Where Am I?
The Cognitive Architecture of Spatial Reorientation

Alexandre Duval

Thesis submitted for the degree of
Doctor of Philosophy
(Philosophy)

The University of Sheffield
Faculty of Arts and Humanities
Department of Philosophy

April 2019

Abstract

Navigation is of fundamental importance to humans, just as it is for other species. And, like most other animal species, we possess a number of distinct navigational processes. This thesis examines navigation, focusing particularly on the widely studied phenomenon of reorientation following disruption to spatial behavior. In typical reorientation experiments, subjects rely on the three-dimensional surface layout of an environment to find a desired goal following disorientation, and they do so to the exclusion of other important spatial cues. An influential explanatory framework aims to account for such findings by holding that subjects possess a modular mechanism known as *the geometric module*, which only operates on geometric information about three-dimensional extended surfaces. This thesis provides a sustained defense of this framework and develops a new type of geometric-module theory of reorientation. I begin by making the case that, if the general geometric-module framework is right, it has deep implications for two foundational debates in philosophy of psychology: the debate about the nature of mental representations and the debate about the structure of the mind. I then address the two most pressing challenges against the framework. The first challenge comes from what I call *the explanatory inflexibility objection*, which holds that the geometric-module framework simply does not have the required flexibility to deal with evidence that non-geometric cues can affect subjects' search behavior in some experimental contexts. The second challenge arises from an alternative explanatory framework, the *view-matching framework*, which aims to explain subjects' behavior in reorientation experiments by appealing to *snapshots*, stored representations of the subjects' two-dimensional retinal stimulation at specific locations. In answering these two challenges, I put forward a new type of geometric-module theory which has stronger implications for debates in philosophy of psychology than standard geometric-module models.

Word count (main text and footnotes): 55,965

Acknowledgments

Many people have helped me get to a point in my life where I am ready to submit a PhD thesis. First, I would like to thank my parents for all the support they have given me in my studies at the PhD-level, but also at any level below that going back to kindergarten. I certainly couldn’t have gotten to this point without them. My sister, my brother, my brother-in-law, my sister-in-law and my aunt Pierrette have also played an important part in my intellectual and academic development, and they have helped in numerous ways since I started the PhD. I would also like to thank them.

I owe an enormous debt of gratitude to my two supervisors, Steve Laurence and Luca Barlassina. They provided constant encouragement and invaluable feedback at many difficult stages in the writing of this thesis and the academic publication process. There is no telling just how much their comments have improved the dissertation that follows. It has been a great pleasure to have them as my supervisors. I am also very grateful to Steve, my primary supervisor, for the inordinate amount of time he spent poring over drafts of mine.

I would like to thank three anonymous reviewers for the journal *Cognition* as well as one of its editors for extremely helpful feedback on a paper published in the journal which comprise most of the material included in Chapters 3, 4 and 5. Their comments on the paper have led to substantial additions and improvements in these chapters as a knock-on effect.

On a different note, I am extremely grateful to have been part of the postgraduate community of Sheffield University’s Department of Philosophy for four years. I can’t name here every MA and PhD students from the department with whom I’ve had stimulating conversations over the years and who have provided support in some way or other — the list would simply be too long. But I wish to thank everyone who was there at the same time as me and who played a part in making the department the welcoming place that it is. However, I would like to single out the following people for the role they played in my daily academic life for four years: Miklos Kurthy, Francesco Antilici, Isela González Vázquez, Lewis Brooks, Damiano La Manna, Adriana Clavel, Anton Eriksson, Neri Marsili, James Lewis, Ahmad Fattah, María José Pietrini and Giulia Casini. Special thanks to Miklos and Francesco for many highly entertaining conversations about philosophy of psychology, for our reading groups, and for very helpful feedback on various parts of this thesis.

I also wish to thank the following people outside the philosophy department who provided support, encouragement and many engaging conversations at various points in the last four years: Charles Côté-Bouchard, Cyndie Caron, Naïla Payment, Pierre Lavoie, Ryo Yokoe, Alexandre Romano, Louis-Olivier Lortie and Benoit Proulx.

Finally, I would like to thank Professor Kate Jeffery and members of her lab at UCL's Institute of Behavioural Neuroscience for welcoming me in their journal club for a few months in the 2016-2017 academic year. They provided substantive critical feedback on a presentation of an early draft of Chapter 4, as well as answering various questions about their fascinating work on the neurophysiology of spatial navigation. Visiting the lab and attending the journal club proved to be an incredibly motivating and rewarding experience.

The completion of this thesis was supported by the Social Sciences and Humanities Research Council of Canada (through Doctoral Fellowship #752-2014-1195) as well as a Faculty Scholarship from the University of Sheffield.

Contents

Chapter 1

Why Philosophers Should Care About Reorientation Experiments and the Hypothesis of a Geometric Module

1. Spatial Navigation and Philosophy of Psychology

Consider the following feats. A bee, which was captured on its way back to the hive and suddenly released in an environment located much further away in a different direction, heads straight towards the hive a few seconds following its release (Cheeseman et al., 2014). A chimpanzee, which has seen an experimenter hide pieces of fruit at 18 different locations following a zigzagging path in a large enclosure, recovers each of the 18 pieces of food when later released in the enclosure (Menzel, 1973). Moreover, it does so by following a different and often more efficient path than the experimenter took. A human being, who has spent most of her day working in a previously unknown building chooses a very efficient walking path to return home even though she has no visual access to her home from that building.

Spatial navigation encompasses all these types of feats, and much more. As I will use the expression, *spatial navigation* is the process by which an agent moves its body through one or more environments in an (often successful) attempt to get to a desired object or practically relevant location, such as a nest, a food source, a water source or a hiding place. Of course, agents sometimes fail to reach a specific desired object or location. We have all experienced, or at least heard of someone, getting lost in the woods while hiking. But despite occasional navigational failures, human and non-human animals are remarkably efficient at moving their body through physical space to get what they want or where they want to be. And this remarkable efficiency tends to be forgotten in everyday life because we often know immediately and intuitively how to behave in order to reach a certain desired location. As soon as we reflect on it, we know which direction to go in order to get back home from work, or to find the bathroom in our house.

But the efficiency with which humans and animals navigate the world raises a number of highly complex scientific questions, many of which have important implications for foundational issues in philosophy of psychology. In fact, I think that we can even make a case that investigating human and non-human animals' ability to navigate the world — and in particular investigating this ability through the assessment of theoretical models developed in reaction to specific experimental work on spatial navigation — is uniquely suited to make progress on *two foundational debates* in philosophy of psychology.

The first debate presupposes the truth of *representationalism*, the thesis that human and non-human minds contain internal representations — namely, mental objects or events which bear semantic relations to the world, mental objects or events that aim to capture through their own structure various aspects of the world. Philosophers have proposed various high-level, theoretical arguments either in favor or against this thesis over the last fifty years or so (e.g., Churchland, 1981; Dennett, 1987; Egan, 2011; Fodor, 1981). But, if we assume the truth of representationalism, a new and important set of questions arises. Such questions constitute the heart of the first debate that we will consider:

> **The nature of mental representations** — What kind of properties do human and non-human animals' mental representations have? What sort of information about the world do these representations encode? And how do they encode it?

The most popular way of approaching these questions in recent years has been to ask how far away from sensory input — how abstract — representations are in their content. At one extreme, some people hold that mental representations only encode low-level perceptual properties, like *being red* or *occupying 50-degree arc in my visual field*, along with motor properties (Barsalou, 1999; Prinz, 2002). On their view, higher cognitive capacities, like doing arithmetic or making up your mind on how to vote on an upcoming election, are only subserved by perceptual and motor representations. At the other extreme, some philosophers hold that human minds contain, along with such lower-lever representations, mental representations that encode highly abstract properties of the world as such (Fodor, 1975, 2008), properties like *being a group that contains three entities*, *being a good person*, and *being my home*. Some researchers even hold that most non-human species possess representations that encode abstract properties as such (Gallistel, 1990), though which properties those are exactly remains highly contentious.

Few philosophers have discussed the relevance of research on spatial navigation to the debate about the nature of mental representations (for some exceptions, see Carruthers, 2015; Rescorla, 2009). But it is clear that the research directly pertains to the debate. Psychologists, neuroscientists and roboticists who work on spatial navigation have provided a large number of theoretical models to account for various aspects of efficient spatial navigation, and perhaps the biggest fault line among the models turns on the issue of whether a given model holds that agents' navigational decisions are based on mental representations that encode metric properties and relations of salient objects and surfaces in visited environments. Models reliant on representations that encode such metric properties and relations (e.g., Gallistel, 1990; O'Keefe & Nadel, 1978) directly support the second position just outlined. Metric properties and relations of surfaces are quite abstract. They include the length, width, height and curvature of surfaces, as well as the angles and distances between distinct surfaces or the boundaries at which they meet. The property of *being a surface of length L*, for example, cannot be represented directly by the retinal stimulation or any other type of sensory input. Representations that encode metric properties and relations of salient objects and surfaces in detail would reproduce many aspects of the physical maps that humans use for navigation.

Now consider models that reject such representations. Models of that type almost always emphasize the importance of appealing to *snapshots*, stored representations of the subjects' two-dimensional retinal stimulation at specific locations, to account for various empirical results about spatial navigation (e.g., Cartwright & Collet, 1987; Cruse & Wehner, 2011; Möller & Vardy, 2006). The biggest challenge for such models is to make sense of how subjects can be so efficient at navigation — to explain (e.g.) how animals that do not have direct visual access to a known goal location generally manage to return to it when they so desire, often following paths they have never gone on before. This may sound like an impossible challenge. Yet, proponents of such models have been remarkably ingenious in recent years at developing methods to address it. Of course, they need to appeal to more than *just snapshots* to do so. But they have argued, powerfully in many cases, that we can account for extant empirical results without positing representations that encode bona fide metric properties of objects and surfaces as such. So, if these models turn out to best describe the cognitive processes involved in spatial navigation, they would provide inductive support for the view that animals' representations only encode low-level perceptual and motor properties.

Thus, in either case (i.e., regardless of which of the two types of models wins the day), research on spatial navigation will have important implications for the debate about the nature of mental representations.

Now, the second foundational debate that we will consider centers on the following questions:

> **The structure of the mind** — How are human and non-human animal minds structured? Are they composed of a multitude of independent cognitive processes that operate according to different functions and internal rules? Or are they, rather, composed of a few types of processes that can flexibly deal with information from a large variety of cognitive domains and that cooperate extensively to bring an agent to act in the world?

These questions have been at the center of philosophical inquiry in one form or other for an extremely long time. Plato, for example, took a stance on these important issues in *The Republic* by proposing a 'tripartite' account of the mind according to which human minds are composed of three main faculties: reason, spirit and appetite. In its modern incarnations, however, the debate about the structure of the mind has focused largely on the concept of *modularity* that arose from the work of Jerry Fodor. In his first book on the topic (Fodor, 1983), he held that a given a cognitive mechanism is *modular* just in case it possesses nine specific properties. Here, I will follow the subsequent literature, as well as Fodor's own later work (Fodor, 2000), by focusing on the two of the nine properties previously identified: *domain-specificity* and *encapsulation*.[1] On the relevant notion of modularity that arises from this subsequent literature and that will serve as the basis of the discussion in this thesis, a cognitive mechanism is *modular* (to some interesting degree) just in case it is both domain-specific and encapsulated (to some interesting degree).

Consider sound-based language perception for example, a capacity which Fodor (1983) takes as his main case study. Now, a cognitive mechanism is *domain-specific* to the

---

[1] See Samuels (2006, 2011) for important discussions of the other seven properties identified by Fodor and of why these two properties are the most theoretically interesting.

extent that it can only deal with very specific type of task or input.[2] According to Fodor, humans possess a mechanism for sound-based language perception which is highly domain-specific because it only deals with the problem of inferring phonological properties of certain noises perceived by the agent (namely, verbal utterances). A cognitive mechanism is *encapsulated*, then, to the extent that it only has access to a limited amount of information contained in the agent's mind once it is turned on. On Fodor's view, the mechanism for sound-based language perception is highly encapsulated because it ignores most of the agent's stored information about the world when it estimates the phonological properties of utterances. It only relies on some specialized, internal rules and information about how certain sequences of sound frequencies correspond to specific phonemes. The mechanism performs its operations without paying attention to the agent's information about (e.g.) the current temperature, the time of day, or even about the speaker's larger communicative intentions.

Issues pertaining to language perception aside, the notion of modularity can be used to formulate two radically opposed positions on the second debate, the debate about the structure of the mind. At one end of the spectrum, we have the *massive modularity thesis* — the view that perceptual and higher cognitive processes are realized largely by a significant number of mechanisms that are modular to a strong degree (e.g., Sperber, 1994, 2002; Carruthers, 2006b). Higher cognitive processes are processes that deal with belief formation and action planning. Proponents of the thesis regard higher cognition as structured in terms of a set of relatively independent systems for dealing with distinct aspects of the world (e.g., social exchanges and commitments among conspecifics, emotions and mental states of conspecifics, nutritional value of various food types, spatial arrangement of objects and surfaces).

At the opposite end of the spectrum, we have the view that perceptual and higher cognitive processes are realized by a few domain-general and/or unencapsulated mechanisms. This is a view that is espoused, for example, by radical proponents of positions that fall under labels like 'empiricism', 'associationism', 'connectionism', 'Bayesianism' and 'neural reuse' (e.g., Anderson, 2014; Buller, 2005; Prinz, 2006). For example, in line with this view, some proponents of Bayesianism hold that the same

---

[2] See, for example, Samuels (2006) and Carruthers (2006a, ch. 1) for a more precise, formal approach to defining *domain-specificity*. The intuitive characterization provided here will be sufficient for our purposes.

mathematical principles drawn from Bayes’ work on probability apply across all of perception and higher cognition and that they allow integration of information at multiple levels (Clark, 2016).

Importantly, the two foundational debates just presented are conceptually perpendicular. It is true that people who defend the existence of abstract, amodal representations often endorse some form of massive modularity (Carruthers, Sperber; Tooby & Cosmides, 1992), while people who reject the existence of such representations often tend to see the mind as underpinned by a few domain-general and/or unencapsulated mechanisms (Barsalou, Prinz). But this alignment is in no way mandatory. Fodor (1983, 2000) is the most obvious example here, as he defends the existence of representations with highly abstract content, yet vigorously opposes the massive modularity thesis. It is also possible in theory to defend massive modularity while denying the existence of full-on abstract representations. This could be done, for example, by holding that perceptual systems are modular and that higher cognition mostly involves the re-purposing of perceptual systems to solve new tasks, as radical empiricists of a certain ilk may want to do.

The contemporary debate about the structure of the mind has produced a number of wide-ranging and often largely theoretical arguments, most of which are explicitly framed around the goal of either supporting or undermining some version of the massive modularity thesis. These arguments appeal to a host of considerations to make their case: cognitive flexibility in scientific thinking (Fodor 1983, 2000); how to solve the frame problem in cognitive science (Shanahan & Baars, 2005); selection pressures for specialized psychological mechanisms (Carruthers, 2006b; Tooby & Cosmides, 1992); neural plasticity and connectivity (Anderson, 2010; Buller, 2005). However, many people have voiced concerns in the last few years that these kinds of wide-ranging, largely theoretical arguments are highly inconclusive one way or the other and that they lead to a sort of a stalemate (e.g., Colombo, 2013; Fuller & Samuels, 2014; Samuels, 2011).

There is another way to proceed in the controversy over massive modularity which isn’t liable to these concerns, however. It stems from the fact that, over the years, proponents of massive modularity have endorsed and touted a number of hypotheses about the existence of modules underpinning specific higher cognitive capacities. Three hypotheses have attracted the most attention in that respect: the hypothesis of a module for detecting cheaters in social exchanges (Cosmides, 1989; Cosmides, Barrett, & Tooby, 2010), the hypothesis of a module for face recognition (Kanwisher, McDermott, & Chun, 1997; Kanwisher & Yovel, 2006) and the hypothesis of modules for attributing mental states to other agents (Carruthers, 2013; Leslie, 1994). These hypotheses are based in good part on circumscribed findings originating from specific experimental methodologies, not on wide-ranging and largely theoretical arguments. So, a good way to make progress on the massive modularity controversy consists in carefully assessing the viability of hypotheses like these, namely hypotheses which (i) are widely endorsed by proponents of the massive modularity thesis; (ii) are linked to a specific cognitive capacity; (iii) arise from the interpretation of results related to a proprietary experimental methodology; (iv) are relatively specific about the nature of the representations and principles on which the purported modules operate. On the one hand, if it turns out that hypotheses that satisfy these four properties hold up to scrutiny, the empirical findings in question will provide inductive support for the massive modularity thesis. On the other hand, if hypotheses that satisfy these properties don’t hold up to scrutiny — i.e., if it turns out that one can adequately account for the empirical findings related to these hypotheses by appeal to domain-general and/or unencapsulated mechanisms —, then the findings may end up providing inductive evidence against massive modularity.

So, how have the three hypotheses just proposed fared since they were first put forward? I think that it is fair to say that the jury is still out on the three of them. At the moment, they all face specific objections dealing with their own proprietary experimental methodology for uncovering the purported module. However, I submit that one of the most convincing and detailed cases for a module in higher cognition — perhaps *the most* convincing case — comes from a fourth hypothesis, a hypothesis pertaining to work on spatial navigation. To see what that case is exactly though, we need to discuss a specific type of navigation experiment. We turn to it now.

2. The Reorientation Task

Many species can find their way back to important locations — their nest, a shelter, a food source — when their exploratory and foraging activities are disrupted in various ways, as we saw in the case of displaced bees above. In an attempt to study how animals achieve such feats, the psychologist Ken Cheng (1986) developed an influential experimental paradigm known as *the reorientation task*. He discovered that rats rely on the three-dimensional surface layout of an environment to return to a desired goal following disorientation — and yet they often ignore other important spatial cues in the same context. Similar patterns have then been observed with a wide variety of experimental subjects, including human infants (Hermer & Spelke, 1996), human adults under cognitive-load conditions (Hermer-Vazquez, Spelke, & Katsnelson, 1999), non-human primates (Gouteux, Thinus-Blanc, & Vauclair, 2001), other mammals (Lee, Tucci, Sovrano, & Vallortigara, 2015), fish (Sovrano, Bisazza, & Vallortigara, 2003), birds (Vallortigara, Zanforlin, & Pasti, 1990) and even insects (Sovrano, Potrich, & Vallortigara, 2013; Wystrach & Beugnon, 2009).

In a typical reorientation task, an animal is first incited to search for a food reward in a rectangular enclosure with distinctively colored walls and/or containing panels with distinctive two-dimensional patterns placed in the corners (Figure 1). Upon finding the reward, the animal is removed from the enclosure to undergo a disorientation procedure, which consists in some form of rotation without visual input. Finally, it is put back in the original enclosure, where various aspects of its search behavior are recorded.

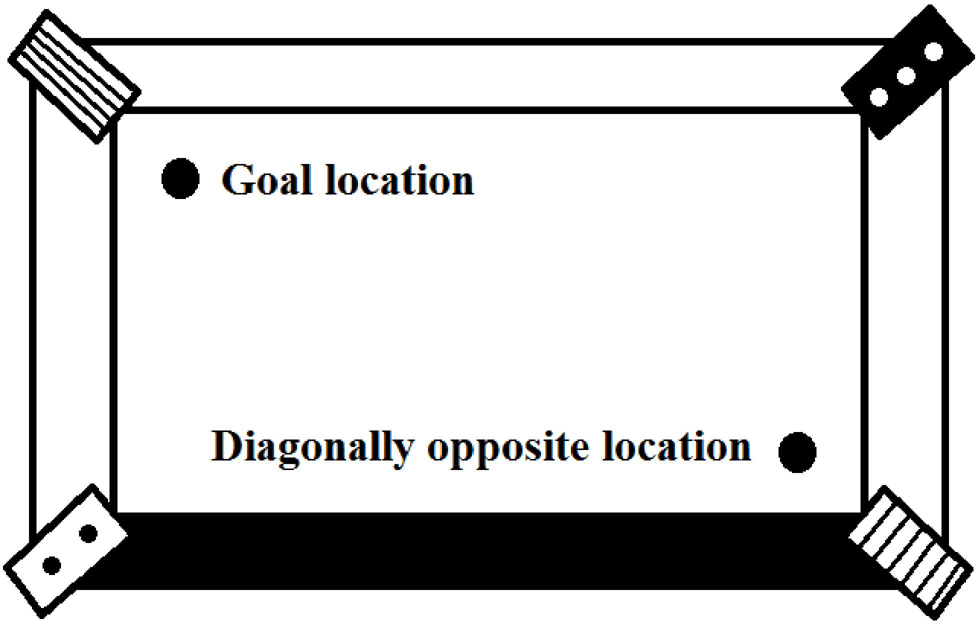

Figure 1. Enclosure of a typical reorientation task as seen from above. It has three white walls, one black wall and panels with distinctive two-dimensional patterns in the four corners.

In many experiments of this form, subjects search significantly more often at two locations than anywhere else in the enclosure without significantly favoring one over the other: (1) where they originally found the reward, and (2) at the equivalent location in the diagonally opposite corner. (For ease of exposition, I will refer to these two locations as the *diagonally adequate* locations, and to the other possible search locations as the *diagonally inadequate* locations.) That is so despite the fact that subjects could in theory use the salient colors or two-dimensional patterns on the walls to infer which of those two locations is the goal location.[3]

Two terms of art from the ensuing literature are helpful to describe such results: 'geometric cues' and 'featural cues'. *Geometric cues* pertain to the metric properties and relations of *three-dimensional extended surfaces* — surfaces like walls, floors, ceilings, cliffs, hills and valleys. For example, the length of a wall, the angle of a tilted floor and the concavity of a hillside would count as geometric cues as understood here. *Featural cues*, on the other hand, include isolated three-dimensional objects (e.g., chairs, columns, small rocks, isolated trees) as well as colors, two-dimensional patterns and textures on three-dimensional extended surfaces. With this new terminology, we can summarize the results of many reorientation experiments as follows: subjects usually rely on the geometric cues of an enclosure to search for the goal upon re-entry, and they often ignore featural cues.[4]

Note also that researchers who perform reorientation tasks carefully choose which type of distinctive featural cues they put in the corners or on the walls of the experimental enclosure. They choose featural cues which they know, based on prior experimental findings, that their subjects can easily detect in a variety of conditions other than reorientation tasks. So, we cannot simply brush off the results of reorientation tasks by saying that the subjects' visual apparatus does not have the precision or capacity to detect the relevant cues. Something more is going on. It presumably has something to do with how information about geometric cues and featural cues is processed and used

---

[3] See Cheng and Newcombe (2005), Vallortigara (2009, 2017) and Cheng, Huttenlocher and Newcombe (2013) for detailed reviews of the experimental literature. Note also that, though human adults follow these patterns in cognitive-load conditions, they go back to the corner where they originally found the reward on nearly every trial in regular conditions (i.e., non-cognitive-load conditions).

[4] The definition of 'geometric cues' given here and used in the reorientation literature is relatively idiosyncratic. On this definition, some metric properties do not count as *geometric cues* — namely, metric properties of objects that are not also three-dimensional extended surfaces. For instance, the height of an isolated object, like a chair, is not a geometric cue, but the height of a large wall is a geometric cue. In this thesis, I always use the expression 'geometric cue' in the sense proposed in this paragraph.

to guide behavior following the disorientation procedure. This is where the notion of modularity comes in.

3. The Geometric-Module Framework and the Two Foundational Debates

In the paper that started the reorientation literature, Cheng (1986) hypothesized that rats possess a modular mechanism that guides their spatial behavior upon recovery of normal perceptual access to their surroundings. He called it *the geometric module*. Importantly, he interpreted the subjects' systematic rotational errors in search behavior as rather direct evidence of the encapsulation of this mechanism. More specifically, he took the errors as showing that the process in charge of guiding subjects' behavior in reorientation tasks completely ignores all information besides geometric information in its operations, that it is *encapsulated* from everything except geometric information. Why else would subjects be ignoring the very useful and salient featural cues close to the goal location to distinguish it from the diagonally opposite location?

Since the publication of Cheng's (1986) seminal paper, many empirical and theoretical papers have been published in favor or against the hypothesis of a geometric module for reorientation. Among them, many theoretical papers raise doubts about Cheng's inference from rotational errors to encapsulation. Papers that explicitly endorse the hypothesis of a geometric module, on the other hand, do not all offer equivalent interpretations of that idea however. So, we should not treat them as proposing a perfectly unified theory with a single canonical formulation. They rather point to a general explanatory framework that comprises multiple theories with similar commitments. I will call it *the geometric-module (GM) framework.* Here, I take C. R. Gallistel's (1990) highly influential theory as the basis for my presentation of the framework, along with Sang Ah Lee and Elizabeth Spelke's (2010a) closely related account.[5]

These authors put forward the following kind of account. When animals first explore an environment, they store a representation of the global geometry of three-dimensional extended surfaces in that environment — henceforth *a geometric representation*. Then,

---

[5] Other important GM models include Wang and Spelke (2002, 2003), Shusterman and Spelke (2005), Cheng (2005), Cheng and Gallistel (2005), Sovrano and Vallortigara (2006), Lee, Sovrano, and Spelke (2012), and Gallistel and Matzel (2013).

later, when they recover normal perceptual access to their surroundings in that same environment following a disruption of some sort, they retrieve that representation. The geometric module then receives and automatically uses the representation to make an estimate of the subject's current heading and precise location within the environment. It computes the subject's heading and location by aligning the retrieved geometric representation to a new geometric representation constructed from current perceptual input. In doing so, it completely ignores featural information despite the fact that the subject itself does notice and register many relevant featural cues. At the end of the alignment process, the geometric module feeds the estimated heading and location to other navigation and motor systems, which in turn use that information to plan a path to the goal location.

According to the GM framework, the geometric module is modular in virtue of the fact that it is domain-specific and encapsulated to a strong degree. First, the mechanism is *domain-specific* because it only deals with problems pertaining to behavior guidance in situations of loss and sudden recovery of perceptual access. Second, it is *encapsulated* in that it operates only on representations of the global geometry of three-dimensional extended surfaces. It systematically ignores featural information as well many other types of information stored by the subject for navigation-related purposes. For example, it has no truck with information about the time of day, the outside temperature, the nutritional value of the last type of food encountered, the subject's level of hunger or thirst, and the level and location of pain stimuli. Even if the subject wanted to conjoin such information with geometric information, the highly specialized nature of the mechanism's internal principles could not enact the integration (Gallistel, 1990).

GM theorists unite in holding that the reason why subjects often make rotational errors in rectangular enclosures is that the geometric module completely ignores all information besides geometric information. So, more than 30 years later, the commitment to encapsulation still plays a central explanatory role in the GM framework. The geometric module only relies on the content of geometric representations to perform its function, and there are two ways of aligning geometric representations of rectangular environments so that their content matches. One way produces the correct estimation of the animal's heading and location in the enclosure. The other yields an estimation of the animal's heading that is off by 180°, and wrongly places the animal in the diagonally opposite corner of where it actually stands.

Since the beginning of the reorientation literature, many people have challenged Cheng's original inference from systematic rotational errors to encapsulation by formulating alternative explanatory frameworks that eschew encapsulated mechanisms. We will consider one such framework at length in this thesis (Chapters 3 and 4). But I want to emphasize here that, despite endorsing Cheng's inference, all GM theorists also explicitly acknowledge that reorientation subjects can rely on featural cues in at least some experimental contexts. In fact, Cheng (1986) himself demonstrated that repeated exposures to an unchanging enclosure with a stable goal location often induce reorientation subjects to start using featural cues close to the goal location to return to the goal location and thus leads to a slow decrease in the number of rotational errors over time.

In order to account for these sorts of results, most GM theorists posit a *beacon-homing process* as distinct from the geometric module. Beacon homing consists simply in moving toward a featural cue in a straight line until the agent has reached or retrieved a sought-after goal previously experienced as being inside or near the cue. It is a simple navigation strategy which requires the use of two representations, a representation of the goal and a representation of the cue, between which a link has been created because of that prior experience. GM theorists conceive of the beacon-homing process as independent from the geometric module and in competition with it to take control in guiding subjects' search behavior in reorientation tasks. It takes control when geometric cues are particularly unhelpful, or when the link between the representation of the goal and the representation of the cue is strong enough. The link itself can be strengthened by repeatedly obtaining the goal when in proximity to the cue.[6]

Various empirical considerations support GM theorists' claim that distinct processes deal with geometric cues on the one hand and featural cues on the other hand. Perhaps the most influential such consideration comes from the results of reorientation experiments that put distinctive featural cues in the four corners of a rectangular

---

[6] This picture does not fit well, however, with the performance of human adults in non-cognitive-load conditions (see fn. 3) because in these conditions they start relying on featural cues on their very first trials. The issue of explaining why human adults' performance changes across cognitive-load and non-cognitive-load conditions is beyond the scope of this thesis. See Section 5 for some details about the relevant literature. But note that the GM framework is committed to the view that the results from the cognitive-load conditions are more central or explanatory basic than the ones obtained in non-cognitive-load conditions.

experimental enclosure during a training phase (as in Figure 1) and then remove the distinctive featural cues from the two diagonally adequate corners during a testing phase. Even though many subjects manage to start going back to the correct corner significantly more often than the diagonally opposite corner in the training phase of such experiments, they do not carry on doing so in the testing phase (Cheng, 1986; Vallortigara et al., 1990). They simply fall back to chance between the two diagonally adequate corners in the testing phase *despite the fact that making use of the distinctive featural cues in the two other corners could easily help subjects figure out which one is the correct corner*. This suggests that subjects cannot integrate featural information with geometric information in reorientation tasks. We thus need at least two distinct processes to explain reliance on each type of cue in different contexts.

I now want to explain how the GM framework bears on the two foundational debates introduced in Section 1.

Consider the debate about the nature of mental representation first. It is one of thing to claim that, through vision, humans experience the world as being three-dimensional, as being composed of three-dimensional objects and surfaces. Everyone in philosophy of mind and psychology accepts this. But it is quite another to claim that, through vision, human minds construct *mental representations* that encode metric properties of three-dimensional extended surfaces as such. This requires accepting the existence of mental representations with highly abstract content. Though this view has been defended by important researchers (Marr, 1982; Pylyhsyn, 2003), it has also been frequently rejected (Gibson, 1979; Warren, 2012). Yet, the GM framework bolsters an even more radical position than the one defended by Marr (1982) in three ways: (1) it supports the view that *non-human animal minds* also possess such representations, since reorientation tasks performed in non-human species like rats and bees have shown that they also make systematic rotational errors in similar contexts; (2) it supports the view that the representations in question capture the metric properties of three-dimensional extended surfaces that have been perceived at one time or another in the current environment, *even those currently outside the visual field*; (3) it supports the view that these representations are held in a memory store for at least a few minutes once the agent leaves the environment, if not much longer.

Now, turn to the debate about the structure of the mind. Though many philosophers would seem willing to grant that perceptual processes are realized in modular mechanisms if push comes to shove, there is far more skepticism and hostility around the idea of positing multiple modules in higher cognition (Fodor, 2000; Sterelny, 2003; Buller, 2005). Yet, the geometric module, if it exists as described by GM theorists, belongs squarely within the realm of higher cognition. It is involved in action guidance (it helps the subjects steer themselves toward a known goal), and it does not take part in any perceptual process. It merely receives already constructed geometric representations from upstream perceptual systems. So, were the GM framework proven right, it would provide an extremely clear case of a module in higher cognition. This might in turn bolster the standing of other specific hypotheses of modules in higher cognition by settling once and for all the question as to whether there could be *any* modules in higher cognition. It would thus strengthen the case for hypotheses of higher cognitive module whose empirical status is not entirely clear, such as the three hypotheses mentioned in Section 1, even in the absence of further evidence specifically related to them — thereby also increasing the overall degree of support for the massive modularity thesis. Moreover, given the level of details with which the GM framework has been assessed, both empirically and theoretically, it might also provide cognitive scientists and philosophers with ideas about how to find and argue for other modules in higher cognition.

In sum, if the GM framework is right, it provides important support for radical positions on both debates discussed in Section 1. But before we can reap such important benefits from the framework, we need to determine whether it is indeed true. This is a large part of what this thesis attempts to do. In the following chapters, I argue that the GM framework provides the best account of reorientation experiments. Many researchers have put forward important challenges to the GM framework in the last fifteen years, either in the form of direct objections to the framework or of alternative explanatory frameworks for reorientation experiments. I argue in this thesis that the most important objection against the GM framework fails to undermine it (Chapter 2) and that we should reject the most influential alternative explanatory framework for reorientation experiments (Chapter 3—5).

4. The Geometric-Module Framework Bears on Many Further Philosophical Issues

We have just seen how the GM framework bears on the two foundational debates in philosophy of psychology identified in Section 1. Making progress on these two debates constitute the main reason why I focus on the GM framework in this thesis, and the two debates thus provide the background against which my defense of the GM framework will be ultimately assessed. However, I want to note here that there are many additional philosophical issues on which the GM framework can bear — and thus additional reasons to be interested in the viability of the framework — even though this thesis won't focus on them. Here are just a few:

4.1. The Nativism/Empiricism Debate

In its contemporary form, the nativism/empiricism debate turns on the nature of the psychological mechanisms that serve as the foundation for the acquisition of further psychological mechanisms, processes or concepts (Margolis & Laurence, 2013; Laurence & Margolis, 2015). Nativism is the view that human minds contain a large variety of domain-specific mechanisms that serve as the foundation for the acquisition of further psychological traits. Empiricism is the view that human minds only comprise a few mechanisms that serve as the basis for the acquisition of further psychological traits, and that many or all of these mechanisms are domain-general. Domain-general mechanisms include, for example, mechanisms that obey principles of associative learning or Bayesian statistical reasoning. Empiricists hold that such mechanisms are involved and explain performance in a wide range of cognitive domains, such as spatial navigation and language acquisition. They are willing to acknowledge the existence of domain-specific mechanisms, like the geometric module, but only on the assumption that these domain-specific mechanisms have been constructed by domain-general learning mechanisms through specific experiences of some sort.

The GM framework bears on the nativism/empiricism debate as follows. Eighteen-month-old human infants display similar behavior to non-human species in reorientation tasks: they go to the diagonally adequate corners significantly more often than the two other corners from the very first trials, yet they don't choose the correct corner significantly above chance as compared to the diagonally opposite one (Hermer &

Spelke, 1996). This means that, if the GM framework is right, the geometric module and the downstream path-planning systems with which it cooperates are operational at the latest around the 18-month mark in humans. This, along with the complexity of the purported function performed by the geometric module (the estimation of the subject's heading and location within the current environment by aligning two geometric representations), strongly suggests that the geometric module and attendant path-planning mechanisms are not acquired as a result of, or constructed from, the operations of a few domain-general learning mechanisms. Domain-general learning mechanisms require extensive feedback based on specific experiences in order to build an even moderately efficient mechanism for a complex task like estimating one's heading and location. The more complex the task, the more time it takes to build the relevant mechanism. Hence, if true, the GM framework would strengthen the case for nativism, at least as applied to spatial navigation abilities.

Note also that domestic chicks choose diagonally adequate corners significantly above chance in reorientation tasks *merely three days after birth* in cases where the only environment they have experienced in their life was a circular one (Chiandetti & Vallortigara, 2008) and in cases where they have never had visual access to three-dimensional extended surfaces in their lives up to that point (Chiandetti, Spelke, & Vallortigara, 2015). Assuming the truth of the GM framework, such results further suggest that the geometric module and downstream path-planning mechanisms may have been part of the psychological acquisition base for a long part of humans' evolutionary lineage.

4.2. Cognitive Foundations of Formal Geometry

Where do geometric concepts like TRIANGLE, SQUARE, LINE, POINT, ANGLE, PARALLEL and CONGRUENT come from? Consider also the various theorems of Euclidean geometry that we encounter early on in school — e.g., that the sum of the interior angles of a triangle is 180 degrees, that squares have four axes of symmetry. Where do intuitions about the validity of the inferences supporting these theorems come from? Everyone involved in the debate that stems from these questions seems to agree that geometric concepts and intuitions come from cognitive mechanisms involved in spatial navigation. Hence, the disagreement turns largely on the nature and origins of those mechanisms. Here again, we have an opposition between nativist-minded philosophers and empiricist-minded philosophers.

The nativist-minded view is that the navigation-involved mechanisms that give rise to these concepts and intuitions are domain-specific (i.e., they deal only with specific issues pertaining to spatial navigation) and they are not acquired from the operations of more domain-general learning mechanisms (cf. Descartes, 1637/2001; Kant, 1787/1998). In other words, such mechanisms are part of the human psychological acquisition base, the set of psychological mechanisms which serve as a foundation for the acquisition of further psychological traits. Such mechanisms might harbor specific geometric concepts or principles as basic components which they routinely rely on to perform their operations. Or, alternatively, they may also contain specialized templates to produce specific geometric concepts and intuitions based on minimal observations of the physical world through navigation.

The empiricist-minded view is that the navigation-involved mechanisms that give rise to geometric concepts and intuitions are domain-general themselves, or that they are acquired largely from the operations of domain-general learning mechanisms. This view holds that experiences of moving through physical space and planning routes plays a large and fundamental role in the acquisition of geometric concepts and intuitions through the operations of domain-general mechanisms that extract specific regularities from these experiences (Berkeley, 1709).

How does the fate of the GM framework relate to this debate about the cognitive foundations of formal geometry? Assuming the geometric module is part of the human acquisition base (as proposed in relation to the previous debate), the GM framework offers important support for the nativist-minded view in this debate. Given the complexity of its purported function, the geometric module arguably relies on a variety of high-level geometric concepts and principles in its operations, perhaps even Euclidean axioms in some form or other. Moreover, it is not altogether implausible to imagine that the geometric module would have found a way, through human evolution, to make these concepts and principles at least partly or indirectly available to a few further systems, such as the language faculty. It would have allowed people with such a cognitive architecture to get a significant head-start when tackling everyday problems that involve verbally describing the sizes and shapes of territories or foraging routes. Of course, exactly how many of our personal-level geometric concepts and intuitions arise from the internal structure of the geometric module would need to be determined through further empirical investigations — and some prominent GM theorists, like Spelke, Lee, & Izard (2010), hold quite a moderate stance on this issue. But, in any case, the GM framework doesn't offer much ground for enthusiasm on the empiricist-minded side of the debate, which predicts an important role for domain-general mechanisms.

4.3. How Practical Reasoning Works

How do agents' beliefs and desires lead them to perform certain actions? For example, how do someone's various beliefs and desires lead them to vote for a certain candidate in a given political election? The general question is a fundamental issue for both philosophers and psychologists, and one for which much less is known than one might believe. Philosophers often tout the so-called *practical inference* — the inference that leads an agent from desiring *that q* and believing *that* $p \rightarrow q$ (where p is something the agent can bring about) to desiring *that p* — as a key explanatory idea when trying to account for practical reasoning. But note that appealing to the practical inference has virtually no explanatory power on its own. That is because it is not much use without a working account of how, for example, specific desires get to guide action at specific times (given that we cannot act on all our desires at once). Or a working account of how the right $p \rightarrow q$ conditional belief is chosen to trigger a practical inference in the first place. Not every such conditional belief will do. I may desire to be rich, and believe that being the president of a multinational corporation would make me rich — yet I may still not desire to be the president of a multinational corporation.

Needless to say, how someone deals with these issues depends heavily on whether or not that person is committed to wide-scale modularity or domain-specificity in higher cognition. A somewhat middle-ground position is that we have a specific practical-reasoning system that queries distinct sets of desire- and belief-generating modules in order to retrieve the strongest desire at a given time (i.e., the one that will guide action at that time) and the conditional belief most relevant to guide the agent to satisfy that desire (Carruthers, 2004). This practical-reasoning system itself performs practical inferences as part of its internalized operations, and it may well count as both domain-general and unencapsulated in an important sense. A more radical possibility is that practical inferences are largely performed within specific modules themselves without the supervision of an overarching domain-general practical-reasoning system (Gallistel, 1980; Carruthers, 2006a). The GM framework provides support for this second, more radical option. On standard GM models, the geometric module starts guiding action in cooperation with downstream path-planning systems as the result of the activation of a certain motivational state — a state that presents the agent as being (e.g.) hungry, thirsty, or seeking an escape from the current environment. In those cases, the downstream path-planning systems calculate a proposed path that would allow the subject to quickly satisfy those goals based on the estimation of the subject's heading and location produced by the geometric module. The geometric module and attendant path-planning systems thus perform a form of practical inference without need of an overarching practical-reasoning system.

4.4. Other Issues

Now, these are only three of the most obvious philosophical issues on which the GM framework bears. But there are many more. I will succinctly mention two additional issues here: (i) the role of language in cognition; (ii) the origin of flexible thinking in humans. Due to evidence that human adults rarely ever make rotational errors in typical reorientation tasks (Hermer & Spelke, 1996), and several results suggesting an important causal role for language systems in their performance (Hermer-Vazquez, Moffet, & Munkholm, 2001; Hermer-Vazquez et al., 1999), the GM framework has been used as a basis to make a widely discussed suggestion pertaining to these two debates. The suggestion is that the language faculty operates as a means to combine information pertaining to modules dedicated to distinct domains through the formulation of specific sentences in inner speech (Carruthers, 2002), thus providing humans with a way to produce cross-domain thoughts inaccessible to non-human animals. The suggestion arose from an analysis of the operations of the geometric module which suggested that the syntax of natural languages could facilitate the combination of some representations produced by geometric module and a distinct module dealing with featural cues. Though this suggestion and its bearing on the two debates just mentioned rightly brought a lot of attention to the reorientation literature in its early years, I will not discuss the suggestion at all in the next chapters. That is because more recent results (Bek, Blades, Siegal, & Varley, 2010; Pyers, Shusterman, Senghas, Spelke, & Emmorey, 2010; Ratliff & Newcombe, 2008a; Shusterman, Lee, & Spelke, 2011) have offered a far more complex picture of the role of language in reorientation tasks than Hermer-Vazquez et al.’s original results suggested, and it is not yet clear how much of the original suggestion can be preserved in light of this new evidence. In any case, it remains clear that investigating how the geometric module cooperates with language systems provides a great way to address the role of language in cognition and how modular mechanisms may participate in flexible thinking.

In sum, the GM framework has clear and important implications in relation to a number of highly important debates in philosophy of psychology and cognitive science, besides the two foundational debates identified in Section 1.

5. Structure of the Thesis

Let us recap what we saw in the previous sections. Research on spatial navigation is advantageously positioned to make progress on the two foundational debates identified in Section 1. And within research on spatial navigation, a specific class of experiments, reorientation experiments, are particularly well-suited for the task, insofar as they give us a way to assess the hypothesis of a geometric module for spatial reorientation. So, this gives philosophers of psychology a strong incentive to pay attention to the reorientation literature and the fate of the GM framework. In this vein, the main goal of the thesis is to defend the GM framework from two main challenges. More specifically, I argue that the most important objection against it fails to undermine the GM framework, and that we should reject the most influential alternative explanatory framework that aims to account for reorientation experiments.

Here is how I will proceed to make my case. **Chapter 2** considers the objection in question, which I call *the explanatory inflexibility objection*. The objection focuses on a specific type of reorientation experiments: transformational tasks. Those are reorientation tasks which involve a significant transformation of the experimental enclosure following a training phase. We find significant variation in whether and how subjects rely on featural cues in these experiments, and the objection holds that the GM framework does not have the required flexibility to explain this variation. I propose to answer that objection by positing the existence of a mechanism whose job is *to select* the relevant geometric representation at the beginning of a trial in order to feed it to the geometric module. This mechanism’s normal operations are disrupted by the transformations performed in reorientation tasks.

The next three chapters of the thesis tackle the debate between the geometric-module framework and the most influential alternative explanatory framework that aims to account for the results of reorientation tasks: *the view-matching (VM) framework*. The latter framework aims to explain systematic rotational errors by appealing to stored representations of the subject's two-dimensional retinal stimulation at specific locations of the environment. These representations are called *snapshots*. According to the VM framework, a view-matching process compares previously recorded snapshots to the current retinal stimulation in order to guide the animal back to the goal location. VM theories reject the idea that reorientation subjects store information about geometric cues as such, or that they possess a module in higher cognition for spatial navigation.

The remaining chapters make three main contributions to the debate between the two frameworks, one in each chapter. **Chapter 3** puts forward a detailed argument that, contrary to what many GM and VM theorists alike suppose, existing evidence does not clearly favor either framework over the other. Rather, I argue that the debate has reached something of an impasse because each framework can provide systematic explanations for the type of empirical results seen as favoring the other framework. **Chapter 4** develops a new problem — *the representation selection problem* — that offers the prospect of breaking the impasse by introducing a new type of explanatory consideration that theories of both types must address. The representation selection problem requires explaining how subjects can reliably select the relevant representation with which they initiate the reorientation process. It arises because reorientation subjects have representations from multiple environments in memory while undergoing a trial, be they geometric representations or snapshots. I argue that VM theories do not have the resources to properly address this problem. **Chapter 5** then develops a new GM account which can provide a natural response to the representation selection problem.

Overall, the arguments in this thesis will suggest that the GM framework is the best framework to account for the results of reorientation experiments, and that it therefore offers important support for the existence of mental representations that encode high-level abstract properties of the world, as well as for the massive modularity thesis.

6. Three Caveats

Before we move on to Chapter 2, I have three methodological caveats.

First, note that I will take as a working assumption the truth of representationalism — the view that human and non-human minds contain internal representations of aspects of the world — throughout this thesis. More specifically, I will presuppose that subjects acquire full-on mental representations of salient environments they visit during reorientation experiments and that they rely on such representations to guide their behavior, whether they be geometric representations, snapshots or some third kind of representation. One of the reasons for making this assumption is practical: Even though reorientation studies have given rise to large empirical and theoretical literatures, we have yet to see a serious, detailed anti-representationalist account of reorientation results. So, it is not possible for me to discuss how anti-representaionalists may decide to handle the type of considerations that I present in this thesis. Some anti-representationalists may be tempted to argue that VM theories can or should be re-cast in anti-representationalist terms given the theories' minimal representational commitments. But even if such a move is possible in principle, it is not worth investigating here. That is because, if my arguments involving the representation selection problem in Chapters 4 and 5 are right, it will follow that the VM framework is in deep trouble precisely because it eschews representations that encode high-level abstract information about visited environments, most notably information about geometric cues as such. So, precisely what could make VM theories appealing to the anti-representationalists is also what renders them inadequate. For the same reason, the arguments of Chapters 4 and 5 will also vindicate my decision to take the truth of representationalism as a working assumption.

Second, note that I do not want to imply, through the use of expressions like 'theories of reorientation' and 'framework of reorientation', that there is a well-circumscribed cognitive process which is always triggered by reorientation experiments and only by reorientation experiments to which one might unambiguously refer with the expression 'spatial reorientation'. In fact, I will argue explicitly in Chapter 3 that, at least from the perspective of the GM and the VM frameworks, it is virtually guaranteed that multiple distinct navigation processes compete and manage to take control of reorientation subjects' behavior in different contexts. Thus, I use the expression 'theories of reorientation' in a deflationist way to pick out various theoretical models of the *behavioral results obtained in reorientation experiments* (i.e., experiments that closely follow Cheng's 1986 experimental design). Many such models will or should posit multiple, competing processes.

Third, I want to emphasize that, contrary to what is sometimes implied, the debate between the GM framework and the VM framework does not boil down to a simple question of whether we need to posit allocentric representations, as opposed to egocentric representations, in order to account for reorientation results. According to usual definitions, *allocentric representations* are representations whose references frames are anchored to objects external to the agent, like a specific rock or tree in my backyard. *Egocentric representation*s, on the other hand, are representations whose reference frames are anchored to the subject whose representations those are. Many people take the geometric representations posited by GM models to be allocentric (presumably on the ground that the reference frame of a geometric representation must be anchored to three-dimensional extended surfaces in the environment) and the very different type of representation at the heart of VM models, snapshots, as egocentric (presumably on the ground that the reference frame of a snapshot must be anchored to the subject's retina). And many researchers further seem to infer from this that all there is to the debate between GM and VM theorists is the issue of whether the representations that underlie subjects' performance in reorientation tasks are allocentric or egocentric.

So why doesn't the debate between GM and VM theorists amount to this simple issue? To begin with, some theories of reorientation universally categorized as belonging to the GM framework and that accept its core explanatory strategy (e.g., Wang & Spelke 2002) maintain that the geometric module's geometric representations are egocentric and that reorientation subjects never employ allocentric representations (see also Lee et al., 2012). According to these theories, reorientation subjects' geometric representations encode important properties of three-dimensional extended surfaces as such, but the representations' reference frame is not anchored to the surfaces themselves. Though I happen not to accept Wang and Spelke's argument for that view, I also do not want or need to officially commit myself here to a specific position in the complex debate to which their argument has given rise (see, e.g., Burgess, 2006). More importantly, I believe that the distinction between allocentric and egocentric representations is not all that clear, appearances to the contrary notwithstanding. I do not have space here to give a detailed argument to that effect, but suffice it to say that I think that it is far clearer to talk directly about the kind of representations posited by a given theory, what the content of these representations is, and how the mechanisms posited by the theory makes use of a given representation through its content. As clear as it may seem, the issue of whether a representation's reference frame *is anchored to the subject or rather to an external object* is often indeterminate from the perspective of various theories, and not just due to underspecification. For that reason, I will also avoid using expressions like 'allocentric representations' and 'egocentric representations' in what follows.

Chapter 2

Answering the Explanatory Inflexibility Objection Against the Geometric-Module Framework

1. Introduction

In this chapter, I will discuss the most important objection against the geometric-module (GM) framework. It focuses on a class of experiments that occupy a central place in the reorientation literature: transformational tasks. Transformational tasks are reorientation tasks which involve a significant transformation of the experimental environment following a training phase. One of Cheng's (1986) experiment is of that type. Rats were first trained in a rectangular enclosure with distinct and stable panels in each of the four corners. The rats eventually started relying on the panels to choose the correct corner in a reliable way. Then, in one of the test phases, Cheng moved all the panels by one corner in a clockwise manner. At this point, there were no corners that were both geometrically appropriate and featurally appropriate from the perspective of the training phase. The rats could choose a corner by relying either on the geometric cues or on the featural cues they had seen in the training phase, but not both simultaneously. After this transformation, the rats' performance became unsystematic: they chose the geometrically appropriate corners on 49% of the test trials, and they chose the featurally appropriate corner on 47% of the test trials (Cheng 1986, Figure 4C).

The interest in transformational tasks has grown in good part from a wish to stress the limitations of the GM framework (see, e.g., Cheng, 2008; Graham, Good, McGregor, & Pearce, 2006; Gray, Bloomfield, Ferrey, Spetch, & Sturdy, 2005; Pearce, Graham, Good, Jones, & McGregor, 2006). I want to discuss a specific objection to the GM framework that has been raised based on this kind of task, which I call *the explanatory inflexibility objection*. Here it is in schematic form:

> Subjects sometimes rely on geometric information (e.g., Cheng, 1986, Experiment 2) and they sometimes rely on featural information (e.g., Vallortigara et al., 1990, Experiment 4) when they choose a corner in the test phase of transformational tasks — and sometimes they display unsystematic

> corner choices (e.g., Cheng, 1986, Experiment 3). Any viable theory of reorientation must be able to explain this variability in the results of transformational tasks. Unfortunately, the GM framework does not have the "required flexibility" (Twyman & Newcombe, 2010, p. 1330) to do so (see also Cheng et al., 2013). Therefore, it is not a viable framework of reorientation. Or, at the very least, its explanatory ineptitude counts as a serious strike against it.

I will argue in the chapter that this objection does not threaten the GM framework. It is true that there is a lot of variation in the results of transformational tasks. But it is false to say that the GM framework does not have the required flexibility to explain this variation. I will make a case for this view as follows. In Section 2, I present the explanatory inflexibility objection in more detail by discussing the results of a typical transformational task. In Section 3, I provide empirical evidence to the effect that reorientation subjects possess a long-term database that simultaneously stores geometric representations from multiple environments. This allows me to sketch out a general strategy for explaining the results of problematic transformational experiments from the perspective of GM theorists: When faced with results of transformational tasks that seem problematic from the perspective of the GM framework, appeal to the hypothesis that the transformation of the experimental enclosure prevents the selection mechanism — the mechanism that feeds geometric representations from the geometric-representation database to the geometric module — from working properly. In Section 4, I identify seven issues on which future GM models will need to take a stance in order to give more detailed explanations and predictions of the results of transformational tasks. In Section 5, I then show that, even without taking a stance on these seven issues now, it is clear that the general explanatory strategy just outlined provides significant leeway for GM theorists to explain the results of transformational tasks that seem to contradict the predictions of the GM framework.

So, my aim in this chapter is mostly negative: I want to reject the explanatory inflexibility objection against the GM framework. However, analyzing this objection will also allow us to see various novel ways of developing GM theories that have not been discussed in the literature on spatial reorientation. Thus, this chapter will also present positive proposals about the nature and the structure of the geometric module and related mechanisms.

2. The Explanatory Inflexibility Objection

In this section, I will present the explanatory inflexibility objection in more detail. I will focus on Twyman and Newcombe's (2010) formulation of the objection which is the clearest version in the literature. A good place to start is by clarifying the notion of cue competition, which plays a key role in their formulation. *Cue competition* is a phenomenon by which the control of a cue over behavior interferes in some way with the control of another cue over behavior in a given task. Twyman and Newcombe make reference to three types of cue competition in their argument: overshadowing, blocking and potentiation.

Overshadowing occurs when responding to a certain stimulus A is lessened by the fact that it has been simultaneously paired with the unconditioned stimulus along with another stimulus B. Had stimulus A been paired individually with the unconditioned stimulus, the subject's conditioned response to A would have been stronger. *Blocking* occurs when a previous history of pairing between a certain stimulus B and the unconditioned stimulus lessens or prevents a subject from developing a conditioned response to a new stimulus A. *Potentiation* is the opposite of overshadowing. It occurs when responding to a certain stimulus A is heightened by the fact that it has been simultaneously paired with the unconditioned stimulus along with another stimulus B. Had stimulus A been paired individually with the unconditioned stimulus, the subject's conditioned response to A would have been weaker. The question as to when and why certain training contexts give rise to potentiation rather than overshadowing (and vice versa) is a vexed problem for theories of associative learning. (See, for instance, Durlach and Rescorla, 1980.) What seems clear, however, is that cases of potentiation happen much less often than those of overshadowing.

In order to get a better grasp of these notions and of transformational tasks, let us take a look at one of Pearce et al.'s experiments (Pearce et al., 2006, Experiment 2), perhaps the most influential such experiment. Pearce et al. claim that this experiment demonstrates that information about the color the walls of an enclosure sometimes *overshadows* learning about the shape of the enclosure in the reorientation process. This paper, along with some other papers by Pearce and colleagues, sparked a flurry of reorientation experiments assessing cue competition (see Pearce, 2009 for review).

Pearce et al.'s (2006) experiments were conducted with rats, and they were inspired by an influential experimental paradigm using water mazes to test navigation abilities (Morris, 1981). On any given trial, the subject is put in an enclosure positioned in a large circular pool. Experiment 2 had one experimental group and two control groups: the control-white group and the control-black-and-white group. The single difference between the three groups pertained to the color of the walls in the enclosure. During the training sessions of the experimental group, the two long walls of the enclosure were black and the two short walls were white; for the control-white group, all the walls were white (see Figure 2); for the control-black-and-white group, the color of the opposing walls often changed from black to white (for half of the trials, the long walls were black and the short walls were white, and for the other half, the long walls were white and the short walls were black).

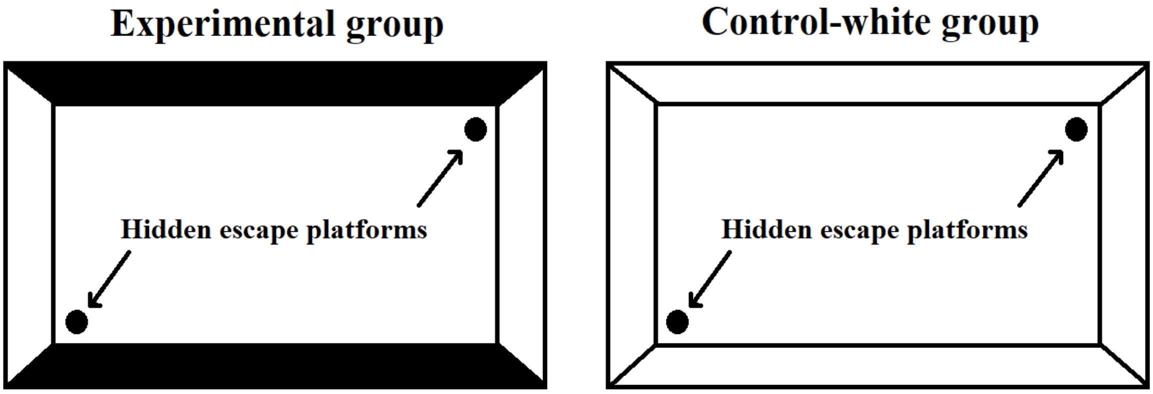

Figure 2. The enclosures used during the training sessions of Pearce et al.'s (2006) Experiment 2 as seen from above.

On training trials, the experimenters placed two hidden platforms in the water not too far from the surface. For half the subjects, these platforms were located in the two corners with a long wall on the left and a short wall on the right for the duration of the training phase (as in Figure 2). For the other half, the platforms were located in the corners with opposite geometric properties for the duration of the training phase. A given trial ended when the subject had remained for more than 30 seconds on one of the two platforms, at which time the experimenters would retrieve it from the enclosure. The water was slightly colder than what rats usually like, so the cold water motivated them to escape from the pool by finding the platforms as soon as they were put in the enclosure. Each daily training session consisted of four trials, and the pool was rotated by an angle of 90, 180 or 270 degrees every time the group completed one trial (the angle was chosen randomly between these three values).

Pearce et al. used two distinct measures of success to determine whether the rats were learning to use the shape of the enclosure to locate the hidden platform during the training sessions: the mean escape latency (the average time required for subjects to first put their paws on the platform) and the mean percentage of trials on which subjects entered the diagonally adequate corners before the diagonally inadequate corners. Pearce et al. found that all three groups improved quite rapidly according to both measures. The mean escape latency started around 65 seconds for the three groups, and it steadily went down until it reached asymptote at around 20 seconds. The mean percentage of diagonally adequate first corner choice started around 50% for all groups and increased until it reached asymptote at around 95% at the end of the training phase.

There was only one test trial. In that trial, the four walls of the rectangular enclosure were white for all three groups. The experimenters also removed the hidden platforms from the enclosure. They then filmed the behavior of each rat in the pool over a period of 60 seconds. Since there was no hidden platform for subjects to find, the experimenters had to discard the first measure of success (mean escape latency). However, instead of relying on the second measure to determine the success of the rats in the test trial, Pearce et al. opted for a distinct measure of success borrowed from Morris (1981): the mean percentage of time *spent over the 60-second test trial* in the diagonally adequate quadrants. By this third measure, the three groups scored differently during the test trial: the experimental group’s performance was not significantly different from chance, whereas subjects in the two control groups (especially the control-white group) spent significantly more time in the diagonally adequate quadrants. Using these results, Pearce et al. concluded that learning the shape of the rectangular enclosure had been overshadowed by the color of the walls in the experimental group. Pearce et al. (2006) also claim to have found instances of each of the two other types of cue competition — blocking and potentiation — in some of their additional experiments.

I have already mentioned that the explanatory inflexibility objection draws on the complex variation in the behavioral patterns found in transformational tasks as its evidential basis. The magnitude of this variation is already apparent in Pearce et al.’s (2006) experiments. For instance, the results of Pearce et al.’s Experiment 1 suggests that if you perform the exact same transformational task as in Experiment 2 with a *kite-shaped* enclosure as opposed to a *rectangular* enclosure, then the color of the walls will *potentiate* the shape of the enclosure (rather than *overshadow* it). More specifically, they found that rats that had been trained in a kite-shaped enclosure with two white walls and two black walls spent more time in the geometrically appropriate corners on the test trial in an all-white kite-shaped enclosure than rats that had been trained in the all-white kite-shaped pool from the very start of the experiment. Or consider Cheng's (1986) Experiment 2. It is a transformational task which differs from Pearce et al.'s Experiment 2 only in that it occurred in a food-search paradigm (rather than a pool-escape paradigm) and that it had many test trials (as opposed to only one). For the last test phase of the experiment, Cheng did not find any evidence of overshadowing, potentiation or blocking of geometric cues. In fact, Cheng's four rats relied on geometric cues for every single trial in a series of 50 consecutive trials following the removal of featural cues. These examples show that there is a lot of variability pertaining to cue competition during reorientation tasks. In some contexts, we see one form of cue competition whereas, in apparently similar contexts, we see another form of cue competition — or no cue competition at all.

We now have all the elements required to present the explanatory inflexibility objection as raised by Twyman and Newcombe (2010). After presenting the results of various experiments purporting to demonstrate the phenomenon of cue competition in reorientation tasks, Twyman and Newcombe provide just about the clearest formulation of the objection in the literature: "This variation [i.e., the absence and presence of overshadowing, blocking and potentiation in various experimental contexts involving transformations of the experimental enclosure] is one of the facts in this literature that a successful theory should be able to explain, and modularity theory [i.e., the GM framework] clearly does not have the required flexibility" (p. 1330). And they seem to infer from this purported inflexibility that GM theories are not viable theories of reorientation. Or, at the very least, that GM theories' explanatory ineptitude regarding the variation in the results of transformational tasks counts as a serious strike against them.

It is not entirely clear to me why proponents of the explanatory inflexibility objection think that the GM framework does not have the required flexibility to explain the variation in the results of transformational tasks. Some of them seem to believe that GM theorists are committed to the view that the geometric module completely bridges the gap between perception and action in reorientation tasks in such a way that any spatial movement that subjects make on a given trial reveals the module's internal operations. But there are two natural ways in which GM theorists can reject this view. First, they can hold that reorientation subjects possess parallel navigation processes or mechanisms that compete with the geometric module to guide subjects' search behavior in reorientation tasks, like a beacon-homing process (see Chapter 1:3). Second, GM theorists can maintain that reorientation subjects possess higher cognitive mechanisms (i.e., mechanisms that are not merely perceptual) upstream from the geometric module whose operations and output have important downstream effects on the geometric module. Some GM theorists have endorsed the first move as a response to the explanatory inflexibility objection (e.g., Lee & Spelke, 2010a), but nobody has ever pursued or even recognized the possibility of the second move.

I fully agree that we should posit parallel mechanisms that compete with the geometric module to guide spatial behavior in reorientation experiments, as I argue at length in Chapter 3. However, for various reasons that I cannot go into here, I suspect that the first move does not constitute the best way of *addressing* the explanatory inflexibility objection because it does not deliver the required explanatory latitude to deal with the whole range of results in transformational tasks. The second move appears far more promising to me in that respect. What I do in the rest of this chapter is to make a case that this move is empirically well-motivated. I will seek to enrich the GM framework by, among other things, considering empirical data about long-term success on reorientation tasks. I will take the following questions as my starting point: If there is a geometric module, where does it get its geometric representations from? Does it interact in important ways with other upstream components or mechanisms involved in navigation?

3. Introducing the Selection Interference Strategy

In this section, I will provide the schema of an explanation for the results of various transformational tasks. This schema is compatible with all the main tenets of the GM framework, and it significantly extends it. To present that schema, I will pursue a two-step strategy, which I will call *the selection interference strategy.* First, I will provide empirical evidence to support the view that the geometric representations on which the geometric module operates are stored in a long-term database that simultaneously stores geometric representations from many distinct environments. Second, I will argue that

GM theorists can explain the results of many transformational tasks by appeal to the hypothesis that the transformation of the experimental enclosure during the test phase prevents the selection mechanism — the mechanism that sends the required geometric representation to the geometric module at the outset of a trial — from working properly. In many cases, the selection mechanism does not treat the transformed enclosure in the test trials as the same environment as the one observed during training trials.

3.1. First Step: Defending Two Claims about the Geometric-Representation Database

The first thing to note here is that the GM framework is committed to the existence of a database from which the geometric module receives geometric representations of previously encountered environments. As Gallistel (1990, p. 208) puts it, the geometric module's function consists in "comput[ing] congruence between perceived shape and remembered shape [of the environment]" in order to determine the subject's heading and location after it recovers normal perceptual access to its surroundings. The representation of the perceived shape of the environment comes more or less directly from perceptual systems (the visual system, in most cases). But where does the representation of the *remembered* shape come from? It must be that the representation of the remembered shape of the environment is kept in some form of storage system before being sent to the geometric module at the beginning of a reorientation trial. Why? Because, as GM theorists conceive of it, the representation of the remembered shape was constructed on a prior visit to the environment, when the subject knew its heading and location. The representation of the remembered shape has to be withheld by some kind of storage mechanism while waiting for a representation of the currently perceived shape of the environment to be generated by perceptual systems. In other words, the geometric module needs an associated database for geometric representations.

With this in mind, we can now distinguish two claims contained in the initial description of the first step of the selection interference strategy:

> **Long-term storage**: The geometric-representation database can hold geometric representations of previously encountered environments for many hours, if not many days.

> **Storage of multiple representations**: The geometric-representation database can simultaneously store geometric representations from many distinct environments.

These two claims are empirically dissociable. The first claim makes an estimate of the maximum length of time during which the associated database can keep geometric information about one environment which the subject has previously visited. According to it, the geometric-representation database will retain, at least in some contexts, the geometric representation from a specific environment for a long period of time after the subject loses normal perceptual access to that environment. Here is a rather straightforward way in which a researcher could test the claim of long-term storage. The researcher starts by training some animals to head for a specific corner in a given enclosure devoid of featural cues, over a large series of trials occurring in a short period of time. She then keeps the subjects in a different room for a long period of time — say around 20 hours. Following this 20-hour break, she puts them back in the enclosure for a single reorientation trial. If the subjects reliably choose the diagonally adequate corners on this trial, it would strongly support the claim of long-term storage within the context of the GM framework. Unfortunately, no empirical evidence of this kind has been published in the reorientation literature yet. Reorientation researchers never report and analyze data specifically about subjects' performance on *the very first trial* after a large time gap. They nearly always average subjects' corner choices across multiple trials on a given day of training or testing.

The second claim concerns the number of geometric representations which can be held in the database at once. I take it as a non-optional commitment of the GM framework that the geometric-representation database holds at least *one* representation at any given time in typical reorientation experiments, namely a representation of the experimental enclosure as last observed. That being said, we can easily imagine that, in many contexts, the database keeps more than one geometric representation at a given time. For instance, it might be the case that the geometric-representation database holds three geometric representations for the whole length of any reorientation experiment: one representation of the experimental enclosure, one of the intertrial waiting cage, and one of their home cage (which is often located in a different room).

One of the best ways of empirically assessing the claim of multiple-environment storage is through reorientation tasks that involve distinct experimental enclosures which subjects visit in alternation. Very few papers provide data from such reorientation experiments. To my knowledge, there are only two such papers (Horne et al., 2010; Julian, Keinath, Muzzio, & Epstein, 2015). Consider Horne et al. (2010). In an experiment designed to test the effects of lesions to the perirhinal cortex on navigational and object-recognition abilities, Horne et al. ran two groups of rats — one group who had sham lesions and another group who had lesions to the perirhinal cortex — in a reorientation task employing a water tank with a hidden escape platform. The reorientation task included, in effect, two distinct enclosures visited in alternation: one white rectangular enclosure and one black enclosure of the same size. Importantly, the correct corner, the corner harboring the hidden escape platform, had opposite geometrical properties in the two enclosures: the correct corner always had a long wall on the left and a short wall on the right in one of the two enclosures, and it had the opposite properties in the other enclosure. On test trials where the escape platform was removed from the tank, both groups of rats spent significantly more time in the diagonally adequate corners than the diagonally inadequate corners *in both enclosures*. The fact that subjects in *the sham group* performed as they did, in particular, strongly suggest that rats' geometric-representation database can simultaneously hold two distinct geometric representations.

Given the importance of the claim of long-term storage and the claim of multiple-environment storage for the selection interference strategy, I will now give some indirect evidence which, I believe, strongly supports both of them. However, the empirical dissociability of those two claims means that I cannot defend them in the exact same way. So, I will provide a separate analysis for each claim.

When it comes to the claim of long-term storage, it is worth thinking about the results of reorientation tasks where the goal remains in the same location within a single non-changing experimental enclosure over multiple days of trials for a given subject. In many such tasks, the subjects' choice of geometrically appropriate corners improves over that period of stability despite large time gaps between the last trial of a given day and the first trial of the next day (on the order of 19 hours, often more).[7] From the

---

[7] For relevant experiments, see Pearce et al.'s (2006) Experiment 2, Vallortigara et al.'s (1990) Experiment 1, and Sovrano, Bisazza and Vallortigara's (2002) Experiment 1.

perspective of the GM framework, it is hard to explain this improvement over large time gaps without presupposing the truth of the claim of long-term storage.

Now, I have already noted above that the results of reorientation experiments where subjects visit two distinct experimental enclosures in alternation (e.g., Horne et al., 2010) offer direct support for the claim of multiple-environment storage. But I will appeal here to a specific class of navigation experiments called *latent-learning experiments* to strengthen my case for the claim. Latent-learning experiments are not reorientation experiments as reorientation experiments are usually understood (i.e., latent-learning experiments do not usually follow Cheng's 1986 experimental design), so it will give another dimension to the case for the claim of multiple-environment storage.

A famous latent-learning experiment was conducted by Tolman and Honzik (1930). It involved three groups of rats learning how to navigate in a maze. Rats in the first control group never received or saw any reward while in the maze, rats in the second control group always found a reward at the end of the maze, and rats in the experimental group got a reward when reaching the end of the maze *but only from the 11th day of training onward*. There are three facts about this experiment that are noteworthy from the perspective of my argument. First, the experimental group ran the maze faster than the second control group from the 12th day onward. Second, we have strong reasons to believe that, in latent-learning experiments of this type, subjects are either already oriented at the beginning of a given trial, or that they reorient themselves quickly and then start strolling in the maze. Third, Tolman and Honzik used a maze in which there are no distinctive two-dimensional patterns or colors on the walls, or even smells. In other words, there were no distinctive featural cues on which the rats could rely to get to the end of the maze.

Why are these facts relevant to assessing the claim of multiple-environment storage? It is natural for GM theorists to subscribe to the view that the geometric representations underlying the performance of rats in the third group, the experimental group, are the very same representations as those fed to the geometric module for reorientation (i.e., they are stored in the geometric-representation database). The explanation is simple: for more-or-less explicit evolutionary and computational reasons about the costliness of duplicating information, GM theorists presuppose that the geometric representations fed to the geometric module are also used in guiding oriented spatial behavior, potentially by systems other than the geometric module. In particular, GM theorists presuppose that the geometric representations stored in the database are used in cases where the animal is already oriented and the only reliable information to which it has access about the location of various target objects is the geometry of extended surfaces around the animal. And this fits perfectly with the description of latent-learning experiments, which involve mazes with a very complex geometry of three-dimensional extended surfaces, and no relevant featural cues. It is possible in theory that subjects in *the second group* weren't relying on any kind of geometric representations to perform as they did. Subjects in the second group could have relied on a learned sequence of response-based patterns (e.g. something that could be described by verbal instructions such as "turn left at the first corner, turn right at the second corner, turn right again at the third corner and then you will get to the goal") to get to the reward quickly. But, given Tolman and Honzik's experimental design, subjects in the *third group* could not have learned any such sequence to get to the reward quickly *because there was no such reward to reach on the first 10 days*. For subjects in the third group, there is no obvious way of quickly getting to the reward on the 12th day without relying on the maze's complex geometric cues in some important way. So, for GM theorists, the best explanation is that subjects in the third group possess geometric representations of the whole maze, and that those are stored in their geometric-representation database.

Now, for people who take geometric representations to underlie subjects' performance in latent-learning experiments (as GM theorists should), there are good reasons to hold that these subjects' geometric-representation database contains more than one geometric representation at once. Consider the third group in Tolman and Honzik's experiment again. Tolman and Honzik's favored explanation for that group's fast performance from day 12 onward — and the one to which GM theorists should subscribe given what we said in the previous paragraph — is that the rats in that group begin acquiring a relatively fine-grained geometric representation of the maze from the very start of the experiment. However, notice here that this explanation presumes that the rats in the third group begin acquiring a well-defined geometric representation from the very start of the experiment *despite the fact there are no direct practical advantages in doing so at that moment*. At the start of the experiment, the maze does not contain any source of food or water, no potential location for a home base, no potential prey, threat or predator, no potential source of danger if the rats stay too long in the environment.

Notice also that the maze is not the only environment to which the rats have perceptual access during the whole duration of the experiment. After every trial, they are brought back to their home cage (which contains food and water). Moreover, it is practically beneficial for rats to acquire a geometric representation of the home cage, if only to help them find the food and water faster when they are brought back to their cage. In these circumstances — there is no direct practical advantage in acquiring a representation of the maze, and there is a practical advantage in acquiring a representation of the home cage, and the subjects acquire a well-defined representation of the maze —, people who adopt Tolman and Honzik's explanation should accept that the rats acquire geometric representations of both the maze and their home cage if they are given enough time. The supposition that they only acquire a geometric representation of the maze — but not of their home cage — despite the asymmetry in practical usefulness and length of time spent in both environments is extremely doubtful. More generally, what latent-learning experiments suggest, from the perspective of the GM framework, is that subjects start encoding a fine-grained geometric representation as soon as they are put in a new environment regardless of whether they have practical incentives to do so, and that the geometric-representation database can contain representations of at least two distinct environments simultaneously.

3.2. Second Step: Transformational Tasks Can Interfere with the Selection Mechanism

We have just adduced evidence for the view that the geometric-representation database can contain more than one geometric representation at once. Thus, there must be a process by which one geometric representation is chosen from all those available in the database at the beginning of the reorientation process in order to initiate the congruence computation. I call the mechanism that performs this process *the selection mechanism*.

A fundamental question about the selection mechanism is: How does it operate? In particular, what kind of information does it use to select the relevant geometric representation at the outset of a trial? I will provide a more detailed discussion of some options about how to deal with this issue in the next section. But for now, I want to argue that the selection mechanism gives at least some modest role to featural cues in its operations. More specifically, that the selection mechanism relies on what I will call *feature maps*. A feature map, as I will employ the term in this thesis, is a representation of a set of featural cues identified in a given environment along with their location in it. Each feature map is indexed to a corresponding geometric representation of the same environment. In feature maps, featural cues are associated with an address label which specifies the location of each cue in terms of its distance and/or extension with respect to the three-dimensional extended surfaces of the environment. For example, the feature map of the training enclosure of the experimental group in Figure 2 encodes the presence of the color black, and it associates that color with an address label that specifies that it completely covers the two *long* walls. The same feature map also encodes the presence of the color white, and it associates the color white with an address label that specifies that it completely covers the two *short* walls. Feature maps may also include the identity and location of other types of featural cues besides colors and two-dimensional patterns on walls, like prominent isolated objects or possibly even salient smells.

It might be thought that the GM framework categorically rejects the idea that subjects of reorientation tasks acquire anything like feature maps. But that is not the case at all. For various reasons, some highly influential GM theorists (e.g., Cheng, 1986, 2005; Gallistel, 1990; Gallistel & Matzel, 2013) hold that subjects store feature maps of the experimental enclosure in reorientation tasks along with geometric representations of the enclosure. As long as the geometric module itself is not taken to operate on or retrieve any information from feature maps, the commitment to feature maps is compatible with the GM framework.

So, what is the evidence that the selection mechanism relies on feature maps? It comes from a set of reorientation experiments involving multiple experimental enclosures. Consider first the following experiment by Julian et al. (2015). They alternately exposed mice to two white rectangular enclosures, each with a set of three salient vertical black stripes (Figure 3). The enclosures were built of the same material and they had the same dimensions. Importantly, the vertical black stripes were on one of the *short* walls in one of the two enclosures, and vertical black stripes covering the same width were on one of the *long* walls in the other enclosure. Moreover, the correct corner had opposite geometric properties in the two enclosures: in one enclosure, the correct corner had a short wall on the left and a long wall on the right; in the other enclosure, the opposite situation obtained. In this experiment, the subjects chose the diagonally adequate corners significantly more often than the diagonally inadequate corners in both enclosures, and they did not choose the goal corner significantly more often than the diagonally opposite corner in either enclosure.

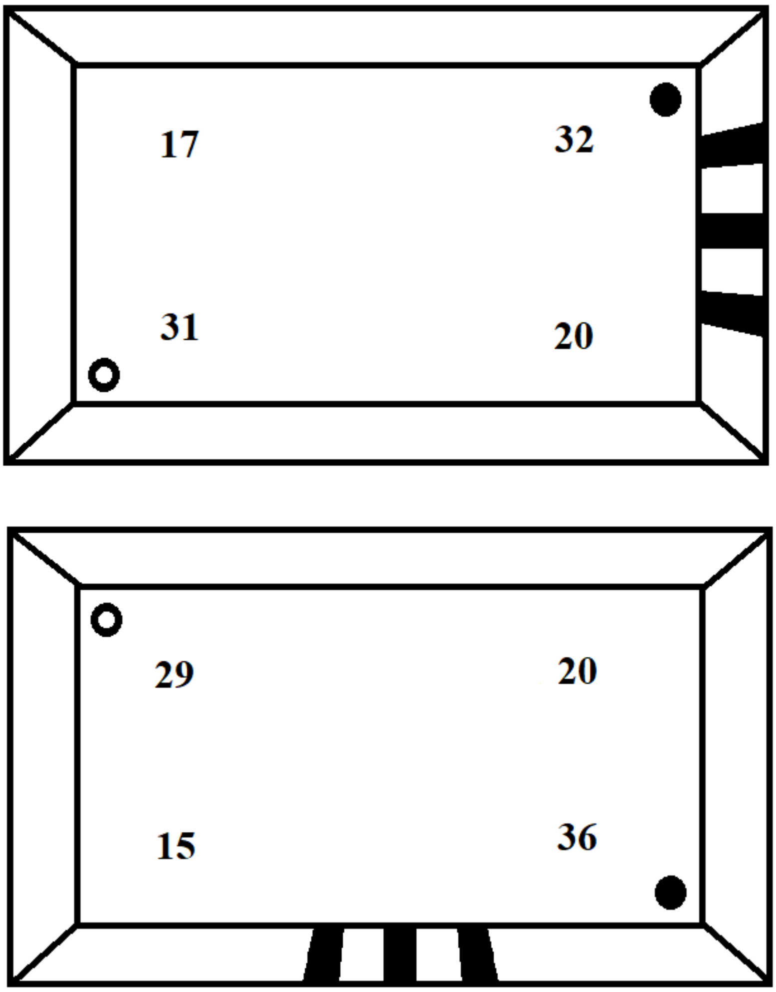

Figure 3. The enclosures used in the experiment by Julian et al. (2015) described in the main text. Black dots indicate the location of the hidden food source in a given enclosure (the goal location). White dots represent the location of the diagonally opposite corner. Numbers in each corner indicate the percentage of trials where the subjects first dug in that corner. (Based on Figure 3 in Julian et al. 2015.)

From the perspective of the GM framework, this means that Julian et al.’s subjects reliably selected the geometric representation from the current enclosure when returning to it. To see why, recall first that GM theorists hold that, though other processes can take control of behavior in various contexts (like the beacon-homing process, see Chapter 1:3), the process responsible for rotational errors is the geometric module. Given that subjects perform many rotational errors in the experiment, this means that GM theorists must hold that the geometric module is responsible for guiding subjects’ behavior on most trials. Yet, the geometric module only pays attention to geometric cues to guide behavior. Thus, given that the goal corner had opposite left-right relations with the short walls and long walls in the two experimental enclosures, had the geometric module operated on the geometric representation from the other experimental enclosure it would have led subjects to search for goal in the diagonally inadequate corners of the current experimental enclosure. So, unless the selection mechanism reliably sent a representation from the current experimental enclosure to the geometric module, subjects would not have searched for the desired goal mostly in the diagonally adequate corners, which they did.

This observation is important because it gives us a way of prodding the selection mechanism's operating principles. The only obviously apparent aspect that differentiate visits to the two enclosures was the location of the vertical stripes with respect to the walls. (Subjects couldn't have simply tracked the order of presentation of the enclosures to select the relevant representation. Julian et al. varied the order of presentation of the two enclosures across daily sessions. On one day, they would start with the enclosure where the black stripes were on the short wall. On the next, they would start with the enclosure where the black stripes were on the long wall.) So, this tells us that the selection mechanism likely relies on the *location of featural cues with respect to three-dimensional extended surfaces* in order to choose the right geometric representation at the outset of a trial.

Yet, this still does not give us direct support for the view that the selection mechanism uses *feature maps* to select the relevant geometric representation. To reach that conclusion, we also need evidence that *the identity of featural cues* plays a role in the selection process. And the experiment we have just discussed does not bear on this issue because both enclosures contained all the same featural cues, despite the fact that some were not in the same position. Julian et al.'s subjects could have managed to reliably select the relevant representation based only on recorded information about the location of distinctive featural cues in the current environment, without recording or relying on information about the identity of the cue beyond the fact that it is a distinctive featural cue that stands out in the environment.

But, fortunately, there is evidence from other experiments that subjects do rely on the identity of featural cues for selection purposes. Consider in particular the Horne et al. study discussed in the previous section. For reasons similar to the ones given with respect to the Julian et al. experiment, we should assume that Horne et al.'s subjects reliably selected the geometric representation from the current enclosure to initiate the reorientation process when returning to that enclosure. Yet, the only factor that differentiated the two rectangular enclosures was the color of their walls. Otherwise, the two enclosures had the exact same dimensions and shape, were made of the same material, and were placed in the exact same location of the experimental room for a given trial. (Horne et al. also changed the order of presentation between the enclosures at different times in the experiment, so it can’t be that the subjects were just tracking the order such that one representation got associated in the subjects’ mind to odd-numbered trials and the other representation got associated to even-numbered trials.) So, the foregoing strongly suggests the selection mechanism operates both on the location of featural cues with respect to three-dimensional extended surfaces *and* the identity of featural cues. And the best way of explaining this, from the perspective of the GM framework, is that subjects have acquired a feature map from each visited enclosure in such experiments and that they sometimes use them to select the corresponding geometric representation.

From this, we can then infer that reorientation subjects possess, along with the geometric-representation database, a *feature-map database* that holds multiple feature maps of distinct environments at once. In at least some contexts, the selection mechanism makes use of representations within the feature-map database to help it pick out the best associated geometric representation in the geometric-representation database. In this way, the two databases participate in the selection process. We can also conclude that the feature-map database can hold feature maps for at least around 23 hours in rodents, given that there is at least 23 hours before the last trial of a given day and the first trial of the next day in the Julian et al. and Horne et al. experiments and that all reported behavioral measures reported in the two papers suggest that subjects chose diagonally adequate corners just as frequently on the first trial of a day as on the other trials.

We are finally in a position to formulate more clearly the central idea behind the selection interference strategy as follows. Any transformational task introduces a modification in the experimental enclosure before the test trials. However, such a modification can easily interfere with the normal operations of the selection mechanism. This is especially likely given that many transformational tasks involve moving salient featural cues within the experimental enclosure or even removing them entirely from the enclosure, and that the selection mechanism pays attention to feature maps which are precisely meant to capture the location and identity of featural cues. Transformations of the experimental enclosure can thus lead to many different types of outcomes that differ from the typical predictions that reorientation researchers attribute to the GM framework about transformational tasks. One possibility is that a given modification to the enclosure leads the selection mechanism to pick out a representation of an environment other than the experimental enclosure (like the representation of the home cage, for example) during the test trials. In this case, the geometric module would likely make a wrong estimate of the agent's heading and location based on the geometry of that other environment. The agent is therefore likely to look at the wrong location for a specific goal.

Another possibility is that the selection mechanism simply fails at singling out even one geometric representation following the changes to the experimental enclosure because no stored representation adequately matches the currently perceived environment well enough according to the selection mechanism's internal criteria. Thus, the geometric module would not even be triggered following the subject's recovery of normal perceptual access in the changed enclosure. Another mechanism would presumably take control of spatial behavior at this point, a mechanism which may well make use of featural information to guide the subject's behavior.

Yet another possibility is that the selection mechanism manages to pick the relevant representation and send it along to the geometric module, but that in doing so the selection mechanism also notifies other navigation processes and systems, through a specific type of signal, that the current known environment has been changed. The signal may thus lead subjects to (e.g.) wander the whole environment and pay closer attention to salient cues that have been changed, be they geometric or featural cues.

Obviously, the foregoing analysis still leaves us short of a full explanation — let alone a prediction — of the variation in the results of transformational reorientation tasks. What is important to realize at this point is that, in order to provide a full explanation of the variation in the results of transformational tasks within the GM framework, some empirically motivated decisions must be taken about how to further develop GM theories. We will discuss the most relevant issues in that regard in the next section.

4. Seven Open Issues That Matter for Interpreting Transformational Tasks

In this section, I will go over seven open issues from within the GM framework that are relevant to explaining the results of transformational tasks by appeal to the selection interference strategy. I cover one issue per subsection below. The first two issues are explicitly discussed in the reorientation literature. The third one is rarely discussed. The last four issues have been ignored entirely because reorientation researchers never explicitly address the possibility of selection interference. I will spend more time on each of these last four given that they are particularly relevant in trying to explain the results of transformational tasks.

4.1. What is the Content of Geometric Representations?

What is the content of the geometric representations stored in the geometric-representation database? Those representations must encode geometric information about three-dimensional extended surfaces — that is where the name of the GM framework comes from —, but how detailed does that encoded information have to be? There are various possibilities. In his early work, Cheng (1986) simply assumes that the whole shape of the extended surfaces in a given environment is encoded in great detail in the corresponding geometric representation. On his view, the representation registers that shape in all its curves, angles, gaps, sides and the concavity/convexity of its various sides. In later work, however, Cheng (2005) suggests that we should conceive of geometric representations as encoding very minimal geometric information, only what are called *shape parameters*. Those parameters are "like summary statistics or abstracts" (p. 8) of the shape of extended surfaces in the agent's environment. For instance, Cheng (2005) discusses the possibility that only the axes of symmetry of the shape of extended surfaces are encoded when an agent wanders through a given environment. Any more detailed geometric information — like the number of sides in the shape, their length, the angles at which they meet or their curvature — is simply not encoded in the geometric representations. (See Cheng 2005, Section 3, for two other ways of conceiving of shape parameters.)

4.2. How is the Alignment of Geometric Representations Achieved?

A major commitment of the GM framework is that the geometric module's role in reorientation tasks is to estimate the agent's heading and/or location by aligning two representations, the representation of the currently perceived environment and the representation of the remembered environment. But how is the alignment procedure achieved? There are different proposals in the literature about this. Following Gallistel (1990), most GM theorists explicitly assume that the process unfolds by comparing shape parameters extracted from both of these geometric representations. One possibility, for instance, is that the congruence-finding process operates by rotating the symmetry axes of the remembered shape so that they coincide with the symmetry axes of the currently perceived shape — in the way in which you can align two congruent rectangles by rotating them until the pairs of symmetry axes of one of them coincides with the pairs of symmetry axes of the other. All the current versions of the GM framework which address this issue explicitly presuppose that the congruence-finding process makes use of shape parameters, though they often avoid committing themselves on which kind of shape parameters they use. Note also that, even if the congruence-finding process works with shape parameters, it doesn't automatically entail that the geometric representations stored in the database only record shape parameters (and that they cannot record the shape of extended surfaces in great detail, as implied by Cheng 1986). All that Gallistel's suggestion about the use of shape parameters entails is that, if those geometric representations encode very detailed information about the shape of extended surfaces, shape parameters must be extracted from this detailed encoding before the congruence-finding process begins.

4.3. Interactions of the Geometric Module with Other Navigation Mechanisms

How does the geometric module interact with other cognitive systems and how far removed is it from motor systems? Opponents of the GM framework often assume more or less explicitly that the geometric module's output goes directly to the motor systems and thus has direct control over behavior whenever subjects recover normal perceptual access to their surroundings in a reorientation trial. This assumption is wrong for two main reasons. First, as GM theorists conceive of it, the geometric module's role in reorientation is to estimate the agent's heading and location after recovering normal perceptual access to his surroundings. In other words, the geometric module's function in reorientation *is to make available a representation of the agent's heading and location to specific downstream cognitive systems controlling spatial behavior after recovery of normal perceptual access*. It thus only has influence over spatial behavior through its impact on systems that make use of the geometric module's estimation of the agent’s heading and location to govern spatial behavior after reorientation. It has no direct control over motor systems beyond that. Second, as we noted in Chapter 1:3, many GM theorists explicitly acknowledge that reorientation subjects rely on a beacon-homing process sensitive to featural cues to return to the correct corner significantly more often than any other corner in rectangular enclosures in some specific contexts. So, at the very least there must be some sort of mechanism that oversees the competition between the beacon-homing process and the geometric-module-based goal-finding process to determine which one of the two processes will take control of the subject’s behavior in a specific context.

Related to these two points, the third issue concerns largely two things: (1) the nature of the interactions between the geometric module and downstream mechanisms whose role is to determine whether to look for a certain goal based on the geometric module’s estimation of the subject’s heading and location within the current environment and, if so, what path to choose; (2) the nature of the processes that compete with the geometric module to guide behavior and of the overseeing mechanism that governs the competition between them.

4.4. In Which Conditions Are Geometric Representations and Feature Maps Discarded?

I noted above that geometric representations and feature maps are held in memory through their respective databases. But there are many theoretical options about the conditions in which these two databases *discard* their own representations. Here is a non-exhaustive list:

Time limit: Representations contained in either the geometric-representation or feature-map database are discarded after a fixed period of time following the moment when they were first recorded or last selected. We saw in Section 3 that various empirical findings suggest that the geometric-representation database of rats, chicks and splitfins can, in some contexts, keep a geometric representation for at least 19 hours after they last selected it. Similar reasoning led us to conclude that the feature-map database of rats and mice can hold a feature map for at least 23 hours after they last used it for selection purposes. Hence, it could be the case that, for these species, all geometric representations and feature maps are kept in their respective databases for a fixed period of time without being selected — 24 hours, say —, but that they are automatically discarded after that.

Size limit: The two databases have a size limit. One possibility is that the limit pertains to the *number* of geometric representations or feature maps that they contain. For example, we could imagine that representations of both types are recorded from the last six environments visited by the subject. Then, as soon as the agent enters a new environment and encodes the corresponding geometric representation and feature map, the oldest of the six geometric representations and the oldest of the six feature maps in the two databases are discarded to make place for the new geometric representation and the new feature map. Another possibility is that the limit stems from the *quantity* of information that both databases maintain at a given time. In that scenario, there is no precise maximum number of geometric representations or feature maps which can be kept in the databases at once. The maximum number of representations held at once in either database depends on the complexity of the geometric representations or feature maps being recorded as the subject wanders in the world. As their complexity increases, fewer representations can be kept at a given time.

Usefulness considerations: Geometric representations and feature maps are kept in storage based mainly on usefulness considerations, such as whether they depict environments where food, water, shelter, prey, predators, conspecifics and/or any kind of threat or reward are found. If none of these elements are discovered in a given environment, then the corresponding geometric representation and feature map are discarded much faster from their respective databases. We could also imagine that, if a given geometric representation or feature map is recorded close to the nest, then it is kept in storage longer. If it is recorded further away from the nest, it is discarded faster.

Interactions between time, size and usefulness considerations: Geometric representations and feature maps are discarded based on the time since they were first recorded (or last selected), on the database's size limit as well as usefulness considerations. One version of this suggestion might claim, for instance, that each database has two subcomponents — one for representations that depict environments where food, water, shelter and threats are found, and another database for representations that do not depict such environments —, and that these two databases have a different size and/or time limit.

4.5. How Does the Selection Mechanism Work?

The selection mechanism's function is to choose a geometric representation from all those available in the geometric-representation database and then feed it to the geometric module so that the geometric module aligns it to the representation of the currently perceived environment. We have seen in Section 3.2 that there is strong evidence, from experiments like those of Horne et al. and Julian et al., for the view the selection mechanism gives a role to feature maps in its operations. But this does not mean that the selection mechanism can only take into account information contained in feature maps to pick out a geometric representation. Though it is a possibility, many other types of information could also be used. We can thus imagine various general proposals about the principles guiding the selection mechanism:

Feature maps alone: The selection mechanism chooses one geometric representation only in virtue of how well its associated feature map matches the identity and location of featural cues in the current environment. This might seem like a simple option. But, in fact, it raises many questions because there are infinitely many ways of measuring the closeness between a feature map's depiction of featural cues and the featural cues in the currently perceived environment. It is possible to measure the closeness, for instance, by focusing only on some of the following categories of featural cues, but not others: (i) the color of three-dimensional extended surfaces; (ii) the two-dimensional shape patterns on three-dimensional extended surfaces; (iii) the smell surrounding extended surfaces; (iv) the dimensions and shape of isolated objects; (v) the color of isolated objects; (vi) the smell of isolated objects. Even when we consider only this non-exhaustive list of types of featural cues, there is a lot of latitude concerning how each type could be taken into account by the selection mechanism. For example, the selection mechanism could, in theory, look only for feature maps which depict extended surfaces with colors that match almost perfectly the colors of the extended surfaces currently perceived — or use a lower standard of color resemblance. If it simultaneously uses two or more categories of featural cues, it could weigh discrepancies among them in a great many ways.

Feature maps and geometric information: The selection mechanism chooses one geometric representation as a function of how closely its depiction of the global geometry of three-dimensional extended surfaces matches the geometric cues in the current environment and of how well its associated feature map matches the current environment. For instance, it could be that the selection mechanism looks primarily for a high level of matching related to geometric cues, relying on feature maps only as tie-breakers if more than one geometric representation has a high level of matching with the geometric cues in the current environment. Many other schemes are also possible, such as ones that give priority to feature maps over geometric information instead. Even putting issues of priority aside, there are many possible ways of spelling out this option because assessing the level of matching of a given geometric representation with the geometric cues in the current environment can be done in different ways. It can be done, for example, by tallying differences in many subtle details of the represented three-dimensional extended surfaces — such as various aspects of their curvature and height, say — on the assumption that geometric representations are highly detailed in their content (see first issue above). Or it can be done by comparing very simple shape parameters pertaining to the current environment and the stored geometric representations, parameters which abstract away from much of the details of the geometry of three-dimensional extended surfaces. Note also that some transformational tasks make important changes to the shape of the experimental enclosure for the test trials (Pearce, Good, Jones, & McGregor, 2004; Tommasi & Polli, 2004). Thus, if this option is right, such changes could cause selection failure in a variety of contexts. That would happen because, in those cases, no stored geometric representation matches the geometric cues in the current environment well enough to be selected, and the selection mechanism simply gives up on its task.

Feature maps and contextual information: The selection mechanism picks out one geometric representation as a function of how well its associated feature map matches the current environment as well as contextual cues that have nothing to do with the currently perceived environment: for example, the time elapsed since the agent was last disoriented, the environment in which the agent resided just before losing normal perceptual access to its surroundings, the time of the day when the agent lost perceptual access to its surroundings, how the subject lost perceptual access to its surroundings (e.g., sudden versus not sudden). For instance, the selection mechanism could work in such a way that (i) when the subject loses perceptual access to its surroundings in the morning, it seeks a geometric representation of a cage that the subject has seen most often *in the morning* and whose feature map matches the current environment well, and that (ii) when the subject loses perceptual access in the afternoon, it looks for a geometric representation of a cage that the subject has seen most often *in the afternoon* and whose feature map matches the current environment well. Of course, it is also theoretically possible that the selection mechanism uses all these different types of contextual cues simultaneously to choose one geometric representation, or only some of these contextual cues in some contexts.

Feature maps, geometric information and contextual information: The selection mechanism chooses one geometric representation by paying attention to the three previous types of information. Unsurprisingly, there are even more ways of conceiving of the operations of the selection mechanism under this option. On top of all the parameters described in the three previous paragraphs, we can distinguish three general suggestions here: (i) the selection mechanism prioritizes one type of cues (featural, geometric or contextual) over the two others so that this type always wins over the two others; (ii) the mechanism weighs one type of cues above the others yet this type does not always win over the others; (iii) the mechanism treats all three types more or less equally. As an instance of the first suggestion, we could imagine that the selection mechanism always begins by seeking one geometric representation whose depiction of the geometry of extended surfaces closely matches the geometry of extended surfaces in the currently perceived environment. However, if it finds two distinct geometric representations whose depiction of the geometry of extended surfaces matches roughly equally well the currently perceived environment, then it chooses between the two representations in virtue of featural and contextual cues. We can also imagine more static or dynamic versions of these frameworks. For example, on a dynamic version of suggestion (i), the selection mechanism may use statistical information about the prior reliability of each of the three types of cues in various spatial endeavors in order to prioritize the most reliable of the three types of cues.

4.6. What Happens When There is Selection Failure?

What happens if the selection mechanism cannot find even a single geometric representation in the database by the usual means after the agent recovers normal perceptual access to its surroundings? Which cognitive mechanisms guide or take control of the agent’s spatial behavior then? How do they operate? Consider the following example. In Pearce et al.’s (2006) Experiment 1, half of the subjects were trained in a kite-shaped enclosure. Though Pearce et al. do not mention anything to that effect, it is rather likely that the rats had never seen an enclosure with that particular shape before. Their rats were naive subjects (i.e., they had never participated in an experiment before this one), and lab rats are generally raised in rectangular or square-shaped cages. They probably hadn't seen any environment distinct from their home cages before the experiment started. So, if the rats' selection mechanism paid attention to the geometry of extended surfaces, it may have been unable to find even a single geometric representation whose depiction of the geometry of extended surfaces approximately matched the perceived kite-shaped environment the very first time they were put into this new apparatus. What happens then? We know that the rats did not just freeze on that first trial. They started exploring the enclosure. If there was indeed selection failure, which navigation process controlled their spatial behavior? What kind of principles did that process operate on?

Again, we can imagine various options pertaining to this issue. Here is a representative, yet non-exhaustive list:

New selection attempt: When there is no match in the geometric-representation database, the geometric module induces a second, more relaxed selection attempt. Rather than relying on a stringent selection principle (e.g., *the geometry of three-dimensional extended surfaces in the current environment must match perfectly the selected geometric representation's depiction of the geometry of extended surfaces*), the selection mechanism seeks a representation that approximately matches the current environment along some dimensions (context, geometry, features, or any combination of them). The geometric module is then activated with that approximately-fitting geometric representation to make an estimate of the agent's heading and location. Spatial navigation then resumes as in normal contexts, with a potentially wrong estimate of the agent's heading and location. Thus, the subject might head directly for a specific location in the environment and start looking around that location for a certain goal — a goal it expects to find in that location based on its potentially wrong estimate of its current location and heading.

Beacon-homing process: When the selection process fails, the geometric module is bypassed altogether and the subject relies on a beacon-homing process to guide its behavior (see Chapter 1:3 for a more detailed description of this process). This process motivates the agent to move toward a single featural cue in the current environment, based on that cue’s past reliability in finding a specific desired object or location. For instance, if a certain color, two-dimensional pattern or smell has always been found close to some significant location, it can become the target of the beacon-homing process. The process then guides the subject’s behavior by motivating it to move in a straight line toward the relevant featural cue. In doing so, the process ignores geometric cues altogether, or even the presence of other types of featural cues in the environment.

View-matching process: When the selection process fails, the geometric module is bypassed and a view-matching system takes over. This view-matching system operates on snapshots — stored representations of the subject’s two-dimensional retinal stimulation at specific locations of the environment — to guide behavior. One of those snapshots might, for instance, capture all the significant edges in the agent’s visual field when it is facing a source of food from a north-east direction. (This view-matching option is inspired by view-matching theories of reorientation. See Chapter 3 for a detailed description of such theories.) There are two ways of thinking about how this view-matching system could govern spatial behavior following selection failure. On the one hand, the view-matching system might simply make an estimate of the agent’s heading and location (as on the view-matching theory proposed by Sheynikhovich, Chavarriaga, Strösslin, Arleo and Gerstner, 2009). Once this is done, the type of path-planning mechanism on which the geometric module normally relies to guide behavior based on an estimation of the subject’s heading and location (see third issue) takes control of the subject’s spatial movements. Spatial navigation then resumes as in normal reorientation tasks, but with a potentially ill-fitting estimate of the agent’s heading and location. On the other hand, the view-matching system might directly guide the subject’s spatial movements. It could do so, for example, by directly motivating the agent to move to places which minimize the discrepancy between its current retinal stimulation and one snapshot retrieved from memory in order find a specific goal associated with that snapshot (as on the view-matching theory proposed by Stürzl, Cheung, Cheng and Zeil, 2008). The view-matching system would thus eschew any attempt at estimating the agent’s heading and location in the current environment.

This option re-conceptualizes view-matching theories of reorientation as theories concerning what happens when the agent is in what it considers to be a new environment and it cannot rely on the geometric module’s usual methods to establish the location and heading within the current environment. On that option, the view-matching system would work as a kind of backup system: It may not be as efficient as the geometric module to help subjects return to known goal locations, but it is better than nothing. When re-conceived as theories of what happens when the geometric-representation selection process fails, I see no reasons why GM theorists could not endorse these view-matching models. The GM framework is committed to the existence of a mechanism that relies only on the geometric information about three-dimensional extended surfaces in its operations (the geometric module), but it says nothing about the kind of information used by distinct, parallel systems that might guide subjects' behavior when that mechanism is not activated.

Exploration process: In a large variety of experimental conditions, when non-human animals are released in an environment that differs significantly from any that they have encountered before, they display a specific type of behavior known as *exploratory behavior* (O'Keefe & Nadel, 1978). Typical exploratory behavior consists in moving around the environment to sample large areas of it, often spending time around salient cues to look at them, sniff them and touch them. A surprising amount of research has been done concerning the characteristics of exploratory behavior in new environments (see Thompson, Berkowitz, & Clark, 2018, for a review), tackling issues such as what paths subjects choose to sample the environment (the chosen paths have important random aspects, but subjects are not just moving in a random direction after every new step), whether exploring subjects spend more time sampling boundaries like walls and cliff edges than other parts of the environment (they seem to), and how they approach specific isolated cues. It seems likely that selection failure in the test trials of a transformational task could trigger, at least in some contexts, an exploration process. After all, selection failure consists, in an important sense, in the realization that you are in a new environment. And the main function of exploratory behavior on most models of animal spatial navigation is precisely to get the animal to move around so that it can acquire information about new environments.

Whether reorientation subjects undergo exploration phases in the test trials of transformational tasks matters a lot when it comes to the interpretation of transformational tasks because reorientation researchers assess subjects' performance based on behavioral measures such as the following: where subjects dig first upon entering the environment, which corner they head toward upon being released, which corner they spend more time in overall. So, for example, it might be that reorientation subjects spend a lot of time near one corner that is closest to a salient featural cue in the test trials of a given transformational task because the selection mechanism has failed to choose one geometric representation and the subjects want to ascertain the properties of this cue as part of an exploration phase of the environment. Such exploratory behavior would not undermine the GM framework in the slightest, though it may go against an overly simplistic interpretation of the GM framework which takes the framework as committed to the view that reorientation subjects' behavior must always, or almost always, be guided by geometric cues alone, even in transformational tasks.

Personal-level problem solving relying on episodic memory and conscious, procedural strategies (humans only): It is sometimes said or assumed that the main cognitive processes involved in reorientation operate at the unconscious level (cf. Lee & Spelke 2010a, p. 174). This is certainly a reasonable hypothesis from the perspective of people defending any of the well-known theories of reorientation. (It would appear implausible for GM theorists to claim that humans consciously extract shape parameters from two distinct geometric representations in effecting reorientation, or for view-matching theorists to hold that we consciously use a method of vector descent to minimize the discrepancies between the current view and a remembered view. When I try to find my bearings after recovering normal perceptual access to my surroundings, I do not have any such conscious experiences. I assume it is the same for other humans.) I also presume that the first three options presented in relation to this issue are plausible only as descriptions of what happens at the unconscious level once the selection process fails. But, in trying to tackle the sixth issue, GM theorists could plausibly defend the view that (a) if the selection mechanism fails at finding even one geometric representation, the agent then immediately becomes consciously aware that he is in a novel environment; and that (b) this triggers other types of potentially non-spatial — even somewhat domain-general — cognitive processes. For instance, agents, especially human adults, might rely on consciously accessible experiences or knowledge to make their spatial decisions following a failed attempt at selecting a geometric representation in their geometric-representation database. Suppose that I fail to recognize where I am once I step out of the elevator in the building where I work. The lobby room in which I currently am is like no other room I have seen in my life: it contains various idiosyncratic geometric and featural cues which I have never encountered before. I might then consciously recall that I was told that there would be major renovations on our floor over Christmas. I thus conclude that I am actually on the right floor, the floor where I meant to be. I then also remember that the location I am aiming to reach is roughly to the north of the elevator, and I walk north as a result. One way in which this sort of navigational decision-making could happen is if some episodic-memory-involving process directly provides the geometric module with the relevant geometric representation — thereby completely bypassing the selection mechanism —, and the geometric module then performs its normal function with that representation. Or consider the following scenario: I find myself completely lost when I visit a new city. Then, I consciously remember that my friend told me to turn left at the intersection of 1st Avenue and 2nd Street and head toward a blue building. I turn left at the relevant intersection as a result of what she told me, and I head for that blue building. In a situation like this, we cannot explain what happens by appeal to an episodic-memory-involving process that picks out the relevant geometric representation. *That is because there is no geometric representation to be picked out the first place.* I simply never had the relevant geometric representation in my database because I have never been in that city before. The best we can say is that the interaction between my episodic memory and the language faculty allowed me to re-construct and follow a conscious, response-based navigation strategy to find the desired goal without triggering the geometric module at all.

Of course, there is a danger in adopting this fifth option in that it risks implicitly committing GM theorists to the view that some highly flexible, domain-general and unencapsulated higher cognitive mechanisms pick up the slack if the selection mechanism cannot do its job normally, thus undermining the case for the massive modularity thesis. Endorsing and spelling out this option in more detail therefore constitute a risky enterprise for GM theorists, at least for those sympathetic to the massive modularity thesis. On the other hand, I assume that most people will have experienced something like the two scenarios just described. So, it is not clear that GM theorists can deny the possibility of this option outright.

A combination of all this: As you might imagine, yet another option is that two, three, four or five of the mechanisms just described compete or cooperate in various ways following selection failure. I will not go into details about how these mechanisms might compete or cooperate, or how their output might be weighted against each other by an overarching mechanism in such a case. One possibility which I find attractive for human adults, however, is that the mechanism for conscious choice competes or cooperates with one of the other four unconscious mechanism proposed above — at least in some cases. This would appear to make sense of two seemingly opposed phenomena about human adult navigation: (1) in some contexts, people claim to have had various conscious experiences that helped them find their way after getting lost, and their description of how they managed to find their way seems to match their behavior (e.g., Marie claims to have used a blue building to find her bearing following the conscious remembering of her friend’s advice, and she really did visually track a blue building while moving around); (2) in some contexts, people find their way and they cannot explain how they found their way, or the explanation they provide does not match their behavior (e.g., Mark claims to have relied on a certain landmark to find his way, but in fact he did not even see that landmark until he was at the desired location).

4.7. What Happens When the Selection Mechanism Triggers the Environmental-Change Signal?

I mentioned in Section 3.2 that it is possible that, in some situations, the selection mechanism manages to pick the relevant geometric representation and send it along to the geometric module, but that in doing so it triggers a signal to the effect that the current known environment has been changed to some important degree. Call this signal *the environmental-change signal*. Positing the existence of such a signal may seem like a far-fetched move, brought about by desperation to find an answer to the explanatory inflexibility objection. But, in fact, there are strong independent considerations, both empirical and theoretical in nature, for the existence of such a signal (O’Keefe & Nadel, 1978; Poulter, Hartley, & Lever, 2018).

I will mention one empirical consideration here, one that strikes me as particularly telling. It pertains to electrical oscillatory patterns in the hippocampus of mammals as they explore known environments. In normal contexts, when mammals move around an environment, the local field potential of their hippocampus oscillate around a specific frequency band. The hippocampal oscillatory pattern within that specific band concurrent with the subject moving is known as *the theta rhythm*. In rats, the band in question is roughly between 7 Hz and 11 Hz. Importantly, when rats have spent a significant amount of time over a number of visits in a given environment that has not changed in visual appearance, the frequency of the theta rhythm increases slightly over a period of few days (while remaining within the 7-11 Hz range), so that at the end of period it is significantly higher than at the beginning. And it turns out that, in some circumstances, if experimenters change some important aspects of the experimental enclosure while leaving that enclosure in the same overall context (i.e., in the same larger, experimental room), the subjects' theta frequency drops significantly below the initial frequency on the very first visit to the enclosure (Jeewajee, Lever, Burton, O'Keefe, & Burgess, 2008).

Jeewajee et al. interpreted their results as showing, among other things, that rats can distinguish situations where they are currently in a *known, recognized environment whose properties have changed in some important way* from situations where they are in a *completely new, unrecognized environment*. Jeewajee et al. controlled for a large variety of factors that could have explained why the theta frequency dropped below the initial frequency, and they plausibly argue that the relevant factor is the fact that a known, recognized environment has been changed. If subjects had simply treated the changed experimental enclosure as a new environment, the frequency of the theta rhythm should have been similar to the initial frequency at the beginning of the experiment, which it was not.

Now, I am not claiming that sudden drops in theta frequency themselves constitute the environmental-change signal. I am claiming, rather, that variations in the theta frequency across testing conditions such as Jeewajee et al.'s provide *indirect* evidence of the existence of such a signal, which may or may not itself be neurally realized in sudden drops in the theta frequency. It is also natural to think that the selection mechanism would be responsible for triggering the environmental-change signal because, from the perspective of the GM framework, the process of selecting the relevant geometric representation to feed to the geometric module consists in many ways in recognizing which environment you are in.

On this view, the selection mechanism would decide in some cases to alert other navigation processes through the environmental-change signal that changes to the environment have been noticed with respect to prior visits. This would happen in cases where the selection mechanism has found a geometric representation which matches the current environment well enough to be selected and sent to the geometric module, but where the fit is not that great either (as determined by internal criteria used by the selection mechanism to determine whether a given geometric representation is a good fit for the current environment, see fifth issue above). The environment-change signal would thus imply or convey to other navigation processes that the output of the geometric module's estimation of the subject's heading and location may not be as reliable as usual to find a desired object or practically relevant location.

Part of the function of the signal may thus be precisely to activate navigation processes that would not normally guide behavior when the geometric module has been triggered by the reception of a geometric representation from the selection mechanism. For example, production of the environmental-change signal could make it more likely that subjects will rely on one of the following processes (discussed in the sixth issue) rather than the geometric module to guide behavior in reorientation tasks: the beacon-homing process, a view-matching process or an exploration process. In this case, subjects may not go at all to the geometrically appropriate corners. Whether any of these processes is activated following the environmental-change signal, and if so which one, could also be modulated by the history of reliability of each process, as well as by various contextual and emotional factors. After all, various contextual and emotional factors affect subjects' tendency to display exploratory behavior in other types of navigation experiments (O'Keefe & Nadel, 1978; Thompson et al., 2018).

5. The Flexibility of the Selection Interference Strategy

It should be clear by now that the GM framework makes much clearer predictions regarding search behavior in non-transformational reorientation tasks than in transformational tasks. For non-transformational tasks, it says that the subjects' spatial behavior will initially be guided by geometric cues only, unless subjects have been exposed to stable and distinctive featural cues in the geometrically appropriate corners for a long time (in which case, subjects will rely on a beacon-homing process that pays attention to featural cues to return to the correct corner) — barring exceptional circumstances or the interference of other cognitive systems. For transformational tasks, a given GM theory cannot yield any specific prediction at the moment without precise commitments on all seven issues.

How can we make progress on these issues? Some of them could be tackled relatively directly with experimental designs that do not steer far from those of typical reorientation tasks. For example, in Section 3.1, I have suggested a type of reorientation experiment that would allow us to test for the existence of a time limit pertaining to the geometric-representation database. In Chapter 5, I will also make altogether new predictions in the reorientation literature based on specific commitments about the content of geometric representations. Performing the experiments proposed in Section 3.1 and Chapter 5 would thus constitute a first step in tackling the fourth issue and the first issue, respectively. We may well need to appeal to more complex analyses and theoretical considerations to settle the other issues.

We should note one thing here, however. Regardless of exactly which options appear most plausible in the end regarding each of these seven issues, the GM framework will still possess a good amount of flexibility to explain the results of transformational tasks by means of the selection interference strategy. In order to explain why subjects relied on geometric cues instead of featural cues in one specific reorientation experiment (or vice versa), the GM theorist will be able to appeal to many details of the experiment (e.g., the time gap between training trials and test trials, the extent of the transformation of the enclosure between training trials and test trials, the location of goals in other environments than the experimental enclosure whose geometric representations may have been wrongly selected) or invoke specific rules governing the activation of various processes which may guide subjects' search behavior following selection failure or the triggering of the environmental-change signal (e.g., the beacon-homing process, a view-matching process, an exploration process). More generally, we can already distinguish four general moves that the selection interference strategy makes available to GM theorists in order to provide potential explanations for the results of transformational tasks:

> **Absence of the relevant representation** — GM theorists can claim that the geometric representation of the experimental enclosure recorded during the training trials of a given transformational task was discarded before the test trial for various reasons pertaining to time, size or contextual factors (see fourth issue above). Therefore, the selection mechanism could not pick that representation for the test trial. In this case, either the representation of another environment was selected, or the geometric module was bypassed altogether by the processes described in relation to the sixth issue above.

Selection of the wrong representation — GM theorists can maintain that, even though the geometric representation of the experimental enclosure acquired during the training trials of a given transformational task was still in geometric-representation database for the test trial, the selection mechanism picked out a representation that is distinct from the one that was recorded during training trials. The selection error was caused by the transformation of the enclosure between the training trials and the test trial. By relying on that other geometric representation, the subject made a wrong assessment of its own heading and location as well as the location of various goals. The subject's behavior was therefore not in line with what a more standard interpretation of the GM framework suggests the subject should do. The subject did not choose the geometrically appropriate corners from the perspective of the training trials.

Selection failure — GM theorists can claim that, even though the geometric representation acquired during the training trials was not discarded before the test trial, the selection mechanism did not manage to send *any* geometric representation to the geometric module on the test trial because no geometric representation had a good enough fit with the current environment to be singled out according to some internal criteria of good match used by the selection mechanism (as described in relation to the fifth issue). In that case, the geometric module was bypassed by one or many of the processes described in the discussion of the sixth issue.

Issue of the environmental-change signal — GM theorists can hold that, even though the relevant geometric representation was still in the geometric-representation database and the selection mechanism *did* manage to send it along to the geometric module, the selection mechanism also triggered the environmental-change signal. The signal alerted other navigation processes and systems of extensive changes to the current environment with respect to prior visits to the same environment. This had some knock-on effects that go against how a more standard interpretation of the GM framework suggests the subject should behave (see seventh issue). In particular, it made the subject much more likely to rely on the beacon-homing process, a view-matching process or an exploration process, though animals may still rely on the geometric module in some contexts despite the production of the signal.

Let us consider an example: the experimental group in Pearce et al.’s (2006) experiment discussed in Section 2. For the training trials, two of the four walls were black, the two others were white (Figure 2, left-hand side). Pearce et al. then used an all-white rectangular enclosure of the same size for the test trial (as in Figure 2, right-hand side). GM theorists may well want to claim that this change in the color of the walls interfered with the rats’ selection mechanism. But which of these four moves is most plausible? The first move appears implausible because the test trial happened only a few minutes after the last training trials (so the geometric representation of the experimental enclosure should not have been discarded due to the time elapsed since it was last selected), and the subjects did not visit any other environment besides the home cage and the waiting cage during the experiment (so the geometric representation of the experimental enclosure should not have been discarded due to size issues, given that results from the Horne et al. and Julian et al. experiments suggest that rodents’ geometric-representation database can easily maintain the geometric representations of *two distinct experimental enclosures* in similar contexts). It also seems highly unlikely that the geometric representation of the experimental enclosure would have been discarded due to lack of usefulness. After all, the experimental enclosure contained two practically relevant locations throughout the training phase of the experiment (the locations where the two escape platforms were hidden), and a given training trial would not end until subjects had found one of the two platforms.

This leaves the second, third and fourth moves. They seem about equally plausible given what we know about the experiment itself. Pearce et al. transformed the experimental enclosure for the test trial by changing the color of the two long walls so that the experimental enclosure was suddenly all white. Given that the selection mechanism pays attention to feature maps, as we established in Section 3.2, and that feature maps are precisely meant to capture things such as the color of specific walls, this suggests that the selection mechanism would not have been able to find a geometric representation whose associated feature map was a perfect fit for the current environment. This could thus easily have led to one of the outcomes described in the second, third and fourth moves — depending on exactly which option is taken regarding the fifth issue described in Section 4 as well as various details of the experiment, such as the exact dimensions of the subjects’ home cage and waiting cage and the nature of the featural cues each cage contained.

It is also possible that there are important individual differences about the threshold for selection failure (as opposed to mere production of the environmental-change signal) in such a way that no single move correctly applies to *all subjects* in some transformational tasks.[8] This would offer one more lever by which GM theorists can deal with transformational tasks whose results seem hard to summarize by sentences like "Subjects relied mainly on geometric cues" or "Subjects relied mainly on featural cues", like the transformational task discussed in the first paragraph of this chapter.

Overall, the foregoing discussion shows that the selection interference strategy gives GM theorists significant leeway to explain various dimensions of variations in the results of transformational tasks.

6. Conclusion

In this chapter, I proposed a general strategy to address the most important objection against the GM framework. The selection interference strategy, as I call it, shows that GM theorists have an adequate set of tools for dealing with the variability in the results of transformational tasks. Given the empirical backing for the existence of a geometric-representation database, a feature-map database and a selection mechanism, the strategy provides an empirically well-motivated response to the explanatory inflexibility objection. GM theorists can make four different moves to account for results showing that subjects do not go frequently to the geometrically appropriate corners in a given transformational task: they can claim (i) absence of the relevant representation, (ii) selection of the wrong representation, (iii) selection failure, or (iv) production of the environmental-change signal. Moreover, without more information about which options GM theorists have (or are likely) to take on the seven open issues highlighted in Section 4, reorientation researchers cannot appeal to the results of a few transformational tasks in order to adjudicate the dispute between the GM framework and other explanatory frameworks that aim to deal with the results of reorientation experiments.

---

[8] Some neurophysiological studies suggest the existence of individual differences about the threshold for selection failure. See, in particular, Colgin et al.'s (2010) discussion of why Wills, Lever, Cacucci, Burgess and O'Keefe (2005) rejected some of the subjects from their global-remapping study, along with the analysis of the relevance of global-remapping studies to claims about representation selection provided in Chapter 5:4 of this thesis.

Relatedly, spelling out the selection interference strategy has allowed us to see that there are many different ways of developing the GM framework beyond its central tenets. In fact, transformational tasks themselves could prove useful within the GM framework for settling the seven issues. These experiments provide indirect data about how the selection mechanism works and about what happens when that mechanism fails to pick even one environmental representation by the usual means. Of course, a given transformational task alone will not be enough to motivate a specific option on the fifth or sixth issue, say. The same result can often be explained either by claiming selection of the wrong representation or by claiming selection failure. But the results of many transformational tasks taken together — along with the results of other types of reorientation and navigation experiments — might put strong constraints on how GM theorists can develop their models.

In the next three chapters, I will discuss the debate between the GM framework and another highly influential framework of reorientation to make a claim for the overall superiority of the GM framework over the latter framework. On my way there, I will take a stance on many of the seven issues discussed in Section 4 (most notably: the first, second, third and fifth issues). My commitments in this regard are not in any sense inevitable or mandatory for GM theorists seeing as I appeal mostly to high-level, defeasible considerations to justify my choices. But I hope that they will allow us to see what a well-motivated, fleshed out GM account could eventually look like.

Chapter 3

The Debate Between the Geometric-Module and the View-Matching Frameworks has Reached an Impasse

1. Introduction

Besides the GM framework, another important explanatory framework that aims to account for the results of reorientation experiments is *the view-matching (VM) framework*. It aims to explain reorientation results by appealing to stored representations of the subject's two-dimensional retinal stimulation at specific locations of the environment. These representations are called *snapshots*. According to VM theories, a view-matching process compares previously recorded snapshots to the current retinal stimulation in order to guide the animal back to the goal location in a reorientation trial.

The next three chapters investigate the debate between the GM framework and the VM framework.[9] I discuss this debate at length for two main reasons. First, the VM framework is in an important sense the most radical alternative to the GM framework in the reorientation literature. It eschews high-level abstract representations (e.g., geometric representations) and domain-specific mechanisms (e.g., the geometric module). Instead, it gives a major explanatory role to low-level perceptual representations (snapshots) and a seemingly highly domain-general mechanism (the mechanism in charge of the view-matching process). So, if the VM framework is right, it threatens to show that reorientation results do not provide any support whatsoever for the two views identified in Chapter 1 in relation to the GM framework: the view that human and non-human animals' representations encode highly abstract properties of the world and the massive modularity thesis. The second main reason for discussing the debate at length is that the VM framework has been the most influential framework in the reorientation literature along with the GM framework since its inception. It has led to a variety of well-specified models (e.g., Stürzl et al., 2008; Wystrach, Cheng, Sosa, & Beugnon, 2011), some of which cover a wide-range of behavioral and

---

[9] See Tommasi, Chiandetti, Pecchia, Sovrano, and Vallortigara (2012) and Cheng et al. (2013) for reviews of the debate.

neurophysiological findings about spatial navigation along with reorientation results (Sheynikhovich et al., 2009) in ways that go far beyond what other frameworks provide, even including the GM framework. The VM framework has also been used to provide quantitative modeling and simulations of various important studies (Cheung, Stürzl, Zeil, & Cheng, 2008; Dittmar, Stürzl, Baird, Boeddeker, & Egelhaaf, 2010; Dittmar, Stürzl, Jetzschke, Mertes, & Boeddeker, 2014).

The overarching aim of the next three chapters is to provide a new powerful argument for the GM framework in the debate between the two frameworks. This will pass through the development of a new problem — *the representation selection problem* — that applies to both frameworks. But, before introducing this problem, I will make the case in the current chapter that, contrary to what many GM and VM theorists suppose, existing evidence does not clearly favor either framework over the other. More specifically, I will argue in this chapter that the debate between the two frameworks has reached an impasse because each framework can provide systematic explanations for the type of empirical results seen as favoring the other framework. Here is how I will proceed. In Section 2, I give a sketch of the explanatory strategy pertaining to the VM framework by presenting one influential VM account. In Sections 3 and 4, I discuss in turn the main behavioral findings in the debate and the main neuroscientific findings, arguing that each side can explain the relevant findings. The overall analysis will highlight the importance of appealing to new considerations to break the impasse, like the considerations that I introduce in relation to the representation selection problem in the next chapter.

2. The View-Matching Framework

Just as in the case of the GM framework, papers that appeal to the notion of *view matching* to explain reorientation results do not propose a perfectly unified theory. They rather put forward a family of related theories that make up a general explanatory framework, *the view-matching (VM) framework*. It rests on a core idea shared by all theories, namely that snapshots play a crucial role in guiding subjects' behavior in reorientation tasks.

Snapshots are representations that encode some aspects of the intensity, color and/or sharp discontinuities in light patterns on the retina at a specific time. Suppose, for example, that someone faces a blue desk in front of a red wall. A snapshot from her perspective would encode various aspects of her retinal stimulation produced by the desk and the wall — but nothing about the desk and the wall as a three-dimensional object and surface respectively. The snapshot may include, on the other hand, the specification of an encompassing patch of red around a smaller, trapezoid-shaped patch of blue. It could also encode higher-order information about the retinal stimulation, such as optic flow patterns at the time of recording (Dittmar et al., 2010).

How are snapshots supposed to aid in reorientation tasks? The VM framework holds that a view-matching process guides the subjects' spatial behavior by comparing previously recorded snapshots to the current retinal stimulation. In this section, I will present what I call *the movement-based model* (Cheung et al., 2008; Stürzl et al., 2008) as an illustration of a VM theory. This model represents the main way in which the VM framework has been developed and defended in recent years.

According to the movement-based model, the first operation performed by subjects upon recovery of perceptual access to their surroundings consists in retrieving a specific snapshot from memory, a snapshot *of the current environment as seen from the goal location* (Figure 4B). The subjects then deploy the following strategy: move in such a way as to *minimize the discrepancy* between that snapshot that was retrieved from memory and the current retinal stimulation. On this view, the subjects do not represent or explicitly compute the distance and the direction of the goal location ahead of their spatial movements in the environment. On the contrary, the view-matching process operates through a sort of feedback mechanism. Subjects move in a series of short steps, each of which induces a reassessment of the best direction of movement for the next step by taking into account the new visual input. They stop moving when they reach a position at which taking a new step in any direction would increase the discrepancy between the snapshot of the goal location and the current retinal input.

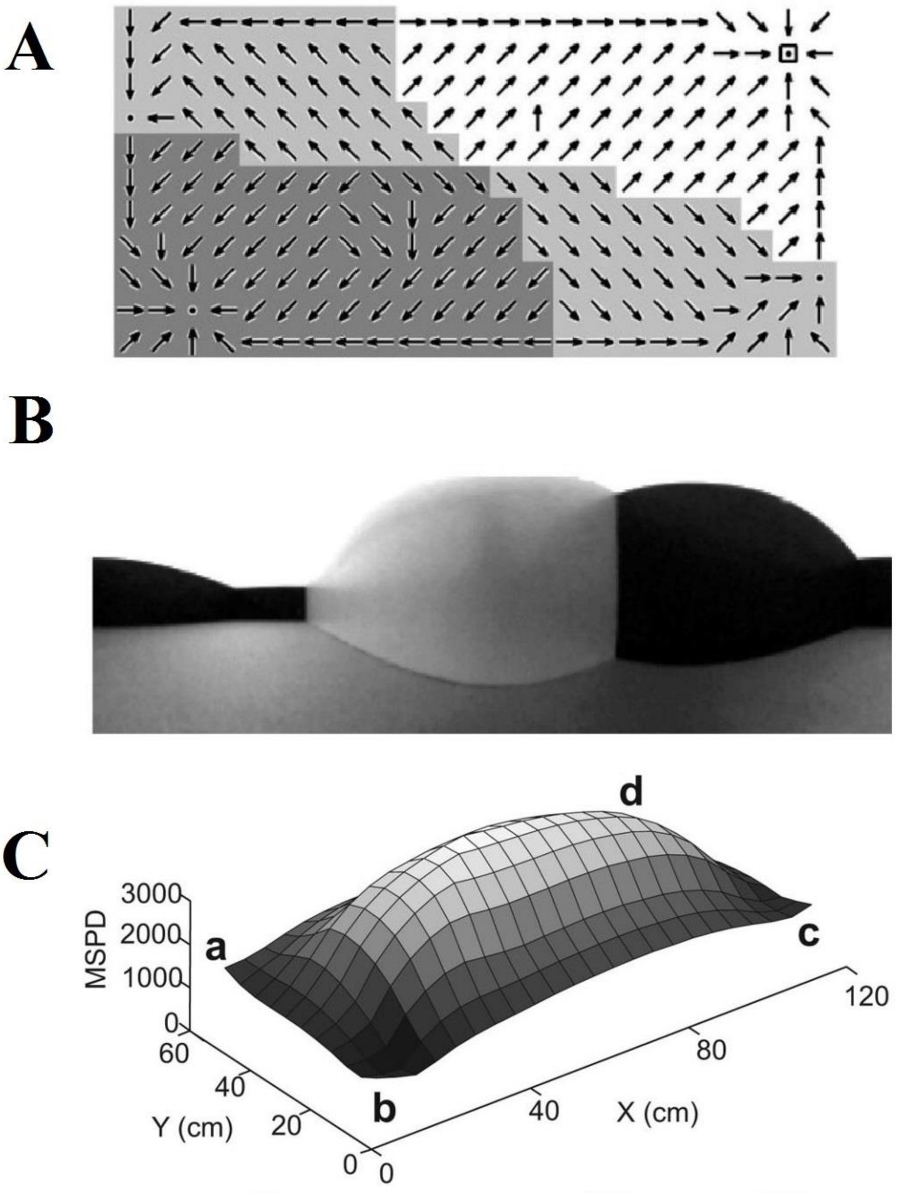

Figure 4. (A) A map of the direction that minimizes discrepancy between the image of the goal location and the current image as a function of the simulated animal’s location. (B) Image of the goal location in a camera-based simulation of Cheng’s (1986) experiment. The enclosure has three black walls and one white wall. The goal corner has a white wall on the left and a black wall on the right. (C) Amount of discrepancy (Z-axis) between the image of the goal location and the current image as a function of location in the enclosure (X-axis and Y-axis). Highest values of discrepancy are represented in white, whereas the lowest values are represented in black. (The three figures were taken from Stürzl, W., Cheung, A., Cheng, K., & Zeil, J. (2008). The information content of panoramic images I: The rotational errors and the similarity of views in rectangular experimental arenas. *Journal of Experimental Psychology: Animal Behavior Processes*, 34(1), 1–14. Copyright 2008 by the American Psychological Association. Reprinted with permission.)

Discrepancy minimization works as follows. A retrieved snapshot comprises a set of points for which various low-level visual properties are encoded, like luminosity, color, edge orientation, and potentially higher-level properties like optic flow patterns. The view-matching process compares the snapshot with the current retinal input *point by point*. It could start by considering, say, the left uppermost point in the snapshot and calculate a local discrepancy value based on differences with the properties recorded for the corresponding point in the retinal stimulation. The view-matching process then repeats this for every pair of corresponding points until it can sum up all the individual points' discrepancy value to provide an overall discrepancy assessment for that specific viewpoint. An overall discrepancy assessment is then calculated in the same way for every other visual perspective in a sample of perspectives that covers about 360 degrees. The subject then makes one step in the direction that produces the lowest overall discrepancy assessment. And the cycle starts again.

Proponents of the movement-based model have performed many virtual and robotic simulations to provide a proof of concept. These simulations have shown that, on a variety of computational schemes for encoding visual information in snapshots and performing point-by-point comparison, agents would end up searching for the goal at one of the two diagonally adequate corners if they started near the middle of a rectangular enclosure, with no significant preference for one corner over the other.

Consider Figure 4A. It provides an overhead representation of the trajectories of an animal in a simulation of one of Cheng's (1986) experiments by VM theorists (Stürzl et al., 2008). This simulation involves a moving camera that aims to capture the perspective of an animal exploring the experimental enclosure. The simulated animal follows a discrepancy-minimizing strategy as it is trying to match the picture in Figure 4B with its current retinal projection. It will pursue different paths depending on its initial location, as evidenced by the fact that following the arrows on the map will sometimes lead to different end points when you start in different locations. The animal also often changes direction, and the enclosure contains different areas where it can get stuck. These are called *catchment areas*. They are the areas which lead simulated agents to that same end point if they follow a discrepancy-minimizing strategy. Figure 4A depicts four catchment areas, whereas Figure 4C shows image discrepancy as a function of location in that simulation. As evidence of the success of their simulations of Cheng's (1986) experiment, Stürzl et al. note that the two deepest local discrepancy minima correspond to the subjects' most frequent search locations in those experiments (e.g., corners b and d in Figure 4C) and that these two local minima give rise to wide-ranging catchment areas that encompass or come near the center of the enclosure.

A basic commitment of all VM theories is that the view-matching process does not operate on information about geometric cues as such. Rather, geometric cues influence subjects' search behavior in reorientation tasks because the view-matching process is particularly sensitive to retinal information which happens to track geometric cues in those experiments. For example, some VM theorists claim that the brightness or color contrast between light coming from the walls of the enclosure and light arriving from the ceiling of the experimental room produces two large regions with highly distinct brightness or color levels in the retinal stimulation. Of course, the location and shape of these two regions differ significantly when the animal looks toward the diagonally *adequate* corners as opposed to the diagonally *inadequate* corners. And this in turn strongly affects the view-matching process. On the other hand, subjects often ignore featural cues in reorientation tasks because the featural cues employed in those experiments yield circumscribed regions with distinct brightness or color levels in the retinal stimulation. According to VM theorists' analysis, the regions are too small, or not distinct enough, to sway the view-matching process.

Notice the stark difference with the explanation of reorientation results provided by GM theories. The latter appeal to representations that capture information that often cannot be perceived in one look. On many GM accounts, geometric representations encode metric properties of the overhead projection of all the three-dimensional extended surfaces in the environment. Constructing such a representation requires many complex computations, such as estimating the position and orientation of various extended surfaces in a common reference frame. It unfolds over time as the subject stitches together information from multiple perspectives. Local properties of the retinal stimulation, such as color and edge orientation at specific points in the visual field, play a crucial role in this construction process. But they do not directly participate in the subsequent alignment process performed by the geometric module. And the geometric module does not follow anything close to a point-by-point comparison procedure either. According to Gallistel, for example, it rather proceeds by the translation and rotation of shape parameters, parameters that summarize the global shape of three-dimensional extended surfaces as seen from an overhead perspective (see Chapter 2:4.2).

3. Behavioral Findings from Reorientation Experiments

In this section and the next one, I will discuss some of the most relevant empirical findings in the debate between the two frameworks. Each of the findings considered here is widely seen as offering important support for one framework over the other. But I will argue, with respect to each finding, that proponents of the other framework can provide a systematic and similarly plausible explanation for that finding. There is no doubt that the empirical results in question each constrain theory-building in important ways, but it is my contention that they do not clearly favor one framework over the other. That each framework can systematically explain the other framework's prized findings in this way will suggest that the debate has reached something of an impasse.

I will consider in this section three *behavioral findings* perceived as most directly relevant to the debate. In the next section, I will turn to the *neuroscientific findings* which many see as most germane to the debate.

3.1. Reliance on Subtle Geometric Cues

Sang Ah Lee and colleagues have developed an important class of reorientation experiments in order to compare GM and VM theories (Lee & Spelke, 2010a, 2011; Lee, Spelke, & Vallortigara, 2012; Lee, Winkler-Rhoades, & Spelke, 2012). They employed one of two distinct types of enclosures in their reorientation tasks. The first type consists of enclosures with subtle geometric cues, such as a rectangular frame made of 2-cm-high white rods placed on a light gray floor. The second type comprises enclosures with comparably salient featural cues, such as contrasting floor coloring with the same shape as the frame just described. Just as GM accounts predict, and in direct opposition to what extant VM accounts predict, their subjects relied on the subtle geometric cues in the first category of enclosures in an attempt to find the goal, and they ignored the featural cues in the second category of enclosures. For that reason, many see Lee et al.'s results as offering a powerful blow in favor of the GM framework against the VM framework.

I agree that such results fit well with GM accounts while severely undermining extant VM accounts. But I think that VM theorists have the resources to explain such results within the VM framework if they go beyond extant models (see also Pecchia & Vallortigara, 2012; Wystrach & Graham 2012). One way to do so would be to claim that the view-matching process gives more weight to the two following types of visual properties in its operations: (i) shading that arises due to the curvature of three-dimensional extended surfaces present in the visual field; (ii) visual edges that arise due to the fact that two extended surfaces meet to form a physical edge or due to the fact that a given extended surface stands in front of another surface or object. More specifically, it could be the case that, when the view-matching process compares the retrieved snapshot and the retinal stimulation point by point in order to yield an overall discrepancy assessment, it multiplies the local discrepancy value at a given point by a large positive constant whenever the properties encoded at that point pertain to category (i) or (ii). Such a scheme would allow reorientation subjects to rely on subtle perturbations in three-dimensional extended surfaces to guide their behavior without representing the surfaces' geometric structure as such.

Consider the experiment described above. The 2-cm-high rods produce visual edges because the rods stand in front of a tiny portion of the floor from various perspectives. So, these visual edges fall into category (ii). And, on the account under consideration, it means that the local discrepancy value at points in the visual field that correspond to those edges will get multiplied by a large positive constant. Hence, the view-matching process should produce a high overall discrepancy assessment when facing diagonally inadequate locations in enclosures of the first type because the edges in the current retinal stimulation will not match those of the retrieved snapshot from that perspective. This should then help subjects stay away from those locations and move toward the diagonally adequate locations.

Now, the contrasting floor coloring used in the second case produces visual edges in roughly the same location in the visual field as the rods do when the subjects face corresponding locations in the enclosure with rods. But since those visual edges do *not* arise due to the fact that some three-dimensional surface stands in front of another or because two surfaces meet to form a physical edge — the floor is completely flat after all —, the view-matching process that I have proposed on behalf of VM theorists should not give extra weight to the local discrepancy value pertaining to those edges. The process should thus produce a significantly lower overall discrepancy assessment when facing the diagonally inadequate locations than in enclosures of the first type, making it less likely that the visual edges could sway the subject toward the diagonally adequate locations.

Of course, adopting such an account would force VM theorists to admit the existence of additional perceptual mechanisms: mechanisms for singling out three-dimensional extended surfaces in the visual field; mechanisms that identify shading caused by the curvature of those surfaces as opposed to changes in light intensity caused by other factors; mechanisms for distinguishing visual edges that pertain to category (ii), as opposed to visual edges that merely arise due to two-dimensional patterns on flat surfaces. But, importantly, the existence of such mechanisms is compatible with the VM theorists' claim that subjects do *not* extract information about the geometry of three-dimensional extended surfaces as such to guide their search behavior. The mechanisms do not require the extraction of this information to perform their attributed function, let alone of shape parameters of the entire environment as seen from an overhead perspective. For example, mechanisms of the first class might rely on visual edges, stereopsis, occlusion patterns and optic flow to identify the location and visual boundaries of the three-dimensional elements currently in the visual field and then on the elements' visual size to identify those that count as three-dimensional extended surfaces as opposed to three-dimensional isolated objects.

Finally, note that VM theorists can co-opt a part of GM theorists' evolutionary rationale to motivate an account like this. GM theorists have often defended their models by appeal to an evolutionary argument (Gallistel, 1990) to the effect that selection pressures would favor the emergence of the geometric module because the global geometry of three-dimensional extended surfaces are highly stable and reliable for estimating one's heading and location. VM theorists could thus maintain that, although it would have been very useful for any species to possess a geometric module, this required too big a jump in the complexity of the mechanisms used for spatial navigation. So, evolutionary pressures have led instead to the acquisition of the next best thing: a view-matching process that gives more weight to aspects of the retinal stimulation that reliably track the curvature of, and visual contours afforded by, three-dimensional extended surfaces. In sum, by adopting the new proposal made here about the inner workings of the view-matching process at the heart of their accounts, VM theorists can provide a similarly plausible explanation of the results presented above. Therefore, neither framework receives more empirical support from them.

3.2. Beacon Homing

It has been known since the beginning of the reorientation literature that repeated exposures to an unchanging enclosure with a stable goal location often induce subjects to start using featural cues to return to the goal location and therefore leads to a decrease in the number of rotational errors over time (Cheng, 1986). In order to account for this, many GM theorists posit a beacon-homing process (e.g., Gallistel, 1990; Lee, Shusterman, & Spelke, 2006). As noted in Chapter 1:3, *beacon homing* consists in moving toward a featural cue in a straight line until the subject has retrieved or reached a desired goal previously experienced as being inside or near the cue. It is a simple navigation strategy which requires the use of a representation of the goal and a representation of the cue, between which a link has been formed because of prior experiences.

On the models defended by these GM theorists, the beacon-homing process is independent from the geometric module and even competes with it to take control in guiding reorientation subjects' search behavior. It takes control when geometric cues are particularly unhelpful, or when the link between the representation of the cue and the representation of the goal is strong enough. The link itself can be strengthened by repeatedly obtaining the desired goal when in proximity to the cue. VM theorists, on the other hand, have rejected the existence of a beacon-homing process on the ground that view-matching processes are indirectly sensitive to both geometric cues and featural cues, and therefore that no additional process is necessary to account for the occasional reliance on featural cues.

However, over time, various reorientation experiments have come to vindicate GM theorists' commitment to a beacon-homing process. The most striking such experiments are ones with geometrically uninformative environments where subjects look with similarly high frequency at two corners equally close to a salient featural cue experienced in the past as being in the vicinity of the goal *even though those two corners offer completely different visual perspectives* (e.g., Lee, Ferrari, Vallortigara, & Sovrano, 2015; Lee, Tucci, & Vallortigara, 2017; Lee, Vallortigara, Ruga, & Sovrano, 2012). Were subjects guided by other types of processes, such a view-matching process, they would not search at both corners with a high frequency, let alone at a similar frequency. Beacon homing, on the other hand, leads subjects to search for a desired goal based only on its observed proximity to an associated cue, and thus uniquely accounts for the patterns obtained.

Because GM theorists have been defending the existence of a beacon-homing process for years whereas VM theorists have generally rejected it, many reorientation researchers have taken the results of such experiments as supporting the GM framework over the VM framework. However, I do not think that this is right conclusion to draw. I agree that these results strongly support the existence of a beacon-homing process, but I think that the issue of whether reorientation subjects possess a beacon-homing process is largely irrelevant in the debate between the two frameworks. I submit that the best way of capturing the core issues in the debate turns on the two following claims: (1) reorientation subjects store information about the global geometry of the three-dimensional extended surfaces provided by the experimental enclosure as such; (2) there is a modular mechanism in higher cognition that operates on that information alone, often guides search behavior, and is the main cause of rotational errors in typical reorientation tasks. The GM framework's central commitment is that those two claims are true. The VM framework's central commitment is that those two claims are false and that rotational errors are generally caused by a view-matching process instead.

From that perspective, nothing prevents VM theorists from simply admitting the existence of an auxiliary beacon-homing process that guides behavior in some specific contexts. By hypothesis, the beacon-homing process operates on featural information, not on geometric information. So, admitting its existence does not make either claim more or less plausible. It thus also follows that VM theorists can provide an equally good explanation of the results described above and that neither framework receives more support from such experiments alone.

3.3. Reliance on Complex Sets of Featural Cues

Many reorientation experiments are seen as providing support for VM theories over GM theories. The most relevant are arguably reorientation tasks in which subjects rely on a complex set or gradient of featural cues at some remove from the goal in order to return to it (e.g., Dittmar et al., 2014; Gillner, Weiß, & Mallot, 2008; Pecchia, Gagliardo, & Vallortigara, 2011; Pecchia & Vallortigara, 2010). For example, in Gillner et al.'s experiment, human subjects used a large-scale color gradient on the wall of a circular room in order to return to a location they had reached at the end of a previous foray in the room. Such experiments are seen as supporting VM framework over GM because, in them, subjects rely on some general configuration of cues best captured by snapshots (as evidenced, in the case of the Dittmar et al. and Gillner et al. experiments, by quantitative modeling of search patterns) and because neither the geometric module or a beacon-homing process are suited to account for reliance on more than one featural cue at a given time, let alone a gradient of color with no clear demarcating boundaries.

I agree that such results fit well with the VM framework and severely undermine current GM models. But, fortunately, nothing prevents GM theorists from admitting the existence an additional process, a view-matching process, in order to account for such results. The claims identified as central to GM framework in Section 3.2 do not require denying the existence of view-matching processes. So, much as proponents of both frameworks can posit an auxiliary beacon-homing process that guide subjects toward individual featural cues when they are perceived as reliably located close to the goal, GM theorists may posit an auxiliary view-matching process that helps subjects make use of complex sets of featural cues at some remove from the goal when matching the visual perspective they offer to the current retinal stimulation reliably leads back to that goal. GM theorists could motivate such a move by pointing out that two distinct types of backup processes with different operating principles are better than just one. Hence, GM theorists can provide an equally satisfying explanation of the experiments described above by acknowledging the existence of a backup view-matching process. Therefore, these experiments alone do not favor one framework over the other.

3.4. Where Do the Behavioral Findings Leave Us?

We have seen above that, though behavioral results of different types spell trouble for extant GM and VM models, no specific result offers unequivocal evidence for either of the two general frameworks all on its own. Proponents of both frameworks can adopt systematic explanatory strategies to deal with each relevant result. VM theorists can account for the reliance on subtle geometric cues by holding that reorientation subjects' view-matching process gives extra weight to relevant aspects of the visual appearance of three-dimensional extended surfaces. VM theorists can also account for the strong evidence in favor of beacon homing by simply acknowledging the existence of a beacon-homing process. GM theorists, on the other hand, can explain the reliance on complex set of featural cues by positing an auxiliary view-matching process that takes control of subjects' search behavior in specific conditions.

Note also that, given the type and the strength of the evidence involved in each case, these moves do not seem optional. This means that, in order to account for all the findings discussed above, GM theorists must posit three processes: the process performed by the geometric module, the beacon-homing process, and a backup view-matching process. VM theorists, for their part, must posit two processes: a view-matching process that guides behavior in a large variety of contexts and the beacon-homing process.

From this, it follows that the VM framework may hold a slight advantage over the GM framework from the sum of the behavioral evidence collected to date. First, given that GM theorists must posit a larger number of processes, this likely means that they will also have to adopt more complex rules governing the competition between the processes they posit than VM theorists. Second, the geometric module is hypothesized to perform far more complex operations than any other process VM theorists must accept. Both considerations thus suggest that VM theorists can present overall simpler theories in order to account for all the behavioral results obtained to date, and that they may hold a slight edge for that reason. That being said, the sort of modest simplicity advantage involved here could not provide decisive support for one framework over the other (Sober, 2015). So, we do not yet have behavioral results of a type that could prise the two frameworks apart.[10]

4. Neuroscientific Findings

In the last ten years or so, various neuroscientific findings have been proposed as bases for arguments in favor of the GM framework (e.g., Cheng & Newcombe, 2005; Jeffery, 2010; Lee & Spelke, 2010b; Gallistel & Matzel, 2013). I will argue here that four neuroscientific findings often considered as the most promising in this regard do not clearly favor either framework.

---

[10] I have omitted some important behavioral findings from the discussion above because they seem to raise equally complex issues for both frameworks. For example, Waller, Loomis, Golledge and Beall (2000), along with Mallot and Lancier (2018), give strong evidence that human subjects can extract distance information between a goal and multiple isolated objects in order to return to it. This raises serious issues for VM theorists given their commitment against the extraction of geometric information from the environment. And it raises serious questions for GM theorists given that nearly all current GM accounts reject the idea that subjects store geometric information about *isolated objects*, which those accounts count as featural cues.

4.1. Brain Regions Involved in Processing Information about Three-Dimensional Extended Surfaces

Functional imaging experiments show that three regions in the human brain become much more active when subjects see images containing large-scale extended surfaces, like images of buildings or landscapes, than when they view pictures of specific isolated objects without background (Epstein, Patai, Julian, & Spiers, 2017; Epstein & Vass, 2014). They are: the parahippocampal place area, the occipital place area and the retrosplenial cortex. The parahippocampal place area even shows a noticeable increase of activity when participants see images that include vertical extended surfaces which barely protrude above the ground in comparison to otherwise similar pictures without vertical surfaces (Ferrara & Park, 2016).

Following the discovery of the parahippocampal place area as a region sensitive to extended surfaces, some authors tentatively suggested that the region might comprise the geometric module in humans (e.g., Cheng & Newcombe, 2005). But the view that this region, or the two other regions for that matter, constitute the geometric module has now been largely rejected (Cheng et al., 2013). Arguably the most problematic issue with this view is that by hypothesis the geometric module ignores all featural information, and yet all three regions have since been shown to display high activation when subjects see isolated objects in some specific contexts, such as when those objects have a large visual size or when they have been experienced as navigationally relevant in the past (Janzen & van Turennout, 2004; Troiani, Stigliani, Smith, & Epstein, 2014).

If there is strong evidence against the view that one of these three regions constitute the geometric module, then why do reorientation researchers often take the functional imaging work cited above as relevant to assessing the two frameworks? Many seem to endorse the following argument. The results from this wide-ranging literature strongly suggests that that the human visual system treats three-dimensional extended surfaces as a special class of objects in an attempt to extract information about them specifically (see, e.g., Park, Brady, Greene, & Oliva, 2011). Indeed, various control conditions from these experiments show that three-dimensional extended surfaces activate the regions precisely in virtue of being three-dimensional extended surfaces, and not because of some further fact about them, like their visual size or experienced navigational relevance. Yet, all GM models predict the existence of mechanisms dedicated to localizing three-dimensional extended surfaces in the visual field and extracting information about them specifically. VM theorists, on the other hand, seem at a loss to explain why three-dimensional extended surfaces receive special treatment in this way, especially given their frequent claims that unanalyzed or relatively low-level visual information contained within snapshots is all that is needed or used for navigation purposes.

This is a powerful argument for GM models over extant VM models. But I believe that VM theorists have a plausible answer to it if they go beyond extant models. The discussion in Section 3.1 implies that VM theorists can and should admit that the view-matching process at the heart of their accounts gives more weight to specific *visual* properties pertaining to three-dimensional extended surfaces, such as the shading caused by the surfaces' curvature, or the visual edges caused by these surfaces' physical edges or position with respect to other objects. This move is compatible with the VM framework because the framework only rejects the idea that subjects record *geometric* information about such surfaces, not specific *visual* information related to them. In line with this, VM theorists could say that part of the three regions' function is to extract or operate on specific visual information about three-dimensional extended surfaces in order to feed or implement various processes, such as the view-matching process guiding reorientation subjects' search patterns. This sort of view predicts that images containing three-dimensional extended surfaces should produce high activation in the three regions as opposed to control images, as observed, because the mere presence of such surfaces in the images coerces the regions into performing or receiving an analysis of the surfaces' visual properties. Hence, VM theorists can provide similarly plausible explanations of experiments that strongly implicate the three regions in the processing of information about three-dimensional extended surfaces, and so their results do not clearly favor one framework over the other.

4.2. Global Remapping of Place Cells in Enclosures with Different Shapes

Place cells are neurons that fire when the subject represents itself as being in a specific location in an environment. They are found in the mammalian hippocampus in large numbers (Ekstrom et al., 2003; O'Keefe & Nadel, 1978). A given place cell might fire, for example, when the subject is standing around 25 cm north of a specific wall and 60 cm east of another particular wall in an experimental enclosure regardless of the subject's heading. In fact, the majority of place cells fire at high rate around only one location in a given enclosed environment. That location is the cell's *firing field*.

Global remapping is an important phenomenon involving place cells that many have claimed supports the existence of geometric representations (e.g., Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002). In global remapping, the firing fields of active place cells change drastically with respect to each other in two contexts. When this happens, two place cells whose firing fields used to be about 45 cm from each other might now fire around the same location. Or one of them may not have a clear firing field anymore. Consider, for example, an experiment designed by Colgin et al. (2010). Colgin et al. trained rats to forage for food in two distinctively shaped enclosures, one square-shaped and the other circle-shaped. Both enclosures were made of the same material, and they were connected by an alleyway that allowed subjects to go from one to the other at their leisure. After some time, Colgin et al. removed the alleyway and started exposing the subjects to each enclosure in alternation. Measurements of place cells' activity over multiple trials showed that their firing fields changed drastically with respect to each other *across enclosures*, but stayed remarkably stable *within the same enclosure*. The subject's place cells thus displayed global remapping between the two enclosures. In fact, further experiments suggest that exposure to distinctively shaped enclosures reliably produces global remapping across the enclosures (Lever et al., 2002; Mayer, Bhushan, Vallortigara, & Lee, 2018; Muller & Kubie, 1987; Wills, et al., 2005), though chances increase significantly if subjects first have the opportunity to walk by themselves between them (Colgin et al., 2010).

Why have such results been taken support the existence of geometric representations? The fact that place cells maintain stable firing-fields with respect to each other when subjects return to similarly shaped enclosures cries out for explanation. A large set of place cells could not systematically adopt a similar configuration by random chance on multiple visits over many days. Arguably, the best explanation for this is that subjects keep long-term representations of the global geometry of the environments they visit, and that those representations also play an important role in driving or modulating place-cell activity when subjects switch between environments (cf. Jeffery, Anderson, Hayman, & Chakraborty, 2004).

This reasoning has in turn been taken to support for GM theories against VM theories because the latter reject the existence of geometric representations. But, much as we have seen in the previous subsections, VM theorists can seek alternative explanations for such results. In particular, they can insist that, without extracting the enclosures' global geometric structure as such, the visual system feeds information to place cells in a way that allows them to indirectly track the shape of the enclosures visited. After all, the variation in the shape of the enclosures create important differences in the retinal stimulation, differences on which the subjects' visual system could easily pick up to drive place cells' activity. Moreover, the idea that view-based navigation processes can track geometric cues without operating on information about these cues as such already plays a central role in VM theorists' explanation of reorientation performance. The response considered here thus applies VM theorists' basic explanatory strategy for reorientation tasks to the phenomenon of global remapping across distinctively shaped environments. So, though GM theorists may have a more elegant explanation of the results discussed above, VM theorists have the resources to explain the phenomenon of global remapping in place cells across distinctively shaped enclosures. It follows that this phenomenon taken alone does not provide unequivocal support for either framework.

4.3. The Discovery of Boundary Vector Cells

Boundary vector cells (BVCs) are neurons that fire when there is a boundary, such as a vertical extended surface, at a certain fixed distance in a specific direction from the subject. For example, such a cell might fire when there is a wall standing about 15 centimeters away from the subject in a northward direction. BVCs have been found in the rat subiculum (Barry et al., 2006; Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009).

Many reorientation researchers think that BVCs represent the perfect candidate for the neural substrate of the geometric module (e.g., Cheng & Newcombe, 2005). It is easy to see why: BVCs directly track geometric information, namely the *distance* and direction of *three-dimensional extended surfaces* from the subject. It has also been discovered that BVCs display similar activity in the presence of a sudden drop at their preferred firing distance and direction as with respect to vertical extended surfaces (Stewart, Jeewajee, Wills, Burgess, & Lever, 2014). And this has motivated extending the original characterization of boundaries to include sudden drops, such as cliffs at the edge of a table or a land promontory. Moreover, the color and texture of a vertical surface does not significantly affect BVCs' firing rate or their preferred distance and direction (Lever et al., 2009).

One possible reason why researchers believe that the discovery of BVCs provides strong backing to the GM framework is that, because BVCs constitute the perfect candidate for the neural substrate of the geometric module, their discovery ipso facto bolsters the GM framework. Unfortunately, however, this reason does not hold up to scrutiny. While I agree that BVCs are a good candidate for constituting part of the neural substrate of the geometric module *once you accept its existence*, subjects must possess far more than just BVCs for the GM framework to be vindicated. A single BVCs' activity indicates the current presence of a boundary at a very specific point in space. The cell's activity thus constitutes, at best, a transient and very local representation of the environment — a far cry from the long-term representations of the global geometry of three-dimensional surfaces at the heart of the GM framework. For there to be a BVC-based mechanism that counts as the geometric module, subjects would need, at the very least, neural machinery that can pool together the activity of multiple BVCs whose preferred firing distance and direction range over a wide selection of values. And there is no direct evidence for such machinery at the moment. The only source of support for this view comes from the so-called *BVC model of place-cell firing* (Barry et al., 2006), which posits something like the requisite neural organization. But given the existence of alternative theoretical accounts of place-cell firing (Krupic, Bauza, Burton, Lever, & O'Keefe, 2014; Widloski & Fiete, 2014), we cannot reach any justified conclusion at the moment about whether BVCs' activity gets pooled in the relevant way.[11]

Another possible reason why researchers see the discovery of BVCs as offering empirical backing to the GM framework is that their discovery shows the need to posit geometric representations of the relevant type — i.e., long-term representations of the global geometry of three-dimensional extended surfaces — *upstream* from BVCs

---

[11] In fact, even assuming it wins the day over other accounts, how the BVC model gets filled out matters for the viability of the idea of a BVC-based mechanism that counts as the geometric module. For example, this idea requires that BVCs not interact too closely with cells that track featural cues. Yet, perhaps the best developed version of the BVC model (Byrne, Becker & Burgess, 2007) posits direct connections from BVCs to cells that track featural cues, and vice versa.

themselves. On this view, we cannot make sense of BVCs' firing patterns without assuming that their activity is modulated by prior geometric representations. Unfortunately, this reason does not hold up to scrutiny either. Computational modeling suggests that optic flow could allow subjects to continually estimate their distance from a boundary in a way that explains BVCs' firing patterns without relying on any such prior representation (Raudies & Hasselmo, 2012). In fact, an account of BVCs' firing patterns based on optic flow dovetails with one highly plausible functional explanation as to why animals have BVCs. This explanation holds that the role of BVCs consists in providing online information about the position of boundaries currently surrounding the subject in order to help it to avoid running straight into or over a boundary when moving around. Such information could also guide the agent in choosing paths that steer clear of potentially untraversable boundaries when planning a path ahead of time.[12] Hence, given the availability to VM theorists of plausible accounts of BVCs' firing patterns and of their role in navigation that eschew the postulation of geometric representations of the relevant type, the discovery of BVCs does not clearly favor the GM framework over VM framework.

4.4. Place Cells and Grid Cells Maintain Stable Firing Fields in the Dark

Neurophysiologists have discovered many types of spatially-tuned cells in the mammalian hippocampal formation besides place cells and boundary vector cells. A particularly important class of spatially-tuned cells is the class of *grid cells*. They are cells that fire in a hexagonal pattern in any environment. Their firing fields are similar to those of place cells, but grid cells usually have numerous fields in even relatively small environments. Grid cells have been found in large numbers in the entorhinal cortex of mammals (Doeller, Barry, & Burgess, 2010; Hafting, Fyhn, Molden, Moser, & Moser, 2005), a region which provides an important source of input to the hippocampus, where place cells are housed.

Now, an important empirical finding that has generally been seen as favoring the GM framework over the VM framework is that both place and grid cells continue firing when subjects are in complete darkness. It is well-known that individual place cells and

---

[12] See Solstad, Boccara, Kropff, Moser and Moser (2008) and Stewart et al. (2014) for related suggestions regarding the role of boundary-related cells in navigation.

grid cells studied *in normal lighting conditions* usually maintain stable firing fields for long periods of time in a given enclosure, often over multiple visits. But what various experiments further suggest is that both place cells and grid cells can also maintain the *same* stable fields in darkness — provided that the beginning of a trial in the known enclosure was under normal lighting conditions (Hafting et al., 2005; Jeffery, Donnett, Burgess, & O'Keefe, 1997; Quirk, Muller, & Kubie, 1990).

How can that be possible from perspective of the VM framework, a framework which gives prominent role to *visual* information in the form of snapshots to explain various aspects of search behavior in navigation experiments? Struck by these results, some prominent GM theorists (e.g., Lee & Spelke, 2010a; Gallistel & Matzel, 2013) have taken the stability of place and grid fields in darkness to undermine the VM framework.

I agree that such results are problematic for extant VM models insofar as many such models only focus on the task of explaining search behavior in reorientation tasks *by appeal to visual input*. The movement-based model, for example, makes no mention of other mechanisms besides the view-matching system that performs the discrepancy-minimization procedure described in Section 2. And, since the discrepancy-minimization procedure operates solely on visual input, VM theorists cannot appeal to it to handle the results just discussed.

However, I think that VM theorists have the resources to explain such results if they are willing to posit additional navigation processes besides the view-matching process. One process which seems particularly useful to explain such results is *path integration*. Path integration is the process of estimating one's own distance and direction from a previously visited point in a given environment, without resorting to known landmarks or cues. It is driven in large part by idiothetic information, such as optic flow, vestibular input, motor efference copy and proprioception. Acknowledging the existence of this process does not require positing geometric representations in the sense relevant for the debate between GM and VM theorists (i.e., representations of the *global geometry of three-dimensional extended surfaces as such*). Path integration merely requires that the subject records the distance between two points in an environment, as well as the orientation of one of the two points in egocentric space. These two points may well be located far away from extended surfaces.

Though some researchers seem tempted to do so, I believe that it would be counterproductive to characterize the VM framework in a way that prevents VM theorists, as a matter of definition, from positing a process like path integration. First, there is overwhelming evidence for the ability to perform path integration in a large number of species, including all those studied in reorientation tasks (Etienne & Jeffery, 2004; Gallistel, 1990; Mittelstaedt & Mittelstaedt, 1980). So, any framework that would deny its existence would be a non-starter. Second, many prominent VM theorists (e.g., Ken Cheng, Allen Cheung, Antoine Wystrach) explicitly acknowledge and discuss the need to posit a path-integration process in other publications. Third, one important neurocomputational model of spatial navigation (Sheynikhovich et al., 2009) explicitly appeals to both snapshots and a path-integration process to explain results from the reorientation literature as well as the literature on place cells. I believe that it is important that this model be categorized as belonging to the VM framework because it categorically rejects the two following ideas: (i) reorientation subjects store representations about the global geometry of three-dimensional extended surfaces as such; (ii) the navigation systems responsible for rotational errors in search behavior in reorientation tasks are modular. As highlighted in Section 3.2, these ideas have been at the center of the debate between proponents of the GM framework and the VM framework since the beginning.

Now, if VM theorists acknowledge the existence of a path-integration process, then they may well have the resources to explain how place cells and grid cells maintain stable firing fields in the dark. In particular, once they acknowledge the existence of that process, VM theorists can reject the idea that the relationship between the view-matching process and spatially-tuned cells is a unidimensional one, such that the sudden loss of visual input results in an instantaneous decrease in the cells' activity. More specifically, they can claim that (i) the activity of place cells and grid cells is driven by a mix of visual input and idiothetic input while the animal has visual access to its surroundings; and that (ii) their activity is driven mainly by idiothetic input while the animal is in the dark.[13] In this vein, Sheynikhovich et al. (2009) explicitly posit two main sources of information driving place-cell and grid-cell activity: idiothetic input and view-based information conveyed through the matching of snapshots with the

---

[13] This is now a standard view at the heart of most neurophysiological models of place-cell and grid-cell firing (see, e.g., Widloski & Fiete, 2014).

current retinal input. Sheynikhovich et al. do not explicitly talk about what happens when the subject suddenly finds itself in the dark, but there is no reason why they could not also endorse (ii).

At this point, someone might retort that what causes trouble for VM theorists is not just the fact that place cells and grid cells maintain stable firing fields in the dark. What causes trouble for VM theorists is the fact that the firing fields remain stable in the dark *over long periods of time by relying on tactile information acquired when the subject touches the walls of the enclosure with its paws or whiskers*. Indeed, various empirical results and theoretical modeling suggest that subjects do not and could not maintain stable firing fields in the dark for more than 20-30 seconds without getting tactile feedback from the walls (Cheung, Ball, Milford, Wyeth, & Wiles, 2012; Zhang, Schonfeld, Wiskott, & Manahan-Vaughan, 2014). The path-integration process accumulates error too quickly when left to its own devices. And yet the firing fields of both place cells and grid cells remain stable over long minutes in the dark (Hafting et al., 2005; Save, Cressant, Thinus-Blanc, & Poucet, 1998; Save, Nerad, & Poucet, 2000). This strongly suggests that tactile feedback plays a role in maintaining the place and grid fields' stability. And the challenge here is for VM theorists to explain how that can be the case.

Consider first what GM theorists would say about this. GM theorists, because they are committed to geometric representations, can provide a natural explanation of how that can be the case. They can say that tactile feedback helps maintaining stable firing fields because it interacts with the geometric representation of the current environment. The geometric representation allows the subject to continually update its estimate of how far it is from a given wall at any given time. When the subject touches a wall with its paws or whiskers which is closer or further away than expected according to the geometric representation, tactile information thus helps the subject to correct for the drift in the place-cell and grid-cell signals.

It might seem that VM theorists are in deep trouble here precisely because they cannot avail themselves of a similar explanation. But, in fact, that is not right. They can provide an alternative and well-motivated explanation of long-term field stability through tactile feedback. Here is how it might go. When a subject explores a new environment, connections from border cells to grid cells that fire in the same region of space are strengthened through a form of Hebbian learning.[14] Then, when the animal loses visual access to its surroundings, the strengthened connections between a given border cell and its associated grid cells allow to correct for drift in the grid-cell signal. More precisely, when a border cell starts firing because of tactile input, it sends a powerful signal to its paired grid cells in order to increase the probability that those cells will fire themselves when the subject is close to the boundary. This coerces each associated grid cell to fire when the subject comes back to some of the cell's original firing fields along the boundary. Once grid-cell activity is corrected, grid cells themselves could then help place cells fall back into place (given the direct connections from grid cells in the entorhinal cortex to place cells in the hippocampus).[15]

On that kind of view, animals do not possess a full-fledged unified geometric representation of the environment because they cannot retrieve, or compute over, metric relations among extended surfaces as such. No single mechanism can operate on all these associative links as though they are one single representation. Rather, a myriad of independent associations between grid cells and border cells allow for error correction in the dark. The whole process unfolds in a piecemeal fashion, with a limited number of independent border cells exerting pressure on grid cells to fall back into place at any given time. Hence, if VM theorists adopt this sort of view, they could provide a well-motivated account of the results cited above showing that place cells and grid cells maintain stable firing fields in the dark for long periods of time by relying on tactile input — and they could do so without implicitly committing themselves to the existence of representations of the global geometry of three-dimensional extended surfaces as such. This means, in turn, that the results discussed in this section do not clearly favor the GM framework over the VM framework.

---

[14] Border cells are cells that fire when the animal is near a border in a certain direction. They have been recently discovered in the rat entorhinal cortex (Solstad, Boccara, Kropff, Moser, & Moser, 2008), in the same region where grid cells are housed. Various considerations suggest that border cells may be driven by simple optic-flow input when the subject has visual access to its surroundings (Raudies & Hasselmo, 2012), along with the head-direction signal. Tactile feedback, along with the head-direction signal, may also drive their firing patterns in the dark. Hence, there is no direct need to posit an underlying geometric representation to account for the firing patterns of border cells in the kind of experiments discussed in this section.

[15] A few empirical and theoretical papers suggest and provide at least partial support for various aspects of this kind explanation of grid-cell stability through tactile feedback (e.g., Cheung et al., 2012; Hardcastle, Ganguli, & Giocomo, 2015; Krupic, Bauza, Burton, Barry, & O'Keefe, 2015; Krupic et al., 2014; Zhang et al., 2014).

4.5. Where do the Neuroscientific Findings Leave Us?

In the foregoing discussion, I have pushed back against claims that specific neuroscientific findings provide unequivocal support for the GM framework over the VM framework. First, I argued that VM theorists can account for the involvement of three brain regions in the processing of information specifically about three-dimensional extended surfaces by claiming that the regions extract or operate on *visual* information pertaining to the surfaces as opposed to *geometric* information. Second, I pointed out that VM theorists can plausibly account for the phenomenon of global remapping across distinctively shaped enclosures by applying to it the basic explanatory strategy they deploy for reorientation tasks. More specifically, they can say that the visual system feeds information to place cells in a way that indirectly tracks the shape of the environments. Third, I argued that the discovery of BVCs does not provide clear evidence for the existence of either the geometric module or the geometric representations on which it operates, and that VM theorists can account for the role of BVCs in navigation by holding that they participate in movement control and path planning. Fourth, I noted that VM theorists can explain how place and grid cells maintain stable firing fields in the dark by holding that the activity of both types of cells is driven largely by idiothetic input in the absence of normal visual access to the surrounding environment.

Overall, GM theorists may hold a slight edge over VM theorists from the sum of the neuroscientific evidence collected to date because GM theorists can offer a more elegant explanation of the phenomenon of global remapping across distinctively shaped enclosures. But the sort of modest elegance advantage involved here could not provide decisive support for one framework over the other (Sober, 2015). Therefore, we do not yet have neuroscientific results of a type that could prise the two frameworks apart.[16]

---

[16] It might seem as though I have left out an important neuroscientific finding: In a groundbreaking study, Keinath, Julian, Epstein and Muzzio (2017) discovered that place cells fired either in their original firing field or in the rotationally equivalent location in the diagonally opposite corner in reorientation tasks with rectangular enclosures containing salient featural cues. These results have been largely viewed as striking a major blow for the GM framework over the VM framework. However, though important, I do not think that these results directly bear on the debate between the two frameworks. The results do fit GM theories well, but precisely this pattern of results was explicitly predicted by a major VM account years ago (see Sheynikhovich et al. 2009, simulation 3).

5. Conclusion

Many empirical findings have been claimed to have clear import in the debate between the GM and the VM framework. But we have seen above that neither the most relevant behavioral findings (Section 3) nor the most relevant neuroscientific findings (Section 4) provide unequivocal support for one framework over the other. The two frameworks thus receive equivalent degrees of empirical confirmation at the moment. Moreover, because each framework can offer *systematic* explanations for all the other framework's prized findings, there is no reason to suppose that obtaining further empirical data within the current experimental paradigms will change that assessment. Hence, it seems that we have reached an impasse in the debate between the two frameworks.

In what follows, I argue that a new explanatory problem offers the prospect of breaking the impasse in favor of the GM framework. The problem focuses on what I will call the *main reorientation process*: the process posited by each framework to explain why reorientation subjects frequently perform rotational errors in typical rectangular enclosures. For GM theorists, this is the process performed by the geometric module and that eventuates in an estimation of the subject's heading and location. For VM theorists, it is a view-matching process. The main reorientation process stands in contrast to the auxiliary processes that both GM and VM theorists may posit to account for various other findings (e.g., beacon homing; see Sections 3.2 and 3.3). It is the central locus of disagreement in the debate between the two frameworks.

Chapter 4

Where the View-Matching Framework Founders: The Representation Selection Problem

1. Introduction

We have seen in the previous chapter that there is something of an impasse in the current debate between GM and VM theorists because each camp can provide systematic alternative explanations for the findings perceived as favoring the other framework. The analysis of the previous chapter did not show, however, that the debate is intractable. Not by any means. It simply suggests that we should look for new kinds of considerations in order to make progress in the debate. This is what I do in this chapter. This chapter has two main aims. The first one is to develop an important new problem for researchers aiming to explain the results of reorientation experiments — *the representation selection problem*. The problem introduces a new type of explanatory consideration that theories of reorientation must address. The second aim is to argue that this problem causes serious trouble for the VM framework. Assessing how GM theories fare with respect to the representation selection problem will then be the task of next chapter. There, I will show that GM models, suitably modified, can address this new explanatory challenge.

This is the plan for the current chapter. Section 2 presents the representation selection problem itself. The next two sections explain why VM theorists do not have the resources to address it in relation to two most influential VM models on the market. Section 3 does so with respect to the movement-based model presented in the previous chapter. Section 4 does so with respect to a distinct and impressively detailed model (Sheynikhovich et al., 2009) which we have not discussed at length yet. I call it the *heading-based model*.

2. The Representation Selection Problem

The main reorientation process, as conceived by either GM and VM theorists, must make use of a representation of the current environment in order to perform its operations efficiently. However, subjects in reorientation experiments are typically exposed to multiple environments, each with important goal locations, in the minutes and hours preceding experimental trials. This generates the following problem: How do subjects reliably select the relevant representation which they use to initiate the process? This is the representation selection problem.

Some researchers have indirectly touched on issues of representation selection in relation to reorientation experiments before (Gallistel, 1990; Wang & Spelke, 2002; Jeffery, 2010; Julian et al., 2015). But no one has noticed that they have far-reaching consequences pertaining to theory choice and theory building, as I argue in the rest of this thesis.[17]

Let us start unpacking this. In nearly all reorientation experiments with non-human animals, subjects alternate between at least three environments: the experimental enclosure, a waiting cage, and a home cage where they spend a large part of their day. Moreover, subjects frequently return to one goal location in the experimental enclosure (a food source or a hidden escape platform) as well one or more goal locations in the home cage (e.g., a water source and a food source). Human subjects, for their part, visit numerous environments in the hours before they enter the experimental enclosure for a trial. Many of those environments (e.g., their bedroom, their kitchen, their backyards, the waiting room in the laboratory) contain practically relevant objects and locations, and the subjects spend far more time there than in the experimental enclosure.

All of this means that non-human and human subjects have reorientation-relevant representations from multiple distinct environments in memory while undergoing a trial, be they geometric representations or snapshots. It also entails that subjects need to select one such representation with which they will initiate the main reorientation process at the beginning of a new trial. What gives the representation selection problem its force, then, is that we have strong reasons to believe that reorientation subjects *reliably select* representations *from the current environment* (i.e., the experimental enclosure) in order to initiate the main reorientation process in a large variety of

---

[17] Though I favor a unified treatment across a very broad range of species, note that my primary focus in what follows will be on mammals because the most influential VM models have been proposed for them in the first instance (Cheung et al., 2008; Sheynikhovich et al., 2009; Stürzl et al., 2008) and because mammals have been at the center of research on the reorientation task and on spatial navigation for years.

experimental contexts.

Consider, for example, reorientation tasks where subjects look for the goal in an experimental enclosure that remains perceptually indistinguishable from trial to trial — henceforth *classical reorientation tasks*. That subjects in these tasks reliably initiate the main reorientation process with representations from the experimental enclosure follows from the fact that they search for the goal in the diagonally adequate locations much more frequently than elsewhere on any given trial in these tasks. The position and direction of practically relevant objects or locations with respect to specific cues (geometric, featural or otherwise) in previously visited environments does not bear any systematic relationship to the position and direction of the goal with respect to similar cues in the experimental enclosure. For example, a rat's water source might be in a corner with a long wall on the left and short wall on the right in its home cage, whereas the goal location in the experimental enclosure occupies a corner with a short wall on the left and a long wall on the right. Moreover, experimenters always counterbalance the position of the goal in the experimental enclosure across subjects, and sometimes within subjects across trials, in a way that further ensures the absence of a systematic relationship between cues to the goal location in the experimental enclosure and previously visited environments. So, unless subjects in these tasks reliably initiate the main reorientation process with representations from the experimental enclosure, they would not search frequently in the diagonally adequate locations there, which they do.

Now consider a different type of reorientation task where subjects alternate between two perceptually distinct and non-changing experimental enclosures from the start. Highly illuminating studies of that kind come from Julian et al. (2015; see Figure 5). In their first experiment, mice were alternatively exposed to two white rectangular enclosures with black stripes. The enclosures were of the same size and built of the same material. However, whereas one had *vertical* black stripes on one of its short walls, the other had *horizontal* black stripes on one of its short walls. Importantly, the goal corner had different geometric properties in the two enclosures: in one enclosure, the goal corner had a long wall on the left and a short wall on the right; in the other enclosure, it was the opposite. Their second experiment had a similar design, except that this time what distinguished the two enclosures was the location of the vertical black stripes. Otherwise identical vertical black stripes were displayed on one of the *short* walls in one enclosure and on one of the *long* walls in the other enclosure.

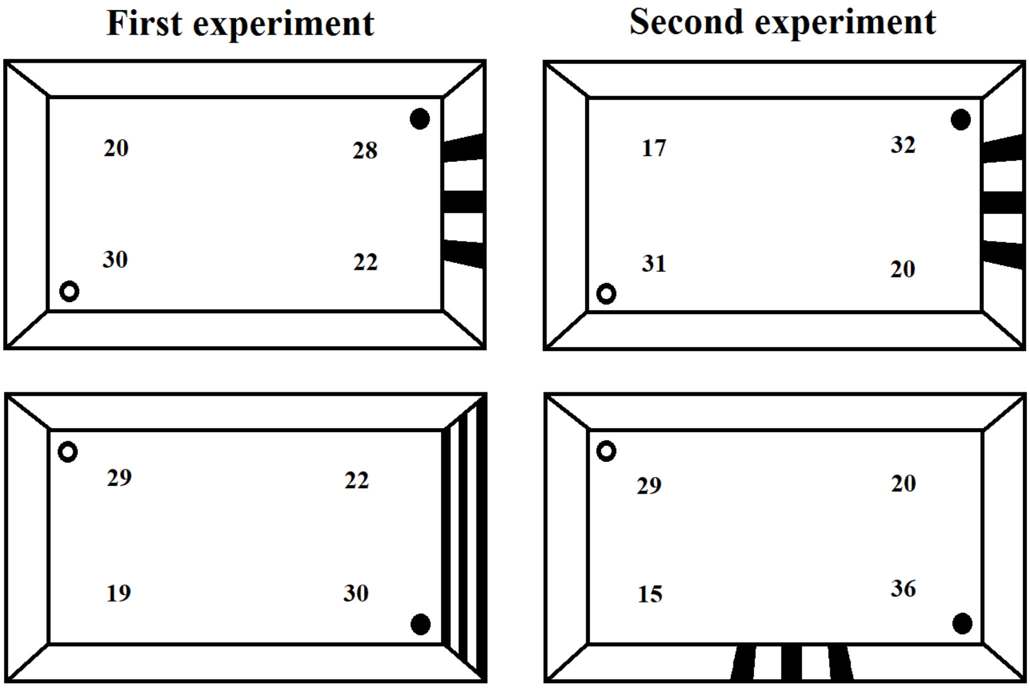

Figure 5. The setup for the two main experiments in Julian et al. (2015). Black dots indicate the location of the hidden food source in a given enclosure (the goal location). White dots represent the location of the diagonally opposite corner. Numbers in each corner indicate the percentage of trials where the subjects first dug in that corner. (Based on Figures 2 and 3 in Julian et al., 2015.)

What matters for our purposes is this: in both experiments, the mice chose the diagonally adequate corners significantly more often than the diagonally inadequate corners in the *two* enclosures. Or consider an experiment performed by Horne et al. (2010) where rats alternated between a white rectangular enclosure and a black rectangular enclosure of the same size. The goal corner had opposite left-right relations with the long and short walls in the two enclosures, and the rats spent significantly more time in the diagonally adequate corners in the two enclosures during test trials. The behavior of Julian et al.'s and Horne et al.'s subjects strongly suggest that they reliably initiated the main reorientation process with representations from the *current experimental enclosure* in these experiments.[18]

---

[18] To see why, recall first that both GM and VM theorists must hold that that an enclosure's geometric cues has a greater impact on the main reorientation process than do featural cues in order to explain the prevalence of rotational errors (Chapter 3:5). Thus, given that the goal corner had opposite left-right relations with the long walls and short walls in the two enclosures in these experiments, representations from the *other* experimental enclosure would have led the subjects to search for goal in the diagonally *inadequate* corners of the *current* experimental enclosure. Moreover, for reasons highlighted above, representations from other environments, like the home cage or the waiting cage, would not have led subjects to search frequently in the diagonally adequate corners of the current experimental enclosure. So, unless Julian et al.'s and

We will see in what follows that it is no trivial task to explain why subjects are so good at selection — to explain, in particular, how subjects reliably select a representation from the current experimental enclosure at the outset of a trial in both classical reorientation experiments and experiments involving two non-changing experimental enclosures. But, first, it may seem that there is a simple procedural solution to the representation selection problem. For instance, one of the two following procedural principles might appear do the trick:

> **Most Recent Environment** — At the beginning of a trial, select a representation from the last environment visited.

Strictly speaking, for non-human animals, the last environment visited before a trial is nearly always the waiting cage. This is a small box with opaque walls where subjects spend a few minutes before and between visits to the experimental enclosure. It is itself located some distance away from the experimental enclosure, and experimenters usually cover it with a lid to prevent light from coming in. The waiting cage is also where subjects are when the disorientation procedure happens. (In many experiments with non-human animals, the disorientation procedure consists in rotating the waiting cage while the subject is inside. In other experiments with non-human animals, the procedure consists rather in rotating the experimental enclosure while the subject rests in the unmoving waiting cage.) So, in a strict sense of the expression ‘last environment visited’, this principle wrongly predicts that non-human subjects will frequently choose a representation from the waiting cage at the outset of a trial.[19] But suppose we use the expression to mean rather *the last environment visited before the disorientation*

Horne et al.’s subjects reliably initiated the main reorientation process with representations from the current experimental enclosure, they would not have searched mostly in the diagonally adequate locations there, which they did.

[19] Exactly what counts as a *trial* varies across two important categories of reorientation tasks. *Reference memory tasks* are a type of classical reorientation task where the location of the goal in the experimental enclosure remains the same throughout the experiment for a given subject. In these experiments (and also in experiments like the Julian et al. and Horne et al. studies), a trial is just any exposure to an experimental enclosure during which the subject is incited to search for the goal. It is always preceded by the disorientation procedure. *Working memory tasks* are a type of classical reorientation task where the location of the goal changes after every sequence of two exposures to the experimental enclosure for a given subject. The first exposure is meant to give the subject time to find the new location of the goal. Once it has, the experimenter performs the disorientation procedure. Following this, the subject recovers normal perceptual access to its surroundings in the experimental enclosure, where aspects of its search behavior are recorded. For working memory tasks, a trial is each *second exposure* to the experimental enclosure in these two-exposure sequences.

procedure. Under this interpretation, the principle still fails to explain reliable selection in experiments involving two experimental enclosures. In the Julian et al. experiments, for example, any subject underwent four trials per day over multiple days, came back to its home cage at the end of a given day, and never visited the same experimental enclosure on two consecutive trials. On the proposed interpretation, the principle therefore wrongly predicts that Julian et al.'s subjects should have chosen a representation from their home cage or the other experimental enclosure at the outset of a trial. Similarly, half of Horne et al.'s subjects visited the other experimental enclosure just before the disorientation procedure on a crucial test trial. Thus, this principle wrongly predicts that subjects in this group should have chosen a representation from the other enclosure on that test trial.

> **Most Recent Goal Location** — Select the representation linked to the most recently encountered goal location.

This principle fails with respect to reorientation tasks involving two non-changing experimental enclosures for similar reasons. For example, it wrongly predicts that, for a given trial in the Julian et al. experiments, the subject should choose a representation from another environment than the current experimental enclosure on more than half of the trials because, given the structure of the experiments (see previous paragraph), the last time it looked for a goal was either in the home cage or in the other experimental enclosure. And the subject chose the diagonally adequate corners above chance in the other enclosure. The principle also wrongly predicts that, on the crucial test trial performed by Horne et al. alluded in the previous paragraph, the same half of the subjects should have selected a representation from the other experimental enclosure because the last goal that they encountered before the trial was the escape platform from the other enclosure.

We can also readily reject any selection principle which rely on the order of presentation of environments. Non-human subjects often alternate between the home cage (HC), the waiting cage (WC) and the experimental enclosure (EE) in a complex sequence determined by the experimenters. For example, in reorientation studies designed by John Pearce and his colleagues (see, e.g., Pearce et al., 2006), subjects are exposed to these three enclosures in the following order on most days: HC, WC, EE, WC, EE, WC, EE, WC, EE, WC, HC. Julian et al., for their part, changed the order of presentations of the two experimental enclosures from one day to the next in their experiments.

Could subjects simply be initiating the reorientation process with the representation that has most often led to successful goal finding in the past? Unfortunately, no, because the representation with the most successful history will undoubtedly come from some environment other than the current experimental enclosure. For non-human subjects, the environment in question will be the home cage. Non-human subjects come back to their home cage between trials or daily sessions, and they must return to specific goal locations in it in order to get water and/or food for example. They presumably go back to known goal locations at a higher rate there than in the experimental enclosure given the importance of these goals and the fact that subjects spend the vast majority of their time there on any given day. Moreover, animals usually get acquainted with the home cage for days before the beginning of a reorientation experiment. This gives them the opportunity to build up a long track record of successful use for representations pertaining to each goal there.

A specific type of ethological study called a *displacement study* also provides evidence against a large variety of procedural principles that rely on the frequency, recency, or success with which the subject has chosen a particular representation. In such studies, subjects often show a high rate of return to their nest or home range when humans intentionally move them, without visual input, to new locations that lie a few hundred meters or a few kilometers away (see Papi, 1992 for review). The high rate of return strongly suggests that subjects store reorientation-relevant representations from many environments located far from their nest, and that they reliably initiate the main reorientation process with a representation from the environment of release upon recovery of perceptual access to their surroundings. This in turn seems to rule out any selection principle that relies on frequency, recency or success because, in these studies, experimenters often release subjects outside the area where they spend the majority of their time.

At this point, some reorientation researchers might be tempted to invoke widely accepted, independent navigation processes in order to tackle the representation selection problem. For example, it is highly plausible that, when storing a new reorientation-relevant representation in memory, subjects keep a record of the path integration coordinates where the representation was registered. So, maybe subjects could use the coordinates provided by the path-integration process to select the right representation upon recovery of perceptual access? Though this idea may help explain reliable representation selection for different types of experiments, it cannot do so for reorientation tasks. Reorientation researchers perform the disorientation procedure precisely in order to knock out path integration. Moreover, following the disorientation procedure with non-human animals, experimenters often transport the subjects from the waiting cage to the experimental enclosure in their hands and in complete darkness (e.g., Cheng, 1986). This prevents subjects from being able to infer, even in theory, the distance and direction along which they are being transported from the waiting cage to the experimental enclosure.

We can also reject the idea of appealing to a beacon-homing process (as described in Chapter 3:3.2) to deal with selection issues. That process is inflexible and simply guides the subject toward a known featural cue, whereas what we are trying to explain is how the main reorientation process — the process mainly responsible for subjects' frequent rotational errors *despite the presence of salient featural cues* close to the goal location in rectangular enclosures — is reliably initiated with the right representation. A closely related suggestion might go as follows:

> **Use Associated Featural Information** — Suppose that, when storing a new reorientation-relevant representation in memory, subjects keep a record of some of the featural cues that are currently within their visual field, or that are currently detected through other sense modalities like smell or touch. Then, following recovery of perceptual access, they select the representation whose associated featural information most closely matches some aspect(s) of the current environment.

This principle may be able to account for reliable selection in some experiments (e.g., Horne et al., 2010). But, unfortunately, it cannot deal with important cases. Consider Julian et al.'s second experiment. They controlled for all the featural cues in both enclosures — except, of course, for the location of the vertical stripes. The walls and the floor were made of the same material in both enclosures, had the same odor, the same color, and both enclosures were located behind the very same curtains. The vertical stripes also had the same color and the same width in both enclosures. So, this principle wrongly predicts that the subjects should choose randomly between representations from the two experimental enclosures.

Here is another intuitively plausible procedural response to the representation selection problem. Suppose that, whenever subjects store a new reorientation-relevant representation in memory, they associate to that representation information about the type of goal just found (e.g., food, drinkable water, hidden escape platform). Then, following recovery of perceptual access, subjects select a representation associated with the most pressing goal. Unfortunately, however, this strategy is not viable either. In many classical reorientation tasks, subjects have access to food in their home cage, and the goal in the experimental enclosure is also food (e.g., Cheng, 1986; Vallortigara et al., 1990). So, this principle wrongly predicts that subjects in these experiments should randomly choose between snapshots from the home cage and the experimental enclosure. It also wrongly predicts that, in experiments like the Julian et al. and Horne et al. studies, subjects should choose randomly between the representations from the two experimental enclosures because the two enclosures hold the same type of goal.

Overall, the foregoing strongly suggests that there is no simple procedural solution to the representation selection problem. Instead, a given theory of reorientation will need to appeal to information encoded within the reorientation-relevant representations themselves and/or special information stored along with them. In the next section, we will consider the prospects for the movement-based model of the VM framework in this regard.

3. The Movement-Based Model Does Not Have the Resources to Handle the Representation Selection Problem

In the rest of this chapter, I argue that the VM framework does not have the resources to deal with the representation selection problem. More specifically, that it does not have the resources to explain how reorientation subjects reliably initiate the view-matching process with snapshots from the current experimental enclosure upon recovery of perceptual access at the beginning of a trial. I make my case with respect to the movement-based model in this section and the heading-based model in the next.

What kind of selection principle might proponents of the movement-based model appeal to in order to tackle the representation selection problem? We saw in the previous section that any plausible theory of reorientation will need to employ an information-based selection principle rather than a procedural principle. Therefore, to give a sense of the difficulty that the problem poses for the movement-based model, I will consider here a representative sample of information-based principles that they might adopt. Let us start with the following principle:

> **Use Depth Information** — *Depth information* is information about the subject's distance from some salient object or surfaces currently in the visual field. Suppose that subjects attach depth information to each new snapshot stored in memory (cf. Wystrach & Graham, 2012). Then, following recovery of perceptual access, they select the snapshot whose associated depth information most closely matches the current environment.

We can reject this principle summarily. It wrongly predicts that, in experiments like the Julian et al. and Horne et al. studies, subjects should randomly choose between snapshots from the two experimental enclosures on any trial because the two enclosures have the same size and shape. Moreover, if VM theorists assume that subjects systematically record any depth information in this way, they risk implicitly smuggling into their accounts the assumption that subjects rely on geometric representations to select the relevant snapshot.

> **Use Best Match to Retinal Stimulation** — Select the snapshot that has the highest level of matching with the retinal stimulation at the outset of a trial.

I will discuss this principle in more detail because important models of snapshot-based navigation developed outside the reorientation literature assume that initiating view-matching processes with the snapshot that has the highest level of matching can help subjects find their bearing when their spatial behavior is disrupted in various ways (e.g., Cartwright & Collett, 1987; Franz, Schölkopf, Mallot, & Bülthoff, 1998). And, in fact, this principle may be able to explain reliable selection in some classical reorientation experiments whose home cages and waiting cages differ significantly from the experimental enclosure in terms of the shape, size and the composition of their walls, because such differences increase the chances that the snapshot from the experimental enclosure would reliably yield the highest level of matching at the outset of a trial.

Unfortunately, however, this principle makes the wrong predictions about some crucial cases. Consider, in particular, Julian et al.'s second experiment. The walls of both enclosures have the same geometric properties and the enclosures contain the same featural cues, so this principle wrongly predicts that subjects should choose randomly between snapshots from the two enclosures. The fact that the vertical stripes occupy different locations in both enclosures is not enough to yield a higher level of matching for the snapshot that was originally taken in the current experimental enclosure than for the snapshot from the other experimental enclosure. Why? First, subjects begin their trials far from the goal location in both enclosures, and there is bound to be a high level of mismatch whenever an agent is far from the location where a snapshot was originally recorded. Second, VM theorists must hold that an enclosure's geometric cues have a much greater impact on the assessment of the level of matching than do circumscribed featural cues in order to explain the prevalence of rotational errors (see Chapter 3:2) and both enclosures have the same geometric cues.

This principle also fails to account for the results of an extremely important neurophysiological study: Wills et al. (2005). In the first phase of that experiment, rats were repeatedly exposed to two enclosures where they foraged for food. The first one, which Wills et al. called *the morph box*, was made of several narrow rectangular sections covered by brown tape. The angles between the sections could be changed, and Wills et al. originally gave the morph box the shape of a square. The second enclosure was made of one smooth piece of wood shaped as a circle and painted white. It turns out that the subjects' place cells displayed *global remapping* across enclosures (see Chapter 3:4.2 for a description of what global remapping involves), but kept a similar configuration when subjects returned to the same enclosure. The best explanation of these results from the perspective of the VM framework combines these three claims: (i) subjects acquired snapshots from both environments early on in the experiment; (ii) subjects selected snapshots pertaining to the current environment upon returning to it; (iii) the selected snapshots triggered the relevant place-cell configuration.[20] Now comes

---

[20] The cells' configuration could not have been determined by (a) *distinct path-integration coordinates associated with each enclosure* because both enclosures were alternatively positioned in the same location of the very same experimental room; by (b) *the recognition of distinct featural cues* because place cells adopted a configuration close to the one in the circular wooden enclosure when they were later placed in the original morph box shaped as a circle, despite the radical difference in featural cues (see main text for details); or by (c) *the*

the important part: following a few days of trials, Wills et al. discarded the circle-shaped wooden enclosure. After that, they exposed the rats in alternation to the morph box shaped as a square and the morph box shaped as a circle. Importantly, though *the circle-shaped morph box* and the now discarded circular wooden enclosure were both circular, the visual perspective from within *the circle-shaped morph box* resembled the perspective from within *the square-shaped morph box* far more than the perspective from within the now discarded circular wooden enclosure. (After all, the morph box's color and texture differed significantly from the color and texture of the wooden enclosure, and the joints between the rectangular sections in the morph box projected similar vertical edges on the retina independently of the box's shape.) Hence, the selection principle considered here predicts that the subjects visiting the *circle-shaped* morph box should select snapshots originally recorded in the *square-shaped* morph box. This, in turn, entails that place cells should adopt a configuration close to the one from the square-shaped enclosure, or at least very different to the one from the discarded circular wooden enclosure. On the contrary, the cells' firing fields in the circle-shaped morph box closely matched those in the circular wooden enclosure.

We can summarize why the Wills et al. study causes trouble for proponents of the movement-based model, and VM theorists more generally, as follows. Granted, the variations in *geometric cues* between *the circle-shaped morph box* and *the square-shaped morph box* do create important differences in the retinal stimulation, differences on which the subjects' visual system could easily pick up to determine place cells' configurations without employing underlying geometric representations (as noted in Chapter 3:4.2). However, those differences are nowhere near in scale to the differences in the retinal simulation caused by the variations in *featural cues* between the *circle-shaped morph box* and the *circular wooden enclosure.* Yet, geometric cues alone determine place cells' configuration in the circle-shaped morph box. It is hard to make sense of this without admitting that some mechanisms extract the information about geometric cues as such from the retinal stimulation, and use that information to drive place cells' firing patterns.

---

recognition of the general depth of the visual scene because the cells' configuration across a sequence of intermediary octagon-shaped enclosures was only predicted by subtle changes in the global geometry of the enclosures, changes that roughly preserved the distance from the center of the box to the closest and most salient wall.

Now, VM theorists may want to adopt a strategy similar to the one outlined in Chapter 3:3.1 as a reply. More specifically, they may want to turn to the following selection principle:

> **Use Best Match to Stimulation (with emphasis on visual properties that track geometric cues)** — Select the snapshot that has the highest level of matching with the current retinal stimulation at the outset of a trial, with extra weight given to (i) shading due to the curvature of three-dimensional extended surfaces and to (ii) visual edges that arise due to occlusion by those surfaces or due to two such surfaces meeting together to form a physical edge.

As plausible as the move proposed in Chapter 3:3.1 may be regarding the experiments discussed there, this principle does not constitute a viable response to the representation selection problem. The most basic issue with it is that, in order to deal with the two types of experiments just highlighted, VM theorists would have to make opposite assumptions about the extent of the additional weight given to the relevant visual properties that track geometric cues. On the one hand, in order to accommodate place cells' patterns across distinctively shaped boxes in the type of study performed by the Wills et al., VM theorists would need to assume that these properties receive an extremely *high* additional weight. On the other hand, if VM theorists want to have any hope of accounting for snapshot selection in studies like the Julian et al. experiments, they would need to assume that these properties receive a relatively *low* additional weight. Otherwise, the impact of geometric cues on the assessment of the level of matching will automatically swamp the impact of featural cues. And the two enclosures employed in Julian et al.'s experiments only differ in terms of the identity or the location of the featural cues that they contain.

The last principle that we will consider does not aim to describe *how* subjects go about selecting the relevant snapshot with which they initiate the main reorientation process. Rather, it is a principle about *when to trigger* the selection process itself:

> **Mismatch Threshold** — If the mismatch between the selected snapshot and the current retinal stimulation goes higher than a certain critical threshold, then initiate the selection of a new snapshot.

There is no doubt that the notion of a mismatch threshold is useful in dealing with various aspects of spatial navigation and reorientation performance from the perspective of snapshot-based approaches (Franz et al., 1998; Wystrach, 2009). However, this principle does not offer a plausible response to representation selection problem because it simply cannot make the requisite explanatory difference with the type of experimental work considered above.

Let us begin with the Wills et al. study. The issue it raises for VM theorists is to explain why snapshots from the circular wooden enclosure get selected upon entering the circle-shaped morph box even though snapshots from the square-shaped morph box contain much more similar visual information. Appealing to the notion of a mismatch threshold does nothing to alleviate this issue. If anything, it makes it worse. If snapshots from the squared-shaped morph box get discarded because they yield a high level of mismatch with the current retinal stimulation when subjects are suddenly exposed to the circle-shaped morph box, snapshots from the circular wooden enclosure should get discarded even faster in the same context. Hence, proponents of the movement-based model cannot solve the issue raised by the Wills et al.'s study by appeal to this principle.

Turn to the Julian et al. experiments now. Note first that proponents of the movement-based model must accept the following idea: the view-matching process tolerates the level of mismatch (i.e., does not trigger the selection of a new snapshot) when subjects face the rotational-error corner in the current experimental enclosure in the Julian et al. experiments. Without this commitment, they could not explain why reorientation subjects perform an equivalent number of rotational errors and correct choices in the experiments. But now suppose, for the sake of argument, that one of Julian et al.'s subjects selects the snapshot from the *other* experimental enclosure at the outset of a trial. The view-matching process initiated with this snapshot should then lead the subject to move toward one of the diagonally *inadequate* corners in the current enclosure. (After all, the diagonally *inadequate* corners in the current enclosure share the same geometric properties with the diagonally *adequate* corners within the environment from which the snapshot originates, and VM theorists hold that geometric cues have much larger impact than featural cues on the view-matching process in typical rectangular enclosures). But then the problem is this: the two diagonally inadequate corners in the current experimental enclosure produce an equivalent level of mismatch with respect to the selected snapshot as the rotational-error corner in the *other* experimental enclosure because those corners have the same geometric properties and roughly the same featural properties. Given the commitment highlighted above, this means that, as the subject enters one of the diagonally inadequate corners in the current enclosure, the level of mismatch should stay below the threshold, thus preventing the search for a new snapshot. Hence, if the animal starts with a snapshot from the other experimental enclosure, the mismatch threshold will not trigger a new selection. So, we still need an explanation as to why subjects reliably select a snapshot from the current experimental enclosure at the beginning of a trial as opposed to one from the other experimental enclosure.

We have now assessed and rejected four selection principles against the backdrop of one particular VM model, the movement-based model. Importantly, the analysis above makes only minimal presuppositions about how snapshots guide search behavior in reorientation tasks, and therefore the points made above generalize to other VM accounts. For example, variations on the movement-based model that eschew the discrepancy minimization procedure at the heart of the model (see Chapter 3:2) in favor of a slightly different approach to movement guidance (e.g., Wystrach et al., 2011) make the same problematic predictions with respect to these four selection principles for similar reasons. One VM model requires a longer discussion, however, because it departs much much more significantly from the movement-based model. How this specific model fares with respect to the representation selection problem constitute the topic of the next section.

To summarize: the representation selection problem constitutes a major problem for VM theories, and the principles I have considered are both representative and arguably the most promising candidates given the resources available in the VM framework. The principles fail to explain how subjects go about selecting snapshots from the relevant environment at the outset of a reorientation trial because in one crucial type of case the relevant environment differs from another environment in virtue of the location of featural cues with respect to the environments’ global geometry in ways not adequately captured by either depth information or by visual information (e.g., Julian et al., 2015) and because in another crucial type of case the relevant environment resembles another environment in virtue of its global geometry in ways that go against matching by visual information (e.g., Wills et al., 2005).

4. The Heading-Based Model Does Not Have the Resources to Handle the Representation Selection Problem Either

Most extant VM models resemble the movement-based model in important ways, but the impressively detailed model developed by Sheynikhovich et al. (2009) stands apart from others. I call it *the heading-based model*. In this section, I will argue that this distinctive model does not have the resources to tackle the representation selection problem any more than the movement-based model does.

There are three main differences between the movement-based model and the heading-based model that are relevant for our purposes.[21] First, the latter model holds that subjects store snapshots taken from a variety of locations and heading directions in each environment they explore, and not just from the goal locations. Registration works as follows: if, at any given moment during exploration, there are less than a fixed number of stored snapshots that achieve a certain predetermined threshold of matching with the retinal simulation, then the current snapshot outputted by the visual system is stored in memory. This involves producing a new neuron which becomes associated with the snapshot — producing what they call a *view cell*. The snapshot thus stored is also indexed to the current heading estimation. In Sheynikhovich et al.'s simulations, exploration of a new environment leads to the storage of hundreds of snapshots even in relatively small enclosures. Second, on the heading-based model, view matching does not directly affect spatial movements. The main task of the view-matching system consists in estimating the subject's heading. That estimated value then helps other navigation and motor systems to plan the subject's path toward the goal location. Third, the view-matching system does not operate on a single, specific snapshot. Rather, it

---

[21] Sheynikhovich et al. posit two distinct pathways for view-based navigation: the locale pathway and the taxon pathway. I use the expression *the heading-based model* to refer to their account of the *locale* pathway only. The *taxon* pathway, as they conceive of it, implements a less flexible view-matching process than the one posited for either the locale pathway or the movement-based model described in Chapter 3:2. Their simulations of that process can only predict the results of reorientation tasks in which the start location remains identical throughout the whole experiment, and such tasks constitutes only a subset of all reorientation experiments performed. And, just like other simulations performed by VM theorists, the simulations of that process do not incorporate visits to other environments, like the home cage, between trials. It is therefore very unlikely that Sheynikhovich et al. could appeal to the taxon pathway account for the experimental results discussed in Sections 2 and 3 of the current chapter. Note also that, on Sheynikhovich et al.'s view, only the locale pathway can influence place-cell activity. Hence, only operations within the locale pathway could account for the results of Wills et al. (2005) that played an important role in the previous section.

makes indiscriminate use of *all the stored snapshots* in order to compute the estimated heading, regardless of where they were recorded.

Here is a rough sketch of how it works. Upon recovery of normal perceptual access, the view-matching system first calculates, *for each previously recorded snapshot*, the angular head movement that *would produce* the best match between the current retinal stimulation and that snapshot. The system relies on computational principles that predict how much the current retinal input would change if the agent were to turn its head by a specific angle. The calculation does not involve or trigger any spatial behavior or actual movements of the head. The view-matching system then uses the hypothetical angular head movement just calculated to yield an estimation of the current heading *with respect to that specific snapshot* by subtracting the hypothetical angle from the snapshot's associated heading value. A toy example will help here. Suppose that a snapshot was originally stored with an associated heading of 122 degrees clockwise from a reference direction. Once the animal returns to the environment, the system calculates the angular head movement that would produce the best match with that snapshot. Let us say that it is 55 degrees clockwise in that case. Then the view-matching system makes an estimate of the subject's heading *with respect to that specific snapshot* as follows: 122 - 55 = 67 degrees clockwise from the reference direction.

This purely cognitive operation is repeated for each previously registered snapshot, and the heading value that comes up most often is then chosen as the subject's final heading estimation, a value that influences various downstream processes. For example, if 67 degrees clockwise is the value that comes up most often in our toy example, it is chosen as the final heading estimation. Following this, a specific subset of all view cells becomes active: the cells whose associated heading value approximates the final heading estimation and whose associated snapshot closely matches the current retinal input. In our toy example, this means that the view cells whose associated heading is closest to 67 degrees clockwise and whose associated snapshot closely match the current retinal input will start firing. The activation of these view cells then leads to a cascading sequence of activation that eventuates in the triggering of place cells (a well-studied category of neurons in the mammalian hippocampus which fire when the subject represent itself in a specific location, see Chapter 3:4.2). Only place cells which have frequently been activated at the same time as the currently active view cells on previous visits will get so triggered. The firing of place cells then guides spatial behavior toward the goal location through learned associations between place cells and hypothesized *action cells* whose activation lead the subjects to rotate their body to a predetermined orientation and then move forward.

We do not need to delve deeper into the model's complex account of the neural pathway from view cells to place cells and of the way in which place cells guide behavior toward a known goal location. For our purposes, it suffices to note that the model's explanatory power rests in large part on the idea that the final heading estimation produced by the view-matching system strongly constraints all downstream processes. For example, Sheynikhovich et al. rely on this idea to explain why subjects of reorientation tasks often make rotational errors. According to their analysis, many snapshots registered in rectangular environments produce a heading value which is either close to the agent's actual heading, or off by about 180 degrees — for the kind of reasons outlined in Chapter 3:2 about geometric cues having stronger impact on view matching than featural cues. This often causes the final heading estimation produced by the view-matching system to be off by about 180 degrees as a result. When that happens, view cells trigger an incorrect location signal in place cells, which then gives rise to goal-search behavior in the rotationally equivalent corner.

Now, in their writings, proponents of the heading-based model do not explicitly posit a mechanism that segregate snapshots from distinct environments in memory. But the model has to be supplemented with such a mechanism. To begin with, note that snapshots and their associated heading from one environment do not bear any systematic relation to the visual perspectives and associated heading from another environment. From the fact that a snapshot registered in environment E1 with associated heading Y would display maximal matching with the current retinal input if the agent rotated its head by X degrees clockwise in the current environment E2, *we can conclude nothing about the current heading in environment E2*. The current heading could be close to (Y - X), but it could also lie far away from this value. For example, a rat might have recorded a snapshot of a corner of its home cage with an associated heading of 38 degrees clockwise from a reference direction, and yet that snapshot only produces high matching when the rat is facing a completely different direction in the experimental enclosure (144 degrees clockwise from the reference direction, say) because this is the only direction in the experimental enclosure where the shape of extended surfaces and the featural cues affect the retinal input in a way similar to the corner in the home cage.

If the rat then uses that snapshot from the home cage to calculate the current heading in the experimental enclosure, the view-matching system will yield a heading value which is off by 106 degrees clockwise.

From the fact that snapshots and their associated heading from one environment do not bear any systematic relation to the visual perspectives of another environment, we can infer that the view-matching system cannot literally use *all* of a subject’s stored snapshots to estimate the subject’s heading upon returning to a known environment in reorientation tasks. Recall that the final heading estimation is simply the value that comes up most often when comparing every snapshot to the current retinal stimulation. Of course, in a case where only a few mistaken heading values are added to the mix, the view-matching process will rarely output a final heading estimation that is significantly off the mark. Correct heading values will simply drown out the bad ones. But, unfortunately for proponents of the heading-based model, they cannot claim that reorientation tasks are cases like that. The heading-based model is committed to a profligate snapshot storage strategy. It holds that, whenever there are fewer than 15 snapshots that provide high matching with the current retinal input during exploration, a new snapshot is registered. This storage strategy leads to the registration of hundreds, if not thousands, of snapshots from environments like the waiting cage and the home cage. Thus, the numerous snapshots from these two environments would detract the view-matching process when the agent recovers normal perceptual access to its surroundings in the experimental enclosure, leading the process to yield a random value as the final heading estimation. This in turn would lead the subject to look for the desired goal in random locations. But, again, this behavior is not observed: subjects do not look for the desired goal in random locations. In all reorientation tasks where geometric and featural cues remain stable over time, subjects choose the diagonally adequate locations significantly more often than chance.[22]

---

[22] Someone might retort that, since bad heading estimations are randomly distributed in a 360-degree span — whereas the correct ones converge on the subject’s actual heading, or its actual heading *minus 180 degrees* —, the bad ones will not drown out the correct ones in the production of the final heading estimation. That reply won’t do for two reasons. First, given the rate and range of errors in the estimation of each snapshot-specific heading estimation reported in Sheynikhovich et al. (2009, Figure 12A, environment B-II), it is extremely likely that it is already only a tiny minority of the snapshots originally recorded in an experimental enclosure that yield correct heading values — namely, those snapshots that were recorded close to where the animal finds itself when it recovers perceptual access to its surroundings or close to the rotationally equivalent location in the diagonally opposite corner. Second, this reply cannot account for the fact that the subjects of the Julian et al. and Horne et al. experiments search for

Proponents of the heading-based model must therefore acknowledge that subjects select a *specific subset* of all the snapshots available in memory in order to initiate the view-matching process. The most promising option for them is to say that subjects store snapshots from the same environment together with tag concepts like ENVIRONMENT1, ENVIRONMENT2, ENVIRONMENT3. Subjects then select the relevant *tagged subset* when they return to a known environment.

So, the representation selection problem arises for the heading-based model in the following form: its proponents must explain how reorientation reliably subjects select the relevant *subset* of all recorded snapshots in order to perform the process of heading estimation at the outset of a trial. Snapshots from environments other than the current one would lead the view-matching system to make an incorrect heading estimation, except by chance. But, unfortunately, since this model has the same limited representational resources as the movement-based model, it cannot put the four selection principles discussed in the previous section to any better use in explaining how subjects reliably select the right subset at the beginning of the trial, the subset containing snapshots from the relevant environment. Just to give one example: the principle *Use Best Match to Retinal Stimulation* combined with this model cannot account for how Wills et al.'s subjects go about selecting the relevant subset because stored snapshots from the circular wooden enclosure would have a lower level of matching with the retinal stimulation in the circle-shaped morph box than stored snapshots from the square-shaped morph box, and the principle therefore wrongly predicts that subjects should select a subset containing snapshots from the square-shaped morph box.

Since the four principles we saw in the last section are the most promising from the perspective of the VM framework at large, it follows that the heading-based model is in no better position to handle the representation selection problem than the movement-based model.

---

food in two different diagonals in two enclosures with the same global geometry and orientation in an experimental room. If the subjects relied on *all of their stored snapshots* to estimate their heading in both enclosures, they would choose the same two diagonals in both enclosures or behave randomly in both. According to Sheynikhovich et al.'s simulations, geometry always has a much stronger impact than featural cues on the final heading estimation in the type of experimental enclosure used in those tasks.

5. Conclusion

I started this chapter by developing a new problem in the reorientation literature: the representation selection problem. It is the problem of explaining how reorientation subjects manage to reliably select a representation from the current environment at the outset of a trial. It arises because we have strong empirical reasons to believe, from the perspective of both GM and VM theorists, that reorientation subjects frequently initiate the main reorientation process with a representation of the current environment, even though they spend a significant amount of time in other environments (e.g., their home cage, a waiting cage) in the minutes and hours that precede a given trial. I went on to argue that the problem cannot be solved by appealing to procedural selection principles (such as *Most Recent Environment*), and so reorientation researchers will need to adopt information-based selection principles tailored to their preferred explanatory framework to provide a satisfactory response to the representation selection problem. I then rejected four information-based selection principles from the perspective of the two most prominent VM models: the movement-based model and the heading-based model. Given the resources available to VM theorists, the information-based principles I discussed were both representative and arguably the most promising candidates. To repeat, the principles fail to explain how subjects go about picking snapshots from the relevant environment at the beginning of a reorientation trial because in one crucial type of case the current environment differs from another environment in virtue of the location of featural cues with respect to the environments' global geometry in ways not properly captured by either visual or depth information (e.g., Julian et al., 2015) and because in another crucial type of case the relevant environment resembles another environment in virtue of its global geometry in ways that go against matching by visual information (e.g., Wills et al., 2005). The foregoing analysis therefore strongly suggests that the overall VM framework does not have the resources to handle the representation selection problem.

Chapter 5

How Geometric-Module Theorists Can Solve the Representation Selection Problem

1. Introduction

The previous chapter made the case that the VM the framework does not have the resources to solve the representation selection problem. The next question is: How does the GM framework fare with respect to the representation selection problem? We should note at the outset that, as things stand, GM theories cannot solve it either. That is because all extant GM models remain entirely silent about how subjects select the relevant geometric representation at the beginning of a trial. The one possible exception to this is Wang and Spelke's (2002) model, which one may interpret as holding that subjects choose the relevant geometric representation by relying on a stored snapshot linked to it. On that interpretation of their model, the snapshot whose content most closely matches the current retinal stimulation triggers the selection of its paired geometric representation. Unfortunately, this account of representation selection does not work for the reasons highlighted in Chapter 4 in relation to the principle *Use Best Match to Retinal Simulation*.

Though GM theories cannot solve the representation selection problem in their current form, I will argue in this chapter that there is a natural extension that can be made to GM theories that allows them to solve it. Showing this involves many steps, however. In Section 2, I provide a recap of the GM framework's core explanatory strategy, adding new details from Gallistel's highly influential GM account. In Section 3, I try to motivate a new type of extended GM account by appeal to evolutionary considerations pertaining to efficient spatial navigation, showing how to expand standard accounts like Gallistel's. The crucial move consists in positing a separate selection mechanism that gives an important role to geometric information in its operations. Sections 4 and 5 take a very specific stance, on behalf of this new account, on two important questions about which I had remained uncommitted in Chapter 2: What is the content of geometric representations which the geometric module receives from the selection mechanism? How does the selection mechanism perform its operations? In Section 6, I draw together various strands of arguments from the previous sections to show that my new account can offer a satisfying answer to the representation selection problem. This opens the door in Section 7 for predictions about multiple-enclosure reorientation experiments based on this new account. In Section 8, I provide an alternative explanation for an influential set of reorientation experiments whose results may seem to contradict an important element of my new account. Section 9 then considers and rejects the possibility for VM theorists to similarly posit a separate selection mechanism in order to address the representation selection problem. Finally, in Section 10, I point out that this new account entails that there are two distinct pathways through which featural information can affect reorientation subjects' performance.

2. Recap on the Geometric-Module Framework and Gallistel's Views on Alignment

In Chapter 1, I took Gallistel's (1990) influential GM model as the basis for my presentation of the GM framework, in combination with Lee and Spelke's (2010a) theory. Here, I will present Gallistel's ideas about the nature of the alignment process performed by the geometric module in more detail because they will prove useful for the rest of the discussion in this chapter.

First, recall that Gallistel, Lee and Spelke adopt the following kind of GM theory. When animals first explore an environment, they store a geometric representation, a representation of the global geometry of three-dimensional extended surfaces in that environment. Then, later, when they recover normal perceptual access to their surroundings in the same environment following a disruption of some sort, they retrieve that representation, which is fed to the geometric module. The geometric module then employs the representation to make an estimate of the subject's heading and location within the environment. It computes the subject's heading and location by *aligning* the retrieved geometric representation to a new geometric representation constructed from current perceptual input. It completely ignores featural information in its operations despite the fact that the animals themselves notice and register many salient featural cues in the environment. At the end of the alignment process, the geometric module sends the estimated heading and location to downstream path-planning mechanisms, which in turn use that information design a specific route to the goal location and motivate the animals to follow it.

At the heart of this account lies the following idea: that the geometric module performs its function by *aligning* two geometric representations. But how does that work exactly? Gallistel proposes that alignment proceeds in virtue of the matching of mathematical parameters that provide a summary of the global shape of three-dimensional extended surfaces. He focuses in particular on *principal axes*. The *major* principal axis is the line that goes along the length of a figure and separates it into two roughly equal areas, whereas the *minor* principal axis is the line that goes along the figure's width and separates it into two roughly equal areas (see Cheng, 2005 for formal characterization). In the case of a rectangle, the major axis passes through the middle of its short walls, while its minor axis passes through the middle of its long walls.

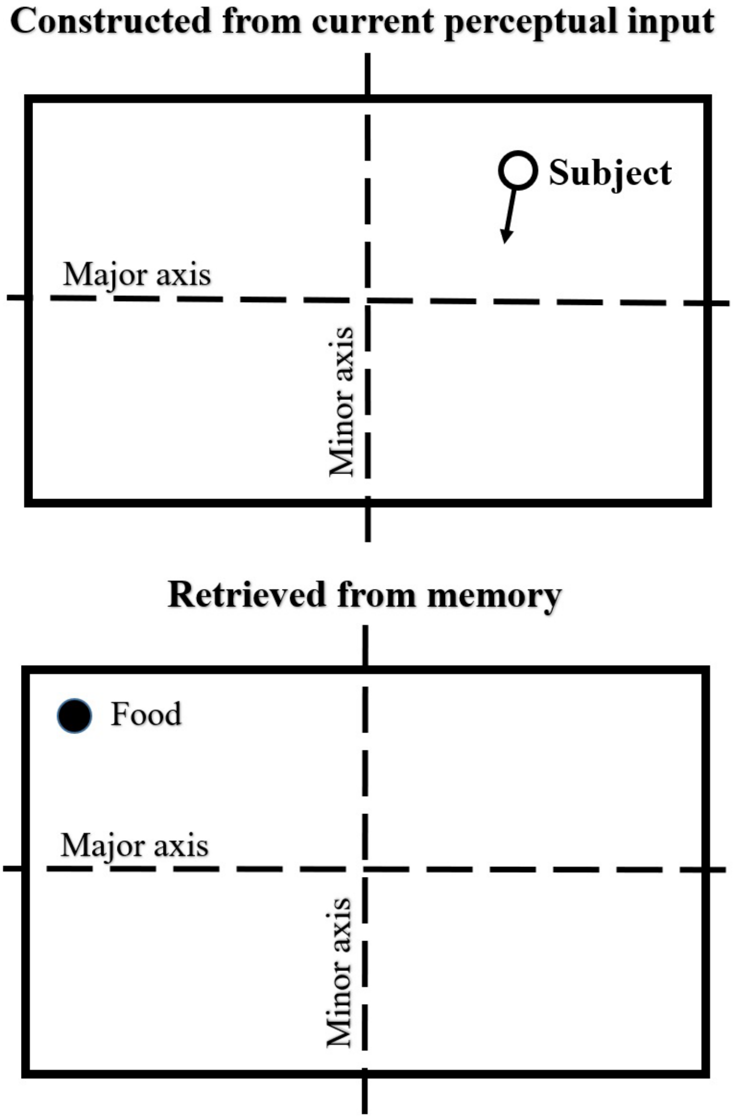

Figure 6. Principal axes from the geometric representation constructed from the current perceptual input (top) and principal axes from the geometric representation retrieved from memory (bottom). The open circle and connected arrow respectively represent the subject's estimated location and its estimated heading with respect to the representation.

The account under consideration holds that the geometric module aligns representations by matching two sets of principal axes. One set comes from the geometric representation constructed from the current perceptual input, the other set stems from the geometric representation retrieved from memory (Figure 6). The axes pertain to a two-dimensional figure — the two-dimensional figure formed by the overhead projection of the three-dimensional extended surfaces of the environment represented. (Many reorientation tasks employ a *rectangular* enclosure, so the figured formed by overhead projection in that case is, unsurprisingly, a rectangle.) The matching of axes itself unfolds in two steps: first, translate one set of major-minor axes such that the point where they intersect is on top of the point where the other set of axes intersect; second, rotate one set of major-minor axes until both major axes are on top of each other.

We do not need to delve further into the specifics here as this sketch is only intended to provide an illustration of this type of account. As Cheng and Gallistel (2005) point out, different GM accounts could employ many other alignment schemes based on shape parameters, some of which assume far more detailed parameters.[23] It suffices to note here that alignment by the matching of shape parameters allows the subject to estimate its own heading and location *with respect to the reference frame of the representation retrieved from memory*, the same frame of reference previously used to encode the location of the goal in the environment (Figure 6, bottom).

Details of the alignment procedure aside, GM theorists believe that the reason why subjects perform rotational errors in rectangular enclosures is that the geometric module completely ignores featural information. The module only relies on the content of geometric representations, and there are two ways of aligning geometric representations of rectangular environments so that their content matches. One way produces the correct estimation of the subjects' heading and location in the enclosure. The other yields an estimation of the subjects' heading that is off by 180º, and wrongly places them in the diagonally opposite corner of where they actually stand.

---

[23] See Kelly, Chiandetti, and Vallortigara (2011) and Sturz, Gurley, and Bodily (2011) for examples of papers that investigate the use of parameter-based alignment schemes in reorientation tasks.

3. Introducing the GM+SM Model

With these clarifications out of the way, we can now come back to the representation selection problem itself. I have already laid out in Chapter 2 some theoretical options about how to deal with issues of representation selection within the GM framework (see in particular section 4.5 therein). However, I did not take a clear stance on any of the options I described. I want to do things differently in the current chapter. I start by asking what an *efficient selection process* should look like from the perspective of GM theorists. This will lead me to adopt a far more committed and detailed theory of reorientation from within the GM framework than what I proposed in Chapter 2. This is what is needed to offer an adequate response to the representation selection problem.

So, what kind of information should an efficient selection process allied to the geometric module appeal to in its operations? On the new account that I propose here, subjects select the relevant geometric representation primarily by appeal to geometric cues themselves, though featural cues play a surprising, auxiliary role in the selection process. I motivate this new account by extending, with a slight twist, the logic of a popular argument for the existence of the geometric module.

On most GM theories, the function of the geometric module is to produce an estimate of the subject's heading and/or precise location within an environment upon recovery of perceptual access to its surroundings. Thus, as part of an evolutionary argument for the GM framework (Gallistel, 1990), many GM theorists emphasize the reliability of geometric cues for the estimation of heading and location in natural environments. For one thing, the shape of three-dimensional extended surfaces in the wild remains extremely stable over time. Cliffs, hills and plateaus barely change on a 100-year scale, barring rare geological events like landslides. For another, three-dimensional extended surfaces in the wild form highly irregular geometric shapes such that there is often only one way of aligning them with a stored geometric representation. So, this means that the process of estimating the subject's heading and location by alignment of geometric cues will very rarely be confounded by symmetrical environments in the wild as rectangular enclosures often confound the process because there are two ways of aligning their geometric cues to a stored representation.

This is good as far it goes. But the GM theorists' evolutionary argument does not explicitly address a crucial aspect of efficient navigation: representation selection. For the geometric module's estimation of heading and location to be reliable, it must have been computed in relation to a geometric representation of the *current* environment. Fortunately, just as geometric cues in the surrounding environment provide a highly reliable way of estimating your heading and location within it, they also offer a highly reliable means of singling out the right representation in the first place. Along with their extreme stability, three-dimensional extended surfaces in the wild form such irregular shapes that two distinct environments rarely possess the same global geometric properties. Even two valleys hardly ever have the same height, width or curvature. Thus, insofar as there were significant evolutionary pressures for efficient navigation upon sudden recovery of perceptual access, geometric cues should also play a dominant role in the selection process. More precisely, we should expect animals to follow a selection principle similar to this one: *Choose the stored geometric representation whose content most closely matches the content of the geometric representation constructed from current perceptual input*.

This does not constitute the whole story though. When we push our analysis of the factors contributing to efficient navigation further, we find an important asymmetry regarding the reliability of *featural* cues. On the one hand, featural cues taken individually are *extremely poor* indicators of heading and location within a known environment. If I visually match a featural cue to one I have seen before and I accurately estimate my distance D from it, this only tells me that I am likely to be located somewhere on a circle of radius D centered on its previously recorded location. And it tells me nothing about my heading. I need to rely on two or more visual featural cues to infer my current heading and precise location. Moreover, relying on multiple featural cues in this way can lead to substantial errors if any one of them has been misidentified or slightly displaced, as often happens with isolated objects like rocks or fallen trees. Smells, which also constitute featural cues, provide even less precise information than visually perceptible cues. Unsurprisingly, all of this provides support for the encapsulation from featural information of the process that estimates the subject's heading and location upon recovery of perceptual access to its surroundings.

On the other hand, however, featural cues taken individually constitute *moderately* reliable indicators of which environment someone is in. Perceptually matching with a high level of certainty a current featural cue to one I have perceived before provides defeasible evidence that I am back in the environment where I last perceived that cue. It does not matter much if only one featural cue gets matched, as opposed to two or three cues. Or if that cue has been moved slightly. Or if the cue is olfactory, as opposed to visual or tactile. The moderate reliability of featural cues in that context suggests a role for featural cues as *tie-breakers* in the selection process. For example, suppose that, upon recovery of perceptual access, two or more geometric representations fare approximately equally well in terms of how their content matches the global geometry of the surrounding environment. That would happen, for instance, if an agent visited two clearings with similar diameter, say, or two flat open fields. In that case, if she recorded some featural information from each environment and attached it to each environment's geometric representation, she could rely on that information to pick out the one geometric representation that pertains to the current environment.

Hence, the foregoing recommends the following commitment: subjects acquire and store a *feature map* for each environment they visit, and they index that map to the corresponding geometric representation of the environment.[24] A feature map, as I use the term here, is a representation of a set of featural cues identified in a given environment along with their location in it. In feature maps, featural cues are associated with an address label which specifies the location of each cue in terms of its distance and/or extension with respect to the three-dimensional extended surfaces of the environment. For example, the feature map of a white rectangular enclosure encodes the presence of the color white, and it associates that color with an address label that specifies that it completely covers all four walls.

The asymmetry regarding the reliability of featural cues for the estimation of the subject's heading and location versus representation selection also supports positing an additional component in charge of representation selection distinct from and external to the geometric module itself. I call it the *selection mechanism*. It implements the following selection principle:

> **Geometry First, Feature Maps Second** — Choose the stored geometric representation whose content best matches the content of the geometric

---

[24] Some prominent GM theorists (Cheng, 1986, 2005; Gallistel, 1990; Gallistel & Matzel, 2013) already hold this commitment, in one form or other, for independent reasons.

> representation constructed from the current perceptual input. If multiple representations match it roughly equally well, choose the geometric representation whose corresponding feature map best matches the feature map constructed from the current perceptual input.

The selection mechanism feeds the chosen geometric representation to the geometric module, and the latter then computes the subject's heading and location with it. Unlike the selection mechanism, the geometric module is not sensitive at all to featural information.[25] Information about the position of the goal linked to the selected geometric representation also helps further navigation and motor systems to plan a path to the estimated position of the goal in the environment.

I assume in what follows that the selection mechanism is modular in virtue of being domain-specific and encapsulated, though this does not constitute an essential tenet of the account developed here. Why should one make this assumption? First, the idiosyncratic nature of the principle that the selection mechanism implements, as well as the complex nature of the representations it deals with, suggests that it operates on the basis of specialized internal rules. This in turn increases the chance that it is highly *domain-specific* by virtue of dealing only with one task: the task of selecting a representation from the current environment and sending it to other navigation systems at various critical moments. Second, we may expect, for efficiency reasons, the selection mechanism to systematically ignore all other types of information besides geometric representations and feature maps. If so, that makes it *encapsulated* to a strong degree.

It may sound paradoxical to say that the selection mechanism is encapsulated even though it calls on feature maps. But note that just because the geometric module itself is encapsulated in virtue of ignoring all information distinct from geometric information it does not mean that this constitutes the only way for a mechanism to be encapsulated. Any component counts as encapsulated to the extent that architectural constraints prevent it from accessing large parts of the information contained in the mind (Fodor,

---

[25] The commitment to distinct mechanisms for selection on the one hand and the estimation of heading and location on the other hand is inspired by Julian et al.'s claim that we should posit distinct systems for place recognition and heading retrieval. It also gets support from the same considerations about separate modifiability that Julian et al. put forward for their claim.

1983, 2000). And human and non-human animals register and maintain an enormous amount of information about themselves and the world at any given time. From that perspective, geometric representations and their corresponding feature maps constitute a very restricted input class, only moderately larger than the geometric module's own extremely restricted input class. Notice also that the selection principle proposed above merely requires the selection mechanism to use featural information bound within a feature map, and thus indexed to a geometric representation through an address label. So, the mechanism may not even be able to operate on *naked* featural information, on featural information without an address label, as input.

It also follows that the new GM account presented here constitutes a moderate GM theory in that it acknowledges that subjects integrate geometric and featural information within some navigation systems outside the geometric module (as Cheng 1986, 2005 does). But, if proponents of this account further assume that the selection mechanism is modular, as proposed, it will also be more radical than other GM theories in the sense that it will posit *two* distinct modular components operating on geometric information, not just one. The first component is the geometric module itself, and as discussed above it always ignores featural information in its operations. The second component is the *selection mechanism (SM).* It implements the selection principle just described, is external to the geometric module and relies on both geometric representations and featural maps. Given the name of both components, I call this new account *the GM+SM model*.

(I want to emphasize here that I will not be using the expression 'GM+SM model' to refer to just any GM account that accepts the existence of a selection mechanism that feeds geometric representations to the geometric module. Following the arguments of Chapter 2, I take it as a given that any GM theorist needs to accept the existence of a mechanism whose role is to select the relevant geometric representation at the outset of a reorientation trial in order to send it to the geometric module. Rather, I will use the expression 'GM+SM model' to refer to the specific type of GM model that I develop in this chapter. This model is committed to a very specific selection principle, *Geometry First, Feature Maps Second*, as well as many specific claims about the content of geometric representations and the computational structure of the selection mechanism, as we will see in the next sections. It is also committed to the modularity of the selection mechanism in a way that the framework developed in Chapter 2 was not.)

It should be easy to see, at an intuitive level, how the GM+SM model can solve the representation selection problem. On the one hand, the selection mechanism gives priority to geometric cues pertaining to the shape and size of three-dimensional extended surfaces, and the experimental enclosure differs significantly in shape and size from the home cage and the waiting cage. So, the mechanism should reliably single out the representation of the experimental enclosure over those of the home cage or waiting cage at the outset of a reorientation trial. On the other hand, feature maps encode information about the identity and location of featural cues, and the selection mechanism treats feature maps as tie-breakers. Thus, the selection mechanism should reliably pick out the representation from the *current* experimental enclosure in reorientation experiments where there are two experimental enclosures that differ only in the identity or location of featural cues they contain (e.g., Julian et al., 2015).

Of course, this provides only an informal explanation of representation selection. In order to give a more detailed answer to the representation selection problem, we need to address the two following questions: What is the content of geometric representations? What are the computational algorithms by which the selection mechanism implements the high-level selection principle *Geometry First, Feature Maps Second*? I deal with these questions in turn in Sections 4 and 5. I then return to the representation selection problem in Section 6.

4. The Content of Geometric Representations

As noted in Section 2, some GM theorists hold that geometric representations include the *principal axes* of the overhead projection of the three-dimensional extended surfaces in an environment. On one common interpretation, this means that geometric representations encode information about the point of intersection of the axes as well as their relative length.

I agree that geometric representations contain some such information about shape parameters, but I will argue here that it is extremely unlikely that this information exhausts the content of geometric representations. In fact, my main proposal in what follows is that geometric representations also encode detailed *localized* information about the geometry of three-dimensional extended surfaces in an environment.

Localized information is information which can be associated to a specific location in space. For instance, information to the effect that there is a vertically extended surface at a certain location in an environment counts as localized information, even if the surface itself will generally extend significantly beyond that point.

The motivation for that proposal is that the best way to exploit the global geometric structure of three-dimensional extended surfaces in an environment for selection purposes is to pay attention to exactly how that global geometric structure is realized by the surfaces at a myriad of individual points in the environment. A selection mechanism that operated only on non-localized information would be severely impacted in its ability to select a representation from the current environment, and it would thus likely be selected against from an evolutionary perspective. Non-localized information includes, for example, the perimeter determined by vertically extended surfaces in an environment, and the relative and absolute length of the major and the minor axes contained within those surfaces. Many natural environments share one or many such properties without having much else in common.

This main proposal also receives important empirical support from a follow-up condition of the Colgin et al. global-remapping experiment described in Chapter 3:4.2. In the initial condition, subjects were exposed in alternation to a square-shaped enclosure and a circle-shaped enclosure on a large number of trials. In the follow-up condition, the same subjects were then exposed to a sequence of six enclosures made of the same material whose global shape gradually morphed from the square-shaped enclosure to circle-shaped enclosure (see Figure 7). Colgin et al. made the following observations in that context: subjects' place cells show highly similar configuration to the square-shaped enclosure on visits to the 1:7 and 2:6 octagonal enclosures; place cells show highly similar configuration to the circle-shaped enclosure on visits to the 3:5 and 4:4 octagonal enclosures; place cells display global remapping for each enclosure from the first set with respect to each enclosure from the second set and vice versa (see Wills et al., 2005 for similar results). The best explanation of these results from the perspective of GM theorists involves three claims: (i) place-cell activity across the six enclosures is driven or modulated by only two geometric representations that all subjects possess, a representation of the square-shaped enclosure and a representation of the circle-shaped enclosure; (ii) the subjects' selection mechanism chooses the representation of the square-shaped enclosure when they come back to the 1:7 or 2:6

enclosures; (iii) their selection mechanism chooses the representation of the circle-shaped enclosure when they come back to the 3:5 or 4:4 enclosures.

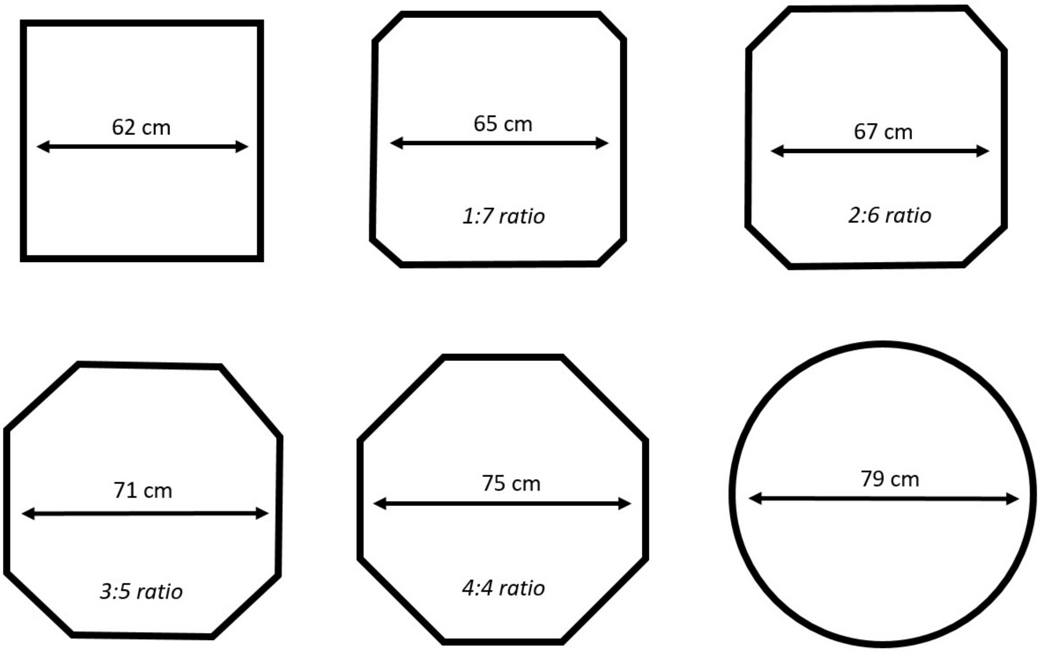

Figure 7. The six enclosures used in Colgin et al.'s follow-up condition. The octagonal enclosures' ratio number represents the number of 7.5-cm-wide sections that make up their short walls as compared to their long walls. For example, the 1:7 enclosure's short walls are each made up of *one* section, whereas its long walls are each made up of *seven* sections.

Why do these claims matter? They matter because they undermine the most plausible accounts of the selection mechanism that appeal only to non-localized information in order to explain its operations. First, because all six enclosures have the same perimeter by design, the subjects' selection mechanism could not have been merely relying on the *perimeter provided by vertically extended surfaces* to pick out the relevant representation. Second, their selection mechanism could not have been merely relying on the *relative length of principal axes* to choose the relevant representation either. Each pair of principal axes for the six enclosures have equal length (e.g., the two principal axes of the square-shaped enclosure are both 62-cm long). Third, it is highly unlikely that the subjects' selection mechanism merely used the *absolute length of principal axes* to select the relevant representation. The principal axes of the 3:5 enclosure (which are 71-cm long) are about equally close in length to those of the square-shaped enclosure and the circle-shaped enclosure (62 cm and 79 cm respectively), and subjects are bound to make small errors in the measurement of the axes' absolute length which, in many cases, would make their estimated length closer to those of the square-shaped enclosure by a few centimeters. So, if all the selection mechanism took into account was the axes' absolute length, subjects' place cells should have adopted the configuration typical of the square-shaped enclosure on close to half the sessions when the subjects found themselves in the 3:5 enclosure. But, by all measures available, they almost never did.

It is not clear what other type of non-localized information could play a role in the operations of the selection mechanism in a way that explains the changes in place-cell activity in Colgin et al.'s follow-up condition. Thus, Colgin et al.'s findings naturally lead us to the view that the selection mechanism must at least give some role to *localized* information about the geometry of three-dimensional extended surfaces in its operations, and thus support the main proposal presented above.

What would an account of the content of geometric representations that fits with this main proposal look like? I will present here one among many potential such accounts, mostly in order to provide a proof of concept. Figuring out which one of those is the correct account would require discussion of many more empirical results than I can address here or that currently exists. So, the account of content that I present here may well turn out to be false in the long run, and the GM+SM model is not committed to the details of this account. But for what follows it will help to have a concrete example of an account that appeals to localized information.

On the account I have mind, geometric representations primarily encode information about the boundaries created by vertically extended surfaces meeting horizontal surfaces (e.g., bottom of walls) and by sudden drops over an edge (e.g., a cliff). However, rather than encoding the complex shape formed by the boundaries in all their details, geometric representations encode an approximation of that shape in the form of linked, oriented *boundary segments* as seen from an overhead perspective.[26] Those segments have a set length (say, around 5 cm for rodents), and they can take one of many orientations in a 360-degree span. Most importantly, the position and orientation of each individual boundary segment is encoded with respect to the geometric representation's reference frame and coordinate system. See Figure 8 for a visual depiction of the geometric representation of Colgin et al.'s 4:4 octagonal enclosure on that account of content.

---

[26] This account of content is directly inspired by Byrne et al.'s (2007) model of place-cell firing. However, the GM+SM model does not adopt Byrne et al.'s claim that subjects often register a large number of geometric representations in a given environment, where each representation is anchored at a slightly different origin point.

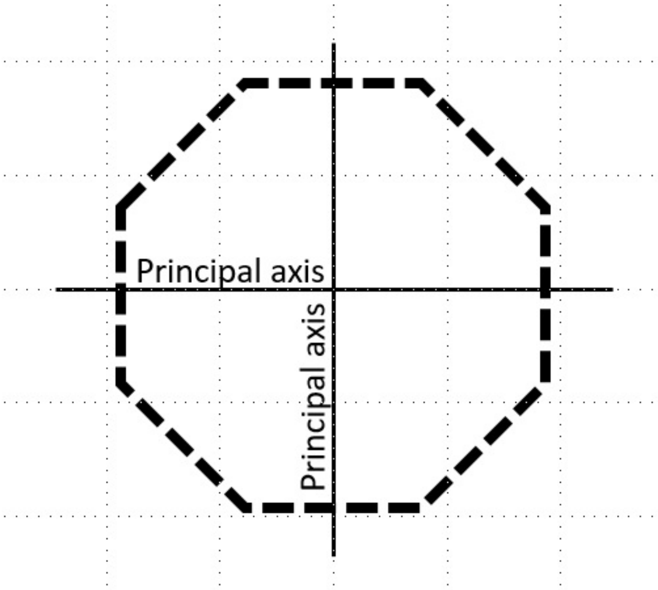

Figure 8. Visual depiction of a geometric representation from the 4:4 enclosure according to the account of content discussed in the main text.

Relatedly, this account of content may help to explain the Colgin et al. findings. We find much less overlap in the position and orientation of boundary segments when we superimpose an overhead projection of the 4:4 enclosure upon an overhead projection of the square-shaped enclosure (Figure 9, left) than when we superimpose it upon an overhead projection of the circle-shaped enclosure (Figure 9, right). So, assuming that the selection mechanism chooses the relevant representation based on the local conformity in the position and orientation of boundary segments, this account appears to predict that the selection mechanism will systematically choose the representation of the circle-shaped enclosure when the subject visits the 4:4 enclosure as opposed to the representation of the square-shaped enclosure.

Note, finally, that the inclusion of detailed localized information in geometric representations does not preclude the need for shape parameters. Quite the opposite. Any reasonable account of the content of geometric representations should hold that they include shape parameters, such as principal axes (as in Figures 8 and 9). Shape parameters allow the geometric module to perform a swift and efficient alignment procedure for purposes of estimating the subject’s heading and location (Gallistel, 1990). Nothing we saw above detracts from this. Moreover, the use of shape parameters simplifies the selection process itself, as I explain in the next section.

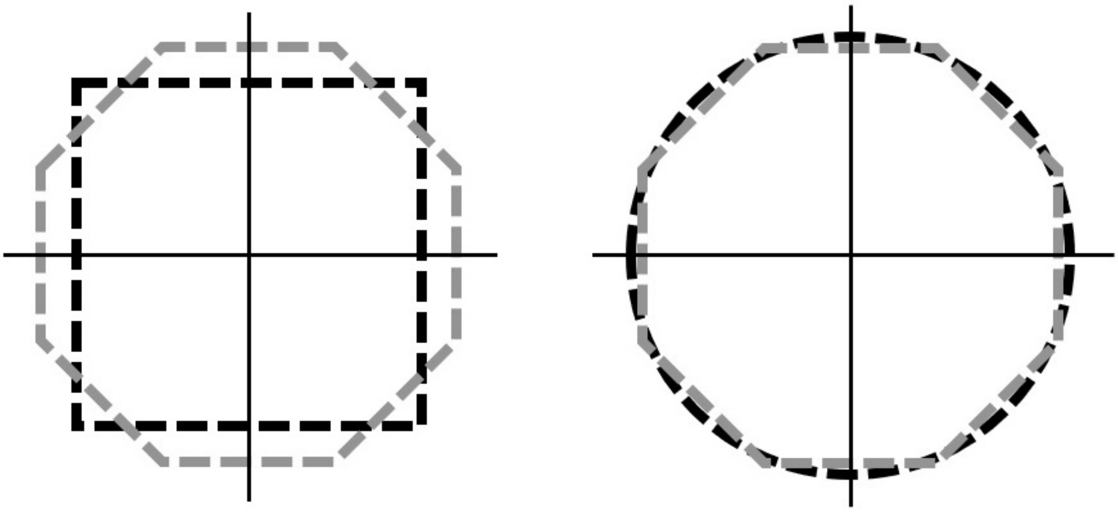

Figure 9. Visual depiction of a geometric representation from the 4:4 enclosure (in gray) when superimposed upon a representation from the square-shaped enclosure (in black, left) and upon a representation from the circle-shaped enclosure (in black, right).

5. The General Computational Structure of the Selection Mechanism

Since we are interested in a proof of concept, the details of the computational structure of the selection mechanism are not directly relevant. Instead, I will provide here an overview of the general computational structure of the mechanism. Many different types of computational algorithms could allow the mechanism to do an efficient job, and figuring out which one actually underlies its operations depends on many further considerations. Moreover, a general characterization will be sufficient to see why the GM+SM model offers significant headway on the representation selection problem. I also want something general in order to see how the GM+SM model could fit with a wide range of accounts of the content of geometric representations and of how the geometric module itself operates.

I start by appealing to simple efficiency considerations in order to sketch the mechanism's basic operations. First, we can expect the selection mechanism to use heuristics to initially narrow down the number of stored geometric representations that will subsequently go through a more thorough comparison process with the geometric representation constructed from current perceptual input — henceforth the *current representation*. This would help to reduce the computational costs of the selection process. One plausible such heuristic, for example, may consist in rejecting stored representations for which the sum of the absolute length of their principal axes falls outside a plausible range of the sum of the absolute length of the principal axes of the current representation.

Second, we can expect the selection mechanism to produce a *geometry-matching index* for each remaining stored geometric representation as compared to the current representation. If one of the stored geometric representation has a significantly higher index than other representations with respect to the current representation, the selection mechanism should then send that representation to the geometric module. Calculating such an index is important because there will often be small discrepancies in the content of geometric representations constructed in the same environment at different times due to unavoidable noise in the estimation of geometric properties. The selection mechanism thus needs to maximize the degree of matching with the current representation rather than seek a perfect fit, and this in turn depends on the use of a matching index calculated according to fixed rules.

Third, in cases where two or more geometric representations have a higher matching index than other representations, but not significantly more than each other, the selection mechanism should fall back on the *feature-matching index* of these representations as a tie-breaker. The selection mechanism computes a geometric representation's feature-matching index by comparing its associated feature map to the feature map associated to the current representation. Again, calculating such an index is important because there will rarely be a perfect fit among feature maps constructed in the same environment at different times. The selection mechanism should then send the geometric representation with the highest feature-matching index to the geometric module.

Fourth, if no stored geometric representation scores higher than a certain fixed minimal value for the geometry-matching index, the selection mechanism should send a signal to other navigation systems to the effect that the subject is currently in a new environment. If there is a similar threshold for the feature-matching index, however, we should expect it to be much lower given the much greater variability of featural cues over time in natural environments.

Making claims about the selection mechanism's operations beyond these four points becomes trickier. In what follows, I draw out the implications of the argument made in

Section 4 to the effect that geometric representations encode localized information about extended surfaces. A corollary of that argument is that the selection mechanism should calculate the geometry-matching index of a given stored geometric representation by comparing its localized information with the current representation's localized information at a large number of points once the two representations have been *aligned* one to another. Aligning two representations consists in putting their reference frame in correspondence, and it requires the use of shape parameters. As explained in Section 2, when done using principal axes, alignment consists of matching the principal axes of the two representations in the following way: first, translate one set of major-minor axes such that the point where they intersect is on top of the point where the other set of axes intersect; second, rotate one set of major-minor axes until both major axes are on top of each other.

How does the selection mechanism calculate the geometry-matching index by comparing localized information once the two representations are aligned? This will depend heavily on the exact content and format of geometric representations, and thus cannot be decided without appeal to auxiliary assumptions. Consider, for example, the account of the content of geometric representations presented in Section 4 and illustrated in Figures 8 and 9. On that account, we can reasonably hold that the selection mechanism calculates the geometry-matching index as follows. For every boundary segment in the current representation, the mechanism identifies the closest boundary segment in the stored representation and estimates the distance between the two segments and the difference in their respective orientation. It then computes a *local conformity value* at that coordinate in such a way that the closer the two segments' respective position and orientation are to one another, the higher that value is. Once this has been done for a large enough number of boundary segments in the current representation, the selection mechanism tallies all the local conformity values to produce the geometry-matching index.[27]

What about the calculation of the feature-matching index? Feature maps are built on the reference frame provided by their associated geometric representations. Therefore, the

---

[27] Suppose that the alignment procedure identifies more than one possible way of aligning the two representations, as with representations of symmetrical environments. What then? In that case, the selection mechanism should first calculate the geometry-matching index for each possible way of aligning the representations, and then pick the highest value as the *overall* geometry-matching index.

selection mechanism may use the following general strategy for the computation of the index: align their associated geometric representations, then produce local conformity values by comparing the identity of the featural cues at corresponding points in the feature maps, and finally tally the local conformity values. This strategy would allow the selection mechanism to be sensitive to both the identity and location of featural cues in its operations.

6. How the GM+SM Model Can Solve the Representation Selection Problem

Sections 4 and 5 have provided us with more detailed characterizations of the content of geometric representations and of the computational structure of the selection mechanism. So, we are now in a good position to see how the GM+SM model can solve the representation selection problem.

One important aspect of the representation selection problem is to explain how reorientation subjects reliably select the representation from the experimental enclosure at the outset of a trial as opposed to a representation from the home cage or the waiting cage. The GM+SM model tackles that aspect as follows. The experimental enclosure often differs substantially in shape and/or size from the home cage and the waiting cage, generally covering more than twice the area of either cage. Moreover, when the selection mechanism compares two representations from environments that differ significantly in shape and/or size, it will calculate a low geometry-matching index. That is because the locations where one representation indicates the presence of vertically extended surfaces will generally be far from the locations where the other representation indicates the presence of vertically extended surfaces once the two representations have been aligned to one another (as Figure 9, left, illustrates in the case of environments that differ in shape). And this will negatively impact assessments of local conformity between the two representations. On the other hand, when the selection mechanism compares two representations from environments of similar shape and size, it will calculate a much higher geometry-matching index precisely because the locations where each representation indicates the presence of vertically extended surfaces will be very close to each other (as in Figure 9, right). For these reasons, the selection mechanism will calculate a much higher geometry-matching index for the representation of the experimental enclosure over the representations of the other two cages when the subject finds itself in the experimental enclosure. Thus, the selection mechanism will select the former over the latter at the outset of a reorientation trial.

The other main aspect of the representation selection problem consists in explaining how reorientation subjects select a stored representation from the *current* experimental enclosure at the outset of a trial when they have been visiting in alternation two experimental enclosures of the same shape and size which differ in terms of the identity or location of featural cues that they contain (e.g., Julian et al., 2015). The GM+SM model handles that aspect as follows. Because the experimental enclosures have the same shape and size, their respective stored geometric representations will yield a similar geometry-matching index as compared to the current representation. Therefore, the selection mechanism will fall back on their feature-matching index as a tie-breaker. The selection mechanism should calculate a fairly high feature-matching index for the stored geometric representation from the *current* enclosure because of the high level of local conformity in featural cues at every point in its associated feature map as compared to the current representation's feature map. In contrast, the mechanism should produce a significantly lower feature-matching index for the stored geometric representation from the other enclosure because of the drastic discordance in at least a few points in its associated feature map. The difference in the value of the two representations' feature-matching index will then lead the selection mechanism to pick out the geometric representation from the *current* experimental enclosure over the representation from the other enclosure.

Hence, the GM+SM model offers a natural extension to standard GM accounts that allows them to deal with the two main aspects of the representation selection problem.

7. Predictions about Multiple-Enclosure Reorientation Experiments

Based on the response to representation selection problem just outlined, the GM+SM model makes some relatively strong predictions about a specific class of reorientation experiments: reorientation experiments where subjects are exposed in alternation to multiple experimental enclosures with different goal locations. On the other hand, other GM theories and VM theories cannot make predictions about such experiments because, without a specific potential solution to the representation selection problem — something no extant GM or VM account has ever provided —, they are in a very important sense incomplete. So, the predictions described here are distinctive of the GM+SM model.

An important type of prediction of the GM+SM model stems from the claim, defended in 8.4, that the selection mechanism should send to the geometric module representations from environments with similar shape and size as the current environment at the beginning of a trial. On that basis, the model predicts that, when exposed in alternation to two experimental enclosures with different shapes (e.g., a 60cm-by-120cm rectangular enclosure, and a triangular enclosure with a base of 60 cm and two long walls of 120 cm), subjects in a proper motivational state will search at geometrically adequate corners significantly above chance in both enclosures. We should also expect that, when exposed in alternation to two rectangular environments that are scaled versions of each other (e.g., a 60cm-by-120cm enclosure and a 90cm-by-180cm enclosure), subjects in a proper motivational state will go back to diagonally adequate corners significantly above chance in both enclosures.

The evolutionary analysis developed in Section 3 in order to motivate the GM+SM model also makes some predictions, though in a subtler way. That analysis primarily rests on the claim that evolutionary pressures have led to the acquisition of a mechanism for representation selection that gives a dominant role to geometric cues because of the cues' reliability for selection. If that claim is right, it stands to reason that this mechanism should receive and operate on geometric information that tracks, at least indirectly, each category of geometric cues that has the following three properties: (i) cues of that category are stable over time; (ii) they vary significantly across natural environments; (iii) they are relatively easy to exploit for selection purposes from a computational perspective (or they correlate with other geometric cues that are). This, in turn, suggests that any category of geometric cues that satisfy these three properties might have an impact on subjects' performance in multiple-enclosure reorientation experiments.

For example, we can make a good case that *the height of vertically extended surfaces* possesses these three properties. First, the height of specific vertically extended surfaces does not change overnight in the wild. Cliffs, mounds and large rocks maintain their height for long periods of time. Second, the height of vertically extended surfaces varies widely across locations and environments, from a few centimeters to a few kilometers in some contexts. Third, height cues appear relatively easy to exploit from a computational perspective. On the one hand, geometric representations could encode the height of vertically extended surfaces by simply appending to each location in their reference frame a value indicating the height reached by the vertically extended surface at that coordinate, if there is any. On the other hand, the selection mechanism could easily take those values into account in its operations by treating it as another dimension of assessment when comparing the localized information at corresponding points in aligned representations. Assuming that height cues possess these three properties, the GM+SM model then predicts that, when exposed in alternation to two rectangular enclosures that have some walls whose height differ, subjects will search at the diagonally adequate corners significantly above chance in the two environments, even if the walls closest to the goal location in the two enclosures have the same height.

In short, though framed at a high-level, the GM+SM model has relatively strong commitments about the use of the shape and size of environments, as well as the height of vertically extended surfaces, for selection purposes. This, in turn, hints at specific patterns of behavior in multiple-enclosure reorientation experiments. Other GM and VM theories cannot make predictions about such experiments because, without a proposed solution to the representation selection problem, they remain in an important sense incomplete.

8. What Happens in Enclosures Composed of Fragmented Surfaces?

In this section, I want to address a potential challenge to the GM+SM model that arises from an influential set of empirical results obtained by Lee, Sovrano and Spelke (2012). Lee et al. performed reorientation experiments whose results they interpreted as supporting a type of GM model on which geometric representations explicitly encode the distance and sense properties of the goal location related to three-dimensional extended surfaces, but nothing else — where *sense* relates to the distinction between left and right, between (e.g.) a goal location being to the left of a specific wall or to its right. On Lee et al.'s view, it could be, for example, that what is encoded in a given geometric representation of a rectangular enclosure is merely the fact *that the goal is at the corner on the left of an extended surface which is closest from the center of the environment*. They took their results to show, in particular, that reorientation subjects' geometric representations *do not* encode, either explicitly or implicitly, angles at which three-dimensional extended surfaces meet or the length of three-dimensional extended surfaces themselves. They do not encode, for example, the length of the long walls of a rectangular enclosure or the angles at which a short wall meets a long wall.

If Lee et al.'s interpretation of their results is right, their experiments would therefore seem to completely undermine my claim that geometric representations encode a lot of details about the global geometry of three-dimensional extended surfaces in a given environment. And recall that this claim played an important role in the argument in Section 6 to the effect that the GM+SM model can solve to representation selection problem. Moreover, though I have not made this claim explicitly yet, it is virtually guaranteed that, on the GM+SM model, geometric representations encode information that tracks at least indirectly the length of three-dimensional extended surfaces in a visited environment. That is because localized information about the geometry of three-dimensional extended surfaces encoded with respect to a reference frame from an overhead perspective is bound to indirectly track properties like the length of three-dimensional extended surfaces. Consider, for example, the toy account of the content of geometric representations presented in Section 4. On that account, the length of a specific wall is encoded implicitly by the number of boundary segments that reside on a straight line where the wall is located.

So, it better be, for the sake of the GM+SM model, that there is something wrong with Lee et al.'s interpretation of their experiments. This is what I argue here. I do believe that their results are important and revelatory about how best to develop GM accounts — but not in the way and for the reasons suggested by Lee et al. To see why that is the case, we need to look carefully at Lee et al.'s experimental methodology.

Here is what they did. Lee et al. performed a number of reorientation experiments with two-year-old children in rhombus-shaped enclosures composed of continuous extended surfaces (Figure 10A), fragmented straight surfaces (Figure 10B), or fragmented corner-forming surfaces (Figure 10C). Subjects were incited to try to find a desired object hidden in one of four containers. For each of these three conditions, the hiding containers were either located in the four corners of the rhombus or in the middle of the four sides of the rhombus (whether there was a surface close-by or not). It turns out that subjects in the first and second conditions were able to rely on the three-dimensional surface layout to search for the desired object in one of the two geometrically adequate containers — *but not in the third condition, where they were at chance between the four* *hiding containers*.[28] Lee et al. interpreted the results from the third condition as showing that subjects' geometric representations *do not* encode the angles at which three-dimensional surfaces meet or the relative length of the principal axes in an environment. They took the results from the second condition, along with other conditions not mentioned here, as showing that geometric representations encode distance information. (For that condition, they took subjects' representations to encode something like this: "the target location [is] at the wall whose left side is further from the center (or, equivalently, from my position) than its right side" (p. 153) in the case where the hiding containers were located in the middle of the sides of the rhombus.) Finally, they took subjects' incapacity to rely on walls of different lengths in a fragmented square-shaped enclosure (not pictured) as showing that the length of three-dimensional extended surfaces is not encoded in reorientation subjects' geometric representations.

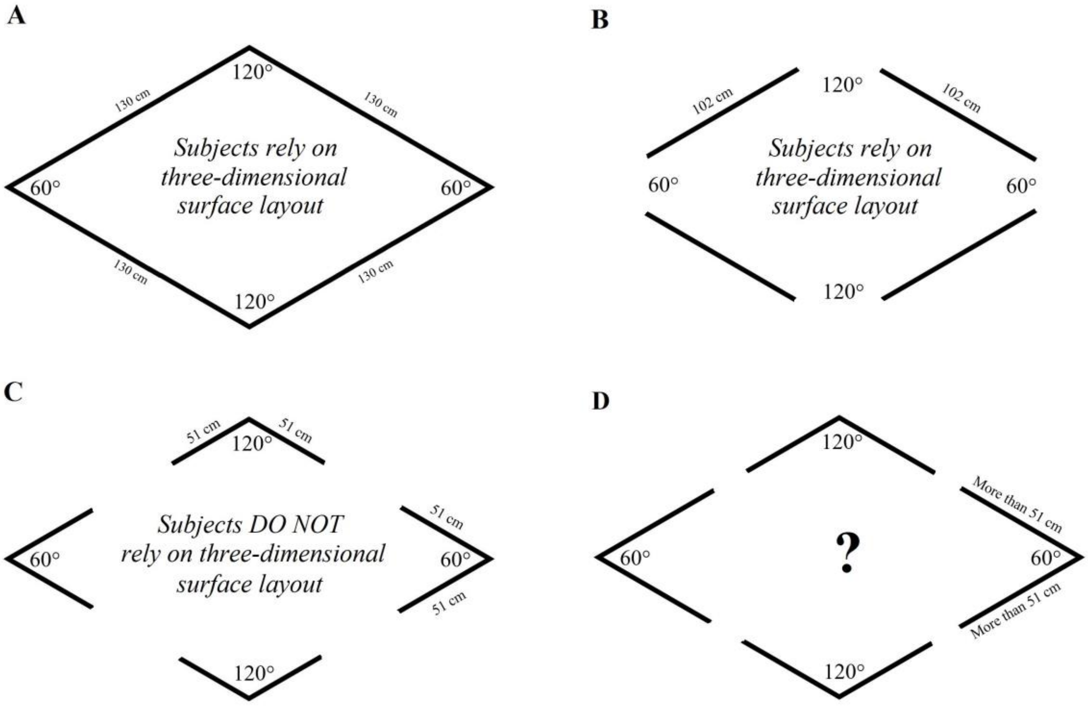

Figure 10. Rhombus-shaped experimental enclosures used by Lee, Sovrano and Spelke (2012): (A) enclosure with continuous surfaces; (B) enclosure with fragmented straight surfaces; (C) enclosure with corner-forming surfaces. The fourth enclosure (D) is an enclosure that could be used for a proposed test of the prediction made by one of the alternative explanations considered in the main text.

What should we make of this? The first thing to note here is that the implicit explanation of the experiments on which Lee et al. rely to make these claims about the

---

[28] The location of the hiding containers (whether the four containers were in the four corners of the rhombus or rather in the middle of the four sides of the rhombus) had no significant effect on search behavior in all three conditions.

content of geometric representations is incomplete as it stands. In order to explain the difference between the results of the second and third conditions, it is not sufficient to say that information about the angles at which three-dimensional extended surfaces meet to form corners is not encoded in geometric representations (but information about distances to three-dimensional extended surfaces is). That is because distance information would be enough to help subjects *in the third condition* to return to the geometrically adequate locations above chance. Subjects in the third condition could return to the geometrically adequate corners if all they encoded was the following geometric information: *that the target location is at the corner-forming surface which is closest/farthest from the center* in the case where the hiding containers are in the corners of the rhombus, or *that the target location is on the left/right from the corner-forming surface which is closest from the center* in the case where the hiding containers are in the middle of the sides of the rhombus.

Hence, it must be that there is something about the fact that the surfaces in the third condition *form corners* (i.e., are not straight) that prevents information about them from making it into subjects' geometric representations and thus eventually be used by the geometric module to perform its function. *However, once you realize that, it becomes clear that Lee et al.'s claim that geometric representations do not encode angle or length information is explanatory otiose.* What really does the work is the claim that information about corner-forming surfaces of the type used in the third condition is not encoded in the subjects' geometric representations, whereas geometric information about straight surfaces of the type used in the second condition *is* encoded in geometric representations. Someone might think that it is a distinction without a difference. But, in fact, this distinction matters quite a lot because it points toward the need for a kind of explanation altogether different from Lee et al.'s preferred explanation.

So, what is it about the fact that the fragmented surfaces in the third condition are not straight which could prevent information them form making into geometric representations?[29] One possibility which I find particularly plausible comes from the

---

[29] The idea that it is just the fact that they are not straight *in and of itself* which prevents information about these surfaces from being encoded in subjects' geometric representations strike me as highly implausible. There are very few perfectly straight three-dimensional extended surfaces in the wild, and so adopting this idea would rob the GM framework of any remotely plausible evolutionary story as to how many species came to possess the geometric module.

idea that there is a *width threshold* that determines when perceptual systems treat something as a three-dimensional *isolated object* versus a three-dimensional *extended surface*. Recall from Chapter 1 that three-dimensional isolated objects count as featural cues and therefore information about them do not make into geometric representations. Now, there must be something in virtue of which some elements in the visual field get treated as three-dimensional *extended surfaces* and others get treated as three-dimensional *isolated objects*. Being a three-dimensional extended surface (as opposed to a three-dimensional isolated object) is not an objective property in the world when conceived as an all-or-nothing characteristic: There is a whole spectrum from highly isolated objects (e.g., an atom in the void) to massively extended ones (e.g., the sun). But, on GM models, information about a given three-dimensional object either makes into a given subject's geometric representation of the current environment, or it does not. It is all or nothing. So, perceptual systems' decision to classify objects as belonging to one class or the other must be based on some more or less arbitrary, fixed internal criterion.

This is where the idea of positing a *width threshold* might be useful. On the kind of account I have in mind, what prevents geometric information about the corner-forming surfaces in the third condition from being recorded in geometric representations is that these surfaces *fall below a certain width threshold*, and so do not get treated as three-dimensional extended surfaces at all. They get treated like three-dimensional *isolated objects*, and thus information about them does make it into subjects' geometric representations. For example, it could be that, for humans, objects that are more than, say, 95-cm wide when facing them get treated as three-dimensional extended surfaces, and that objects that are less than 95-cm when facing them get treated as three-dimensional isolated objects. The corner-forming surfaces in the third condition are 88-cm wide (for obtuse corners) and 51-cm wide (for acute corners) when the subject is standing in front of them, so they would fall below this specific width threshold, whereas the 102-cm-wide straight surfaces in the second condition would fall just above the threshold. Importantly, on this view, geometric representations could still be extremely detailed (as proposed in Section 4, and as needed for my argument in Section 6). It is only that three-dimensional objects need to reach a certain width for (potentially highly detailed, localized) geometric information about them to make it into subjects' geometric representations.

In sum, once we realize that what really does the required explanatory work with respect to Lee et al.’s results is the assumption that information about the corner-forming surfaces in the third condition do not make it into geometric representations (but information about the straight surfaces in the second condition do), it follows that Lee et al.’s results do not threaten the GM+SM model’s commitment to the view that geometric representations encode highly detailed, localized geometric information about three-dimensional extended surfaces as such.

Of course, there are other possible explanations as to why information about the corner-forming surfaces in the third condition do not get into subjects’ geometric representations. But the one I have just proposed seems to me to be particularly well-motivated. Note also that, if this explanation is right, it gives us the following prediction (Figure 10D): in a fragmented rhombus-shaped enclosure with corner-forming surfaces which exceed a certain yet unspecified width threshold (somewhere between 88 cm and 102 cm for 2-year-old human children), subjects will return to the two geometrically adequate locations significantly above chance.

9. Couldn’t View-Matching Theorists Appeal to a Separate and Similarly Flexible Selection Mechanism?

In this section, I want to consider one potential objection to the overall analysis proposed in this chapter and previous one. Recall that I argued in last chapter that the VM framework does not have the resources to handle the representation selection problem. Then, I pointed out in the current chapter that GM theories could not tackle the problem in their current form either. However, I went on to propose a new GM account built around a new hypothesized component, *the selection mechanism*, in order to tackle the problem. This component makes use of both geometric cues and featural cues according to complex internal rules to pick out the relevant geometric representation.

At this point, one may wonder: Couldn’t proponents of the VM framework deal with the representation selection problem by similarly positing a separate selection mechanism that operates on multiple types of cues according to complex rules? This is a good question. Unfortunately, the answer is that VM theorists cannot incorporate into their accounts a selection mechanism that delivers what they need without incurring substantial costs. To see why, note that, in order to get any explanatory benefits from the idea of positing a separate selection mechanism in dealing with experiments discussed in Chapter 4 and Section 4 of the current chapter, proponents of the VM framework would have to claim that the selection mechanism operates on global geometric representations combined with feature maps in order to select the relevant snapshot. In other words, they would need to adopt a *hybrid* account, the best version of which goes as follows. When subjects explore a new environment with goal locations, they memorize representations of three types: global geometric representations, feature maps and snapshots. A global geometric representation and a snapshot that are registered in the same environment, or at the same time, get linked together. Later, upon returning to a known environment, the selection mechanism picks out a stored global geometric representation using the principle *Geometry First, Feature Maps Second*. Following that, the agent initiates the view-matching process described by VM theorists with a snapshot previously linked to the selected global geometric representation.

But such a hybrid account carries substantial costs indeed for theorists sympathetic to the VM framework. First, it forces them to accept the view that animals construct representations of the global geometry of three-dimensional surfaces as such and that these representations play a major role in the cognitive processes involved in reorientation tasks, in direct opposition to one of the main motivations behind the VM framework. In fact, on a natural way of understating the debate between the two frameworks (see Chapter 3:3.2), this view directly contradicts the central commitment of the VM framework. Second, on threat of simply collapsing into a GM theory, such an account must assume that, though subjects systematically store global geometric representations to help with the snapshot-selection process, they *never* directly use these representations to guide their search behavior when looking for known goal locations. That is what the selected snapshot is for. But it is not clear why animals would *not* avail themselves of these geometric representations to return to known goal locations, when geometric cues are at least as useful for finding such locations as they are for selection purposes. Moreover, the global geometry of three-dimensional extended surfaces is a far more reliable guide to known goal locations than the visual appearance of the area or extended surfaces surrounding those locations. The visual appearance of any area or surface in the wild changes based on the time of day, the season, and the subject's exact location. These two considerations thus seem to rule the possibility of adopting the sort of hybrid account just sketched, thereby showing that VM theorists cannot answer the representation selection problem by simply positing a separate selection mechanism.

10. Two Pathways Dealing with Featural Information

I end this chapter by drawing an important implication from the GM+SM model about the use of featural information in reorientation tasks. If the GM+SM model is right, it means that there are at least *two distinct pathways* through which information about featural cues as such can affect reorientation performance: one that involves the beacon-homing process (see Chapter 3:3.2) and the other centered around the selection mechanism. It might be tempting to suppose that a single mechanism can subserve the use of featural information for beacon homing and representation selection. However, the beacon-homing process and the selection mechanism are hypothesized to have vastly different properties and functions. The former operates in parallel to the geometric module and competes with it to guide search behavior. It is also reinforcement-dependent and based in the striatum, and it completely ignores geometric information in its operations (Doeller & Burgess 2008; Doeller, King, & Burgess, 2008; Lee et al., 2017). The latter gives a dominant role to geometric information in its operations, cooperates closely with the geometric module, is reinforcement-independent and cannot guide behavior in any direct way. Moreover, the fact that striatal lesions do not negatively impact geometry-based search behavior in navigation experiments similar to reorientation tasks (McDonald & White, 1994; see also Lee et al., 2017) suggests that the selection mechanism is not based in the striatum at all. If it was, striatal lesions should prevent it from doing its job and the geometric module would not receive any input representation from it, arguably preventing geometry-based search behavior.

A significant amount of research in the reorientation literature has focused on the question of whether the impact of featural cues on reorientation performance in specific contexts merely arises due to a beacon-homing process or to a more complex process that integrates featural and geometric information. GM theorists have maintained that it is due only to beacon homing (Lee, Shusterman, & Spelke, 2006), whereas many influential reorientation researchers have chosen the latter option as part of an argument against GM accounts (Pearce, 2009; Twyman & Newcombe, 2010). The GM+SM model has the potential to turn the specific framing of that dispute on its head. It suggests that the impact of featural cues on performance is due both to a beacon-homing process *and* to a process that integrates geometric information and featural information. But the model does so in a way that respects the core tenets of the GM framework.

11. Conclusion

The goal of this chapter was to show that the GM theorists can provide a satisfying solution to the representation selection problem. For this purpose, I have developed a new type of extended GM theory, which I call the GM+SM model. It posits a separate selection mechanism that gives an important role to geometric cues in its operations. I have shown that this model can provide a response to the problem by appealing to the localized information contained in geometric representations to distinguish environments which have different sizes and shapes (such as the home cage versus the experimental enclosure), as well as the information about the identity and location of featural cues contained in feature maps.

Conclusion

The geometric-module (GM) framework is an explanatory framework that aims to account for the results of a specific type of experiments about spatial navigation: reorientation experiments. Though it went largely unchallenged from its initial formulation in 1986 until about 2005, it has since faced a number of important challenges. The main goal of this thesis was to defend the GM framework from two such pressing challenges in order to use this framework to make some headway regarding two foundational debates in philosophy of psychology, the debate about the nature of mental representations and the debate about the structure of the mind. I will start by summarizing how this defense proceeded. I will then point out that a number of novel ideas that I put forward in this thesis to defend the GM framework in fact strengthens the support that it offers for rather radical positions on these two debates: the view that human and non-human animals' representations encode high-level abstract properties about the world and the massive modularity thesis.

Here is how the defense of the GM framework went. In Chapter 2, I tackled the most important objection to the GM framework, the explanatory inflexibility objection. It holds that the GM framework cannot explain substantive variations in whether and how subjects rely on featural cues in transformational tasks because the framework does not have the explanatory latitude to do so. My response to this objection consisted in developing *the selection interference strategy.* There are very strong empirical reasons to believe that reorientation subjects can hold in a long-term memory store geometric representations and feature maps from multiple distinct environments, and that the selection mechanism makes use of feature maps to select the relevant geometric representation at the outset of a trial. Thus, changes to the shape of the enclosure and/or the position and nature of featural cues within the experimental enclosure across training and testing phases of transformational tasks may well prevent the selection mechanism from operating normally in a variety of contexts. This provides GM theorists with four general moves for dealing with variations in the results of transformational tasks: claiming (1) that the geometric representation of the experimental enclosure recorded during the training trials was discarded before the testing phase due to one of a variety of potential factors related to proper maintenance of the memory store; (2) that, even though the geometric representation acquired during the training trials is still in the memory store, the selection mechanism picked out a representation that is distinct from the one of the experimental enclosure that was recorded during training trials; (3) that, even though the representation of the experimental enclosure is still in the memory store, the selection mechanism did not manage to pick out even one representation in the trials of the testing phase (4) that, even though the relevant geometric representation is still in the memory store and the selection mechanism *did* manage to feed it to the geometric module, the selection mechanism triggered the environmental-change signal as it did so. Transformational tasks can lead to each of these four situations depending on their exact methodology, and each type of situation can have effects on search behavior that seem at odds, at first sight, with predictions of GM models. Hence, the selection interference strategy shows that the GM framework can provide an empirically well-motivated response to the explanatory inflexibility objection. It also points toward the importance for GM theorists to take a stance on seven substantive issues in the future (such as how to think about the content of geometric representations or the internal operations of the selection mechanism) to provide more detailed explanations of the behavioral patterns observed in transformational tasks.

In the three remaining chapters, I then turned to arguably the biggest and most complex challenge against the GM framework. That challenge arises from the view-matching (VM) framework of spatial reorientation. The VM framework was the second detailed explanatory framework to emerge from the reorientation literature, and it aimed to break the hold which the GM framework had on the literature (Cheng, 2008). It rejects the idea of positing a modular mechanism in higher cognition that only operates on geometric information to account for subjects' search behavior. The framework's explanatory strategy depends rather on positing a view-matching system that compares recorded snapshots of the experimental enclosure with the current retinal stimulation to return to the goal location. The GM and VM frameworks now represent the two most influential theoretical approaches to explaining the results of reorientation tasks.

We can summarize the dialectic of the last three chapters as follows. To settle the debate between the two frameworks, reorientation researchers have focused until now on a specific set of behavioral and neuroscientific findings. The former pertain to how specific types of arrangements of geometric or featural cues affect search behavior in non-transformational one-enclosure reorientation experiments. The latter relate to the discovery of brain regions and of spatially-tuned cells whose activity demonstrate a sensitivity to properties of three-dimensional extended surfaces. The first main contribution of the thesis to the debate was an extended argument that these findings do not clearly favor one framework over the other because each framework has access to systematic explanatory strategies to deal with each of them (Chapter 3). The second main contribution of the thesis to the debate was the development of an important new problem — *the representation selection problem* — that offers the prospect of breaking the impasse between these two frameworks by introducing a new type of explanatory consideration that both frameworks must address (Chapter 4). I argued that VM theories do not have the resources to solve that problem, and that GM theories cannot address the problem as they stand. But the third main contribution was the development of a new GM account that can provide a satisfying solution to this problem (Chapter 5). The account specifies a second module involved in reorientation tasks, and it sketches this second module's internal principles as well as its interactions with the geometric module as understood in classic GM models (Gallistel, 1990; Lee & Spelke, 2010a). I finally pointed out that a similar move on the part of VM theorists of positing a separate selection mechanism operating on multiple types of cues would force VM theorists to endorse a hybrid account with severe costs. Overall, this shows that the GM framework is superior to the VM framework.

So, together, the core chapters of this thesis, Chapters 2 to 5, addressed the two most pressing challenges against the GM framework. At the risking of stating the painfully obvious, the arguments put forward in this thesis appeal to explanatory considerations about representation selection at many crucial junctures. Issues of representation selection had not been developed or addressed in any explicit form in the reorientation literature before. Nearly all theoretical and empirical research pertaining to the literature until now has focused on accounting for very specific aspects of search behavior in reorientation tasks with the eventual hope of giving precise quantitative modeling for things such as the exact path taken to the goal location in a given trial, the exact percentage of digs in each of the corners of an experimental enclosure. Researchers have thus ignored issues of representation selection — by, in many cases, simply assuming that a representation from the current environment has been selected at the outset of a trial. This constitutes a major explanatory gap in all current models of reorientation on the market, regardless of which larger explanatory framework they belong to. And I hope I have made a convincing case that addressing these issues head-on can transform the polemic about which explanatory framework of reorientation we should favor.

Now, I claimed in Chapter 1 that, if the GM framework is true, it allows to make important progress on two foundational debates in philosophy of psychology by supporting rather radical positions on each of them. Do some of the novel ideas put forward in this thesis to defend the GM framework force us to reassess the type of support that the GM framework provide for these specific positions? I think they do, in that they show that the GM framework provides even *stronger* support for these positions than what philosophers and psychologists usually envisage.

Consider first the debate about the nature of mental representations. I mentioned in Chapter 1 that the GM framework bolsters the view that human and non-human minds contain mental representations that encode highly abstract information about the world. It does so, I pointed out, because it entails that various species across the animal kingdom, including humans, possess representations that encode some metric properties and relations of salient surfaces in visited environments as such. Metric properties and relations are very far from low-level sensory properties. But the arguments proposed in Chapter 5 suggest that a GM account properly developed to deal with the representation selection problem should go one step further. It should posit the existence of *highly detailed* geometric representations of visited environments. The resulting representations' content goes far beyond the content that some extant GM models attribute to subjects' geometric representations.

To see why, note that some extant GM theories are quite minimalist regarding the content of geometric representations. Lee, Sovrano and Spelke's (2012) model, for instance, hold that geometric representations explicitly encode the distance and sense properties of the goal location related to three-dimensional extended surfaces, nothing more — where *sense* pertains to the distinction between left and right. On their view, it could be that what is encoded in a given geometric representation of a rectangular enclosure is merely the fact *that the goal is at the corner on the right of an extended surface which is farthest from the center of the environment*.

The new type of GM account that I proposed in Chapter 5 presupposes the existence of representations that encode much more complex metric properties and relations than that. It holds that subjects' geometric representations encode shape parameters along with detailed localized information about the global geometry of three-dimensional extended surfaces. On the one hand, shape parameters are extracted from the global geometry of the three-dimensional extended surfaces as seen from an overhead perspective. They are not present in the physical environment itself. They constitute highly abstract, summary properties of the distribution of three-dimensional extended surfaces present in the environment. On the other hand, storing localized information about surfaces in a given representation requires an explicit reference frame, which in turn entails the encoding of the distance between every pair of represented localized elements in a map-like way. Hence, the ideas put forward in Chapter 5 provide even stronger support for the view that human and non-human animals' representations encode highly abstract information about the world.

Now turn to the debate about the structure of the mind. I said in Chapter 1 that the GM framework constitutes one of the best empirically motivated cases for a specific module in higher cognition, and thus buttresses one of the most controversial aspects of the massive modularity thesis, namely the commitment to the existence of a variety of modules in higher cognition. I then argued in Chapter 5 that careful consideration of the representation selection problem from the perspective of the GM framework requires positing a new cognitive component distinct from the geometric module which is likely modular too: the selection mechanism. The resulting account, the GM+SM model, provides even stronger support the massive modularity hypothesis by suggesting the existence of *two* domain-specific and encapsulated mechanisms in higher cognition, rather than just one.

The views developed in this thesis also strengthen the support that the GM framework is seen as offering for specific positions in other debates in philosophy of psychology. I cannot make this case at length here. But consider the nativism/empiricism debate for a salient example. As we saw in Chapter 1, humans and non-human species choose the diagonally adequate corners in a rectangular enclosure significantly above chance even at a very early age (at 18 months in humans, Hermer & Spelke, 1996). We noted that this bolsters the case for nativism because, given the complexity and idiosyncrasy of the function performed by the geometric module on standard GM accounts, it appears highly unlikely that subjects could acquire the geometric module by an early age from domain-general learning mechanisms. But now the arguments proposed in this thesis further entail that, for the geometric module to work properly, a further component must operate properly as well: the selection mechanism. The geometric module cannot do its job if it does not receive the relevant representation from the selection mechanism at the outset of a trial. Yet, the selection mechanism's function and internal principles seem just as complex and idiosyncratic as the geometric module's. What are the chances that reorientation subjects reliably acquire both the geometric module and the selection mechanism through domain-general learning mechanisms by the 18-month mark (in humans) or the 3-day mark (in domestic chicks)? Thus, the new type of model defended in Chapter 5 significantly strengthens the case for nativism that originated from standard GM accounts.

In sum, the type of GM model developed in this thesis offers even stronger support than typical GM models for relatively radical positions in several debates in philosophy of psychology. So, if my proposals are right, they have important philosophical implications.

Let me finish by highlighting some issues that would warrant further investigation in order to make the strongest possible case for the GM framework, and to see its implications for philosophical debates more clearly. The first and perhaps most important issue is just figuring out more precisely how the selection mechanism works, and the exact content of geometric representations received by the geometric module from the selection mechanism. I have made proposals with respect to this issue in Chapter 5, but they are tentative. And even if they are right, they are very much the beginning of the inquiry. Determining the exact content of geometric representations will require more empirical work from the type of multiple-enclosure reorientation tasks proposed in Chapter 5 to test the predictions of the new type of GM model that I put forward, as well as global-remapping studies of the sort developed by Wills et al. (2005) and Colgin et al. (2010) that played an important role in the analysis of Chapters 4 and 5. Moreover, I have identified many important questions in Chapter 2 in relation to the selection mechanism that deserve more theoretical scrutiny, such as the following: What happens when the selection mechanism cannot find a single geometric representation in memory whose content fits the current environment well enough to be selected? What kind of alternative navigation processes take control of subjects' search behavior in that case?

A second important issue is to make a case that the GM framework is better than other explanatory frameworks besides the VM framework. The two best known alternative frameworks are *the adaptive-combination framework* (Newcombe & Ratliff, 2007; Xu, Regier, & Newcombe, 2017) and *the associative framework* (Miller & Shettleworth, 2007, 2013). Issues of representation selection do not arise in as dramatic a form for them as they do for the VM framework because both frameworks can in principle accept that subjects store a detailed global geometric representation and a corresponding feature map for each visited environment (or even a type of compound representation which includes both detailed geometric information and featural information). In other words, because they do not have to deal with the severe representational constraints of VM models, theories belonging to either of these two frameworks seem less vulnerable to the representation selection problem. But I do think that both frameworks fail significantly in other regards, and I hope to make this case explicitly in future research.

A third issue is to provide an adequate answer to an influential objection against the GM framework that arises from the so-called *size effect*. It turns out that subjects from various species make relatively few rotational errors when there are distinctive featural cues close to the goal location in *large* rectangular enclosures. In large rectangular enclosures, as opposed to small ones, subjects search for the goal at the correct location significantly more often than at any other location in most contexts (Learmonth, Newcombe, Sheridan, & Jones, 2008; Sovrano, Bisazza, & Vallortigara, 2007; though see Lee & Vallortigara, 2015). This is the size effect. Many reorientation researchers hold that the size effect causes a major problem for the GM framework because, on the typical evolutionary rationale given as a motivation for the GM framework (Gallistel, 1990; see also Chapter 5:3), the geometric module would not be expected to guide the agent's behavior only in small spaces. Geometric cues are just as useful for guiding search behavior through estimations of the subject's heading and location in large spaces than in small spaces. So, we should expect the module to be active and perform the same function in both large and small environments.

I cannot give a full response to this objection here, but I will just note the following: that I strongly advise against biting the bullet by granting that the geometric module is only active in small environments. GM theorists should look for other explanations of the size effect. Here might be the beginning of one. Many important experiments that provide evidence of the existence of the size effect are in fact transformational tasks (e.g., Chiandetti, Regolin, Sovrano, & Vallortigara, 2007; Ratliff & Newcombe, 2008b).

In the relevant studies, experimenters make important changes to the geometry of the enclosure and/or their featural cues following a training phase. However, such transformational tasks may well lead to the phenomenon of selection interference that I outlined in Chapter 2 as deeply problematic when drawing on the results of transformational tasks to argue against the GM framework. So, the appeal to the size effect may not remain as a powerful objection against the GM framework after such experiments have been reinterpreted in light of the selection interference strategy developed in Chapter 2.

References

Anderson, M. L. (2010). Neural reuse: a fundamental organizational principle of the brain. *The Behavioral and Brain Sciences*, *33*(4), 245-266.

Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. Cambridge, MA: MIT Press.

Barry, C., & Burgess, N. (2007). Learning in a geometric model of place cell firing. *Hippocampus, 17*(9), 786–800.

Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffery, K. J., Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences, 17*(1–2), 71–97.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*(4), 577-609.

Bek, J., Blades, M., Siegal, M., & Varley, R. (2010). Language and spatial Reorientation: Evidence from severe aphasia. *Journal of Experimental Psychology: Learning Memory and Cognition*, *36*(3), 646–658.

Berkeley, G. (1709). Essay on a new theory of vision. In D. Clarke (ed.), *Berkeley: Philosophical Writings.* Cambridge: Cambridge University Press, 2008.

Buller, D. (2005). *Adapting minds: Evolutionary psychology and the persistent quest for human nature*. Cambridge, MA: MIT Press.

Burgess, N. (2006). Spatial memory: how egocentric and allocentric combine. *Trends in Cognitive Sciences, 10*(12), 551–557.

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: A neural model of spatial memory and imagery. *Psychological Review, 114*(2), 340–375.

Carruthers, P. (2002). The cognitive functions of language. *The Behavioral and Brain Sciences*, *25*(6), 657-674.

Carruthers, P. (2004). Practical reasoning in a modular mind. *Mind and Language*, *19*(3), 259–278.

Carruthers, P. (2006a). *The architecture of the mind*. Oxford: Oxford University Press.

Carruthers, P. (2006b). The case for massively modular models of mind. In R. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 1–21). Oxford: Blackwell Publishing.

Carruthers, P. (2013). Mindreading in infancy. *Mind and Language*, *28*(2), 141–172.

Carruthers, P. (2015). *The centered mind: What the science of working memory shows us about the nature of human thought*. Oxford: Oxford University Press.

Cartwright, B. A., & Collett, T. S. (1987). Landmark maps for honeybees. *Biological*

Cybernetics, *57*(1–2), 85–93.

Cheeseman, J. F., Millar, C. D., Greggers, U., Lehmann, K., Pawley, M. D. M., Gallistel, C. R., Warman, G. R., Menzel, R. (2014). Way-finding in displaced clock-shifted bees proves bees use a cognitive map. *Proceedings of the National Academy of Sciences*, *111*(24), 8949–8954.

Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, *23*(2), 149–178.

Cheng, K. (2005). Reflections on geometry and navigation. *Connection Science*, *17*(1–2), 5–21.

Cheng, K. (2008). Whither geometry? Troubles of the geometric module. *Trends in Cognitive Sciences*, *12*(9), 355–361.

Cheng, K., & Gallistel, C. R. (2005). Shape parameters explain data from spatial transformations: Comment on Pearce et al. (2004) and Tommasi & Polli (2004). *Journal of Experimental Psychology: Animal Behavior Processes*, *31*(2), 254–259.

Cheng, K., Huttenlocher, J., & Newcombe, N. S. (2013). 25 years of research on the use of geometry in spatial reorientation: A current theoretical perspective. *Psychonomic Bulletin & Review*, *20*(6), 1033–1054.

Cheng, K., & Newcombe, N. S. (2005). Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychonomic Bulletin & Review*, *12*(1), 1–23.

Cheung, A., Stürzl, W., Zeil, J., & Cheng, K. (2008). The information content of panoramic images II: View-based navigation in nonrectangular experimental arenas. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(1), 15–30.

Chiandetti, C., Regolin, L., Sovrano, V. A., & Vallortigara, G. (2007). Spatial reorientation: The effects of space size on the encoding of landmark and geometry information. *Animal Cognition*, *10*(2), 159–168.

Chiandetti, C., Spelke, E. S., & Vallortigara, G. (2015). Inexperienced newborn chicks use geometry to spontaneously reorient to an artificial social partner. *Developmental Science*, *18*(6), 972–978.

Chiandetti, C., & Vallortigara, G. (2008). Is there an innate geometric module? Effects of experience with angular geometric cues on spatial re-orientation based on the shape of the environment. *Animal Cognition*, *11*(1), 139–146.

Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, *78*(2), 67–90.

Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.

Colgin, L. L., Leutgeb, S., Jezek, K., Leutgeb, J. K., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2010). Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *Journal of Neurophysiology, 104*(1), 35–50.

Colombo, M. (2013). Moving forward (and beyond) the modularity debate: A network perspective. *Philosophy of Science, 80*(3), 356–377.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition, 31*, 187–276.

Cosmides, L., Barrett, H. C., & Tooby, J. (2010). Adaptive specializations, social exchange, and the evolution of human intelligence. *Proceedings of the National Academy of Sciences, 107*(Supplement 2), 9007–9014.

Cruse, H., & Wehner, R. (2011). No need for a cognitive map: Decentralized memory for insect navigation. *PLoS Computational Biology*, *7*(3), e1002009.

Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.

Descartes, R. (1637/2001). The optics. In P. J. Olscamp (Ed. and Trans.), *Discourse on method, optics, geometry and meteorology.* Indianapolis, In: Hackett.

Dittmar, L., Stürzl, W., Baird, E., Boeddeker, N., & Egelhaaf, M. (2010). Goal seeking in honeybees: Matching of optic flow snapshots? *Journal of Experimental Biology, 213*(17), 2913–2923.

Dittmar, L., Stürzl, W., Jetzschke, S., Mertes, M., & Boeddeker, N. (2014). Out of the box: How bees orient in an ambiguous environment. *Animal Behaviour, 89,* 13–21.

Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature, 463*(7281), 657–661.

Doeller, C. F., & Burgess, N. (2008). Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. *Proceedings of the National Academy of Sciences, 105*(15), 5909–5914.

Doeller, C. F., King, J. A., & Burgess, N. (2008). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences, 105*(15), 5915–5920.

Durlach, P. J., & Rescorla, R. A. (1980). Potentiation rather than overshadowing in flavor-aversion learning: an analysis in terms of within-compound associations. *Journal of Experimental Psychology. Animal Behavior Processes*, *6*(2), 175–187.

Egan, F. (2011). Representationalism. In E. Margolis, R. Samuels, & S. P. Stich (Eds.), *Oxford handbook of philosophy of cognitive science* (pp. 250–272). New York: Oxford University Press.

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature, 425*(6954), 184–188.

Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience, 20*(11), 1504–1513.

Epstein, R. A., & Vass, L. K. (2014). Neural systems for landmark-based wayfinding in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences, 369*(1635), 20120533.

Etienne, A. S., & Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus, 14*(2), 180–192.

Ferrara, K., & Park, S. (2016). Neural representation of scene boundaries. *Neuropsychologia, 89*, 180–190.

Fodor, J. A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.

Fodor, J. A. (1981). *Representations: Philosophical essays on the foundations of cognitive science*. Cambridge, MA: MIT Press.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Fodor, J. A. (2000). *The mind doesn't work that way*. Cambridge, MA: MIT Press.

Fodor, J. A. (2008). *LOT 2: The language of thought revisited*. New York: Oxford University Press.

Franz, M. O., Schölkopf, B., Mallot, H. A., & Bülthoff, H. H. (1998). Learning view graphs for robot navigation. *Autonomous Robots, 5*(1), 111–125.

Fuller, T., & Samuels, R. (2014). Scientific inference and ordinary cognition: Fodor on holism and cognitive architecture. *Mind and Language*, *29*(2), 201–237.

Gallistel, C. R. (1980). *The organization of action*. Hillsdale, NJ: Lawrence Erlbaum.

Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.

Gallistel, C. R., & Matzel, L. D. (2013). The neuroscience of learning: Beyond the Hebbian synapse. *Annual Review of Psychology*, *64*, 169–200.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Gillner, S., Weiß, A. M., & Mallot, H. A. (2008). Visual homing in the absence of feature-based landmark information. *Cognition*, *109*(1), 105–122.

Gouteux, S., Thinus-Blanc, C., & Vauclair, J. (2001). Rhesus monkeys use geometric and nongeometric information during a reorientation task. *Journal of Experimental Psychology: General*, *130*(3), 505–519.

Graham, M., Good, M. A., McGregor, A., & Pearce, J. M. (2006). Spatial learning based on the shape of the environment is influenced by properties of the objects forming the shape. *Journal of Experimental Psychology: Animal Behavior Processes*, *32*(1), 44–59.

Gray, E. R., Bloomfield, L. L., Ferrey, A., Spetch, M. L., & Sturdy, C. B. (2005).

Spatial encoding in mountain chickadees: features overshadow geometry. *Biology Letters*, *1*(May), 314–317.

Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.

Hardcastle, K., Ganguli, S., & Giocomo, L. M. (2015). Environmental boundaries as an error correction mechanism for grid cells. *Neuron*, *86*(3), 827–839.

Hermer, L., & Spelke, E. (1996). Modularity and development: The case of spatial reorientation. *Cognition*, *61*(3), 195–232.

Hermer-Vazquez, L., Moffet, A., & Munkholm, P. (2001). Language, space, and the development of cognitive flexibility in humans: The case of two spatial memory tasks. *Cognition*, *79*(3), 263–299.

Hermer-Vazquez, L., Spelke, E. S., & Katsnelson, A. S. (1999). Sources of flexibility in human cognition: dual-task studies of space and language. *Cognitive Psychology*, *39*(1), 3–36.

Horne, M. R., Iordanova, M. D., Albasser, M. M., Aggleton, J. P., Honey, R. C., & Pearce, J. M. (2010). Lesions of the perirhinal cortex do not impair integration of visual and geometric information in rats. *Behavioral Neuroscience, 124*(3), 311–320.

Janzen, G., & van Turennout, M. (2004). Selective neural representation of objects relevant for navigation. *Nature Neuroscience, 7*(6), 673–677.

Jeewajee, A., Lever, C., Burton, S., O’Keefe, J., & Burgess, N. (2008). Environmental novelty is signaled by reduction of the hippocampal theta frequency. *Hippocampus*, *18*(4), 340–348.

Jeffery, K. J. (2010). Theoretical accounts of spatial learning: A neurobiological view. *Quarterly Journal of Experimental Psychology, 63*(9), 1683–1699.

Jeffery, K. J., Anderson, M. I., Hayman, R., & Chakraborty, S. (2004). A proposed architecture for the neural representation of spatial context. *Neuroscience and Biobehavioral Reviews, 28*(2), 201–218.

Jeffery, K. J., Donnett, J. G., Burgess, N., & O’Keefe, J. M. (1997). Directional control of hippocampal place fields. *Experimental Brain Research*, *117*(1), 131–142.

Julian, J. B., Keinath, A. T., Muzzio, I. A., & Epstein, R. A. (2015). Place recognition and heading retrieval are mediated by dissociable cognitive systems in mice. *Proceedings of the National Academy of Sciences*, *112*(20), 6503–6508.

Kant, I. (1787/1998). *Critique of pure reason* (P. Guyer and A. Wood, Trans.). Cambridge, MA: Cambridge University Press.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience, 17*(11), 4302–4311.

Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *361*(1476), 2109–2128.

Keinath, A. T., Julian, J. B., Epstein, R. A., & Muzzio, I. A. (2017). Environmental geometry aligns the hippocampal map during spatial reorientation. *Current Biology*, *27*(3), 309–317.

Kelly, D. M., Chiandetti, C., & Vallortigara, G. (2011). Re-orienting in space: Do animals use global or local geometry strategies? *Biology Letters*, *7*(3), 372–375.

Krupic, J., Bauza, M., Burton, S., Barry, C., & O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature*, *518*(7538), 232–235.

Krupic, J., Bauza, M., Burton, S., Lever, C., & O'Keefe, J. (2014). How environment geometry affects grid cell symmetry and what we can learn from it. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1635), 20130188.

Laurence, S., & Margolis, E. (2015). Concept nativism and neural plasticity. In E. Margolis, & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 117-147). Cambridge, MA: MIT Press.

Learmonth, A. E., Newcombe, N. S., Sheridan, N., & Jones, M. (2008). Why size counts: Children's spatial reorientation in large and small enclosures. *Developmental Science*, *11*(3), 414–426.

Lee, S. A., Ferrari, A., Vallortigara, G., & Sovrano, V. A. (2015). Boundary primacy in spatial mapping: Evidence from zebrafish (Danio rerio). *Behavioural Processes*, *119*, 116–122.

Lee, S. A., Shusterman, A., & Spelke, E. S. (2006). Reorientation and landmark-guided search by young children: Evidence for two systems. *Psychological Science*, *17*(7), 577–582.

Lee, S. A., Sovrano, V. A., & Spelke, E. S. (2012). Navigation as a source of geometric knowledge: Young children's use of length, angle, distance, and direction in a reorientation task. *Cognition*, *123*(1), 144–161.

Lee, S. A., & Spelke, E. S. (2010a). A modular geometric mechanism for reorientation in children. *Cognitive Psychology*, *61*(2), 152–176.

Lee, S. A., & Spelke, E. S. (2010b). Two systems of spatial representation underlying navigation. *Experimental Brain Research*, *206*(2), 179–188.

Lee, S. A., & Spelke, E. S. (2011). Young children reorient by computing layout geometry, not by matching images of the environment. *Psychonomic Bulletin & Review*, *18*(1), 192–198.

Lee, S. A., Spelke, E. S., & Vallortigara, G. (2012). Chicks, like children, spontaneously reorient by three-dimensional environmental geometry, not by image matching. *Biology Letters*, *8*(4), 492–494.

Lee, S. A., Tucci, V., Sovrano, V. A., & Vallortigara, G. (2015). Working memory and reference memory tests of spatial navigation in mice (Mus musculus). *Journal of*

Comparative Psychology, 129(2), 189–197.

Lee, S. A., Tucci, V., & Vallortigara, G. (2017). Spatial impairment and memory in genetic disorders: Insights from mouse models. *Brain Sciences, 7*(2), 17.

Lee, S. A., & Vallortigara, G. (2015). Bumblebees spontaneously map location of conspecific using geometry and features. *Learning and Motivation, 50*, 32–38.

Lee, S. A., Vallortigara, G., Ruga, V., & Sovrano, V. A. (2012). Independent effects of geometry and landmark in a spontaneous reorientation task: A study of two species of fish. *Animal Cognition, 15*(5), 861–870.

Lee, S. A., Winkler-Rhoades, N., & Spelke, E. S. (2012). Spontaneous reorientation is guided by perceived surface distance, not by image matching or comparison. *PLoS ONE, 7*(12), e51373.

Leslie, A. M. (1994). ToMM, ToBy, and Agency: Core architecture and domain specificity. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 119–148). Cambridge: Cambridge University Press.

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience, 29*(31), 9771–9777.

Lever, C., Wills, T., Cacucci, F., Burgess, N., & O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature, 416*(6876), 90–94.

Mallot, H. A., & Lancier, S. (2018). Place recognition from distant landmarks: Human performance and maximum likelihood model. *Biological Cybernetics, 112*(4), 291–303.

Margolis, E., & Laurence, S. (2013). In defense of nativism. *Philosophical Studies, 165*(2), 693–718.

Marr, D. (1982). *Vision: A computational view*. San Francisco: Freeman Press.

Mayer, U., Bhushan, R., Vallortigara, G., & Lee, S. A. (2018). Representation of environmental shape in the hippocampus of domestic chicks (Gallus gallus). *Brain Structure and Function, 223*(2), 941–953.

McDonald, R. J., & White, N. M. (1994). Parallel information processing in the water maze: Evidence for independent memory systems involving dorsal striatum and hippocampus. *Behavioral and Neural Biology, 61*, 260–270.

Menzel, E. W. (1973). Chimpanzee spatial memory organization. *Science, 182*(4115), 943–945.

Miller, N. Y., & Shettleworth, S. J. (2007). Learning about environmental geometry: an associative model. *Journal of Experimental Psychology. Animal Behavior Processes, 33*(3), 191–212.

Miller, N. Y., & Shettleworth, S. J. (2013). Associative models of instrumental learning:

A response to Dupuis and Dawson. *Journal of Experimental Psychology. Animal Behavior Processes*, *39*(3), 287–293.

Mittelstaedt, M. L., & Mittelstaedt, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften*, *67*(11), 566–567.

Möller, R., & Vardy, A. (2006). Local visual homing by matched-filter descent in image distances. *Biological Cybernetics*, *95*(5), 413–430.

Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, *12*(2), 239–260.

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, *7*(7), 1951–1968.

Newcombe, N. S., & Ratliff, K. R. (2007). Explaining the development of spatial reorientation: Modularity-plus-language versus the emergence of adaptive combination. In J. Plumer, & J. Spencer (Eds.), *The emerging spatial mind: Mapping the nature of change in spatial cognitive development* (pp. 53–76). New York: Oxford University Press.

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

Papi, F. (1992). *Animal homing*. London: Chapman & Hall.

Park, S., Brady, T. F., Greene, M. R., & Oliva, A. (2011). Disentangling scene content from spatial boundary: Complementary roles for the parahippocampal place area and lateral occipital complex in representing real-world scenes. *Journal of Neuroscience*, *31*(4), 1333–1340.

Pearce, J. M. (2009). The 36th Sir Frederick Bartlett Lecture: An associative analysis of spatial learning. *Quarterly Journal of Experimental Psychology*, *62*(9), 1665–1684.

Pearce, J. M., Good, M. A., Jones, P. M., & McGregor, A. (2004). Transfer of spatial behavior between different environments: Implications for theories of spatial learning and for the role of the hippocampus in spatial learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *30*(2), 135–147.

Pearce, J. M., Graham, M., Good, M. A., Jones, P. M., & McGregor, A. (2006). Potentiation, overshadowing, and blocking of spatial learning based on the shape of the environment. *Journal of Experimental Psychology: Animal Behavior Processes*, *32*(3), 201–214.

Pecchia, T., Gagliardo, A., & Vallortigara, G. (2011). Stable panoramic views facilitate snap-shot like memories for spatial reorientation in homing pigeons. *PLoS ONE*, *6*(7), e22657.

Pecchia, T., & Vallortigara, G. (2010). View-based strategy for reorientation by geometry. *Journal of Experimental Biology*, *213*(17), 2987–2996.

Pecchia, T., & Vallortigara, G. (2012). Spatial reorientation by geometry with freestanding objects and extended surfaces: A unifying view. *Proceedings of the Royal Society B: Biological Sciences, 279*(1736), 2228–2236.

Poulter, S., Hartley, T., & Lever, C. (2018). The neurobiology of mammalian navigation. *Current Biology*, *28*(17), R1023–R1042.

Prinz, J. (2002). *Furnishing the mind: concepts and their perceptual basis*. Cambridge, MA: MIT Press.

Prinz, J. J. (2006). Is the mind really modular? In R. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 22–36). Oxford: Blackwell Publishing.

Pyers, J. E., Shusterman, A., Senghas, A., Spelke, E. S., & Emmorey, K. (2010). Evidence from an emerging sign language reveals that language supports spatial cognition. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(27), 12116–12120.

Pylyshyn, Z. W. (2003). *Seeing and visualizing: It's not what you think*. Cambridge, MA: MIT Press.

Quirk, G. J., Muller, R. U., & Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience, 10*(6), 2008–2017.

Ratliff, K. R., & Newcombe, N. S. (2008a). Is language necessary for human spatial reorientation? Reconsidering evidence from dual task paradigms. *Cognitive Psychology*, *56*(2), 142–163.

Ratliff, K. R., & Newcombe, N. S. (2008b). Reorienting when cues conflict: Evidence for an adaptive-combination view. *Psychological Science*, *19*(12), 1301–1307.

Raudies, F., & Hasselmo, M. E. (2012). Modeling boundary vector cell firing given optic flow as a cue. *PLoS Computational Biology*, *8*(6), e1002553.

Rescorla, M. (2009). Cognitive maps and the language of thought. *British Journal for the Philosophy of Science*, *60*(2), 377–407.

Samuels, R. (2006). Is the human mind massively modular? In R. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 37–56). Oxford: Blackwell Publishing.

Samuels, R. (2011). Massive modularity. In E. Margolis, R. Samuels, & S. P. Stich (Eds.), *Oxford handbook of philosophy of cognitive science* (pp. 60–91). New York: Oxford University Press.

Save, E., Cressant, A., Thinus-Blanc, C., & Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. *The Journal of Neuroscience*, *18*(5), 1818–1826.

Save, E., Nerad, L., & Poucet, B. (2000). Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus*, *10*(1), 64–76.

Shanahan, M., & Baars, B. (2005). Applying global workspace theory to the frame problem. *Cognition*, *98*(2), 157–176.

Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., & Gerstner, W. (2009). Is there a geometric module for spatial orientation? Insights from a rodent navigation model. *Psychological Review, 116*(3), 540–566.

Shusterman, A., Lee, S. A., & Spelke, E. S. (2011). Cognitive effects of language on human navigation. *Cognition*, *120*(2), 186–201.

Shusterman, A., & Spelke, E. S. (2005). Language and the development of spatial reasoning. In P. Carruthers, S. Laurence, & S. P. Stich (Eds.), *The innate mind: structure and contents* (pp. 89–106). New York: Oxford University Press.

Sober, E. (2015). *Ockham's razors: A user's manual*. Cambridge University Press.

Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science, 322*(5909), 1865–1868.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2002). Modularity and spatial reorientation in a simple mind: Encoding of geometric and nongeometric properties of a spatial environment by fish. *Cognition*, *85*(2), 51–59.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2003). Modularity as a fish (Xenotoca eiseni) views it: Conjoining geometric and nongeometric information for spatial reorientation. *Journal of Experimental Psychology: Animal Behavior Processes, 29*(3), 199–210.

Sovrano, V. A., Bisazza, A., & Vallortigara, G. (2007). How fish do geometry in large and in small spaces. *Animal Cognition*, *10*(1), 47–54.

Sovrano, V. A., Potrich, D., & Vallortigara, G. (2013). Learning of geometry and features in bumblebees (Bombus terrestris). *Journal of Comparative Psychology*, *127*(3), 312–318.

Sovrano, V. A., & Vallortigara, G. (2006). Dissecting the geometric module: A sense linkage for metric and landmark information in animals' spatial reorientation. *Psychological Science, 17*(7), 616–621.

Spelke, E., Lee, S. A., & Izard, V. (2010). Beyond core knowledge: Natural geometry. *Cognitive Science, 34*(5), 863–884.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 39–67). Cambridge: Cambridge University Press.

Sperber, D. (2002). In defense of massive modularity. In I. Dupoux (Ed.), *Language, brain, and cognitive development: Essays in honor of Jacques Mehler* (pp. 47–57). Cambridge, MA: MIT Press.

Sterelny, K. (2003). *Thought in a hostile world*. Oxford: Blackwell Publishing.

Stewart, S., Jeewajee, A., Wills, T. J., Burgess, N., & Lever, C. (2014). Boundary coding in the rat subiculum. *Philosophical Transactions of the Royal Society B: Biological Sciences, 369*(1635), 20120514.

Sturz, B. R., Gurley, T., & Bodily, K. D. (2011). Orientation in trapezoid-shaped enclosures: Implications for theoretical accounts of geometry learning. *Journal of Experimental Psychology: Animal Behavior Processes, 37*(2), 246–253.

Stürzl, W., Cheung, A., Cheng, K., & Zeil, J. (2008). The information content of panoramic images I: The rotational errors and the similarity of views in rectangular experimental arenas. *Journal of Experimental Psychology: Animal Behavior Processes*, 34(1), 1–14.

Thompson, S. M., Berkowitz, L. E., & Clark, B. J. (2018). Behavioral and neural subsystems of rodent exploration. *Learning and Motivation*, *61*(April), 3–15.

Tolman, E. C., & Honzik, C. H. (1930). Introduction and removal of reward, and maze performance in rats. *University of California Publications in Psychology, 4*, 257–275.

Tommasi, L., Chiandetti, C., Pecchia, T., Sovrano, V. A., & Vallortigara, G. (2012). From natural geometry to spatial cognition. *Neuroscience and Biobehavioral Reviews, 36*(2), 799–824.

Tommasi, L., & Polli, C. (2004). Representation of two geometric features of the environment in the domestic chick (Gallus gallus). *Animal Cognition*, *7*(1), 53–59.

Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 19–136). New York: Oxford University Press.

Troiani, V., Stigliani, A., Smith, M. E., & Epstein, R. A. (2014). Multiple object properties drive scene-selective regions. *Cerebral Cortex*, 24(4), 883–897.

Twyman, A. D., & Newcombe, N. S. (2010). Five reasons to doubt the existence of a geometric module. *Cognitive Science*, *34*(7), 1315–1356.

Vallortigara, G. (2009). Animals as natural geometers. In L. Tommasi, M. A. Peterson, & L. Nadel (Eds.), *Cognitive Biology* (pp. 83–104). Cambridge, MA: MIT Press.

Vallortigara, G. (2017). Comparative cognition of number and space: The case of geometry and of the mental number line. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1740), 20170120.

Vallortigara, G., Zanforlin, M., & Pasti, G. (1990). Geometric modules in animals' spatial representations: A test with chicks (Gallus gallus domesticus). *Journal of Comparative Psychology*, *104*(3), 248–254.

Waller, D., Loomis, J. M., Golledge, R. G., & Beall, A. C. (2000). Place learning in humans: The role of distance and direction information. *Spatial Cognition and Computation, 2*(4), 333–354.

Wang, R. F., & Spelke, E. S. (2002). Human spatial representation: Insights from animals. *Trends in Cognitive Sciences, 6*(9), 376–382.

Wang, R. F., & Spelke, E. S. (2003). Comparative approaches to human navigation. In

K. J. Jeffery (Ed.), *The neurobiology of spatial behaviour* (pp. 119–143). Oxford: Oxford University Press.

Warren, W. H. (2012). Does this computational theory solve the right problem? Marr, Gibson, and the goal of vision. *Perception, 41,* 1053–1060.

Widloski, J., & Fiete, I. (2014). How does the brain solve the computational problems of spatial navigation? In D. Derdikman & J. J. Knierim (Eds.), *Space, time and memory in the hippocampal formation* (pp. 373–407). Vienna: Springer Verlag.

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science, 308*(5723), 873–6.

Wystrach, A. (2009). Ants in rectangular arenas. *Communicative & Integrative Biology, 2*(5), 388–390.

Wystrach, A., & Beugnon, G. (2009). Ants learn geometry and features. *Current Biology, 19*(1), 61–66.

Wystrach, A., Cheng, K., Sosa, S., & Beugnon, G. (2011). Geometry, features, and panoramic views: Ants in rectangular arenas. *Journal of Experimental Psychology: Animal Behavior Processes, 37*(4), 420–435.

Wystrach, A., & Graham, P. (2012). View-based matching can be more than image matching: The importance of considering an animal's perspective. *I-Perception, 3*(8), 547–549.

Xu, Y., Regier, T., & Newcombe, N. S. (2017). An adaptive cue combination model of human spatial reorientation. *Cognition, 163*, 56–66.

Zhang, S., Schonfeld, F., Wiskott, L., & Manahan-Vaughan, D. (2014). Spatial representations of place cells in darkness are supported by path integration and border information. *Frontiers in Behavioral Neuroscience, 8*(June), 1–12.