

Maps in the Head?

Michael Rescorla

Any creature that travels through space needs some ability to navigate. The psychologist Edward Tolman (1948) proposed that rats navigate using *cognitive maps*. His proposal flouted the behaviorist consensus of his day, which sought to explain all mental and behavioral phenomena in terms of stimulus-response associations, without recourse to mental representations. Numerous scientists have subsequently pursued the cognitive map hypothesis as applied to diverse species, with especially notable contributions by O'Keefe and Nadel (1978) and Gallistel (1990). What could it possibly mean to say that an animal has a map inside of its head? And why should we believe any such thing? The present entry will pursue these questions.

Map-based Navigation

Scientists standardly distinguish four main animal navigation strategies. Each strategy is an instance or an analogue of a human navigation strategy:

Beaconing, i.e. travel to a goal using sensory input that emanates from the goal. For example, you might walk towards a nearby tree, using its visually perceived distance to guide your approach. Beaconing has limited utility, because it only helps you travel to a destination that currently impinges on your sensory apparatus.

Route following: We frequently navigate by following a series of instructions (e.g. "Turn left at the fork in the road"). Similarly, many species conform to sensorimotor routines, whereby a

specific sensory stimulation triggers a specific behavioral response. For example, honeybees can store retinal “snapshots” of the environment as seen from various locations, and they can learn to fly in some direction when confronted with a snapshot (Collett and Collett 2002). Honeybees can chain together such sensorimotor routines: an initial stimulus triggers some motor behavior until a new stimulus triggers a new motor behavior, and so on.

Dead reckoning uses self-motion cues to maintain a running record of position. Dead reckoning is ubiquitous among non-human vertebrates and invertebrates (Gallistel 1990: 57-102). Even the humble desert ant has impressive dead reckoning capacities. Using dead reckoning, the desert ant can travel long, circuitous routes and then return directly home along a straight path.

Map-based navigation, i.e. navigation using a cognitive map. This is the most controversial of the four navigation strategies. Scientists continue to debate the extent, if any, to which cognitive maps figure in human and non-human navigation. What counts as “map-like” mental representation also remains unclear.

As Bermúdez (1998: 203-207) and Kitchin (1994) emphasize, a persistent problem in the scientific literature is that different authors use the phrase “cognitive map” in different ways. In (Rescorla, 2009), I distinguished two notable usages. A *cognitive map in the loose sense* is a mental representation that represents geometric aspects of the environment. Those aspects may be *metric* (e.g. distances and angles), *topological* (e.g. connectedness and adjacency), or otherwise. A *cognitive map in the strict sense* is a mental representation that has the same basic representational properties and mechanisms as ordinary concrete maps. A cognitive map in the strict sense has the same representational 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.

Evidence for Cognitive Maps in the Loose Sense

To defend the existence of cognitive maps, scientists usually cite evidence that animals take novel detours and shortcuts. Tolman (1948) argued thusly, based on his study of rats traveling through mazes. Kramer (1957) independently argued along similar lines for pigeons navigating through nature. A recurring experimental paradigm in both laboratory and field work is to displace the animal to an unfamiliar release point within a familiar environment. In many cases, the animal travels directly from the release point to a goal (e.g. the location of some reward). This is advanced as evidence that the animal located itself on a cognitive map and thereby computed a route from the release point to the goal. Researchers have developed these ideas in considerable detail for various species (Gallistel 1990; Jacobs and Menzel 2014; O'Keefe and Nadel 1978; Weiner et al. 2011), including rats (Geva-Sagiv et al. 2015; Morris 1981), pigeons (Bingman 2011), (Schiffner and Wiltschko 2013; Wallraf, 2005), and honeybees (Menzel and Greggers 2015; Cheeseman et al. 2014).

Critics respond that novel routes can often be explained without positing cognitive maps (Bennett 1996; Collett, Chittka, and Collett 2013; Mackintosh 2002). For example, Cheung et al. (2014) claim that any novel routes taken by honeybees may simply reflect route following, in which the bee tries to minimize discrepancies between stored snapshots and its current panoramic view of the environment. Critics also adduce experimental evidence that certain animals, including honeybees (Dyer 1991; Wehner et al. 1990) and rats (Whishaw 1991), sometimes do *not* take shortcuts when doing so would be advantageous. Critics conclude that cognitive maps are not needed to explain whatever novel routes animals actually take.

Despite such controversies, a broad albeit non-unanimous consensus has emerged over the past few decades: many mammals take some novel routes best explained by positing mental representations that represent metric aspects of the environment, including distances between locations (Gallistel and Matzel 2013; Jacobs and Menzel 2014; Weiner et al. 2011). A good illustration is an experiment on golden hamsters performed by Etienne et al. (1998). Hamsters lived in a 180-cm-diameter circular arena containing four identical, symmetrically placed cylinders. They learned that one cylinder, distinguishable from the others only by its location, contained food. Upon being lured to a location within the arena, hamsters could navigate directly towards the food-containing cylinder, even in the dark. How did the hamsters do this? They had no access to sensory cues that might inform beaconing or route following. Dead reckoning surely played a large role, enabling the hamster to estimate its current position. But dead reckoning taken on its own would not enable the hamster to compute a course to the cylinder. Evidently, the hamster mentally represented the cylinder's position, integrated that representation with the deliverances of dead reckoning, and thereby computed a route to the cylinder.

A recent field illustration of map-based mammalian navigation features the Egyptian fruit bat (Tsoar et al. 2011). Using GPS technology, researchers tracked flight paths of bats. When displaced to a desert area 44 km outside their normal flight range, bats showed a remarkable ability to navigate directly to one of two goals: a familiar feeding site or the home cave, depending on whether they were hungry. This astonishing feat cannot be explained in terms of beaconing (distinctive sensory cues to the goal were not available), dead reckoning (bats were transported inside a cloth bag, so they could not dead reckon), or route following (the release point was far outside the bats' visually familiar area, so it is not plausible that they had acquired suitable sensorimotor routines). Tellingly, bats were initially quite disoriented when released

inside a large crater, but they became well-oriented upon exiting the crater. The bats apparently determined their current position using visual landmarks that were only visible upon exiting the crater (e.g. city lights). On that basis, they computed a route to the goal. An explanation along these lines presupposes that bats have a large-scale representation of landmark locations.

Localization and Mapping

How do animals construct and update cognitive maps? The answer depends heavily on psychological, physiological, and environmental details for each species. We are only beginning to understand these matters. However, some general features of map-based mammalian navigation are relatively well-established.

Any animal that represents the spatial layout of its environment must have at its disposal mental *coordinates* that represent locations in the environment (Gallistel and Matzel 2013). Researchers standardly distinguish between *allocentric* and *egocentric* coordinates. Allocentric coordinates are anchored to the external environment (e.g. the sun or the animal's home). Egocentric coordinates are anchored to the creature's body. Map-based navigation uses allocentric rather than egocentric coordinates, because it requires representations of landmark position that remain relatively stable as the animal moves.

We know through introspection that humans perceive the egocentric distances and directions of objects. There is also experimental evidence that many animal species, from insects to mammals, perceive egocentric distance and direction (Kral 2003). Mammalian navigation draws crucially upon perceptual estimates of egocentric position:

- *Localization*, i.e. estimation of one's own allocentric position. Dead reckoning is a widely employed localization strategy. However, dead reckoning is fallible and noisy,

rendering it unreliable over long periods. Accordingly, many species employ an additional localization strategy called *piloting* (Gallistel 1990). Piloting estimates current position by observing landmarks whose positions are represented on the allocentric cognitive map. Given the egocentric positions of sufficiently many landmarks, and given the allocentric positions of those landmarks, it is basic trigonometry to compute one's current position.¹

- *Mapping*. Perception supplies egocentric estimates of landmark positions, and dead reckoning supplies an allocentric estimate of one's position and orientation. Combining these estimates, one can form allocentric estimates of landmark positions. One thereby converts egocentric spatial representations into an allocentric cognitive map. Gallistel (1990) reviews evidence that map-construction along these lines occurs in various species.

Localization and mapping deploy *coordinate transformations* between egocentric spatial representations and allocentric representations. A coordinate transformation converts a representation in one coordinate system into a representation in a different coordinate system.

Mapping relies on dead reckoning, which becomes increasingly unreliable as uncorrected errors accumulate. Piloting can correct those errors, but piloting presupposes an allocentric cognitive map. For that reason, localization and mapping are intertwined. The animal must estimate its current allocentric position while *simultaneously* estimating allocentric landmark positions. To solve this simultaneous estimation problem, the animal must somehow integrate the dead reckoning estimate with perceptually-based egocentric position estimates and with any past allocentric landmark position estimates.

An analogous estimation problem arises in *robotics*, where it is called the *simultaneous localization and mapping problem* (SLAM). An autonomous navigating robot must estimate its own position along with the positions of salient landmarks. The most successful robotics solution is grounded in *Bayesian decision theory*, a mathematical theory of reasoning and decision-making under uncertainty. On a Bayesian approach, the robot maintains a probability distribution over possible maps of the environment, using self-motion cues and sensory input to update probabilities as it travels through space. Bayesian robotic navigation algorithms have achieved notable success (Thrun, Burgard, and Fox 2005). Given how well Bayesian solutions to SLAM work within robotics, it is natural to conjecture that some animals use Bayesian inference when navigating (Gallistel, 2008; Rescorla 2009). Scientists have recently begun offering Bayesian models of animal navigation (Cheng et al. 2007; Cheung et al. 2012; Madl et al. 2014; Madl et al. 2016; Penny, Zeidman, and Burgess 2013). The models look promising, although this research program is still in its infancy.

Neurophysiological Underpinnings

How are cognitive maps realized in the brain? What neural processes implement mapping, localization, and route planning? While we do not have complete answers to these questions, we know a lot about the neural states and processes that underlie mammalian map-based navigation.

O'Keefe and Dostrovsky (1971) discovered that the rat hippocampus contains *place cells*, each responding selectively to a specific spatial location. On that basis, O'Keefe and Nadel (1978) proposed that the hippocampus provides the neural substrate for cognitive mapping. This work generated a huge surge of interest in cognitive maps, especially among neuroscientists.

Ensuing research discovered several other notable cells (Moser, Kripff, and Moser 2008):

- Several areas in the rat brain contain *head direction cells* (Taube 2007). A head direction cell fires when the rat's head is at a certain angle with respect to an external reference direction.
- The rat entorhinal cortex contains *grid cells* (Hafting et al. 2005), each responding selectively to multiple spatial locations in the available environment. The locations where a cell fires form a periodic grid that covers the environment. Different cells generate grids with different scales and different orientations. Metric information about the physical environment can be extracted from the firing patterns of grid cells (Moser and Moser 2008).
- The rat entorhinal cortex contains *border cells* (Solstad et al. 2008), each of which fires when the rat is near a border oriented in a certain direction.

Edvard Moser, May-Britt Moser, and John O'Keefe shared the 2014 Nobel Prize in Physiology or Medicine for their work in this area.

Neuroscientists have developed detailed mathematical models describing how place cells, grid cells, and other such cells support navigation (e.g. Bush et al. 2015; Cheng and Frank, 2011; McNaughton, et al. 2006; Solstad, Moser, and Einvell 2006). In many cases, the models are reasonably well-integrated with cognitive-level theories that allude to cognitive maps, dead reckoning, localization, mapping, path-planning, coordinate transformations, probability distributions, etc. However, just as we are far from completely understanding the mental processes through which mammals form, update, and deploy their cognitive maps, so are we far from completely understanding the neural implementation of those mental processes.

Cartographic Representation

The phrase “cognitive map” naturally suggests that these mental representations resemble ordinary concrete maps in important respects. To what extent, and in what ways, do cognitive maps resemble concrete maps? Are they cognitive maps in the *strict* sense?

Although many aspects of cartographic representation remain ill-understood, we can isolate four important properties of the concrete maps employed within human society:

- (1) *Maps represent geometric aspects of physical space.* A map represents the layout of entities in space. The map thereby represents geometric relations among those entities. Maps vary in precisely which geometric relations they represent. City maps represent metric structure, while subway maps only represent topological structure.
- (2) *Maps have veridicality-conditions.* A map is evaluable as veridical or non-veridical, depending on how the world is. The map is veridical only when it correctly represents geometric relations among entities. Thus, it is veridical under certain conditions, non-veridical under others.
- (3) *Maps have geometric structure.* A map does not merely *represent* geometric structure. The map itself is geometrically structured. For example, a city map has metric structure. A map’s geometric structure is representationally significant, as clause (4) elucidates.
- (4) *A map is veridical only if it replicates salient geometric aspects of the region that it represents.* Informally, a map purports to replicate relevant geometric aspects of physical space. More formally, a map is veridical only if there exists a structure-preserving function from the map to the region that it represents. For example, a city map is veridical only when distances on the map are proportional to distances in the physical environment.

I do not advance (1)-(4) as a finished theory of cartographic representation, but rather as a springboard for further inquiry. For present purposes, the key point is that a mental representation should share properties (1)-(4) to the extent that it counts as a cognitive map *in the strict sense*. Does animal navigation feature mental representations with properties (1)-(4)?

We have already canvassed evidence that mammalian navigation uses cognitive maps in the loose sense, i.e. mental representations with property (1). However, this commonality is less impressive than it may initially appear, because it hinges on the unexplicated term “represent.” Philosophers and psychologists have proposed many different theories of representation (e.g. Burge 2010; Davidson 2001; Fodor 1990; Gallistel 1990; Millikan 1984), and the theories vary wildly in how much is required for one entity to “represent” another. Saying without further elucidation that a map “represents” geometric aspects of the environment does not tell us much.

In effect, (2) provides one way of glossing (1). An advantage of (2) over (1) is that (2) uses the relatively well-understood notion *veridicality-condition*, which has long been a major focus of philosophical research into representation. Many important mental states have veridicality-conditions. To illustrate:

- Beliefs are evaluable as true or false. For example, my belief *that Barack Obama is president* is true iff Barack Obama is president. So beliefs have truth-conditions.
- Desires are evaluable as fulfilled or unfulfilled. For example, my desire *to eat chocolate* is fulfilled only if I eat chocolate. So desires have fulfillment-conditions.
- Perceptual states are evaluable as accurate or inaccurate. For example, my perceptual experience *as of a red sphere located a certain egocentric distance from me* is accurate only if a red sphere is located a certain egocentric distance from me. So perceptual states have accuracy-conditions.

Truth, fulfillment, and accuracy are species of veridicality. So beliefs, desires, and perceptual states all have veridicality-conditions. As (2) asserts, concrete maps also have veridicality-conditions. If you do not recognize that concrete maps may be veridical or non-veridical (that a map may or may not *correctly* represent the region that it represents), then you have missed a fundamental aspect of our navigational and cartographic practices.

Do *cognitive* maps have veridicality-conditions? It is far from clear how to answer this question. The strategy I will now pursue is to reflect on the role played by veridicality-conditions within psychological explanation.

The Explanatory Role of Veridicality-conditions

Intentional explanation is explanation that cites veridicality-conditions or representational properties that contribute to veridicality-conditions. The most familiar example is *folk psychology*: our everyday practice of citing beliefs, desires, and other mental states to explain mental and behavioral outcomes. Folk psychology routinely identifies mental states through their veridicality-conditions. For example, we might identify a belief as the belief *that Obama is president*, thereby specifying a condition that must obtain for the belief to be true. Or we might identify a desire as a desire *to eat chocolate*, thereby specifying a condition that must obtain for the desire to be fulfilled.

Taking inspiration from folk psychology, cognitive science offers numerous intentional explanations. For example, *perceptual psychology* studies how the perceptual system transits from proximal sensory stimulations (e.g. retinal stimulations) to perceptual states that estimate shapes, sizes, colors, locations, and other observable properties. A perceptual state is veridical only if perceived objects have the estimated shapes, sizes, colors, locations, and other such

properties. The science identifies perceptual states through representational properties that contribute to veridicality-conditions --- e.g. through specific shapes, sizes, color, and locations estimated by the perceptual system (Burge, 2010; Rescorla 2015). Intentional explanations of perception have proved enormously fruitful, illuminating a wide range of perceptual phenomena.

Does cognitive science offer successful intentional explanations of animal navigation?

While there is room for healthy debate here, my own view is that intentional discourse contributes serious explanatory value at least when applied to *mammalian* map-based navigation. Scientific research into mammalian navigation hinges upon a straightforward thought: mammalian cognitive maps are *estimates*. They estimate geometric aspects of the environment, including the spatial layout of landmarks. An estimate is evaluable as veridical or non-veridical. Cognitive science identifies mammalian cognitive maps at least partly through their veridicality-conditions, i.e. through the conditions that they estimate as obtaining. By identifying cognitive maps in this way, the science delineates systematic patterns of interaction between allocentric cognitive maps, egocentric perceptual states, and actions.

To illustrate, consider *coordinate transformations* between allocentric and egocentric representations. As we have seen, these coordinate transformations underwrite mammalian localization and mapping. They also underwrite the interface between cognitive maps and action: to travel towards a goal, the animal often converts its allocentric representation of the goal into an egocentric representation with immediate consequences for action (Gallistel 1999). Overall, coordinate transformations figure pivotally in scientific theorizing about mammalian navigation (Madl et al., 2015; Shadmehr and Mussa-Ivaldi 2012: 35-66; Wilber et al. 2014), including some impressively detailed computational models (Byrne and Becker 2007; Sheynikhovich et al. 2009; Touretzky and Redish 1996).

A coordinate transformation *preserves veridicality* when it carries veridical representations into veridical representations. Virtually all scientific treatments presume that mammalian coordinate transformations typically preserve veridicality, at least approximately. Given a veridical allocentric cognitive map, the relevant coordinate transformations typically yield veridical (or approximately veridical) egocentric representations of landmark positions. Given veridical egocentric representations of landmark positions, and given a veridical estimate of one's own allocentric position and orientation, the relevant coordinate transformations typically yield veridical (or approximately veridical) allocentric representations of landmark positions. Approximate veridicality-preservation is a core presupposition of scientific research into mammalian navigation, including the aforementioned computational models. This core presupposition, although not often made explicit, guides the construction of detailed theories describing how cognitive maps interact with perception and action. It also helps us explain the extraordinary success with which mammals navigate. Veridical egocentric perceptual estimates, combined with veridical estimates of the animal's position and heading, tend to cause veridical allocentric maps. Veridical allocentric cognitive maps tend to cause veridical egocentric representations, which in turn tend to cause successful actions.²

Researchers have developed this explanatory strategy with increasing experimental and theoretical sophistication over ensuing decades. The strategy presupposes that cognitive maps have veridicality-conditions. After all, a coordinate transformation can only preserve veridicality if the representations over which it operates have veridicality-conditions.

I favor a broadly *scientific realist* viewpoint: explanatory success is a *prima facie* guide to truth. From a scientific realist viewpoint, successful intentional explanation provides reason to attribute veridicality-conditions. For example, the explanatory success of perceptual psychology

provides reason to attribute veridicality-conditions to perceptual states (Burge 2010; Rescorla 2015). Likewise, successful intentional explanations of mammalian navigation provide reason to attribute veridicality-conditions to mammalian cognitive maps. I conclude that (2) applies to mammalian cognitive maps.³

Bayesian models of mammalian navigation provide further evidence for this conclusion. The basic idea behind Bayesian models is that the navigational system maintains a probability distribution over a *hypothesis space*. Each hypothesis represents some aspect of the spatial environment. One such hypothesis might represent that a certain landmark has a certain allocentric location. Another hypothesis might represent that the animal itself has a certain allocentric location. Hypotheses of this kind are incorporated into cognitive maps, which estimate overall spatial layout. The probability assigned to a cognitive map is determined by the probabilities assigned to component hypotheses. The navigational system regularly updates its probabilities in light of perceptual input and self-motion cues. In this manner, localization and mapping become exercises in statistical inference.

How should we understand the “hypotheses” to which probabilities get assigned? Current Bayesian models identify the hypotheses through representational properties that contribute to veridicality-conditions. For example, when we identify a hypothesis *as* representing that a landmark has a certain allocentric location, we cite a condition that must be satisfied for the overall cognitive map to be veridical: that the landmark has the hypothesized location. We thereby identify the hypothesis in intentional terms. Bayesian models describe how probabilities over hypotheses *as identified in intentional terms* change in light of perceptual input and self-motion cues. The science presupposes that mammalian navigation deploys cognitive maps with veridicality-conditions, and it describes probabilistic inference over hypotheses *identified by how*

the hypotheses contribute to cognitive maps' veridicality-conditions. Hence, the science presuppose that (2) applies to mammalian cognitive maps. The success of the Bayesian research program provides further reason to attribute veridicality-conditions to mammalian cognitive maps. As the research program accrues more explanatory success, the case for an intentional analysis of mammalian navigation should grow commensurately stronger.

Geometrically Structured Mental Representations?

I now consider the representational *format* of cognitive maps. Do they have representationally significant geometric structure? More precisely, do they share properties (3) and (4) with ordinary concrete maps?

Even if we grant that an animal mentally represents geometric structure, why should we hold that the animal uses geometrically structured mental representations? What would it even mean to ascribe geometric structure to a mental representation? Pylyshyn (2007: 80-81) warns against the *intentional fallacy* --- the fallacy of confusing properties of a representation with properties of what it represents. Mental representations of color are not colored. Mental representations of loudness are not loud. Why should mental representations of geometric structure be geometrically structured? Surveying a range of navigational behaviors, Pylyshyn concludes (2007: 178): “however impressive these behaviors may be, and even when they reveal something about the content of the representation (what information must have been encoded), they reveal little about the form of the representation involved that makes it maplike.”

Any theorist who posits cognitive maps in the strict sense must answer Pylyshyn's challenge. Note furthermore that cognitive maps do not seem to have literal spatial structure in the brain. In particular, nearby place cells do not correspond to nearby locations in physical

space. Thus, any satisfying theory of geometrically structured cognitive maps must articulate a notion of “geometric structure” much more abstract than literal spatial structure in the brain.

In this connection, it is helpful to recall the highly abstract character of modern mathematical geometry. The standard modern procedure is to isolate axioms of geometric structure, such as metric or topological structure. For example, 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 real numbers such that, for all elements a and b in X

$$d(a, b) \geq 0$$

$$d(a, a) = 0$$

$$d(a, b) = d(b, a)$$

$$d(a, c) \leq d(a, b) + d(b, c).$$

A metric space may be composed of any entities whatsoever. What matters is not the set X itself, but rather the relations between X 's elements. Moral: *any* entities may be enveloped within a metric structure.

In principle, then, it makes sense to talk about geometric structure over the mental coordinates that appear on a cognitive map. Indeed, if C is a set of mental coordinates, then there are infinitely many metric spaces (C, d) . Obviously, most of these metric spaces are irrelevant to the animal's navigation and hold no interest for cognitive science. Can we isolate some useful sense in which the animal's psychology instantiates a geometric structure over C ? If so, does the resulting geometric structure contribute to veridicality-conditions as (4) dictates?

Several authors have explored how something like (3) and (4) might be true of cognitive maps (e.g. Brecht et al. 2014; Heck 2007; Muller, Stead, and Pach 1996; Rescorla 2009; Shea 2014; Terrazas et al. 2005). The basic idea behind most treatments is that functionally significant

neural or psychological relations among mental coordinates induce geometric structure over the cognitive map, where this structure represents geometric relations in physical space. For example, Shea (2014) suggests that place cells may have a co-activation structure that represents proximity relations in physical space. An important task for future scientific and philosophical research is to investigate suggestions along these lines. Doing so should illuminate whether, and in what sense, cognitive maps have representationally significant geometric structure.

Conclusion

Cognitive maps figure pivotally in navigation across a range of species. Numerous navigational phenomena are difficult or impossible to explain unless we posit cognitive maps in the loose sense. Animal navigation therefore provides strong evidence for a broadly representationalist approach to psychology. A vast interdisciplinary literature spanning many decades provides great insight into the nature of cognitive maps, their neurophysiological underpinnings, and the psychological processes in which they participate. We understand quite a bit about cognitive maps, as compared with most other mental representations posited by philosophers and scientists. Nevertheless, numerous questions remain about their format, content, psychological role, and neural basis. This entry will have served its purpose if you feel moved to investigate further.

Notes

¹ Honeybees can perceptually estimate the egocentric distances and directions of landmarks. As Burge (2010: 508) emphasizes, the resulting perceptual estimates do not appear to exert much impact upon honeybee localization. Honeybee localization seems to operate primarily through dead reckoning, with periodic resets of the odometer when the bee encounters a familiar landmark (Srinivasan 2011).

² Do coordinate transformations between egocentric and allocentric representations play a significant role in honeybee navigation? The answer is unclear. By comparison with scientific theorizing about mammalian navigation, scientific theorizing about honeybee navigation assigns relatively little weight to coordinate transformations. For

example, as mentioned in note 1, honeybees do not seem to localize based upon egocentric perceptually-based representations of landmark distances and directions. Thus, my argument in the main text does not readily generalize from mammals to honeybees. In general, it is unclear whether attribution of veridicality-conditions adds explanatory value to the scientific study of honeybee navigation (Burge 2010: 509-514; Rescorla 2013).

³ Philosophers sometimes suggest that non-intentional discourse can reproduce any explanatory benefits afforded by intentional explanation (Field 2001; Stich 1983). They claim that we can eliminate intentional locutions from cognitive science, without explanatory loss. In (Rescorla 2015), I argue that such are claims implausible when applied to intentional explanations of human perception. I think they are also implausible when applied to intentional explanations of mammalian map-based navigation. For present purposes, I must leave my assessment undefended.

References

- Bennett, A. (1996) “Do Animals Have Cognitive Maps?”, *The Journal of Experimental Biology* 199, pp. 219-224.
- Bingman, V. (2011) “Making the Case for the Intelligence of Avian Navigation,” in R. Menzel and J. Fischer (eds.) *Animal Thinking*, Cambridge: MIT Press.
- Brecht, M., Ray, S., Burgalossi, A., Tang, Q., Schmidt, H., and Naumann, R. (2014) “An Isomorphic Mapping Hypothesis of the Grid Representation.” *Philosophical Transactions of the Royal Society B* 369, pp. 1-10.
- Burge, T. (2010) *Origins of Objectivity*, Oxford: Oxford University Press.
- Bush, D., Barry, C., Manson, D., and Burgess, N. 2015. “Using Grid Cells for Navigation.” *Neuron* 87, pp. 507-520.
- 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. 340-375.
- Cheeseman, J., Millar, C., Greggers, U., Lehmann, K., Pawley, M., Gallistel, C. R., Warman, G., and Menzel, R. (2014) “Way-finding in Displaced Clock-shifted Bees

-
- Proves Bees Use a Cognitive Map,” *Proceedings of the National Academy of Sciences* 111, pp. 8949-8954.
- Cheng, K., Shettleworth, S., Huttenlocher, J., and Rieser, J. (2007) “Bayesian Integration of Spatial Information,” *Psychological Bulletin* 13, pp. 625-637.
- Cheng, S., and Frank, L. (2011) “The Structure of Networks that Produce the Transformation from Grid Cells to Place Cells.” *Neuroscience* 197, pp. 293-306.
- Cheung, A., Collett, M., Collett, T., Dewar, A., Dyer, F., Graham, P., Mangan, M., Narendra, A., Philippides, A., Stürzl, W., Webb, B., Wystrach, A., and Zeif, J. (2014) “Still No Convincing Evidence for Cognitive Map Use by Honeybees,” *Proceedings of the National Academy of Sciences* 111, pp. e4397-e4397.
- Cheung, A., Ball, D., Milford, M., Wyeth, G., and Wiles, J. (2012) “Maintaining a Cognitive Map in Darkness: the Need to Fuse Boundary Knowledge with Path Integration,” *PLoS Computational Biology* 8, p. e1002651.
- Collett, T.S., and Collett, M. (2002) “Memory Use in Insect Visual Navigation,” *Nature Reviews Neuroscience* 3, pp. 542-552.
- Collett, M., Chittka, L., and Collett, T. (2013) “Spatial Memory in Insect Navigation,” *Current Biology* 23, pp. R789-R800.
- Davidson, D. 2001. *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press.
- Dyer, F. (1991) “Bees Acquire Route-based Memories but not Cognitive Maps in a Familiar Landscape,” *Animal Behavior* 41, pp. 239–246.
- Etienne, A., Maurer, R., Berlie, J., Reverdin, B., Rowe, T., Georgakopoulos, J., and Séguinot, V. (1998) “Navigation through Vector Addition,” *Nature* 396, pp. 161–164.
- Field, H. (2001) *Truth and the Absence of Fact*, Oxford: Oxford University Press.

-
- Fodor, J. (1990) *A Theory of Content and Other Essays*, Cambridge: MIT Press.
- Gallistel, C. R. (1990) *The Organization of Learning*, Cambridge: MIT Press.
- . (1999) "Coordinate Transformations in the Genesis of Action," in B. Bly and D. Rumelhart (eds.) *Cognitive Science: Handbook of Perception and Cognition*, 2nd ed., New York: Academic 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.
- Gallistel, C. R., and Matzel, L. (2011) "The Neuroscience of Learning: Beyond the Hebbian Synapse," *Annual Review of Psychology* 64, pp. 169-200.
- Geva-Sagiv, M., Las, L., Yovel, Y., and Ulanovsky, N. (2015) "Spatial Cognition in Bats and Rats: From Sensory Acquisition to Multiscale Maps and Navigation," *Nature Reviews Neuroscience* 16, pp. 94-108.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. (2005) "Microstructure of a Spatial Map in the Entorhinal Cortex," *Nature* 436, pp. 801-6.
- Heck, R. (2007) "Are There Different Kinds of Content?," in B. McLaughlin and J. Cohen (eds.) *Contemporary Debates in Philosophy of Mind*, Malden: Blackwell.
- Jacobs, L., and Menzel, R. (2014) "Navigation Outside of the Box: What the Lab Can Learn from the Field and What the Field Can Learn from the Lab," *Movement Ecology* 2, pp. 1-22.
- Kitchin, R. (1994) "Cognitive Maps: What Are They and Why Study Them?," *Journal of Environmental Psychology* 14, pp. 1-19.
- Kral, K. (2003) "Behavioral-Analytic Studies of the Role of Head Movements in Depth

-
- Perception in Insects, Birds, and Mammals,” *Behavioral Processes* 64, pp. 1-12.
- Kramer, G. (1957) “Experiments in Bird Orientation and their Interpretation,” *Ibis* 99, pp. 196-227.
- McNaughton, B., Battaglia, F., Moser, E., and Moser, M.-B. (2006) “Path Integration and the Neural Basis of the ‘Cognitive Map’.” *Nature Reviews Neuroscience* 7, pp. 663-678.
- Mackintosh, N. J. (2002) “Do Not Ask Whether They Have a Cognitive Map, but How They Find their Way about,” *Psicológica* 23, pp. 165-185.
- Madl, T., Chen, K., Montaldi, D., and Trapp, R. (2015) “Computational Cognitive Models of Spatial Memory in Navigation Space: A Review,” *Neural Networks* 65, pp. 18-43.
- Madl, T., Franklin, S., Chen, K., Montaldi, D., and Trapp, R. (2014) “Bayesian Integration of Information in Hippocampal Place Cells,” *PloS One* 9, pp. e89762.
- . (2016) “Towards Real-World Capable Spatial Memory in the LIDA Architecture.” *Biologically Inspired Cognitive Architectures* 16, 87-104.
- Menzel, R., and Greggers, U. (2015) “The Memory Structure of Navigation in Honeybees.” *Journal of Comparative Physiology A* 201, pp. 547-561.
- Millikan, R. (1984) *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.
- Morris, R. (1981) “Spatial Localization Does Not Require the Presence of Local Cues.” *Learning and Motivation* 12, pp. 239-260.
- Moser, E., Kripff, E., and Moser, M.-B. (2008) “Place cells, Grid Cells, and the Brain’s Spatial Representation System,” *Annual. Reviews Neuroscience* 31, pp. 69–89.
- Moser, E., and Moser, M.-B. (2008) “A Metric for Space.” *Hippocampus* 18, pp. 1142-1156.
- Muller, R., Stead, M., and Pach, J. (1996) “The Hippocampus as a Cognitive Graph.” *The*

-
- Journal of General Physiology* 107, pp. 663-694.
- O'Keefe, J., and Dostrovsky, J. (1971) "The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely-moving Rat," *Brain Research* 34, pp. 171–175.
- O'Keefe, J., and Nadel, L. (1978) *The Hippocampus as a Cognitive Map*, Oxford: Clarendon University Press.
- Penny, W., Zeidman, P., and Burgess, N. (2013) "Forward and Backward Inference in Spatial Cognition," *PLoS Computational Biology* 9, pp. e1003383.
- Pylyshyn, Z. (2007) *Things and Places*, Cambridge: MIT Press.
- Rescorla, M. (2009) "Cognitive Maps and the Language of Thought," *The British Journal for the Philosophy of Science* 60, pp. 377-407.
- . (2013) "Millikan on Honeybee Navigation and Communication," in D. Ryder, J. Kingsbury, and K. Williford (eds.) *Millikan and Her Critics*, Malden: Wiley-Blackwell.
- . (2015) "Bayesian Perceptual Psychology," in M. Matthen (ed.) *The Oxford Handbook of the Philosophy of Perception*, Oxford: Oxford University Press.
- Schiffner, I., and Wiltschko, R. (2013) "Development of the Navigational System in Homing Pigeons: Increase in Complexity of the Navigational Map," *Journal of Experimental Biology* 216, pp. 2675–2681.
- Shadmehr, R., and Mussa-Ivaldi, S. (2012) *Biological Learning and Control*, Cambridge: MIT Press.
- Shea, N. (2014) "Exploitable Isomorphism and Structural Representation." *Proceedings of the Aristotelian Society* 114, pp. 123-144.
- Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arlea, A., and Gerstner, W. (2009) "Is There a Geometric Module for Spatial Orientation? Insights From a Rodent Navigation Model,"

-
- Psychological Review* 116, pp. 540-566.
- Solstad, T., Boccara, C., Kropff, E., Moser, M.-B., and Moser, E. (2008) "Representation of Geometric Borders in the Entorhinal Cortex." *Science* 322, pp. 1865-1868.
- Solstad, T., Moser, E., and Einevoll, G. (2006) "From Grid Cells to Place Cells: A Mathematical Model." *Hippocampus* 16, pp. 1026-1031.
- Srinivasan, M. (2011) "Honeybees as a Model for the Study of Visually Guided Flight, Navigation, and Biologically Inspired Robotics," *Physiological Review* 91, pp. 413-460.
- Stich, S. (1983) *From Folk Psychology to Cognitive Science*, Cambridge: MIT Press.
- Taube, J. (2007) "The Head Direction Signal: Origins and Sensory-Motor Integration," *Annual Review of Neuroscience* 30, pp. 181-207.
- Terrazas, A., Krause, M., Lipa, P., Gothard, K., Barnes, C., and McNaughton, B. (2005) "Self-Motion and the Hippocampal Spatial Metric." *The Journal of Neuroscience* 25, pp. 8085-8096.
- Thrun, S., Burgard, W., and Fox, D. (2005) *Probabilistic Robotics*, Cambridge: MIT Press.
- Tolman, E. (1948) "Cognitive Maps in Rats and Men," *Psychological Review* 55, pp. 189-208.
- Touretzky, D., and Redish, A. D. (1996) "A Theory of Rodent Navigation Based on Interacting Representations of Space," *Hippocampus* 6, pp. 247-70.
- Tsoar, A., Nathan, R., Bartan, Y., Vyssotski, A., Dell'Omo, G., Ulanovsky, N. (2011) "Large-scale Navigational Map in a Mammal," *Proceedings of the National Academy of Sciences* 108, pp. E718-E724.
- Wallraf, H. (2005) *Avian Navigation: Pigeon Homing as a Paradigm*, New York: Springer.
- Wehner, R., Sleuler, S., Nievergelt, C., Sha, D. (1990) "Bees Navigate by Using Vectors and Routes rather than Maps," *Naturwissenschaften* 10, pp. 479-482.

-
- Weiner, J., Shettleworth, S., Bingman, V., Cheng, K., Healy, S., Jacobs, L., Jeffery, K., Mallot, H., Menzel, R., and Newcombe, N. (2011) “Animal Navigation: A Synthesis,” in R. Menzel and J. Fischer (eds.) *Animal Thinking*, Cambridge: MIT Press.
- Whishaw, I. (1991) “Latent Learning in a Swimming Pool Place Task by Rats: Evidence for the Use of Associative and Not Cognitive Mapping Processes,” *The Quarterly Journal of Experimental Psychology* 43, pp. 83-103.
- Wilber, A., Clark, B., Forster, T., Tatsuno, M., and McNaughton, B. (2014) “Interaction of Egocentric and World-Centered Reference Frames in the Rat Posterior Parietal Cortex,” *The Journal of Neuroscience* 34, pp. 5431-5446.

Further reading

John O’Keefe and Lynn Nadel’s seminal *The Hippocampus as Cognitive Map* (1978) is still well worth reading, especially for its comparison of map-based navigation with other navigational strategies. C. R. Gallistel’s classic *The Organization of Learning* (1990) remains an outstanding introduction to cognitive maps. For a more recent overview, see (Gallistel and Matzel, 2013). Madl et al. (2015) survey numerous computational models of navigation. For an accessible introduction to grid cells, see Edvard Moser and May-Britt Moser’s “Understanding the Cortex through Grid Cells,” in *The Future of the Brain*, eds. Gary Marcus and Jeremy Freedman (Princeton: Princeton University Press, 2014). Neurophysiological models of mammalian navigation are helpfully discussed in Talfan Evans, Andrej Bicanski, Daniel Bush, and Neil Burgess’s “How Environment and Self-Motion Combine in Neural Representations of

Space,” *The Journal of Physiology* (forthcoming) and Lisa Giacomo, May-Britt Moser, and Edvard Moser’s “Computational Models of Grid Cells,” *Neuron* 71 (2011): 589-603.

Cognitive maps have attracted attention from philosophers, although not as much attention as they deserve. Gareth Evans’s unfinished masterpiece *The Varieties of Reference* (Oxford: Clarendon Press, 1982) argues that cognitive maps undergird fundamental aspects of human thought. Chapter 10 of Tyler Burge’s landmark *Origins of Objectivity* (2010) analyzes cognitive maps from a representationalist perspective, with particular emphasis on relations to perceptual representation. Other notable philosophical treatments include José Luis Bermúdez’s *The Paradox of Self-Consciousness* (Cambridge: MIT Press, 1998), Elisabeth Camp’s “Thinking with Maps,” *Philosophical Perspectives* 21 (2007): 145-82, and John Campbell’s *Past, Space, and Self* (Cambridge: MIT Press, 1994).

Bayesian modeling of the mind has recently emerged as an impressive scientific research program, especially as applied to perception and motor control. Many high quality introductions are available, such as Nick Chater and Michael Oaksford’s anthology *The Probabilistic Mind* (Oxford: Oxford University Press, 2008), David Knill and Whitman Richards’s anthology *Perception as Bayesian Inference* (Cambridge: Cambridge University Press, 1996), and Daniel Wolpert’s survey article “Probabilistic Models in Human Sensorimotor Control,” *Human Movement Science* 26 (2007): 511-524. In (Rescorla 2015), I argue that Bayesian perceptual psychology offers intentional explanations. I argue along similar lines for Bayesian models of motor control in “Bayesian Sensorimotor Psychology,” *Mind and Language* 31 (2016): 3-36.