periodic distance function
e	Percentage of winning cells in \max_e
\mathbf{G}	Grid cell activity
$\hat{\mathbf{G}}$	Grid cell activity estimated from place cells
g_{ij}	Activity of grid cell ij
\mathbf{H}	Cartesian coordinate transformation matrix
I	Weight function synaptic strength constant
ij	Grid cell at position (i, j) in the grid
\max_e	Winner-take-all nonlinearity
n^2	Number of cells per module, arranged in a $n \times n$ square
\mathbf{P}	Place cell activity
\mathbf{R}_β	Grid orientation rotation matrix
\mathbf{s}_{pq}	Shifts around the standard rhomboid
T	Weight function global inhibition constant
t	Time step
\mathcal{U}	Uniform distribution
$\mathbf{v}(t)$	True velocity at time point t

Abbreviations and symbols

$\hat{\mathbf{v}}(t)$	Estimated velocity at time point t
\mathbf{W}_g^p	Weights from grid to place cells
w_k^{ijl}	Connection weights from cell ij to cell kl
\mathbf{x}_{ij}	Cartesian coordinates of grid cell ij
x_{\max}	Highest activity input in \max_e
α	Gain parameter defining grid spacing
β	Grid orientation
Γ	Gateway identifier
ϵ	Estimation error
σ	Weight function Gaussian width

[A4] Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments

d_{ij}	Distance from i to j
E	Set of edges
e_{ij}	Edge from v_i to v_j
G	Graph
i, j, k	Places in the maze
n	Number of vertices
p, q	Embedding solution enumerators
T	Triplet of neighboring places
V	Set of vertices
v_i	Vertex at place i
X	Set of vertex positions
x_i	Position of v_i
α_{ijk}	Heading change at j when moving from i to k
λ_1, λ_2	Stress function weight constants

Summary

Decades of research into the neural representation of physical space have uncovered a complex and distributed network of specialized cells in the mammalian brain. It is now clear that space is represented in some form, but the realization remains debated. Accordingly, the overall aim of my thesis is to further the understanding of the neural representation of space, the *cognitive map*, with the aid of theoretical computational modeling (as opposed to data-driven modeling). It consists of four separate publications which approach the problem from different but complementing perspectives:

The first two publications consider goal-directed navigation with topological graph models, which encode the environment as a state-action graph of local positions connected by simple movement instructions. Graph models are often less constrained than coordinate-based metric maps and offer a variety of computational advantages; for example, graph search algorithms may be used to derive optimal routes between arbitrary positions. In the first model, places are encoded by population codes of low-level image features. For goal-directed navigation, a set of simultaneous paths is obtained between the start and goal populations and the final trajectory follows the population average. This makes route following more robust and circumvents problems related to place recognition. The second model proposes a hierarchical place graph which subdivides the known environment into well-defined regions. The region knowledge is included in the graph as superordinate nodes. During wayfinding, these nodes distort the resulting paths in a way that matches region-related biases observed in human navigation experiments.

The third publication also considers region coding but focuses on more concrete biological implementation in the form of place cell and grid cell activity. As opposed to unique nodes in a graph, place cells may express multiple firing fields in different contexts or regions. This phenomenon is known as “remapping” and may be fundamental to the encoding region knowledge. The dynamics are modeled in a joint attractor neural network of place and grid cells: Whenever a virtual agent moves

Summary

into another region, the context changes and the model remaps the cell activity to an associated pattern from memory. The model is able to replicate experimental findings in a series of mazes and may therefore be an explanation for the observed activity in the biological brain.

The fourth publication again returns to graph models, joining the debate on the fundamental structure of the cognitive map: The internal representation of space has often been argued to either take the form of a non-metric topological graph or a Euclidean metric map in which places are assigned specific coordinates. While the Euclidean map is more powerful, human navigation in experiments often strongly deviates from a (correct) metric prediction, which has been taken as an argument for the non-metric alternative. However, it may also be possible to find an alternative metric explanation to the non-metric graphs by embedding the latter into metric space. The method is shown with a specific non-Euclidean example environment where it can explain subject behavior equally well to the purely non-metric graph, and it is argued that it is therefore a better model for spatial knowledge.

Beyond the individual results, the thesis discusses the commonalities of the models and how they compare to current research on the cognitive map. I also consider how the findings may be combined into more complex models to further the understanding of the cognitive neuroscience of space.

Zusammenfassung

In den letzten Jahrzehnten hat die Forschung nach der Frage, wie Raum im Gehirn repräsentiert wird, ein weit verzweigtes Netzwerk von spezialisierten Zellen aufgedeckt. Es ist nun klar, dass Räumlichkeit auf irgendeine Art repräsentiert sein muss, aber die genaue Umsetzung wird nach wie vor debattiert. Folgerichtig liegt das übergeordnete Ziel meiner Dissertation darin, das Verständnis von der neuronalen Repräsentation, der *Kognitiven Karte*, mithilfe von theoretischer Computermodellierung (im Gegensatz zu datengetriebener Modellierung) zu erweitern. Die Arbeit setzt sich aus vier Publikationen zusammen, die das Problem aus verschiedenen, aber miteinander kompatiblen Richtungen angehen:

In den ersten beiden Publikationen geht es um zielgerichtete Navigation durch topologische Graphen, in denen die erkundete Umgebung als Netzwerk aus lokalen Positionen und sie verbindenden Handlungen dargestellt wird. Im Gegensatz zu Koordinaten-basierten metrischen Karten sind Graphenmodelle weniger gebunden und haben verschiedene Vorteile wie z.B. Algorithmen, die garantiert optimale Pfade finden. Im ersten Modell sind Orte durch Populationen von einfachen Bildfeatures im Graphen gespeichert. Für die Navigation werden dann mehrere Pfade gleichzeitig zwischen Start- und Zielpopulationen berechnet und die schlussendliche Route folgt dem Durchschnitt der Pfade. Diese Methode macht die Wegsuche robuster und umgeht das Problem, Orte entlang der Route wiedererkennen zu müssen.

In der zweiten Publikation wird ein hierarchisches Graphenmodell vorgeschlagen, bei dem die Umgebung in mehrere Regionen unterteilt ist. Das Regionenwissen ist ebenfalls als übergeordnete Knoten im Graphen gespeichert. Diese Struktur führt bei der Wegsuche dazu, dass die berechneten Routen verzerrt sind, was mit dem Verhalten von menschlichen Probanden in Navigationsstudien übereinstimmt.

In der dritten Publikation geht es auch um Regionen, der Fokus liegt aber auf der konkreten biologischen Umsetzung in Form von *Place Cell* und *Grid Cell*-Aktivität. Im Gegensatz zu einzigartigen Ortsknoten im Graphen zeigen Place Cells multiple Feuerfelder in verschiedenen Regionen oder Kontexten. Dieses Phänomen wird als

Remapping bezeichnet und könnte der Mechanismus hinter Regionenwissen sein. Wir modellieren das Phänomen mithilfe eines Attraktor-Netzwerks aus Place- und Grid Cells: Immer, wenn sich der virtuelle Agent des Modells von einer Region in eine andere bewegt, verändert sich der Kontext und die Zellaktivität springt zu einem anderen Attraktor, was zu einem Remapping führt. Das Modell kann die Zellaktivität von Tieren in mehreren Experimentalumgebungen replizieren und ist daher eine plausible Erklärung für die Vorgänge im biologischen Gehirn.

In der vierten Publikation geht es um den Vergleich von Graphen- und Kartenmodellen als fundamentale Struktur der kognitiven Karte. Im Speziellen geht es bei dieser Debatte um die Unterscheidung zwischen nicht-metrischen Graphen und metrischen euklidischen Karten; euklidische Karten sind zwar mächtiger als die Alternative, aber menschliche Probanden neigen dazu, Fehler zu machen, die stark von einer metrischen Vorhersage abweichen. Deshalb wird häufig argumentiert, dass nicht-metrische Modelle das Verhalten besser erklären können. Wir schlagen eine alternative metrische Erklärung für die nichtmetrischen Graphen vor, indem wir die Graphen im metrischen Raum einbetten. Die Methode wird in einer bestimmten nicht-euklidischen Beispielumgebung gezeigt, in der sie Versuchspersonenverhalten genauso gut vorhersagen kann, wie ein nichtmetrischer Graph. Wir argumentieren daher, dass unser Modell ein besseres Modell für Raumrepräsentation sein könnte.

Zusätzlich zu den Einzelergebnissen diskutiere ich außerdem die Gemeinsamkeiten der Modelle und wie sie in den derzeitigen Stand der Forschung zur kognitiven Karte passen. Darüber hinaus erörtere ich, wie die Ergebnisse zu komplexeren Modellen vereint werden könnten, um unser Bild der Raumkognition zu erweitern.

List of publications

This thesis is comprised of four separate publications. The first three have been accepted and published, and the fourth – *Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments* – is currently in preparation. The codes **A1-4** refer to the full texts in the appendix and are used throughout the thesis.

- A1** Hanspeter A Mallot et al. (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17. Reprinted with permission.
DOI: 10.1007/978-3-030-57983-8_1
- A2** Matthias Brucklacher et al. (2021). “Hierarchical Planning in Multilayered State-Action Networks.” *ESANN 2021 proceedings*. Reprinted with permission.
DOI: 10.14428/esann/2021.ES2021-45
- A3** Tristan Baumann and Hanspeter A Mallot (2023a). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239. Reprinted with permission.
DOI: 10.1016/j.neunet.2022.10.019
- A4** Tristan Baumann and Hanspeter A Mallot (2023b). “Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments”. Under review.
DOI: 10.1101/2023.06.09.544331

Authorship

Nr.	Accepted publication yes/no	List of authors ¹	Position of candidate in list of authors	Scientific ideas by the candidate (%)	Data generation by the candidate (%)	Analysis and interpretation by the candidate (%)	Paper writing done by the candidate (%)
1	yes	M, E, Ba	3	30	100	100	50
2	yes	Br, M, Ba	3	20	0	0	20
3	yes	Ba, M	1	70	100	70	70
4	no	Ba, M	1	60	100	80	90

¹ Authors: **Ba** Tristan Baumann, **Br** Matthias Brucklacher, **E** Gerrit A Ecke, **M** Hanspeter A Mallot.

Table 1: Personal contribution to the individual publications in approximate percentages. The table is the same as in the separate form *Erklärung zum Anteil an gemeinschaftlichen Veröffentlichungen* (Declaration for the contribution to collaborative publications).

[A1] Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations (Mallot, Ecke, and Baumann, 2020)

The basic idea of feature-guided navigation dates back to my bachelor’s and master’s theses (*3D Navigation using Microsnapshots*, 2016 and *Microsnapshot Navigation*, 2019) under the supervision of Hanspeter Mallot and Gerrit Ecke. The underlying theory and concept behind the algorithm were worked out in collaboration with my supervisors, and I coded the simulation and performed data analysis. The publication is a summary of my master’s thesis with a shifted focus on population coding, written in equal parts by Hanspeter Mallot and me. The text was also published as part of Gerrit Ecke’s doctoral thesis (Ecke, 2022).

[A2] Hierarchical Planning in Multilayered State-Action Networks (Brucklacher, Mallot, and Baumann, 2021)

The concept of the study was jointly developed in the lab seminar with the aim to improve the plausibility and performance of wayfinding in large-scale graphs like the microsnapshot graph in [A1] Mallot et al. (2020). The realization of the project, i.e., the implementation of the algorithm, simulation, and data analysis were performed by Matthias Brucklacher for his master’s thesis (*Topological Navigation with Hierarchically Structured Spiking Neural Networks*, 2020) under the supervision of Hanspeter Mallot and in collaboration with me. Matthias Brucklacher also wrote the publication, which is a summary of his master’s thesis, for the *ESANN 2021* conference.

My contribution to the work was mostly conceptual. In addition to general discussion, I contributed to various specifics of the implementation, especially the parts related to graph theory, and I wrote a part of the introduction.

[A3] Gateway identity and spatial remapping in a combined grid and place cell attractor (Baumann and Mallot, 2023a)

I initially started the project as a more biologically plausible implementation of the previous graph-based models ([A1] Mallot et al., 2020 or [A2] Brucklacher et al., 2021) on simulated place cells, which then turned into a grid and place cell attractor model for remapping in multi-compartment environments.

I devised the initial concept, wrote the computational models and simulations, and collected and analyzed the data. Hanspeter Mallot helped in refining the theory and interpreting the results. He also contributed to the mathematical description of the model and wrote the abstract and parts of the introduction. The text was otherwise written by me and refined in collaboration.

[A4] Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments (Baumann and Mallot, 2023b)

The project started with the bachelor’s thesis of Katja Körner (*Microsnapshot Navigation in a Non-Euclidean Environment – Embedding of Feature-Graphs*, 2020) under the supervision of Hanspeter Mallot and me. In her thesis, Katja Körner at-

tempted to find plausible embeddings of a microsnapshot graph ([A1] Mallot et al., 2020) using multidimensional scaling (MDS), with mixed results. The theory was refined by Hanspeter Mallot and me, and I reattempted the embedding with a less complex graph and an optimization-based approach instead of MDS. The implementation of the algorithm, computational modeling, data generation, and writing were all done by me with feedback from Hanspeter Mallot. The results were discussed collaboratively in the lab seminar.

Bibliography

- Baumann, Tristan and Hanspeter A Mallot (2023a). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239.
- Baumann, Tristan and Hanspeter A Mallot (2023b). “Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments”. Under review.
- Brucklacher, Matthias, Hanspeter A Mallot, and Tristan Baumann (2021). “Hierarchical Planning in Multilayered State-Action Networks.” *ESANN 2021 proceedings*.
- Ecke, Gerrit (2022). “On the relationship between neuronal codes and mental models”. PhD thesis. Universität Tübingen.
- Mallot, Hanspeter A, Gerrit A Ecke, and Tristan Baumann (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17.

Chapter 1

Introduction

In the late 17th century, famous physicist Isaac Newton originally defined the concept of absolute, observer-independent physical space as a theoretical basis for his *Principia*. The question whether such a space exists and if it can be perceived *as such* kicked off a centuries-long and still ongoing philosophical debate involving famous names like Berkeley, Hume, or Kant (O’Keefe and Nadel, 1978).

The argument revolves around the problem that, by definition, an observer-independent space cannot be observed. What remains is the subjective experience, that is, the perception and internal representation of space and its structure. Points of dispute are, for example, if this subjective experience of space is learned or innate, how it is understood in relation to the world or the observer, or if and how it follows mathematical and geometrical rules. Regardless of its philosophical status, if space is processed by our cognitive apparatus, its experience becomes an issue of cognitive and computational psychology and may be studied with empirical and theoretical methods (O’Keefe and Nadel, 1978; Mallot, 2024).

The modern research field that studies how space is perceived and processed is called “spatial cognition” and was arguably established by John O’Keefe and Lynn Nadel with the publication of their book *The Hippocampus as a Cognitive Map* in 1978. The lab had famously discovered place-encoding cells in the rat hippocampus a few years earlier (O’Keefe and Dostrovsky, 1971; O’Keefe, 1976). For the discovery of these “place cells”, John O’Keefe was awarded the 2014 Nobel prize in Physiology or Medicine together with Edvard Moser and May-Britt Moser for their discovery of the related position-encoding “grid cells”.

Today, 45 years later, *The Hippocampus as a Cognitive Map* remains the theoretical pillar on which nearly all subsequent study of spatial cognition and the hippocampal formation rests. Key questions in the field concern the acquisition of

spatial knowledge, its properties, how we use it, how the understanding of space changes over the life span, and how it varies from individual to individual (Nadel, 2013; Warren et al., 2017).

The perception and utilization of space is closely interlinked with many other cognitive domains including vision, object recognition, social cognition, language, and the understanding of events, actions, and causality. Among these, spatial cognition is the most widespread and is considered to be the evolutionary predecessor of the other domains (Mallot, 2024; Peer et al., 2021).

In this thesis, I present four computational models of spatial cognition which revolve around an internal representation of space, the so-called “cognitive map”. In the following, this concept is introduced, as well as the related topics navigation, wayfinding, hierarchy, the involved brain areas and neurons, remapping, and the general organization of the spatial knowledge.

1.1 Navigation: from routes to maps

“Navigation is the process of determining and maintaining a course or trajectory from one place to another.”

– Gallistel, 1990, p. 35.

This statement concisely describes navigation as a process with four parts: a start, a goal, the planning of a route, and finally performing the selected behavior. In the course of evolution, organisms have developed a multitude of strategies to deal with each part in various ways. Generalizing broadly, the strategies may be divided into two frameworks depending on their complexity (O’Keefe and Nadel, 1978, [A1] Mallot et al., 2020):

- (1) Stimulus-response association, and
- (2) the representation of places and their relations in a cognitive map.

1.1.1 Stimulus-response association

The term *stimulus-response association* or *recognition-triggered response* describes an innate or learned reflexive or reflex-like behavior that is triggered upon the recognition of a specific cue (Trullier et al., 1997; [A1] Mallot et al., 2020). In the literature, this framework has also been called “means-end-relations” (Tolman, 1932), “taxon

system” (O’Keefe and Nadel, 1978) or “control laws” (Kuipers, 1978). In a stimulus-response association, behavior is inflexible and entirely depends on the particular cues that may be encountered by the animal. These include, for example, phototaxis (orienting towards light), chemotaxis (following a chemical gradient), or thigmotaxis (moving along a wall) (Mallot, 2024). The approach of a target landmark, which involves orienting the body towards an observable goal, is referred to as “guidance”, “piloting”, or “homing” if the landmark is the animal’s home location (O’Keefe and Nadel, 1978; Trullier et al., 1997; Franz and Mallot, 2000; Eichenbaum, 2017). In a more general theoretical context, the stimulus-response association may be considered a state-action pair, which can account for more behavior than just navigation.

It directly follows that by learning stimulus-response-stimulus triplets, individual actions may be chained together to more complex sequences (e.g., Collett et al., 1998; Collett and Collett, 2002). Such a series that leads the animal from one place to the next until a distant goal is reached is called a “route”. Routes allow the animal to reach places that are not immediately observable, but the sequences alone are generally considered inflexible because they depend on the correct ordering of events from start to goal (O’Keefe and Nadel, 1978; Eichenbaum, 2017; Mallot, 2024). For example, route knowledge usually does not include detours or alternatives if a passage is blocked.

The stimulus-response and route frameworks may be maximized by considering overlapping population codes of states and routes and linking them up to topological graphs from which novel routes can be inferred from the combination of known route segments (Trullier et al., 1997; Mallot and Basten, 2009; Madl et al., 2015; Eichenbaum, 2017; [A1] Mallot et al., 2020). Graph models are again introduced in Section 1.4.

However, contrary to the beliefs of early behaviorist psychology, a range of behaviors have been observed that are ill-explained by simple stimulus-response associations. During navigation, animals dynamically weigh multiple inputs and desires which may lead to different behaviors in the same situation, such as hunger, thirst, fear, or curiosity. They have explicit place memories from which optimal strategies may be derived and can suddenly change their behavior if relevant information is acquired. In this sense, the behavior is not based on the encountered stimuli alone but also depends on the internal states of the animal, such as goals, desires, beliefs, percepts, and so on. The framework that accommodates the inner state and beliefs of the animal is called the “cognitive map” (Mallot, 2024).

1.1.2 The cognitive map

The internal representation of places in a relational structure is called the “cognitive map”. As opposed to mere route knowledge, places in the cognitive map are systematically related to each other by a set of spatial transformation rules (O’Keefe and Nadel, 1978). The idea is generally attributed to Edward Tolman, who observed latent learning without reward in rodents, which could not be explained with the then prevalent behaviorist theory, i.e., reward-based stimulus-response associations:

“We believe that in the course of learning something like a field map of the environment gets established in the rat’s brain [...] and it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release.”

– Tolman, 1948, p. 192.

Whereas a route specifies a predefined sequence of actions from start to goal, a map specifies none of these. Instead, actions and behavior may directly be derived from places in the map and it can be used to navigate from any arbitrary place in the map to any other, including novel shortcuts and detours (Tolman, 1948; O’Keefe and Nadel, 1978; Gallistel, 1990; Trullier et al., 1997; Mallot, 2024; [A4] Baumann and Mallot, 2023b). In O’Keefe and Nadel (1978), this system is called “locale” and is declarative knowledge (“knowing that”), in contrast to the “knowing how” of the taxon system (Squire and Knowlton, 1995). The cognitive map is the most general and flexible navigation framework, and it is now generally accepted that many animals, but definitely humans, have access to this representation (Gallistel, 1990; Nadel, 2013; Warren, 2019; [A4] Baumann and Mallot, 2023b).

A key concept related to (metric) cognitive maps are spatial reference frames. A map may generally be defined in an *egocentric* or *allocentric* fashion. In an egocentric map, places and landmarks are defined relative to the observer’s fixed body axes, i.e., an object is to the front or behind, or left or right of the observer. For an illustration of the reference directions, see Fig. 1.1. Whenever the observer moves, the positions of all objects in the map shift instead. In an allocentric map, places and landmarks have static coordinates and relations may, for example, be expressed with cardinal directions like north or west. The position of the observer in the map is dynamically updated and relative directions need to be computed from the observer’s orientation and position in the map (Moser et al., 2017; Mallot, 2024).




Figure 1.1: Reference frames. The relative distance and direction between animal and goal depends on the choice of allocentric (the coordinate origin) or egocentric (the rat) reference frame. Here, “North” is defined in the same allocentric coordinate system as the coordinate origin but drawn on the rat for illustrative purposes. It is not a separate reference. Figure redrawn from Nyberg et al. (2022) with permission.

The orientation of the observer and map are especially important for successful navigation: If distances are over- or underestimated, the animal can simply stop or continue moving forward relative to the actual distance, but if directions are wrong, the goal will be missed entirely. Therefore, it is imperative to have access to a reference direction. This reference is in principle arbitrary and can be maintained by path integration or may be inferred from landmark positions or compass-like senses. Path integration describes the ability to track one’s own position by accumulating traveled direction and distance relative to a starting point. It can be performed even without external information (e.g., in the dark) by relying on inertial information, but is susceptible to the accumulation of errors (Eichenbaum, 2017; Epstein et al., 2017; Peer et al., 2021).

Relying on fixed external references such as the directions of landmarks, compass-like information like the bearing of the sun, the stars, or geomagnetism, is much more robust and powerful, and many species have developed the ability to use this information to travel vast distances (e.g., Wiltschko and Wiltschko, 2005; Müller and Wehner, 2007).

The necessity to track and relate spatial information to these external references makes the formation and maintenance of a map much more complex than simple route knowledge (Eichenbaum, 2017). However, once formed, the information content of a map is much higher than route knowledge, because it describes all possible routes in the map, and each new addition to the map is automatically related to all other places, albeit at a computational cost. Map knowledge is also more robust

than the strict sequences in route knowledge: The position of a place or landmark is independent of others, and it may be removed without impacting the remainder (O’Keefe and Nadel, 1978).

Of course, both types of navigation can coexist, and are likely used by many animals, including humans. In rodent experiments, it has been shown that the availability of extra-maze references, together with a large number of choices, favors the use of a map-based strategy, while the absence of cues and minimal choices favor landmark-based homing. This is supported by findings that the different types of navigation utilize different regions of the brain, and inhibiting one region will lead to the expression of the other behavior in conflicting situations (McDonald and White, 1994; McDonald et al., 2004).

1.2 The neural basis of the cognitive map

The cognitive map is built from individual experiences of the environment. However, some of its properties may also depend on the scaffolding of the map, that is, specific features of the underlying brain circuits. As such, to fully understand navigation, it is necessary to analyze the physiology involved in the cognitive map.

45 years ago, O’Keefe and Nadel (1978) reviewed the effects of hippocampal damage on the performance of rodents in a variety of navigation tasks and concluded that the hippocampus is generally necessary for navigation, but only in tasks that require a cognitive map. Animals with damaged or lesioned hippocampus instead relied on route following and recognition-triggered behavior, which have been shown to be supported by non-hippocampal areas like the dorsal striatum instead (O’Keefe and Nadel, 1978; Morris et al., 1982; McDonald and White, 1994; Eichenbaum, 2017).

Accordingly, O’Keefe and Nadel suggested the hippocampus as the locus of the cognitive map. Their claim was supported by the earlier discovery of the so-called “place cells” in the dentate gyrus and the CA1 and CA4 subfields (*cornu ammonis*, usually abbreviated) of the rat dorsal hippocampus: These cells were observed to fire whenever the rat was at a specific position or combined position and orientation within the testing environment (O’Keefe and Dostrovsky, 1971; O’Keefe, 1976). The preferred firing area of a place cell is known as its “place field”; for an example, see Fig. 1.2a. The place fields of multiple cells overlap, and the entire environment is represented by a population code of place cells (Wilson and McNaughton, 1993).

Place cells encode the animal’s location in an allocentric reference frame and

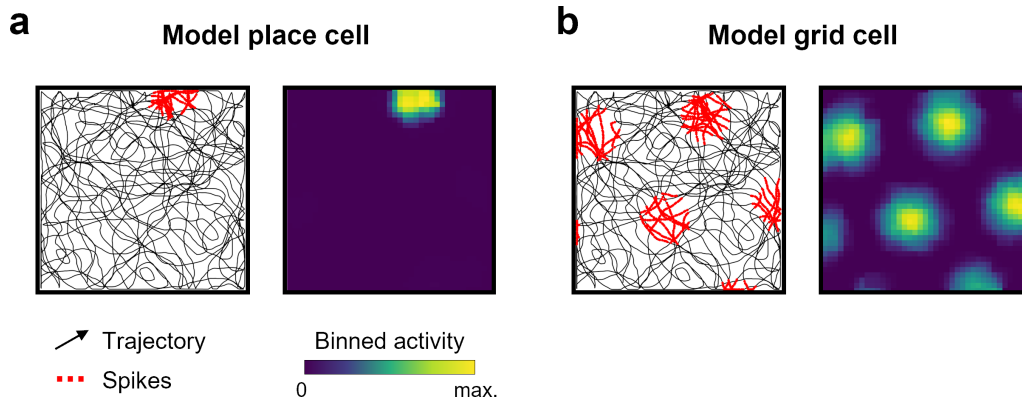


Figure 1.2: Example firing fields of (a) a model hippocampal place cell and (b) an entorhinal grid cell, based on the same trajectory. Spikes of the target cell are recorded while the animal explores the environment. The accumulated spikes per bin can then be plotted as an activity heatmap. The plots here are schematic examples generated by the model described in [A3] Baumann and Mallot (2023a) and are not based on actual recordings.

are anchored to local landmarks and geometry (O’Keefe, 1976). Their place fields rotate, shift, stretch, and deform in response to corresponding changes to local cues and the enclosure walls, but only if the animal has learned the undistorted version beforehand (Muller and Kubie, 1987; O’Keefe and Burgess, 1996). However, the cells do not require the presence of these cues and continue to fire at the same location in the dark or if the cues are removed (O’Keefe and Nadel, 1978; Moser et al., 2017).

Since their discovery in rodents, the existence of place cells has been confirmed in many other species, including (but not limited to) non-human primates (Rolls and O’Mara, 1995), bats (Yartsev and Ulanovsky, 2013), and humans (Ekstrom et al., 2003). Consistent with the cognitive map theory, fMRI recordings in humans have also found activity in the hippocampus that corresponds to the distance between locations (Morgan et al., 2011; Howard et al., 2014; Deuker et al., 2016; Epstein et al., 2017). Note that it is generally assumed that the hippocampal formation serves the same broader function in all mammalian species and findings in one species can be transferred to another; however, the question whether the same underlying navigational systems are available to humans and other animals, is still unresolved (Epstein et al., 2017).

Following the discovery of place cells, a number of other cell types supporting the cognitive map hypothesis have been found in the hippocampal formation, particularly in the entorhinal cortex, presubiculum, subiculum, and dorsal hippocampus (Moser et al., 2017; Nyberg et al., 2022). Key among these were the compass-like

head direction cells which are tuned to allocentric directions (Taube et al., 1990), border or boundary vector cells which signal the presence and distance of nearby walls (Solstad et al., 2008; Savelli et al., 2008; Lever et al., 2009), speed cells tuned to running speed (Kropff et al., 2015), discrete object-encoding cells (Deshmukh and Knierim, 2011), and metric distance-encoding grid cells (Fyhn et al., 2004; Hafting et al., 2005; Doeller et al., 2010). Among these, many cells are also conjunctive, expressing multiple types of information at the same time (Solstad et al., 2008; Kropff et al., 2015; Hardcastle et al., 2017).

Especially the grid cells received much attention due to their striking resemblance to a coordinate system. Located in the (medial) entorhinal cortex, these cells fire periodically in a way that tiles the environment into a continuous 2D hexagonal grid with consistent spacing and orientation (Fig. 1.2b), and multiple grid cells are organized in modules of similar firing properties. Like place cells, grid cells always fire at the same positions, regardless of orientation or running speed, implying that grid cells have continuous access to information about traveled distance and direction. Consequently, they are considered a mechanism for path integration (McNaughton et al., 2006; Peer et al., 2021).

Like place fields, the firing fields of grid cells also stretch, deform, and rotate if the learned environment is changed accordingly (Krupic et al., 2015), suggesting that the activity may be linked to external cues in a way that can override blind path integration (Moser et al., 2017). Head direction and border cells have been suggested as candidates for relating the grid cells to the fixed environments (McNaughton et al., 2006; Epstein et al., 2017; Peer et al., 2021).

In addition, cell activity is highly context-dependent: When the rat moves or is moved from one enclosure to another, place fields unpredictably change to a new, uncorrelated pattern (Lever et al., 2002; Leutgeb et al., 2005; Leutgeb et al., 2007; Colgin et al., 2008; Julian et al., 2018). If the animal later returns to the first enclosure, the original pattern can be observed again, and the process is repeatable. The place cells express two orthogonal place codes, one for each context. This phenomenon is known as (global) remapping. It can also be observed in passive animals when cues in the environment are changed to resemble another environment (Bostock et al., 1991; Lever et al., 2002; Leutgeb et al., 2005), and in dependence of the current task, goal, or recent history (Latuske et al., 2018; Keinath et al., 2020). The addition or manipulation of barriers within an environment does not cause global place cell remapping (Muller and Kubie, 1987; Duvelle et al., 2021; Widloski and Foster, 2022), although single place fields may be affected.

Remapping also occurs in grid cells, although it behaves slightly different: When the context changes, the firing peaks of different grid modules have been observed to shift to a new, uncorrelated pattern while scale and relative orientation remain unchanged (Fyhn et al., 2007; Marozzi et al., 2015). Due to functional and anatomical projections from the entorhinal cortex to the hippocampus, grid cell realignment has previously been suggested as a possible candidate for causing place cell remapping, but not exclusively (Monaco and Abbott, 2011; Bush et al., 2014; Moser et al., 2017).

Taken together, the findings suggest that the cognitive map is instantiated by a network of entorhinal and hippocampal neurons which encode position, distance, and orientation with sufficient accuracy and flexibility to enable a dynamic representation of the animal’s location (Epstein et al., 2017; Moser et al., 2017; Nyberg et al., 2022). In addition, the hippocampal formation creates unique representations for different tasks, regions, or contexts which are specific to that cognitive structure (Eichenbaum, 2017; Latuske et al., 2018; Widloski and Foster, 2022).

1.3 Wayfinding and regions

To guide navigation, the cognitive map must, at the very least, enable the animal to derive the current direction and distance to its goal. In its simplest form, this information may correspond to a Euclidean straight line, but many environments and behaviors require more circuitous paths and may allow for detours and shortcuts (Nyberg et al., 2022).

While the exact mechanisms for wayfinding remain unknown, experimental work has made a variety of discoveries that provide insight into how the hippocampal formation might be involved (Nyberg et al., 2022). Evidence for distance coding corresponding to both Euclidean straight-line and path distance has been found in humans (Howard et al., 2014; Brunec et al., 2017) and bats (Sarel et al., 2017); the latter study also found evidence for goal-direction-encoding cells in bat CA1. Goal-directed fMRI correlates were also measured in the human entorhinal and subicular regions (Chadwick et al., 2015).

Recent models propose that the grid cells may be involved in obtaining this information. For example, by comparing the shift in grid cell population codes at the current and goal locations, a goal-directed vector could be derived (Bush et al., 2015; Banino et al., 2018). Alternatively, the grid cells could directly probe different directions, virtually moving the activity ahead of the animal until the goal is found

(Erdem and Hasselmo, 2012). If the cognitive map contains information about the connectedness of locations, such as the adjacency of place fields, a gradient could be spread throughout the map. This gradient could then guide the animal towards the goal and even around obstacles (Erdem and Hasselmo, 2012; Ponulak and Hopfield, 2013; [A2] Brucklacher et al., 2021), similar to breadth-first search in graph theory.

An additional factor for wayfinding, especially over large distances, is the subdivision of space into separate regions or hubs that hierarchically structure large spaces to simplify search and to support planning. Specifically, the term “regionalization” is used in this context to refer to the observation that humans cluster locations and landmarks hierarchically on the basis of spatial and non-spatial attributes like neighborhood, shared boundaries, and semantics (Noack et al., 2017; Mallot, 2024).

Overall, navigation is influenced by the shape and position of superordinate regions, navigational bottlenecks like gateways or boundaries between regions, and nested structures like rooms within buildings (Mallot, 2024). For example, directions and distances between places in different regions are biased towards the center of gravity of their superordinate regions from memory (Stevens and Coupe, 1978), and distance estimates and response time increase over region boundaries (McNamara, 1986; Peer and Epstein, 2021). Regional knowledge is quickly acquired, and subjects preferably perform routes that cross the fewest region boundaries or reach the goal region as fast as possible (Wiener and Mallot, 2003). The behavior may be explained by a “fine-to-coarse” planning strategy, where more distant places are represented at a lower resolution, and only the current region is fully represented; details about the later path are worked out as the corresponding regions are entered (Wiener and Mallot, 2003). In [A3] Baumann and Mallot (2023a) we propose that this mechanism may be realized by global remapping in different contexts.

Environmental hierarchy and regionalization are less explored in non-human animals, but studies of rodent behavior in hierarchically structured mazes indicate that the animals also act efficiently in accordance with a hierarchical task structure (Roberts, 1979; Fountain and Rowan, 1995).

1.4 Structure and content of the cognitive map

The cognitive map is a theoretical construct to explain spatial processing in the brain, particularly in the hippocampal formation. It assumes that the animal stores information about the environment in a simplified form and that this information

corresponds to the environment to at least some degree. Due to preconceptions about maps, one might be led to believe that we carry around an actual two-dimensional map in our heads, and introspection might not necessarily disagree. However, this is not the intent behind the term. The cognitive map is not meant as a picture which looks like the content it represents, but rather as an information structure from which map-like images can be reconstructed and navigational behavior generated. Like the intuitive image one would have of a map, this structure has classically been assumed to be Euclidean metric; that is, places and landmarks are embedded in Euclidean space (O’Keefe and Nadel, 1978; Gallistel, 1990; McNaughton et al., 2006; Nadel, 2013). Euclidean space is a continuous metric space defined by (usually two or three) coordinate axes, and relationships between points can be expressed in terms of vectors, distances, and angles (Peer et al., 2021).

The creation of a Euclidean metric map requires the ability to determine the absolute position and orientation of one’s own body and all encountered landmarks relative to the coordinate origin. This process can happen iteratively and is the central topic of *simultaneous localization and mapping* (SLAM) in mobile robotics (e.g., Durrant-Whyte and Bailey, 2006).

Once properly formed, a Euclidean metric map is a powerful computational tool for navigation, because every position in the map is automatically related to every other position (Fig. 1.3). Trajectories between arbitrary locations can directly be inferred from the map and do not need to be learned individually. This includes the ability to navigate to novel places and to find shortcuts or generate detours and allows the structure to store an immense amount of data ((Nadel, 2013); [A4] Baumann and Mallot, 2023b).

However, it turns out that human performance in navigational tasks is often much worse than what a single coherent Euclidean map would predict. Direction and distance estimates and shortcuts show systematic errors that go against geometric postulates like the triangle inequality, which a Euclidean map should in principle adhere to (Tversky, 1992; Foo et al., 2005; Warren et al., 2017; Meilinger et al., 2018; Warren, 2019). The findings imply that space is encoded much less precisely than what would be expected from a coordinate-based metric map.

Therefore, a less restrictive alternative is often proposed: Spatial knowledge may take the form of a non-metric graph (also often called *topological* graph because it reflects a coarser topology rather than an exact map), which consists of nodes describing places or possible states and edges describing adjacency, paths, actions, or transition probabilities (Fig. 1.3). Navigation works similar to route knowledge

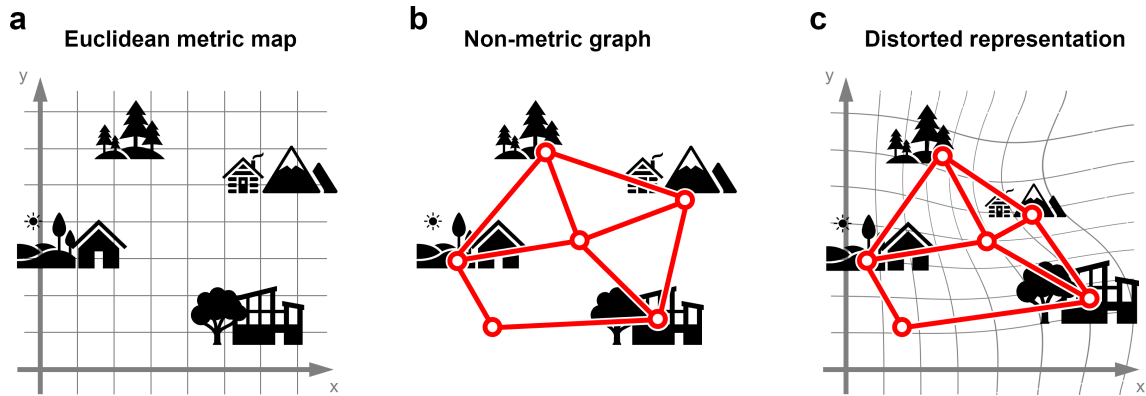


Figure 1.3: Maps and graphs. **(a)** Euclidean metric map with ground truth positions. All places are represented by specific unique coordinates and vectors to arbitrary positions can be calculated. **(b)** Non-metric graph. In the graph, only specific places and the relations between them are encoded, but no other positions. Additional connections and the location of other places cannot be directly inferred from the graph, but it may be possible to guess them; for example, if the graph is labeled with distance and direction information, direct shortcuts between distant nodes can be obtained from vector addition (Warren, 2019). **(c)** There is no guarantee that the cognitive map matches the ground truth. Instead, it may be distorted, e.g. due to measurement error. Both models can account for such a distortion, but in the metric map, conflicts might eventually arise between local position information and the stored representation.

in that each trajectory is a unique state-action sequence (Kuipers, 1978; Kuipers, 2000; Mallot and Basten, 2009; Stachenfeld et al., 2017). However, unlike simpler stimulus-response-based behavior, a graph can be much more complex; states may have more than one associated action and novel sequences can be derived from the graph at any point.

The discrete states make the graph model much less robust than a metric map: If place recognition fails along the way, the route cannot be continued and the animal needs to return to a place it recognizes and start anew. Robustness of navigation thus depends foremost on the invariance of place recognition, which may improve with larger catchment areas (like place fields) or population-coded places and routes ([A1] Mallot et al., 2020). To allow for detours and shortcuts, the graph may be labeled with local vector information like the distance and direction to neighboring places. By adding the vector information along a path in the graph, the direct connection between distant points can be inferred. Compared to the metric map, the individual labels are still independent from each other and need not be consistent or adhere to metric postulates. Therefore, the so labeled graphs may be better at explaining

large deviations from the Euclidean metric observed in human navigation (Warren et al., 2017; Warren, 2019).

What the graph model gains in explanatory power and simplistic elegance, it in turn loses in computational and storage efficiency. Each addition to the graph requires the explicit encoding of at least one, if not multiple links to its neighbors, resulting in an enormous number of pairwise links in a well-known environment (Nadel, 2013).

The two model classes, metric map and non-metric graph, are often considered in opposition (Chrastil and Warren, 2014; Warren, 2019; Peer et al., 2021). However, finding definite evidence for either model is not easy, because they tend to be correlated in most regular environments, and both path-based and metric information can be derived from either model. It has also been argued that both models may be available to the brain in a complementary fashion, for example, with Euclidean metric immediate environments connected by a hierarchical region graph (Kuipers, 1982; Meilinger, 2008). Still, it remains open to debate whether spatial knowledge is Euclidean metric, graph-based, or a combination of both (Peer et al., 2021).

1.5 Aim of this thesis

The interaction with space is one of the fundamental abilities that most if not all living beings developed over their evolutionary history. The ability to move and navigate is more specialized but has still resulted in an immense breadth of behaviors and adaptations with varying complexity. Research on the neural representation of space in mammals has uncovered a highly developed system of specialized cell types situated in the hippocampus, entorhinal cortex, prefrontal cortex, and more. The characteristic firing patterns of these cells have prompted investigators to search for the neural mechanisms underlying both single cell activity and the overall network with behavioral studies, neuroimaging, and theoretical modeling.

The considerable difference between measurable single-cell activities in restricted environments and the resulting complex behaviors makes the cognitive neuroscience of space a highly conceptual field. In fact, the cognitive map, the theoretical neural representation of space, is an interpretation of what investigators *expect* the combined system to do but not a measurable quantity (so far). Accordingly, theoretical modeling is needed to combine the findings from electrophysiology and behavioral studies to further our understanding of spatial cognition and the underlying systems.

In the following, I present four theoretical computational models for the cognitive map, which contribute to different facets like wayfinding, hierarchy, and the structural and functional organization of the system. Roughly following the order of the introduction, the works focus on various aspects related to the population coding of positions and routes, effects of region boundaries, place and grid cell remapping, and graph versus map-based representations. Beyond the individual results, I also discuss which aspects of the models best match our current knowledge of the cognitive map, the functional implications, and ways to combine the findings into more complex models. Note that a more detailed summary of the results is given in the next chapter.

- In [A1] Mallot et al. (2020), we consider a parsimonious model for the cognitive map based on simple visual features. Sets of these features are distinct enough that positions and routes can robustly be represented by population coding. By following “votes” towards the next position along the route, a virtual agent is able to navigate a large-scale environment without further structural or behavioral constraints.
- In [A2] Brucklacher et al. (2021), we propose a neural network for hierarchical wayfinding based on superordinate region nodes. Activity is spread throughout the network to create a gradient that leads a virtual agent towards the goal. Importantly, this gradient is influenced by the superordinate region nodes, which enables the model to replicate region biases observed in human navigation experiments.

Importantly, these graph models operate under the assumption that each place in the environment corresponds to a specific unique node (or population of nodes) in the graph. Hippocampal place cells, on the other hand, remap between contexts and express multiple firing fields in larger environments (see Section 1.2).

- To account for this phenomenon, in [A3] Baumann and Mallot (2023a), we propose a model for place and grid cell remapping based on the context change at region boundaries. Whenever a region is entered, associated place cell patterns are reactivated from memory and the system changes to the region-specific place code by relying on attractor dynamics. The model matches remapping behavior and cell activity observed in multi-compartment environments and predicts a regionalized hierarchical cognitive map.

- Lastly, in [A4] Baumann and Mallot (2023b), we join the longstanding debate on the functional organization of the cognitive map: Fundamentally, the representation of space in the brain must deviate from real space due to restrictions that come with successive local egocentric measurements. These deviations are especially apparent in non-Euclidean environments and have been used to argue that human spatial knowledge is mostly non-metric and therefore best explained by a non-metric graph model. We argue for a metric alternative, a distorted Euclidean embedding, which can explain human navigational performance in a virtual reality (VR) experiment equally well.

Bibliography

- Banino, Andrea et al. (2018). “Vector-based navigation using grid-like representations in artificial agents”. *Nature* 557.7705, pp. 429–433.
- Baumann, Tristan and Hanspeter A Mallot (2023a). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239.
- Baumann, Tristan and Hanspeter A Mallot (2023b). “Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments”. Under review.
- Bostock, Elizabeth, Robert U Muller, and John L Kubie (1991). “Experience-dependent modifications of hippocampal place cell firing”. *Hippocampus* 1.2, pp. 193–205.
- Brucklacher, Matthias, Hanspeter A Mallot, and Tristan Baumann (2021). “Hierarchical Planning in Multilayered State-Action Networks.” *ESANN 2021 proceedings*.
- Brunec, Iva K et al. (2017). “Contracted time and expanded space: The impact of circumnavigation on judgements of space and time”. *Cognition* 166, pp. 425–432.
- Bush, Daniel, Caswell Barry, and Neil Burgess (2014). “What do grid cells contribute to place cell firing?” *Trends in neurosciences* 37.3, pp. 136–145.
- Bush, Daniel et al. (2015). “Using grid cells for navigation”. *Neuron* 87.3, pp. 507–520.
- Chadwick, Martin J et al. (2015). “A goal direction signal in the human entorhinal/subicular region”. *Current Biology* 25.1, pp. 87–92.
- Chrastil, Elizabeth R and William H Warren (2014). “From cognitive maps to cognitive graphs”. *PloS one* 9.11, e112544.
- Colgin, Laura Lee, Edvard I Moser, and May-Britt Moser (2008). “Understanding memory through hippocampal remapping”. *Trends in neurosciences* 31.9, pp. 469–477.
- Collett, Matthew et al. (1998). “Local and global vectors in desert ant navigation”. *Nature* 394.6690, pp. 269–272.

- Collett, Thomas S and Matthew Collett (2002). “Memory use in insect visual navigation”. *Nature Reviews Neuroscience* 3.7, pp. 542–552.
- Deshmukh, Sachin S and James J Knierim (2011). “Representation of non-spatial and spatial information in the lateral entorhinal cortex”. *Frontiers in behavioral neuroscience* 5, p. 69.
- Deuker, Lorena et al. (2016). “An event map of memory space in the hippocampus”. *elife* 5, e16534.
- Doeller, Christian F, Caswell Barry, and Neil Burgess (2010). “Evidence for grid cells in a human memory network”. *Nature* 463.7281, pp. 657–661.
- Durrant-Whyte, Hugh and Tim Bailey (2006). “Simultaneous localization and mapping: part I”. *IEEE robotics & automation magazine* 13.2, pp. 99–110.
- Duvelle, Éléonore et al. (2021). “Hippocampal place cells encode global location but not connectivity in a complex space”. *Current Biology* 31.6, pp. 1221–1233.
- Eichenbaum, Howard (2017). “The role of the hippocampus in navigation is memory”. *Journal of neurophysiology* 117.4, pp. 1785–1796.
- Ekstrom, Arne D et al. (2003). “Cellular networks underlying human spatial navigation”. *Nature* 425.6954, pp. 184–188.
- Epstein, Russell A et al. (2017). “The cognitive map in humans: spatial navigation and beyond”. *Nature neuroscience* 20.11, pp. 1504–1513.
- Erdem, Uğur M and Michael Hasselmo (2012). “A goal-directed spatial navigation model using forward trajectory planning based on grid cells”. *European Journal of Neuroscience* 35.6, pp. 916–931.
- Foo, Patrick et al. (2005). “Do humans integrate routes into a cognitive map? Map-versus landmark-based navigation of novel shortcuts.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31.2, p. 195.
- Fountain, Stephen B and James D Rowan (1995). “Coding of hierarchical versus linear pattern structure in rats and humans.” *Journal of Experimental Psychology: Animal Behavior Processes* 21.3, p. 187.
- Franz, Matthias O and Hanspeter A Mallot (2000). “Biomimetic robot navigation”. *Robotics and autonomous Systems* 30.1-2, pp. 133–153.
- Fyhn, Marianne et al. (2007). “Hippocampal remapping and grid realignment in entorhinal cortex”. *Nature* 446.7132, pp. 190–194.
- Fyhn, Marianne et al. (2004). “Spatial representation in the entorhinal cortex”. *Science* 305.5688, pp. 1258–1264.
- Gallistel, Charles R (1990). *The organization of learning*. The MIT Press.
- Hafting, Torkel et al. (2005). “Microstructure of a spatial map in the entorhinal cortex”. *Nature* 436.7052, pp. 801–806.
- Hardcastle, Kiah et al. (2017). “A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex”. *Neuron* 94.2, pp. 375–387.
- Howard, Lorelei R et al. (2014). “The hippocampus and entorhinal cortex encode the path and Euclidean distances to goals during navigation”. *Current Biology* 24.12, pp. 1331–1340.
- Julian, Joshua B et al. (2018). “The neurocognitive basis of spatial reorientation”. *Current Biology* 28.17, R1059–R1073.

- Keinath, Alexandra T et al. (2020). “DG–CA3 circuitry mediates hippocampal representations of latent information”. *Nature communications* 11.1, p. 3026.
- Kropff, Emilio et al. (2015). “Speed cells in the medial entorhinal cortex”. *Nature* 523.7561, pp. 419–424.
- Krupic, Julija et al. (2015). “Grid cell symmetry is shaped by environmental geometry”. *Nature* 518.7538, pp. 232–235.
- Kuipers, Benjamin (1978). “Modeling spatial knowledge”. *Cognitive Science* 2.2, pp. 129–153.
- Kuipers, Benjamin (1982). “The ‘map in the head’ metaphor”. *Environment and behavior* 14.2, pp. 202–220.
- Kuipers, Benjamin (2000). “The spatial semantic hierarchy”. *Artificial intelligence* 119.1-2, pp. 191–233.
- Latuske, Patrick et al. (2018). “Hippocampal remapping and its entorhinal origin”. *Frontiers in behavioral neuroscience* 11, p. 253.
- Leutgeb, Jill K et al. (2007). “Pattern separation in the dentate gyrus and CA3 of the hippocampus”. *science* 315.5814, pp. 961–966.
- Leutgeb, Stefan et al. (2005). “Independent codes for spatial and episodic memory in hippocampal neuronal ensembles”. *Science* 309.5734, pp. 619–623.
- Lever, Colin et al. (2009). “Boundary vector cells in the subiculum of the hippocampal formation”. *Journal of Neuroscience* 29.31, pp. 9771–9777.
- Lever, Colin et al. (2002). “Long-term plasticity in hippocampal place-cell representation of environmental geometry”. *Nature* 416.6876, pp. 90–94.
- Madl, Tamas et al. (2015). “Computational cognitive models of spatial memory in navigation space: A review”. *Neural Networks* 65, pp. 18–43.
- Mallot, Hanspeter A (2024). *From Geometry to Behavior: An Introduction to Spatial Cognition*. MIT Press.
- Mallot, Hanspeter A and Kai Basten (2009). “Embodied spatial cognition: Biological and artificial systems”. *Image and Vision Computing* 27.11, pp. 1658–1670.
- Mallot, Hanspeter A, Gerrit A Ecke, and Tristan Baumann (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17.
- Marozzi, Elizabeth et al. (2015). “Purely translational realignment in grid cell firing patterns following nonmetric context change”. *Cerebral Cortex* 25.11, pp. 4619–4627.
- McDonald, Robert J, Bryan D Devan, and Nancy S Hong (2004). “Multiple memory systems: the power of interactions”. *Neurobiology of learning and memory* 82.3, pp. 333–346.
- McDonald, Robert J and Norman M White (1994). “Parallel information processing in the water maze: evidence for independent memory systems involving dorsal striatum and hippocampus”. *Behavioral and neural biology* 61.3, pp. 260–270.
- McNamara, Timothy P (1986). “Mental representations of spatial relations”. *Cognitive psychology* 18.1, pp. 87–121.
- McNaughton, Bruce L et al. (2006). “Path integration and the neural basis of the ‘cognitive map’”. *Nature Reviews Neuroscience* 7.8, pp. 663–678.

- Meilinger, Tobias (2008). “The network of reference frames theory: A synthesis of graphs and cognitive maps”. *Spatial Cognition VI*. Springer, pp. 344–360.
- Meilinger, Tobias et al. (2018). “Humans construct survey estimates on the fly from a compartmentalised representation of the navigated environment”. *Spatial Cognition XI*. Springer, pp. 15–26.
- Monaco, Joseph D and Larry F Abbott (2011). “Modular realignment of entorhinal grid cell activity as a basis for hippocampal remapping”. *Journal of Neuroscience* 31.25, pp. 9414–9425.
- Morgan, Lindsay K et al. (2011). “Distances between real-world locations are represented in the human hippocampus”. *Journal of Neuroscience* 31.4, pp. 1238–1245.
- Morris, Richard GM et al. (1982). “Place navigation impaired in rats with hippocampal lesions”. *Nature* 297.5868, pp. 681–683.
- Moser, Edvard I, May-Britt Moser, and Bruce L McNaughton (2017). “Spatial representation in the hippocampal formation: a history”. *Nature neuroscience* 20.11, pp. 1448–1464.
- Muller, Robert U and John L Kubie (1987). “The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells”. *Journal of Neuroscience* 7.7, pp. 1951–1968.
- Müller, Martin and Rüdiger Wehner (2007). “Wind and sky as compass cues in desert ant navigation”. *Naturwissenschaften* 94, pp. 589–594.
- Nadel, Lynn (2013). “Cognitive maps”. *Handbook of spatial cognition*, pp. 155–171.
- Noack, Hannes et al. (2017). “Sleep enhances knowledge of routes and regions in spatial environments”. *Learning & Memory* 24.3, pp. 140–144.
- Nyberg, Nils et al. (2022). “Spatial goal coding in the hippocampal formation”. *Neuron*.
- O’Keefe, John (1976). “Place units in the hippocampus of the freely moving rat”. *Experimental neurology* 51.1, pp. 78–109.
- O’Keefe, John and Neil Burgess (1996). “Geometric determinants of the place fields of hippocampal neurons”. *Nature* 381.6581, pp. 425–428.
- O’Keefe, John and Jonathan Dostrovsky (1971). “The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat.” *Brain research*.
- O’Keefe, John and Lynn Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- Peer, Michael and Russell A Epstein (2021). “The human brain uses spatial schemas to represent segmented environments”. *Current Biology* 31.21, pp. 4677–4688.
- Peer, Michael et al. (2021). “Structuring knowledge with cognitive maps and cognitive graphs”. *Trends in cognitive sciences* 25.1, pp. 37–54.
- Ponulak, Filip and John J Hopfield (2013). “Rapid, parallel path planning by propagating wavefronts of spiking neural activity”. *Frontiers in computational neuroscience* 7, p. 98.
- Roberts, William A (1979). “Spatial memory in the rat on a hierarchical maze”. *Learning and Motivation* 10.2, pp. 117–140.

- Rolls, Edmund T and Shane M O'Mara (1995). "View-responsive neurons in the primate hippocampal complex". *Hippocampus* 5.5, pp. 409–424.
- Sarel, Ayelet et al. (2017). "Vectorial representation of spatial goals in the hippocampus of bats". *Science* 355.6321, pp. 176–180.
- Savelli, Francesco, D Yoganarasimha, and James J Knierim (2008). "Influence of boundary removal on the spatial representations of the medial entorhinal cortex". *Hippocampus* 18.12, pp. 1270–1282.
- Solstad, Trygve et al. (2008). "Representation of geometric borders in the entorhinal cortex". *Science* 322.5909, pp. 1865–1868.
- Squire, Larry R. and Barbara J. Knowlton (1995). "Memory, hippocampus, and brain systems". *The cognitive neurosciences*. Ed. by Michael S. Gazzaniga. Cambridge, MA: The MIT Press. Chap. 53, pp. 825–837.
- Stachenfeld, Kimberly L, Matthew M Botvinick, and Samuel J Gershman (2017). "The hippocampus as a predictive map". *Nature neuroscience* 20.11, pp. 1643–1653.
- Stevens, Albert and Patty Coupe (1978). "Distortions in judged spatial relations". *Cognitive psychology* 10.4, pp. 422–437.
- Taube, Jeffrey S, Robert U Muller, and James B Ranck (1990). "Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis". *Journal of Neuroscience* 10.2, pp. 420–435.
- Tolman, Edward C (1948). "Cognitive maps in rats and men." *Psychological review* 55.4, p. 189.
- Tolman, Edward C (1932). *Purposive Behavior in Animals and Men*. New York: The Century Co.
- Trullier, Olivier et al. (1997). "Biologically based artificial navigation systems: Review and prospects". *Progress in neurobiology* 51.5, pp. 483–544.
- Tversky, Barbara (1992). "Distortions in cognitive maps". *Geoforum* 23.2, pp. 131–138.
- Warren, William H (2019). "Non-euclidean navigation". *Journal of Experimental Biology* 222.Suppl_1, jeb187971.
- Warren, William H et al. (2017). "Wormholes in virtual space: From cognitive maps to cognitive graphs". *Cognition* 166, pp. 152–163.
- Widloski, John and David J Foster (2022). "Flexible rerouting of hippocampal replay sequences around changing barriers in the absence of global place field remapping". *Neuron* 110.9, pp. 1547–1558.
- Wiener, Jan M and Hanspeter A Mallot (2003). "'Fine-to-coarse' route planning and navigation in regionalized environments". *Spatial cognition and computation* 3.4, pp. 331–358.
- Wilson, Matthew A and Bruce L McNaughton (1993). "Dynamics of the hippocampal ensemble code for space". *Science* 261.5124, pp. 1055–1058.
- Wiltschko, Wolfgang and Roswitha Wiltschko (2005). "Magnetic orientation and magnetoreception in birds and other animals". *Journal of comparative physiology A* 191, pp. 675–693.

Yartsev, Michael M. and Nachum Ulanovsky (2013). “Representation of Three-Dimensional Space in the Hippocampus of Flying Bats”. *Science* 340.6130, pp. 367–372.

Chapter 2

Results

This chapter gives a brief summary of the findings of each publication. The full results including relevant figures and interpretations can be found in the appendix. Further details are also given in Chapter 3: Discussion.

2.1 [A1] Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations (Mallot et al., 2020)

We presented a parsimonious view-based graph model for spatial navigation, which relies on simple image features for place recognition and series of recognition-triggered responses for wayfinding but does not require explicit metric or coordinate-based information. Invariant encoding of places and routes is achieved via *dual population coding*: Each position is defined by a set of simple but descriptive image features (SURF, Bay et al., 2008) with overlapping catchment areas, connected to neighboring features as nodes in a large graph. For wayfinding, multiple parallel routes are calculated between the local features and the features at the goal; the goal can then be reached by following the set of routes by combining them in a voting scheme.

The algorithm was tested in a virtual street network by repeatedly guiding a virtual agent from start to goal in real-time. The algorithm performed well; routes were only slightly longer than optimal straight-line trajectories and performance was not hindered by aliases in the feature graph, which led to occasional impossible connections. As a consequence of the population coding of routes, the resulting trajectories were not paths in the graph but a metric average of multiple simultaneous routes.

The algorithm also sometimes guided the agent along alternative roads over multiple repeats of the same trial if the routes were of similar length.

We concluded that navigation based on stimulus-response association can successfully be combined with spatial population coding and graph knowledge. The method requires only simple image processing and avoids more complex problems like the selection of optimal views or learning a consistent metric map. The trade-off is that the redundancy encoded in the graph comes with higher memory and processing requirements.

2.2 [A2] Hierarchical Planning in Multilayered State-Action Networks (Brucklacher et al., 2021)

Complex environments can be subdivided into multiple meaningful regions which simplify path planning and have a measurable effect on human navigation (e.g., Wiener and Mallot, 2003; Hochmair et al., 2008). To simulate the effects of regionalization on wayfinding, we created a series of spiking neural networks representing a low-level place graph connected to a simpler high-level region graph. We then investigated the influence of this hierarchical structure on planning speed and the resulting route choice in comparison to non-hierarchical single-level models. Directed routes were obtained by initiating a wave of spiking activity at the goal location which propagated through the network, resulting in a goal-directed gradient.

We found that the hierarchical structure drastically sped up the wayfinding process and that the algorithm chose biased routes in line with human navigational behavior: Rather than selecting a straight-line trajectory like the non-hierarchical single-level model, the multi-level model found routes that preferentially entered the goal region as quickly as possible, which also has been observed in human navigation (Wiener and Mallot, 2003). The additional region nodes only slightly increase the size of the graph compared to the non-hierarchical version but lead to better results. We concluded that the model may be a more biological plausible cognitive map than a model without region representation.

2.3 [A3] Gateway identity and spatial remapping in a combined grid and place cell attractor (Baumann and Mallot, 2023a)

Place and grid cells have been observed to remap to uncorrelated but specific firing patterns in different compartments of the experimental maze. The cell behavior cannot be explained from path integration, i.e., grid cell activity, and local sensory input alone. We proposed a model of place and grid cells in which remapping is triggered by context changes, in this case when moving from one compartment to another. In the model, each entrance to a region (the “gateway”) is associated with a stored place cell population code, which is reinstated whenever the region is entered. Attractor dynamics then also cause the grid cells to remap, and local path integration may resume.

We let a virtual agent explore multi-compartment environments and compared the simulated cell behavior to results from rodent research. Importantly, gateways in the model are linked to local cues and can therefore be confused if two regions look similar enough. This will lead to expression of the same firing fields in visually identical rooms, in line with measurements from rodent place and grid cells in multi-compartment environments (e.g., Fuhs et al., 2005; Carpenter et al., 2015; Grieves et al., 2016). By making the gateways direction-selective, the model was also able to replicate directionally selective firing fields which have been commonly observed in narrow corridors.

Based on the results, we argued that remapping is a fundamental property of cognitive maps, indicating that the map is divided into multiple regions. That is, the world is at least divided into similar and dissimilar parts, and only the current context is represented in full; however, because animals can clearly execute trajectories into other regions, some sort of higher-level representation must exist. The theory is possibly supported by the existence of region-encoding cells in the perirhinal or medial prefrontal cortices (cf. Hyman et al., 2012; Bos et al., 2017).

We also proposed that the model dynamics might be realized by hypothetical “gate cells”, which would learn a place cell pattern via neuronal plasticity and reactivate that pattern whenever the relevant stimuli occurred. A candidate for such a cell type may exist in the retrosplenial cortex (Jacob et al., 2017), which has been suggested to be involved in the recognition of context changes (Pothuizen et al., 2008; Wesierska et al., 2009).

2.4 [A4] Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments (Baumann and Mallot, 2023b)

Literature on the representational structure of the cognitive map, i.e., what form spatial information might take in the brain, generally distinguishes between Euclidean metric and non-metric graph-based representations. This long-standing debate in part arose due to the observation that human behavior in many cases strongly deviates from what a Euclidean metric map would predict. Non-metric models can often better explain the deviations but sacrifice useful metric properties in turn.

We focused on a specific example, a navigation experiment by Warren et al. (2017), in which human participants explored a non-Euclidean environment in VR. Participant estimates were strongly biased towards the non-Euclidean distortions, which led the authors to conclude that similarly, only a non-Euclidean graph-based cognitive map could explain the findings. We considered another alternative, a systematically distorted Euclidean metric map, obtained by embedding the non-metric graph into 2D coordinates.

We compared the models on the same dataset as Warren et al. (2017) and found that the distorted map predicted the behavior equally well, even though a Euclidean metric map can in principle not represent the non-Euclidean environment without errors. Our results at least partially disagree with the original study, and we argue that the embedded graph may be a better model for the cognitive map, because it utilizes the additional information conveyed by the metric constraints rather than discarding it. Our findings also support the possibility of combined models with both metric and non-metric properties (e.g., Peer et al., 2021).

Bibliography

- Baumann, Tristan and Hanspeter A Mallot (2023a). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239.
- Baumann, Tristan and Hanspeter A Mallot (2023b). “Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments”. Under review.
- Bay, Herbert et al. (2008). “Speeded-up robust features (SURF)”. *Computer vision and image understanding* 110.3, pp. 346–359.

- Bos, Jeroen J et al. (2017). “Perirhinal firing patterns are sustained across large spatial segments of the task environment”. *Nature communications* 8.1, p. 15602.
- Brucklacher, Matthias, Hanspeter A Mallot, and Tristan Baumann (2021). “Hierarchical Planning in Multilayered State-Action Networks.” *ESANN 2021 proceedings*.
- Carpenter, Francis et al. (2015). “Grid cells form a global representation of connected environments”. *Current Biology* 25.9, pp. 1176–1182.
- Fuhs, Mark C et al. (2005). “Influence of path integration versus environmental orientation on place cell remapping between visually identical environments”. *Journal of neurophysiology* 94.4, pp. 2603–2616.
- Grieves, Roddy M et al. (2016). “Place field repetition and spatial learning in a multicompartment environment”. *Hippocampus* 26.1, pp. 118–134.
- Hochmair, Hartwig H, Simon J Büchner, and Christoph Hölscher (2008). “Impact of regionalization and detour on ad-hoc path choice”. *Spatial Cognition & Computation* 8.3, pp. 167–192.
- Hyman, James M et al. (2012). “Contextual encoding by ensembles of medial prefrontal cortex neurons”. *Proceedings of the National Academy of Sciences* 109.13, pp. 5086–5091.
- Jacob, Pierre-Yves et al. (2017). “An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex”. *Nature neuroscience* 20.2, pp. 173–175.
- Mallot, Hanspeter A, Gerrit A Ecke, and Tristan Baumann (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17.
- Peer, Michael et al. (2021). “Structuring knowledge with cognitive maps and cognitive graphs”. *Trends in cognitive sciences* 25.1, pp. 37–54.
- Pothuizen, Helen HJ, John P Aggleton, and Seralynne D Vann (2008). “Do rats with retrosplenial cortex lesions lack direction?” *European Journal of Neuroscience* 28.12, pp. 2486–2498.
- Warren, William H et al. (2017). “Wormholes in virtual space: From cognitive maps to cognitive graphs”. *Cognition* 166, pp. 152–163.
- Wesierska, Malgorzata, Iwona Adamska, and Monika Malinowska (2009). “Retrosplenial cortex lesion affected segregation of spatial information in place avoidance task in the rat”. *Neurobiology of learning and memory* 91.1, pp. 41–49.
- Wiener, Jan M and Hanspeter A Mallot (2003). “‘Fine-to-coarse’ route planning and navigation in regionalized environments”. *Spatial cognition and computation* 3.4, pp. 331–358.

Chapter 3

Discussion

This thesis considers specificities of spatial cognition and the cognitive map through a lens of computational modeling. The findings of each included publication are already discussed at length in their respective sections in the appendix; therefore, in the following overall discussion, the focus lies not so much on the individual results, but rather on the additional knowledge that can be gained from the combination of the individual findings, how the different models may complement each other, or where they might disagree. I specifically discuss how the models and the cognitive map in general may be used for wayfinding, navigation, and the hierarchical segmentation of complex large-scale environments and conceptual spaces. Many of the points offer potential future avenues of research for theoretical computational modeling or behavioral and electrophysiological neuroscience.

Ever since the discovery of place cells and related parahippocampal cell types, the field of spatial cognition has advanced at a tremendous rate. The cells have become one of the most powerful tools in understanding the processes behind spatial mapping and the cognitive system and navigation may become one of the first brain functions that are fully unraveled. With this in mind, the models discussed here successfully describe a variety of phenomena and findings from long-distance navigation, regionalization, spatial hierarchy, and remapping – but there are also many particularities that they do not cover. Humans and rodents do not have a perfect 3D reconstruction of the environment in their head, but there is also more structure to their knowledge than simple images or rote memorization. As always with extremes, the true cognitive map likely lies somewhere in between.

3.1 Population coding of routes and wayfinding

To model navigation, increasingly complex systems can be created that account for mechanisms like invariant landmark of place recognition, strategic selection of way-points and reference frames, the embedding of disparate places in a coherent metric map, or hierarchical structures. However, following *Occam's razor*, models always come with a trade-off between power and simplicity, i.e., how many functions or how much data they can explain with the fewest parameters possible. It is therefore of interest to consider the more parsimonious models for spatial navigation.

In state-action graphs describing navigation (e.g., Kuipers, 1978; Franz et al., 1998; Kuipers, 2000; Warren et al., 2017; Warren, 2019), a place in the environment usually corresponds to a single node in the graph, such that a unique state-action scheme can control each navigational step. Without further specification of places, this type of encoding is very fragile: if a place is not recognized along the way, navigation may fail, possibly catastrophically. On the other hand, improving place recognition or, for example, selecting which places or landmarks make good choices for nodes in the graph, leads to much more complicated models.

In [A1] Mallot et al. (2020), we argued that, from an evolutionary point of view, robust wayfinding should be possible even with only rudimentary place recognition and distributed place representations. Accordingly, we proposed a (in some respects) minimal view graph model (cf. Franz et al., 1998), a type of state-action network where views are associated with movement instructions towards the next place along a route.

The model differs from the classical state-action graphs by using population codes for both place and route definitions: At any point in time, sets of descriptive low-resolution image patches like corners or edges (SURF, Bay et al., 2008) are extracted from the scene. The features are intended as a rudimentary place recognition system that eschews the need for more complex image segmentation or landmark or object recognition methods. While the features are relatively unique, they are prone to aliasing, i.e., there may be other places in the environment which contain the same features, especially in large graphs.

Each SURF feature has a specific area from which it can be detected, its *catchment area*. A specific position in the environment was thus defined by the population code of features with overlapping catchment areas, and the entire environment was mapped by a graph of these features. Invariant place recognition was then gained by defining a place similarity threshold function which allowed for the graded recog-

nition of places from parts of the corresponding population code. The system was intended to function similar to how place cells may encode each position with a population code of overlapping place fields (Wilson and McNaughton, 1993), but directly based on visual input.

For wayfinding, multiple non-overlapping routes were obtained from breadth-first search (Dijkstra’s algorithm) throughout the graph, and movement instructions were assigned to each feature along the routes. As a consequence of the overlapping place code, action selection was also based on a set of instructions at each point along the way, and the resulting movement was decided by a voting scheme. Because each feature had a specific catchment area and many features were part of a route, this population coding system led to route representation with a certain breadth, allowing for trajectories which are not part of the graph itself.

In testing, the algorithm efficiently and robustly guided an agent around a large virtual environment, finding routes that were only 10% to 20% longer than optimal straight-line trajectories. The method successfully combined spatial population coding with graph-based representation and guidance. However, what the model gained in simplicity and robustness, it lost in computational efficiency: Over time, the indiscriminate addition of newly detected features led to massive graphs which were costly to store and access.

With [A2] Brucklacher et al. (2021), we proposed an alternative navigation scheme based on a hierarchical state-action network in multi-region environments. The states at the lower level of the hierarchy were modeled as neurons that each represented a specific place in the environment, comparable to place cells (but without place fields or population codes). Connections between states were bidirectional and represented vectors from one place to the other. For wayfinding, the vertices corresponding to the goal were activated and propagated a wave of neuronal activity throughout the network, which strengthened goal-directed connections (cf. Ponulak and Hopfield, 2013). The wave passed through the entire map and formed a gradient vector field pointing back to its origin, which could be followed to reach the goal.

Gradient-based navigation is comparable to the population coding of routes with a certain breadth in the sense that deviating from the center of the route will not necessarily result in failure or a restart of the wayfinding process; at least to a certain extent, the correct path remains available. However, as opposed to a set of routes, the gradient vector field extends over the entire environment and can always lead the agent towards the goal barring local minima.

In addition, the graph was outfitted with higher-level region nodes that were con-

nected to the lower level in a way that allowed the activity to spread to distant regions faster. The so defined regions also distorted the vector field in a way that allowed the model to replicate navigation behavior of humans in similar environments. For example, when choosing among equally long paths to the goal, subjects have been observed to prefer routes that crossed less region borders and routes that entered the goal region as quickly as possible (Wiener and Mallot, 2003; Hochmair et al., 2008).

Both models describe plausible methods to derive goal-directed routes on the level of place cells. Both models are also robust, because the solutions they give allow for deviations from the direct route to the goal without immediate failure. Aside from these commonalities, however, the models function very differently and make different predictions on how a biological implementation, i.e., the cognitive map, may be involved in wayfinding.

3.1.1 The role of preplays and graph search

Biological evidence for the encoding of future trajectories has been found in a place cell firing behavior known as “(p)replay”: during rest or pauses in navigation, for example, at decision points in mazes, temporally compressed sequences of place cell activity can be measured that depict paths ahead of the animal. These trajectories can be completely novel and pass through unexplored portions of the environment, suggesting that they are a neural mechanism for prospective planning and not just route memory (Johnson and Redish, 2007; Dragoi and Tonegawa, 2011; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015).

Preplays differ from the sets of paths in [A1] Mallot et al. (2020) in some respects. While the sequence depicted by a preplay may also have a certain breadth due to the spatial extent of the involved place fields, there is no indication that this is a result of multiple parallel paths like in the wayfinding algorithms used in our model. Furthermore, preplay events do not always occur (or are at least not always measured) and do not necessarily correspond to the animal’s future path. The role of the sequences in wayfinding remains inconclusive (Nyberg et al., 2022).

Research suggests that hippocampal preplay may, for example, only aid planning under specific circumstances, such as when the task places a high cognitive or mnemonic demand on the animal, but not if the task is simple (Pfeiffer and Foster, 2013; Xu et al., 2019). The sequences may also depict routes to places that are to be avoided or to alternative goals. Accordingly, hippocampal preplays may be

related to memory consolidation by actively maintaining a representation of less-visited places (Nyberg et al., 2022). This is similar to a suggested role of the medial prefrontal cortex (mPFC): In the region, temporally compressed activity sequences occur simultaneous to hippocampal preplays (Tang et al., 2021; Nyberg et al., 2022). The difference seems to be that while the hippocampus alternates between different possible routes during deliberation, the mPFC maintains and predicts the actual upcoming choice. In this sense, both regions are likely involved in a cooperative interaction for wayfinding, memory consolidation and decision making (Tang et al., 2021).

In any case, preplays seem to be deliberate virtual paths through the mental representation of space and, as such, place a strong demand on models of the cognitive map that intend to be biologically plausible. This is exemplified in a critical difference between graph-theoretical models (including [A1] Mallot et al., 2020 and [A2] (Brucklacher et al., 2021)) and preplays in biology: There seems to be no search (in the graph-theoretical sense) prior to the preplay, or it at least cannot be detected in the place cell activity. That is, there is no expanding wave or functionally similar breadth-first search throughout the biological network (Pfeiffer, 2020). Rather, the hippocampus appears to be already aware of the upcoming route and goal choices before the actual preplay occurs. Preplays are therefore either not indicative of a path planning process at all, or another mechanism must exist. Note that the oft-cited “sweep” ahead of the animal in Johnson and Redish (2007) describes the linear activation sequence of place cells in the replay and is not meant to imply a searching behavior.

As already mentioned in the introduction, the involvement of grid cells may be a possible solution for this conundrum. Recent evidence suggests that grid cells are directly involved in place cell replays: During the events, grid and place cells show increased coordination (Ólafsdóttir et al., 2016; Yamamoto and Tonegawa, 2017) and inhibiting direct input from medial entorhinal cortex to CA1 reduces the length of preplays in awake animals, suggesting that the input is necessary for proper route simulation (Yamamoto and Tonegawa, 2017; Nyberg et al., 2022). In agreement, fMRI recordings from human entorhinal cortex show increased grid cell-like activity during imagined navigation (Bellmund et al., 2016).

If the position encoded by the grid cells could, for example, be virtually moved while the animal remains in place, it may be possible for the brain to directly calculate goal-directed vectors from the difference in population codes for the current location and the goal (Bush et al., 2015; Banino et al., 2018). If the goal is too far

away or unknown, the virtual movement could instead be used to probe different directions originating from the animal’s current position; the so generated activity could then be evaluated to find positions closer to the goal (Erdem and Hasselmo, 2012). In line with this hypothesis, grid cells in the superficial layers of the medial entorhinal cortex can also replay trajectories without accompanying place cell activity (O’Neill et al., 2017), but whether this means that the areas and cell types play a separate role in wayfinding remains to be determined (Nyberg et al., 2022).

In [A3] Baumann and Mallot (2023a), we suggested that grid and place cells are functionally connected to form a combined stable attractor in which changes to the activity of one cell type will directly influence the other. The combined attractor model was intended as an explanation for remapping, but it could also be used to simulate preplays, for example, with directional probes (Fig. 3.1). In line with the suggested grid cell mechanisms, the model would predict sequences of place cell activity from virtual grid cell movements and could thereby explain spontaneous place cell preplays without preceding search.

3.2 Regions and remapping

In human spatial memory and path planning tasks, differences in distance estimation, recall speed, and preferred routes arise depending on region partition and transitions between regions (Stevens and Coupe, 1978; McNamara, 1986; McNamara et al., 1989; Wiener and Mallot, 2003; Hochmair et al., 2008; Schick et al., 2019). Note that the definition of what constitutes a region is often not clear-cut; coherent areas separated by physical barriers are obvious examples, but regions may also overlap or have different definitions depending on the involved task or context. In any case, navigation is influenced by this hierarchical structure, and it is therefore likely that the cognitive map is similarly hierarchically organized into regions (Julian et al., 2018).

In [A3] Baumann and Mallot (2023a) we argued that the remapping observed in place and grid cells (e.g., Fyhn et al., 2007), i.e., the expression of independent firing fields in different contexts, may be the mechanism behind this hierarchical organization. That is, the world is divided into regions of different abstract contexts, which are dynamically expressed by the place and grid cell population patterns. Remapping naturally occurs at region boundaries, for example, when the animal moves from one compartment to another, but it does not occur within a coherent

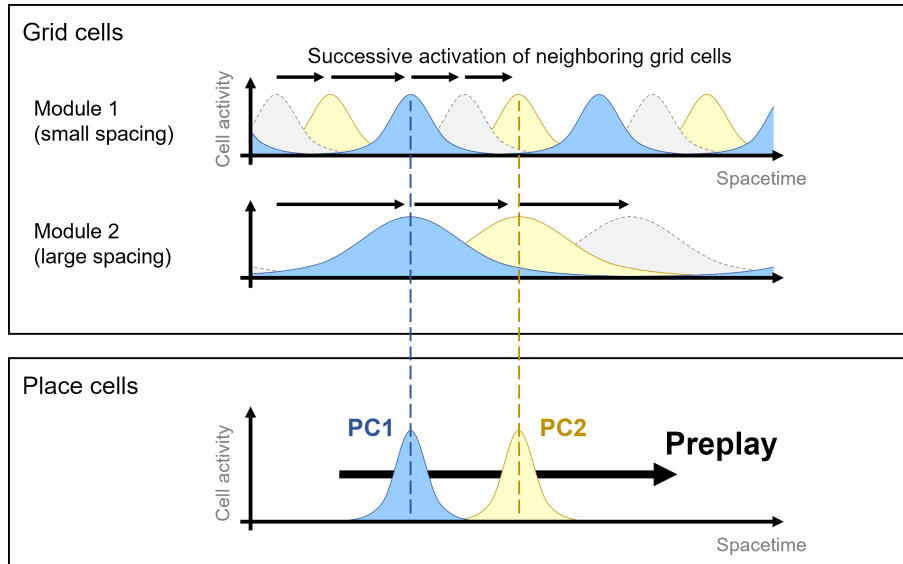


Figure 3.1: Grid cell forward probes could generate preplays even after remapping. The figure shows a simplified 1-dimensional example of the place cell model of grid cell summation, as described in [A3] Baumann and Mallot (2023a). That is, place cell formation is a result of grid cell peak coincidence. This system is able to generate place cell preplays by successively activating neighboring grid cells, which will lead to place cell discharge (PC1, PC2) whenever activity peaks in the two grid cell modules coincide. Importantly, the system only needs to know the correct grid cell sequence (e.g., the toroidal surface model in Guanella et al., 2007), but nothing about the relationship between place cell firing fields, and is therefore compatible with remapping: Grid cells remap module-wise by shifting the positions of cell firing peaks, but the ordering of cells within a module remains the same (Marozzi et al., 2015). Therefore, the same sequence of cells can still be activated after remapping, but other peaks will coincide, leading to a different place cell preplay in the new context. This example is based on a similar model for goal-directed navigation proposed in Erdem and Hasselmo (2012).

region. Due to reuse of the same cells in different population codes in other contexts, only the current context, i.e., the current bounded region, can be fully represented at a time. At region boundaries, the local context change causes the cells to remap, thus representing the new region in the cognitive map (Klukas et al., 2021). In this way, remapping should necessarily lead to a hierarchical representation of the world.

To capture these properties, we designed an attractor neural network of place and grid cell firing. In the model, the context is realized as a database of region transitions with associated place cell patterns: Whenever the simulated animal moves from one compartment to another, the corresponding pattern is reactivated via top-down input, and both place and grid cells remap concurrently due to reciprocal

connections. Because the place code in the model is limited to the current region and only changes when the animal moves to another region, transitions between regions must be an especially important navigational bottleneck; and unless the animal can *virtually* remap, wayfinding should be limited to the current region and nearby region transitions. Indeed, preplays often occur at decision points in mazes (Johnson and Redish, 2007; Xu et al., 2019), which may correspond to context or region boundaries, but it is unknown whether these can cross the boundaries or how remapping is involved.

The limitations raise the question of how routes may be planned into other, currently unavailable regions in the first place. In [A2] Brucklacher et al. (2021), we presented a hierarchical wayfinding algorithm relying on a graph with representative region nodes. In the model, the entire graph was available at all times and remapping between regions was not considered. To include these considerations, the model could be changed to only represent one region at a time while other regions would remain collapsed until they are entered (Fig. 3.2). Cells at the region boundary would then directly connect to the neighboring higher-level region nodes, making them subgoals for long-distance navigation akin to the gateways in [A3] Baumann and Mallot (2023a). As in [A2] Brucklacher et al. (2021), the so constructed cognitive map should lead to navigation in accordance with the hierarchical structure presented by the region nodes. For example, routes might preferably enter the goal region as soon as possible and cross the least amount of region boundaries regardless of actual region size, in line with human behavior (Wiener and Mallot, 2003; Hochmair et al., 2008; Schick et al., 2019).

Still, the complete remapping between contexts remains puzzling: It is clearly not the case that details about another context are inaccessible outside of that context. We can easily visualize familiar places and routes, and we are usually not surprised when we travel to other locations. The system must have a predictive component and some information about distances, directions, and spatial extent must be available in advance.

A possible explanation may be that this information is instead available at lower resolutions or reduced level of detail. Other regions may be represented by simplified spatial schemata or “gists” (Farzanfar et al., 2023) describing only abstract relations or the commonalities between multiple environments. A gist of space is a representation with reduced detail, just like the gist of an experience is meant as a brief summary of a particular memory. The integration of multiple schemata might allow for the independent representation of other contexts; without such a system, an al-




Figure 3.2: Hierarchical graph with region transitions. **(a)** A graph of two known regions with their corresponding higher-level region nodes R1, R2. The white nodes are transitions between the two regions. As in [A2] Brucklacher et al. (2021), the region nodes are bidirectionally connected to each lower-level node (edges omitted for clarity). **(b)** The lower level of the graph is collapsed and only the current region remains available for navigation. Due to the connection between the transition nodes (white) and the goal region (R2), they become important subgoals along the way. Once the goal region is entered, navigation can proceed as normal. In this scheme, wayfinding would function as in [A2] Brucklacher et al. (2021), but the system would allow for remapping and the reuse of place cells to generate the place code in other regions.

gorithm purely based on sequences of discrete regions alone might be too restrictive to explain the full experience of (human) navigation.

3.2.1 Neurological evidence for region coding

In the brain, candidates for region or broader context representation exist in various areas. In humans, the prefrontal cortex has been suggested as a locus for representing this hierarchical structure (Balaguer et al., 2016; Schapiro et al., 2016), in line with the general involvement of the area in complex planning tasks (Spiers and Gilbert, 2015; Epstein et al., 2017). Cells in the medial prefrontal cortex show different firing patterns depending on time, region, and task (Hyman et al., 2012), although the patterns are less distinct than different hippocampal place codes.

In the rodent perirhinal cortex, neurons have been found that continuously fire while the animal stays within a specific region like an arm of the experimental maze (Bos et al., 2017). Functional connectivity suggests that the postrhinal cortex might also provide navigational context to the hippocampus via the MEC, but further research, for example, from lesion studies, is required to confirm its involvement in remapping (Ho and Burwell, 2014; Julian et al., 2018).

The retrosplenial cortex is likely also involved in the encoding of regions and context: The area is reciprocally connected to the hippocampus and parahippocampal region and plays a role in spatial working memory (Vann et al., 2009) and the representation of situated contexts like rooms within a larger building (Kim and Maguire, 2018). Lesions in the area impair the ability to deal with conflicting spatial cues and to account for context changes like switching from light to darkness or distal and local cues (Pothuizen et al., 2008; Vann et al., 2009). Interestingly, some cells in the retrosplenial cortex have direction-specific local firing fields at doorways between experimental rooms (Jacob et al., 2017), and are therefore candidates for the theoretical gateway units proposed in [A3] Baumann and Mallot (2023a).

3.3 The form of spatial knowledge

The cognitive map models discussed in the previous sections are a mixture of non-metric graph-based and metric map-based approaches to spatial representation. In general, metric maps are more focused on the representation of geometry and the embedding of places and landmarks in a shared coordinate system, while graph-based models focus on relations between places and state-action-state transitions. Both types of representations have their advantages and disadvantages; for a summary, see Section 1.4 in the introduction or Peer et al. (2021) for a recent review.

Regardless of computational or functional advantages, the question may be asked whether the information represented in the human cognitive map is more graph- or map-like. The form of the spatial representation in the brain has been debated since the very beginnings of the study of spatial cognition, dating back to Edward Tolman’s “means-end-field” (Tolman (1932), a state-action graph) and O’Keefe and Nadel (1978) arguing for Euclidean metric properties to represent a similarly Euclidean metric world.

In general, graph models are not metric and thus less constrained than metric maps: Spatial relations between encoded places are independent and can describe arbitrary information that need not reflect any real-world properties. For example, if distance or direction between nodes are encoded (the “labeled graph”, cf. Warren, 2019), the information can violate metric postulates like the triangle inequality. In a consistent metric map, on the other hand, stored information needs to adhere to the metric and cannot be arbitrary. Under the metric scheme, these constraints greatly reduce the number of possible configurations that the cognitive map can take, and

compromises must be made to account for subjective measurement errors (O’Keefe and Nadel, 1978).

This possible mismatch between the subjective experience of space and the corresponding map may be used to distinguish the models to reveal the underlying structure of spatial knowledge: In a recent navigation study by Warren et al. (2017), human participants explored a VR maze by physically walking around a large room while wearing head-mounted displays. Importantly, the maze contained two seamless non-Euclidean wormholes connecting distant parts. These wormholes were intended to create a mismatch between the subjective experience of the maze, i.e., the local position information from path integration in the real physical room, and the presented virtual reality. After learning the maze, participants had to estimate directions to the remembered positions of different landmarks.

The authors found that the estimates were systematically biased towards the wormholes in a way that best matched vectors along the shortest path (through the wormholes) in the maze, but not the shortest distance in Euclidean metric space, i.e., a straight line in the physical room. Forward and backward estimates also differed on the same landmark pair, i.e., the estimated direction from landmark a to b was not the inverse of the estimated direction from b to a . The authors concluded that the results are incompatible with a Euclidean metric map and are better explained by a non-metric graph labeled with local vector information.

In [A4] Baumann and Mallot (2023b) we considered another alternative, a distorted metric map of the maze that would minimize the difference to the subjective experience of the participants in the Warren et al. (2017) experiment. We were able to obtain such a map by embedding the proposed labeled graph into 2D Euclidean coordinates via the minimization of a stress function. We compared the models on the original dataset and found that our model, the “embedded graph”, predicted the directional estimates equally well, thereby refuting the conclusion in Warren et al. (2017). At least from this study, it can therefore *not* be concluded that the cognitive map takes the form of a non-metric graph rather than a metric map, and the question remains open.

In support of metric models like the embedded graph, we argued that the main difference between the two model classes lies in the treatment of repeated measurements as familiarity with the environment increases: In the beginning, the lack of constraints in the non-metric graph may be computationally advantageous, because spatial knowledge can immediately be stored, while the estimation of coordinates for the metric embedding remains imprecise with only few measurements. However, as

more information about the environment is acquired, this dynamic changes: In the non-metric labeled graph, improving distance and angle information will only ever improve a single label without exploiting the constraints that these measurements might impose on adjacent labels. In the metric embedding, on the other hand, the update of one estimate will also improve the estimate of other adjacent places if not the entire map. Since place learning is usually not finished after a single pass, the embedding is therefore the preferable representation. The main advantage of a metric cognitive map is thus not its resemblance to the physical world but the possibility to integrate repeated measurements into a consolidated structure ([A4] Baumann and Mallot, 2023b).

Of course, the purely non-metric graph and the complete Euclidean metric map are extremes, and mixed models that combine non-metric and metric information are probable and may have useful advantages, like the hierarchically structured map with different types of representation at different levels discussed in the previous section (Section 3.2, see also Couclelis et al. (1987), Kuipers (2000), and Meilinger (2008)). A combined hierarchical model with metric representation at the lowest level but only simplified relational (graph-based) knowledge of more distant regions could be a very efficient way to represent extended space, because only the most immediate surroundings need to be available in detail for navigation.

On the other hand, there is no intrinsic requirement for more abstract superordinate regions to be represented as a non-metric graph. A metric model could also be used for higher-level maps at different resolution levels, much like an actual drawn map; after all, it is unlikely that available metric information would not be used by the brain. Consider, for example, multiple buildings on a campus, each represented as a separate region in the cognitive map. With increased familiarity, general knowledge about the relative distance, scale, and orientation of the buildings would accumulate in addition to the local maps. Over time, the individual buildings could then also be combined in a higher-level metric map of the campus.

Another alternative to the dichotomy of metric and non-metric spatial knowledge is the possibility that the brain is able use either model depending on the circumstance. For example, metric maps may preferably be used in open areas with many possible paths, while graph-based models may be more efficient in environments with few alternative routes like the interior of buildings (Peer et al., 2021). It is also possible that the brain uses both models dynamically and can transform one type of representation into the other. That is, metric maps can be constructed from graph knowledge and vice versa. Such redundant coding would give the cognitive

map increased power and flexibility, but it might be inefficient if one type of coding is enough.

3.3.1 Non-spatial domains

The distinction between metric and non-metric representations is not only relevant to physical space, especially when considering the complex role of the hippocampus and the hippocampal formation. In addition to the encoding of space and position, the hippocampus is also necessary for the formation of long-term memory (Buzsáki and Moser, 2013). Accordingly, both memory and general planning may have evolved from earlier navigation mechanisms, and the brain might treat physical and mental spaces fundamentally the same. The representational structure of the cognitive map, that is map coding, landmark representation, regionalization, and route planning, may therefore be a general system for a wide variety of non-spatial cognitive domains (Tolman, 1948; O’Keefe and Nadel, 1978; Buzsáki and Moser, 2013; Eichenbaum and Cohen, 2014; Constantinescu et al., 2016; Epstein et al., 2017; Eichenbaum, 2017; Bellmund et al., 2018).

For example, in the rodent hippocampus, cells have been identified in addition to place cells that code for odors (Wood et al., 1999), time points (MacDonald et al., 2011; Rubin et al., 2015), and sound frequencies (Aronov et al., 2017) independent of location. In humans, fMRI responses from the hippocampal formation match vector coding in conceptual or abstract spaces: For example, hippocampal activity was found to match interpersonal distances in a social space spanned by the affiliations and social hierarchy of multiple people (Tavares et al., 2015). Similarly, grid-like activity was measured in the entorhinal cortex when human subjects viewed images of birds with variable neck and leg length sourced from an abstract 2D “bird space” spanned by these two parameters. Response increased when the depicted image sequences were aligned to a sixfold rotational symmetry axis, indicative of grid representation (Constantinescu et al., 2016).

The findings suggest that these domains are also encoded as metric, which would allow the brain to represent arbitrary states in these abstract spaces rather than just discrete stimuli or events (Aronov et al., 2017). Wayfinding would correspond to the consideration of future states and state transitions and may therefore play a more general role in planning or prospective thinking (Buzsáki and Moser, 2013; Epstein et al., 2017), and replays of past routes through conceptual spaces might correspond to episodic memory (Eichenbaum, 2017). Consistent with these ideas, damage to

the hippocampus results in global amnesia in both spatial and non-spatial domains (Squire and Knowlton, 1995; Squire, 2004; Eichenbaum, 2017).

Following the preceding discussion, abstract spaces might also be divided into different regions, like, for example, a group of mutual acquaintances in the social space. The same systems that encode hierarchy and higher-level regions in physical space could also represent abstract categories. In the temporal domain, time-encoding cell ensembles in the hippocampus form qualitatively different representation in a manner resembling place cell remapping when the main temporal parameter is altered, (MacDonald et al., 2011; Buzsáki and Moser, 2013). So far, the interaction between region boundaries, gateways, and remapping has been mostly restricted to the spatial domain, but these properties (and in this sense, the models discussed here) may have a much broader application and play a general role in all cognitive systems.

3.4 Open questions and future research

- **Gateway-based coding and rooms with multiple entrances:** In [A3] Baumann and Mallot (2023a), we proposed that the place code within a room is based on remapping at its entrance. The model is consistent with cell recordings from multicompartement mazes, but the rooms usually only have a single entrance. What happens to the place code if the animal enters a previously learned room from another entrance? Will it generate a new place code like in the hairpin maze (e.g., direction-specific firing in Derdikman et al., 2009) or can it recover the original code and its coordinates when the room is recognized? And at what point does that happen?
- **Preplays and remapping:** As discussed above, the interaction between the prediction of upcoming paths in preplays and hippocampal remapping is unexplored. If preplays indeed reflect future paths, then cases must exist where the upcoming path will lead the animal into another region, which is not currently represented in the hippocampal place code (due to remapping). Is virtual remapping possible, or can preplays simply not cross region boundaries?
- **Regionalization and place coding in natural environments:** Fundamentally, the remapping model discussed here is based on the assumption that natural outdoor environments are also subdivided into different regions. That is, remapping should also occur when the animal simply explores its en-

vironment, and not just in clearly defined mazes. There may be fundamental differences between the cognitive maps of wild and laboratory animals due to the limited world of the latter.

- **Refinement of the embedded graph:** So far, the embedded graph model has only been evaluated on a single dataset, and it may very well be unable to match the non-metric graph on other datasets. The big remaining question is whether the same behavior that has been explained by a non-metric graph can also always be explained with a metric embedding.

Bibliography

- Aronov, Dmitriy, Rhino Nevers, and David W Tank (2017). “Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit”. *Nature* 543.7647, pp. 719–722.
- Balaguer, Jan et al. (2016). “Neural mechanisms of hierarchical planning in a virtual subway network”. *Neuron* 90.4, pp. 893–903.
- Banino, Andrea et al. (2018). “Vector-based navigation using grid-like representations in artificial agents”. *Nature* 557.7705, pp. 429–433.
- Baumann, Tristan and Hanspeter A Mallot (2023a). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239.
- Baumann, Tristan and Hanspeter A Mallot (2023b). “Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments”. Under review.
- Bay, Herbert et al. (2008). “Speeded-up robust features (SURF)”. *Computer vision and image understanding* 110.3, pp. 346–359.
- Bellmund, Jacob LS et al. (2016). “Grid-cell representations in mental simulation”. *Elife* 5, e17089.
- Bellmund, Jacob LS et al. (2018). “Navigating cognition: Spatial codes for human thinking”. *Science* 362.6415, eaat6766.
- Bos, Jeroen J et al. (2017). “Perirhinal firing patterns are sustained across large spatial segments of the task environment”. *Nature communications* 8.1, p. 15602.
- Brucklacher, Matthias, Hanspeter A Mallot, and Tristan Baumann (2021). “Hierarchical Planning in Multilayered State-Action Networks.” *ESANN 2021 proceedings*.
- Bush, Daniel et al. (2015). “Using grid cells for navigation”. *Neuron* 87.3, pp. 507–520.
- Buzsáki, György and Edvard I Moser (2013). “Memory, navigation and theta rhythm in the hippocampal-entorhinal system”. *Nature neuroscience* 16.2, pp. 130–138.

- Constantinescu, Alexandra O, Jill X O'Reilly, and Timothy EJ Behrens (2016). "Organizing conceptual knowledge in humans with a gridlike code". *Science* 352.6292, pp. 1464–1468.
- Couclelis, Helen et al. (1987). "Exploring the anchor-point hypothesis of spatial cognition". *Journal of environmental psychology* 7.2, pp. 99–122.
- Derdikman, Dori et al. (2009). "Fragmentation of grid cell maps in a multicompart-ment environment". *Nature neuroscience* 12.10, pp. 1325–1332.
- Dragoi, George and Susumu Tonegawa (2011). "Preplay of future place cell sequences by hippocampal cellular assemblies". *Nature* 469.7330, pp. 397–401.
- Eichenbaum, Howard (2017). "The role of the hippocampus in navigation is memory". *Journal of neurophysiology* 117.4, pp. 1785–1796.
- Eichenbaum, Howard and Neal J Cohen (2014). "Can we reconcile the declarative memory and spatial navigation views on hippocampal function?" *Neuron* 83.4, pp. 764–770.
- Epstein, Russell A et al. (2017). "The cognitive map in humans: spatial navigation and beyond". *Nature neuroscience* 20.11, pp. 1504–1513.
- Erdem, Uğur M and Michael Hasselmo (2012). "A goal-directed spatial navigation model using forward trajectory planning based on grid cells". *European Journal of Neuroscience* 35.6, pp. 916–931.
- Farzanfar, Delaram et al. (2023). "From cognitive maps to spatial schemas". *Nature Reviews Neuroscience* 24.2, pp. 63–79.
- Franz, Matthias O et al. (1998). "Learning View Graphs for Robot Navigation". *Autonomous Robots* 5, pp. 111–125.
- Fyhn, Marianne et al. (2007). "Hippocampal remapping and grid realignment in entorhinal cortex". *Nature* 446.7132, pp. 190–194.
- Guanella, Alexis, Daniel Kiper, and Paul Verschure (2007). "A model of grid cells based on a twisted torus topology". *International journal of neural systems* 17.04, pp. 231–240.
- Ho, Jonathan W and Rebecca D Burwell (2014). "Perirhinal and postrhinal func-tional inputs to the hippocampus". *Space, time and memory in the hippocampal formation*, pp. 55–81.
- Hochmair, Hartwig H, Simon J Büchner, and Christoph Hölscher (2008). "Impact of regionalization and detour on ad-hoc path choice". *Spatial Cognition & Com-putation* 8.3, pp. 167–192.
- Hyman, James M et al. (2012). "Contextual encoding by ensembles of medial pre-frontal cortex neurons". *Proceedings of the National Academy of Sciences* 109.13, pp. 5086–5091.
- Jacob, Pierre-Yves et al. (2017). "An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex". *Nature neuroscience* 20.2, pp. 173–175.
- Johnson, Adam and A David Redish (2007). "Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point". *Journal of Neuro-science* 27.45, pp. 12176–12189.

- Julian, Joshua B et al. (2018). “The neurocognitive basis of spatial reorientation”. *Current Biology* 28.17, R1059–R1073.
- Kim, Misun and Eleanor A Maguire (2018). “Hippocampus, retrosplenial and parahippocampal cortices encode multicompartement 3D space in a hierarchical manner”. *Cerebral Cortex* 28.5, pp. 1898–1909.
- Klukas, Mirko et al. (2021). “Fragmented Spatial Maps from Surprisal: State Abstraction and Efficient Planning”. *bioRxiv*, pp. 2021–10.
- Kuipers, Benjamin (1978). “Modeling spatial knowledge”. *Cognitive Science* 2.2, pp. 129–153.
- Kuipers, Benjamin (2000). “The spatial semantic hierarchy”. *Artificial intelligence* 119.1-2, pp. 191–233.
- MacDonald, Christopher J et al. (2011). “Hippocampal “time cells” bridge the gap in memory for discontinuous events”. *Neuron* 71.4, pp. 737–749.
- Mallot, Hanspeter A, Gerrit A Ecke, and Tristan Baumann (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17.
- Marozzi, Elizabeth et al. (2015). “Purely translational realignment in grid cell firing patterns following nonmetric context change”. *Cerebral Cortex* 25.11, pp. 4619–4627.
- McNamara, Timothy P (1986). “Mental representations of spatial relations”. *Cognitive psychology* 18.1, pp. 87–121.
- McNamara, Timothy P, James K Hardy, and Stephen C Hirtle (1989). “Subjective hierarchies in spatial memory.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15.2, p. 211.
- Meilinger, Tobias (2008). “The network of reference frames theory: A synthesis of graphs and cognitive maps”. *Spatial Cognition VI*. Springer, pp. 344–360.
- Nyberg, Nils et al. (2022). “Spatial goal coding in the hippocampal formation”. *Neuron*.
- O’Keefe, John and Lynn Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- O’Neill, Joseph et al. (2017). “Superficial layers of the medial entorhinal cortex replay independently of the hippocampus”. *Science* 355.6321, pp. 184–188.
- Ólafsdóttir, H Freyja, Francis Carpenter, and Caswell Barry (2016). “Coordinated grid and place cell replay during rest”. *Nature neuroscience* 19.6, pp. 792–794.
- Ólafsdóttir, H Freyja et al. (2015). “Hippocampal place cells construct reward related sequences through unexplored space”. *Elife* 4, e06063.
- Peer, Michael et al. (2021). “Structuring knowledge with cognitive maps and cognitive graphs”. *Trends in cognitive sciences* 25.1, pp. 37–54.
- Pfeiffer, Brad E (2020). “The content of hippocampal “replay””. *Hippocampus* 30.1, pp. 6–18.
- Pfeiffer, Brad E and David J Foster (2013). “Hippocampal place-cell sequences depict future paths to remembered goals”. *Nature* 497.7447, pp. 74–79.

- Ponulak, Filip and John J Hopfield (2013). “Rapid, parallel path planning by propagating wavefronts of spiking neural activity”. *Frontiers in computational neuroscience* 7, p. 98.
- Pothuizen, Helen HJ, John P Aggleton, and Seralynne D Vann (2008). “Do rats with retrosplenial cortex lesions lack direction?” *European Journal of Neuroscience* 28.12, pp. 2486–2498.
- Rubin, Alon et al. (2015). “Hippocampal ensemble dynamics timestamp events in long-term memory”. *elife* 4, e12247.
- Schapiro, Anna C et al. (2016). “Statistical learning of temporal community structure in the hippocampus”. *Hippocampus* 26.1, pp. 3–8.
- Schick, Wiebke et al. (2019). “Language cues in the formation of hierarchical representations of space”. *Spatial Cognition & Computation* 19.3, pp. 252–281.
- Spiers, Hugo J and Sam J Gilbert (2015). “Solving the detour problem in navigation: a model of prefrontal and hippocampal interactions”. *Frontiers in human neuroscience*, p. 125.
- Squire, Larry R (2004). “Memory systems of the brain: a brief history and current perspective”. *Neurobiology of learning and memory* 82.3, pp. 171–177.
- Squire, Larry R. and Barbara J. Knowlton (1995). “Memory, hippocampus, and brain systems”. *The cognitive neurosciences*. Ed. by Michael S. Gazzaniga. Cambridge, MA: The MIT Press. Chap. 53, pp. 825–837.
- Stevens, Albert and Patty Coupe (1978). “Distortions in judged spatial relations”. *Cognitive psychology* 10.4, pp. 422–437.
- Tang, Wenbo, Justin D Shin, and Shantanu P Jadhav (2021). “Multiple time-scales of decision-making in the hippocampus and prefrontal cortex”. *Elife* 10, e66227.
- Tavares, Rita Morais et al. (2015). “A map for social navigation in the human brain”. *Neuron* 87.1, pp. 231–243.
- Tolman, Edward C (1948). “Cognitive maps in rats and men.” *Psychological review* 55.4, p. 189.
- Tolman, Edward C (1932). *Purposive Behavior in Animals and Men*. New York: The Century Co.
- Vann, Seralynne D, John P Aggleton, and Eleanor A Maguire (2009). “What does the retrosplenial cortex do?” *Nature reviews neuroscience* 10.11, pp. 792–802.
- Warren, William H (2019). “Non-euclidean navigation”. *Journal of Experimental Biology* 222.Suppl_1, jeb187971.
- Warren, William H et al. (2017). “Wormholes in virtual space: From cognitive maps to cognitive graphs”. *Cognition* 166, pp. 152–163.
- Wiener, Jan M and Hanspeter A Mallot (2003). “‘Fine-to-coarse’ route planning and navigation in regionalized environments”. *Spatial cognition and computation* 3.4, pp. 331–358.
- Wilson, Matthew A and Bruce L McNaughton (1993). “Dynamics of the hippocampal ensemble code for space”. *Science* 261.5124, pp. 1055–1058.
- Wood, Emma R, Paul A Dudchenko, and Howard Eichenbaum (1999). “The global record of memory in hippocampal neuronal activity”. *Nature* 397.6720, pp. 613–616.

- Xu, Haibing et al. (2019). “Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze”. *Neuron* 101.1, pp. 119–132.
- Yamamoto, Jun and Susumu Tonegawa (2017). “Direct medial entorhinal cortex input to hippocampal CA1 is crucial for extended quiet awake replay”. *Neuron* 96.1, pp. 217–227.

Appendix A1

Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations

Authors: Hanspeter A Mallot, Gerrit A Ecke, and Tristan Baumann, 2020.

Abstract

Topological schemes for navigation from visual snapshots have been based on graphs of panoramic images and action links allowing the transition from one snapshot point to the next; see, for example, Cartwright and Collett (1987) or Franz et al. (1998). These algorithms can only work if at each step a unique snapshot is recognized to which a motion decision is associated. Here, we present a population coding approach in which place is encoded by a population of overlapping “firing fields”, each of which is activated by the recognition of an unspecific “micro-snapshot” (i.e. feature), and associated to a subsequent action. Agent motion is then computed by a voting scheme over all activated snapshot-to-action associations. The algorithm was tested in a large virtual environment (Virtual Tübingen, Van Veen et al., 1998) and shows biologically plausible navigational abilities.

A1.1 Introduction

A1.1.1 Parsimonious Representations of Space

The evolution of spatial cognition in animals started from simple stimulus-response behaviors such as stimulus-driven orienting reactions, and proceeded further by a

number of innovations that include (i) mechanisms for egomotion perception and path integration, (ii) the memorization of stimulus-response pairs composed of a distinguishable landmark and a navigational action (“recognition-triggered response”), (iii) the concatenation of such recognition-triggered responses into chains or routes, and (iv) the linking-up of multiple recognition-triggered responses into networks or graphs in which novel routes can be inferred by the combination of known route segments (Trullier et al., 1997; Mallot and Basten, 2009; Wiener et al., 2011; Madl et al., 2015). In addition, mechanisms for invariant landmark and place recognition, strategic selection of way- or anchor-points, metric embedding of place-graphs, or hierarchical graph structures may improve navigational performance and are thus likely to play a role.

Since many different models can be built on these elements, it is of interest to ask for a minimal or most parsimonious model supporting a given level of behavioral flexibility. In this paper, we address this question for the case of the minimal cognitive architecture supporting way-finding behavior. By a minimal model, we mean a model meeting the following requirements:

1. A minimal model should be close to the evolutionary starting point of stimulus-response schemata;
2. it should require only a small amount of visual invariance in object recognition and therefore work with the rawest possible image information;
3. it should use simple decision processes in path planning such as recognition-triggered responses; and
4. it should not rely on explicit metric information which is hard to obtain.

With these constraints in mind, we present a model for graph-based navigation that marks a lower bound of cognitive complexity required for way-finding and that can be used to study further improvements resulting from additional evolutionary innovations

The model presented in this paper is not primarily about the hippocampal system for place as is known from rodents and some other mammalian groups, although some inspiration has been drawn from these results. Our main interest, however, is a computational theory of navigation based on devices such as snapshots and state-action schemata. Such computational theory will have implications for navigational behavior in insects, mammals, humans, and even robots.




Figure A1.1: Graph-navigation in single-unit and population coding. A: In standard topological navigation, every place is represented by a unique node and route segments are given by the graph-links. The desired trajectory shown in blue is therefore a graph path. B: In dual population coding, a bundle of paths is constructed for a given navigation problem. Three such paths without common nodes (except start and goal) are shown in red, orange, and yellow colors in the figure. The desired trajectory (shown in blue) is then calculated by a voting scheme over the currently visible nodes of all paths. It is generally not a path of the graph.

A1.1.2 Dual Population Coding

Two basic elements of spatial representations are (i) stimulus-response schemata such as Tolman’s (1932) means-ends-relations, O’Keefe and Nadel’s (1978) taxon system, Kuipers’ (1978) control laws, or the place-recognition-triggered response of Trullier et al. (1997), and (ii) the representations of places and place relations such as Cartwright and Collett’s (1982) snapshot-codes for places or O’Keefe and Nadel’s (1978) locale system. The two systems are connected by the role that place recognition takes as a “stimulus” in the stimulus-response schemata involved.

In the classical state-action-approach, it is assumed that each place is a state represented by just one node of a graph (Franz et al., 1998; Kuipers, 1978; Kuipers, 2000; Muller et al., 1996) such that a unique state-action schema will control each navigational step. If place recognition fails, navigation will go wrong. Robustness of navigation therefore depends foremost on the robustness and invariance of place recognition as a prerequisite. Here, we argue that in an evolutionary view of navigation, robust pathfinding should be possible even with rudimentary place recognition and distributed place representations.

Our model differs from standard models of topological navigation (Franz et al.,

1998; Kuipers, 1978; Kuipers, 2000; Muller et al., 1996) in two major respects that can be summarized as “dual population coding” (see Fig. A1.1): first, at any instant in time, many nodes of the graph are activated and encode the agent’s position in a population scheme. This avoids costly selection processes of strategic anchor points and has the additional advantage that the visual cues and recognition processes can be kept simple. Of course, population coding of space is well in line with empirical findings in the place-cell literature (Wilson and McNaughton, 1993). Second, as a consequence of population coding of space, route selection has to be based on many interacting recognition-triggered response schemata, one for each active unit in the population code. This is implemented by a voting scheme where the suggested motion decisions from all active schemata are averaged. The idea of view voting has been suggested earlier for human behavioral data (Mallot and Gillner, 2000). In insect navigation, a similar scheme has been suggested for route following with multiple snapshots by (Baddeley et al., 2012; Differt and Stürzl, 2021; Smith et al., 2007), but unlike our model, these models do not allow for alternative route decisions from a given position. As a result of dual population coding, the trajectory eventually found by the algorithm is not a path of the graph, but a metric average of bundles of many paths connecting individual nodes in the population codes for start and goal.

A1.2 Navigation Algorithm

This section will explain the algorithm in five steps, starting from the initial definition and later matching of features, and proceeding to the learning of graph edges and their directional labels. Once the graph is learned, the voting scheme is applied for pathfinding. All examples shown are taken from a virtual reality implementation where a simulated agent is exploring and later navigating in the “Virtual Tübingen” environment (Van Veen et al., 1998).

A1.2.1 Feature Detection

Micro-snapshots are defined as “upright speeded-up robust features” (U-SURF), as implemented in the OpenCV computer vision library (Bay et al., 2008; Bradski, 2000). SURF finds interest points as intensity blobs by searching local maxima of the determinant of the image Hessian; color information is ignored. Scale invariance is achieved by considering each feature point at its optimal scale. In a second step, a

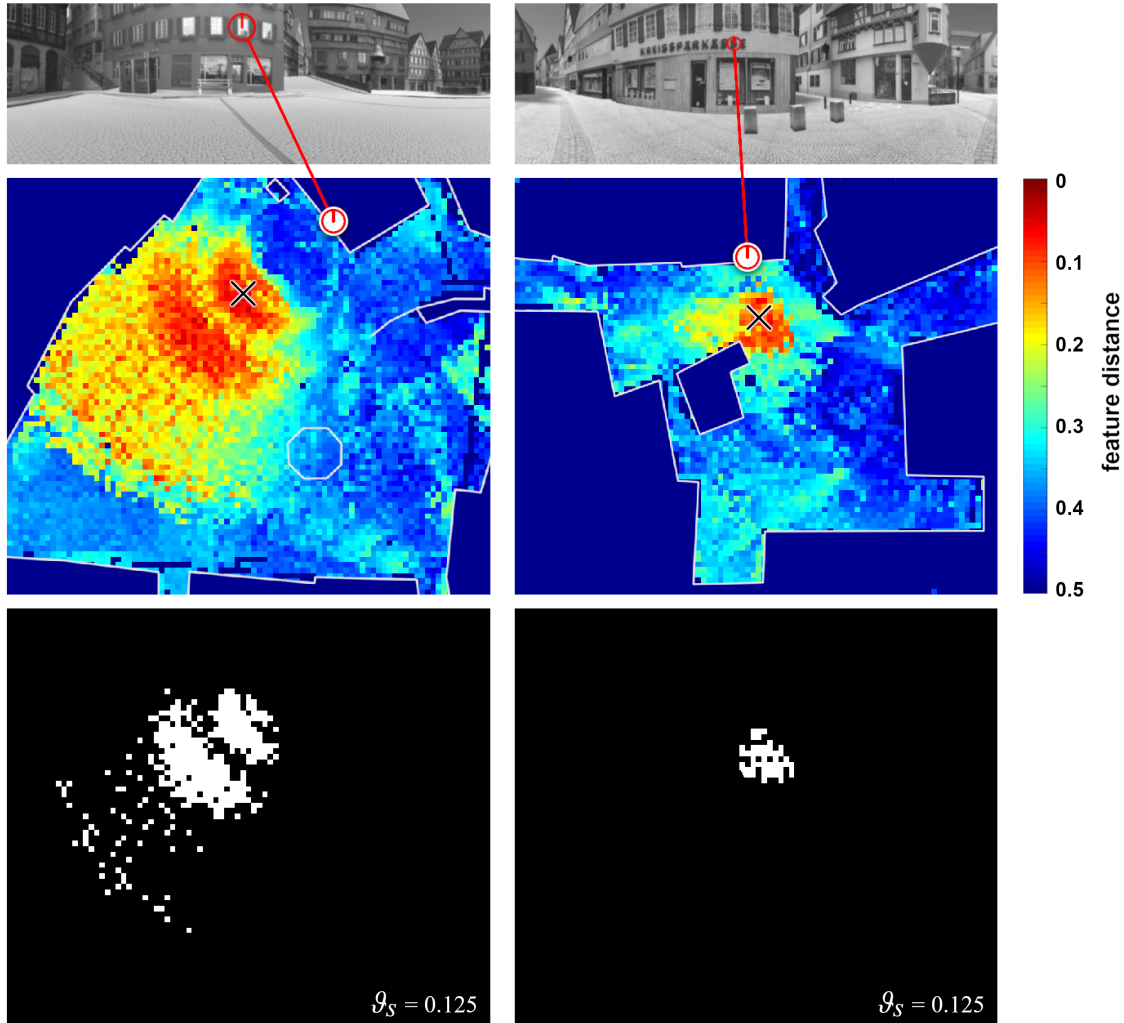


Figure A1.2: Place fields. Top row: two views from a scene with detected features (a window and a letter from a company nameplate). Middle row: local maps of the environment superimposed with the similarity of the reference feature to any feature detected from each position in the map ($\min_i \|\mathbf{d}_{\text{ref}} - \mathbf{d}_{i,x}\|^2$). The black \times marks the position where the reference feature was first detected. Third row: same map with points where one feature was identified with the reference feature, based on both the similarity criterion and the neighborhood consensus criterion. Note that the set of locations is not connected. Also, in the larger open space (left column: Market place), the place fields tend to be larger than in smaller places (right column: Street crossing “Krumme Brücke”).

64-dimensional vector (“descriptor”) is associated with each blob, containing information about image intensity gradients in a small patch around the interest point. The descriptor is used to compare and match features with each other. In U-SURF, it is not assigned a unique orientation and is therefore not rotation invariant. Rotation invariance is not required in our algorithms since the agent is confined to movements in the plane. The number of scale levels was limited to two octaves with two layers each since information about the viewing distance of a feature should not be completely ignored. SURF feature robustness was further increased by considering only features that appear in two successive frames. Whenever we refer to features as extracted from a given frame it means that those features also appeared in the preceding frame.

The features of a frame were ranked according to the value of the determinant of the local Hessian, i.e. their contrast. Then, the features were pruned so that only up to 30 features per frame were used for further analysis. In principle, the amount of features per frame could also be reduced globally by increasing the detection threshold of the SURF method. However, this could potentially lead to situations where no feature would be detected at locations where contrast is low.

We denote the features as f_i and their descriptors as \mathbf{d}_i ; $F = \{f_i | i = 1, \dots, n\}$ is the set of all features stored in the system.

A1.2.2 Feature Matching

Whenever a feature is detected by the U-SURF procedure, it is checked for identity with all stored features in F using two criteria. First, the root mean squared difference between the descriptors of the compared features should be below a threshold ϑ_S . Second, to avoid aliasing in large sets of features, we require that the features share a context of at least ϑ_N other features (neighborhood consensus). To this end, we store for each feature f_i the set N_i of simultaneously visible other features. Two features f_i, f_j are thus identified with each other, if $\|\mathbf{d}_i - \mathbf{d}_j\|^2 < \vartheta_S$ and $|N_i \cap N_j| \geq \vartheta_N$. If an encountered feature is found to be novel, it is included into F .

The threshold for neighborhood consensus, ϑ_N , depends on the total number of features detected in each image. In our simulations, the value was set to $\vartheta_N = 4$ at up to 30 different features per frame. Note that aliasing still occurred occasionally even with expanded feature-neighbor matching (see Sect. A1.2.3 and Fig. A1.4 below). In practice, the algorithm is robust against a small amount of outliers and can find and navigate routes even with faulty map data. See Sects. A1.2.5 and A1.3 for more




Figure A1.3: Graph edge learning. Left: The dotted line shows the trajectory of agent with time steps (small dots) and learning steps (bold dots). Two features f_i and f_j have already been encountered and added to the feature set. The agent is currently moving with heading η in the place field of feature f_i , but outside of the place field of feature f_j . Right: The agent has now passed the overlap zone where both features are detected. At the next learning step, feature f_j is detected but feature f_i has moved out of sight. In this situation, a bidirectional pair of edges a_{ij}, a_{ji} is added to the graph and the edges are labeled with the current heading or its inverse, $\pm\eta$.

details.

Figure A1.2 shows two features in the respective images from the Virtual Tübingen data set. In the second row, the position from which each feature was first defined and added to F is marked by a cross. For all positions in open space, color indicates the similarity of the most similar visible feature with the stored one; dark blue marks locations inside of houses that cannot be entered by the agent. The third row of Fig. A1.2 shows the area from which the feature is detected, using the two-step comparison procedure with similarity of descriptors and neighborhood consensus. It will be called the place field of the feature and roughly corresponds to the catchment areas in snapshot homing or the firing fields of a neuron tuned to the feature.

A1.2.3 Graph Edge Formation

If a feature that was previously visible gets out of sight, the agent must have traveled a path out of the place field of this feature to some point inside the place fields of other features that remain or have become visible. This is the basic idea of learning graph edges in the algorithm. In order to avoid too high densities of graph links, new edges are not stored at every time step, but at a slower pace.

The basic time step of the algorithm is the frame, i.e. the recording of one image; we denote frames by the index t . The frame rate used in the graphics simulations




Figure A1.4: Map of the testing environment “Virtual Tübingen” with view graph. The view graph can be embedded into a map by placing each feature at the agent’s position from where it was first detected, and drawing the graph’s edges between them blue lines). The shown graph completely maps the virtual environment and consists of 222,433 nodes and 3,492,096 edges. Some of the edges connect very distant features (long blue lines crossing the empty white space). These are wrong connections resulting from aliasing.

below is 30 frames per second. Graph learning does not occur at every time step, but only once in a while, when the agent has moved sufficiently far away from the last learning event. Learning steps are counted by a second counter l . In the simulations below, the distance that the agent must have traveled before a new learning step occurs was set to two simulated meters, which corresponds to 30 frames. Note that we use the position ground truth of the VR simulation for stepping l . This can easily be relaxed by some simple path integration algorithm which was, however, not implemented. The time steps at which learning counter l is stepped, are denoted by $t(l)$.

Let F_1 and F_2 be the sets of features visible at two subsequent learning steps l and $l + 1$, respectively. Assume $f_j \in F_1$ and $f_j \notin F_2$. We then add a pair of directed edges a_{jk}, a_{kj} (forward and backward) between f_j and up to three randomly chosen features $f_k \in F_2$ to the graph. The edges are labeled with the current heading or its inverse, respectively (see Fig. A1.3 and next paragraph). The number three of edges created per vanishing feature is chosen to avoid exceedingly high computation costs in later graph search. For the same reason, the upper limit of a node’s degree after repeated visits is set to 100.

An example of the graph after prolonged exploration appears in Fig. A1.4.

A1.2.4 Edge Labeling and Reference Direction

During the entire travel, the agent is estimating and maintaining an allocentric reference direction ν which is initialized to the value $\nu = 0$ at frame 1 (see Fig. A1.5). All other angles are expressed relative to this reference direction, i.e. in an allocentric scheme similar to the head direction in allocentric path integration (Cheung and Vickerstaff, 2010; Taube, 2007). The dependent angles are (i) the current heading angle η_i , (ii) the feature bearings $\hat{\beta}_i$ stored with each feature f_i upon definition of the feature, and (iii) the directional labels of the edges a_{ij}, a_{ji} which are initialized with or against the current heading angle, i.e. $a_{ij} = \eta_t$ or $a_{ji} = \eta_t + \pi$, respectively.

The reference direction is constantly affected by a noise process \mathbf{n} and updated according to the available landmark cues, i.e. the bearings of known features. Let F_t denote the set of known features visible at frame t and $\hat{\beta}_i$ be their stored bearings. The agent compares the current feature bearings with the stored ones and computes the average deviation as a circular mean. Then, the reference direction is updated as

$$\nu_{t+1} = \nu_t + \frac{\lambda}{|F_t|} \text{cmean}_{\{i|f_i \in F_t\}} \left(\hat{\beta}_{t,i} - \beta_{t,i} \right) + \mathbf{n}, \quad (\text{A1.1})$$

where λ is set to 0.05 and the standard deviation of \mathbf{n} is set to $\sigma = 0.025$ rad. The circular mean of a set of angles $\{\gamma_i\}$ is defined as

$$\text{cmean}_A(\gamma_i) := \text{atan2} \left(\sum_{i \in A} \cos \gamma_i, \sum_{i \in A} \sin \gamma_i \right). \quad (\text{A1.2})$$

This updating rule attributes the average bearing error to the reference direction. It can compensate for the noise, but introduces a new type of error if the features are unequally distributed in the image. Assume, for example, that the agent relies only on features on its left. If it moves forward, these features will move further to the left, leading to positive deviations $\hat{\beta}_i - \beta_i$. The algorithm will then assume that the reference direction has turned to the left. As a result, the reference direction in a large environments drifts with the agent's position, as is illustrated in Fig. A1.6. However, in prolonged exploration, the assumed reference directions convergence to a stable, locally consistent distribution over explored space.

In addition, the stored bearings for each feature, $\hat{\beta}_i$ are updated at each learning step at which the feature i is re-detected by the iterative mean:

$$\hat{\beta}_{t(l+1),i} = \frac{c_i}{c_i + 1} \hat{\beta}_{t(l),i} + \frac{1}{c_i + 1} \beta_{t(l+1),i}, \quad (\text{A1.3})$$




Figure A1.5: The head-direction system. A: During the entire simulation, the system is maintaining a reference direction ν which is initialized to the movement direction in the first frame. Heading angle η and feature bearings β_i are always expressed relative to ν . The “true north” direction is known to the virtual reality simulation, but not to the agent. B: If a feature is detected, its stored bearing label $\hat{\beta}_i$ is compared to the actual bearing in the current image, β_i and the reference direction is updated so as to reduce the difference between $\hat{\beta}_i$ and β_i . Of course, this is done for many features simultaneously, as described in Eq. A1.1.

where c_i is a counter stepped at each update and $\beta_{t(l),i}$ is measured relative to the current compass direction ν_t .

Finally, a link a_{ij} may be rediscovered upon a later encounter of the same location. In this case the associated direction label α_{ij} is updated as

$$\alpha_{ij}^{\text{new}} = \frac{c_{ij}}{c_{ij} + 1} \alpha_{ij}^{\text{old}} + \frac{1}{c_{ij} + 1} \eta_t. \quad (\text{A1.4})$$

Again, this can happen only when the slow learning counter l is stepped. As for the bearings, c_{ij} is a counter stepped at every update and the η_t is measured from the current compass direction. Note that the counters c_i and c_{ij} in Eqs. A1.3 and A1.4 can be avoided by replacing the bearing and heading angles by unit vectors and storing sums of these unit vectors as labels. From the accumulated vectors, the angles can then simply be obtained by the atan2 function.

After training, the agent will have built a data structure $\{F, E\}$ where F is a set of features f_i with descriptors \mathbf{d}_i , expected bearings $\hat{\beta}_i$, and feature contexts N_i ; the subset of currently detected features F_i characterizes the position of the agent. R is a set of directed edges a_{ij} with direction labels α_{ij} indicating the direction of movement required to get from the place field of feature i into the place field of feature j . In addition to this stored data, a reference direction ν is maintained as a



Figure A1.6: Compass direction drift over a large explored area. The ν estimate may deviate substantially (over 90°) from its starting value, but remains locally consistent.

working memory and updated from the comparison of known and perceived feature bearings. This latter system models the head-direction systems of rodents and flies (Seelig and Jayaraman, 2015; Taube, 2007). Place recognition is based entirely on the recognition of features and the associated place fields.

A1.2.5 Pathfinding and Voting

The algorithm presented here is able to navigate the mapped environment by using graph search methods on the view graph. It is able to guide an agent to any user-selected goal location, as long as the environment has been explored sufficiently. The goal location is defined by a set of known features, and can for example be provided by an image depicting the goal. The algorithm then calculates multiple non-overlapping paths from features at the agent’s current position to the set of goal features, and uses a voting scheme to obtain navigable trajectories the agent can follow towards the goal location (Fig. A1.7).

In each pathfinding event, for one currently visible feature, the shortest path is found to one of the features in the goal set with Dijkstra’s algorithm (Dijkstra, 2022). Then, the nodes of that path are temporarily removed from the graph, except for the first and last nodes, and the search is repeated for another randomly selected pair of nodes. Due to the node removal, each path will have zero overlap with all previous paths. Still, when represented in a metric map, path trajectories will be similar due to overlapping place fields.

The search terminates when the pair of randomly selected start and goal nodes is




Figure A1.7: Route examples. Two different examples of path bundles (orange) from the agent’s current position (black cross) to a goal locations (set of green dots). The blue background shows the edges of the view graph, as described in Fig. A1.4.

unconnected in the graph lacking the temporarily removed nodes or when an upper limit has been reached. For example, in the tests detailed in the “Evaluation” section below, we used 30 successive Dijkstra searches, but the algorithm regularly found only a lower number of routes (~ 27), depending on the amount of exploration. Note that all edges are considered to have the same length, i.e., Dijkstra paths only differ in the number of edges they traverse.

Once a bundle of paths has been obtained, we could use the initial edge labels α_{ij} to determine the movement direction from the start location. Later however, it is not clear which step of each path applies at each position along the overall travel. Therefore, at each position, we determine the set of currently visible features also included in the present bundle. From each such feature we take the next edge along the respective path and thus obtain a set of movement votes (Fig. A1.8).

Each Dijkstra path p in the bundle P has an ordered set of edges $E_p = \{a_{ij}, a_{jk}, a_{kl}, \dots\}$ and set of nodes $F_p = \{f_i, f_j, f_k, f_l, \dots\}$. At navigation time, the set of visible features is F_t . We now consider the indices of all outgoing edges of currently visible features contained in a path of the bundle, $J_t = \{(i, j) | f_i \in F_t \wedge a_{ij} \in \bigcup_{p \in P} E_p\}$. The set of locally applicable motion directions is then given by $\alpha_{ij} | (i, j) \in J_t$. From these we obtain the movement consensus as

$$\bar{\alpha}_t = \text{cmean}_{J_t} \alpha_{ij}, \quad (\text{A1.5})$$

where cmean is the circular mean as defined in Eq. A1.2.

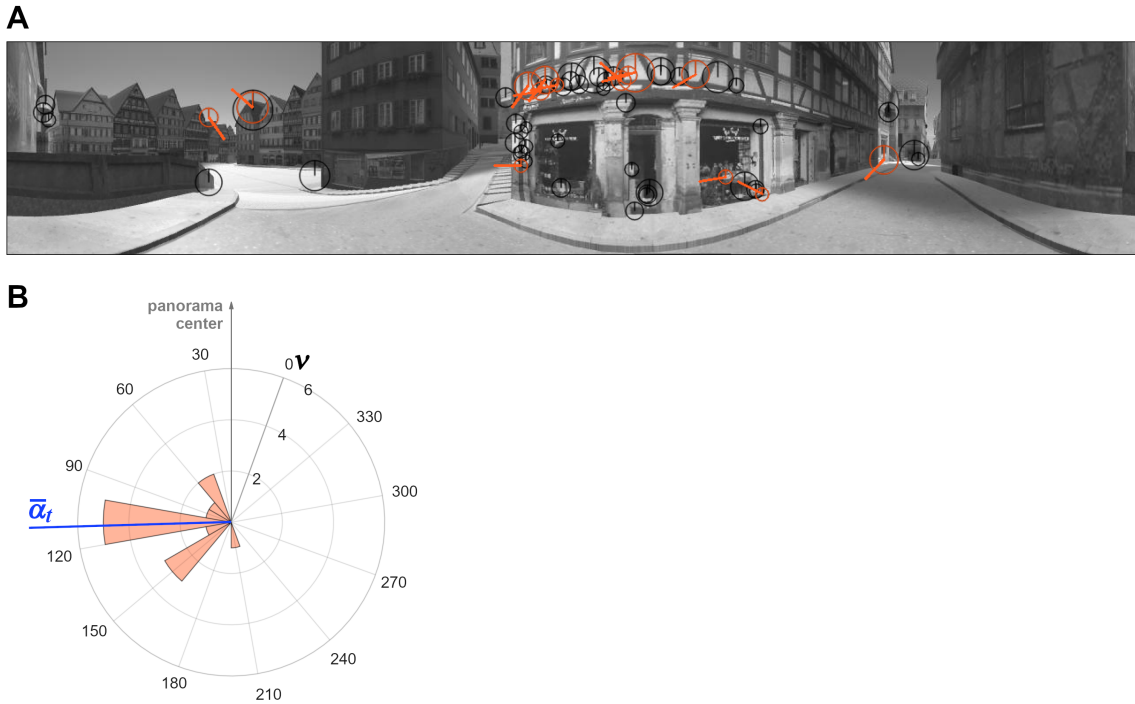


Figure A1.8: Direction voting. 360° panorama frame with detected SURF features (black and orange circles with vertical bar). During pathfinding, movement is derived from features that are also part of the path bundle (orange circles): The thick lines originating from the orange features show their respective movement direction vote relative to 0° (thin vertical stripe). The histogram shows the votes sorted into 10° bins and the mean direction $\bar{\alpha}_t$.

The final heading vector $\boldsymbol{\eta}_{t+1}$ is calculated with stiffness $\kappa \in [0, 1]$ as

$$\boldsymbol{\eta}_{t+1} = \kappa \boldsymbol{\eta}_t + (1 - \kappa)(\cos \bar{\alpha}_t, \sin \bar{\alpha}_t)^\top. \quad (\text{A1.6})$$

This results in a smoothing of the trajectory to reduce sway and reduce corner-cutting behavior; κ was set to 0.7.

Moving into the direction of $\boldsymbol{\eta}_{t+1}$ ideally leads the agent along a route specified by the bundle of paths, where it will continue to encounter labeled nodes. To facilitate this process, during path following, the number of features detected in each frame is doubled, which improves the odds of detecting labeled nodes. If the number of usable nodes, $|J_t|$, drops below a threshold of two, we assume that the agent has diverged from the path. In this case, a new bundle of Dijkstra paths is calculated with the current feature set F_t as a starting point.

A1.3 Evaluation

Our procedure differs in three important respects from standard approaches to graph-based navigation such as (Franz et al., 1998; Kuipers, 1978; Kuipers, 2000; Schölkopf and Mallot, 1995). First, the state of the agent is not characterized by a single node of the graph, but by a set of nodes which is unique for each confusion area. Second, path planning, i.e., the process of generating a path from the map, is not a single graph search but a two-step process involving a bundle of graph paths and a local voting scheme for direction. Third, we do not employ an explicit control law or homing scheme for approaching intermediate goals. Indeed, such intermediate goals are not explicitly used. Rather, the agent proceeds in small, “ballistic” steps.

The algorithm determines if the goal is reached by comparing the set of currently visible features, F_t , to the set of goal features, F_g , and considers the goal to be reached if $|F_t \cap F_g| / |F_g| \geq 0.35$. Note that there may be some offset between the agent’s final position and exact goal location since there is no optimization step in our ballistic procedure.

Relying on the average of a set of movement instructions solves the problem of wrong connections introduced into the graph due to aliasing, if enough alias-free paths are present. A crucial problem of mapping without global metric embedding and only relying on visual similarity is aliasing, the possibility that two features at distinct locations are confused. This can lead to the formation of wrong or impossible connections in the graph, which tend to be shorter than navigable connections (see

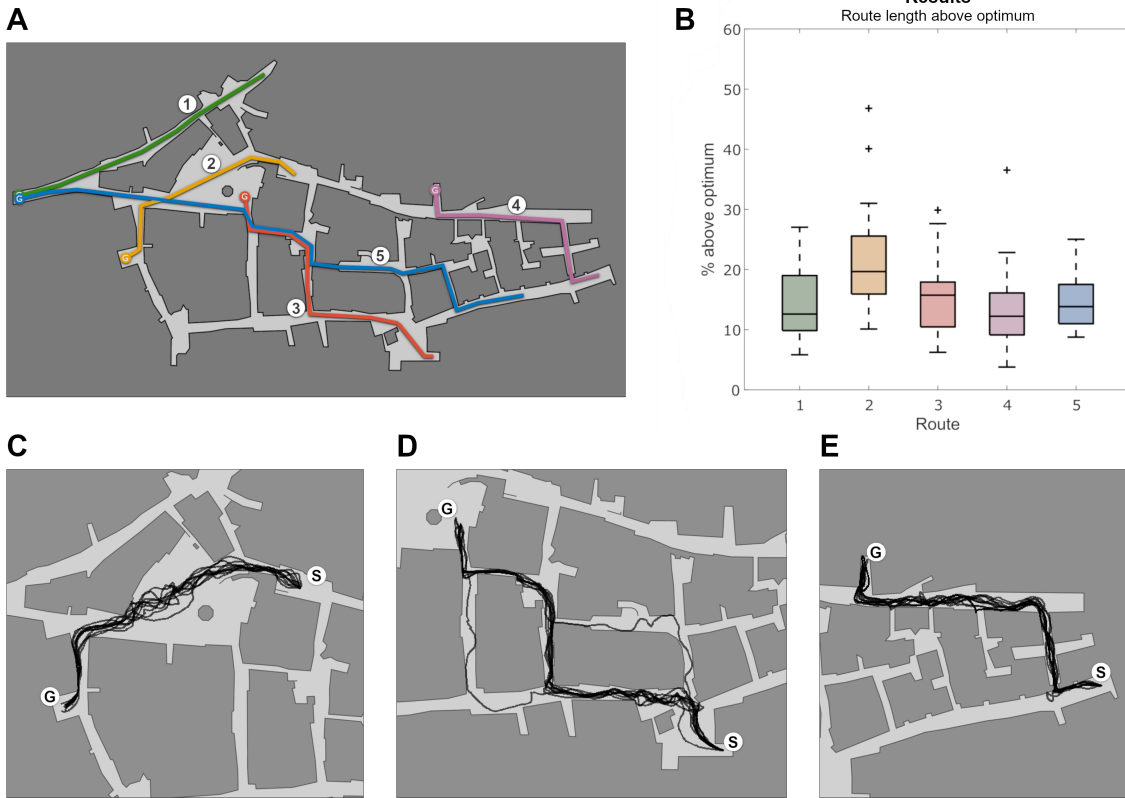


Figure A1.9: Performance in “Virtual Tübingen”. A: Map depicting the evaluation routes. B: Results of the evaluation. The box plot shows route length above an optimal trajectory, as depicted in A. Performance is somewhat worse for route 2, because it traverses a wide open area containing lots of distant landmarks, which are worse for exact navigation. C–E: Ten repetitions of each of the routes 2, 3 and 4. C: repetitions of route 2 show larger variations in open spaces. D: repetitions of route 3 show occasional choice of route alternatives as well as directional sway within single repetitions. E: repetitions of route 4 show the lowest variability among the evaluation routes.

Fig. A1.4). However, as long as a sufficient number of correct edges corresponding to navigable trajectories exist, the votes of the erroneous connections will not cause navigation to fail.

Finally, if the agent is unable to move during navigation, for example due to obstacles or being stuck in a corner, or if no consensus can be found in the set of movement instructions, the bundle of paths is recalculated, which has always solved the problem in our simulation. If the agent ever gets lost, for example because no known features are recognized, it may return to exploration behavior for a short while (e.g., random walk).

The algorithm was tested and evaluated in a virtual environment of the downtown area of Tübingen, Germany (3D model based on (Van Veen et al., 1998)), rendered in the Unity engine (Unity Technologies, 2018). The agent in the virtual environment was equipped with a 360° horizontal FoV and 60° vertical FoV camera projecting to a 1280×240 pixel image. Depending on location, the SURF feature detector would detect some 20 to 350 features per image, which were pruned to a maximum of 30 during exploration and 60 during path following.

When the agent had explored every street of the model at least once, exploration was terminated (compare Fig. A1.4). In the subsequent test phase, five pathfinding tasks were defined as start and goal views. Each task was repeated 20 times, and the traveled distance was measured and compared to that of the shortest possible route (Fig. A1.9A). The algorithm solved the tasks by first selecting features from the start view and estimating the reference direction ν from the features' bearings. Next, 30 Dijkstra paths were calculated, and from these, the overall trajectory was generated as described above.

The trajectories found for a single task may greatly differ between repetitions due to many stochastic influences, such as node selection for start and goal nodes and the noise added to the reference direction update. Further variation is introduced by numerical effects in collision detection and the latency between concurrent components of the programs running the algorithm and simulation. The algorithm may even guide the agent along different roads over multiple trials if they are close in length to the optimal route (see Fig. A1.9D).

The algorithm managed to successfully and efficiently guide the agent from start to goal in all trials with steady directional movement. On average, the agent's routes were only 10% to 20% longer than the optimal routes (Fig. A1.9B). The agent performed better, i.e., the routes were shorter, when they were leading mostly through roads and alleys rather than traversing large open spaces such as the market square.

The algorithm was able to guide the agent even with large drift in the reference direction ν of over 90° offset from the starting ν (see Fig. A1.6).

In conclusion, double population coding successfully combines the idea of spatial population coding with topological navigation by stimulus-response associations. It does not require highly processed input information but works with sequences of raw panoramic snapshots and basic feature extraction. Navigational performance is overall good. With respect to the question of parsimony, systematic work on the number of required features and graph links is required. In future work, the algorithm will be used to model human navigational performances such as the perception of local reference directions (Mou and McNamara, 2002), view voting (Mallot and Gillner, 2000), or local metric in spatial long-term memory (Warren, 2019).

Bibliography

- Baddeley, Bart et al. (2012). “A model of ant route navigation driven by scene familiarity”. *PLoS computational biology* 8.1, e1002336.
- Bay, Herbert et al. (2008). “Speeded-up robust features (SURF)”. *Computer vision and image understanding* 110.3, pp. 346–359.
- Bradski, Gary (2000). “The openCV library.” *Dr. Dobb’s Journal: Software Tools for the Professional Programmer* 25.11, pp. 120–123.
- Cartwright, BA and TS Collett (1982). “How honey bees use landmarks to guide their return to a food source”. *Nature* 295.5850, pp. 560–564.
- Cartwright, BA and TS Collett (1987). “Landmark maps for honeybees”. *Biological cybernetics* 57, pp. 85–93.
- Cheung, Allen and Robert Vickerstaff (2010). “Finding the way with a noisy brain”. *PLoS computational biology* 6.11, e1000992.
- Differt, Dario and Wolfgang Stürzl (2021). “A generalized multi-snapshot model for 3D homing and route following”. *Adaptive Behavior* 29.6, pp. 531–548.
- Dijkstra, Edsger W (2022). “A note on two problems in connexion with graphs”. *Edsger Wybe Dijkstra: His Life, Work, and Legacy*, pp. 287–290.
- Franz, Matthias O et al. (1998). “Learning View Graphs for Robot Navigation”. *Autonomous Robots* 5, pp. 111–125.
- Kuipers, Benjamin (1978). “Modeling spatial knowledge”. *Cognitive Science* 2.2, pp. 129–153.
- Kuipers, Benjamin (2000). “The spatial semantic hierarchy”. *Artificial intelligence* 119.1-2, pp. 191–233.
- Madl, Tamas et al. (2015). “Computational cognitive models of spatial memory in navigation space: A review”. *Neural Networks* 65, pp. 18–43.
- Mallot, Hanspeter A and Kai Basten (2009). “Embodied spatial cognition: Biological and artificial systems”. *Image and Vision Computing* 27.11, pp. 1658–1670.

- Mallot, Hanspeter A, Gerrit A Ecke, and Tristan Baumann (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17.
- Mallot, Hanspeter A and Sabine Gillner (2000). “Route navigating without place recognition: What is recognised in recognition-triggered responses?” *Perception* 29.1, pp. 43–55.
- Mou, Weimin and Timothy P McNamara (2002). “Intrinsic frames of reference in spatial memory.” *Journal of experimental psychology: learning, memory, and cognition* 28.1, p. 162.
- Muller, Robert U, Matt Stead, and Janos Pach (1996). “The hippocampus as a cognitive graph.” *The Journal of general physiology* 107.6, pp. 663–694.
- O’Keefe, John and Lynn Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- Schölkopf, Bernhard and Hanspeter A Mallot (1995). “View-based cognitive mapping and path planning”. *Adaptive Behavior* 3.3, pp. 311–348.
- Seelig, Johannes D and Vivek Jayaraman (2015). “Neural dynamics for landmark orientation and angular path integration”. *Nature* 521.7551, pp. 186–191.
- Smith, Lincoln et al. (2007). “Linked local navigation for visual route guidance”. *Adaptive Behavior* 15.3, pp. 257–271.
- Taube, Jeffrey S (2007). “The head direction signal: origins and sensory-motor integration”. *Annu. Rev. Neurosci.* 30, pp. 181–207.
- Tolman, Edward C (1932). *Purposive Behavior in Animals and Men*. New York: The Century Co.
- Trullier, Olivier et al. (1997). “Biologically based artificial navigation systems: Review and prospects”. *Progress in neurobiology* 51.5, pp. 483–544.
- Unity Technologies (2018). *Unity*. Version 2018.1.5f1. URL: <https://unity.com/>.
- Van Veen, Hendrik AHC et al. (1998). “Navigating through a virtual city: Using virtual reality technology to study human action and perception”. *Future Generation Computer Systems* 14.3-4, pp. 231–242.
- Warren, William H (2019). “Non-euclidean navigation”. *Journal of Experimental Biology* 222.Suppl_1, jeb187971.
- Wiener, Jan M et al. (2011). *Animal Navigation. A Synthesis*. Cambridge: The MIT Press, pp. 51–76.
- Wilson, Matthew A and Bruce L McNaughton (1993). “Dynamics of the hippocampal ensemble code for space”. *Science* 261.5124, pp. 1055–1058.

Appendix A2

Hierarchical Planning in Multilayered State-Action Networks

Authors: Matthias Brucklacher, Hanspeter A Mallot, and Tristan Baumann, 2021.

Abstract

The ability to decompose large tasks into smaller subtasks allows humans to solve complex problems step-by-step. To transfer this ability to an automated system, we propose a spiking neural network inspired by the neurobiological mechanics of spatial cognition to represent space on multiple levels of abstraction. As behavioral experiments suggest that humans integrate spatial knowledge in a graph of places, neurons in the state-action network encode locations while connections between them represent transition actions. In a series of simulation experiments, the influence of hierarchy on planning speed and on the resulting route choice in comparison to single-level models is investigated. We find that the model chooses biased subgoals in line with experiments on human navigation.

A2.1 Introduction

Evidence for the use of hierarchical representations for decomposition of complex tasks is ample and its advantages have frequently been discussed (Wiener and Mallot, 2003; Botvinick et al., 2009; Sutton et al., 1999). One approach to model it, is Hierarchical Reinforcement Learning (Botvinick et al., 2009; Sutton et al., 1999) (HRL). There, memorization of longer state-action sequences with known subgoals, so-called “options”, is introduced. The issue then shifts to how these options are

learned, termed the “option discovery problem” (Botvinick et al., 2009). In this work, a different approach is presented, in which the hierarchical structure of the environment is explicitly represented. Thereby, we address the shortcoming of mechanistic, neural-level models for hierarchical problem-solving. As suggested in previous work on graph-based navigation (Mallot et al., 2020; Kuipers, 2000), the state-action network proposed here is a set of environmental states (the neurons/nodes) connected by transition actions (the synapses/edges). It is then hierarchized by a layer of region neurons to represent extended areas, for example districts within a city, an island surrounded by water or even a whole country. Indeed, evidence for neurons encoding larger navigation segments has been found by Bos et al. (2017).

Human navigation in regionalized environments underlies two strong biases. In an experimental study by Wiener and Mallot (2003), subjects first explored a virtual environment and were then asked to navigate to a given landmark. The environments were divided into multiple regions either by clear boundaries (a river) or common landmark categories (animals, cars). When choosing among equally long paths to the goal, subjects preferred routes that crossed less region borders. Interestingly, they also showed a strong bias to access the region containing the goal as quickly as possible, neglecting alternative routes of equal length. In this paper, a similar setting is used to assess model behavior at region transitions, since these are critical points for hierarchical planning systems. We study interaction of the hierarchical network with a biologically plausible planning mechanism from Ponulak and Hopfield (2013) who applied it to single-level networks that lacked explicit action representation. This planning mechanism can be executed in parallel, in contrast to serial graph searches such as Dijkstra’s algorithm in Mallot et al. (2020). Simulation experiments firstly focus on the influence of hierarchical structure on planning time, as faster planning would be a major advantage for natural and artificial system, and secondly on route choice.

A2.2 Path-planning in neural state-action networks

Each neuron in the state-action network illustrated in Figure A2.1 encodes a memorized state in the environment. This can be a visual feature of an object in the environment, for example a landmark. When the agent recognizes this feature, the respective neuron receives a sensory input current. Together with the global in-

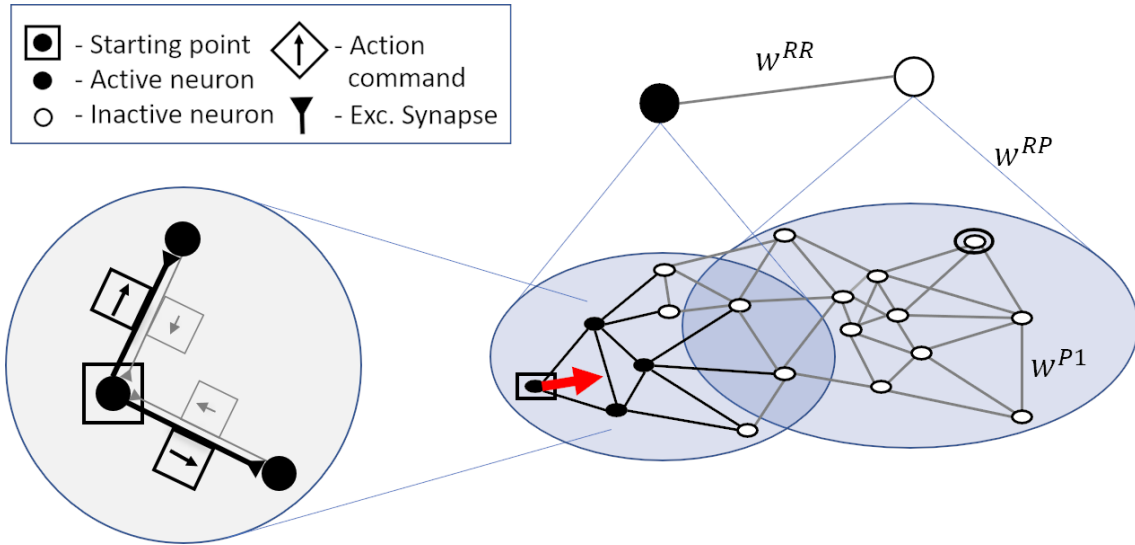


Figure A2.1: Components of the hierarchical state-action network. Left: Each synaptic connection between neurons is linked to a transition action that participates in voting when active. Strong synapses (bold) dominate the voting process and result in motion (red vector on the right) towards the goal marked by a circle. Strengthening of goal-directed connections is a result of the planning process described in the main text. Right: Region neurons (shown enlarged) represent extended areas in space. Each region neuron is bidirectionally connected to each lower-level neuron within its region (connections omitted for clarity).

hibitory current, the position of the agent is then encoded in a localized bump of activity (Ponulak and Hopfield, 2013), a distributed representation. Connections between neurons are bidirectional and represent transition vectors as shown in Figure A2.1. Since this work focuses on planning, we assume the network structure to be given and refer to Mallot et al. (2020) for unsupervised learning of state-action networks.

For path-planning, both sensory input and the global inhibitory system are turned off and the neurons corresponding to the goal location are activated. This leads to a wave of activity propagating throughout the network. As the wave passes, goal-directed connections are strengthened (Ponulak and Hopfield, 2013). This is achieved through the use of a time-dependent update rule for synaptic weights that requires the use of spiking neurons. As neuron k spikes after neuron j , the connection from k to j , w_{kj} , becomes stronger than the opposing connection from j to k :

$$\frac{dw_{kj}(t)}{dt} = \delta(s_k) \cdot A^- \cdot \left(1 - \exp\left(-\frac{s_j}{\tau_0}\right) \right) \exp\left(-\frac{s_j}{\tau_{STDP}}\right) \quad (\text{A2.1})$$

Here, A^- is a constant regulating the strength of the weight update, s_j and s_k denote the time since the last spike of neuron j and k respectively. $\tau_0 \ll \tau_{STDP}$ and τ_{STDP} are time constants. Evidence for such a mechanism that requires only locally available information has been found for instance by Roberts and Leen (2010). After planning, sensory input is turned on again and restores the activity bump at the agent’s current location (black neurons in Figure A2.1). The action to execute is then determined by a voting process, in which actions (vectors in squares in Figure A2.1) are weighted by both synaptic weights and presynaptic neuronal activity. Phrased differently, during navigation the agent follows a synaptic vector field (SVF) assigning a vectorial action to each position.

A2.3 Results

Increased planning speed To test the influence of hierarchical structure upon planning speed, a two-dimensional network of 24×24 neurons is set up. Each neuron is bidirectionally connected to neighbors one (with $w^{P1} = 6$, Figure A2.1) and two edges ($w^{P2} = 1$) away on the grid. Wavefront propagation in this single-level network is compared to hierarchical networks with an additional layer of 2×2 (respectively 3×3 and 4×4) neurons. These region neurons are connected to neighboring region neurons with $w^{RR} = 3$ and to the lower-level with $w^{PR} = 2$. Since it can be assumed

that regions overlap, neurons on borders are connected to region neurons of all adjacent regions (cf. Figure A2.1) with reduced synaptic weights to normalize input from the higher layer.

A wave of spiking activity is initiated by activating the goal neuron in the upper right corner of the network, creating a goal-directed SVF as it propagates (Figures A2.2a-c). In the hierarchical networks, the wave travels on both levels. As shown in Figure A2.2d, addition of the second layer drastically speeds up the planning process. Until all neurons are activated, the single-level network takes 55.6 ms, compared to 32.4 ms of the network with 2×2 regions (not shown), 35.6 ms (3×3) and 38.2 ms for the network with 4×4 regions. This equals a reduction of planning time by up to 42%.

Influence of hierarchy on route choice Following the design of an experiment on human navigation by Wiener and Mallot (2003), we set up an environment of discrete places arranged on a grid and distributed across two islands interconnected by bridges (Figure A2.3). As in Wiener and Mallot (2003), the agent is placed on a crossroad and must navigate to a given place on the opposing island. To determine the influence of hierarchical network structure, we compare two agents: one with a single-level representation of the environment and a second one with an additional region neuron for each of the two islands. Network connectivity follows the previously described simulation.

In the planning phase, the neuron representing the goal is activated. In the single-level network, the resulting activity wave spreads radially from the goal (Figure A2.3a), whereas in the hierarchical network, it propagates quicker inside the region containing the goal neuron (Figure A2.3b) and then slows down at the border. Again, wave propagation is faster in the hierarchical network. The resulting SVFs are oriented towards the goal, but in the hierarchical network, vectors within the left region tend to point stronger to the goal region than in the single-level network. Figures A2.3c) and d) show the resulting trajectories of an agent that is limited to discrete transitions between the places: these are shown as black horizontal or vertical lines whose thickness indicates the frequency with which each step was chosen. The grey band shows the average trajectory calculated from all node-to-node paths. The single-level model chooses paths leading straight to the goal. In contrast, the hierarchical model displays a preference to reach the goal region as quickly as possible as indicated by the bent grey line. Both models reliably reach the goal on an optimal (shortest possible) route: the single-level model in 88% of trials and the hierarchical model in 90%. Note that here all trajectories consisting of rightward

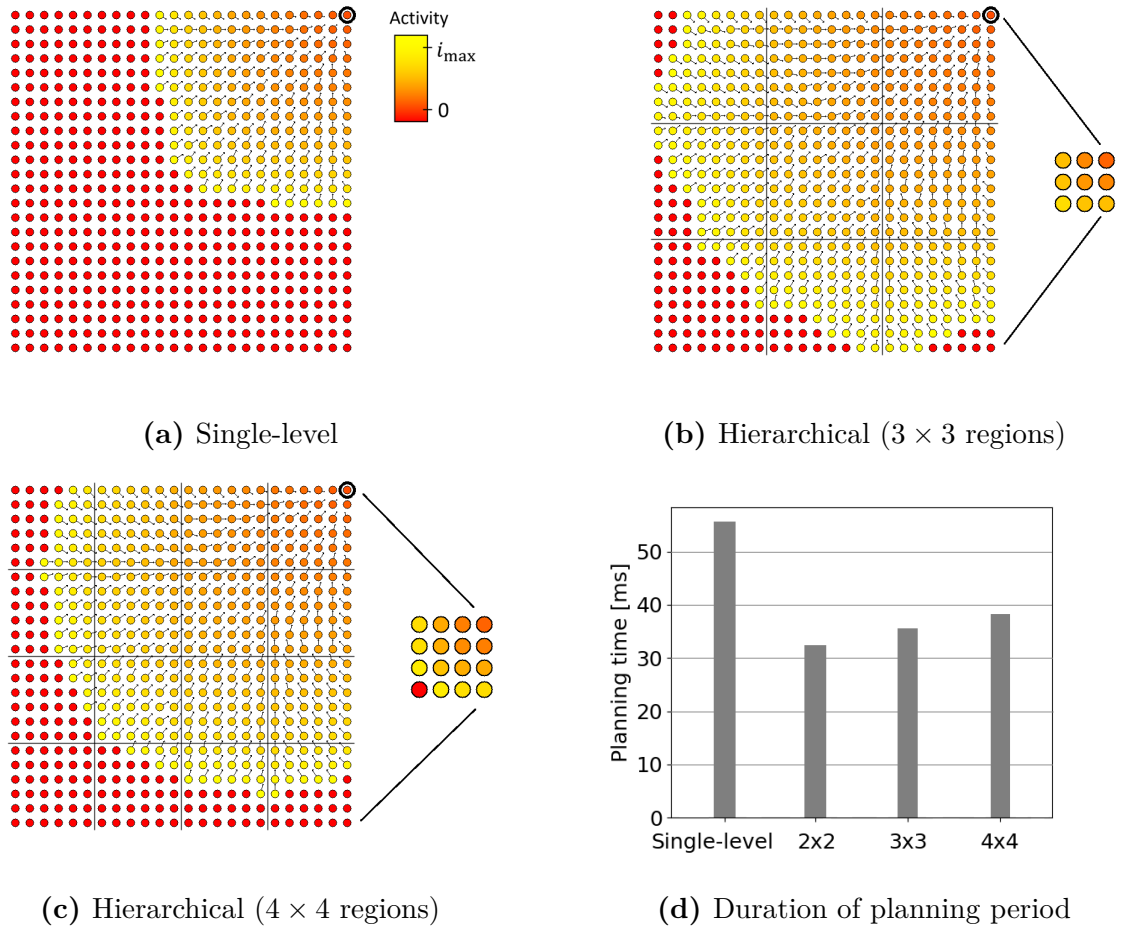


Figure A2.2: Influence of hierarchical structure on planning time. (a): Activity in a single-level network after 30 ms. (b) and (c): same as (a) for hierarchical networks of different region sizes, as illustrated by black lines. Regions neurons are shown in increased size. (d): Influence of region size on planning speed (2×2 corresponds to the largest regions).

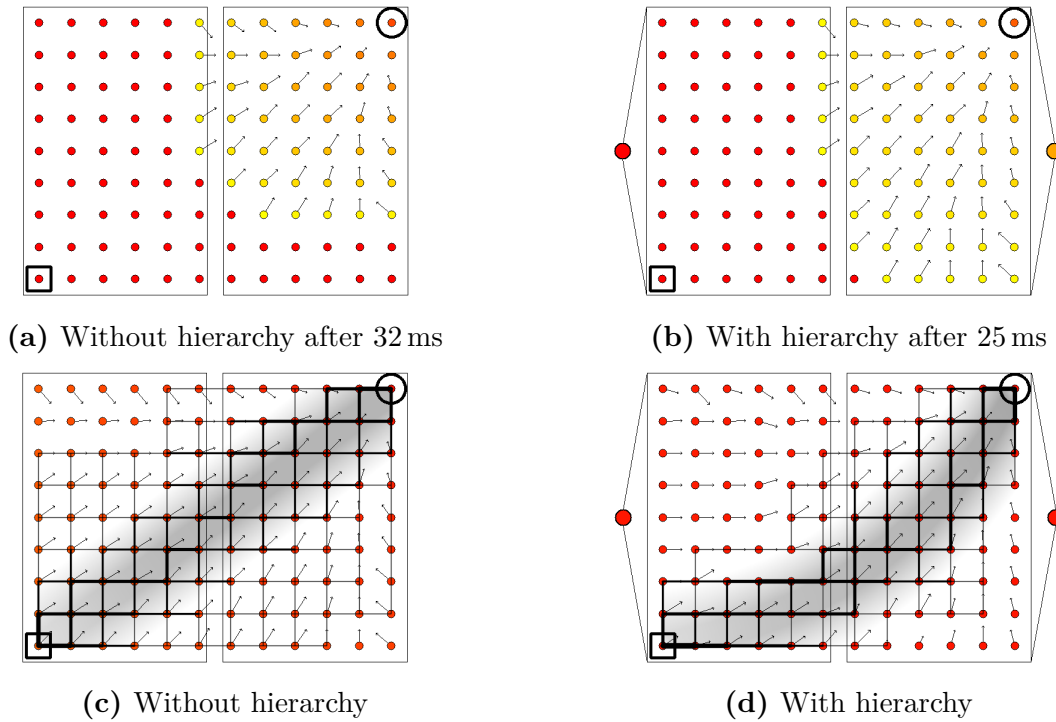


Figure A2.3: Influence of hierarchy on route choice. (a) and (b): Wavefront propagation from the goal marked by the circle in networks without and with hierarchical structure. For the latter, region neurons are shown in increased size. (c) and (d): SVF established by the wave and resulting trajectories. single-level model in 88% of trials and the hierarchical model in 90%. Note that here all trajectories consisting of rightward and upward steps in the figure have the same total length (are optimal).

and upward steps in the figure have the same total length (are optimal).

A2.4 Discussion

Addition of a hierarchically superordinate layer to the network leads to an increase in planning speed as well as behavioral effects in line with studies on human navigation. Compared to their single-level counterparts, the hierarchical networks took much less time for the planning process. The increase in wave velocity resulted from an activity wave in the higher layer that provided synaptic input to the wavefront on the lower level. Larger regions provide an advantage here (Figure A2.2d), but it can be assumed that there is an optimal granularity. As regions become larger, the activity wave on the higher level speeds up, and in the extreme case decays before it can provide input to the lower-level wave.

Simulation experiments on navigational behavior of the model show that a virtual

agent, equipped with neurons to represent regions, set the closest transition point to the goal region as a subgoal for navigation, disregarding equally short alternatives. This effect was caused by the altered shape of the propagating wavefront: Activation of a region neuron provided additional synaptic input to all neurons within that region. This input then led to a quick wave traversal of the goal region while transition to the neighboring region was more costly, slowing the wave down at the borders (Figure A2.3b). From there, it propagated perpendicularly to the borders – in contrast to the wavefront of the single-level networks that spread radially from the goal (Figure A2.3a). The resulting SVF of the hierarchical network was thus biased towards region transitions and with it the resulting path choices. This observation is in line with behavior of human subjects (Wiener and Mallot, 2003) that preferred to reach the goal region as quickly as possible, too. Here, model behavior is in contrast to HRL. If, in HRL, options are chosen purely based on the reward structure of the environment, all paths of equal length will be selected with equal probability, as can also be expected from breadth-first path searches. Compared to HRL, subgoal selection is much more straightforward in our model, where it is simply a result of altered wavefront propagation, whereas in HRL it leads to the option discovery problem (Botvinick et al., 2009). To allow for task-domains beyond navigation, states and actions in the present model would have to be generalized similarly to implementations of HRL (Sutton et al., 1999).

A2.5 Conclusion

The proposed model combines a multilayered state-action network with bio-inspired planning. When investigating model behavior, we found that it selects subgoals in line with human subjects. In combination with the model’s biological plausibility, this poses it as a candidate explanation for hierarchical navigation, although experimental evidence is not conclusive yet. Computationally, the parallel nature of the propagating wavefront in combination with an increased planning speed through hierarchical structure provides a strong advantage over serial graph searches. The network is thus an attractive solution for autonomous vehicles, captures the advantages of hierarchical representations, and can be implemented energy efficiently in spiking neuromorphic hardware.

Bibliography

- Bos, Jeroen J et al. (2017). “Perirhinal firing patterns are sustained across large spatial segments of the task environment”. *Nature communications* 8.1, p. 15602.
- Botvinick, Matthew M, Yael Niv, and Andrew G Barto (2009). “Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective”. *Cognition* 113.3, pp. 262–280.
- Brucklacher, Matthias, Hanspeter A Mallot, and Tristan Baumann (2021). “Hierarchical Planning in Multilayered State-Action Networks.” *ESANN 2021 proceedings*.
- Kuipers, Benjamin (2000). “The spatial semantic hierarchy”. *Artificial intelligence* 119.1-2, pp. 191–233.
- Mallot, Hanspeter A, Gerrit A Ecke, and Tristan Baumann (2020). “Dual Population Coding for Path Planning in Graphs with Overlapping Place Representations”. *Spatial Cognition XII*. Springer, pp. 3–17.
- Ponulak, Filip and John J Hopfield (2013). “Rapid, parallel path planning by propagating wavefronts of spiking neural activity”. *Frontiers in computational neuroscience* 7, p. 98.
- Roberts, Patrick D and Todd K Leen (2010). “Anti-hebbian spike-timing-dependent plasticity and adaptive sensory processing”. *Frontiers in computational neuroscience* 4, p. 156.
- Sutton, Richard S, Doina Precup, and Satinder Singh (1999). “Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning”. *Artificial intelligence* 112.1-2, pp. 181–211.
- Wiener, Jan M and Hanspeter A Mallot (2003). “‘Fine-to-coarse’ route planning and navigation in regionalized environments”. *Spatial cognition and computation* 3.4, pp. 331–358.

Appendix A3

Gateway identity and spatial remapping in a combined grid and place cell attractor

Authors: Tristan Baumann and Hanspeter A Mallot, 2023.

Abstract

The spatial specificities of hippocampal place cells, i.e., their firing fields, are subject to change if the rat enters a new compartment in the experimental maze. This effect is known as remapping. It cannot be explained from path integration (grid cell activity) and local sensory cues alone but requires additional knowledge of the different compartments in the form of context recognition at the gateways between them. Here we present a model for the hippocampal-entorhinal interplay in which the activity of place and grid cells follows a joint attractor dynamic. Place cells depend on the current grid cell activity but can also reset the grid cell activity in the remapping process. Remapping is triggered by the passage through a gateway. When this happens, a previously stored pattern of place cell activity associated with the gateway is reactivated from a “gateway database”. The joint attractor will then reinstate the grid cell pattern that was active when the gateway had first been learned and path integration can proceed from there. The model is tested with various mazes used in the experimental literature and reproduces the published results, and we make predictions for remapping in a new maze type. We propose the involvement of memory in the form of “gate cells” that drive the place cells and with them the joint hippocampal-entorhinal loop into the corresponding attractor whenever a compartment is entered.

A3.1 Introduction

The neural substrate of spatial representation, the *cognitive map*, is generally thought to be found in place cells located in the hippocampus (HC). First described by O’Keefe and Dostrovsky (1971), these cells fire maximally whenever the animal is within a small localized region in space, the cell’s place field (O’Keefe and Nadel, 1978; Moser et al., 2017). The environment is covered by the overlapping place fields of different place cells and the population activity can be used to decode the animal’s position and even future trajectories once the firing fields have been mapped (Wilson and McNaughton, 1993; Dragoi and Tonegawa, 2011; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015; Pfeiffer, 2020).

In the five decades since the discovery of place cells, many other cell types supporting spatial representation have been found, including (but not limited to) head direction cells (Taube et al., 1990a; Taube et al., 1990b), border cells (Solstad et al., 2008), and grid cells (Fyhn et al., 2004; Hafting et al., 2005). Especially grid cells have received much attention: Located in the medial entorhinal cortex (MEC) (Fyhn et al., 2004) and pre- and parasubiculum (Boccarda et al., 2010), these cells fire at regularly spaced intervals, tiling the environment into a regular hexagonal lattice. Grid cells are dorsoventrally organized into discrete modules of constant scale and orientation; between modules, scale is increasing in discrete steps (Hafting et al., 2005; Stensola et al., 2012). Within a module, grid phase and orientation between neighboring cells remains fixed (Fyhn et al., 2007). Due to their regular firing properties and the fact that the pattern persists in complete darkness, grid cells are generally thought to be involved in path integration (Hafting et al., 2005; Fyhn et al., 2007; Rowland et al., 2016; Grieves and Jeffery, 2017).

Grid cells are a common spatially modulated cell in the MEC (Hafting et al., 2005) and form a major input to place cells (Brun et al., 2008; Jeffery, 2011; Rowland et al., 2016). However, silencing input from MEC to CA1 does not eliminate place cell activity (Brun et al., 2008; Bush et al., 2014; Brandon et al., 2014) and in rat pups, place cells form before grid cells (Langston et al., 2010; Wills et al., 2012; Bush et al., 2014). Place cells show relatively stable place fields before the hexagonal grid cell pattern is fully developed, suggesting a strong input from allothetic or external cues such as landmarks or boundaries (Bjerknes et al., 2014; Muessig et al., 2015; Moser et al., 2015; Grieves et al., 2018).

This external input causes context-dependent firing changes: When the animal moves or is moved from one compartment to another, place cell specificities may

undergo complete changes to an independent pattern of firing fields. This process is known as “remapping” (Muller and Kubie, 1987; Lever et al., 2002; Leutgeb et al., 2005; Colgin et al., 2008; Julian et al., 2018): Place cells randomly rearrange their firing fields to new, uncorrelated locations, form additional firing fields or cease firing altogether. When the animal later returns or is returned to the original compartment, the place cells remap to their original pattern (Lever et al., 2002; Leutgeb et al., 2005). A given place cell therefore represents different places at different times or, more correctly, in different “spatial contexts”. Remapping also occurs in grid cells: When the context changes, the firing field locations of different modules may shift but grid scale and relative orientation remain unchanged (Fyhn et al., 2007; Cheng and Frank, 2011; Marozzi et al., 2015). Due to direct anatomical projections from MEC to the hippocampus and the finding that electrical stimulation of MEC grid cells causes place cell remapping (Kanter et al., 2017), grid cell realignment has previously been suggested as a possible candidate for triggering place cell remapping (Monaco and Abbott, 2011; Bush et al., 2014).

An important puzzle piece in understanding the remapping process is the nature of the “context change” by which it is triggered. Remapping can be observed when the animal moves to different compartments within a maze, or when the color and location of cues (Bostock et al., 1991) or the overall shape and texture of the compartment is changed (Lever et al., 2002; Leutgeb et al., 2005; Wills et al., 2005; Colgin et al., 2010; Jeffery, 2011). Non-spatial contexts may lead to a modulation of place cell firing rates (“rate remapping” (Latuske et al., 2018)), as has been shown for changes of the current task or goal (Allen et al., 2012) or in the dependence on recent history (Keinath et al., 2020). For a review of remapping types, see Latuske et al. (2018). Importantly, remapping seemingly does not occur if the context remains unchanged: For example, grid cell patterns are undisturbed within each compartment in Fyhn et al. (2007) and Derdikman et al. (2009) and Carpenter et al. (2015).

This suggests that the current compartment is equivalent to the current context, for as long as no remapping occurs, the rat must be within the same area. Therefore, remapping suggests a regionalization of the rat’s cognitive map, although the rat may not have direct access to that information. In this sense, the pattern of active cells provides both information about the current context and the position of the animal within that context (Latuske et al., 2018). Overall, the findings suggest that remapping is an active process that reflects the recruitment of a unique cognitive map from a “cognitive atlas” (Julian et al., 2018), a putative higher-level hierarchical

structure of spatial representation.

How can the hippocampal-entorhinal circuit acquire, maintain and reactivate a multitude of unique, context-dependent maps in a noisy infinite-context world? Some hints to this question may be hidden in a number of particularities of the remapping process:

1. Similar contexts lead to similar representations.

In visually or geometrically identical but otherwise separate compartments, place and grid cell firing patterns have been observed to repeat (Derdikman et al., 2009; Spiers et al., 2015; Carpenter et al., 2015; Grieves et al., 2016; Harland et al., 2017). That is, the place and grid cell activity at corresponding positions in different rooms is highly correlated. This similarity is behaviorally relevant: In a spatial memory task by Grieves et al. (2016), rats had to memorize odor-coded buried rewards in each of four identical rooms connected by a lateral corridor. The rooms were either arranged in parallel or radially around a curved corridor (Fig. A3.1). The animals learned the task faster and performed significantly better in the distinguishable radial condition, where place cells remapped to different patterns in each room, but failed in the parallel condition, where the patterns repeated (Grieves et al., 2016). Remapping also does not just reflect a direction or mode of movement: Derdikman et al. (2009) only observed remapping in a hairpin maze with actual physical walls, and not in rats that were trained to run comparable zig-zag lines in an open room.




Figure A3.1: Schematic example of findings by Grieves et al. (2016). The place fields of the same place cell are shown in two environments. Firing repeats at corresponding positions in each room in the parallel environment, but not when the rooms are arranged radially.

2. Entrance orientation is sufficient to distinguish contexts.

The superior performance in the radial condition is attributed to the head direction system which is assumed to maintain a global reference while the

animal explores the different compartments (Harland et al., 2017). If the orientation of the compartments is the same and they are entered from the same geocentric direction, the reference is also the same and the rooms are confused. On the other hand, if the compartments extend into different directions, such as in the radial arrangement in Fig. A3.1, the head direction reference may be used to detect this difference and the compartments can be distinguished. If the head direction system is silenced, this ability disappears and the radially arranged rooms also show the same local pattern (Harland et al., 2017). In agreement, the inverse case, where locally and globally identical compartments show different firing patterns, has also been reported (Skaggs and McNaughton, 1998; Tanila, 1999; Grieves et al., 2016; Harland et al., 2017; Keinath et al., 2020). Notably, Keinath et al. (2020) found that the most recent entrance influences the place cell firing rate in a compartment with multiple entrances, even if the firing field positions remain the same. Overall, the distinguishing cues seem to be entrance position and direction relative to the compartment, i.e., also orientation references.

3. Local cues only play a limited role.

Overall, the remapping dynamics suggest a strong bias towards local cues, i.e., the sensory and geometrical information of the current surroundings. These findings led to the formulation of border or boundary vector cell models of place cell firing (Hartley et al., 2000; Barry et al., 2006; Bush et al., 2014; Grieves et al., 2018): These cells fire whenever the animal is located at a preferred distance and angle from a nearby wall and thus reflect the local geometry. Based on the repeating patterns in identical compartments, the models suggest that geometric cues measured by boundary vector cells and mediated by the head direction system are the main determinants for place fields. However, boundary vector cells show no remapping. The cells have been reported to continue firing at the same fixed distance and allocentric direction from boundaries across multiple environments, regardless of context (Lever et al., 2009). This is exemplified in e.g., Spiers et al. (2015), where the change of wall color in one of multiple identically shaped parallel rooms caused a different place cell pattern to emerge in the changed room only. Therefore, remapping dynamics must occur at least partially independent of boundary vector cells, and models based on them require additional context (e.g., “contextual gating” in Grieves et al. (2018)).

4. Place and grid cells remap together.

Remapping appears to happen nearly instantaneously (Jezek et al., 2011) and both cell types remap concurrently (Fyhn et al., 2007) at the maxima of contextual difference, such as gateways between compartments or region transitions (Tanila, 1999; Leutgeb et al., 2005; Derdikman et al., 2009; Carpenter et al., 2015; Grieves et al., 2016). Therefore, the context-defining information must be available immediately at the transition. After remapping, patterns usually remain stable and no further remapping occurs as long as the local context remains the same.

Overall, the findings imply that gateways between compartments play a special role: Distinct place and grid cell patterns arise immediately if the compartments can be distinguished at the entrance. On the other hand, if they cannot be distinguished, the context is the same and the cell patterns reflect that similarity. In this sense, the context of the current environment would be based on the most recent gateway rather than just local cues.

Here, we propose a combined grid-and-place cell attractor model augmented with a place memory that is able to capture these remapping dynamics. The attractor is based on the observed connectivity loop of direct projections from MEC to HC (Brun et al., 2008; Jeffery, 2011; Monaco and Abbott, 2011; Rowland et al., 2016) and recurrent connections back from HC to MEC (Hafting et al., 2005; Boccara et al., 2010; Bush et al., 2014; Rennó-Costa and Tort, 2017): As in other models (Gaussier et al., 2007; Almeida et al., 2009; Cheng and Frank, 2011; Monaco and Abbott, 2011; Lyttle et al., 2013; Li et al., 2020), the activity of multiple grid cell modules is summed to form realistic place fields; in addition, we assume that place cells also connect to grid cells and may thus support the current attractor or drive the system into a new one.

Place cell remapping means that path integration is overruled at certain points, based on local position information at that point. This implies that the recognition of local position information is essential. We specifically look at natural region transitions in the form of gateways between rooms and introduce a mechanism called the “gateway identifier”, which represents context at entrances, including head direction. Each gateway has two gateway identifiers, one on each side.

To obtain repeating place and grid cell patterns in similar environments, the model has access to a “gateway database”, a memory that stores and reactivates place cell patterns associated with the different gateway identifiers. The model exploits

the ability of place cells to remap to arbitrary positions when directly stimulated (Diamantaki et al., 2018): Whenever a known gateway is passed, the associated pattern is reactivated and the place cell activity changes to it. Recurrent input from the place cells to grid cells then causes a shift in the grid cell firing fields until the matching attractor is reached.

Within a region, cell activity is based on path integration only. As such, firing fields are only determined by the local position information encountered at the last gateway and some additional noise. Importantly, entrances to ostensibly different regions may share the same gateway identifier if they share the same local position information, leading to the experimentally observed repetition of activity in similar compartments. Finally, when the system identifies an unknown entrance, it randomly remaps and the new activity pattern is added to the database.

Our model is related to models of working memory by Bouchacourt and Buschman (2019), where representations are maintained through recurring connections between a structured sensory layer (the grid cells) and a random, unstructured layer (the place cells). As a model of working memory, the model is inherently limited in its capacity once representations in the unstructured layer start to overlap. In this sense, the gateway database is a long-term memory that binds working memory patterns to context and is able to reinstantiate these patterns in a top-down manner.

In the following, we show that the model is able to replicate the remapping dynamics observed in multicompartment environments (e.g., Tanila, 1999; Fuhs et al., 2005; Carpenter et al., 2015; Grieves et al., 2016), and we propose a novel multicompartment setup (L-shaped rooms, Virtual environments) to confirm or refute the role of gateways as the main determinants for local context.

A3.2 Materials and Methods

See Fig. A3.2.

A3.2.1 Model overview

The model consists of an attractor neural network with fixed weights, simulating the cognitive map of a rat exploring a series of virtual environments. The cell activity is iteratively simulated in real-time; we denote the discrete time step of a single computational sequence as t .

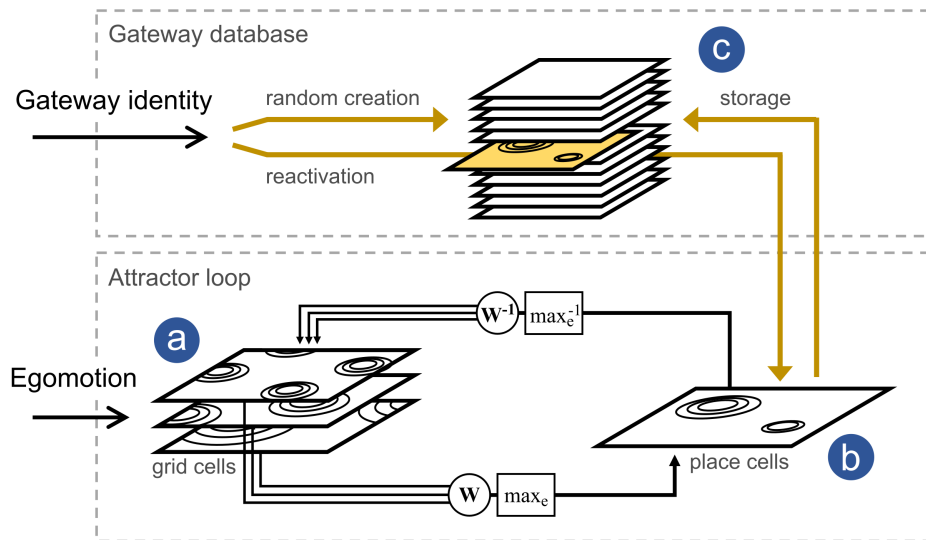


Figure A3.2: Model overview. **(a)** 3×100 grid cells are modeled by moving a stable activity peak over a toroidal surface. The cells form three layers with different input gains, resulting in hexagonal patterns with different scale and orientation. The circles on the layers sketch the firing fields, not the attractor peaks (which are the same for every layer). **(b)** Grid cell activity is combined in a weighted sum and filtered by an (invertible) winner-take-all scheme to form the place cell activity. The place cells in turn activate the grid cell module with the inverse operation, completing the larger attractor loop. **(c)** At region transitions such as room entrances, place cell patterns are memorized so that they may be recalled in the future (yellow arrows). When the virtual rat enters another region, place cell activity is either randomized or a corresponding place cell pattern is reactivated, followed by a shift of the grid cell firing peaks to the best matching attractor.

The core of the model is a grid cell attractor based on (Guanella et al., 2007). Input arrives in the form of an egomotion velocity vector, which shifts grid cell activity peaks to form the well-known hexagonal firing fields. Grid cells are arranged in multiple layers that differ in orientation and gain; from these, place cell activity is obtained by summing over the grid cells in a winner-take-all scheme. We adopted this modeling approach to place cell firing fields from previous studies (Almeida et al., 2009; Monaco and Abbott, 2011; Lyttle et al., 2013; Li et al., 2020) in order to keep the system simple. However, local position information can easily be added as another input to the place cell layer. The place cells then feed back into the grid cell module to complete a larger attractor loop (Fig. A3.2). To explore remapping dynamics, a predefined remapping signal is triggered whenever the virtual rat passes through a gateway and the corresponding gateway identifier is activated. It causes the place cells to either remap to a new, random pattern if the gateway is unknown, or to a previously memorized pattern if the gateway is recognized. By memorizing patterns at the gateways, compartments become associated with specific cell configurations, which remain active for as long as the rat does not move to another compartment.

To distinguish rotated but otherwise identical compartments, head direction information is also implicitly included, but simply assumed to be perfect in its compass-like function. I.e., the virtual rat is always aware of “true north”. This allows us to essentially ignore the virtual rat’s heading and treat the step-wise path integration as a straight-forward x, y -vector addition in a fixed reference frame. Throughout the paper, we use symbolic cardinal directions: North, east, south and west mean up, right, down and left in the figures.

A3.2.2 Grid Cells

Egomotion is processed by the grid cell module, consisting of 300 grid cells divided into three independent granularity layers with increasing scale and varying orientation. Three layers are required for realistic remapping based on independent shifts, as described in (Monaco and Abbott, 2011). In accordance with measurements by Stensola et al. (2012), the grid scale, i.e., the shortest distance between two activity peaks, was set to 40cm for the most fine-grained granularity layer and increased by a factor of 1.42 per layer to 57cm and 80cm for the coarser layers.

In the rodent brain, grid cells in the same granularity layer have been reported to share the same orientation (Stensola et al., 2012), are aligned to nearby walls, and

the orientation remains constant across different contexts (Krupic et al., 2015). In our model, grid orientation was fixed at 8.8° for layers 1 and 2 and 98.8° for layer 3, in accordance with the measurements by Krupic et al. (2015).

Within a layer, a stable activity peak is maintained by attractor dynamics (Guanella et al., 2007). $n^2 = 100$ cells form a topologically organized periodic sheet where neighboring cells have neighboring firing fields (Fig. A3.3a). The cells are connected to all others by short-range activation and long-range inhibition. In order to achieve the hexagonal geometry, the cells are arranged in a rhomboid with angles 60° and 120° , i.e., composed of two equilateral triangles. The rhomboid is the primitive cell of the hexagonal grid. The Cartesian coordinates \mathbf{x}_{ij} of grid cell ij within a rhomboid with lattice constant $1/n$ are

$$\mathbf{x}_{ij} = \frac{1}{2n} \begin{pmatrix} (i+j) \\ \sqrt{3}(j-i) \end{pmatrix} = \frac{1}{2n} \begin{pmatrix} 1 & 1 \\ -\sqrt{3} & \sqrt{3} \end{pmatrix} \begin{pmatrix} i \\ j \end{pmatrix} =: \mathbf{H} \begin{pmatrix} i \\ j \end{pmatrix}. \quad (\text{A3.1})$$

A 5×5 cell example is depicted in Fig. A3.3a. To derive an expression for the hexagonal periodic distance function, we pad a set of eight shifted rhomboids around the standard rhomboid by considering the shifts $\mathbf{s}_{pq} = \mathbf{H} (p, q)^\top$ for $(p, q) \in \{-n, 0, n\}^2$. For given Cartesian points \mathbf{x}, \mathbf{x}' the periodic distance function is

$$d(\mathbf{x}, \mathbf{x}') = \min_{(p,q) \in \{-n,0,n\}^2} \|\mathbf{x} - \mathbf{x}' + \mathbf{s}_{pq}\|, \quad (\text{A3.2})$$

where $\|\cdot\|$ is the Euclidean norm. This description of periodic boundary conditions in a rhomboid is equivalent to the twisted torus approach of Guanella et al. (2007), but can easily be generalized to an arbitrary number of grid cells.

We can now describe the activity of grid cell ij at time step t ,

$$g_{ij}(t) = \sum_{k=1}^n \sum_{l=1}^n g_{kl}(t-1) w_{kl}^{ij}, \quad (\text{A3.3})$$

where the weights w_{kl}^{ij} follow a Gaussian function of the periodic distance defined above,

$$w_{kl}^{ij} = I \exp\left(-\frac{d(\mathbf{x}_{ij}, \mathbf{x}_{kl} + \alpha \mathbf{R}_\beta \widehat{\mathbf{v}}(t))^2}{\sigma^2}\right) - T. \quad (\text{A3.4})$$

The intensity $I = 0.3$ defines the overall synaptic strength, $\sigma = 0.24$ is the width of the Gaussian and $T = 0.05$ is a global inhibition parameter. The scale of the grid is defined by the gain parameter $\alpha \in \{0.05, 0.035, 0.025\}$ corresponding to 40cm, 57cm

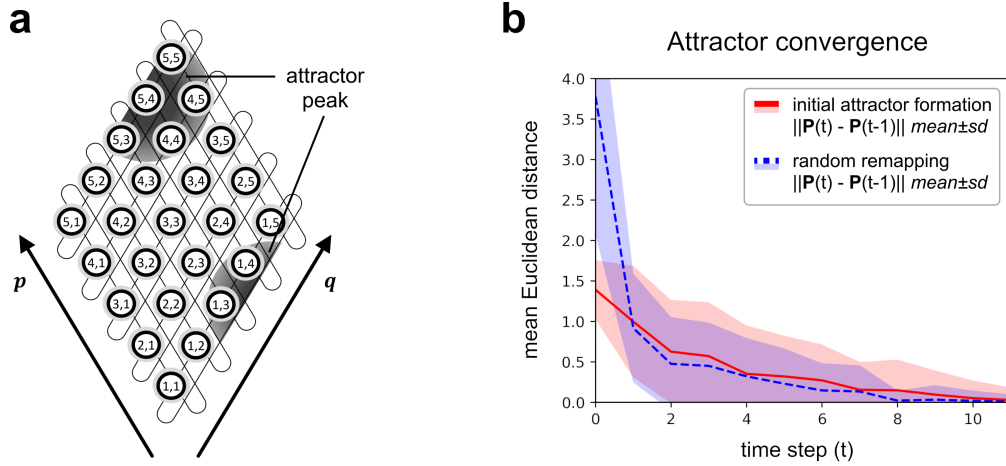


Figure A3.3: (a) Grid cell rhomboid (with only 5×5 cells for clarity). An activity peak (dark gray circle) is centered on cell $g_{5,4}$. The thin lines show periodic boundary conditions along grid cell indexes i, j . The vectors p, q correspond to the side length of the rhomboid and give the shift distance s_{pq} for eq. A3.2. (b) Convergence to an attractor when the model is initialized with zero place cell activity and random grid cell activity (red), or when the model remaps to a random place cell pattern (blue). Average of 100 repetitions.

and 80cm for the three layers. The orientation is given by the rotation matrix \mathbf{R}_β with $\beta \in \{8.8^\circ, 8.8^\circ, 98.8^\circ\}$, respectively. Initially, the cell activity $g_{ij}(t_0)$ is set to a random value drawn from a uniform distribution $\mathcal{U}(0, 0.1)$, but a stable peak is formed in the next time steps (initial attractor formation, Fig. A3.3b).

If the virtual rat moves with velocity $\mathbf{v}(t) = (v_1, v_2)^\top$, synaptic weights in direction of the egomotion estimate $\hat{\mathbf{v}}(t)$ increase, causing the activity peak on the grid cell layer to move by a corresponding amount. The egomotion estimate is obtained from the true velocity by adding multiplicative noise to both components,

$$\hat{\mathbf{v}}(t) = ((1 + \epsilon_1)v_1, (1 + \epsilon_2)v_2)^\top, \quad \epsilon_1, \epsilon_2 \sim \mathcal{U}(-0.25, 0.25). \quad (\text{A3.5})$$

In the sequel, rather than describing the dynamics of individual cells, we combine the activities of all three layers into a single 1×300 grid cell activity matrix

$$\mathbf{G}(t) = [g_1(t), \dots, g_{300}(t)]. \quad (\text{A3.6})$$

For example, the grid cell activity g_{221} now refers to g_{21} in granularity layer 3. This allows us to describe cell and layer-layer interactions in the form of matrix multiplications.

A3.2.3 Place Cells

Place cells form a single layer of 300 cells. In most situations, i.e., as long as no remapping occurs, the place cell activity $\mathbf{P}(t)$ is obtained by summing over grid cell activity with random weights:

$$\mathbf{P}(t) = \max_e (\mathbf{P}(t) \mathbf{W}_g^p), \quad (\text{A3.7})$$

where \mathbf{W}_g^p is a 300×300 synaptic weight matrix drawn from $\mathcal{U}(-1, 1)$. The weight matrix describes the excitatory and inhibitory input from grid to place cells. As in previous studies describing place fields as originating from random summation of grid cell inputs (Almeida et al., 2009; Monaco and Abbott, 2011; Lyttle et al., 2013), a winner-take-all nonlinearity \max_e was used to sparsify place cell activity:

$$y = \max_e(x) = \frac{1}{x_{\max}} \begin{cases} x & \text{if } x > (1 - e)x_{\max}, \\ 0.01 x & \text{otherwise.} \end{cases} \quad (\text{A3.8})$$

x_{\max} is the highest x in the input set and $e = 0.1$ is the percentage of most active cells whose activity is maintained, while all others are attenuated by a factor 0.01.

The winner-take-all scheme restricts cell activity to local maxima, resulting in place fields qualitatively similar to real place cell recordings. Importantly, by only attenuating cells below threshold rather than setting their activity to zero, the entire operation becomes invertible, i.e., it enables place cells to predict grid cells in turn. This property is required for the combined grid-and-place cell attractor loop and remapping dynamics described in the following.

A3.2.4 Combined attractor and inverse place cell - grid cell activation

By linking the place cell activity output back to the grid cells, an attractor is formed. It ensures that a specific place code will always match a specific grid cell arrangement. However, during normal exploration, the place cell input must not hinder the movement of the grid cell peak by the velocity vector. That is, the estimated activity of grid cells from place cell input, $\hat{\mathbf{G}}(t)$ should not differ substantially from the grid cell activity $\mathbf{G}(t)$. The required feedback connections from the place cell layer to the grid cell layers have been shown to exist in the form of recurrent MEC-HC connections (Hafting et al., 2005; Boccara et al., 2010; Bush et al., 2014; Rennó-

Costa and Tort, 2017) and could be learned by standard neural network techniques (Zhang et al., 1997; Linsker, 1997). For the purpose of this study we implemented a straight-forward inversion,

$$\widehat{\mathbf{G}}(t) = \max_e^{-1}(\mathbf{P}(t)) (\mathbf{W}_g^p)^{-1}, \quad (\text{A3.9})$$

where $(\mathbf{W}_g^p)^{-1}$ is the inverse or Moore-Penrose-pseudoinverse of \mathbf{W}_g^p from Eq. A3.7. In our simulations, the random matrices came out with full rank in about 99% of the trials. The nonlinearity $\max_e^{-1}(x)$ is the inverse operation of eq. A3.8,

$$x = \max_e^{-1}(y) = y_{\max} \begin{cases} y & \text{if } y > 1 - e, \\ 100 y & \text{otherwise.} \end{cases} \quad (\text{A3.10})$$

From eqs. A3.7 and A3.9, it is easy to see that $\widehat{\mathbf{G}}(t) \approx \mathbf{G}(t)$ except for a scaling factor. We can therefore stop the intrinsic dynamic of the grid cell module and restart it with the image of the place cell activity $\mathbf{G}(t) := \widehat{\mathbf{G}}(t)$. This will lead to a normal continuation of the grid cell dynamics, as long as $\widehat{\mathbf{G}}(t)$ falls in the basin of attraction of the current grid cell attractor. The intrinsic dynamics of the grid cell module will take care of remaining errors in the inversion process in the next few time steps. The convergence to the attractor after remapping is shown in Fig. A3.3b. The exact inversion of the forward weight matrix and nonlinearity is likely over-specified. The only requirement for a biological implementation would be that the inverse grid to place cell activation would fall into the same attractor basin. Since the cells are active simultaneously, it may be possible to learn the required connections via Hebbian dynamics. We opted for the matrix and nonlinearity inversion out of mathematical convenience and did not attempt to model biological neural dynamics such as synaptic plasticity or noise.

A3.2.5 Remapping and gateway database

Remapping occurs whenever the virtual rat moves from one compartment to another. At the transition, the model is provided a ‘‘gateway identifier’’ Γ , which, in our simulations, is simply a number corresponding to the identity of the local context of the newly entered compartment. The gateway identifier is an auxiliary construct that represents the necessary information that a biological organism would need to recognize a place, e.g., from vision, olfaction or proprioception (Cheng et al., 2013).

Importantly, the gateway identifier also includes head direction information, so that otherwise identical compartments may be distinguished by their orientation, but it does not include path integration or distance information, so that parallel, identical rooms cannot be distinguished by their relative position. In the virtual environments described below, entrances that look the same and are oriented in the same direction share the same gateway identifier (see Fig. A3.4). In the hippocampal-entorhinal network, gate information could be modeled by an additional “gate cell” connected to the place cell layer (see Discussion). This hypothetical cell would activate a specific place cell pattern whenever the room was entered through the corresponding gateway.

In the current implementation of the model, the newly encountered gateway identifier Γ is used to index the gateway database $D = \mathbf{P}_0, \dots, \mathbf{P}_n$. If the database has an entry $\mathbf{P}_\Gamma \in D$, the stored place cell pattern is reactivated by setting $\mathbf{P}(t) := \mathbf{P}_\Gamma$. With the attractor rules described in eq. A3.9, the grid cell attractor will shift to the attractor matching the reactivated place cell pattern. As a side effect, accumulated path integrator noise is also reset when a known gateway is passed.

This mechanism may generate errors at wide gateways that can be passed with different sideways offsets. This problem could be avoided by assuming that the rat shows strict centering behavior at gateways. Here, we measure the sideways offset and add it to the displacement vector $\hat{\mathbf{v}}$ in eq. A3.4. In biology, offset information could for example be provided by border or boundary vector cells, or depth perception.

Resetting the place cell activity and the subsequent convergence of the network to the new grid-and-place cell attractor concludes the remapping process. Grid and place cells have completely changed their firing fields, but the change is consistent in that the same pattern will be activated whenever a given gateway is encountered. Therefore, similar compartments that share the same gateway identifiers will show the same place and grid cell patterns.

The gateway database models the rat’s overall knowledge of the compartments; it is a cognitive atlas in the sense of Julian et al. (2018) that encompasses all context-specific maps. Initially, the gateway database is empty and the model is initialized with random activity. When the rat learns a novel gateway identifier, the current place cell activity pattern is stored. This may happen in three different cases:

1. When entering an unknown room, a novel view will appear. A new gateway identifier is encountered and a random pattern of place cell activity is gener-

ated and stored in the database together with the identifier.

2. When entering a known room, the gateway identifier is recognized and the associated place cell activity pattern is loaded into the place cell layer.
3. New gateway identifiers can also be stored when leaving a room such that the room can be recognized when it is later reentered through the same gate. In order to model direction-specific firing fields observed in linear track environments (McNaughton et al., 1983; Derdikman et al., 2009), this feature of the model is applied only in wide rooms in which the rat can turn around, but not in narrow corridors.

Note that random weight matrices and random initial place cell patterns were chosen as an assumption-free approach to remapping. A more realistic approach would incorporate some sensory information about the compartment, to allow for graded similarity between contexts: Grieves et al. (2016) for example found a decrease in correlation between rotated rooms proportional to the angle between them. Another alternative to random remapping would be to continue the currently active pattern over region boundaries and show no remapping at all. In that case, grid patterns would continue undisturbed across regions.

A3.2.6 Simulation and exploration procedures

For fast parallel computing of the cell matrices, the model was implemented in the machine learning interface TensorFlow (Martín Abadi et al., 2015) running in real-time on a normal notebook PC (CPU: Intel i7-10750H, GPU: NVIDIA GeForce RTX 2070). The model was evaluated in three repetitive multicompartment environments, realized as virtual 3D models in Unity (Unity Technologies, 2021).

The 3D world was only used for collision detection and exploration, not visual input. The gateway identifiers were predefined at locations that match remapping events in experimental findings, and visually similar entrances were assigned the same identifier. All walls had the same thickness (2 cm) and height (11 cm). The environments are described in detail below.

An agent, the “virtual rat”, represented by a 4.8 cm diameter cylinder, explored the environments either by random walking (parallel rooms and L-shaped rooms) or guided by the experimenter (Hairpin maze). The virtual rat moved at a constant speed of 14 cm/s. For the random walk, at each time step t , the movement direction




Figure A3.4: (a) - (d) Virtual environment floor plans used in the simulations. The red dotted lines mark region transitions and the adjacent integers the gateway identifier Γ : Equal numbers indicate the same local context and therefore cause remapping to the same patterns. The asterisk (*) marks the starting position. (b), the hairpin maze, has two starts based on traversal direction. Note that in (c), the gateway identifier of both rooms is the same, $\Gamma = 1$, but in (d), it is different: $\Gamma = 1$ for the southwest and $\Gamma = 4$ for the northeast room.

was rotated by a random angle drawn from $\mathcal{U}(-10^\circ, 10^\circ)$. To avoid exploration deadlocks such as continued movement into a corner, the random walk trajectory reflected off of walls. The hairpin maze environment was explored by hand to obtain complete traversals which are unlikely to result from random walking. Position data (cylinder midpoint) and cell activities were sampled once per time step, which had a duration of approximately 0.1 s.

A3.2.7 Virtual environments

Parallel rooms (Fig. A3.4a). The layout is based on two- or four-room maze experiments where identical rooms are arranged in parallel and are either entered from the same direction (Skaggs and McNaughton, 1998; Fuhs et al., 2005; Spiers et al., 2015; Grieves et al., 2016) or from different directions (Tanila, 1999; Grieves et al., 2016). We combined both layouts in a single two-by-two rooms environment where two sets of parallel rooms are placed on the opposite sides of a central corridor. The rooms are locally indistinguishable except for the position and direction of the entrance. In the figure, the two rooms entered from the north have gateway identifier 1 and the rooms entered from the south gateway identifier 6. The four identifiers of the corridor are all distinguishable by the combination of view and heading.

Hairpin maze (Fig. A3.4b). We created a virtual hairpin maze similar to Derdikman et al. (2009). In hairpin mazes, cell activity has been observed to repeat every second arm and to depend on running direction (Derdikman et al., 2009). Accord-

ingly, we defined the turns between arms as region transitions, with every second arm sharing the same gateway identifier. As described above, place cell patterns in corridors are only memorized when the corridor is entered, and not when it is left. That is, entrances and exits depend on the traversal direction, which will lead to direction-selective firing fields. The maze was traversed on a guided path from the start to the end in separate eastbound and westbound trials.

L-shaped rooms (Fig. A3.4c,d). This novel layout consisted of two adjacent identical L-shaped rooms connected by a single diagonal hallway. In the “same context” condition (Fig. A3.4c), both rooms are entered at the same position (south-east corner) and with the same heading. The rooms are therefore locally indistinguishable and share the same gateway identifier. In the “different context” condition (Fig. A3.4d), one entrance was rotated by 90° so that the rooms are entered with a different heading. Previous studies have suggested that entry direction alone may be sufficient for rats to distinguish otherwise identical contexts (Tanila, 1999; Grieves et al., 2016; Keinath et al., 2020). However, in these experiments, entrance position and bearing remained visible from the entire room and could have been used as landmarks. Models reflecting local cues alone would therefore predict the same activity in the far part of the room in both conditions. In the L-shaped rooms, the entrance is not visible from the far part of the room. Therefore, if contextual information can be maintained in absence of context-defining cues, cell patterns or behavioral measures should differ even if the rooms are completely identical.

A3.2.8 Data analysis

For each environment, place and grid cell activity maps were produced by sorting the localized activity of each cell into 2×2 cm bins, averaged over all time steps. Activity maps were smoothed by a 3×3 bin Gaussian kernel. Due to the radius of the virtual rat and the wall thickness, activities in different compartments were at least three bins apart, and smoothing did not extend across walls.

To assess firing field repetition, Pearson correlations were calculated between different regions of interest on the same cell activity map. Only same-size regions were compared; we did not attempt to correlate scaled cell activities. To account for random variations, each environment was explored ten times with different random number seeds, and the correlations were averaged. Place and grid cell results are reported separately.

Additionally, as a means of confirming the regions of interest as correlation max-

ima, we created autocorrelograms by correlating activity maps with bin-wise shifted versions of themselves.

A3.3 Results

In the following, we report model predictions for a series of environments, in the form of binned place and grid cell activities. In the model, the cells form a combined attractor where the activity of one cell type directly influences the other. Therefore, correlation results for the different cell types, although reported separately, are highly similar.

A3.3.1 Parallel rooms (Fig. A3.5)

We compared place and grid cell activity of matching bins for the four rooms (numbers 1 to 4 in fig. A3.5a) combined into parallel (west: rooms 1 and 3 versus east: 2 and 4), opposed (north: 3 and 4 versus south: 1 and 2) and rotated opposed (north: 3 and 4, rotated 180°, versus south: 1 and 2), rotated by 180° to match entrance direction. All place cells in all repetitions showed space-dependent activity patterns. Consistent with findings from animal research (Tanila, 1999; Fuhs et al., 2005; Spiers et al., 2015; Carpenter et al., 2015; Grieves et al., 2016), we observed highly repetitive activity for the parallel condition (mean east/west correlation 0.75, 95% confidence interval (CI) [0.69, 0.79]), but not in the opposed condition (north/south 0.01, 95% CI [-0.10,0.12]) or rotated opposed condition (north rotated/south -0.01, 95% CI [-0.13,0.10]). Parallel room correlations were significantly higher than those of opposed rooms ($u = 0.0, p < 0.0001$, Mann-Whitney-U-Test, MWU) and rotated opposed rooms ($u = 0.0, p < 0.0001$, MWU).

Grid cell correlations match the findings from place cells. Activity similarly repeated in the parallel condition (east/west 0.84, 95% CI [0.81, 0.87]) and was uncorrelated in the opposed condition (north/south 0.01, 95% CI [-0.10,0.12]) and rotated opposed condition (north rotated/south -0.02, 95% CI [-0.13,0.09]).

Like in the place cells, the parallel room correlations were significantly higher than those of opposed rooms ($u = 0.0, p < 0.0001$, MWU) and rotated opposed rooms ($u = 0.0, p < 0.0001$, MWU), indicating that the difference in activity in the opposed rooms is not just caused by the room orientation. Overall, the results suggest that remapping at room entrances is sufficient to decorrelate cell patterns for different contexts while still producing repetitive activity under similar contexts.

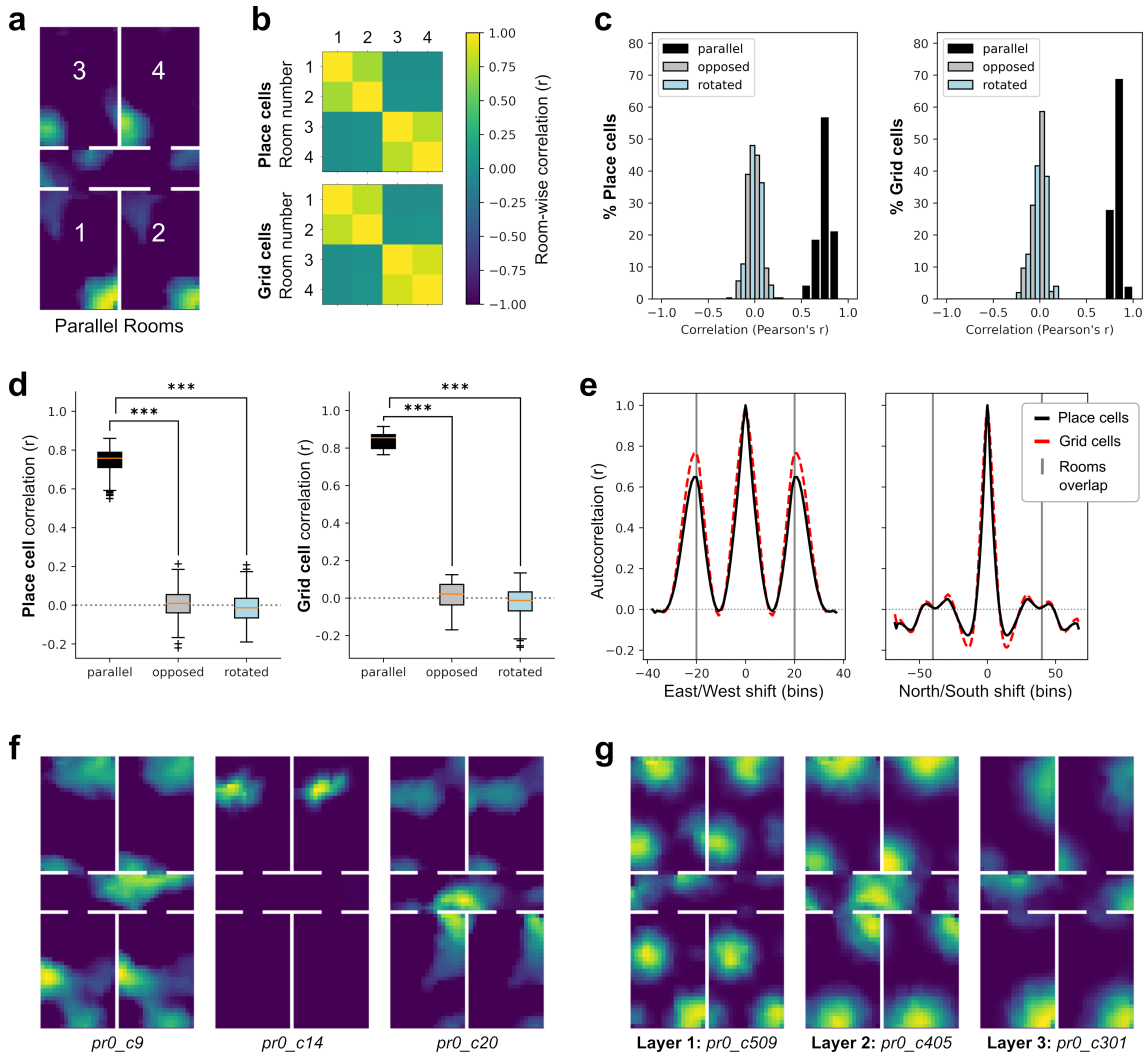


Figure A3.5: Parallel rooms. (a) Example of place cell activity in the different rooms (1-4) and the corridor (middle). (b) Average room-wise correlations for place and grid cells. As expected, the correlation between parallel rooms is high and the correlation between opposed rooms low. (c, d) Place cell (left) and grid cell (right) correlation distribution between parallel, opposed, and rotated opposed rooms. In agreement with the average room-wise correlation matrix in (a), cell correlations cluster around 0 for the opposed condition. The distributions are relatively narrow and don't overlap. (e) Linear autocorrelogram of all cells shifted in bin-wise increments, including activity in the corridor. The vertical gray lines signal shifts where the rooms line up. (f) Examples of 3 place cells, expressing similar place fields in parallel rooms but different place fields in opposed rooms. (g) The same can be seen in grid cell firing fields. The grid structure in the larger layers is not visible because the rooms are shorter than the grid scale. The codes refer to the model repetition (pr0-pr9) and cell identity (0-599).

A3.3.2 Hairpin maze (Fig A3.6)

In the hairpin maze, firing fields have been shown to depend both on the direction of movement and arm identity (Derdikman et al., 2009). Trajectories (Fig. A3.6a) were split into *eastbound* (from arm 1 to arm 10) and *westbound* (from arm 10 to arm 1). The agent entered the initial arm and was guided through the maze to the end of the last arm, where it exited the maze. It then reentered the same arm for a separate trajectory in the opposite direction. For each of the ten random initializations of the model, the maze was traversed twice in either direction for a total of 40 traversals.

Results depended on the overall traversal directions (*eastbound* vs. *westbound*), as well as the local directions within each arm. Remapping occurred at the apices between the arms (Fig. A3.4b). In agreement with animal studies (Derdikman et al., 2009), model place and grid cell firing fields repeated in every second arm.

A clear checkerboard pattern (Fig. A3.6b) emerged from the correlation of place cells in different arms. Excluding the outer arms, arms traversed in the same direction (i.e., odd- or even-numbered arms) showed a significantly higher correlation (eastbound 0.90, 95% CI [0.88, 0.92]; westbound 0.92, 95% CI [0.90, 0.93]) than arms traversed in a different direction (eastbound -0.01, 95% CI [-0.12, 0.11]; westbound -0.02, 95% CI [-0.13, 0.09]; Same direction vs. alternating: eastbound $u = 0.0, p < 0.0001$, westbound $u = 0.0, p < 0.0001$, MWU). Firing fields in the eastbound and westbound traversals were uncorrelated (0.01, 95% CI [-0.11, 0.12]).

Grid cell results were generally similar to place cells. Grid cell activity was also highly correlated in same-direction arms (eastbound 0.92, 95% CI [0.90, 0.94]; westbound 0.94, 95% CI [0.92, 0.95]) but not in alternating arms (eastbound -0.03, 95% CI [-0.14, 0.09]; westbound -0.03, 95% CI [-0.14, 0.08]; Same direction vs. alternating: eastbound $u = 0.0, p < 0.0001$, westbound $u = 0.0, p < 0.0001$, MWU). Again, overall eastbound and westbound traversal correlations were uncorrelated (0.01, 95% CI [-0.11, 0.12]).

Correlation between the outer and inner arms was low on average (eastbound 0.14 95% CI [0.03,0.25], westbound 0.24 95% CI [0.13,0.34]). In the model, non-zero correlations between the outer and inner arms are a result of random remapping and noise, but are not indicative of any corridor similarities. The outer arms were treated as separate regions with their own gateway identifiers. In a biological animal, low but non-zero correlation might reflect a degree of similarity or confusion between the outer and inner arms, but this was not modeled.

In general, the matching correlations in the hairpin maze are higher and the

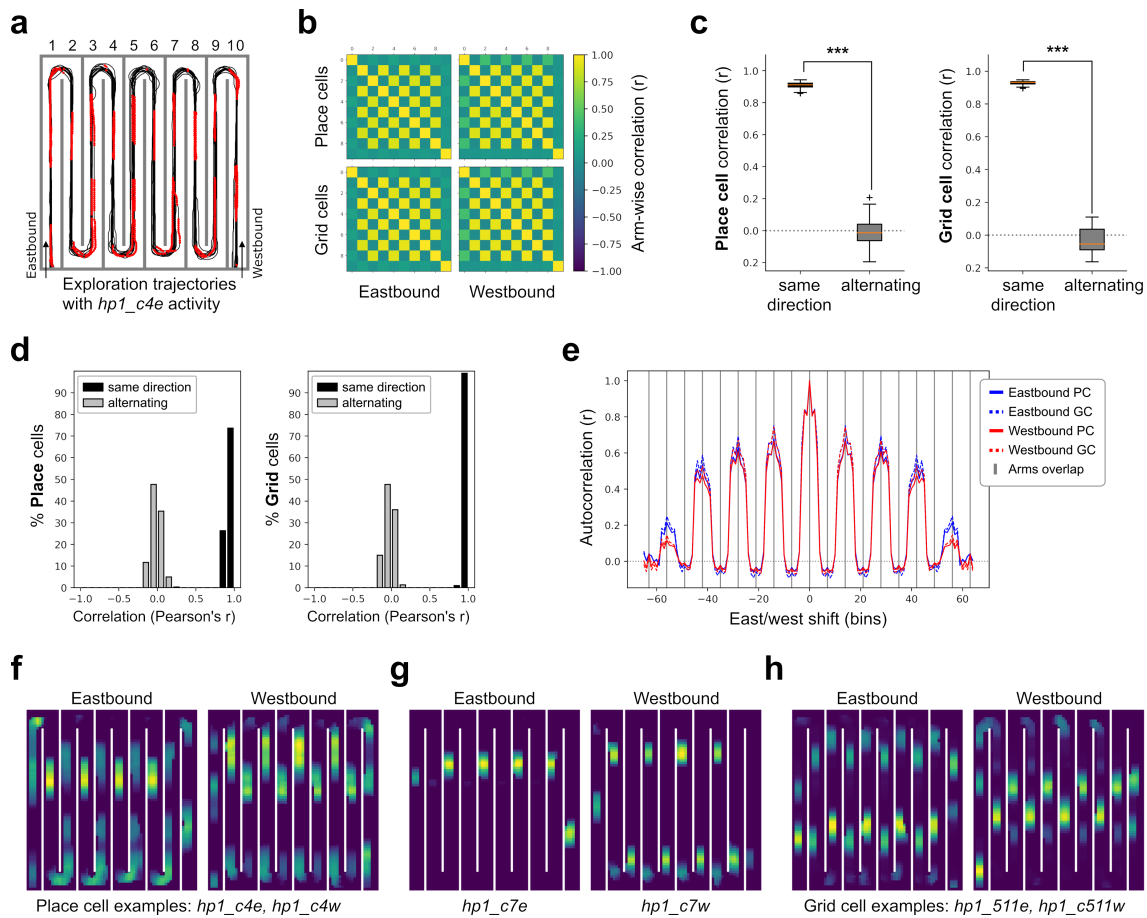


Figure A3.6: Hairpin Maze. **(a)** Exploration of the maze. The red dots mark eastbound cell activity of the cell in (f), where activity $> 10\%$ of max. **(b)** Arm-wise place cell (top) and grid cell (bottom) correlation for eastbound (left) and westbound (right) traversals. **(c, d)** Place cell (left) and grid cell (right) correlation distribution for the inner arms (arms 2 to 9), traversed in the same or different directions. The plots show the combined eastbound and westbound data. **(e)** Linear autocorrelation of all cells, shifted in one-bin increments. The autocorrelation peaks whenever the activity is shifted by two arms. **(f,g)** Example place fields of the same cell for eastbound and westbound traversals. The cells show the characteristic repeating patterns. **(h)** Same as (f,g) but for a grid cell from the smallest layer. The grid pattern is obscured by the narrow maze arms.

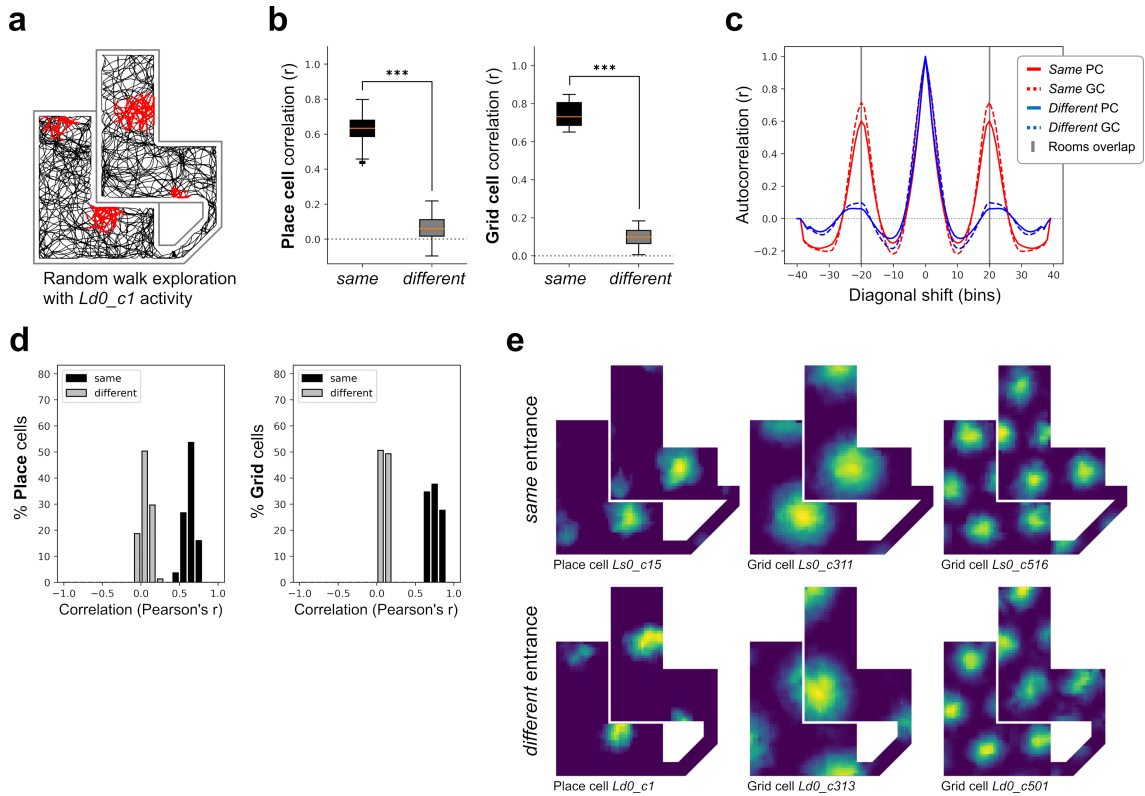


Figure A3.7: L-shaped rooms. (a) Random walk exploration in the *different* entrance condition. The red dots mark eastbound cell activity of the lower left cell in (e), where activity $> 10\%$ of max. (b) Place cell (left) and grid cell (right) correlation for *same* and *different* entrances. (c) Linear autocorrelogram of all cells, shifted diagonally in one-bin increments. The autocorrelation peaks in the *same* condition where the rooms overlap. The autocorrelogram includes firing information from the corridor, which was excluded from the other comparisons. (d) Cell correlations as in (b). (e) Firing field examples, from left to right in both conditions: Place cell, grid cell from the largest granularity layer, grid cell from the smallest granularity layer.

confidence interval is narrower compared to the other two environments. This is likely a result of the manual exploration with frequent region transitions, which caused accumulated path integrator noise to be reset more often.

A3.3.3 L-shaped rooms (Fig. A3.7)

The L-shaped rooms environment was explored in two entrance conditions, *same*, where both rooms were entered from the same direction, and *different*, where both rooms were entered from different directions. Note that the conditions were explored by separate, unpaired models, i.e., cells are not matched between conditions.

As in the other environments with repeating activity, place cell correlation (Fig.

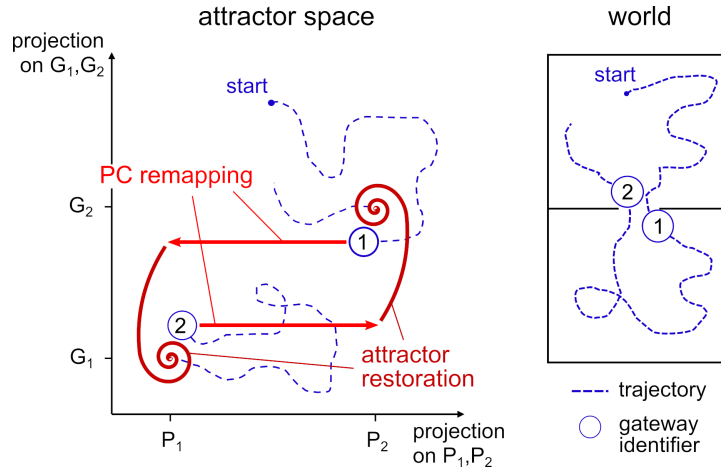


Figure A3.8: Schematic phase diagram of the remapping process. Left: 2-dimensional projection of the joint state space of place and grid cells for two attractors P_1, G_1 and P_2, G_2 . The projection is spanned by the straight lines through P_1, P_2 (x-axis) and through G_1, G_2 (y-axis). Right: Agent trajectory in a two-room maze. For further explanation, see text.

A3.7b) was significantly higher in the *same* condition (0.63, 95% CI [0.56, 0.69]) compared to the *different* condition (0.06, 95% CI [-0.05, 0.17], Place cell *same* vs. *different*: $u = 0.0, p < 0.0001$, MWU).

For grid cells, correlations in *same* (0.74, 95% CI [0.69, 0.79]) were higher than between place cells but similarly uncorrelated in the *different* condition (0.10, 95% CI [-0.01, 0.21]) (Statistic: $u = 0.0, p < 0.0001$, MWU).

Firing patterns repeated when the rooms were indistinguishable at their entrances, but differed otherwise. Compared to the repeating activity conditions in the other two environments, the correlations in the *same* condition were lower. This may be a result of a less frequent path integrator reset, since the L-shaped rooms contained only two region transitions. Therefore, input noise may have played a higher role.

The bend in the rooms does not impact the results. It is not considered a gateway and causes no remapping. Because the model relies on path integration alone, it does not matter whether the entrance remains visible. The L-shaped rooms only serve as an illustration of what sort of activation to expect in a case that has not previously been covered.

A3.4 Discussion

Place and grid cells have been observed to express uncorrelated firing patterns in different contexts (Muller and Kubie, 1987; Lever et al., 2002; Leutgeb et al., 2005; Fyhn et al., 2007; Colgin et al., 2008; Julian et al., 2018). The phenomenon, called “remapping”, has long been understood as a means of spatial pattern separation (Leutgeb et al., 2007) and may reflect a hierarchical organization of spatial representation (Julian et al., 2018). Remapping does not reflect purely local sensory differences, as both separate and identical patterns have been observed in similar contexts (Tanila, 1999; Carpenter et al., 2015; Grieves et al., 2016). However, the dynamics that give rise to the different firing patterns remain unknown.

In this study, we explore the idea that remapping is triggered by context change, defined as a transition between spatial regions such as the compartments of a maze. The mechanism requires a memory of contexts, implemented in our model as the gateway database. Within this memory, each context is represented as a pair consisting of (i) an identifier that can be detected from sensory input such as a snapshot and (ii) the pattern of place cell activity prevailing when the context was first encountered. Remapping occurs when a known gateway identifier is detected, e.g. by snapshot matching. In this case, the ongoing place cell activities are overridden by the stored place cell activity pattern of the recognized gateway. Projections into entorhinal cortex drive the grid cell attractor into its corresponding state and the normal path integration dynamics proceed from there.

The process is schematically illustrated in fig. A3.8: In a two-room environment, the agent starts by exploring the northern room. The accompanying attractor trajectory is shown as a dashed blue line in the left part of the figure. As soon as the agent enters the southern room, the gateway identifier 1 is recognized, the place cells remap to the stored pattern, and the system converges to the closest attractor, (P_1, G_1) . Exploration of the southern room results in continuous attractor shifts until the gateway is passed again (2) and another remapping event occurs.

Importantly, similar regions are assumed to share the same gateway identifier and entering them causes the same pattern to be activated. Thereby, the model was able to generate repetitive place and grid cell firing patterns qualitatively similar to those observed in multicompartiment environments. In the hairpin maze, the model also replicated the directionally selective firing commonly observed in corridors.

Within each region, grid cell activity is solely controlled by noisy path integration. This causes drift over time, which is also reset by remapping. Local cues, landmarks

and geometry do not influence cell activity and were indeed not modeled at all. We thus present an alternative to boundary vector cell models of place cell remapping, where geometric cues measured by boundary vector cells are the main determinants for place fields (Hartley et al., 2000; Barry et al., 2006; Bush et al., 2014; Grieves et al., 2018). In these models, repetitive patterns in similar compartments are a result of locally similar geometry.

With the L-shaped rooms environment, we propose a novel multicompartment setup where the two model types may be compared: At the far part of the rooms, the bend in the L-shape obscures the room entrance which is manipulated in two different contexts. Our model predicts that the cell patterns selected at the entrance persist over the entire region, even in parts of the rooms that are locally indistinguishable. I.e., if firing fields or behavioral measures are different between rooms, the context signal from the gateway identifier is sufficient to explain place cell remapping. If activity instead repeats, local cues are more important than the context at the entrance.

Note that we do not wish to imply that boundary vector cells play no role in place field generation. They provide position encoding independent from path integration (Bicanski and Burgess, 2020) and are able to explain place field deformation when the environment changes shape (O’Keefe and Burgess, 1996; Grieves et al., 2018) or barriers are inserted (Barry et al., 2006; Grieves et al., 2018). However, as mentioned in the introduction, boundary vector cells do not remap (Lever et al., 2009) and would need additional contextual gating to maintain distinct place cell patterns in different contexts.

A3.4.1 Implications for biology

The underlying assumption behind the context-dependent cognitive map is one of hierarchy. This is directly implied by the vaguely defined term “context” itself, but also by the lack of remapping *within* a context; the spatial representation can therefore be at least divided in similar and dissimilar parts. The re-use of place cells in different population codes implies that, within a specific context, information about other contexts is unavailable. However, rats are clearly able to execute trajectories across region boundaries, so some sort of higher-level representation likely exists. Possible candidates may be cell ensembles in the medial prefrontal cortex (Hyman et al., 2012) or neurons in the perirhinal cortex (Bos et al., 2017): The former show different firing patterns depending on region, time and task and could therefore en-

code context (Hyman et al., 2012). The latter are neurons that fire continuously while the rat is within a specific region in space, such as an arm of a maze.

For our model, the representation of other regions is less important than the representation of transitions to these regions, i.e., the gateways. For this, we introduced the notions of “gateway identifier” and “gate cell”, which represent the recognition of the entrance and the shift to the associated cell attractor. The biological gate cell would learn a place cell pattern via neuronal plasticity and reactivate that pattern whenever the relevant stimuli occurred. Outside of the associated places, the cell should remain silent. The activation of place cells from memory is plausible: for example, with replays (Pfeiffer, 2020) and preplays (Dragoi and Tonegawa, 2011; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015), the rat brain has been shown to depict sequences of place cells without the animal physically occupying the associated locations.

The gate cell attractor mechanism is intended as an alternative to contextual gating (M. Hayman and Jeffery, 2008), where the grid cell inputs into the hippocampus are modulated by a contextual signal to create different firing fields. Their model requires the continuous excitation and inhibition of large clusters of cells to produce orthogonal place field patterns. The advantage of our memory-driven gate cell system would be that it is more efficient: relatively few place cells need to be memorized and reactivated a single time to initially drive the attractor. I.e., the gate cell model scales much better with a high number of competing rooms or regions. On the other hand, the gate cell model cannot recover learned cell patterns at arbitrary positions, such as when the animal is placed by an experimenter (discussed below). In a sense, the distinction between the models is one of degree, i.e., how strongly the perception of context is influenced by memory and sensory information, and where exactly the neuronal connections lie.

For gateway-based context, it is of utmost importance for the animal to correctly recognize gateways, or specifically, the new region behind the gateway. A failure to do so could result in catastrophic mislocalization. Clearly, rats are able to detect changes in context, but evidence for an (over-)representation is scarce. Spiers et al. (2015) and Grieves et al. (2018) report the clustering of place fields around doorways to compartments, indicating that the doorways are at least salient landmarks; on the other hand, Duvelle et al. (2021) explicitly did not find such an overrepresentation in an experiment where the passage between compartments was blocked by lockable doors.

After detection, the rat must also be able to store gateways and associated place

cell patterns in long-term memory. A possible candidate may be the retrosplenial cortex (RSC): The rat RSC is reciprocally connected to the hippocampal formation and involved in a variety of spatial tasks (Vann et al., 2009). Lesions of the RSC only cause modest spatial deficits within a context, but impair spatial memory (Vann et al., 2009) and the ability to account for context changes, such as distal cue rotation or switching from light to darkness (Wesierska et al., 2009; Pothuizen et al., 2008).

Around 9% of RSC cells are head direction cells, some of which are additionally tuned to specific movements and locations (Chen et al., 1994; Jacob et al., 2017). Here, of particular interest may be a rare type of RSC cell described in the supplementary material of Jacob et al. (2017): The cells show localized firing at the doorway between two compartments and are tuned to a particular direction. This is precisely the behavior predicted for the gateway identifier units in our model.

Mathematically speaking, the definition and detection of gateways is non-trivial. A biologically plausible solution would be the prediction or tracking of changes geometric and visual information over time; at region transitions, this change is especially high (*surprise*, Butz et al. (2004) and Klukas et al. (2022)). In mobile robotics, the detection of doorways and narrow passages is a well-known problem with a variety of approaches based on cameras and laser range finders (e.g., Anguelov et al. (2004)) and posterior map segmentation (Thrun, 1998).

In humans, the cognitive map is generally thought to be hierarchically organized: in spatial memory and path planning tasks, differences arise in subjective distance estimation, recall speed (McNamara, 1986; McNamara et al., 1989) and preferred routes depending on region partition and transitions (Schick et al., 2019). In rodents, on the other hand, path planning is relatively unexplored. The best candidates are preplays, high-frequency place cell sequences observed in periods of inactivity. The preplays depict future trajectories, sometimes even through unexplored parts of the environment, indicating a planning component (Dragoi and Tonegawa, 2011; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015). However, whether preplays can cross region boundaries into other sub-maps is yet unknown.

A3.4.2 Remapping and place learning

Remapping at gateways is sufficient to replicate the place and grid cell dynamics observed in multicompartments environments, and it is the minimal requirement for distinct patterns in different rooms. However, there are many situations in which remapping is not quite as straight-forward. For example, what happens if the rat is

enters a compartment in darkness, or if it is placed at an arbitrary position by the experimenter? What if the surrounding cues are changed while the animal is within the room? Or what about environments with multiple entrances?

Remapping can also occur if the animal does not pass through a gateway (e.g., Lever et al., 2002; Anderson and Jeffery, 2003; Colgin et al., 2010; Jezek et al., 2011; Marozzi et al., 2015). For example, in Jezek et al. (2011), rats are “teleported” between two environments by the switching of light cues, and their firing fields remap accordingly. Gateways are convenient, but in principle, arbitrary positions could be associated with a place cell pattern much in the same way. These key locations could then be recognized and reactivated whenever the animal would experience environmental contextual change, prompting it to reorient itself. In the model, this type of remapping would be straight-forward to implement, but it would come with increased memory requirements and a high risk of aliasing within a region. Remembering every single location seems plausible for small, well known environments as in Jezek et al. (2011), especially if the animals are rewarded with randomly strewn food, but it seems unlikely that a rat could immediately localize itself when placed in an open field. Alternatively, the distance and direction to a remembered place could be inferred from depth perception or path integration (i.e., information provided by the boundary vector cells (Lever et al., 2009)) and the animal may be able adjust the shift the cell patterns to the correct position. In both cases, it would be more efficient to primarily memorize salient and behaviorally relevant locations, such as gateways.

Similarly, compartments with multiple entrances are no problem if the place cell pattern is memorized when leaving through each of them, meaning that the environment has been explored adequately. Keinath et al. (2020) report a difference in firing rate but not place field position depending on the most recent entrance, i.e., the animals were able to recover the overall place code in a familiar environment. Still, there is a qualitative difference between remapping to a novel pattern on first entry and reproducing a previous pattern. The former requires some sort of gateway representation, and the latter a memory. How representations evolve in an unfamiliar environment with multiple rooms and entrances remains to be explored.

In the behavioral experiment by Grieves et al. (2016), rats eventually improved in their ability to distinguish multiple identical rooms over days, although they never reached the performance of the control group. Similarly, Carpenter et al. (2015) showed that an initially repeating grid cell pattern between two identical, parallel rooms changed into one continuous pattern over both rooms after prolonged experi-

ence. The findings indicate that the identical room problem is not unsurmountable, even if it takes a lot of time. Interestingly, one interpretation of the findings by Carpenter et al. (2015) is that, instead of learning to differentiate the two rooms as different regions, the animals combined them into one larger region with no further remapping at the gateways. This explanation nicely fits to the notion of gateways based on surprise (e.g., Butz et al., 2004; Klukas et al., 2022). After all, the entrances are not surprising if the animals know the rooms in and out. Our model currently does not simulate long-term learning dynamics and changing gateway and region definitions - it doesn't even detect and recognize gateways automatically. However, the addition of such functions in future iterations is feasible.

A3.4.3 Remapping without the MEC

Evidence points towards grid cells in the MEC as an important driving force behind hippocampal place cell remapping (Monaco and Abbott, 2011; Bush et al., 2014; Kanter et al., 2017): Grid cell remap concurrently with place cells, the MEC directly projects to the hippocampus (Fyhn et al., 2007), and electrical stimulation of the area leads to hippocampal remapping (Kanter et al., 2017). In accordance, place cell activity and remapping in our model directly arise from grid cell input. However, place cells must be at least partially independent from MEC input: Their activity persists and the cells continue to remap in different contexts, even if grid cell firing is disrupted (Brandon et al., 2014) or the entire MEC is lesioned (Schlesiger et al., 2018).

In our model, place cell activity is sustained by grid cell input and would therefore immediately cease if the input was severed. To maintain activity in the absence of grid cell input, the model place cells would either need to be augmented with recurrent weights or input from another source. This presents an opportunity to reconcile our model with boundary vector cell models (Hartley et al., 2000; Barry et al., 2006; Bush et al., 2014; Grieves et al., 2018), where local geometry, i.e., the distance to walls, determines place fields. A combined model could continue to function with interrupted path integration. Perhaps, a combined model with context-dependent grid cells and context-independent boundary vector cells could also explain the phenomenon of *partial* remapping Latuske et al. (2018), where only some firing fields remap, while others stay at the same location.

Interestingly, place fields that are newly formed while grid cells are interrupted, remain the same when grid cell activity is later restored (Brandon et al., 2014). This

is consistent with the proposed feedback connections from place cells to grid cells in our model; once grid cell activity is restored, the newly formed place fields should move the grid cell activity to the closest attractor.

In conclusion, we suggest that remapping is a mechanism subserving the integration of the systems for path integration and place recognition, i.e., O’Keefe and Nadel’s (1978) Locale and Taxon systems. The cognitive map is thereby structured into a set of contexts or local charts, memorized by gateway identifiers. Within the charts, navigation is based on local metrics. On a coarser scale, the map is likely organized by the adjacencies of the charts, i.e., some sort of graph structure.

Acknowledgment

This work was carried out at the Department of Biology of the University of Tübingen, Tübingen, Germany. This research did not receive any specific grant from funding agencies in the public, commercial, or non-profit sectors.

Bibliography

- Allen, Kevin et al. (Oct. 2012). “Hippocampal Place Cells Can Encode Multiple Trial-Dependent Features through Rate Remapping”. *Journal of Neuroscience* 32.42, pp. 14752–14766.
- Almeida, Licurgo de, Marco Idiart, and John E. Lisman (June 2009). “The Input–Output Transformation of the Hippocampal Granule Cells: From Grid Cells to Place Fields”. *Journal of Neuroscience* 29.23, pp. 7504–7512.
- Anderson, Michael I. and Kathryn J. Jeffery (Oct. 2003). “Heterogeneous Modulation of Place Cell Firing by Changes in Context”. *Journal of Neuroscience* 23.26, pp. 8827–8835.
- Anguelov, D. et al. (Apr. 2004). “Detecting and Modeling Doors with Mobile Robots”. *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA ’04. 2004.* Vol. 4, 3777–3784 Vol.4.
- Barry, Caswell et al. (2006). “The Boundary Vector Cell Model of Place Cell Firing and Spatial Memory”. *Reviews in the Neurosciences* 17.1-2, pp. 71–97.
- Baumann, Tristan and Hanspeter A Mallot (2023). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239.
- Bicanski, Andrej and Neil Burgess (Sept. 2020). “Neuronal Vector Coding in Spatial Cognition”. *Nature Reviews Neuroscience* 21.9, pp. 453–470.
- Bjerknes, Tale L., Edvard I. Moser, and May-Britt Moser (Apr. 2014). “Representation of Geometric Borders in the Developing Rat”. *Neuron* 82.1, pp. 71–78.

- Boccaro, Charlotte N. et al. (Aug. 2010). “Grid Cells in Pre- and Parasubiculum”. *Nature Neuroscience* 13.8, pp. 987–994.
- Bos, Jeroen J. et al. (May 2017). “Perirhinal Firing Patterns Are Sustained across Large Spatial Segments of the Task Environment”. *Nature Communications* 8.1, p. 15602.
- Bostock, Elizabeth, Robert U. Muller, and John L. Kubie (1991). “Experience-Dependent Modifications of Hippocampal Place Cell Firing”. *Hippocampus* 1.2, pp. 193–205.
- Bouchacourt, Flora and Timothy J. Buschman (July 2019). “A Flexible Model of Working Memory”. *Neuron* 103.1, 147–160.e8.
- Brandon, Mark P. et al. (May 2014). “New and Distinct Hippocampal Place Codes Are Generated in a New Environment during Septal Inactivation”. *Neuron* 82.4, pp. 789–796.
- Brun, Vegard Heimly et al. (Jan. 2008). “Impaired Spatial Representation in CA1 after Lesion of Direct Input from Entorhinal Cortex”. *Neuron* 57.2, pp. 290–302.
- Bush, Daniel, Caswell Barry, and Neil Burgess (Mar. 2014). “What Do Grid Cells Contribute to Place Cell Firing?” *Trends in Neurosciences* 37.3, pp. 136–145.
- Butz, Martin V., Samarth Swarup, and David E. Goldberg (2004). “Effective Online Detection of Task-Independent Landmarks”. *Online Proceedings for the ICML’04 Workshop on Predictive Representations of World Knowledge*.
- Carpenter, Francis et al. (May 2015). “Grid Cells Form a Global Representation of Connected Environments”. *Current Biology* 25.9, pp. 1176–1182.
- Chen, Longtang L. et al. (Sept. 1994). “Head-Direction Cells in the Rat Posterior Cortex”. *Experimental Brain Research* 101.1, pp. 8–23.
- Cheng, Ken, Janellen Huttenlocher, and Nora S. Newcombe (Dec. 2013). “25 Years of Research on the Use of Geometry in Spatial Reorientation: A Current Theoretical Perspective”. *Psychonomic Bulletin & Review* 20.6, pp. 1033–1054.
- Cheng, Sen and Loren M Frank (2011). “The structure of networks that produce the transformation from grid cells to place cells”. *Neuroscience* 197, pp. 293–306.
- Colgin, Laura L. et al. (July 2010). “Attractor-Map Versus Autoassociation Based Attractor Dynamics in the Hippocampal Network”. *Journal of Neurophysiology* 104.1, pp. 35–50.
- Colgin, Laura Lee, Edvard I. Moser, and May-Britt Moser (Sept. 2008). “Understanding Memory through Hippocampal Remapping”. *Trends in Neurosciences* 31.9, pp. 469–477.
- Derdikman, Dori et al. (Oct. 2009). “Fragmentation of Grid Cell Maps in a Multi-compartment Environment”. *Nature Neuroscience* 12.10, pp. 1325–1332.
- Diamantaki, Maria et al. (Apr. 2018). “Manipulating Hippocampal Place Cell Activity by Single-Cell Stimulation in Freely Moving Mice”. *Cell Reports* 23.1, pp. 32–38.
- Dragoi, George and Susumu Tonegawa (Jan. 2011). “Preplay of Future Place Cell Sequences by Hippocampal Cellular Assemblies”. *Nature* 469.7330, pp. 397–401.

- Duvelle, Éléonore et al. (Mar. 2021). “Hippocampal Place Cells Encode Global Location but Not Connectivity in a Complex Space”. *Current Biology* 31.6, 1221–1233.e9.
- Fuhs, Mark C. et al. (Oct. 2005). “Influence of Path Integration Versus Environmental Orientation on Place Cell Remapping Between Visually Identical Environments”. *Journal of Neurophysiology* 94.4, pp. 2603–2616.
- Fyhn, Marianne et al. (Mar. 2007). “Hippocampal Remapping and Grid Realignment in Entorhinal Cortex”. *Nature* 446.7132, pp. 190–194.
- Fyhn, Marianne et al. (2004). “Spatial representation in the entorhinal cortex”. *Science* 305.5688, pp. 1258–1264.
- Gaussier, Ph et al. (2007). “A model of grid cells involving extra hippocampal path integration, and the hippocampal loop”. *Journal of integrative neuroscience* 6.03, pp. 447–476.
- Grieves, Roddy M, Éléonore Duvelle, and Paul A Dudchenko (July 2018). “A Boundary Vector Cell Model of Place Field Repetition”. *Spatial Cognition & Computation* 18.3, pp. 217–256.
- Grieves, Roddy M. and Kate J. Jeffery (Feb. 2017). “The Representation of Space in the Brain”. *Behavioural Processes* 135, pp. 113–131.
- Grieves, Roddy M. et al. (2016). “Place Field Repetition and Spatial Learning in a Multicompartment Environment”. *Hippocampus* 26.1, pp. 118–134.
- Guanella, Alexis, Daniel Kiper, and Paul Verschure (Aug. 2007). “A Model of Grid Cells Based on a Twisted Torus Topology”. *International Journal of Neural Systems* 17.04, pp. 231–240.
- Hafting, Torkel et al. (Aug. 2005). “Microstructure of a Spatial Map in the Entorhinal Cortex”. *Nature* 436.7052, pp. 801–806.
- Harland, Bruce et al. (Sept. 2017). “Lesions of the Head Direction Cell System Increase Hippocampal Place Field Repetition”. *Current Biology* 27.17, 2706–2712.e2.
- Hartley, Tom et al. (2000). “Modeling Place Fields in Terms of the Cortical Inputs to the Hippocampus”. *Hippocampus* 10.4, pp. 369–379.
- Hyman, James M. et al. (Mar. 2012). “Contextual Encoding by Ensembles of Medial Prefrontal Cortex Neurons”. *Proceedings of the National Academy of Sciences* 109.13, pp. 5086–5091.
- Jacob, Pierre-Yves et al. (Feb. 2017). “An Independent, Landmark-Dominated Head-Direction Signal in Dysgranular Retrosplenial Cortex”. *Nature Neuroscience* 20.2, pp. 173–175.
- Jeffery, Kathryn J. (Nov. 2011). “Place Cells, Grid Cells, Attractors, and Remapping”. *Neural Plasticity* 2011, e182602.
- Jezek, Karel et al. (2011). “Theta-paced flickering between place-cell maps in the hippocampus”. *Nature* 478.7368, pp. 246–249.
- Julian, Joshua B. et al. (Sept. 2018). “The Neurocognitive Basis of Spatial Reorientation”. *Current Biology* 28.17, R1059–R1073.

- Kanter, Benjamin R. et al. (Mar. 2017). “A Novel Mechanism for the Grid-to-Place Cell Transformation Revealed by Transgenic Depolarization of Medial Entorhinal Cortex Layer II”. *Neuron* 93.6, 1480–1492.e6.
- Keinath, Alexandra T. et al. (June 2020). “DG–CA3 Circuitry Mediates Hippocampal Representations of Latent Information”. *Nature Communications* 11.1, p. 3026.
- Klukas, Mirko et al. (Jan. 2022). *Fragmented Spatial Maps from Surprisal: State Abstraction and Efficient Planning*. Preprint.
- Krupic, Julija et al. (Feb. 2015). “Grid Cell Symmetry Is Shaped by Environmental Geometry”. *Nature* 518.7538, pp. 232–235.
- Langston, Rosamund F. et al. (June 2010). “Development of the Spatial Representation System in the Rat”. *Science* 328.5985, pp. 1576–1580.
- Latuske, Patrick et al. (2018). “Hippocampal Remapping and Its Entorhinal Origin”. *Frontiers in Behavioral Neuroscience* 11.
- Leutgeb, Jill K. et al. (Feb. 2007). “Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus”. *Science* 315.5814, pp. 961–966.
- Leutgeb, Stefan et al. (July 2005). “Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles”. *Science* 309.5734, pp. 619–623.
- Lever, Colin et al. (Aug. 2009). “Boundary Vector Cells in the Subiculum of the Hippocampal Formation”. *Journal of Neuroscience* 29.31, pp. 9771–9777.
- Lever, Colin et al. (Mar. 2002). “Long-Term Plasticity in Hippocampal Place-Cell Representation of Environmental Geometry”. *Nature* 416.6876, pp. 90–94.
- Li, Tianyi, Angelo Arleo, and Denis Sheynikhovich (Jan. 2020). “Modeling Place Cells and Grid Cells in Multi-Compartment Environments: Entorhinal–Hippocampal Loop as a Multisensory Integration Circuit”. *Neural Networks* 121, pp. 37–51.
- Linsker, Ralph (Nov. 1997). “A Local Learning Rule That Enables Information Maximization for Arbitrary Input Distributions”. *Neural Computation* 9.8, pp. 1661–1665.
- Lyttle, David et al. (2013). “Spatial Scale and Place Field Stability in a Grid-to-Place Cell Model of the Dorsoventral Axis of the Hippocampus”. *Hippocampus* 23.8, pp. 729–744.
- M. Hayman, Robin and Kathryn J Jeffery (2008). “How heterogeneous place cell responding arises from homogeneous grids—a contextual gating hypothesis”. *Hippocampus* 18.12, pp. 1301–1313.
- Marozzi, Elizabeth et al. (2015). “Purely translational realignment in grid cell firing patterns following nonmetric context change”. *Cerebral Cortex* 25.11, pp. 4619–4627.
- Martín Abadi et al. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. URL: <https://www.tensorflow.org/>.
- McNamara, Timothy P (Jan. 1986). “Mental Representations of Spatial Relations”. *Cognitive Psychology* 18.1, pp. 87–121.

- McNamara, Timothy P., James K. Hardy, and Stephen C. Hirtle (1989). “Subjective Hierarchies in Spatial Memory”. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15.2, pp. 211–227.
- McNaughton, B. L., C. A. Barnes, and J. O’Keefe (Sept. 1983). “The Contributions of Position, Direction, and Velocity to Single Unit Activity in the Hippocampus of Freely-Moving Rats”. *Experimental Brain Research* 52.1, pp. 41–49.
- Monaco, Joseph D. and Larry F. Abbott (June 2011). “Modular Realignment of Entorhinal Grid Cell Activity as a Basis for Hippocampal Remapping”. *Journal of Neuroscience* 31.25, pp. 9414–9425.
- Moser, Edvard I., May-Britt Moser, and Bruce L. McNaughton (Nov. 2017). “Spatial Representation in the Hippocampal Formation: A History”. *Nature Neuroscience* 20.11, pp. 1448–1464.
- Moser, May-Britt, David C. Rowland, and Edvard I. Moser (Feb. 2015). “Place Cells, Grid Cells, and Memory”. *Cold Spring Harbor Perspectives in Biology* 7.2, a021808.
- Muessig, Laurenz et al. (June 2015). “A Developmental Switch in Place Cell Accuracy Coincides with Grid Cell Maturation”. *Neuron* 86.5, pp. 1167–1173.
- Muller, R. U. and J. L. Kubie (July 1987). “The Effects of Changes in the Environment on the Spatial Firing of Hippocampal Complex-Spike Cells”. *Journal of Neuroscience* 7.7, pp. 1951–1968.
- O’Keefe, J. and J. Dostrovsky (Nov. 1971). “The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely-Moving Rat”. *Brain Research* 34.1, pp. 171–175.
- O’Keefe, John and Neil Burgess (May 1996). “Geometric Determinants of the Place Fields of Hippocampal Neurons”. *Nature* 381.6581, pp. 425–428.
- O’Keefe, John and Lynn Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford : New York: Clarendon Press ; Oxford University Press.
- Ólafsdóttir, H Freyja et al. (June 2015). “Hippocampal Place Cells Construct Reward Related Sequences through Unexplored Space”. *eLife* 4. Ed. by Howard Eichenbaum, e06063.
- Pfeiffer, Brad E. (2020). “The Content of Hippocampal “Replay””. *Hippocampus* 30.1, pp. 6–18.
- Pfeiffer, Brad E. and David J. Foster (May 2013). “Hippocampal Place-Cell Sequences Depict Future Paths to Remembered Goals”. *Nature* 497.7447, pp. 74–79.
- Pothuizen, Helen H. J., John P. Aggleton, and Seralynne D. Vann (2008). “Do Rats with Retrosplenial Cortex Lesions Lack Direction?” *European Journal of Neuroscience* 28.12, pp. 2486–2498.
- Rennó-Costa, César and Adriano B. L. Tort (Aug. 2017). “Place and Grid Cells in a Loop: Implications for Memory Function and Spatial Coding”. *Journal of Neuroscience* 37.34, pp. 8062–8076.
- Rowland, David C. et al. (2016). “Ten Years of Grid Cells”. *Annual Review of Neuroscience* 39.1, pp. 19–40.

- Schick, Wiebke et al. (July 2019). “Language Cues in the Formation of Hierarchical Representations of Space”. *Spatial Cognition & Computation* 19.3, pp. 252–281.
- Schlesinger, Magdalene I. et al. (Mar. 2018). “Hippocampal Global Remapping Can Occur without Input from the Medial Entorhinal Cortex”. *Cell Reports* 22.12, pp. 3152–3159.
- Skaggs, William E. and Bruce L. McNaughton (Oct. 1998). “Spatial Firing Properties of Hippocampal CA1 Populations in an Environment Containing Two Visually Identical Regions”. *Journal of Neuroscience* 18.20, pp. 8455–8466.
- Solstad, Trygve et al. (Dec. 2008). “Representation of Geometric Borders in the Entorhinal Cortex”. *Science* 322.5909, pp. 1865–1868.
- Spiers, Hugo J. et al. (Jan. 2015). “Place Field Repetition and Purely Local Remapping in a Multicompartment Environment”. *Cerebral Cortex* 25.1, pp. 10–25.
- Stensola, Hanne et al. (Dec. 2012). “The Entorhinal Grid Map Is Discretized”. *Nature* 492.7427, pp. 72–78.
- Tanila, Heikki (1999). “Hippocampal Place Cells Can Develop Distinct Representations of Two Visually Identical Environments”. *Hippocampus* 9.3, pp. 235–246.
- Taube, J. S., R. U. Muller, and J. B. Ranck (Feb. 1990a). “Head-Direction Cells Recorded from the Postsubiculum in Freely Moving Rats. I. Description and Quantitative Analysis”. *Journal of Neuroscience* 10.2, pp. 420–435.
- Taube, J. S., R. U. Muller, and J. B. Ranck (Feb. 1990b). “Head-Direction Cells Recorded from the Postsubiculum in Freely Moving Rats. II. Effects of Environmental Manipulations”. *Journal of Neuroscience* 10.2, pp. 436–447.
- Thrun, Sebastian (Feb. 1998). “Learning Metric-Topological Maps for Indoor Mobile Robot Navigation”. *Artificial Intelligence* 99.1, pp. 21–71.
- Unity Technologies (Apr. 15, 2021). *Unity*. Version 2020.3.4f1. URL: <https://unity.com/>.
- Vann, Seralynne D., John P. Aggleton, and Eleanor A. Maguire (Nov. 2009). “What Does the Retrosplenial Cortex Do?” *Nature Reviews Neuroscience* 10.11, pp. 792–802.
- Wesierska, Malgorzata, Iwona Adamska, and Monika Malinowska (Jan. 2009). “Retrosplenial Cortex Lesion Affected Segregation of Spatial Information in Place Avoidance Task in the Rat”. *Neurobiology of Learning and Memory* 91.1, pp. 41–49.
- Wills, Thomas, Caswell Barry, and Francesca Cacucci (2012). “The Abrupt Development of Adult-like Grid Cell Firing in the Medial Entorhinal Cortex”. *Frontiers in Neural Circuits* 6.
- Wills, Tom J. et al. (May 2005). “Attractor Dynamics in the Hippocampal Representation of the Local Environment”. *Science* 308.5723, pp. 873–876.
- Wilson, Matthew A. and Bruce L. McNaughton (Aug. 1993). “Dynamics of the Hippocampal Ensemble Code for Space”. *Science* 261.5124, pp. 1055–1058.
- Zhang, K., G. Ganis, and M.I. Sereno (June 1997). “Anti-Hebbian Synapses as a Linear Equation Solver”. *Proceedings of International Conference on Neural Networks (ICNN’97)*. Vol. 1, 387–389 vol.1.

Appendix A4

Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments

Authors: Tristan Baumann and Hanspeter A Mallot, 2023.

Abstract

The structure of the internal representation of surrounding space, the so-called *cognitive map*, has long been debated. A Euclidean metric map is the most straightforward hypothesis, but human navigation has been shown to systematically deviate from the Euclidean ground truth. Vector navigation based on non-metric models can better explain the observed behavior, but also discards useful geometric properties such as fast shortcut estimation and cue integration.

Here, we propose another alternative, a Euclidean metric map that is systematically distorted to account for the observed behavior. The map is found by embedding the non-metric model, a labeled graph, into 2D Euclidean coordinates. We compared these two models using human data from Warren et al. (2017), where participants had to navigate and learn a non-Euclidean maze (i.e., with Wormholes) and perform direct shortcuts between different locations. Even though the Euclidean embedding cannot correctly represent the non-Euclidean environment, both models predicted the data equally well. We argue that the so embedded graph naturally arises from integrating the local position information into a metric framework, which makes the model more powerful and robust than the non-metric alternative. It may therefore be a better model for the human cognitive map.

A4.1 Introduction

A4.1.1 The cognitive map

The spatial long-term memory contains representations of places, landmarks, and local views. A sequence of navigational actions connecting these representations is called a route and animals with such route knowledge are able to navigate between known places by following these routes (Collett et al., 1998; Collett and Collett, 2002; Warren, 2019; Mallot, 2024). If knowledge about many different items, places, and routes is integrated and novel routes and shortcuts can be inferred from previously learned route segments, the representation is called a map (Tolman, 1948; O’Keefe and Nadel, 1978; Gallistel, 1990; Trullier et al., 1997; Mallot, 2024). The cognitive map is thus a form of declarative memory in the sense that it characterizes “knowing what” or “knowing where” as opposed to the non-declarative “knowing how” of routes or guidance information (O’Keefe and Nadel, 1978; Squire and Knowlton, 1995).

A cognitive map is the most general form of spatial long-term memory, and it is believed that many animals, including humans, have access to this representation (Gallistel, 1990; Nadel, 2013; Warren, 2019). This is exemplified by the existence of neural correlates of position, the place cells (O’Keefe and Nadel, 1978; Rolls and O’Mara, 1995; Ekstrom et al., 2003; Yartsev and Ulanovsky, 2013), which encode the position of the animal within the current context via population activity.

The intuition of an internal map is relatively straight-forward, because it matches maps encountered in everyday life: In general, such maps may be broadly characterized by two frameworks: Euclidean metric maps and topological graphs. Euclidean metric maps, such as a bird’s eye view of a city or a satellite image, assign unique coordinates to each position that approximate the real-world geometry by preserving the metric relationships between positions. Topological graphs, such as a subway or bus chart or an instruction manual, describe states and possible actions that lead from one state to another, rather than geometry.

The metric framework (Fig. A4.1C) is considerably better suited to explain environments with a Euclidean geometric structure, and, based on the Kantian notion of an *a priori* assumption of absolute external space (Kant, 1781), it has often been argued that the cognitive map must likewise be Euclidean metric to capture these properties (O’Keefe and Nadel, 1978; Gallistel, 1990; McNaughton et al., 2006; Nadel, 2013). This theory is supported by the existence of grid cells in the entorhinal

cortex, which are believed to encode metric path integration information (Hafting et al., 2005; McNaughton et al., 2006; Peer et al., 2021).

The notion of an absolute Euclidean metric may be challenged, for example, by pointing out that the intuition of straight lines on a curved surface (or any surface that is not a plane) are actually geodesics and not true straight lines in an Euclidean sense (Helmholtz (1876) and Mallot (2024), cf. Indow (1999)). But even an approximately Euclidean or non-Euclidean metric map may be advantageous, since geometric relationships between places are preserved in a highly efficient manner. That is, distances, routes, and shortcuts can be directly inferred from the map and need not be memorized individually. This property enables metric maps to store an immense amount of data, making them powerful informational tools (Nadel, 2013).

However, results from navigation experiments often disagree with the Euclidean metric map hypothesis: Human performance in shortcut or triangle completion tasks, which are often taken as evidence for an Euclidean representation, is highly unreliable with angular errors of over $\pm 90^\circ$ and angular standard deviations between $25^\circ - 45^\circ$ (Foo et al., 2005; Ishikawa and Montello, 2006; Chrastil and Warren, 2013; Warren, 2019). The Euclidean metric postulates are often violated and angle and distance estimations are systematically biased by features of the environment such as landmarks, junctions or region boundaries (Byrne, 1979; McNamara, 1986; Sadalla and Montello, 1989; Tversky, 1992; Warren, 2019; Kim and Doeller, 2022), or the number and recency of preceding turns (Brunec et al., 2017; Meilinger et al., 2018; Peer et al., 2021). In rats, place cells have been shown to stretch and shear following room deformation, while still preserving topological information about the environment (O’Keefe and Burgess, 1996; Dabaghian et al., 2014). These results imply that space is encoded much worse than what a precise metric map would predict.

As an alternative, the comparatively weaker class of topological graphs is often proposed. The environment is expressed through neighborhood or adjacency relations, forming a network of places as graph vertices and paths or actions connecting them as edges (Gillner and Mallot, 1998; Kuipers, 1978; Mallot and Basten, 2009; Warren, 2019; Peer et al., 2021). The graph may be labeled with pairwise distance or angle measurements, but this information need not adhere to the metric postulates and is therefore not metric (Fig. A4.1A). Still, shortcuts and novel routes can be derived via vector addition of the labels along paths in the graph; indeed, Warren et al. (2017) suggest vector addition based on labeled graphs best explains human performance in navigation experiments.

Nevertheless, poor navigational performance, biases, and large errors are not

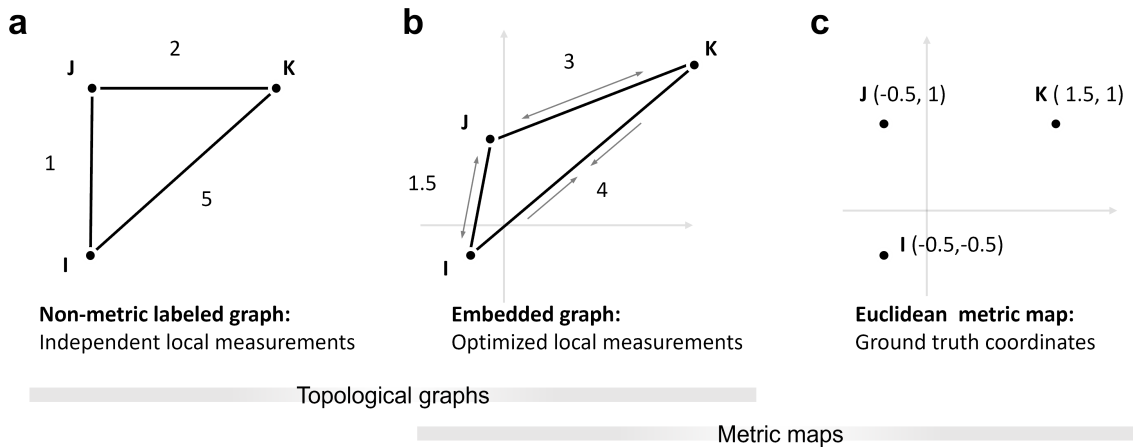


Figure A4.1: Cognitive map hypotheses. **(a)** Non-metric topological graph, labeled with distance measurements. The labels are independent of each other and do not need to adhere to the triangle inequality. **(b)** Embedded graph from (a). To find a Euclidean embedding, the distance labels need to be adjusted to create a valid configuration, for example, by stretching or compressing the edges or “wiggling” on the vertices until the difference between map and labels is minimized. As opposed to the non-metric labeled graph, changes to one label will therefore influence others. **(c)** Euclidean metric map. Places are directly assigned coordinates based on their position in the world. Over time, the coordinates may be refined by repeated measurements and the map will approach the Euclidean ground truth. The same can be expected from the embedded graph optimization if the labels are refined.

enough to completely rule out a Euclidean metric representation because the map may be systematically biased or distorted to large degrees while still being metric (Warren, 2019). Overall, the distance errors in metrically embedded maps will be smaller than in labeled graphs where distance labels are independent. Metric embedding is thus a means for efficiently exploiting all available distance information.

A4.1.2 Distorted maps and non-Euclidean environments

Each individual cognitive representation will generally be different due to acquisition order, biases, and accumulation of measurement errors. One possible advantage of map-like representations in spatial memory is the mutual refinement of (possibly conflicting) local position information over time: As the agent explores the environment, it will repeatedly obtain distance and angle measurements of connections between the known places or graph states.

With a topological graph, repeated measurements of the same information could be used to create more precise labels by averaging. However, the labels will always remain independent of labels corresponding to adjacent connections and, in a triangle, might persistently violate the triangle inequality which defines a mathematical metric (Fig. A4.1a). In the following, we therefore refer to this representation as the *non-metric labeled graph*.

Additional precision can only be gained if repeated measurements of one connection will also improve estimates along other connections in the graph. This may, for example, be achieved by metric embedding (Fig. A4.1b). Since the acquisition of spatial memory is not complete after a single pass through the environment but relies on the consolidation of many local measurements, metric embedding seems to be a natural method for continuous integration of local information. In this sense, cue integration might be the main reason for organizing spatial representations in a metric framework.

If the measured labels are not perfect, Euclidean metric embedding can only approximate the true Euclidean metric relations and will result in a distorted depiction. The so *embedded graph* could therefore be an alternative metric explanation for the large deviations in human navigation, as opposed to the non-metric labeled graph.

In regular environments, differences between a non-metric labeled graph, an embedded graph, and a Euclidean metric map will be minimal, because the models are likely to approach the same underlying ground truth as measurements are refined. Therefore, cases need to be considered in which the models would make different

predictions. With the advent of immersive virtual reality, a unique opportunity has opened up to present non-Euclidean environments, thus dissociating presented metric information from the underlying true Euclidean positions (Zetsche et al., 2009; Kluss et al., 2015; Warren et al., 2017; Widdowson and Wang, 2023). The non-Euclidean manipulations have been shown to heavily influence navigation but are usually not noticed by the subjects (Zetsche et al., 2009; Warren et al., 2017).

A4.1.3 Evidence from wormhole experiments

In the following, we focus on a specific example, Warren et al. (2017), because the experiment offers an excellent setup to investigate the hypotheses with respect to systematic distortion and the data are available online.

Warren et al. (2017) presented participants with a non-Euclidean environment and argued that, if the cognitive map is Euclidean metric, participants should have greater difficulties in learning the non-Euclidean environment compared to control, because mismatches between the cognitive map and the environment should occur. On the other hand, a non-metric graph should have no such issues.

Using head-mounted display virtual reality, Warren et al. created a hedge maze augmented with two invisible wormholes. The wormholes functioned as instant seamless teleportation and 90° rotation between different parts of the maze while participants continued to walk normally in the real-life room, therefore creating a mismatch between maze position and path integration information. Interestingly, only one participant reported noticing any kind of spatial anomaly in the maze.

Participants had to memorize object positions within the maze and were later asked to walk direct shortcuts between them. For this, the participants were moved to a starting object and had time to orient themselves. Then, the walls of the maze disappeared, and the participants had to walk to the presumed position of a target object. The initial angles of the subjects' trajectories were measured and used as directional estimates to compare the non-metric labeled graph and undistorted Euclidean map models.

Warren et al. found that directional estimates were heavily distorted towards the wormholes. This is predicted by vector addition along the shortest path on a labeled graph but not by straight lines in Euclidean ground truth coordinates. The authors thus rejected the Euclidean map hypothesis in favor of the non-metric labeled graph, arguing that only a non-Euclidean structure could explain the observed results (Warren et al., 2017; Warren, 2019). A distorted Euclidean map was briefly

considered but rejected on the basis that such a map “must still satisfy the metric postulates [...] in the inertial coordinate system” (Warren et al., 2017). However, the metric embedding is not a simple averaging of the path integration coordinates (Warren’s “inertial coordinates”) but the result of an optimization that also takes into account the other connections in the graph. The representation of a goal will therefore not necessarily end up in the middle between the path integration vectors obtained along two different paths but may be closer to one or the other.

We reexamined the data used in Warren et al. (2017) with respect to the possibility of a distorted Euclidean map. In the following, we show that such a map can be found by first creating a non-metric labeled graph for the maze and then embedding the graph into 2D Euclidean coordinates. This is achieved by the minimization of the angle and distance differences between graph and map, following the method described in Hübner and Mallot (2007) and Mallot (2024) for the embedding of view graphs. In an ordinary Euclidean environment, the embedding will recover the ground truth coordinates, but in a non-Euclidean environment, a residual error between embedding and local measurements must remain. Because of this error, the models should make different predictions, and may be distinguished by comparing their predictions to experimental data. That is, shortcuts derived from the embedded graph should fall somewhere between the shortcuts from the other two models.

However, we found that both models, the non-metric labeled graph and its Euclidean metric embedding, were able to predict the data equally well. Because the embedded graph is a valid Euclidean map, it is better suited for shortcut generation and especially cue integration than the non-metric alternative. We therefore refute the claim by Warren et al. (2017) that their findings cannot be explained by a Euclidean metric map and argue for the embedded graph as a better alternative explanation.

A4.2 Materials and methods

A4.2.1 Data acquisition

The data used here are figures, measurements, and results from Warren et al. (2017). The anonymized per-subject measurements are available as supplementary material online in the Brown University Digital Repository (<http://dx.doi.org/10.7301/Z0JS9NC5>, retrieved in November 2022). The relevant datasets contain measurements of the direction of individual shortcuts between object pairs in the wormhole




Figure A4.2: Maze and shortcut data. **(a)** Layout of the wormhole maze, redrawn from Warren et al. (2017). The yellow arrows show wormhole position and magnitude. Touching one end of the arrow instantly and seamlessly teleported subjects to the other end. **(b, c)** Example directional estimates for object pairs from the “Route-finding and shortcuts” dataset (Experiment 1 in Warren et al. (2017)) and the “Rips and folds” dataset (Experiment 2 in Warren et al. (2017)). The thin arrows show the Euclidean ground truth direction between objects, the short dotted lines the corresponding subject estimates, and the thick solid line the average subject estimate. The length of the estimates has been normalized and does not reflect walked distance. In (c), the colors indicate different goals.

maze, given as angular difference between the estimate and the straight-line direction in Euclidean ground truth coordinates. We estimated these coordinates from pixel positions in Fig. 2B in Warren et al. (2017) (Fig. A4.3A) and transformed the subject estimates into global angles (i.e., increasing counterclockwise from the positive x-axis or east). The layout of the maze and example subject estimates are shown in Fig. A4.2.

Warren et al. (2017) measured direction estimates in two separate experiments, one to investigate shortcuts (Dataset “Route-finding and shortcuts”, see Fig. A4.2B) and one to investigate the ordinal reversal of landmark positions (Dataset “Rips and folds”, see Fig. A4.2C). “Route-finding and shortcuts” contains directional estimates of 10 subjects (5M, 5F) for four pairs of objects for a total of $10 \times 4 \times 2$ (bidirectional) = 80 measurements. “Rips and folds” contains directional estimates of 11 subjects (9M, 2F) for eight starting locations and three targets each for a total of $11 \times 8 \times 3 = 264$ measurements. For the purpose of this study, both datasets were treated the same but were evaluated separately; this was done for direct comparison and to avoid bias because some subjects may have participated in both studies. For further information about the participants, hardware, and experimental setup please refer to Warren et al. (2017).




Figure A4.3: Graph creation. (a) Vertex positions in the maze. Their pixel coordinates were considered the Euclidean ground truth for the model. The maze was partitioned into straight segments and corners, and one vertex was placed per corner. Two vertices, 12 and 35, were only used in a control graph without wormholes. (b) The corresponding topological graph with edges through wormholes (red dotted lines). The graph was then labeled with local distance and angle measurements based on the ground truth, except for the wormhole edges, which were manually adjusted to reflect the locally distorted topology instead. Note that the distance along the wormhole edges is shortened but not zero.

A4.2.2 Graph and map setup

Plausible Euclidean embeddings were found in two steps: First, a topological graph of the maze was created and labeled with the veridical distances and angles. This graph was also used to derive predictions for the non-metric labeled graph hypothesis, i.e., vector addition of the labels along the shortest paths. Next, the graph was embedded into 2D Euclidean coordinates by iterative minimization of a stress function (Hübner and Mallot, 2007; Mallot, 2024) describing the difference between the coordinates and the local labels.

In general, the creation of the topological graph is a non-trivial problem with a possibly infinite set of solutions consisting of any number of vertices, edges and measurements along the maze. Therefore, good solutions have to be guessed. Because the wormhole maze consisted of well-defined straight segments and corners, we created the graph by placing one vertex per corner and one edge per straight segment (Fig. A4.3B). Formally, we define the graph $G = \{V, E\}$ as a set of n vertices $V = \{v_1, \dots, v_n\}$ corresponding to places in the maze and edges $E = \{e_{ij}, e_{jk}, \dots\}$

describing maze arms connecting the places v_i to v_j and v_j to v_k .

The algorithm for metric embedding is based on local distance and turning information only, without the assumption of a global reference direction (e.g., north). It is therefore based on triplets of neighboring places $T = \{(i, j, k)\}$, i.e., places that can be visited in sequence. For each triplet, the distances d_{ij} , d_{jk} and the turning angle α_{ijk} were measured and added as labels to the topological graph. d_{ij} and d_{jk} describe the distances between places i, j and j, k and α_{ijk} the heading change at j when moving from i to k . All labels were taken from the required egomotion steps such that labels around wormholes differed from the Euclidean ones. The same labeled graph was used for datasets from all subjects.

From the graph, a 2D Euclidean embedding $X = \{(\mathbf{x}_1, \dots, \mathbf{x}_n)\}$ of the n vertices was derived by minimizing the following stress function: The algorithm considers all measured triplets of neighboring places $T = \{(i, j, k)\}$ and their related distance and angle measurements $(d_{ij}, d_{jk}, \alpha_{ijk})$. Each place may appear many times as part of different triplets, and forward-backward movements of the form (i, j, i) are also considered (with $\alpha_{ijk} = 180^\circ$). The stress function can then be written as

$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_{(i,j,k) \in T} \lambda_1 [((\mathbf{x}_j - \mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_k)) - d_{ij}d_{jk}\cos\alpha_{ijk}]^2 + \lambda_2 [((\mathbf{x}_j - \mathbf{x}_i) \otimes (\mathbf{x}_j - \mathbf{x}_k)) - d_{ij}d_{jk}\sin\alpha_{ijk}]^2. \quad (\text{A4.1})$$

here, (\cdot) denotes the dot product and (\otimes) the third component of the cross product, $(\mathbf{a} \otimes \mathbf{b}) := a_1b_2 - a_2b_1$, which is twice the area of the triangle (i, j, k) . The constants λ_1, λ_2 can be used to weigh the components based on their variances (Mallot, 2024); we chose $\lambda_1 = \lambda_2 = 1$.

Finding an embedding that minimizes this stress function is a nonlinear optimization problem. Solutions may, for example, be found with iterative numerical approximations like Newton's method. We used the quasi-Newton method Sequential Least Squares Programming (SLSQP), as implemented in the *SciPy 1.10 optimize* Python library (Virtanen et al., 2020), credited to (Kraft, 1988).

The resulting embedding will be a Euclidean metric map of the graph's vertices with an arbitrary global orientation, but it is not a complete distorted map of the wormhole maze in the sense that it only assigns coordinates to the vertices but not to other places. The distorted position of other places may be found by adding them as additional vertices to the graph before embedding or by interpolation. Nevertheless,

the embedding is sufficient to derive directional predictions.

A4.2.3 Model comparison and data analysis

Next, the non-metric labeled graph and its Euclidean embedding were used to derive predictions about shortcut directions between object pairs. For the non-metric labeled graph, predictions were obtained by finding the shortest path between start and target object using Dijkstra’s algorithm, as implemented in the *NetworkX 3.0* Python library (Hagberg et al., 2008). Along the path, the angles and distances were summed up to a vector, and the global direction of the resultant vector relative to the ground truth coordinates was considered the final shortcut prediction. In the embedded graph, shortcuts were simply the straight lines from start to target objects.

The predictions of the two graph models were compared to the subject data and the prediction error was measured. Because the embedded graph has no defined reference direction, subject estimates had to be considered relative to a local reference. We used the respective local angle between the starting arm and measurement or prediction, which is independent of the reference direction.

For each model, the mean prediction errors and between-subject angular deviation were calculated for the group, and the within-subject angular deviation separately for each participant. The errors were compared with the two-sample Watson-Williams F-test for circular data (Batschelet, 1981), as implemented in the *PyCircStat* Python library (Berens and Sinz, 2022). The null hypothesis assumes that the samples come from underlying distributions with the same mean (Batschelet, 1981), i.e., that the models explain the subject data equally well; note that this does not mean that the models make the same prediction. Cohen’s d was used as a measure for effect size. All statistical tests were two-tailed with $\alpha = 0.05$.

A4.3 Results

A4.3.1 Embeddings

The numerical optimization method may find different local minima. Which solution is found depends on the starting point in the solution space, i.e., the initial vertex positions $X = \{(\mathbf{x}_1, \dots, \mathbf{x}_n)\}$. We restarted the optimization procedure 1000 times with random initial vertex positions $X \sim \mathcal{U}_2(0, 20)$ and found two local minima with




Figure A4.4: Embedded graph. **(a, b)** References for comparison. The unembedded graph (as in Fig. A4.3) and an another embedding which was also found by the optimization method. The second embedding performed worse on the subject data and was not further used. **(b)** The embedded graph, i.e., the labeled graph with the vertices at coordinates that minimize the difference between map and labels. The orientation of the embeddings is arbitrary; here, they were rotated so that the edge (2, 3) is horizontal. The red dotted lines show the edges that pass through wormholes. **(d)** Sketch of the distorted wormhole maze according to the embedding in (c). Edges that cross each other in the embedded graph could, for example, be realized through multi-level paths.

stress values $f(X_1) = 450.68$ (Fig. A4.4c) and $f(X_2) = 367.36$ (Fig. A4.4b). In the following, we report results from the first embedding, which resulted in better fits to the subject data.

A4.3.2 Dataset 1: Route-finding and shortcuts

We derived shortcut predictions from the non-metric labeled and embedded graph models and compared the predictions to human shortcut estimates from Warren et al. (2017), dataset “Route-finding and shortcuts” (Fig. A4.5A-C). The resulting angular prediction error was measured. Rayleigh tests on error direction revealed non-uniform distributions, $z(10) = 9.59, p < .001$ for the non-metric labeled graph and $z(10) = 9.57, p < .001$ for the embedding.

The non-metric labeled graph model showed an average angular error of -12.4° with an angular deviation (AD) of 11.76° and the embedding an average error of -15.26° , $AD = 11.98^\circ$. This difference was not significant ($F(1, 18) = 0.2, p = .63$) with a small effect size ($d = .22$). I.e., the shortcut directions derived from the graph model were not significantly closer to the subject data than the shortcut directions derived from the embedding or vice versa.

The within-subject angular deviation of the errors was fairly high but also similar for both models, with an average of 29.75° for the graph model and 32.15° for the embedding. Statistical comparison ($F(1, 18) = 0.6, p = .42, d = .51$) again revealed no significant difference.

A4.3.3 Dataset 2: Rips and folds

For the purpose of this study, the “Rips and folds” dataset was treated the same as the “Route-finding and shortcuts” dataset, with the only difference being the number of participants (11 vs. 10 in dataset 1) and estimates per participant (24 vs. 8 in dataset 1). The datasets were analyzed separately for the sake of comparison.

We again compared prediction errors of the non-metric labeled graph model and its Euclidean embedding (Fig. A4.5D-F). Rayleigh test on error direction revealed non-uniform distributions, $z(11) = 10.78, p < .001$ for the non-metric labeled graph and $z(11) = 10.77, p < .001$ for the embedded graph. The non-metric labeled graph showed an average angular error of 5.68° , $AD = 8.12^\circ$, and the embedded graph an error of 2.37° , $AD = 9.35^\circ$. This difference was again not significant ($F(1, 20) = 0.72, p = .41$) with a small effect size ($d = .39$). Within-subject angular deviation

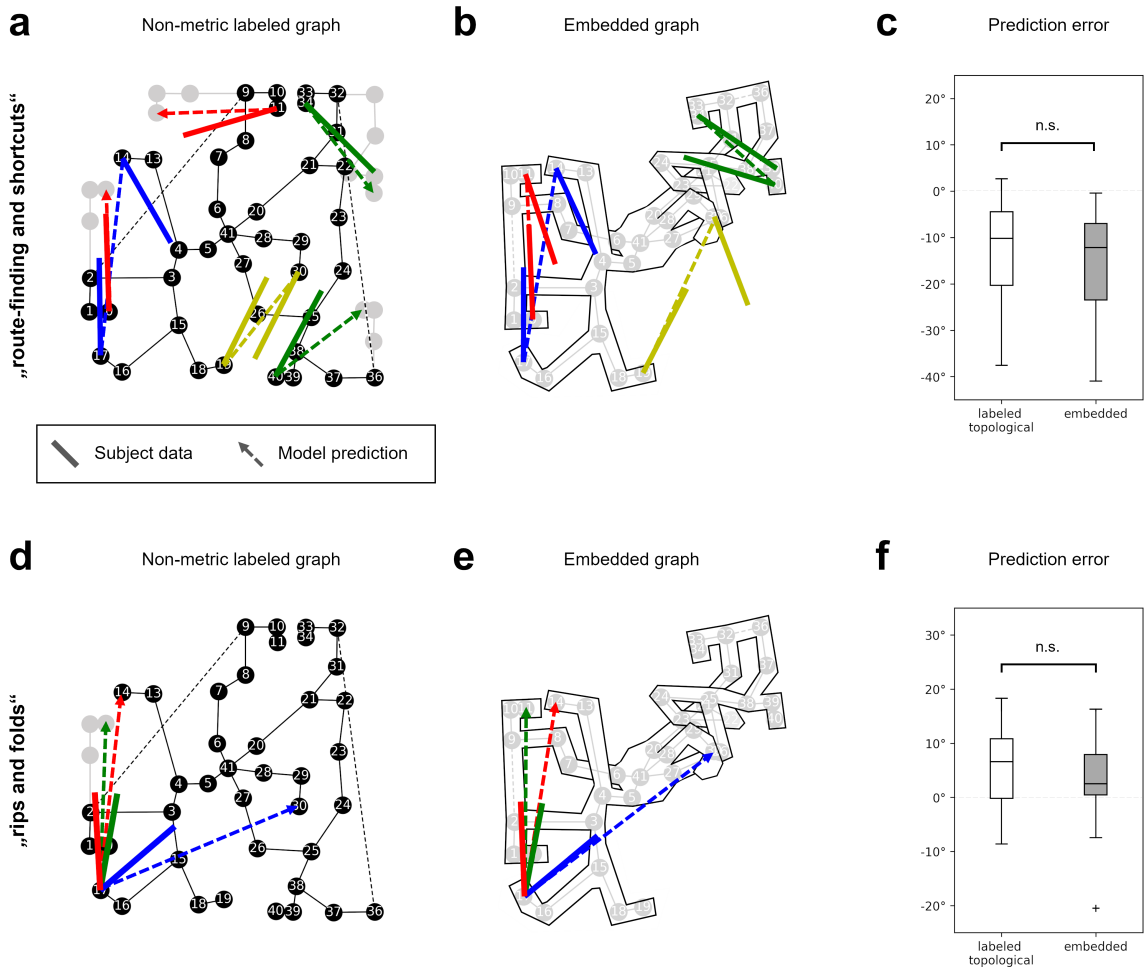


Figure A4.5: Results. (a-c): Dataset “route-finding and shortcuts”, (d-f): Dataset “rips and folds”. **(a)** Shortcut predictions of the non-metric labeled graph (dotted lines) and average subject estimates (solid lines), plotted on ground truth coordinates. The gray vertices show how the graph would continue on routes through wormholes. **(b)** Shortcut predictions of the embedded graph, lines as in (a). The subject estimates were rotated to match the local orientation of the originating maze arm. **(c)** Distribution of the prediction error. The difference between the models is not significant, i.e., they predict the data equally well. **(d)** Example shortcut predictions (dotted lines) and subject estimates (solid lines) for three of the 24 object pairs in the “rips and folds” dataset. **(e)** Shortcut predictions of the embedded graph for the same object pairs as in (d). **(f)** Distribution of the prediction error. The difference between the models is also not significant on this dataset.

of the errors was also high, with an average of 42.36° for the non-metric labeled graph model and 33.77° for the embedding. This difference was trending towards significance ($F(1, 20) = 4.07, p = .057$) with a large effect ($d = .85$).

Although dataset 2 contained many more measurements than dataset 1, there was still no significant difference between the prediction errors, i.e., the models again predicted the data equally well, with the embedding possibly capturing the within-subject variation better.

A4.4 Discussion

Using subject data from Warren et al. (2017), we compared two cognitive map models, the non-metric labeled graph and the embedded graph. We found both models predicted the data equally well, i.e., both models made prediction errors with a similar magnitude and distribution. The embedding may possibly be somewhat better at predicting the within-subject angular deviation in the rips and folds dataset, but the results did not pass the selected significance threshold at $\alpha = .05$. Given the data, we therefore found insufficient evidence to reject the null hypothesis.

Due to the non-Euclidean property of the environment, a perfect Euclidean embedding does not exist and a difference between the models must remain. It is therefore surprising that it did not lead to significantly different prediction errors.

In the original study, subjects explored the environment by walking continuous paths and thereby obtained information not only about the place-to-place distances and turns but also about the overall connectivity of the network. The conclusions drawn in Warren et al. (2017) imply that this network information is not used for the shortcut task, which is thought to be solved by vector addition along the direct path only. Here, we showed that the behavioral data are also consistent with the idea of consolidating both distance and network information in a metrically embedded graph. We thus refute the conclusion in Warren et al. (2017): it is not necessary to discard Euclidean metric properties and to reduce the representation to a non-metric framework in order to explain the observed behavior.

The main difference between the vector navigation in the labeled graph and the embedded graph suggested here lies in the treatment of repeated distance and angle measurements during prolonged navigation. Repeated measurements might simply be used to improve the estimates of distances and angles for individual labels without exploiting the constraints that these measurements impose on adjacent labels

and indeed on the entire graph. Metric embedding, in contrast, allows to make use of these constraints such that improved estimates of one edge will lead to better distance and angle estimates everywhere. In this view, the main advantage of having a metrically embedded representation of space is not so much its resemblance to a geographic map, but the possibility to integrate local and repeated measurements into a consolidated structure. The result is still a graph, but with metrically embedded vertices from which directions can be derived directly without the “mental path integration” procedure suggested by Warren et al. (2017).

Note that even an optimal metric embedding is not necessarily equivalent to the Euclidean ground truth; cognitive space is not natural space, and the internal representation may still be systematically distorted, even under normal Euclidean circumstances. This might explain poor navigational performance even after prolonged exposure to the environment (e.g., Ishikawa and Montello (2006)).

Non-metric topological and metrically embedded information may also coexist. Combined models have previously been proposed, for example, for different levels of spatial hierarchy (Couclelis et al., 1987; Meilinger, 2008), where the local Euclidean structure of individual places or regions is known but higher-level relations between different regions are encoded as a graph. For example, a local plaza may be well-represented by a Euclidean metric map, but directions to other places within the city may only be memorized as a sequence of turns. In the context of this present study, this relates to the problem of what constitutes a vertex of the graph. In our simulation, we placed vertices at all corners of the maze, but other choices are possible. A neural network model assuming metric representations within small regions and categorical knowledge of these regions themselves has been presented by Baumann and Mallot (2023a).

Topological and metric information may also be used under different environmental constraints or at different stages of exploration and familiarization (Peer et al., 2021). Initially, the environment may be encoded in terms of adjacency relations and individual routes, which then over time is consolidated in an encompassing map as the amount of information increases. This scenario is supported by reports that grid cell firing fields are initially anchored by the walls of individual compartments, but with experience extend across boundaries to encompass a larger space (Carpenter et al., 2015; Wernle et al., 2018). The embedding algorithm presented here may also be considered a support, because it describes a transformation of local position information under topological constraints into a Euclidean metric map.

Compliance with Ethical Standards

Disclosure of potential conflicts of interest: The authors declare no conflicts of interest.

Data availability statement: Data will be made available upon request.

Author contribution: **TB** Conceptualization, Formal analysis, Software, Writing; **HAM** Conceptualization, Writing.

Funding information: The research reported in this paper was carried out at the Department of Biology of the University of Tübingen. This research did not receive any specific grant from funding agencies in the public, commercial, or non-profit sectors.

Bibliography

- Batschelet, Edward (1981). *Circular Statistics in Biology*. New York, NY USA: Academic Press Inc., pp. 95–98.
- Baumann, Tristan and Hanspeter A Mallot (2023a). “Gateway identity and spatial remapping in a combined grid and place cell attractor”. *Neural Networks* 157, pp. 226–239.
- Baumann, Tristan and Hanspeter A Mallot (2023b). “Metric information in cognitive maps: Euclidean embedding of non-Euclidean environments”. Under review.
- Berens, Phillipp and Fabian Sinz (2022). *PyCircStat: circular statistics with Python*. <https://github.com/circstat/pycircstat>. Accessed: 2023-03-02.
- Brunec, Iva K et al. (2017). “Contracted time and expanded space: The impact of circumnavigation on judgements of space and time”. *Cognition* 166, pp. 425–432.
- Byrne, Roger W (1979). “Memory for urban geography”. *The Quarterly Journal of Experimental Psychology* 31.1, pp. 147–154.
- Carpenter, Francis et al. (2015). “Grid cells form a global representation of connected environments”. *Current Biology* 25.9, pp. 1176–1182.
- Chrastil, Elizabeth R and William H Warren (2013). “Active and passive spatial learning in human navigation: acquisition of survey knowledge.” *Journal of experimental psychology: learning, memory, and cognition* 39.5, p. 1520.
- Collett, Matthew et al. (1998). “Local and global vectors in desert ant navigation”. *Nature* 394.6690, pp. 269–272.
- Collett, Thomas S and Matthew Collett (2002). “Memory use in insect visual navigation”. *Nature Reviews Neuroscience* 3.7, pp. 542–552.

- Couclelis, Helen et al. (1987). “Exploring the anchor-point hypothesis of spatial cognition”. *Journal of environmental psychology* 7.2, pp. 99–122.
- Dabaghian, Yuri, Vicky L Brandt, and Loren M Frank (2014). “Reconceiving the hippocampal map as a topological template”. *Elife* 3, e03476.
- Ekstrom, Arne D et al. (2003). “Cellular networks underlying human spatial navigation”. *Nature* 425.6954, pp. 184–188.
- Foo, Patrick et al. (2005). “Do humans integrate routes into a cognitive map? Map-versus landmark-based navigation of novel shortcuts.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31.2, p. 195.
- Gallistel, Charles R (1990). *The organization of learning*. The MIT Press.
- Gillner, Sabine and Hanspeter A Mallot (1998). “Navigation and acquisition of spatial knowledge in a virtual maze”. *Journal of cognitive neuroscience* 10.4, pp. 445–463.
- Hafting, Torkel et al. (2005). “Microstructure of a spatial map in the entorhinal cortex”. *Nature* 436.7052, pp. 801–806.
- Hagberg, Aric A., Daniel A. Schult, and Pieter J. Swart (2008). “Exploring Network Structure, Dynamics, and Function using NetworkX”. *Proceedings of the 7th Python in Science Conference*. Pasadena, CA USA, pp. 11–15.
- Helmholtz, Hermann von (1876). “The origin and meaning of geometrical axioms”. *Mind* 1.3, pp. 301–321.
- Hübner, Wolfgang and Hanspeter A Mallot (2007). “Metric embedding of view-graphs: A vision and odometry-based approach to cognitive mapping”. *Autonomous Robots* 23, pp. 183–196.
- Indow, Tarow (1999). “Global structure of visual space as a united entity”. *Mathematical Social Sciences* 38.3, pp. 377–392.
- Ishikawa, Toru and Daniel R Montello (2006). “Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places”. *Cognitive psychology* 52.2, pp. 93–129.
- Kant, Immanuel (1781). *Kritik der reinen Vernunft*. Riga: Johann Friedrich Hartknoch.
- Kim, Misun and Christian F Doeller (2022). “Adaptive cognitive maps for curved surfaces in the 3D world”. *Cognition* 225, p. 105126.
- Kluss, Thorsten et al. (2015). “Representation of impossible worlds in the cognitive map”. *Cognitive processing* 16, pp. 271–276.
- Kraft, Dieter (1988). “A software package for sequential quadratic programming”. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*.
- Kuipers, Benjamin (1978). “Modeling spatial knowledge”. *Cognitive Science* 2.2, pp. 129–153.
- Mallot, Hanspeter A (2024). *From Geometry to Behavior: An Introduction to Spatial Cognition*. MIT Press.
- Mallot, Hanspeter A and Kai Basten (2009). “Embodied spatial cognition: Biological and artificial systems”. *Image and Vision Computing* 27.11, pp. 1658–1670.

- McNamara, Timothy P (1986). “Mental representations of spatial relations”. *Cognitive psychology* 18.1, pp. 87–121.
- McNaughton, Bruce L et al. (2006). “Path integration and the neural basis of the ‘cognitive map’”. *Nature Reviews Neuroscience* 7.8, pp. 663–678.
- Meilinger, Tobias (2008). “The network of reference frames theory: A synthesis of graphs and cognitive maps”. *Spatial Cognition VI*. Springer, pp. 344–360.
- Meilinger, Tobias et al. (2018). “Humans construct survey estimates on the fly from a compartmentalised representation of the navigated environment”. *Spatial Cognition XI*. Springer, pp. 15–26.
- Nadel, Lynn (2013). “Cognitive maps”. *Handbook of spatial cognition*, pp. 155–171.
- O’Keefe, John and Neil Burgess (1996). “Geometric determinants of the place fields of hippocampal neurons”. *Nature* 381.6581, pp. 425–428.
- O’Keefe, John and Lynn Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- Peer, Michael et al. (2021). “Structuring knowledge with cognitive maps and cognitive graphs”. *Trends in cognitive sciences* 25.1, pp. 37–54.
- Rolls, Edmund T and Shane M O’Mara (1995). “View-responsive neurons in the primate hippocampal complex”. *Hippocampus* 5.5, pp. 409–424.
- Sadalla, Edward K and Daniel R Montello (1989). “Remembering changes in direction”. *Environment and Behavior* 21.3, pp. 346–363.
- Squire, Larry R. and Barbara J. Knowlton (1995). “Memory, hippocampus, and brain systems”. *The cognitive neurosciences*. Ed. by Michael S. Gazzaniga. Cambridge, MA: The MIT Press. Chap. 53, pp. 825–837.
- Tolman, Edward C (1948). “Cognitive maps in rats and men.” *Psychological review* 55.4, p. 189.
- Trullier, Olivier et al. (1997). “Biologically based artificial navigation systems: Review and prospects”. *Progress in neurobiology* 51.5, pp. 483–544.
- Tversky, Barbara (1992). “Distortions in cognitive maps”. *Geoforum* 23.2, pp. 131–138.
- Virtanen, Pauli et al. (2020). “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. *Nature Methods* 17, pp. 261–272.
- Warren, William H (2019). “Non-euclidean navigation”. *Journal of Experimental Biology* 222.Suppl_1, jeb187971.
- Warren, William H et al. (2017). “Wormholes in virtual space: From cognitive maps to cognitive graphs”. *Cognition* 166, pp. 152–163.
- Wernle, Tanja et al. (2018). “Integration of grid maps in merged environments”. *Nature neuroscience* 21.1, pp. 92–101.
- Widdowson, Christopher and Ranxiao Frances Wang (2023). “Human spatial learning strategies in wormhole virtual environments”. *Spatial Cognition & Computation* 23.2, pp. 112–131.
- Yartsev, Michael M. and Nachum Ulanovsky (2013). “Representation of Three-Dimensional Space in the Hippocampus of Flying Bats”. *Science* 340.6130, pp. 367–372.

Zetsche, Christoph et al. (2009). “Representation of space: Image-like or sensorimotor?” *Spatial Vision* 22.5, pp. 409–424.