This is a penultimate draft of a chapter to appear in the edited volume Neuroscience & Philosophy (De Brigard and Sinnott Armstrong, Eds.), MIT Press. Please cite the published version.

Memory Structure and Cognitive Maps

Sarah Robins[1], Sara Aronowitz[1], Arjen Stolk

Abstract

A common way to understand memory structures in the cognitive sciences is as a *cognitive map*. Cognitive maps are representational systems organized by dimensions shared with physical space. The appeal to these maps begins literally: as an account of how spatial information is represented and used to inform spatial navigation. Invocations of cognitive maps, however, are often more ambitious; cognitive maps are meant to scale up and provide the basis for our more sophisticated memory capacities. The extension is not meant to be metaphorical, but the way in which these richer mental structures are supposed to remain map-like is rarely made explicit. Here we investigate this missing link, asking *How do cognitive maps represent non-spatial information?* We begin with a survey of foundational work on spatial cognitive maps and then provide a comparative review of alternative, non-spatial representational structures. We then turn to several cutting-edge projects that are engaged in the task of scaling up cognitive maps so as to accommodate non-spatial information: first, on the *spatial-isometric approach*, encoding content that is non-spatial but in some sense isomorphic to spatial content; second, on the *abstraction approach*, encoding content that is an abstraction over first-order spatial information; and third, on the *embedding approach*, embedding non-spatial information within a spatial context, a prominent example being the Method-of-Loci. Putting these cases alongside one another reveals the variety of options available for building cognitive maps, and the distinctive limitations of each. We conclude by reflecting on where these results take us in terms of understanding the place of cognitive maps in memory.

1 Both authors contributed equally

1. Introduction

Over the course of any given day, we are exposed to vast amounts of information, and yet our memory systems are capable of encoding and later retrieving this information. This would be difficult, if not impossible, unless the stored information were structured - that is, organized across various dimensions such as space, time, and semantic content. The use of structure to facilitate effective retrieval can be thought of as a general mnemonic strategy, both in terms of the subpersonal processes that organize the memory system and in terms of the personal-level strategies that we can use to intentionally facilitate recall of particular pieces of information (Aronowitz, 2018). Cognitive scientists interested in memory have thus long been focused on investigations of memory structure. How do we organize information and experiences so as to make subsequent retrieval possible?

A common way to conceptualize memory structures in the cognitive sciences is as a *cognitive map.* Cognitive maps, in the most literal sense, are mental representations which are structured in a way that reflects the features of real space and which aid in navigation. Grounding the structure of memory systems in this basic and general ability that is conserved across a wide range of species has obvious appeal. Cognitive maps thus offer hope of theoretical and interspecies unity, as well as the opportunity to learn more about the structure of human memory by investigating the neural systems and behavior of model organisms like rats and mice, where more extensive and precise interventions are available.

Cognitive maps also present a puzzle. The appeal to these maps begins literally: as an account of how spatial information is represented. Their intended use, however, is more ambitious; cognitive maps are meant to scale up and provide the basis for our more sophisticated memory capacities (e.g., Bellmund 2018). Our memory systems, as well as those of animals, surely represent a variety of non-spatial information. And at least in us, some of this information is richly conceptual and linguistic. The extension is not meant to be metaphorical, but the sense in which these richer mental structures are supposed to remain map-like is rarely made explicit. How precisely is this process of scaling up meant to go? How do cognitive maps represent non-spatial information? There are a range of ways that generalization and abstraction could occur, each of which comes with a unique set of empirical consequences and a distinct view of mental

representation and memory structure. Each, too, comes with a set of particular concerns and challenges. Our aim in this chapter is not to defend any particular view, but instead to provide a framework for exploring the available options.

This project is important for the neuroscience of memory because clarifying what cognitive maps represent and why has consequences for the methodology of identifying cognitive maps, the relationship between kinds of information in memory, and the relationship between memory and other forms of cognition. From a philosophical perspective, thinking carefully about cognitive maps is a window into understanding the nature of mental representation in memory and cognition more broadly. It would be an understatement to say that the nature of perceptual representations has attracted serious philosophical interest -- and yet, the corresponding question in the philosophy of memory remains understudied. We also hope that this chapter can shed light on a debate about map-like forms of representation more generally (e.g. Camp 2007, 2018; Rescorla 2009).

A few caveats: the aim of this chapter is to understand what cognitive maps are and how they are incorporated into memory research. As such, we will not start by defining a cognitive map. Instead, we'll consider empirical work that appeals to this concept, taking note of definitions given by others along the way, attempting to derive a working definition that fits at least the majority of this research. We do not intend our review of this empirical work to be exhaustive. When determining what to include, our primary focus is on the views of cognitive maps that have been developed into accounts of memory structure. We recognize, but do not discuss, the extensive literature on cognitive maps as competing models of spatial navigation and representation in animal cognition (see Bermudez 1998 and Rescorla 2017 for reviews).

We begin, in Section 2, with a survey of two traditions: first, foundational work on cognitive maps that assumes these maps represent information in a spatial structure. Second, a review of alternative, non-spatial representational structures. From the former, we identify a set of themes widely shared by proponents of cognitive maps. From the latter, we extract general lessons for accounts of cognitive structure. With this background, in Section 3, we turn to several cutting-edge projects that are engaged in the task of scaling up cognitive maps so as to accommodate non-spatial information. These projects each do so in interestingly different ways. Some kinds of spatial information may also be represented in a map-like form because they are organized

along dimensions that are substantially analogous to spatial information. In other cases, non-spatial information is represented as an abstraction from spatial information. And still other cognitive maps embed non-spatial information in a spatial structure. Putting these cases alongside one another reveals the variety of options available for building cognitive maps, and the distinctive limitations of each. We conclude by reflecting on where these results take us in terms of understanding the place of cognitive maps in memory.

2. Foundational Work on Cognitive Structures

A. Cognitive Maps as Spatial Structures

Thinking of cognitive structures in terms of cognitive maps has a long history in psychology and neuroscience. The view began as an explanation of maze running abilities in rats and, over time, has developed and changed as it has been used to capture a range of activities, from semantic knowledge structures to the navigational expertise of London taxi drivers (Collins & Loftus, 1975; Maguire, Frackowiak, & Frith, 1997). Throughout, theorists have aimed to make connections between these abilities in experimental animals and humans, but have offered subtly different accounts of why these maps exist, what's essential to their structure, and how the extension from basic neural structure to broader human competences is characterized.

Tolman (1948) is often identified as originating the idea of a cognitive map in this literature (although O'Keefe and Nadel also find historical inspiration in the work of Gulliver (1908)). Tolman's account of the cognitive map emerged from his work on maze running and spatial learning in rats - the dominant method and experimental framework in early 20th century psychology. For Tolman, cognitive maps were part of an argument that explaining the navigational abilities of rats required more cognitive, representational structure than was allowed for by the stimulus-response approach, which was dominant at the time. Specifically, Tolman documented rats' ability to learn shortcuts in mazes -- an ability inexplicable in terms of the animal's learned association with particular places in the maze as individual stimuli. Tolman further observed that rats were capable of latent or non-reinforced learning. That is, rats that were simply allowed to explore mazes while fully fed -- not receiving nor wanting any reinforcement

for their exploration -- were able to learn routes through the maze. In order to explain this behavior, Tolman argued, the rat needed to be credited with the possession of a cognitive map that provided a "field map of the environment" (1948: p. 192). Although Tolman's evidence was based in the maze running behavior of rats, he intended the notion of cognitive map appealed to in explaining this behavior to be applicable to a much wider range of cognitive creatures. Indeed, his 1948 paper was titled "Cognitive Maps in Rats and Men." The paper even concludes with a few pages of speculation on how particular features of human personality and social organization may be explicable within this framework.

Tolman's initial proposal was solidified into a theory of neural structure with the publication of O'Keefe and Nadel's *Hippocampus as a Cognitive Map* (1978). For O'Keefe & Nadel, cognitive maps were not simply a general framework for thinking about mental structure; the cognitive map was posed as a theory of hippocampal function. It was the first such proposal, and a highly systematic one, which helps to explain both the initial excitement about the idea and its lasting influence. There are further differences between Tolman's use of the term and use in O'Keefe and Nadel's framework, which serves as the basis of "Cognitive Map Theory" as its now understood. First, O'Keefe and Nadel take the notion of a map far more literally than Tolman. The claim is not that the information-processing of the hippocampus can be understood as map-like or spatial-ish, but that these cognitive maps are inherently spatial. In putting forward their theory, they make continued, explicit appeal to the Kantian idea of spatial structures as an organizing feature of cognition. These spatial maps are considered innate structures endemic to all cognitive creatures. Second, the extension of these maps to humans is not a metaphorical abstraction from the idea of a spatial map, but is instead characterized as an expansion of the kind of inputs that the spatial system can incorporate and process. This is best illustrated by their account of cognitive maps in humans: "the left hippocampus in humans functions in *semantic mapping*, while the right hippocampus retains the spatial mapping function seen in infra-humans. On this view, species differences in hippocampal function reflect changes in the inputs to the mapping system, rather than major changes in its mode of operation" (1978: p. 3).

The centerpiece of cognitive map theory is the discovery of place cells (O'Keefe & Dostrovsky, 1971): neurons in the hippocampus that fire preferentially—exhibiting a burst of action potentials—in response to a specific location in the organism's

environment. When a rat navigates a maze, for example, some place cells fire at the beginning of the maze, others at the first fork, still others at the second fork, and so on. These place cells are organized topographically, so that their collective firing pattern reflects the rat's route. After the maze has been run, the pattern is rehearsed, establishing a "map" that allows the rat to navigate this environment more easily the next time it is encountered.

The discovery of grid cells further enriches our understanding of the maps created by the hippocampal system (Hafting et al., 2005). Grid cells are found in the medial entorhinal cortex and, in contrast to hippocampal place cells, fire at multiple regularly-spaced locations in the environment. Seen over the rat's trajectory, the spatial firing patterns of these cells provide a grid-like representation of the organism's environment. Other cells select for additional elements of the map - e.g., cells that track objects, landmarks, and other agents (Hoydal et al 2018), head direction cells that fire selectively based on the way the organism's head is oriented relative to its route (Taube, Muller, & Rank 1990), and those that encode information about the distance to borders and edges (Solstad et al, 2008).

In our brief survey of this work, we want to highlight two important features of this literature as we see it. First, even though work on cognitive maps and memory structure has been done mostly with rats and mice and has focused on low-level neural structure, the intent beyond this idea has always been to make claims about the role of such maps in cognitive creatures more general. That is, the aim was not simply to move away from overly simplistic stimulus-response models of non-human animal learning, but to think about the cognitive structures available to these non-human animals in terms of a framework that would encompass cognitive processes and cognitive creatures more generally. How the cross-species and beyond-spatial generalizations of the framework are envisioned differs across particular accounts and interpretations of cognitive maps.

Second, Cognitive Map Theory remains influential and controversial. The framework continues to serve as a serious guide to inquiry into neural structure, especially with regards to the hippocampus (Bellmund et al. 2018. The view also serves as a steady target for alternative conceptions of neural structure and cognitive processing. Many of these criticisms involve claims and/or evidence that the information represented in these "maps" is non-spatial. That such interpretations of the content of

cognitive maps are available is certain. Whether this sense of cognitive maps is an extension of the original framework or an objection to it is more contentious. Answering this question depends on asking, first, how the notion of "map" should be understood. We address this in Section 3.

B. Non-Spatial Cognitive Structures

Cognitive psychology has, since its beginnings, been interested in how humans and other cognitive creatures organize their vast amounts of knowledge so as to support efficient and effective search and retrieval. Although some of these cognitive structures are referred to as "maps," in such cases the term is being stretched to non-standard or metaphorical use. In this section, we'll survey some of the foundational work on structures in cognitive science that do not seem to be map-like and do not primarily encode spatial content.

We'll start with *emergent conceptual structures*. These are ways of organizing and relating information that emerge from amassing overlapping conceptual content, and acquire their particular structures from patterns in the accumulated information. For example, this type of memory structure is often thought to support language comprehension and production. Adele Goldberg (2018) presents a view of language intended to explain the differences between the following kinds of minimal pair:

1) I'll cry myself to sleep
2) I'll cry myself asleep (p57)

The former sentence is perfectly felicitous, whereas native speakers judge the latter to be odd. Perhaps (2) is odd because it is novel or unusual, but as Goldberg notes, we are perfectly happy with unusual sentences like this:
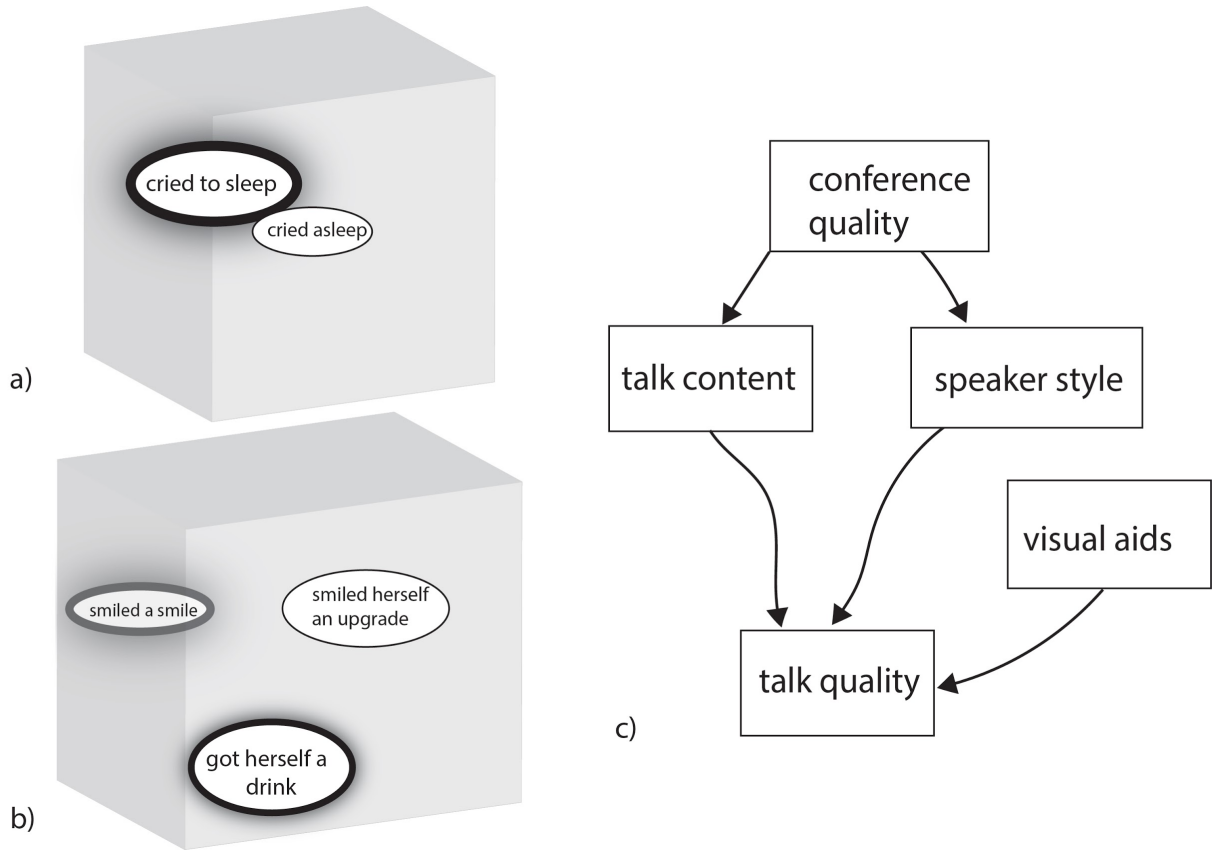
3) She'd smiled herself an upgrade (p76)

Goldberg explains what's special about (2) by appealing to the role of long-term memory organization. On her account, we encode much of the language we hear in a high-dimensional conceptual space that is structured by syntactic form, meaning (in context), phonetic features, and so on. Since there are systematic relationships between many of these features, over time clusters emerge. Bits of language that are encountered more frequently are selectively strengthened, whereas the connection between a word or phrase and its initial context is weakened if it then fails to occur in similar contexts. We have only noisy, implicit access to this space. Thus, the problem with (2) is that it is

close to a stronger competitor, (1). Conversely, (3), while unusual, does not trigger us to recall a more common alternative formulation, and so we judge it to be felicitous.

This example helps us extract several key features of an emergent conceptual structure. Over time, in the case of language, certain features of language stand out as principal components: for instance, in English, sentences with the form of the double-object construction (e.g. she (x) passed him (y) something (z)) almost always have the meaning that x causes y to receive z. This is a regularity that relates sentential form to semantic content, and crucially, Goldberg argues that this regularity arises in memory without any need to learn an explicit rule or start off with innate knowledge. Instead, the emergent conceptual structure allows for all kinds of regularities between a wide variety of features to be learned over sufficient exposure. Such a structure must therefore be (a) *highly dimensional,* in order to catch the relevant features, (b) *associative,* in order to flexibly relate these features without a prior model, and (c) *content-addressable,* in order to utilize stored information to produce responses efficiently.

A second kind of non-spatial structure is a *graphical model,* a family of models within which we'll focus on *Bayesian networks.* In the case we just considered, the regularity between the double-object construction and the type of causal-agential content was represented as a clustering or association. We neither represent the syntactic structure as dependent on the semantic content nor vice-versa; the association is not specific enough to represent anything more than organization in terms of similarity. But in a Bayesian network, relationships between features are represented as a set of conditional (in)dependencies. For example, I might encode information about academic lectures as in figure 1c.

[Figure 1: Conceptual Structures: 1a and 1b represent highly-dimensional spaces as described by Goldberg (2018) which combine semantic information with other dimensions such as syntax and phonetics with dark lines representing phrase frequency. We find "cried asleep" to be more awkward than "smiled herself an upgrade" because of its nearness to a common alternative. 1c depicts a graphical model with nodes representing event features, and edges representing probabilistic dependencies.]

In this model, the nodes are random variables and the edges represent dependencies (sets of local probability models). This allows us to assume that a node is independent of any other node, conditional on its parents: for instance, conditional on *Conference Quality*, *Speaker Style* is independent of *Talk Content*. Notice that in the above graph, an edge connects *Conference Quality* to *Talk Content* but it seems unlikely that *Conference Quality* is a cause of *Talk Content*. A narrower class of these models, causal Bayesian networks, interprets dependence and independence causally; consequently, laws of causality can be applied to the graph structures, such as

transitivity, asymmetry, and non-reflexivity. As is, this graphical representation is completely equivalent to an enumeration of the local probability models; however, when treated as a representation structure, the graphical representation can have properties not shared by the set of local models. For instance, we might search the graphical representation with an algorithm designed specifically for search in a graph, which would produce different behavior than search over other forms of representing the same information (e.g. Dechter and Mateescu, 2007). Whether or not it is true that conceptual knowledge is in fact *represented with*, in addition to *representable by*, graphical models, this hypothesis provides an interesting model of non-spatial mental structures.

Glymour (2001) analyzes a wide swath of human cognition in terms of causal Bayesian networks. These representations have three functions on his account: control, prediction, and discovery. To varying degrees, these functions could be fulfilled just as well no matter the format of the probabilistic and causal information. However, the graphical format is significant as soon as the thinker employing the models is not perfect. Graphical representations figure directly in heuristics and inductive biases, such as a preference for an explanation that appeals to fewer causes (Lombrozo, 2007). Graphical representations allow simple access to points of potential intervention (Gopnik et al, 2004). As we noted above, we can define distinctive algorithms for search over graphical representations, and both noise and lesions to the model will operate differently depending on representational format.

Thus our second class of non-spatial models, causal Bayesian networks, is used to represent all kinds of causal knowledge. These representations function to identify interventions, predict new outcomes, and enumerate new possible theories. Bayesian networks are well-suited to performing these functions because they organize information according to principles (that is, the principles of causation) that (a) apply to the entire domain and (b) align with our interests in manipulating the environment.

Emergent conceptual structures and causal Bayesian networks are both structures that have been posited as operant in memory. Neither of these structures is in any notable way spatial or especially suited for spatial information. Both of these structures are functional: they are thought to have certain features that map on to computational advantages for the thinkers who employ them. Emergent conceptual structures are more flexible than causal Bayesian networks, since the latter can only

represent causal relationships whereas the former can represent non-causal associations. Correspondingly, causal Bayesian networks can express a more complex set of relationships within the causal domain, differentiating cause and effect, and identifying potentially complex causal patterns. Emergent conceptual structures represent many features in exactly one relationship: similarity.

Considering these two structures leaves us with a few key take-aways. First, even models of abstract, domain-bridging, and perhaps distinctively human knowledge, cognitive structures are still thought to be tailored to particular functions. Second, there seems to be a trade-off between the generality of a representation (i.e. the kinds of features it could in principle represent) and its inferential power (i.e. the conclusions that can be derived from the connections among representational sub-units). When data and processing limitations are held fixed, we could either utilize a structure with more flexible (and hence weak) connections, or with less flexible (but more powerful) connecting links. This idea is fairly intuitive, following from a more general connection between flexibility and informativeness. In the case of emergent conceptual structures, we saw an advantage of flexibility at work: Goldberg's model allows speakers to track not just semantic or syntactic patterns separately, but to combine all the information we have about a string of language and thereby to learn patterns of association that cross-cut traditional linguistic categories. Causal Bayesian networks displayed one of the advantages of inferential power: by representing causal structures in graphs, we made the task of determining points of intervention vastly easier. These cases offer helpful comparisons for considering how to manage these trade-offs in characterizing the functions of spatial structures that serve as the basis for cognitive maps.

3. Cognitive Maps and Non-Spatial Information

The foregoing sections divided recent and historical work on memory structures into two categories: spatial (or spatially-grounded) cognitive maps and non-spatial cognitive structures. In this section, we'll look at how the line between them can be blurred, such that cognitive maps might be used to encode less obviously spatial information. Specifically, we ask how minimal the spatial format can be while still

leaving us with a meaningful notion of a cognitive map as a particular kind of functional memory structure.

To do so, we require a more extensive understanding of the basic notion of a map from which the idea of a cognitive map is generated. There is no canonical account of cartographic representation available, but we can provide a sketch by building off of a set of features proposed by Rescorla (2017). For Rescorla, maps 1) represent geometric aspects of physical space[2], 2) have veridicality conditions, 3) have geometric structure, and 4) are veridical only if they replicate salient geometric aspects of the region being represented. Cognitive maps are then maps in a strict sense when they consist of mental representations with these properties.

However, the definition proposed by Rescorla will not capture crucial elements of the spatial maps we've already discussed, since his definition focuses solely on synchronic, intrinsic features[3]. The kind of cognitive maps we've surveyed are also used in navigation, and interpreted and updated accordingly[4]. This addition is crucial: a representation that has all the right internal properties but is never used in navigation is not really a map -- and likewise with one that does not even potentially keep step with changes of information about the environment. Combining this functional role with Rescorla's conditions also lets us derive a fairly distinctive feature of maps, both cognitive and otherwise: we often update a piece of a map, such as a representation of the rooms on my floor, without even assessing a possible re-mapping of global relations, such as the distance between my room and Samarkand. We'll call this feature *locality.*

When extending the notion of cognitive maps to non-spatial information, we relax the definition to capture a more general (or perhaps analogical) sense of map. Most directly, 1) will always be false because the information being represented is not spatial. This will require, in turn, changes to how the veridicality conditions in 4) are understood, and what navigation might mean.

---

2 Dabaghian et al (2014) argue that hippocampal maps represent topological (i.e. ordinal) features of space rather than geometric properties such as absolute distances and angles. We suspect this difference with Rescorla is at least partly terminological. Thus we take (1) to be satisfied by the model proposed by Dabaghian et al.
3 We can see several potential ways to derive updating behaviors from Rescorla's conditions - however, the same cannot be done for navigation, since it's clearly possible for a creature which does not behave at all, let alone navigate, to have a cognitive map on his definition.
4 Camp (2007) focuses more on these dynamic factors - on her view, a map is a representational system with a semi-compositional structure that determines what we can infer, how maps can be assembled, and how updating works.

In this section, we consider three ways of extending the cognitive map. Each involves a distinct way of modifying the first condition on cartographic representations - i.e., what is being represented : first, encoding content that is non-spatial but in some sense isomorphic to spatial content; second, encoding content that is an abstraction over first-order spatial information; and third, embedding non-spatial information within a spatial context.

Before we begin, a caveat. This section is interested in representations of non-spatial information that are in some sense utilizing the map-like representations traditionally associated with space (for a description of these parameters see O'Keefe 1991). This neither entails nor follows from a relationship between the *neural-level realizers* of spatial and non-spatial representations. Howard Eichenbaum and associated scholars have long tried to challenge the cognitive map picture by pointing to non-spatial uses of neural resources thought to be part of cognitive maps (see also MacDonald et al, 2011; Aronov et al 2017; Wood et al., 1999). Thus this line of research is not a case of extending the cognitive map to encompass non-spatial information, so long as the claim is about a shared neural substrate rather than a shared representational structure.

A. Spatial-isomorphic information

Spatial-isomorphic content is a kind of content that is structured according to dimensions that functionally correspond to spatial dimensions. By functional correspondence, we mean that the regularities, limitations, and inference patterns that we commonly apply to spatial dimensions will for the most part apply to these non-spatial dimensions. For example, (Euclidian) spatial distance is symmetric -- if my office is ten feet from the coffee machine, then the coffee machine is ten feet from my office[5]. Spatial-isomorphic content, since its dimensions functionally correspond to spatial dimensions, will tend to have a "distance-like" measure that is symmetric in the same way. It seems reasonable that, were we to have a dedicated cognitive mapping system for dealing with spatial content, this system might also be used for dealing with spatial-isomorphic content.

---

5 Interestingly, path representations are not always symmetric - see Kuipers (1982) for a theoretical computer science perspective on how this asymmetry interacts with the "map" metaphor.

Constantinescu et al. (2016) offer some preliminary evidence that some of the core processes for representing spatial maps can be used for spatial-isomorphic content. Unlike related work on spatial representations of non-spatial information (e.g. Tavares et al. 2015), the authors went beyond neurally co-locating spatial and non-spatial activity. Instead, they focused on a signature of spatial representation: coding of a space into a hexagonal lattice, such that the rate of cell firing corresponds to the orientation of movement relative to the orientation of the lattice. Because strongest firing occurs at 60 degree increments in orientation, the 360 degrees of phase space are divided into six identical regions, giving rise to the lattice's hexagonal symmetry[6]. The authors looked for this hexagonal symmetry as a mark of what are sometimes called "human grid cells". Unlike the grid cells discussed in Section 2a, these neurons are not thought to be restricted to regions of the medial temporal lobe but instead are thought to occur throughout (some of) the brain regions that also form the Default Mode Network, including ventromedial prefrontal and posterior cingulate cortex. Still, previous work has associated these more distributed cells with spatial representations (e.g. Doeller et al, 2010). Rather than a strictly spatial task, Constantinescu et al. taught participants a pattern of association between the appearance of a bird and a set of symbolic cues. The bird figure varied according to neck height and leg lengths, which allowed for a representation of possible bird figures in a two-dimensional space structured by these two features. The bird-cue relationships were chosen so that each cue picked out a single region of this "bird space". The authors indeed found hexagonally symmetric responses (measured in fMRI) in a variety of Default Mode brain regions that seemed to correspond to hexagonal, grid-like representations of "bird space".

The bird space used in this study was spatial-isomorphic, since it was structured according to two dimensions (neck height, and leg length) that could be used to carve up a feature space with several space-like functional dimensions: it was a two-dimensional Euclidean space, with distance and orientation operating just as they would in a real space. Intuitively, the bird space is space-like in that it articulates a "conceptual space", but also space-like in that neck height and leg length are themselves literally spatial dimensions. However, the design of this study allows Constantinescu et

6 However, recent work in rodents (e.g. Stensola et al. 2015) has found a variety of cases where hexagonal symmetry in grid cells is distorted.

al. to differentiate between these two spatial aspects of the stimulus: because the "bird space" and the bird's features in regular space are two distinct spaces, moving through each would produce different patterns of symmetrical activation. Since the stimuli were carefully chosen to avoid passing through phase space in the same way, the observed symmetries in fMRI signal should not reflect the bird's position in visual space.

The use of hexagonal coding itself, if the authors are correct, suggests a second kind of isomorphism. Hexagonal coding is thought to be optimal for spatial representation in particular; Mathis et al. (2015), for example, present an optimal model that ranks hexagonal coding highest for spatial resolution in the two-dimensional plane. In ordinary space, we don't normally privilege one dimension over another: that is, the north-south axis is not in general more informative than the east-west axis. This allows us to value spatial resolution uniformly across the plane. But we do typically privilege those two axes over the up-down axis in navigation. These two features must be assumed in order to show the hexagonal lattice is optimal in the spatial domain. Neither feature needs to obtain in conceptual space. For instance, resolution in the bird-neck-height dimension may be more valuable than information in the bird-leg-length dimension. Were this to be true, the hexagonal symmetries observed by Constantinescu et al. would reflect a sub-optimal representation. And so we can conclude that the use of a hexagonal symmetry code either reflects (i) a genuine isomorphism between the conceptual space and real space, or (ii) a representational choice that favors spatial-isomorphism over customization to the optimal division of conceptual space.

Another kind of spatial isomorphism centers around temporal rather than conceptual structure. Researchers commonly motivate the division of a temporal sequence into parts by analogy with the division of a spatial layout into parts. For instance, Zacks and Swallow (2007) write: "For quite a while, psychologists have known that in order to recognize or understand an object people often segment it into its spatial parts (e.g., Biederman, 1987). A new body of research has shown that just as segmenting in *space* is important for understanding objects, segmenting in *time* is important for understanding events." This literature on event segmentation asks how and why we draw boundaries between events. While Zacks and Swallow take the process of segmentation to be somewhat automatic, DuBrow et al. (2017) present contrasting evidence suggesting that segmentation can be active, abrupt, and driven by top-down goals.

Is this use of space merely a helpful metaphor, or is event structure genuinely spatial-isomorphic? One genuine isomorphism comes from the local structure of both representations. That is, a ubiquitous feature of cognitive maps is their locality. While I have a clear idea of how things in my apartment complex are oriented, and a good idea of how things in the Philadelphia Museum of Art are oriented, I do not necessarily have a joint map that neatly connects the two. Kuipers (1982;2007) views this as a key starting assumption of cognitive maps even in machines: breaking a map into smaller, local maps allows the agent to remain non-committal about global connections. This locality of representation seems to hold for temporal segmentation as well. Upon hearing a story, I might build a temporal "map" of the events of my friend's adventure last week without forming any particular representation of how the details of the events she is describing fit into a temporal sequence of my own schedule last week. Locality naturally arises from the use of schemas; Baldassano et al. (2018) found that temporal boundaries in event schemas across different kinds of events had a common neural signature, provided they shared an abstract schematic structure. That is, schematic representations impose a local structure relative to the device of the schema itself (e.g. from when you enter a restaurant to when you pay the check). Anchoring event segmentation in local (temporal) structure, then, creates an abstract isomorphism with spatial maps, which are anchored in local (spatial) structures.

We could point to a long tradition locating isomorphisms between space and time, tracing at least back to Kant (1781/1787, A33/B49-50). The strength of this tradition, however, is a double-edged sword; the abundance of spatial language used in our everyday talk about time makes it hard to genuinely conceive of the capacity to represent space and the capacity to represent time as distinct. The question of isomorphism between space and time may, from this perspective, be ill-formed if the two capacities are more than accidentally linked to one another.

In summary, one way to extend the core notion of a cognitive map to non-spatial information is to treat the non-spatial information as spatial-isomorphic. These expansions are most efficient in cases where the non-spatial domain has significant regularities that mirror regularities that compose our representations of space, such as a symmetrical distance measure, roughly equal value assigned to discriminability among the dimensions on a two-dimensional plane, and a representation of related "spaces" that can be composed locally and independently.

B. Abstractions over spatial information

Another way to extend cognitive map theory - integrating work on neural-level spatial maps and cognitive-level structure - is to explore ways in which the neural systems that support cognitive maps can process and represent abstractions from spatial information. Here we consider two kinds of spatial abstraction: first, a structure where the abstraction itself is still isomorphic to the lower-order representation of space, and second, abstractions over space that are no longer spatial-isomorphic.

Michael Hasselmo (2011) has used cognitive map theory - with its place, grid, and head direction cells - to build a map-based account of episodic memory. In keeping with key themes of cognitive map theory, Hasselmo's theory is derived largely from work with rats, but is meant to provide an account of episodic memory that can scale to humans. His book-length articulation of the view is entitled *How We Remember* and the "we" here encompasses all mammals with similar hippocampal structure. Critical to Hasselmo's particular version of cognitive map theory is the idea that the hippocampus (and surrounding structures) are a phase coding mechanism, where the map-making activity of place and grid cells is integrated into maps of the environment at multiple spatial and temporal scales. That is, the hippocampus produces a series of cognitive maps, in Rescorla's (2009) loose sense, representing the environment in more or less detail, as a function of the scale imposed by the cells from which they are activated. Together, these maps represent a particular event or experience, serving as the content of an episodic memory. To support the idea of multiple maps, Hasselmo incorporates neurons from the entire hippocampus into his model, rather than focusing primarily on the dorsal portions of the hippocampus as is common in much of the literature.

Hasselmo argues that neurons across the hippocampus share the mapping function; the differences between dorsal and ventral neurons is a matter of the size of their receptive fields, not their general function. As one moves across the hippocampus, from dorsal to ventral, the receptive field size of the neurons increases. This increase in receptive field size results in a comparable increase in the scalar proportion of the map. Maps featuring place cells with the smallest receptive fields represent the organism's immediate surroundings in detail, whereas larger maps are produced by place cells with larger receptive fields, situating the experience within its (increasingly broad) spatial and temporal context. Importantly, the broadest 'maps' may remain spatial in only the

loosest or most metaphorical sense, situating the event within a social, conceptual, or experiential context.

The result is a mechanism that produces representations rich enough to support episodic remembering. The existence of multiple maps allows for a single episode to be recorded at several "scales of experience" (Hasselmo, 2008), capturing the episode as occurring not only at a particular place and time but as associated with various objects and events. For example, consider my episodic memory of walking from my campus office to the University library to return a library book this morning. On Hasselmo's view, the representation of this episode is a conjoined set of maps of the event at different scales of experience. We can think of the smallest-scale map of the event in terms of traditional cognitive map approaches - as an allocentric map of the campus, along with my route from my office to the library. But other maps associated with this episode will represent this event at different spatial, temporal, and contextual scales. The more abstract spatial maps may represent campus with relation to the part of town, city, state or continent in which I live. More abstract temporal maps will represent my route through this map as part of my schedule for the day, or schedule for the week, or activities characteristic of this time in the academic year. Further contextual maps will also be available, where the items represented in the map situate the landmarks along the route on different contextual scales - e.g., this trip to the library as a stage in a particular research project; trees along this route at this time of the year; campus construction at this time, etc.

Hasselmo's model proposes that the cognitive map system can process increasingly abstract characterizations of space and time that can then serve as the content for more elaborate and higher-order episodic memories. However, his hierarchical picture would seem to preserve some degree of structural similarity between levels. On the other hand, Behrens and colleagues (2018) also provide an account of abstraction from first-order spatial information, in the form of eigenvectors corresponding to transformations between first-order spatial (and non-spatial) environments. Unlike a hierarchical, nested representation of experience, an eigenvector is an abstraction that does not share a structure with its first-order counterparts. Eigenvectors fall into a broader class discussed by Behrens et al. (2018) including inductive biases and factorizations -- these are all features applying to a set of

environments or state spaces that aid in learning, but seem to require additional representational resources.

The authors argue for a common set of abstractive capacities operating over both spatial and non-spatial representation, which would utilize higher-order features to drive first-order prediction and planning. Presumably, whatever representational resources would be needed to supplement first-order maps with these higher-order features must be integrated tightly with the first-order maps themselves -- Behrens et al. provide a few suggestions, but how this integration might work is still very much an open question.

C. Embedding non-spatial information in a spatial format

A third way to expand spatial representations would be to keep the spatial structure intact and embed non-spatial information within that structure. The most prominent example of memory success, the Method-of-Loci (MoL), can be understood as organizing information in this way. In the MoL, subjects typically memorize a list of unstructured items such as random words or phone numbers by imagining these items in a structured environment such as a familiar childhood walk through the neighborhood. The MoL thus appears to be a way of using a useful feature of spatial memory to store non-spatial content. The explicit process in the MoL involves two stages: a strategy for encoding items through visualization, and a strategy for retrieving items through a parallel visualization. For example, I would encode a list by imagining a walk through my childhood home, and then later recall the items by imagining walking through the home again and picking up the items.

Questions about how and why mnemonic structure works have received very little attention from memory theorists and scientists. The scant evidence that exists is, however, suggestive and intriguing and invites us to ask more detailed questions about memory structure. Both the testimony of memory champions and some preliminary studies of expert performance reveal that success in the use of mnemonics does not require any particular level of intelligence or distinct cognitive skills. Instead, the key to success using mnemonic techniques is simply practice (Ericsson, 2003; Wilding & Valentine, 1997). These behavioral reports are complemented by neuroimaging studies indicating that those who use these techniques regularly differ from controls in functional, but not structural brain features (Maguire et al., 2003; Raz et al., 2009). A recent and intriguing paper showed that, after only 6 weeks of training, cognitively

matched novices exhibited the same functional changes seen in memory champions (Dresler et al., 2017). Similarly, Yoon and colleagues (2018) have just shown that a person trained to increase their digit span to over 100 digits thirty years ago has retained many of the associated skills despite undertaking no training in the meantime.

In the majority of these cases, the items to be memorized do not have an interesting non-spatial structure themselves (e.g., digit span, presentation order of a shuffled deck of cards). However, looking more closely at the history of this mnemonic technique reveals that it has also been used for semantically structured information. In *De Oratore*, Cicero recommended the technique for memorizing speeches, lyric poems, and the like. The Classicist Minchin (2001) argued that Homer used this method to compose and perform the *Iliad* and *Odyssey*. In the Middle Ages, it was common for monks to use this technique to memorize the Bible's 150 Psalms, as well as long passages from other scholarly texts (Carruthers 2008). This is continued in some forms of contemporary use, as when students use such techniques to help them remember conceptually blocks of information (Kerr & Neisser, 1983; Roediger, 1980; Wang & Thomas, 2000). The ability to achieve additional mnemonic gains by translating information that is already semantically or conceptually structured into a spatial format suggests that there is something about spatial structure *in particular* that is useful for memory retrieval.

We have surveyed three ways in which spatial maps might be "scaled up" to accommodate non-spatial information. First, spatial structures might be re-purposed to represent non-spatial content with a suitably isomorphic structure: for instance, a conceptual "space" can be modeled as a two-dimensional Euclidean plane. Second, spatial structures might be used to represent abstractions over spatial information, whether in the form of higher-order but still spatially structured representations, or with summary statistics that are to some degree represented jointly with traditional maps. Third, non-spatial information might be embedded in a spatial map, overriding or supplementing non-spatial structure with an exogenous spatial structure.

Taking stock of these results in light of the two traditions discussed in Section 2, we can draw two key conclusions. First, following on Tolman's original essay, the idea of seeing how far the basic components of the cognitive map can be stretched to fit non-spatial content is still very much an open research question. Second, the trade-off

between flexibility and inferential power that we observed in the case of conceptual representations characterizes the question of expanding the cognitive map as well. The more we stretch the notion, the more we reduce inferential power in favor of flexibility.

4. Why extend the cognitive map?

The work that we have surveyed in this essay testifies to a persistent interest in understanding the edges of applicability of the concept of a cognitive map. Given the difficulties of extending the cognitive map concept to non-spatial information, why are researchers pursuing this program? What benefits come from thinking about the structure of mental representations within this framework?

One reason to see how far the cognitive map concept stretches ties cognitive functions to evolutionary history. Neuroscientists have long argued for functional and anatomical homologies across a broad range of mammalian species (e.g., Clark & Squire, 2013, but see also Zhao 2018). As a theory of hippocampal function, cognitive maps offer the potential for an explanation of navigational and cognitive systems across all species in which these basic anatomical structures are preserved. A cognitive map, in the core spatial context, links a somewhat general capacity for long-term spatial memory to a more obviously adaptive capacity for spatial navigation. Likewise, a model on which some kinds of conceptual memory utilize a cognitive map provides a potential basis for understanding the evolution of conceptual memory. Viewed from this perspective, the cognitive map project embodies an explanatory aim: not just to have accurate models of mental faculties, but to find models that explain how a creature with our lineage could have developed these faculties. Under this aim, the three types of extension we discussed have divergent implications. Spatial-isomorphic extension is a way of re-using structures that were tailored for a particular spatial purpose for new tasks that happen to have similar structure. Extension by abstraction, on the other hand, is a richer ability applied to the very same core domain. In comparison with these two, embedding non-spatial content in space appears more like a culturally developed trick that makes clever use of evolutionarily adapted structures. In short, there is a story to tell for each way the cognitive map could have become extended, but there are important differences in how this story goes in each case.

Another reason to persist with cognitive maps may be a general commitment to parsimony across cognitive systems. An enriched and extended notion of cognitive maps could provide the representational basis for a broad range of cognitive processes - providing a competitor to the logically structured language of thought, as developed primarily by Fodor (1975, 1987) and Pylyshyn (1984)(see Rescorla 2009 for such a proposal). The more forms of memory, and cognition more broadly, that can be understood from within this basic framework, the more streamlined and unified our understanding of cognitive systems becomes. This idea of parsimony as part of the explanatory aim in understanding the mind is notably different from a concern with parsimony within the mind as a matter of representational efficiency, though both might motivate an interest in conserving structures across cognitive domains.

As we noted in Section 2b, however, parsimony has consequences for the tradeoff between flexibility and informativeness. That is, the more cognitive subsystems employ the same structures, the more flexibly (and thereby less informative) the structures must be. Again, much depends on which form of generalizing to the non-spatial we adopt. Under spatial-isomorphic extension, there is a genuine form of parsimony in applying similar structures to isomorphic domains. But the utility of this approach will be limited to the number of domains that share these structures. There is little reason to hope or expect that *all* domains will be spatial-isomorphic in the requisite ways. The same argument might be harder to leverage for the other two kinds of extension. The discussion of the Hasselmo and Behrens et al. models revealed that abstractions of spatial information might be spatial-isomorphic, but need not be. Embedding non-spatial information in maps likewise would not be predicted by the idea that related functions should in general utilize related structures. Instead, to motivate embeddings of this kind, we'd need a much stronger notion of parsimony: for example, that an explanation with fewer structures should be preferred even when it explains across distinct functions. This stronger form of parsimony would of course undermine the previous rationale, under which we prefer a one-to-one fit of structure to function, since it encodes a preference for a one-to-many fit.

In the cognitive science tradition, memory has long served as one of the best candidates for inter-level integration. It is often, literally, the textbook model of how the various cognitive sciences can all contribute to a shared understanding of a particular cognitive process (e.g., Bermudez, 2014). Similarly, philosophers of neuroscience use

memory as a case study in mechanistic levels of explanation (Craver, 2007). The appeal of cognitive map theory can be understood as a commitment to this aspirational model. The origins of cognitive map theory reflect this - O'Keefe and Nadel were inspired by the relatively simultaneous discovery of place cells, LTP and other cellular and molecular processes in the hippocampus and the observation of the patient H.M.'s loss of episodic memory in response to hippocampal damage during neurosurgery. The promise of inter-level integration relies on a shared understanding of the cognitive process at each of the cascading levels in the model. Each version of cognitive map theory's extension offers a different account of the system whose levels of understanding and inquiry are being integrated. On the spatial-isomorphic interpretation of the extension, the system is one for representing regular and symmetric structures, spatial or otherwise. On the abstraction approach, spatial navigation is the core cognitive function or ability of this inter-level system. On the embedding approach, it's declarative remembering that is central to the inter-level project. In short, inter-level integration in cognitive science requires an understanding of the cognitive ability/system/function of interest, from which the process of decomposition and unification of distinct methodological approaches can begin. Such integration may be possible on each of these three forms of cognitive map theory, but they will each offer a distinct account of the cognitive ability being integrated.

These three reasons *why* one might be interested in expanding the cognitive map to non-spatial domains thus push us in different directions when it comes to *how* to extend the map. But beyond that, they also reveal a puzzle inherent in locating the cognitive map in memory. That is, to what extent is this information structure a proper part of memory? Given that alternative, non-spatial cognitive structures - like those explored in Section 2b - can also be understood as accounts of memory structure, the need to defend the reliance on cognitive maps, however extended, becomes more critical.

5. Conclusion

What is a "cognitive map"? This concept is pervasive but hard to define comprehensively: at its core, a cognitive map represents the environment by taking advantage of the structure of space. As a metaphor, it offers a way of unifying our understanding of the representational structures endemic to cognition, giving a sense of

shared function across a range of species and abilities. Insofar as cognitive maps remain an active research program, more attention should be devoted to the conceptual space between the literal and metaphorical versions.

The concept of a cognitive map is deeply intertwined with contemporary research on memory. In this chapter, we've surveyed a series of recent attempts to extend the concept to cover representations that are not obviously spatial. These projects, seen in the light of historical developments in the theory of spatial memory and cognitive structure, reveal both the ambitions and limitations of research into general-purpose representations.

References

Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature, 543*, 719-722.

Aronowitz, S. (2018). Memory is a modeling system. *Mind & Language.* doi:10.1111/mila.12220

Baldassano, C., Hasson, U., & Norman, K. A. (2018). Representation of real-world event schemas during narrative perception. *Journal of Neuroscience*, *38*(45), 9689-9699.

Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, *100*(2), 490-509.

Bellmund, J. L. S., Gardenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating

cognition: spatial codes for human thinking. *Science, 362.*

Bermudez, J. L. (1998). *The Paradox of Self Consciousness.* Cambridge, MA: MIT

Press.

Bermudez, J. L. (2014). *Cognitive Science: An Introduction to the Science of the Mind,*

2nd edition. Cambridge University Press.

Camp, E. (2007). Thinking With Maps*. *Philosophical Perspectives, 21*(1), 145-182.

doi:10.1111/j.1520-8583.2007.00124.x

Camp, E. (2018). Why Maps are Not Propositional. *Oxford Scholarship Online.*

doi:10.1093/oso/9780198732570.003.0002

Carruthers, P. (2008). Meta-cognition in animals: a skeptical look. *Mind & Language,*

23(1), 58-89.

Clark, R. E., & Squire, L. R. (2013). Similarity in form and function of the

hippocampus in rodents, monkeys, and humans. *Proceedings of the National*

Academy of Sciences, *110*(Supplement 2), 10365-10370.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic

processing. *Psychological Review, 82*(6), 407-428. doi:10.1037/0033-295x.82.6.407

Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. (2016). Organizing conceptual

knowledge in humans with a gridlike code. *Science*, *352*(6292), 1464-1468.

Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of*

Neuroscience. Oxford University Press.

Dabaghian, Y., Brandt, V. L., & Frank, L. M. (2014). Reconceiving the hippocampal map as a topological template. *eLife*, *3*, e03476. doi:10.7554/eLife.03476

Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, *463*(7281), 657.

Dresler, M., Shirer, W. R., Konrad, B. N., Müller, N. C., Wagner, I. C., Fernández, G., ... & Greicius, M. D. (2017). Mnemonic training reshapes brain networks to support superior memory. *Neuron*, *93*(5), 1227-1235.

Ericsson, K. A. (2003). Exceptional memorizers: made, not born. *Trends in Cognitive Sciences, 7,* 233-235.

Glymour, C. N. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology.* MIT Press.

Goldberg, A. E. (2019). *Explain me this: Creativity, competition, and the partial productivity of constructions.* Princeton University Press.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychological review, 111*(1), 3

Gulliver, F. P. (1908). Orientation of Maps. *Bulletin of the American Geographical Society, 40*(9), 538. doi:10.2307/198344

Hafting, T., Fyhn, M., Molden, S., Moser, M., & Moser, E. I. (2005, 06). Microstructure

of a spatial map in the entorhinal cortex. *Nature, 436*(7052), 801-806.

doi:10.1038/nature03721

Hasselmo, M. E. (2008). Grid cell mechanisms and function: contributions of entorhinal

persistent spiking and phase resetting. *Hippocampus*, *18*(12), 1213-1229.

Hasselmo, M. E. (2011). *How we remember: brain mechanisms of episodic memory.*

MIT press.

Høydal, Ø A., Skytøen, E. R., Moser, M., & Moser, E. I. (2018, 03). Object-vector

coding in the medial entorhinal cortex. doi:10.1101/286286

Kerr, N. H., & Neisser, U. (1983). Mental images of concealed objects: New evidence.

Journal of Experimental Psychology: Learning, Memory, and Cognition, *9*(2), 212.

Kuipers, B. (1982). The" map in the head" metaphor. *Environment and Behavior*, *14*(2),

202-220.

Kuipers, B. (2007). An intellectual history of the spatial semantic hierarchy. In *Robotics

and cognitive approaches to spatial mapping* (pp. 243-264). Springer, Berlin,

Heidelberg.

Lombrozo, T. (2007, 11). Simplicity and probability in causal explanation. *Cognitive

Psychology, 55*(3), 232-257. doi:10.1016/j.cogpsych.2006.09.006

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal

"time cells" bridge the gap in memory for discontiguous events. *Neuron, 71,* 737-749.

Maguire, E. A., Frackowiak, R. S., & Frith, C. D. (1997, 09). Recalling Routes around London: Activation of the Right Hippocampus in Taxi Drivers. *The Journal of Neuroscience, 17*(18), 7103-7110. doi:10.1523/jneurosci.17-18-07103.1997

Maguire, E. A., Valentine, E. R., Wilding, J. M., & Kapur, N. (2003). Routes to remembering: the brains behind superior memory. *Nature Neuroscience, 6,* 90-95.

Mathis, A., Stemmler, M. B., & Herz, A. V. (2015). Probable nature of higher-dimensional symmetries underlying mammalian grid-cell activity patterns. *Elife, 4,* e05979.

Minchin, E. (2001). *Homer and the resources of memory: some applications of cognitive theory to the Iliad and the Odyssey* (p. x). Oxford: Oxford University Press.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map.* Clarendon.

O'Keefe, J., & Dostrovsky, J. (1971, 11). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research, 34*(1), 171-175. doi:10.1016/0006-8993(71)90358-1

Raz, A., Packard, M. G., Alexander, G. M., Buhle, J. T., Zhu, H., Yu, S., et al. (2009). A slice of pi: an exploratory neuroimaging study of digit encoding and retrieval in a superior memorist. *Neurocase, 15,* 361-372.

Rescorla, M. (2009). Cognitive maps and the language of thought. *British Journal for the Philosophy of Science, 60,* 377-407.

Rescorla, M. (2017). Maps in the head? In K. Andrews and J. Beck (eds.) *Routledge Handbook of Philosophy of Animal Minds*, Routledge, pp. 34-45.

Roediger, H. L. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition, 8*(3), 231-246.

Solstad, T., Boccara, C. N., Kropff, E., Moser, M., & Moser, E. I. (2008, 12). Representation of Geometric Borders in the Entorhinal Cortex. *Science, 322*(5909), 1865-1868. doi:10.1126/science.1166466

Stensola, T., Stensola, H., Moser, M. B., & Moser, E. I. (2015). Shearing-induced asymmetry in entorhinal grid cells. *Nature, 518*(7538), 20

Taube, J. S., Muller, R. U., & Rank, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. *Journal of Neuroscience, 10*, 420-435.

Tavares, R. M., Mendelsohn, A., Grossman, Y., Williams, C. H., Shapiro, M., Trope, Y., & Schiller, D. (2015). A map for social navigation in the human brain. *Neuron, 87*(1), 231-243.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review, 55*(4), 189-208. doi:10.1037/h0061626

Wang, A. Y., & Thomas, M. H. (2000). Looking for long-term mnemonic effects on

serial recall: the legacy of Simonides. *American Journal of Psychology, 113,* 331-340.

Wiener, S. I., & Taube, J. S. (2005). *Head direction cells and the neural mechanisms of spatial orientation.* MIT Press.

Wilding, J. M., & Valentine, E. R. (1997). *Superior memory.* Psychology Press.

Wood, E. R., Dudchenko, P. A., & Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature, 397,* 613-616.

Yoon, J. S., Ericsson, K. A., Donatelli, D. (2018). Effects of 30 years of disuse on exceptional memory performance. *Cognitive Science, 42,* 884-903.

Zhao, M. (2018). Human spatial representation: what we cannot learn from studies of rodent navigation. *Journal of Neurophysiology, 120,* 2453-2465.