Cognitive Maps and the Language of Thought

Michael Rescorla

Abstract: Fodor advocates a view of cognitive processes as computations defined over *the language of thought* (or *Mentalese*). Even among those who endorse Mentalese, considerable controversy surrounds its representational format. What semantically relevant structure should scientific psychology attribute to Mentalese symbols? Researchers commonly emphasize *logical* structure, akin to that displayed by predicate calculus sentences. To counteract this tendency, I discuss computational models of navigation drawn from *probabilistic robotics*. These models involve computations defined over *cognitive maps*, which have *geometric* rather than *logical* structure. They thereby demonstrate the possibility of rational cognitive processes in an exclusively non-logical representational medium. Furthermore, they offer much promise for the empirical study of animal navigation.

1 Mental representations

Fodor ([1975], [1987]) and Pylyshyn ([1984, [2003]) espouse a theory of cognition based on two doctrines:

(1) Certain core mental processes studied by scientific psychology are mechanical, rule-governed operations upon symbols. In that sense, the processes are *computational*.

(2) The symbols that figure in computational mental activity have syntactic structure and a compositional semantics.

Both doctrines are popular albeit controversial within philosophy and psychology. Following Fodor, philosophers typically refer to the representational system posited by (2) as the *language of thought*, or *Mentalese*. Even among those who endorse Mentalese, considerable controversy surrounds its representational format. What semantically relevant structure should scientific psychology attribute to Mentalese symbols? How closely do such symbols resemble familiar concrete representations like sentences, pictures, diagrams, or maps?

An extreme view, tracing back at least to William of Ockham, holds that all mental representations operate analogously to sentences. Modern exponents often emphasize the sentential structures studied by formal logic. Many AI researchers, including Genesereth and Nilsson ([1987]) and McCarthy and Hayes ([1969]), pursue a “logicist” agenda that treats the predicate calculus, or a suitably supplemented variant of it, as the primary, paradigmatic, or even exclusive medium of thought. At the opposite extreme, some commentators hold that all mental representation operates pictorially, diagrammatically, or cartographically. This “pictorialist” view, popular among medieval philosophers and the British empiricists, finds such recent advocates as Armstrong ([1973]), Barsalou ([1998]), Braddon-Mitchell and Jackson ([2007]), and Cummins ([1996]). Between the extremes of logicism and pictorialism lies a pluralistic position that embraces both logical and non-logical mental representations, assigning neither explanatory primacy over the other. Johnson-Laird ([2004], p. 187), McDermott ([2001], p. 69), Pinker ([2005], p. 7), Sloman ([1978], pp. 144-76), and many others advocate this pluralistic position. Although Fodor's emphasis upon the "languagelike" character of Mentalese might seem to suggest logicism, he inclines more towards pluralism ([2007], pp. 105-16).

My goal is to clarify the pluralistic viewpoint through detailed philosophical analysis of a particularly instructive case study. Two principal challenges face the pluralistic conception: to provide compelling examples of non-logical mental representation, and to explain how such representations differ in a principled way from those patterned after formal logic. To meet these challenges, I will discuss some models of navigation drawn from psychology and robotics. The models posit representations, *cognitive maps*, whose structure is geometric rather than logical. As I will argue, cognitive maps offer key advantages over the putative examples of non-logical representation more commonly studied by philosophers, such as perception and mental imagery.[1]

The computational models I discuss in §4 are thoroughly "cognitivist," without any hint of behaviorism, associationism, Gibsonianism, or connectionism. Specifically, the models embody a commitment to (1)-(2). Thus, they enshrine the "classical" conception of cognition as rule-governed symbol manipulation. From a connectionist or dynamical systems perspective, the contrast between logicist, pictorialist, and pluralistic theories may seem trifling. From *within* the classical conception, however, the contrast matters a great deal. Logicism would have us ignore an important class of promising computational models. Even from a connectionist or dynamical systems perspective, we require a suitably general understanding of the classical conception so as to assess its strengths and weaknesses.

2 Mental imagery, perception, and cognitive maps

Recent discussion of mental imagery focuses on a series of experimental results due to Shepard and Kosslyn, along with various collaborators. Shepard and Chipman ([1970]) and Kosslyn ([1980]) argue that we can best explain these results by positing an imagistic medium of mental representation. Dennett ([1981]) and Pylyshyn ([1984], [2003]) disagree. For analysis of the debate, see (Block, [1983]; Grush [2004], pp. 393-84; Thomas [2007]; Tye [1991]).

Even overlooking that the imagery debate seems no closer to resolution than it was two decades ago, there are several reasons why studying cognitive maps rather than mental images may yield philosophical dividends. First, evidence for mental imagery depends largely (though not entirely) upon linguistic interactions through which experimenters instruct subjects to perform certain cognitive tasks. This evidence does not readily generalize to non-linguistic creatures. In contrast, overwhelming evidence indicates that even insects perform sophisticated navigational feats. Thus, navigational models enjoy wider applicability than models of mental imagery. Second, navigation is more psychologically fundamental than mental imagery. It is vital for survival and procreation. It is arguably central to anything resembling cognition of a physical world.[2] Third, in contrast with generating, inspecting, or manipulating a mental image, forming and updating a cognitive map is an exercise of rational cognitive faculties. It is a type of, or a lower-level analogue to, belief-fixation. As we will see, it shares many features with abductive inference in science and everyday reasoning. Fourth, we have detailed, mathematically sophisticated models of how animals might perform this particular abduction.

The final two points are especially important. Many philosophers suggest that rational cognition requires a logically structured representational medium. In this vein, Devitt writes that '[f]ormal logic gives us a very good idea of how thinking might proceed if thoughts are represented linguistically… We still have very little idea how thinking could proceed if thoughts were not language-like but, say, map-like' ([2006], pp. 146-47). Similarly, Rey holds that 'anything that is capable of rational thought is capable of making *logical transitions* in thought; i.e. it is psychologically possible that it pass from one thought to another by virtue of logical properties of its thought' ([1995], p. 203). On this basis, he argues that rational thought requires a representational medium subsuming something like the predicate calculus. Pylyshyn ([1984], pp. 195-6) argues along similar lines, albeit conjecturally and in a more empirical spirit.

Computational models of navigation answer the challenge posed by Devitt, Rey, and Pylyshyn. The basic idea behind the models I will discuss is that the subject forms and updates a cognitive map of its surroundings, a map which the subject then exploits to reach goal destinations. As we will see, the proposed mechanisms for updating and exploiting cognitive maps are rational. Yet, at least on the surface, the models do not display the familiar hallmarks of logical form: sentential logical connectives, quantifiers, or even predication. The models thereby provide an existence proof that rational cognitive processes can occur in an exclusively non-logical representational medium.

Besides mental imagery, the most widely discussed putative example of "non-discursive" mental representation is *perception*. Beginning with Evans ([1982]) and Peacocke ([1992]), many philosophers have argued that perceptual experiences have *non-conceptual content*, as opposed to the *conceptual content* exhibited by beliefs and desires. McDowell ([1995]) attacks such arguments, as do many other philosophers.

There are several reasons for shifting attention from perception to cognitive maps. With a few exceptions, such as (Bermudez [1998]; Burge [2003], [2005]; Fodor [2007]; Raftopoulos and Müller [2006]), the voluminous philosophical literature on non-conceptual perceptual content tends to downplay scientific research into perception. Again excluding (Fodor [2007]), debate generally emphasizes *content* rather than the *vehicles* of content. Participants in the debate seldom even mention mental representations. Thus, it is unclear how to bring the debate into contact with our main question: what semantically relevant structure should scientific psychology attribute to Mentalese? Finally, and most importantly, even if we were to conclude that perception involves non-logical mental representations, the same might not apply to central cognitive processes such as belief-fixation and decision-making. Admittedly, the boundary between perception and belief-fixation is vexed. Moreover, a satisfying theory of perception will doubtless treat it as abductive and thus somewhat analogous to belief-fixation. Nevertheless, excessive focus on perception fosters the impression that non-logical mental representations figure mainly in "input" processes. By shifting attention to cognitive maps, I seek to dispel that impression.

3 Cognitive maps in psychology

The term "cognitive map" originated with Tolman ([1948]). In opposition to Hull ([1930]), who tried to explain rat navigation through stimulus-response associations, Tolman suggested that rats mentally represent their surroundings. He argued that only a representational approach could explain how rats take novel detours and shortcuts. Since Tolman's opening salvo, the extent to which animal navigation requires mental representation of space has proven controversial. The cognitive map hypothesis found new popularity with publication of (O'Keefe and Nadel [1978]). More recently, Gallistel ([1990]) argues that animals perform computations over representations of spatial aspects of the environment. Most contemporary approaches fall between Hull's extreme anti-cognitivism and Gallistel's extreme cognitivism. For surveys, see (Redish [1999]; Shettleworth [1998]; Trullier, et al. [1997]).

The scientific literature attaches diverse meanings to the phrase "cognitive map," a diversity catalogued by Bermudez ([1998], pp. 203-7) and Kitchin ([1994]. This diversity occasions frequent conceptual and dialectical confusions. I distinguish three especially important usages. A cognitive map *in the trivial sense* is whatever mental or neural mechanism enables an animal to navigate. On this usage, it is tautologous that animals capable of navigation have cognitive maps. A cognitive map *in the loose sense* is a mental representation that represents geometric aspects of the environment. These aspects might be topological (e.g. connectedness, adjacency, or containment), affine (e.g. collinear or parallel), metric (e.g. distances and angles), and so on. A cognitive map *in the strict sense* is a mental representation that has the same basic representational properties and mechanisms as an ordinary concrete map. A cognitive map in the strict sense has the same type of content or format as a concrete map, while a cognitive map in the loose sense merely encodes the same information, possibly in a different way than a concrete map would encode it.

Terminological convention: when I use the phrase "cognitive map" without further qualification, I mean "cognitive map in the loose sense."

Psychologists disagree about whether various animals have cognitive maps in either the loose or the strict sense. Insect navigation is particularly vexed (Gallistel [1994], [1998]; Wehner [2003]; Menzel, et al. [2005]; Stelerny ([2003]), pp. 41-4). Mammalian navigation is somewhat less vexed. In a famous experiment, Cheng and Gallistel placed a rat in a non-square rectangular box, in one of whose corners the rat discovered food (Cheng [1986]). Cheng and Gallistel removed the rat from the box and disoriented it. When returned to an identical box, the rat usually searched for food either in the correct corner or in the diagonally opposite corner. Two diagonally opposite corners are metrically indiscernible, even though they are metrically discernible from the other two corners. Thus, the rat apparently represents metric features of its surroundings. In general, considerable evidence suggests that all mammals represent metric properties, and hence that they have cognitive maps in the loose sense. For a survey, see (Gallistel [1990]). For a less cognitivist perspective, see (Shettleworth [1998]).

Do animals have cognitive maps in the strict sense? That is not a question I will try to answer. But I will explore some relevant philosophical and scientific issues. In this section and the next, I review pertinent results from psychology and AI robotics, respectively. In §§5-6, I analyze these scientific results from a more philosophical perspective.

Following Levitt and Lawton ([1990]), navigation based upon a metric cognitive map faces the following questions: Where am I? Where are other objects and properties located? How do I get where I'm going? The three cognitive tasks corresponding to these questions are usually called *localization*, *mapping*, and *path-planning*. Localization and mapping are exercises in, or analogues to, belief-fixation. Path-planning is an exercise in, or analogue to, decision-making. I focus on localization and mapping, drawing heavily upon the exposition of (Gallistel [1990]).

The most elementary kind of localization is *dead reckoning* (sometimes also called *path integration* or *odometry*), which determines the creature's position by monitoring its motion through space. It may record velocity and integrate to compute position. It may also record acceleration and compute position by integrating twice. Dead reckoning has played a vital role in human marine navigation for millennia. An enormous literature conclusively demonstrates that even primitive creatures such as ants employ dead reckoning (Gallistel [1990], pp. 57-101; Wittlinger, et al. [2007]). For instance, after foraging explorations, the desert ant can return directly back to the nest with remarkable accuracy, even lacking relevant external cues. The devices employed to detect velocity and acceleration vary across species, but they include optical flow, vestibular signals, proprioception, and motor efferent copy.

Dead reckoning is fallible and noisy. Its errors are cumulative, rendering it unreliable over time. Researchers have explored various corrective strategies, requiring a range of representational resources. Here, I focus on a strategy that Gallistel ([1990]) calls *piloting*, whereby one observes the spatial distribution of salient objects and properties (*landmarks*) relative to oneself, using these observations and prior knowledge of the environment to infer one's location. Like dead reckoning, piloting plays a crucial role in marine navigation. Unlike dead reckoning, piloting requires a representation of geometric features of one's environment: a map in the loose sense. It is straightforward trigonometry to calculate one's *allocentric* position from the *egocentric* positions of sufficiently many appropriately positioned landmarks, taking for granted the landmarks' allocentric locations. (Allocentric coordinate systems are defined relative to the external environment, while egocentric coordinate systems are defined relative to one's own body.) See (Gallistel [1990]) for discussion of piloting and for a survey of evidence that various species engage in it.

Piloting introduces several difficulties, the most fundamental of which is that one can determine correct position through piloting only if one already has a relatively accurate representation of the environment.[3] In general, one may not have such a representation. Creatures often explore new terrain whose features are not known *a priori*. Moreover, the environment can change, so that one's map requires constant updating. Finally, even when moving through a static, familiar environment, a pre-existing map may be incorrect and therefore require emendation. In general, then, creatures localizing themselves cannot simply assume that they have an accurate map.

Theoretically, one can construct an allocentric metric map by combining dead reckoning with egocentric coordinates of landmarks. Elementary vector summation converts these two inputs into allocentric coordinates for landmark (Gallistel [1990], pp. 106-9). But, since dead reckoning is fallible, this is not a reliable procedure. Piloting must intervene to correct dead reckoning's cumulative errors. Thus, localization and mapping are hopelessly intertwined, separable only under special circumstances. Within AI, this conundrum is called the Simultaneous Localization and Mapping (SLAM) problem. It is widely regarded as the most formidable hurdle to building autonomous mobile robots. In many crucial respects, SLAM is a special case of abduction. It generates familiar problems of confirmation holism and underdetermination of theory by evidence.

Few discussions in the psychological literature adequately confront SLAM. Most theories either treat localization relative to a known map or mapping relative to known locations, without even mentioning that there is a problem about simultaneous mapping and localization.

The model of rodent navigation developed in (Touretzky and Redish [1996]; Redish and Touretzky [1996]), while in many respects unusually detailed, is typical in its evasion of SLAM. The model, a hybrid of symbolic and connectionist ideas, distinguishes two phases: *learning* and *recall*. During learning, the model employs path integration to learn egocentric distances and bearings of observed landmarks, as well as retinal angles between pairs of landmarks, as viewed from various locations at various orientations. During recall, the model employs this stored information, coupled with landmark-observation, to correct errors in path integration. The model provides no principled basis for deciding when the animal enters the learning phase and when it enters the recall phase. That is determined exogenously, through *ad hoc* setting of parameters in the connectionist network. Thus, when faced with conflicts between path integration and landmark-observation, the model provides no principled basis for resolving those conflicts by altering the position estimate or by altering the stored information about relevant landmarks. In other words, the model provides no principled solution to SLAM.

The most striking manifestation of this lacuna, emphasized in (Balakrishnan, et al. [1998]), is that the model's learning phase presupposes that dead reckoning is reliable over time, which it is not. Indeed, when employing the model in simulations, Touretzky and Redish implemented its learning phase by exogenously setting dead reckoning coordinates to the correct values. For this reason, as they admit ([1996], pp. 267-9), the model cannot explain how rodents actually map unfamiliar environments. Redish and Touretzky concede that "rodents must have some way to correct for path integration errors simultaneous with tuning their place cells during exploration" ([1996], p. 24). But their model does not illuminate the computations through which rodents accomplish this feat. It does not explain how rodents simultaneously update both the position estimate *and* the map based upon dead reckoning and landmark-observation. For further criticisms of the model, see (Balakrishnan, et al. [1998]).[4]

To explore SLAM more systematically, I turn from psychology to AI robotics. Although roboticists have hardly solved the problem, they have made impressive progress. (Readers less interested in technical details can skim §4 and resume reading carefully in §5.)

4 Cognitive maps in robotics

For the past decade, most robotics research on SLAM has occurred within the framework of Bayesian probability theory. This research harmonizes with the general program, popular in both philosophy of science and cognitive psychology, of handling confirmation and abduction through Bayesian methods. As applied to SLAM, the idea is to encapsulate the robot's beliefs about the environment with a probability distribution defined over the space of possible maps. The robot updates this probability distribution based upon motor commands and sensory inputs. For an overview of the Bayesian paradigm in robot mapping, see (Thrun, et al. [2005]), whose exposition and notation I closely follow. (I am somewhat fussier about use-mention distinctions than this text (p. 153), although I blur them whenever seems appropriate.)

More formally, we may express one version of SLAM as follows. At time t, the robot is given as input $z_{1:t}$, its sensor measurements from times 1 to t, and $u_{1:t}$, its motor commands from times 1 to t. The robot must calculate the posterior probability $p(x_t, m_t \mid z_{1:t}, u_{1:t})$, where $x_t = (x, y, \theta)$ represents the robot's "pose" (its location and its bearing relative to some fixed reference direction) at t, and m_t is a map of the environment at t. Abbreviate "$p(x_t, m_t \mid z_{1:t}, u_{1:t})$" as "$bel(x_t, m_t)$." To compute $bel(x_t, m_t)$, we employ an appropriate application of Bayes's rule. Assume that the environment does not change over time, so that the map m requires no temporal index t. Under this and a few other simplifying assumptions, the update rule is:

$$bel(x_t, m) = \eta\ p(z_t \mid x_t, m) \int p(x_t \mid x_{t-1}, u_t)\, bel(x_{t-1}, m)\, dx_{t-1}$$

where η is a scaling constant. This equation is called the *Bayes filter*. Note that $bel(x_{t\text{-}1}, m)$ represents the robot's beliefs about its position and its environment at time t-1. Thus, the Bayes filter computes the robot's beliefs at time t based upon: its beliefs at time t-1; the probability $p(x_t \mid x_{t\text{-}1}, u_t)$ that motor command u_t will carry the robot from pose $x_{t\text{-}1}$ to pose x_t; and the probability $p(z_t \mid x_t, m)$ that sensor reading z_t will result when the robot is in pose x_t within a world described by map m. Assuming an initial probability distribution $bel(x_0, m)$, the Bayes filter provides a recursive technique for calculating the robot's beliefs as time evolves.

How should the robot compute $p(x_t \mid x_{t\text{-}1}, u_t)$ and $p(z_t \mid x_t, m)$? The first quantity corresponds roughly to dead reckoning. A slight wrinkle is that dead reckoning usually relies upon not just motor commands but also sensor measurements of velocity. In practice, roboticists typically ignore this wrinkle by treating velocity measurements as if they were control signals u_t rather than sensor measurements z_t. Generally speaking, then, the algorithms for calculating $p(x_t \mid x_{t\text{-}1}, u_t)$ recapitulate the kinematic computations underlying dead reckoning, emended to accommodate noise (Thrun, et al. [2005], pp. 117-43). The algorithms for calculating $p(z_t \mid x_t, m)$ usually include trigonometric computations, emended to accommodate noise, that convert the pose estimate and a landmark's estimated allocentric coordinates into the landmark's predicted egocentric coordinates and heading (Thrun, et al. [2005], pp. 176-8).

As Thrun, et al. argue ([2005], p. 9), Bayesianism is well-suited to SLAM. Through the Bayes filter, it seamlessly merges perceptual and motor input into an updated representation of one's pose and of the environment. Assuming Bayesian probability theory, this updated representation is rational. Moreover, because the Bayesian approach traffics in probability distributions, it can explicitly represent uncertainty arising from errors in dead reckoning, sensor noise, underdetermination of theory by evidence, and other factors.

In general, implementing the Bayes filter requires a discrete, tractable approximation. To a large extent, research on SLAM consists in constructing tractable approximations that are relatively accurate. I will not try to survey the overwhelming variety of approximations currently being explored. I focus on one of the most venerable paradigms: EKF SLAM, first applied in (Smith and Cheeseman [1986]). In some respects, this paradigm compares unfavorably to newer techniques like the *particle filter* (Thrun, et al. [2005], pp. 437-83). The philosophical points I make in §§5-6 are readily adapted to those newer techniques. I emphasize EKF SLAM because it is easy to describe in a relatively non-technical way. For a rigorous statement, along with details about successful applications in actual robots, see (Thrun, et al. [2005], pp. 309-32).

We treat maps as composed of point landmarks in the Cartesian plane. More precisely, a map m is a vector ($m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2 \ldots m_{N,x}\ m_{N,y}\ s_N$), where $m_{i,x}$ and $m_{i,y}$ are the allocentric coordinates of the ith landmark, s_i is a "signature" encapsulating some of the landmark's features, and N is the number of landmarks. s_i might be a vector whose elements correspond to observable features like height, color, and so on. We assume that $bel(x_t, m)$ is a multivariate Gaussian distribution. We then represent $bel(x_t, m)$ through two parameters: its mean, μ_t, and its covariance matrix, Σ_t. The mean is a $3N+3$ element vector incorporating the robot's pose x_t and a map m: ($x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2 \ldots m_{N,x}\ m_{N,y}\ s_N$). The covariance matrix is a $3N+3$ by $3N+3$ array whose i-j element represents the covariance between the ith and jth variables.[5]

We need an algorithm for computing (μ_t, Σ_t) from (μ_{t-1}, Σ_{t-1}), that is, for computing $bel(x_t, m)$ from $bel(x_{t-1}, m)$. We can employ a basic statistical tool called the *Extended Kalman filter*, a modern descendant of Gauss's method of least squares. Oversimplifying greatly, the Extended Kalman filter yields an algorithm, EKF SLAM, that divides into two stages. The first stage produces a preliminary pose estimate based upon motor command u_t. This stage, which corresponds to the integral $\int p(x_t \mid x_{t-1}, u_t)\, bel(x_{t-1}, m)\, dx_{t-1}$ from the Bayes filter, incorporates the kinematic calculations from dead reckoning. The second stage, which corresponds in the Bayes filter to multiplying the integral by $p(z_t \mid x_t, m)$, loops through all the landmarks observed at time t. For each observed landmark, the algorithm updates both the pose estimate and the map estimate. In so doing, it incorporates the trigonometric calculations that convert the pose estimate and a landmark's estimated allocentric coordinates into the landmark's predicted egocentric coordinates. Beyond these kinematic and trigonometric calculations, EKF SLAM requires only some elementary calculus and matrix algebra.

One disadvantage of EKF SLAM is that it models all beliefs as Gaussian distributions, which are unimodal. Intuitively, this does not allow the robot to entertain multiple hypotheses. What if the robot cannot identify which landmark it currently perceives? Then the rational probability distribution assigns high probability density to several distinct hypotheses. There are various ways to handle this problem. For instance, a *multi-hypothesis* variant upon EKF SLAM allows probability distributions that are mixtures of several Gaussians.

During the 1980s and early 1990s, mobile robotics research focused mainly on path-planning rather than localization and mapping. This research occurred before probabilistic methods achieved dominance in the study of localization and mapping, so it mainly transpired within a non-probabilistic framework. Non-probabilistic path-planning assumes that the robot has a map of its environment, not a probability distribution over the space of possible maps. The main goal of path-planning is to plot a collision-free course towards some destination. This requires converting a map of the environment into a geometric structure, such as a graph, over which one can perform computationally efficient searches. Non-probabilistic research along these lines has yielded many advances, as detailed in (Choset, et al. [2005]). Path-planning within a probabilistic framework is much less well developed, although (Thrun, et al. [2005], pp. 487-568) reviews some preliminary results.

In general, current autonomous mobile robots work fairly well for indoor, static environments. Outdoor and dynamic environments remain largely unsolved problems. However, the field is witnessing rapid advances, with many important developments in just the past ten years. The high rate of recent progress provides grounds for cautious optimism.

To what extent do models like EKF SLAM illuminate animal navigation? Robotics engineers robots, while psychology provides empirical theories of existing animals. Why should computational models designed by roboticists bear any resemblance to the cognitive processes instantiated by animals? The question is especially pressing given how strikingly robots and biological animals differ in their sensory transducers, motor organs, and internal hardware.

So far, few psychologists studying navigation have emphasized Kalman filters. But Bayesianism in general, and Kalman filters in particular, have found fruitful application within areas such as vision (Rao, et al. [2002]) and postural balance (Kuo [2005]). More relevantly for us, Balakrishnan, et al. ([1999]) argue that the rat's hippocampus instantiates a computational model of SLAM based upon the Kalman filter. They support this hypothesis with neural and behavioral data. Similarly, Gallistel ([2008], pp. 138-9) urges psychologists studying animal navigation to deploy the Bayesian techniques developed within probabilistic robotics. It seems likely that the psychological study of navigation will eventually witness further applications of probabilistic robotics. Despite their many hardware differences, robots and animals must solve the same basic tasks: localization, mapping, and path-planning. Animals may not employ the specific algorithms mentioned above. But those algorithms, and the more general Bayesian framework they presuppose, deserve extended empirical scrutiny. Currently, Bayesianism provides our only remotely adequate approach to SLAM. By default, it is our only viable candidate theory of metric spatial representation.

5 Cognitive maps in the strict sense?

I now want to analyze the models surveyed in §4 from a more overtly philosophical perspective. While the previous two sections mostly contained exposition of scientific theories, this section and the next contain far more contentious philosophical argumentation. I begin by examining the extent to which the models from §4 posit cognitive maps *in the strict sense*: that is, representations with "the same basic representational properties and mechanisms as an ordinary concrete map."

An essential feature of ordinary cartographic representation is *geometric structure*. The most common types of structure are metric, as with city maps, and topological, as with subway maps. At a bare minimum, then, a cognitive map in the strict sense should have geometric structure. This raises the question: what is it for a mental representation to have geometric structure? An analogous question arises for logically structured representations: what is it for a mental representation to have logical structure? But the question may seem more challenging for geometric than for logical structure. No one thinks that we can open up the head to discover a miniature map laid out in physical space.[6] What, then, could it possibly mean to attribute geometric structure to a cognitive map?

This worry reflects an insufficiently abstract conception of geometric structure. On the modern mathematical conception, a geometric structure is a set of objects satisfying certain axioms. For instance, a metric space is an ordered pair (X, d), where X is any set and d is a function from $X \times X$ to the non-negative real numbers such that: $d(a, b) = 0$ iff $a = b$, $d(a, b) = d(b, a)$, and $d(a, c) \leq d(a, b) + d(b, c)$. Under this abstract conception, there is no conceptual bar to instantiating geometric structure within a sufficiently powerful computational system. Any entities, including representations manipulated by a mind or computer, may compose a geometric structure. One need simply ensure that the representations bear appropriate computational relations to one another, relations satisfying the relevant axioms.

Consider a point-landmark map m, supplemented with the robot's pose: $(x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$. Viewed on its own terms, it is just a vector, lacking any geometric structure. However, m participates in a computational model that naturally associates it with a Euclidean distance metric. For instance, during the trigonometric calculations corresponding to $p(z_t\,|\,x_t, m)$, EKF SLAM must compute an estimate for the ith landmark's egocentric coordinates. To do so, it applies the Euclidean distance metric, estimating that landmark i has egocentric distance of $\sqrt{(x-m_{i,x})^2+(y-m_{i,y})^2}$ (Thrun, et al. [200], p. 319). Actually, this description is sloppy, since the map's constituents are syntactic items, not numbers, so that arithmetical operations are not directly defined over them. If $den(a)$ is the denotation of a, then the predicted egocentric distance is $\sqrt{(den(x)-den(m_{i,x}))^2+(den(y)-den(m_{i,y}))^2}$. In this way, EKF SLAM associates m with a mapping $d((a\ b), (c\ d)) = \sqrt{(den(a)-den(c))^2+(den(b)-den(d))^2}$ that satisfies the axioms of a metric space.

What does the requisite "association" between m and d consist in? The most natural suggestion is *functionalist*: the association consists in patterns of computational activity. (Cf. Block [1983], p. 514; Tye [1991], pp. 40-1.) Specifically, EKF SLAM associates m with d by repeatedly deploying d during computations over m. We can imagine computations that employed some other metric, such as the city block metric ($d((a\ b), (c\ d)) = |a-c| + |b-d|$) or the discrete metric ($d(x, y) = 1$ iff $x \neq y$). But those metrics play no role in EKF SLAM. The Euclidean metric does. That is why EKF SLAM associates m with the Euclidean metric rather than some other metric.

Undoubtedly, the notion of "association" would benefit from further elucidation. Even lacking such elucidation, there is a clear sense in which, by performing appropriate computations over *m*, EKF SLAM engenders a distance function over the coordinates of *m*. The function does not measure literal separation in space, since syntactic items are not literally laid out in space. But it is a legitimate distance function nonetheless, for it satisfies the axioms of a metric space. Whether or not we want to say that *m* is geometrically structured in *exactly* the same sense as a concrete map, there is a clear sense in which it is geometrically structured. Similar remarks apply to many other representations deployed in AI, Geographic Information Science, and computational geometry.[7]

Ordinary concrete maps do not just *have* geometric structure. They *represent* geometric structure. The compositional mechanisms through which they do so are not entirely understood. We lack a canonical treatment of quotidian cartographic representation comparable to the familiar semantics of first-order logic. A relatively uncontroversial point, emphasized by Sloman ([1978]) and Casati and Varzi ([1999]), is that a map's geometric structure purports to replicate salient relations between objects represented by the map.[8] More precisely, the map is correct only if its geometric structure replicates salient relations between objects represented by the map. If symbol *A* appears a certain distance on a metric map from symbol *B*, this represents that a certain scaled distance separates the objects denoted by *A* and *B*. If symbol *A* is topologically connected to symbol *B* on a subway map, this represents that the subway system connects the subway stations denoted by *A* and *B*. Of course, a map may *inaccurately* replicate geometric structure. Even then, the map's geometric structure contributes to its representational content by *purporting* to replicate geometric relations among objects. A sentence's components may also bear geometric relations to one another, but the relations do not typically carry any representational import. In "Jack loves Mary," the word "Jack" appears a certain distance to the left of "Mary." This does not signify anything about how Jack and Mary relate to one another, either spatially or otherwise, as evidenced by the fact that one could just as well say "Mary is loved by Jack."

I do not claim that, *in general*, representing geometric structure involves replicating (or purporting to replicate) geometric structure. Representing property *Y* does not usually require mimicking or purporting to mimic property *Y*. The word "green" represents the property of being green, but it is not itself green. Similarly, a sentence can represent geometric structure without itself purporting to replicate that structure. For instance, a very long conjunctive sentence describing a subway might "convey the same information" as a map of the subway, in the sense that the sentence and map are true in the same possible worlds. But the geometric relations between words in the sentence do not carry representational import. In contrast, the map's geometric structure carries representational import. Thus, even though the sentence and the map are true in the same possible worlds, they operate in fundamentally different ways. The map, unlike the sentence, is correct only if it replicates relevant geometric structure.

One might wonder *why* maps are correct only if they replicate relevant geometric structure. Given that property replication is not usually a prerequisite for correct representation, why is it a prerequisite for maps? The answer is that this just seems to be how maps operate. It is part of our pre-theoretic concept *map* that a map's geometric structure contributes to its representational content by purporting to replicate geometric relations between represented entities. The map represents some region of physical space as being such that it shares certain geometric properties with the map, and the map is correct only if the region is as the map represents it as being. We can cite various *benefits* of choosing a representational system that operates in this way, a topic discussed by Sloman ([1978], pp. 168-76). But it is not clear that we can provide a deeper explanation for why a map is correct only it replicates relevant geometric structure. Even lacking a deeper explanation, we can recognize, through reflection on a range of examples, that this is an essential feature of ordinary cartographic representation. A cognitive map in the strict sense should preserve it.

Once again, $m = (x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$ exhibits the desired feature. The fact that m contains (x, y) and $(m_{i,x}\ m_{i,y})$, coupled with the fact that m is associated with the Euclidean metric d, signifies that the robot and landmark i are separated by distance $d((x\ y), (m_{i,x}\ m_{i,y}))$. More generally, m is true only if the robot occupies the indicated position with the indicated heading, and only if each landmark has the egocentric distance and bearing predicted by the map. I codify this intuitive idea more formally in the Appendix, where I provide a semantics for m. Even without referring to this semantics, I think it intuitively clear that m is correct only if it replicates relevant geometric aspects of the region it maps. Note that m might be *useful* for navigation even if it is not *true*. For instance, a map with systematic metric distortions might be better than nothing, especially if it preserves affine or topological properties.

Apparently, then, point-landmark maps display the following two properties: (a) the representation has geometric structure; (b) the representation is correct only if its geometric structure replicates salient geometric relations between entities represented by components of the representation.[9] Should we classify representations satisfying (a) and (b) as cognitive maps in the strict sense? That depends on whether (a) and (b) exhaustively enumerate "the basic representational properties and mechanisms" of concrete maps. We might call such representations "cognitive maps *in the approximate sense*," leaving open whether they are cognitive maps in the strict sense. Certainly, a cognitive map in the approximate sense exhibits striking similarities with concrete maps.

6 Logically structured representations?

The models employed in probabilistic robotics suggest a pluralistic conception of computational psychology embracing non-logical mental representations. For instance, point-landmark map $m = (x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2 \ldots m_{N,x}\ m_{N,y}\ s_N)$ does not overtly feature the basic elements of logical form: sentential logical connectives, quantifiers, or even predication. EKF SLAM employs many intermediate representations while updating its pose and map estimate, as encapsulated by m. On the surface, none of these varied representations exhibit the familiar syntax or semantics of the predicate calculus. They are scalars, vectors, or matrices. A natural conjecture, then, is that animal navigation involves computations defined over cognitive maps that have geometric structure but not logical structure.

One might object to this analysis in several different ways. **First**, one might suggest that, despite initial appearances, cognitive maps in the approximate sense have logical structure. **Second**, one might concede that cognitive maps in the approximate sense lack logical structure but insist that computational models involving them include *additional* representations that possess logical structure. **Third**, one might acknowledge that the foregoing models involve exclusively non-logical representations but insist that the models are *implemented* by a deeper logical level of computation and representation. (Cf. Hayes [1985].) **Fourth**, one might argue that animal navigation is best explained by *alternative* models whose computations are defined over predicate calculus sentences instead of, or in addition to, maps in the approximate sense.

In evaluating the **first objection**, we must distinguish between a mere *list* of representations and a *unified* representation that expresses a proposition. As Fodor and Pylyshyn ([1988], p. 27) put it, ‘a theory of mental representation must distinguish the case when two concepts (e.g. THIS BODY, HEAVY) are merely *simultaneously entertained* from the case where, to put it roughly, the property that one of concepts expresses is predicated of the thing that the other concept denotes (as in the thought: THIS BODY IS HEAVY).' Viewed on its own terms, $m = (x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$ is a mere list, and hence does not involve predication. In itself, it has neither metric nor logical structure. If incorporated into a computational model that appropriately exploits the Euclidean distance metric, *m* has metric structure. But does *m* thereby come to possess *logical* structure, such as predication? For instance, where are the predicates in $(x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$? Is s_1 a predicate, used to attribute some property to the location denoted by $(m_{1,x}\ m_{1,y})$? Or is $(m_{1,x}\ m_{1,y})$ the predicate? Lacking further argument, neither suggestion is very plausible. $(m_{1,x}\ m_{1,y})$ is most naturally viewed as denoting a location in space. s_1 is most naturally viewed as denoting properties of *the landmark itself*, not the landmark's location. One might propose that s_1 is a predicate meaning "is occupied by a landmark with such-and-such properties." Yet what motivates such an ornate proposal, save an antecedent commitment to explicating all mental representation as ultimately logical?

Traditionally, philosophers emphasize two related but distinct rationales for attributing logical form to sentences and thoughts. The first is the need to provide a compositional truth-conditional semantics. The second is the need to provide a systematic theory of logical inference.[10] The only way to evaluate whether the first rationale applies to *m* is to study *m*'s semantics in more formal detail. I undertake this task in the Appendix, where I return a negative verdict. I focus here upon the second rationale.

As Frege demonstrated, by attributing proper logical form to sentences, we can systematically describe many patterns of acceptable inference. One can develop this idea in a syntactic direction, as illustrated by proof theory, or in a semantic direction, as illustrated by

Tarski's theory of logical consequence. The syntactic perspective is more relevant to us. It underlies the doctrine, popular within cognitive science and especially AI, that deductive reasoning is proof-theoretic manipulation of logically structured mental representations.

Does this rationale apply to the computational models from §4? Do the models involve inferential patterns best described by attributing logical form to the relevant representations? No. We search those models in vain for familiar inferential patterns like modus ponens, reductio ad absurdum, or Leibniz's law. Proof theory is a diverse field, encompassing Hilbert-style formal systems, Gentzen's sequent calculus, natural deduction systems, Herbrand's theorem, and so on. *None* of these diverse perspectives inform the syntactic manipulations employed by typical algorithms from probabilistic robotics. For instance, EKF SLAM updates *m* through computations that combine kinematics, trigonometry, matrix algebra, and elementary calculus. The computations are like those that any graphics calculator executes, albeit far more demanding of time and memory. The absence of proof-theoretic machinations is particularly glaring when juxtaposed with the genuinely logic-based AI algorithms surveyed by Russell and Norvig ([2003], chaps. 7-10). Those algorithms are directly based upon proof theory. For example, *resolution* and *unification* convert Herbrand's theorem into a search procedure for quantificational unsatisfiability. Computations from probabilistic robotics do not even feature induction by logical complexity, a basic proof-theoretic device. Thus, we do not need to posit logical form to describe the rational computational processes surveyed in §4.

This argument does not conclusively show that map *m* lacks logical form. But it places the burden of proof squarely upon anyone who claims otherwise.

An important subsidiary point emerges from our discussion of logical inference. *Even if we were to concede that the representations in question have logical form*, computations defined over them do not exploit it through logical inference. The computations are not proof-theoretic transformation of symbols. Thus, even if we concede that the representations are patterned after logic, the computations are not. Of course, the only way to show definitively that the models involve no proof-theoretic manipulation is to provide pseudocode for them. That is also the only way to refute the **second of our four objections**: that the models contain additional representations that are logically structured. Still, I hope I have provided enough detail to render my assessment plausible.

Our **first two objections** may seem to receive support from §4's emphasis on probability. Philosophers often depict probability distributions as defined over propositions or sentences. Typical axiomatizations describe how the probability of a logically complex sentence depends on the probabilities of its components. For instance, standard axioms entail that

(*) $p(\neg F) = 1\text{-}p(F)$.

Doesn't the tight relation between probabilistic inference and the propositional calculus show that probabilistic robotics requires a logically structured representational medium?

One *can* develop probability theory over logically structured representations. But the concept of probability is broader. We require that a probability distribution be associated with representations that depict possible states of the world, not that the representations have logical structure. Notably, Kolmogorov offered an axiomatization that does not presuppose logically structured representations. The axiomatization, which is standard in statistics textbooks, defines a probability space as (Ω, A, p), where Ω is a non-empty set, A is a σ-algebra over Ω (i.e. a set of subsets of Ω that contains Ω and is closed under countable union and complementation in Ω), and p is a probability measure. Specifically, we might take Ω to be a set of (pose, map) ordered pairs. Working within this framework, we can derive the standard elements of probability theory, including Bayes's theorem. In this framework, the clause corresponding to (*) is

(**) $p(\Omega/x) = 1\text{-}p(x)$,

where $x \in A$ and Ω/x is the complement of x in Ω. In contrast to (*), (**) makes no overt appeal to negated representations within p's domain. Of course, *we* use negation to define set-theoretic complementation. That does not entail that p attaches probabilities to negated representations. Apparently, then, we can treat robot navigation as an exercise in probabilistic inference defined over map-like representations, without recourse to logically structured representations.

One might claim that the σ-algebra framework *implicitly* assigns logically structured representations to the probability measure's domain. But that would require extensive argument. It is not a straightforward interpretation of what the framework literally says. For instance, if someone were to suggest that taking set-theoretic complements over a probability space is a disguised way of negating propositions, and hence that (**) is a disguised form of (*), we could illustrate the difference with the following toy example. Suppose we feel certain that object x is a located in one of three locations a, b, c, but we are not sure which. We can represent our belief state with a probability distribution defined over the powerset of $\Omega = \{a, b, c\}$, where each element represents the set of possible worlds in which x inhabits the corresponding location. Then (**) entails that $p(\{a, b\}) = p(\Omega/c) = 1\text{-}p(\{c\})$. Here, $\{a, b\}$ represents the set W_1 of possible worlds where x inhabits either a or b. Note that W_1 is *not* the set W_2 of possible worlds corresponding to the proposition "x is not located at c." A world in which x inhabits some fourth location d belongs to W_2 but not to W_1. Hence, taking complements over a probability space need not correspond to negating propositions. So (**) is not a disguised form of (*).

The **third objection** describes one viable conception of robot navigation. Golog, introduced in (Levesque, et al. [1997]), is a popular robotics programming language similar to Prolog. Both languages describe inferential manipulations performed upon predicate calculus sentences. If one uses Golog to implement the computational models surveyed in §4, then the resulting robot involves two distinct levels of computation and representation: one cartographic (in the approximate sense), the other logical. Thus, animal cognition *might* work as the objection describes.

Crucially, however, it *need* not work that way. For instance, the language CES, introduced by Thrun ([2000]) as an extension of C++, includes primitive terms for representing and manipulating probability distributions. The language facilitates implementation of probabilistic models like those canvassed in §4. Navigational programs written in CES manipulate lists, stacks, queues, trees, and so on. They do not manipulate syntactic items endowed with anything resembling the normal syntax or semantics of the predicate calculus.

In principle, then, computational navigational processes do not require a logical implementation. Whether animal navigation *in fact* has a logical implementation is an empirical question I am not trying to answer.

The **fourth objection** describes a possible class of navigational models. For instance, Shanahan and Witkowski explicitly aim to develop "robot architectures in which logic is the medium of representation and theorem proving is the means of computation" ([2001], p. 19). Shanahan ([1997]) develops an account that represents spatial information through an event calculus couched in a logically structured representational medium. Localization and mapping are treated as abductive inferences performed in this medium. A logic programming language, such as Golog, implements the model. Shanahan's account is purely logicist. Specifically, it eschews cognitive maps in the approximate sense. In contrast, Kuipers ([2000]) and Remolina and Kuipers ([2001]) develop a complex hybrid account involving four different levels of spatial representation, one of which is metric and another topological. The metric level involves cognitive maps in the approximate sense. The topological level includes an axiomatized theory about the causal and topological structure of space, a theory that the navigational system exploits during localization and mapping. The theory's axioms are encoded in a logical medium. Thus, Kuipers's account deploys both geometrically *and* logically structured representations.

I will not try to evaluate whether animal navigation instantiates something closer to a cartographic approach, the logicist approach espoused by Shanahan, or the hybrid approach espoused by Kuipers. I mention a few points just to indicate some complexities. Despite much emphasis on SLAM's abductive character, Shanahan offers only sketchy indications regarding how exactly one might solve it. In particular, he does not explain how to resolve the underdetermination of theory by evidence, a point on which Remolina and Kuipers ([2001]) critique him. Kuipers and his colleagues offer an explicit SLAM algorithm for metric maps (Beeson, et al. [2006]). But the algorithm draws solely upon geometric, non-logical representations. Ultimately, then, it is not clear that any successful existing model of *metric* cognitive maps assigns logical representation a central role. Since one can extract topological maps from metric maps, it is unclear whether logical representations should play any essential role.[11]

My thesis is *not* that a purely cartographic, non-logical theory of mental representation best explains animal navigation. Any such conclusion would be premature. We require extensive theoretical and empirical exploration of the rival approaches. My point is that a purely cartographic analysis of animal navigation presently seems just as promising as Shanahan's logicist approach or Kuipers's hybrid approach.

7 Systematicity and productivity

A famous pair of Fodorian arguments ([1987], pp. 147-53) contends that Mentalese yields the best explanation for two phenomena: *productivity* (thinkers can entertain a potential infinity of distinct thoughts) and *systematicity* (there are systematic connections between which thoughts a thinker can entertain). Do these arguments show, contrary to §6, that animal navigation requires a logically structured representational medium? I claim that they do not. Since Fodor ([1987], p. 148) himself admits that the productivity argument is debatable, I focus upon systematicity. Most of what I say would generalize to handle productivity.

Fodor takes as a datum that thought is systematic: the ability to entertain certain thoughts suffices for the ability to entertain certain other thoughts. For instance, thinkers able to entertain the thought that John loves Mary can also entertain the thought that Mary loves John. As Fodor and Pylyshyn observe, '[i]t's not enough just to stipulate systematicity; one is also required to specify a mechanism that is able to enforce the stipulation' ([1988], p. 50). One must do more than show that systematicity is *consistent* with some preferred psychological theory, such as connectionism. One must show why, based on the theory, we would expect minds to be systematic. Fodor and Pylyshyn argue that the only adequate solution is to treat representational mental states as relations to symbols with a combinatorial syntax and semantics ([1988], pp. 37-50). Thus, explaining systematicity requires us to posit a language of thought.

To assess how this argument bears upon animal navigation, let us distinguish two senses in which mental representations might be "sentential" or "languagelike." (Cf. Camp [2007], p.

152.) Representations are *sentential in the weak sense* if they have syntactic structure and a compositional semantics. They are *sentential in the strong sense* if they feature the basic compositional mechanisms familiar from formal logic: predication, truth-functional connectives, quantification, and so on. Fodor often emphasizes mental sentences in the strong sense, but his official definition of Mentalese ([1987], pp. 134-8) only mentions sententiality in the weak sense. In particular, Fodor's systematicity argument only aims to establish the existence of mental sentences in the weak sense. In his own words, the argument aims to establish only 'the combinatorial structure of thoughts' ([1987], p. 151). The argument does not purport to isolate any particular Mentalese compositional mechanisms. As Block observes, then, systematicity 'may reflect only a combinatorial pictorial system' ([1995], p. 411), rather than mental sentences in the strong sense. Similarly, Braddon-Mitchell and Jackson ([2007], p. 182) note that cartographic representation is systematic and hence that one can explain systematicity by treating representational states as relations to *mental maps*.

Our discussion of point-landmark maps confirms this diagnosis. Any creature able to entertain the map $(x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$ is also able to entertain numerous other maps, such as $(x\ y\ \theta\ m_{2,x}\ m_{2,y}\ s_1\ m_{1,x}\ m_{1,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$. A creature who entertains the first map can "recombine" its elements to form the second, just as a thinker who entertains the Mentalese sentence JOHN LOVES MARY can recombine its elements to entertain MARY LOVES JOHN. Thus, the computational models canvassed in §4 entail a form of systematicity. I conclude that systematicity poses no threat to §6's conclusion: a purely cartographic, non-logical approach to animal navigation is at least as promising as an approach that posits logically structured mental symbols.

Based on their discussion of mental maps, Braddon-Mitchell and Jackson argue that we can dispense with logically structured mental representations, treating *all* representational mental states as relations to mental maps. In response to this kind of view, Fodor ([1991], p. 295) and Rey ([1995], p. 208) observe that a purely pictorial or cartographic representational system cannot express negation or disjunction, let alone quantification. Since humans often entertain negated, disjointed, and quantified thoughts, Fodor and Rey conclude that a pictorial or cartographic representational system cannot exhaust the resources of Mentalese.

I find this response to Braddon-Mitchell and Jackson convincing, *as applied to humans*. The response seems less convincing when applied to non-linguistic creatures, such as bees or even rats. Is it so clear that these creatures entertain thoughts with logical structure? Is it so clear that explaining how these creatures behave requires attributing negated, disjoined, or quantified thoughts to them? I have questioned whether such representational resources are needed for explaining one central cognitive activity: navigation. It begs the question even to assume, as Fodor and Pylyshyn do, that non-linguistic creatures can 'perceive/learn that *aRb*' ([1988], p. 44). The thought that *aRb* involves a two-place predicate *xRy*. As we have seen, the computational models canvassed §4 do not feature predicational structure. For further discussion of these issues, see (Rescorla, [forthcoming a]).

8 Consequences for philosophy and psychology

Since the inception of cognitive science, researchers have debated whether we should countenance "non-logical," "non-discursive," or "non-propositional" representations. What does the present discussion add to this perennial dispute? The most basic contribution is to bring the dispute into contact with an impressive class of psychological theories and computational models whose bearing upon it has hitherto received insufficient emphasis. Philosophers rarely pay detailed attention to cognitive maps as possible instances of non-discursive representation. In many respects, however, maps provide a more instructive contrast with logical representation than phenomena such as mental imagery. A particularly important point concerns the possibility of rational processes in a non-logical representational medium. The models presented in §4 are rational: they convert perceptual and motor input into map-like representations rendered probable by that input. They thereby provide a detailed solution to one version of the problem of abduction. Yet they invoke neither logically structured representations nor proof-theoretic manipulations. Even if the models are not remotely close to empirically correct theories of animal navigation, they demonstrate the *possibility* of rational mechanisms fundamentally distinct from the logic-based machinations so often emphasized within philosophy and AI.

Our discussion also bears upon more empirical concerns. The computational models surveyed in §4 provide a promising avenue for scientific research into animal navigation. They offer the only known theoretical framework remotely adequate for solving metric versions of SLAM. They thereby illustrate how non-logical mental representation has the potential to illuminate a fundamental mode of animal cognition.

9 Appendix: cartographic semantics

In this appendix, I offer a compositional semantics for the point-landmark maps deployed in contemporary robotics. I then investigate the philosophical significance of this semantics.

A compositional semantics must do two things: specify the semantic properties of primitive expressions; and display how the semantic properties of a complex expression depend on its structure and on the semantic properties of its parts. When the complex representation is a point-landmark map $m = (x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2\ \ldots\ m_{N,x}\ m_{N,y}\ s_N)$, the relevant structure is the metric d associated with m, and the relevant parts are the primitive elements x, y, θ, $m_{1,x}$, etc. To specify the semantic properties of these elements, we endow m with a denotation function *den*. This function carries each element x, y, $m_{i,x}\ m_{i,y}$, to a real number. It carries θ to a real number between 0 and 2π. It associates s_i with an ordered pair (u_i, v_i), where u_i is a landmark and v_i is a vector of observable properties. Finally, we associate m with a system of Cartesian coordinates for physical space. (Somewhat artificially, we can regard this coordinate system as denoted by the parentheses appearing in m. Doing so ensures that m's truth-condition is a function solely of its metric structure and the denotations of its parts.) Since $den(s_i)$ is an ordered pair, let $den_1(s_i)$ and $den_2(s_i)$ be its first and second elements, respectively. To display how the semantic properties of m depend on those of its parts, we say that

(SEM) m is true iff: the robot inhabits the location in physical space with Cartesian coordinate $(den(x), den(y))$; the robot's allocentric heading relative to the x-axis of the coordinate system is $den(\theta)$; $den_1(s_i)$ has properties $den_2(s_i)$; the egocentric bearing of $den_1(s_i)$, relative to the robot, is $\text{atan2}(den(m_{i,y}) - den(y), den(m_{i,x}) - den(x)) - den(\theta)$; the distance in physical space between the robot and $den_1(s_i)$, as measured in some fixed standard units, is $d((x\ y), (m_{i,x}\ m_{i,y}))$.[12]

SEM codifies the intuitive idea that m is correct only if it replicates relevant geometric aspects of physical space. Thus, it corroborates my claim from §5 that m is a "cognitive map in the approximate sense." As far as I know, SEM is the first attempt in the literature at providing an explicit semantics for point-landmark maps. However, I would argue that something like SEM implicitly informs much of the literature on robot navigation.

At first blush, SEM may look less like a compositional semantics than like a trivial list of semantic properties. Any such appearance is misleading. SEM does not enumerate the semantic properties of *m*'s elements. That was already achieved by our specification of *den*. Rather, SEM displays how the semantic properties of *m*'s elements, together with *m*'s geometric structure, determine *m*'s truth-condition.

Consider the final clause of SEM: "the spatial distance between the robot and landmark $den_1(s_i)$ is $d((x\ y), (m_{i,x}\ m_{i,y}))$." This clause is not a trivial consequence of specifying *den* and *d*. If we omit it, then the resulting semantics SEM' no longer requires that *m* preserve distances. To illustrate, suppose that:

(i) $m = (x\ y\ \theta\ m_{1,x}\ m_{1,y}\ s_1\ m_{2,x}\ m_{2,y}\ s_2)$

(ii) $d((a\ b), (c\ d)) = \sqrt{(den(a) - den(c))^2 + (den(b) - den(d))^2}$

(iii) $(den(x), den(y)) = (0, 0)$

(iv) $(den(m_{1,x}), den(m_{1,y})) = (1, 0)$

(v) $(den(m_{2,x}), den(m_{2,y})) = (5, 0)$

(vi) $den_1(s_1) = u_1$ and $den_1(s_2) = u_2$

(vii) Relative to our Cartesian coordinates for physical space, the robot is located at the origin, and landmarks u_1 and u_2 are located on the *x*-axis. The actual physical distances from the robot to u_1 and u_2 are 4 and 1000, respectively.

Given (i)-(vi), the final clause of SEM entails that *m* is true only if the distances from the robot to landmarks u_1 and u_2 are 1 and 5, respectively. So (vii) entails that *m* is false. In contrast, SEM' allows *m* to be true, as long as *m* satisfies the other clauses of SEM. Intuitively, SEM rather than SEM' yields the correct verdict. This example illustrates the non-triviality of SEM. It shows that SEM imposes a substantive constraint upon the relation between *m* and physical space, a constraint that does not follow just from specifying *m*'s geometric structure and the denotations of its elements. Intuitively, the constraint is that *m* replicate distances between the robot and each landmark. Similar examples illustrate that the other clauses of SEM are non-trivial.

Readers may still resist calling SEM a genuine semantics, on the grounds that it directly *states* truth-conditions for *m* rather than showing how to *derive* truth-conditions. In contrast, Tarskian semantics employs recursive clauses for the logical connectives (e.g. "$p \vee q$" is true iff *p* is true or *q* is true; "$\neg p$" is true iff *p* is not true). One can iteratively apply these clauses to derive truth-conditions for a potential infinity of sentences. Shouldn't a genuinely compositional semantics proceed in similar fashion?

I respond that SEM does not state our desired truth-condition for *m*. Compare it to the Tarskian clause for atomic predication, which runs, nearly enough for our purposes, as follows:

(*) "*Fa*" is true iff "*F*" is true of *den*("*a*").

Our desired truth-condition for "*Fa*" does not involve the denotation relation, "*F*", or "*a*." We obtain the desired truth-condition only when we enrich (*) by specifying which object "*a*" denotes and which objects "*F*" is true of. Similarly, our desired truth-condition for *m* does not involve *den* or elements of *m*. We obtain the desired truth-condition only when we enrich SEM by specifying a particular *den*. For instance, combining the final clause of SEM with (i)-(vi) yields that *m* is true only if the distances from the robot to landmarks u_1 and u_2 are 1 and 5, respectively. Varying *den* would yield a different truth-condition.

Admittedly, SEM does not iterate in a manner comparable to Tarski's recursive clauses for the logical connectives. But neither does (*). Nevertheless, (*) isolates a compositional mechanism that applies uniformly to arbitrary atomic sentences and denotation relations.

Similarly, SEM isolates a compositional mechanism that applies uniformly to arbitrary point-landmark maps and denotation relations.

The differences between SEM and (*) are as revealing as the similarities. (*) treats predicates as "true of" objects. The "truth of" relation, or its converse "satisfaction," crucially informs Tarski's treatment of complex predicates and objectual quantifiers. As Davidson argues ([2005], pp. 98-163), it is difficult to see how we can handle such locutions without deploying satisfaction. Yet satisfaction plays no role in SEM. The only clause of SEM that even remotely resembles (*) is the third: "$den_1(s_i)$ has properties $den_2(s_i)$." Even this clause differs profoundly from (*). As Davidson emphasizes ([2005], pp. 141-163), (*) decomposes a linguistic expression into two distinct elements ("*a*" and "*F*") with two distinct semantic roles (*denoting some object* versus *being true of certain objects*). SEM imposes no comparable decomposition upon s_i or upon any other element of *m*. SEM does not treat *m* as composed of representational elements that are true of objects. Apparently, then, we can provide a compositional semantics for point-landmark maps without attributing anything like Tarskian predicational structure to them.

In this way, SEM facilitates comparison between cartographic and logical mental representation. SEM describes how *m*'s representational properties depend upon those of its parts, but it eschews the characteristic elements of Tarskian semantics, such as the satisfaction relation or Tarskian inductive clauses governing the logical connectives. Thus, SEM embodies fundamentally different compositional mechanisms than standard Tarskian semantics for the predicate calculus. SEM raises many questions, such as what it is for *m* to be associated with a given metric *d*, a given denotation relation *den*, and a given system of Cartesian coordinates for physical space. I see no reason to think that answering these questions requires the resources of

Tarskian semantics. I conclude that we can provide a satisfying truth-conditional semantics for *m* without attributing logical form, including predicative structure, to it.[13]

As Block ([1981], p. 3) and Pylyshyn ([2003], pp. 281-426) observe, the debate over mental imagery has been plagued by vague, metaphorical, or obscure formulations of the respects in which imagistic and propositional representations have different “formats.” In contrast, our analysis suggests a fairly crisp formulation: cognitive maps in the approximate sense have representationally significant *geometric* structure, as opposed to representationally significant *logical* structure. This contrast is exemplified by SEM. Thus, SEM secures a precise sense in which *m* has a different representational format than a predicate calculus sentence. As far as I know, even the most rigorous analyses of imagistic representation, such as (Kosslyn [1980]) or (Tye [1991]), do not come close to providing a formal semantics that would facilitate an analogous comparison.

Department of Philosophy

University of California

Santa Barbara, CA 93106

rescorla@philosophy.ucsb.edu

Acknowledgments

I am indebted to Jacob Beck, Ned Block, Elisabeth Camp, Warren Goldfarb, Richard Heck, Mary Hegarty, Jack Loomis, Daniel Montello, Charles Parsons, Gualtiero Piccinini, and Aaron Zimmerman for helpful feedback. I also benefited from the comments of three anonymous referees for this journal.

References

Armstrong, D. M. [1973]: *Belief, Truth, and Knowledge*, London: Cambridge University Press.

Balakrishnan, K., Bhatt, R., and Honavar, V. [1998]: 'Spatial Learning and Localization in Animals: A Computational Model and Behavioral Experiments', *Proceedings of the Second European Conference on Cognitive Modeling*, Nottingham: Nottingham University Press.

Balakrishnan, K., Bousquet, O., and Honavar, V. [1999]: 'Spatial Learning and Localization in Rodents: A Computational Model of the Hippocampus and its Implications for Mobile Robots', *Adaptive Behavior*, **7**, pp. 173-216.

Barsalou, L. [1998]: 'Perceptual Symbol Systems', *Behavioral and Brain Sciences*, **22**, pp. 557-660.

Beeson, P., Murarka, A., and Kuipers, B. [2006]: 'Adapting Proposal Distributions for Accurate, Efficient Mobile Robot Localization', *IEEE International Conference on Robotics and Automation*.

Bermudez, J. L. [1998]: *The Paradox of Self-Consciousness*, Cambridge, MA: MIT Press.

---. [2003]: *Thinking without Words*, Oxford: Oxford University Press.

Block, N. [1981]: 'Introduction', in N. Block (*ed.*), *Imagery*, Cambridge, MA: MIT Press.

---. [1983]: 'Mental Pictures and Cognitive Science', *The Philosophical Review*, **92**, pp. 499-541.

---. [1995]: 'The Mind as the Software of the Brain', in E. Smith and D. Osherson (*eds*), *An Invitation to Cognitive Science*, vol. 3, Cambridge, MA: MIT Press.

Braddon-Mitchell, D. and Jackson, F. [2007]: *Philosophy of Mind and Cognition*, 2nd ed., Cambridge, MA: Blackwell.

Brooks, R. [1990]: ‘Elephants Don’t Play Chess’, *Robotics and Autonomous Systems*, **6**, pp. 3-15.

Burge, T. [2003]: ‘Perceptual Entitlement’, *Philosophy and Phenomenological Research*, **67**, pp. 503-48.

---. [2005]: ‘Disjunctivism and Perceptual Psychology’, *Philosophical Topics*, **33**, pp. 1-78.

Byrne, P., Becker, S., and Burgess, N. [2007]: ‘Remembering the Past and Imagining the Future: A Neural Model of Spatial Memory and Imagery’, *Psychological Review*, **114**, pp. 340-75.

Camp, E. [2007]: ‘Thinking with Maps’, *Philosophical Perspectives*, **21**, pp. 145-82.

Campbell, J. [1994]: *Past, Space, and Self*, Cambridge, MA: MIT Press.

Casati, R. and Varzi, A. [1999]: *Parts and Places*, Cambridge, MA: MIT Press.

Choset, H., Lynch, K., Hutchinson, S., Kantor, G., Burgard, W., Kavraki, L., and Thrun, S. [2005]: *Principles of Robot Motion*, Cambridge, MA: MIT Press.

Chrisley, R. [1993]: ‘Connection, Cognitive Maps, and the Development of Objectivity’, *Artificial Intelligence*, **7**, pp. 329-54.

Clark, A. [1997]: *Being There*, Cambridge, MA: MIT Press.

Clark, A., and Grush, R. [1999]: ‘Towards a Cognitive Robotics’, *Adaptive Behavior*, 7, pp. 5-16.

Cummins, R. [1996]: *Representations, Targets, and Attitudes*, Cambridge. MA: MIT Press.

Davidson. D. [2005]: *Truth and Predication*. Cambridge, MA: Harvard University Press.

Dennett, D. [1981]: ‘Two Approaches to Mental Images’, in N. Block, (*ed.*), *Imagery*, Cambridge, MA: MIT Press.

Devitt, M. [2006]: *Ignorance of Language*, Oxford: Clarendon Press.

Eilan, N., McCarthy, R., and Brewer, B. [1999]: *Spatial Representation*, Oxford: Oxford University Press.

Evans, G. [1982]: *The Varieties of Reference*, Oxford: Oxford University Press.

Fodor, J. [1975]: *The Language of Thought*, New York: Thomas Y. Crowell.

---. [1987]: *Psychosemantics*, Cambridge, MA: MIT Press.

---. [1991]: 'Replies', in B. Loewer (*ed.*), *Meaning in Mind*, Cambridge, MA: Blackwell.

---. [2003]: *Hume Variations*, Oxford: Clarendon Press.

---. [2007]: 'The Revenge of the Given', in B. McLaughlin and J. Cohen (*eds*), *Contemporary Debates in Philosophy of Mind*, Malden: Blackwell.

Fodor, J., and Pylyshyn, Z. [1988]: 'Connectionism and Cognitive Architecture: A Critical Analysis', *Cognition*, **28**, pp. 3-71.

Gallistel, C. R. [1990]: *The Organization of Learning*, Cambridge, MA: MIT Press.

---. [1994]: 'Space and Time', in N. J. Makintosh (*ed.*), *Animal Learning and Cognition*, San Diego: Academic Press.

---. [1998]: 'Insect Navigation: Brains as Symbol-Processing Organs', in S. Sternberg and D. Scarborough (*eds.*), *An Invitation to Cognitive Science*, vol. 4, Cambridge, MA: MIT Press.

---. [2008]: 'Dead Reckoning, Cognitive Maps, Animal Navigation, and the Representation of Space: An Introduction', in M. Jeffries and W.-K. Yeap (*eds*), *Robotics and Cognitive Approaches to Spatial Mapping*, Berlin: Springer.

Genesereth, M., and Nilsson, N. [1987]: *Logical Foundations of Artificial Intelligence*, Los Altos: Morgan Kaufmann.

Godfrey-Smith, P. [2006]: Mental Representation, Naturalism, and Teleosemantics, in G.

Macdonald and D. Papineau (*eds*), *Teleosemantics*, Oxford: Oxford University Press.

Grush, R. [2001]: 'Self, World, and Space: On the Meaning and Mechanisms of Egocentric and Allocentric Spatial Representation', *Brain and Mind*, **1**: pp. 59-92.

---. [2004]: 'The Emulation Theory of Representation: Motor Control, Imagery, and Perception', *Behavioral and Brain Sciences*, **27**, pp. 377-42.

Hayes, P. [1979]: 'The Logic of Frames', in D. Metzing (*ed.*), *Frame Conceptions and Text Understanding*, Berlin: de Gruyter.

---. [1985]: 'Some Problems and Non-problems in Representation Theory', in R. J. Brachman and H. J. Levesque (*eds*), *Readings in Knowledge Representation*, Los Altos: Morgan Kaufmann.

Heck, R. [2007]: 'Are There Different Kinds of Content?', in B. McLaughlin and J. Cohen (*eds*), *Contemporary Debates in Philosophy of Mind*, Malden: Blackwell.

Hull, C. [1930]: 'Knowledge and Purpose as Habit Mechanisms', *Psychological Review*, **37**, pp. 511-25.

Johnson-Laird, P. [2004]: 'The History of Mental Models', in K. Manktelow and M. C. Chung, (*eds*), *Psychology of Reasoning: Theoretical and Historical Perspectives*, New York: Psychology Press.

Kitchin, R. [1994]: 'Cognitive Maps: What Are They and Why Study Them?', *Journal of Environmental Psychology*, **14**, pp. 1-19.

Kosslyn, S. [1980]: *Image and Mind*, Cambridge: Harvard University Press.

Kuipers, B. [2000]: 'The Spatial Semantic Hierarchy', *Artificial Intelligence*, **119**, pp. 191-233.

---. [2008]. 'An Intellectual History of the Spatial Semantic Hierarchy', in M. Jefferies and W. Yeap (*eds*), *Robot and Cognitive Approaches to Spatial Mapping*. Oxford: Oxford

University Press.

Kuo, A. [2005]: ‘An Optimal State Estimation Model of Sensory Integration in Human Postural Balance’, *Journal of Neural Engineering*, **2**, pp. 235-49.

Levesque, H. J., Reiter, R., Lespérance, Y., Lin, F., and Scherl, R. [1997]: ‘GOLOG: A Logic Programming Language for Dynamic Domains’, *Journal of Logic Programming*, **31**, pp. 59-84.

Levitt, T.S., and Lawton, D. T. [1990]: ‘Qualitative Navigation for Mobile Robots’, *Artificial Intelligence*, **44**, pp. 305-60.

McCarthy, J., and Hayes, P. [1969]: ‘Some Philosophical Problems from the Standpoint of Artificial Intelligence’, in B. Meltzer and D. Michie (*eds*), *Machine Intelligence 4*, Edinburgh: Edinburgh University Press.

McDermott, D. [2001]: *Mind and Mechanism*, Cambridge, MA: MIT Press.

McDowell, J. [1994]: *Mind and World*, Cambridge, MA: Harvard University Press.

Menzel, R., Gregger, U., Smith, A., Berger , S., Brandt, R., Brunke, S., Bundrock, G., Hülse, S., Plümpe, T., Schaupp, F., Schüttler, E., Stach, S., Stindt, J., Stollhoff, N., and Watzl, S. [2005]: ‘Honey Bees Navigate According to a Map-like Spatial Memory’, *Proceedings of the National Academy of Sciences*, **102**, pp. 3040-5.

Peacocke, C. [1992]: *A Study of Concepts*, Cambridge, MA: MIT Press.

Pinker, S. [2005]: ‘So How Does the Mind Work?’, *Mind and Language*, **20**, pp. 1-24.

Pylyshyn, Z. [1984]: *Computation and Cognition*, Cambridge, MA: MIT Press.

---. [2003]: *Seeing and Visualizing*, Cambridge, MA: MIT Press.

O’Keefe, J., and Nadel, L. [1978]: *The Hippocampus as a Cognitive Map*, Oxford: Clarendon University Press.

Raftopoulos, A. and Müller, V. [2006]: 'The Phenomenal Content of Experience', *Mind and Language*, **21**, pp. 187-219.

Rao, R., Olshausen, B., and Lewicki, M. (*eds.*) [2002]: *Probabilistic Models of the Brain: Perception and Neural Function*, Cambridge, MA: MIT Press.

Redish, A. D. [1999]: *Beyond the Cognitive Map*, Cambridge, MA: MIT Press.

Redish, A. D., and Touretzky, D. [1997]: 'Navigating with Landmarks: Computing Goal Locations from Place Codes', in K. Ikeuchi and M. Veloso (*eds*), *Symbolic Visual Learning*. Oxford: Oxford University Press.

Remolina, E. and Kuipers, B. [2001]: 'A Logical Account of Causal and Topological Maps', *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, San Francisco: Morgan Kaufman.

Rescorla, M. [Forthcoming a]: 'Chrysippus's Dog as a Case Study in Non-linguistic Cognition', in R. Lurz (*ed.*), *Philosophy of Animal Minds*, Cambridge: Cambridge University Press.

---. [Forthcoming b]: 'Predication and Cartographic Representation', *Synthese*.

Rey, G. [1995]: 'A Not "Merely Empirical" Argument for a Language of Thought', *Philosophical Perspectives*, **9**, pp. 201-22.

Russell, S. and Norvig, P. [2003]: *Artificial Intelligence: A Modern Approach*, Upper Saddle River: Prentice Hall.

Shanahan, M. [1997]: 'Noise, Non-determinism, and Spatial Uncertainty', *Proceedings of the American Association for Artificial Intelligence*, pp. 153-58.

Shanahan, M., and Witkowski, M. [2001]: 'High-level Robot Control through Logic', *Intelligent Agents VII*, New York: Springer-Verlag.

Shepard, R. and Chipman, S. [1970]: 'Second-order Isomorphism of Internal Representations:

Shapes of States', *Cognitive Psychology*, **1**, pp. 1-17.

Shettleworth, S. [1998]: *Cognition, Evolution, and Behavior*, Oxford: Oxford University Press.

Sloman, A. [1978]: *The Computer Revolution in Philosophy*, Hassocks: The Harvester Press.

Smith, R. C. and Cheeseman, P. [1986]: 'On the Representation and Estimation of Spatial Uncertainty', *International Journal of Robotics Research*, **5**, pp. 56-68.

Stelerny, K. [2003]: *Thought in a Hostile Word*, Malden: Blackwell.

Strawson, P. [1959]: *Individuals*, London: Routledge.

Thomas, N. [2007]: 'Mental Imagery', in E. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, <plato.stanford.edu>.

Thrun, S. [2000]: 'Towards Programming Tools for Robots that Integrate Probabilistic Computation and Learning': *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco: Institute of Electrical and Electronics Engineers.

---. [2003]: 'Robot Mapping: A Survey', in G. Lakemeyer and B. Nebel (*eds*), *Exploring Artificial Intelligence in the New Millennium*, Boston: Morgan Kaufmann.

Thrun, S., Burgard, W., and Fox, D. [2005]: *Probabilistic Robotics*, Cambridge, MA: MIT Press.

Tolman, E. [1948]: 'Cognitive Maps in Rats and Men', *Psychological Review*, **55**, pp. 189-208.

Tootell, R. B., Silverman, M. S., Switkes, E., and de Valois, R. L. [1982]: 'Deoxyglucose Analysis of Retinotopic Organization in Primate Striate Cortex', *Science*, **218**, pp. 902-4.

Touretzky, D., and Redish, A. D. [1996]: 'A Theory of Rodent Navigation Based on Interacting Representations of Space', *Hippocampus*, **6**, pp. 247-70.

Trullier, O., Wiener, S., Berthoz, A., and Meyer, J.-A. [1997]: 'Biologically Based Artificial Navigation Systems: Review and Prospects', *Progress in Neurobiology*, **51**, pp. 483-544.

Tye, M. [1991]: *The Imagery Debate*, Cambridge, MA: MIT Press.

Wehner, R. [2003]: 'Desert Ant Navigation: How Miniature Brains Solve Complex Tasks', *Journal of Comparative Physiology*, **189**, pp. 579-88.

Wittlinger, M., Wehner, R., and Wolf, H. [2007]: 'The Desert Art Odometer: A Stride Integrator That Accounts for Stride Length and Walking Speed', *Journal of Experimental Biology*, **210**, pp. 198-207.

Notes

[1] For general philosophical discussion of spatial representation, see (Bermudez [1998], [2003]; Campbell [1994]; Eilan, et al. [1999]; Evans [1982]; Godfrey-Smith ([2006])). Heck ([2007]) discusses cognitive maps in relation to the debate over nonconceptual content. Chrisley ([1993]), Clark ([1997]), and Clark and Grush ([1999]) discuss computational models of navigation. To my knowledge, no philosopher has discussed the specific models I emphasize, which deploy Bayesian probability theory defined over cognitive maps. For instance, Clarke ([1997]) emphasizes "behavior-based robotics" (e.g. Brooks [1990]), which eschews "internal models" of the external world.

[2] Cf. (Strawson [1959]; Evans [1982]; Bermudez [1998]; Grush [2001]).

[3] Here are four other difficulties. First, when piloting with respect to particular objects, one must re-identify those objects across time, which the robotics literature dubs the *correspondence problem* or the *data association problem*. Second, and more generally, a distribution of properties within the egocentric reference frame may fit equally well with several hypotheses about one's location, which is known in the psychological literature as *perceptual aliasing*. Third, a position may look different across time due to factors like changing lighting conditions.

This is known as *perceptual variability*. Fourth, perceptual input is noisy, so observations may not match what the map predicts even if the map is correct.

[4] Similar criticisms apply to the neural model of spatial memory presented in (Byrne, et al. [2007]). For example, the model posits an initial "training phase" in which hippocampal neuron firing rates are exogenously set to appropriate values encoding the animal's actual allocentric location (p. 372). Perhaps for this reason, Byrne, et al. emphasize that their model does not seek 'to account for learning in a biologically plausible manner' (p. 351).

[5] Motor command u_t is a vector (v_t ω_t), where v_t is the robot's commanded translational velocity at t and ω_t its commanded angular velocity. Sensor reading z_t is a vector containing a signature, an egocentric distance, and an egocentric bearing for each landmark detected at t.

[6] Neuroscience reveals the presence of "topographic maps" in the brain. For instance, (Tootell, et al. [1982]) shows that early-vision in the macaque involves a retinotopically organized area of the visual cortex. However, few people think that animal navigation, in general, draws upon such a literal spatial map.

[7] Proponents of imagistic representations typically maintain that, although such representations are not literally laid out in physical space, they are somehow "spatial." Kosslyn ([1980]) claims that mental images are laid out in "functional space," a conception that Tye ([1991], pp. 32-45) elucidates and defends. In response, Fodor ([2003], p. 36) asks, "How can there be *spatial* relations among patterns (or whatever) in a space that is itself merely functional?". I think that this question poses a real challenge to Kosslyn and Tye. The notion "spatial relation" seems closely tied to physical space. However, the abstract character of modern mathematical geometry shows that geometric structure is *not* closely tied physical space. That is why I focus in the text

upon *geometric* rather than *spatial* structure. A broadly "functionalist" analysis seems *prima facie* much more plausible for the former than for the latter.

[8] This basic idea here, that certain representations (purport to) replicate structural properties of what they represent, is ancient. Johnson-Laird ([2004]) provides a helpful overview. Gallistel ([1990]) develops the theme in conjunction with cognitive maps.

[9] Many other representations employed in robot navigation share these features. For instance, roboticists often employ *occupancy grids*, which decompose space into small cells and assign each cell a binary value indicating whether it is occupied. Each occupancy grid is endowed with metric structure, and that structure carries representational import, as described by (b).

[10] I neglect arguments from generative linguistics for a level of Logical Form (LF) that represents quantifier scope. These arguments clearly do not apply to animal or robot navigation.

[11] One might regard topological representation as prior to metric representation. Kuipers's model assigns sentential representation an essential role in constructing the topological map. Among other things, then, a satisfactory treatment of these questions must address the relation between topological and metric representation. For one view of that relation, see (Kuipers [2008]).

[12] atan2 is an emendation of the arctan function: atan2(y, x) = arctan(y/x) if x>0; sign(y)(π-arctan($|y/x|$) if x<0; 0 if x=y=0; sign(y)$\pi/2$ if x=0 and y$\neq$ 0. The emendations are to ensure the intuitively correct bearing.

[13] This contrast with Tarskian semantics marks another important commonality with ordinary concrete maps. In (Rescorla [forthcoming b]), I argue that concrete maps do not feature predication, as analyzed by Tarski.