



Mapping content: why cognitive maps are non-conceptual mental states

Arieh Schwartz¹ · Nir Fresco^{2,3}

Received: 15 September 2023 / Accepted: 11 January 2025
© The Author(s) 2025

Abstract

Cognitive maps play a crucial role in mammalian navigation. They provide the organism with information about its own location and the locations of landmarks within known environments. Cognitive maps have yet to receive ample attention in philosophy. In this article, we argue that cognitive maps should not be understood along the lines of conceptual mental states, such as beliefs and desires. They are more plausibly understood to be non-conceptual. We clarify what is at stake in this claim, and offer two empirically-informed arguments in its favor. Both arguments submit that cognitive maps are probably non-conceptual because their representational structure seems to differ from that of conceptual mental states.

Keywords Cognitive maps · Navigation · Non-conceptual content · Belief

1 Introduction

Humans and other mammals navigate familiar environments using internal allocentric representations of space and self-position known as *cognitive maps*. Our knowledge of where we are in the environment, of where landmarks are, of our heading, and of distance relations between locations is furnished by cognitive maps. While

✉ Arieh Schwartz
a.g.schwartz@lse.ac.uk

Nir Fresco
nfresco@bgu.ac.il

¹ Centre for Philosophy of Natural and Social Sciences, London School of Economics and Political Science, London, UK

² Department of Cognitive and Brain Sciences, Ben-Gurion University of the Negev, Be'er Sheva, Israel

³ School of Brain Sciences and Cognition, Ben-Gurion University of the Negev, Be'er Sheva, Israel

the sciences of navigation have seen astonishing, Nobel-Prize-winning advances in recent years, cognitive maps have been paid relatively little philosophical attention.¹ It remains unclear, for example, whether we can understand cognitive maps within the framework of the propositional attitudes that we normally use to rationally explain behavior.

We understand one another, from a rational perspective, by ascribing states with *conceptual content*—propositional attitudes, beliefs and desires principally among them. Ascribing mental states with conceptual content allows us to predict and explain what people will do *as the conclusions of chains of reasoning*. This style of explanation is as powerful as it is ubiquitous. It routinely allows for the successful prediction and explanation of events far-off in space and time; for example, that someone will board a plane, travel across the country, and arrive at the local airport at a specific time on a specific date (Fodor, 1987).

But not all mental states fit within this framework of explanation. There are persuasive arguments that at least some perceptual experiences are *non-conceptual*—are representationally different from conceptual mental states (see, e.g., Evans, 1982; Peacocke, 1992; Crane, 1992; Martin, 1992; Heck, 2000; Dretske, 1995; Fodor, 2007; Roskies, 2008). Views diverge on what exactly makes a mental state conceptual rather than non-conceptual (Bermúdez, 2007). The original debate about whether perceptual experience is non-conceptual turned on whether it has the same representational *structure* as belief; namely, on whether the content of every perceptual experience is (and must be) a proposition structured via concepts (Heck, 2000, 2007). Evans argued that experience has a different kind of content than belief (1982: 227). McDowell argued that experience and belief must have the same sort of content, because experience justifies belief (McDowell, 1994: 160).

Evans proposed that cognitive maps are like perceptual experiences in being non-conceptual (1982, Ch. 6). The main goal of this article is to provide some empirically-informed arguments for this conclusion. Cognitive maps do have representational structure. But their structure is importantly different from that of conceptual mental states, e.g. beliefs.

It is significant that cognitive maps are non-conceptual, for both metaphysical and epistemological reasons. It is important metaphysically, because it raises the question of how cognitive maps interface with conceptual mental states, given their difference in representational structure or format. It is also important metaphysically, because it raises the possibility that we share navigational mental states with creatures who lack our concepts.² Moreover, it is important epistemically, because it entails that

¹ Some important exceptions to this generalization include (Bechtel, 2016; Bermúdez, 2003; Godfrey-Smith, 2013; Heck, 2007; Rescorla, 2009, 2017; Shea, 2018).

² Although we appeal to animal models, this appeal does not already assume that we share navigational mental states with these creatures for two reasons. First, we do not simply assume that the same *neurological mechanisms*—place cells, grid cells, and head direction cells—are at work in both humans and non-human animals. The scientific consensus that these mechanisms are shared across species is based on direct experimental evidence. Second, in addition to employing this common machinery, human navigational mental states might call on conceptual operations not at work in non-human animals. Thanks to an anonymous reviewer for prompting clarification about this.

cognitive maps cannot support or undermine beliefs in the same way beliefs support or undermine each other (e.g., deductively).

This article has the following structure. Section 2 provides a relatively thorough introduction to science of navigation, in the hope of serving as a resource for future philosophical work. Section 3 explains the distinction between conceptual and non-conceptual mental states. Section 4 evaluates a set of arguments advanced by Heck (2007) that cognitive maps are indeed non-conceptual. To our knowledge, Heck provides the only explicit arguments for this conclusion. Section 5 develops two novel arguments for this view. Section 6 presents a taxonomy of information that may shed light on the relationships between conceptual and non-conceptual mental states. Section 7 concludes.

2 What are cognitive maps?

Cognitive maps were initially proposed by psychologists in the 1940s to explain the observed capacity of animals to take novel short-cuts through mazes (Tolman, 1948). Over the last fifty years, neuroscientists have increasingly uncovered the cellular components of cognitive maps, which turn out to be shared by a wide variety of species. Cognitive maps are supported by several specialized cell types found in the hippocampal formation, including grid cells in the entorhinal cortex, head direction (HD) cells in the subiculum, and place cells in the hippocampus proper. Celebrating such work, the 2014 Nobel Prize in physiology and medicine was awarded to John O’Keefe, and May-Britt and Edvard Moser, for the discovery of place cells and grid cells, respectively. As cognitive maps are empirical posits, philosophical discussion of cognitive maps should attend as closely as possible to the relevant scientific work. The next subsection explains the initial postulation of cognitive maps, subsequent discoveries of place cells, grid cells, and HD cells, and sketches current thinking about how the operations of these cells support navigation.³ (Readers already familiar with this work may skip ahead to Sect. 3.)

2.1 Cognitive maps and the science of navigation

Tolman and colleagues conducted rodent maze-running experiments in the 1940’s that would undermine the reigning behaviorist assumption that all behavior consists in stimulus–response pairings established through conditioning.⁴ In one key experi-

³While the term ‘cognitive map’ sometimes refers specifically to the activity of place cells in the hippocampus proper, this discussion treats place cells, grid cells, and HD cells as co-equal parts of the same functional unit, a broader cognitive map system. For cognitive maps in the narrow sense cannot be formed or used reliably in navigation without the operation of grid cells and HD cells. Sometimes, grid cells and HD cells are also said to compose an independent *path integration system*, since they can support deductive (dead) reckoning, AKA path integration, even in unfamiliar (unmapped) environments. Path integration is navigation to a goal by keeping track of distance and direction to be traveled, rather than locations of landmarks in objective space. It will not matter if the reader chooses to reserve the term ‘cognitive map’ for the activities of place cells specifically, so long as it is remembered that the arguments in later sections also apply to representations supported by grid cells and HD cells.

⁴This section is indebted to the excellent text, *Human Spatial Navigation*, by Ekstrom et al. (2018).

ment, for example, Tolman and colleagues observed that the greatest proportion of rats trained to a particular maze would, when their accustomed route had been blocked, and many novel paths were added to the maze, choose the unfamiliar path that lead directly to where the reward had been located during the training phase. These results could not be explained by appeal to conditioning, for the simple reason that the rats had never previously encountered the shortcut they took (Tolman et al., 1946). In other experiments using a cross-shaped maze, Tolman and colleagues altered the start locations of trained rats, who nonetheless navigated directly to the goal location. This showed that the rats' knowledge of the reward's location was allocentric rather than egocentric. Tolman described the contrast between his position and that of the behaviorist picturesquely.

[T]he central office [of the rat] is far more like a map control room than it is like an old-fashioned telephone exchange. The stimuli, which are allowed in, are not connected by just simple one-to-one switches to the outgoing responses. Rather, the incoming impulses are usually worked over and elaborated into a tentative, cognitive-like map of the environment. And it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release. (Tolman, 1948: 193)

Positing stored map-like representations of the environment is strongly motivated by the observation of shortcutting behavior. It would be more than twenty years before the machinery responsible for coding this spatial information could be located neurally, through *single-unit recording*.

Single-unit recording is an invasive and direct measure of individual neurons' spiking activity. Microelectrodes thin enough to pass between neurons are precisely driven into place right outside the cell bodies of neurons, from where they can record the cells' electrical discharges. Using this method, O'Keefe and Destrovsky (1971) recorded from cells in the dorsal hippocampus of a rat trying to determine what stimuli excited them. They discovered eight cells with a unique response profile.

These 8 units... did not fire unless the animal was in a moderate state of arousal, [and] was situated in the correct part of the testing platform... These 8 units then appear to have preferred spatial orientations... These findings suggest that the hippocampus provides the rest of the brain with a spatial reference map. (O'Keefe & Destrovsky, 1971: 72-74)

Cells of this sort came to be known as "place cells" because of their tendency to fire only when the animal occupied a particular location (or place field). With enough place cells tuned to particular locations, the brain can register where in allocentric space the organism is, and can index where in the environment important landmarks are, such as sources of food, etc. (O'Keefe & Nadel, 1978).

What stimuli support the establishment and maintenance of these cells' place fields? Place fields can be anchored relative to distal cues in different sensory modalities (visual, auditory, olfactory, etc.) (O'Keefe & Conway, 1976). Importantly, manipula-

tion of distal cues results in precisely corresponding changes in cells' place fields. For example, Muller and Kubie (1987) found that rotating the position of a white place card on the wall of a cylindrical environment produced corresponding angular rotations of the place fields of place cells in trained rats. Enlargement of an environment caused corresponding enlargements of the place fields of place cells in trained rats. Imposition of a barrier within a familiar environment caused the removal of the place field in which the barrier had been placed, even though the barrier occupied only a small portion of the place field. The adaptability of place cells' place fields suits the activity of these cells to capture what is unique to each environment.

Recent evidence suggests that cognitive maps are also used in route planning. While running, a rat's place cells tend to fire only when the rat occupies their place fields. Single-unit recordings indicate, however, that during pauses in navigation, ensembles of place cells fire, in temporally-compressed sequence, whose place fields form possible future trajectories (Dragoi & Tonegawa, 2011; Foster, 2017). Researchers hypothesize that this neural activity, known as *preplay*, is a way of simulating possible routes for the organism to take (see Shea, 2018, ch. 5 for discussion).

For obvious ethical reasons, it is usually impossible to obtain single-unit recordings from human subjects. Sometimes, however, such recordings are a necessary part of preparation for brain surgery, (e.g. to localize where epileptic patients' seizures originate). And in such cases, patients may consent to further testing. The first unit recordings of place cells in humans were reported by Ekstrom et al. (2003). These researchers recorded from the hippocampi of seven epileptic patients playing a virtual-reality taxi-driving game. They discovered 31 neurons that fired selectively to locations in the virtual environment.

More recently, human fMRI studies have reported evidence of hippocampal place activity. Unlike the orientation detectors in V1, place cells that are next to each other in the brain need not have proximate place fields (more on this shortly). This lack of topological organization initially stood in the way of detecting place activity using fMRI, since the smallest *voxel* or unit of the brain scan would contain many place cells with distant place fields (Ekstrom et al., 2018). However, using *multivariate pattern analysis*, a method of comparing distributed activation patterns across multiple voxels, Hassabis et al. (2009) were able to accurately decode from fMRI images which of four locations subjects occupied within a virtual environment.

Although "the cognitive map does share some of the qualities of a cartographic map in that it must represent relative directions and distances of important landmarks" there are also important differences between them (Ekstrom et al., 2018: 10). Place cells are not organized topologically. It is basically random which place cells will be tuned to which place fields (O'Keefe & Nadel, 1978).⁵ By contrast, cartographic maps *are* topologically organized. Proximate parts of a cartographic map vehicle represent proximate places in the mapped area. Relatedly, that a place cell's place field is in the South-West corner of one environment, for example, carries no implications about where the cell's place field will be in a different environment, or indeed

⁵That place cells lack topological organization does not, however, preclude relationships between spatial relations in the world from being represented by analogous relations of *co-activation* between place cells (see Shea, 2018). Thanks to an anonymous reviewer for prompting this clarification.

whether the cell will be active in this other environment at all. Place cells *remap* to different locations in different environments. Moreover, whereas cartographic maps can represent environments ranging from the indefinitely large to the indefinitely small, cognitive maps are specific to smaller environments; rooms, neighborhoods, but not the world (Ekstrom et al., 2018: 47–48).

Lying under the cerebral cortex, deep near the center of the brain, the hippocampus does not receive sensory information directly. It was initially unclear how place cells' place fields could be established, and how distance relations could be computed. The entorhinal cortex is adjacent to the hippocampus and connects it to sensory areas. The discovery of cells in the entorhinal cortex that have uniform grids of place fields was a major turning-point in the science. It is now believed that grid cells together with head direction cells provide input to, "the hippocampus [which] combines this information with landmark-related representation to produce the cognitive map" (Ekstrom et al., 2018: 61).

In 2005, members of the Moser lab recorded cells from the entorhinal cortex of rats foraging in a large cylindrical environment (2 m in diameter) (Hafting et al., 2005). They found that each neuron recorded did not have just one place field, like a place cell, but a regularly spaced and oriented grid-like pattern of firing fields which extended over the whole environment.

In every isolated principle neuron, the firing field formed a grid of regularly tessellating triangles spanning the whole recording surface... When the environment was expanded, the number of activity nodes increased, but their density remained constant... Although spacing, orientation and field size were almost invariant at individual recording locations, spacing and field size increased with distance from the postrhinal border, resulting in more dispersed fields at more ventral electrode positions... Although grids of neighboring cells had similar spacing, field size and orientation, their phases (the vertex locations) were apparently not related. (Hafting et al., 2005: 801-802)

The cellular machinery Hafting et al. uncovered suggests a spatial matrix, or rather a collection of matrices that are imposed on the environment. There are a small number of distinct layers (or *modules*) of grid cells, that each have a uniform scale. The ratio between the scales of modules is constant. The scale of each module is approximately 1.4 times the scale of the module directly above it (Stensola et al., 2012).

Importantly, grids are imposed immediately in novel environments, and no environmental changes affect their spacing (Hafting et al., 2005). These facts strongly suggest that the grid cells play a role in computing distance relations, and inform the establishment of place cells' unique place fields. There is no consensus yet about the mechanisms that establish place cells' place fields. But M.B. Moser and colleagues propose that the operative mechanism might involve the grids of different modules rotating independently across different environments. Independent rotations of modules' grids would provide a means of encoding a vast number of unique combinations, much like the independently rotating dials on a combination lock (Rowland & Moser, 2014; Sugar & Moser, 2019).

Grid cells are not unique to rats. They have been recorded in phylogenetically distant mammals, from Egyptian fruit bats (Yartsev et al., 2011), to humans (Jacobs et al., 2013). In their human recordings, Jacobs and colleagues observed not only the characteristic grid-like activity, but, “these patterns exhibited sixfold (60°) symmetry, indicating that the locations at which these cells activated were arranged in a triangular grid, similar to patterns observed in rodents” (Jacobs et al., 2013).

Navigation using a cognitive map requires some way to keep track of facing direction. Jim Ranck (1984) had discovered certain neurons in the postsubiculum of rats that fired only when the rats’ heads were pointed in specific directions. Ranck named these compass-like neurons head direction cells. HD cells have since been found in several other limbic areas (Taube, 2007). There is no topological relationship between the firing ranges of adjacent head direction cells (Ibid.). While unit recordings of HD cells in humans have not yet been reported, human fMRI studies have shown head direction activity in the same areas of the brain that are known to contain head direction cells in rats (Shine et al., 2016).

In summary, place-based navigation is centrally supported by the operation of three cell types: place cells, grid cells, and HD cells.⁶ The phase constancy of grid cell modules through changes in environment suits grid cell activity to capture objective distance relationships that remain constant across environments. The activity of grid cell modules helps establish the sensitivity of particular place cells to particular locations. The changeability of place cells’ place fields helps them capture the distinctive features of each environment. There is evidence for the operation of place cells, grid cells, and HD cells across diverse mammalian species. Organisms use the distance information from grid modules, location information from place-cells, and heading information from HD cells to navigate to goal locations.

3 The distinction between conceptual and non-conceptual mental states

This section explains (1) why it is generally agreed that conceptual mental states are propositional attitudes; (2) Evans’s way of distinguishing between conceptual and non-conceptual mental states, known as the *content view* (Heck, 2000); (3) a more neutral way of drawing the distinction, known as the *state view* (Heck, 2000); and (4) points of agreement between the content and state views on which the arguments of this article depend.

3.1 Why conceptual mental states are propositional attitudes

The central question in the debate over whether perceptual experience has conceptual content is whether perceptual experience is a propositional attitude. Is McDowell right that, “[i]n experience one takes in, for instance sees, *that things are thus and so*.

⁶Other cell types also contribute, such as boundary cells, speed cells, goal cells, value cells, etc. (Ekstrom et al., 2018). But to keep this section manageable, we have restricted our attention to the three most widely discussed spatial cell types in the hippocampal formation.

That is the sort of thing one can also, for instance, judge” (1994: 9, original emphasis)? Whether experience is a propositional attitude is an interesting and important question in its own right. But many participants in the debate care about this issue because they believe that, while a mental state (or brain state) may have causal powers no matter what kind of content it has, it can only provide the subject with a reason (or justification) if it has the same sort of content belief has, i.e., propositional content (McDowell, 1994; Brewer, 2005; see Heck, 2000 for dissent). On this sort of view, a basic perceptual judgment can only be justified if it directly *imports* (Heck, 2007) some of the content of a perceptual experience.

3.2 The content view

The distinction between conceptual and non-conceptual mental states was introduced by Evans (1982) (Cf. Dretske, 1969). Evans held a particular view about what it takes for a mental state to have conceptual content, i.e. to be a propositional attitude, or what he called a *Thought*. Evans held that every Thought is constituted through the use of separate, indefinitely reusable concepts (conceptual abilities). One concept presents the entity that the mental state is about, and a separate concept presents a feature. Thus, “Thoughts are essentially structured” (Evans, 1982: 102). They are structured out of separate and indefinitely reusable concepts combined in a predicative fashion. On this view, the content view (Heck, 2000), the contents of experiences are metaphysically different from the contents of Thoughts. Conceptual contents are structured from discrete concepts, predicatively combined. Experiences (their contents) are not structured from concepts in this way. Heck’s arguments that cognitive maps are non-conceptual (which we will discuss in Sect. 4) assume the content view.

3.3 The state view

There is a more neutral way of drawing the distinction between conceptual and non-conceptual mental states, known as the state view (Heck, 2000; Cf. Crane, 1992). The state view is more neutral than the content view because it is silent on the metaphysics of conceptual contents (i.e. propositions). The state view begins from the idea that, whatever the metaphysical nature of conceptual contents (propositions), a subject can only occupy a conceptual mental *state*, such as belief, if she grasps the proposition in question entirely using concepts (deployed predicatively, and also logically, if relevant). For example, to believe/hope/fear/etc. that, ‘Professor Jones is sick’, the subject must grasp the person using her PROFESSOR JONES concept, and grasp this person’s health-status, using her SICKNESS concept, so as to apply sickness to Jones.⁷ Non-conceptual mental states, by contrast, are characterized by the fact that they present the world in ways the subject does not entirely grasp via concepts (Cussins, 1990; Martin, 1992; Crane, 1992, 2012). Since grasping a content via concepts requires *possession* of the relevant concepts, state view theorists often argue that experience is non-conceptual by arguing that experience can present the world in ways the subject lacks relevant concepts for (e.g., Martin, 1992).

⁷ Single quotes enclose propositions. Small caps name concepts.

3.4 Points of agreement between the content view and state view

The content view and the state view agree that to be in a conceptual mental state, to stand in a propositional attitude, the subject must grasp the proposition in question entirely using concepts, (combined in the relevant fashion). The views differ because the content view makes the additional assumption that conceptual contents, i.e. propositions, are constituted by the relevant complexes of concepts (Heck, 2000: 485).

Here is the point our later arguments will rest on. When a subject occupies a conceptual mental state about a worldly relation, (for example a spatial relation), she will grasp this relation using a distinct concept. Both the state view and the content view agree about this. By contrast, there is another species of representation in which (some) worldly relations are represented *not* by discrete representational elements, such as concepts, but by structural *relations* between representational elements. We will follow Ramsey (2007) in calling this S-representation. S-representations represent via isomorphism (or homomorphism). Entities are represented because they map to elements in the representation, but (some) relations *between* entities are represented because they map to structural relations *between* their corresponding representation elements (cf. Neander, 2017: ch. 8; Shea, 2018, ch. 5; Cummins, 1996; Ramsey, 2007: ch. 3; Burge, 2018). We submit that grasping a worldly relation by means of a representational relation, (i.e. by using an S-representation), is not the same as grasping it via a distinct representation element, such as a concept. Thus, if cognitive maps are S-representations, not all their content is grasped via concepts. Thus, if cognitive maps are S-representations, they are non-propositional and non-conceptual.⁸ Moreover, (as will become relevant in Sects. 5 and 5.1), if a representation has structure, and yet is *holistic* in the sense that (some) relations between entities are represented automatically via representation of the entities, then the relevant representation is an S-representation.⁹

4 Heck on cognitive maps

Heck (2007) provides three ingenious arguments that cognitive maps are non-conceptual mental states. Heck's purpose is to advance a broader argument by analogy that perceptual experiences are non-conceptual because they are relevantly similar to cognitive maps (Heck, 2007: 128). We agree with Heck's proposal that cognitive maps are non-conceptual, but we question the conclusiveness of their arguments. In broad strokes, these arguments seem inconclusive because they tend to assume that cognitive maps are like everyday cartographic maps. For example, when introducing cognitive maps, Heck says simply that,

⁸Thanks to an anonymous reviewer for helpful suggestions about this section.

⁹The alternative way for a representation to be holistic is for it to lack structure entirely (Fodor, 2007). Photographs might be like this, (but for serious objections, see Burge [2018]). But cognitive maps seem to have structure in any event (see Sect. 4).

Each of us has a mental map of our surroundings that places locations we encounter relative to other, known locations. Now, cognitive maps are obviously representational, and the term ‘map’ is used here because the representations in question are thought to be very much like more familiar sorts of maps. That is to say: We have and employ a mode of storing information about topographic features of our environment that is... unified and, one might say, organic. (Heck, 2007: 125)

If cognitive maps are genuinely like everyday cartographic maps, then they are S-representations, and are non-conceptual. Genuine maps represent (some) worldly relations because relations in the map mirror them. But the premise that cognitive maps are relevantly map-like needs to be supported. We will argue that cognitive maps are indeed S-representations in Sect. 5.

Heck endorses the content view (see Sect. 3) of the distinction between conceptual and non-conceptual mental states. Channeling Evans, they write,

The ability to think that a is F must decompose into the abilities to think of a and to think of a thing as F , abilities that are sufficiently distinct that one’s being able to think that a is F may be *explained by* one’s being able to think of a and one’s being able to think of a thing as F . What I am suggesting is thus that the claim that beliefs have conceptual content should be understood as the claim that the contents of belief are structured in this sense. (Heck, 2007: 123, original emphasis)

This quote expresses commitment to the content view. It claims that conceptual *content* is internally structured from distinct concepts. This is not merely a claim about the prerequisites a subject must satisfy to be in a conceptual type of mental *state*.

For ease of exposition, let us call Heck’s three arguments: (1) the argument from indeterminacy; (2) the argument from failures of composition; and (3) the argument from differential revision. We will discuss each in turn.

4.1 From indeterminacy

The argument from indeterminacy trades on the intimate connection between mental states with conceptual content, like beliefs, and the sentences that can be used to express their content. Suppose, for example, that a subject believes that, ‘The leaf is red’. There is no genuine indeterminacy about which sentences express the belief. It is precisely expressed by, “The leaf is red”, “La feuille est rouge”, etc. Arguably, *all* conceptual mental states are like belief in this way. There is no genuine indeterminacy about which sentences express them. The argument from indeterminacy then proceeds as follows.

If one wanted to represent the content of a map as a structured proposition, what structured proposition would it be? *The only plausible answer would seem to be that the content of the map is given by a complete description of the relationships it indicates...* It is an objection [to this view] that there is no unique

such description. At the very least, there will be a question of how to order the who-knows-how-many conjuncts that would occur in it... There is no unique structured proposition that gives the content of a map because there is no such structure in the map; a map lacks the structure present in a verbal description of what it represents. (Heck, 2007: 125-126, added emphasis)

The first step of the argument is this idea about the connection between conceptual contents (propositions) and sentences. There is no genuine indeterminacy about which sentences express the contents of conceptual mental states. If the content of a cognitive map were conceptual, it too would be precisely expressible by at least one (possibly very long) sentence. There would be no indeterminacy about which sentences express a cognitive map. But there *does* seem to be indeterminacy about which sentences would express a cognitive map. Thus, the contents of cognitive maps are not structured propositions. They are non-conceptual.

Let us grant that there is never indeterminacy about which sentences express conceptual mental states. The important question is whether there is more indeterminacy about which sentences express cognitive maps. One might think there is, if one thought that cognitive maps lacked any structure; namely, that every part of the cognitive map contributed to the content of the whole cognitive map in the same way. Heck is indeed of this opinion, writing that cognitive maps do “not decompose in any determinate way into parts. Cognitive maps, that is to say, are icons in Fodor’s sense” (Heck, 2007: 125; Cf. Fodor, 2007). If that were true, then there would be arbitrarily many different equally good linguistic descriptions of any given cognitive map, since cognitive maps would lack any real structural divisions into parts that adequate linguistic expressions of the cognitive maps would have to reflect.

Even assuming that cognitive maps are like everyday cartographic maps, however, we have no reason to assume they lack structure.¹⁰ For, as Camp (2007, 2018) has persuasively argued, cartographic maps really do have representational structure; they have recurrent constituents, put together according to systematic combinatorial operations (see also Clarke, 2023; Burge, 2018). To be sure, cartographic maps differ from sentences because they are S-representations. But both kinds of representations have structure. So it cannot be argued that there is more indeterminacy about which sentences express cognitive maps on the assumption that cognitive maps, being like cartographic maps, lack structure.

Heck suggests that the only reasonable candidate for a structured proposition that gives the content of a cognitive map would be, “a complete description of the relationships it indicates” (126). Heck’s suggestion that there could be a complete description of the relationships indicated by a cognitive map seems to suggest that cognitive maps do have real parts and structure. It is unclear what “complete” could mean otherwise. But, Heck claims, there would still be indeterminacy about which complete descriptions express a cognitive map, because there would at least “be a question about how to order the who-knows-how-many conjuncts [of the description that would uniquely specify the cognitive map]” (126). Thus, there is indeterminacy

¹⁰Thanks to an anonymous reviewer for helpful comments on this point.

about which sentences express a cognitive map, since representations with different orderings of conjuncts would be equally good candidates.

We doubt whether representations with different orderings of conjuncts need to differ on which proposition(s) they express. Of course, sentences with different orderings of conjuncts can be used to *communicate* different things, for example, different sequences of events. For example, “John drove home and got drunk” can be used to convey a different sequence of events than, “John got drunk and drove home.” But such ordering effects are usually taken to be pragmatic rather than semantic because they are cancellable (Grice, 1961).

In sum, if we assume with Heck that cognitive maps are much like cartographic maps, then then they have real structure (Camp, 2007). Thus, we should not base an inference that it is indeterminate which sentences express cognitive maps on the assumption that cognitive maps are Fodorian icons. Moreover, if we assume, with Heck, that it makes sense to suppose there could be a complete linguistic description of the relationships in a cognitive map, this would only leave in question the order in which the relationships are described. But if cognitive maps have real structure, then different orderings in which relationships are described need not be taken to reflect differences in the contents of the descriptions, since ordering effects are pragmatic. On the whole then, we seem to lack conclusive reasons why it should be indeterminate which sentences express cognitive maps.

4.2 From failures of composition

The argument from failures of composition is based on the idea that conceptual contents, propositions, have a structure that allows them to be combined by Boolean operators: and, or, not, if–then, if-and-only-if. This is a core feature of conceptual contents, since it is only because conceptual contents combine via logical operators that they can be steps in any meaningful deductive reasoning.¹¹ Figuring in chains of such reasoning is the main theoretical purpose for which subjects are attributed conceptual contents. In contrast to propositional representations, which are logically-combinable,

One cannot, for example, form arbitrary Boolean combinations of maps: There is no map that is the negation of my cognitive map of Boston; there is no map that is the disjunction of my map and my wife’s; and so forth. If the content of a cognitive map is a structured proposition, why shouldn’t there be maps with such contents? Why can’t the negations of the atomic formulae that figure in the content of a map also figure in its content? Why can’t these formulae be disjoined? (Heck, 2007: 126)

How should we evaluate this argument? It is true that everyday cartographic maps we are familiar with do not allow for representation of abstract logical relationships, like the negation of a content, or the disjunction of one content with another. But as we have already seen (Sect. 2), there are important differences between the external

¹¹ Tautologies, which do not require logical operators, are not meaningful.

maps that cartographers produce, and the cognitive maps that our brains produce. Moreover, some external map systems are possible which *do* allow for the representation of abstract logical information (Camp, 2007; Clarke, 2022), (while remaining genuinely map-like, insofar as they represent worldly spatial relations via analog relations between representational constituents). We might call these latter representational systems *map-extensions*. In short, one might worry that the argument from failures of composition only works on the assumption that cognitive maps are like familiar paper maps, rather than being map-extensions in the relevant sense. But how might cognitive maps be map-extensions, structurally suited to represent logical relationships?

Map-extensions are like ordinary maps, except that they involve additional principles of composition and use. As Camp (2007) suggests, one such representational system might represent negation, more specifically that there are *no* items at a particular location in the map, by that location's being colored white. Cognitive maps could conceivably operate in the same way, if the relative absence of firing from a (tuned-up) place cell represented that the animal is *not* located in the cell's place field. As Camp further suggests, a map-extension might involve further abstract elements that represent disjunctive information. For example, if the label "food" were to alternate flashing in two different locations in a map, this might represent that food is either in one location or the other. A comparison could be drawn to the phenomenon in cognitive map-based navigation known as virtual trial and error (Nadel, 2013). When an animal is moving around a familiar environment, place cells will tend to fire only when the animal occupies their place fields. During pauses in navigation, however, the animal may look one direction and then another, and place cells will fire, in temporally-compressed sequence, that correspond to places within possible trajectories for the animal to take from its current position; as if the animal were virtually trying out different routes to a goal. Perhaps the content of these representations is disjunctive, that the organism could take one route or another.

We are not claiming that cognitive maps really do represent negation and disjunction in these ways, but only that they might. Heck seems to assume that cognitive maps are like everyday cartographic maps, rather than being map-extensions. But this assumption stands in need of further support since there are map-extensions which do have the relevant sort of logical structure.¹²

¹²Thanks to a reviewer for raising the following worry. "It's one thing for a map to represent that an object is not present in a particular location; that's plausibly possible. But it's another thing to take the negation of a whole map. Is this possible? Heck explicitly claims that it isn't, and I don't see how the points in the paper contradict this." We argue that it is nomologically possible for cognitive maps to be map-extensions whose contents comprise constituents that stand in logical relations, (e.g., of disjunction, or negation). If so, then they have the kind of content can, ("can" in the sense of metaphysical possibility), take wide-scope negation. Of course, real life organisms may *contingently* lack a mechanism that serves to negate entire cognitive map contents. (It is unclear what use organisms would have for such a mechanism.) But this lack would be a contingent fact, not a reflection of the *kind* of content cognitive maps have. The contents of cognitive maps would be structured propositions no less open to negation than the contents of sentences, whether or not entire cognitive map contents ever actually get negated.

4.3 From differential revision

The argument from differential revision is difficult to interpret. We will quote the main parts of the argument below, and then discuss one possible interpretation of it.

Suppose, for example, that my map had previously located an object o at location l . Now here I am at l , and o is not to be found; instead, u is there. What to do? It is clear enough what to say if we restrict our attention to the construction of a representation: I should remove the ‘marker’ that indicates o from its position on the map and put a ‘marker’ representing u there; I can then either put the o -marker somewhere else on the map or just leave it off. But if we wish to regard this transformation as a *rational* one... then we must also be able to describe it in terms of *content*, that is, to describe it not just as a change in a representation but as a change in what is represented: previously, I had taken my environment to be thus-and-so; now I take it to be so-and-thus. The relation between the contents of my maps before and after this change cannot naturally be described in terms of... structured propositions... the problem is that moving the o -marker, for example, does not simply change where o is represented as located; o was also located in relation to other objects, and many of those relations— though not necessarily all of them— will have to be changed as well. It is thus not as simple as swapping one conjunct for another: the sorts of changes involved will be on a much larger scale, and— or so I am suggesting— the nature of those changes can only seem obscure so long as one insists upon describing them as if one were describing a change in belief. (Heck, 2007: 126-127, original emphasis)

One interpretation of the argument is along the following lines. Heck explicitly assumes that cognitive maps play a role in the subject’s reasoning about the environment (2007: 126). Given cognitive maps’ presumed role in reasoning, Heck suggests that we need an explanation of how relationships between cognitive maps and other mental states can be big ‘R’ Rational: rational in terms the *representational natures* of the states involved.¹³ Heck provides an example in which they acquire the perceptual belief that a novel entity is in the location that their cognitive map represents a different entity as being. Heck argues that if cognitive maps had propositional structure, i.e. conceptual content, then it would be “obscure” how all of the revisions that would have to be made to their cognitive map in light of the novel belief would be Rational. This obscurity would be due to the fact that entities and relations are represented independently or separately in propositionally-structured representations. Of course,

¹³ Heck seems to have something like *deductive rationality* in mind; i.e. the sort of rationality which makes transitions between representations normatively good (truth-preserving) in virtue of the structures of the representations involved. But the sort of reasoning/rationality Heck is talking about cannot *just be* deductive, or more broadly structural, since Heck thinks that cognitive maps are Fodorian icons which lack genuine structure. Nevertheless, if appropriate transitions between cognitive maps and other mental states are guaranteed because cognitive maps are *holistic* (see main text), this is tantamount to their being Rational, since the guarantee would follow from the *kind* of representations cognitive maps are, their representational nature, so to speak.

if cognitive maps were conceptual mental states *and* the subject possessed relevant update rules, she could presumably deduce all the appropriate revisions to represented relations. But it is hard to see what the relevant update rules would be. On the other hand, the hypothesis that cognitive maps are holistic (Sect. 3), insofar as they represent relations together with entities, provides a straightforward explanation of how relevant revisions to one's cognitive map could be Rational. Namely, by removing one entity from the cognitive map, and adding a different entity, appropriate revisions of the relations between entities in the map would be guaranteed. If cognitive maps were holistic, appropriate revisions to relations between entities would be guaranteed, because relations between entities would be represented via *representation of the entities* (Camp, 2018). In sum, Heck seems to conclude, the hypothesis that cognitive maps are holistic provides a better explanation of the role they play in reasoning than the hypothesis that they are structured propositions. Thus, other things being equal, cognitive maps are probably holistic. And if cognitive maps are holistic, then it follows that they are not propositionally-structured, i.e. non-conceptual (see Sect. 3).

Let us assume that this interpretation of the argument is roughly correct. We agree with Heck that if cognitive maps are holistic, (perhaps because they are icons, as Heck believes, or more plausibly because they are S-representations [see Sects. 5 and 5.1]), then they are non-propositional, and non-conceptual. But we do not think the argument from differential revision does enough to establish that cognitive maps are indeed holistic.

The main problem with the argument is that the *explanandum* is open to doubt. Heck seems to assume without argument that cognitive maps play a role in the subject's thinking which is big-r Rational. We provide reasons to doubt this assumption in Sects. 5.1 and 6. The transitions between cognitive maps and other mental states may very well be rational only in the more modest reliabilist sense that they tend to produce true beliefs. To the extent that the *explanandum* is open to doubt, an inference to its best explanation is unconvincing.

Despite our reservations about the broader argument, Heck makes a very important point; namely, if cognitive maps were holistic, revisions to them would be much cheaper, computationally, than if cognitive maps were conceptual. If cognitive maps were holistic, all of the revisions to relations between entities in the cognitive map would come for free if relevant revisions get made to the representations of the entities themselves (Camp, 2018). Whereas, in a system of propositional representations, these revisions to represented relations between entities would have to be deduced individually. We will appeal to this important difference in efficiency of revision in the *argument from robustness* (Sect. 5).

5 The argument from robustness

The argument from robustness motivates the idea that cognitive maps are holistic. If cognitive maps are holistic mental representations, they are S-representations. They represent some relations via structural relations rather than concepts. Thus, in grasping a content via a holistic cognitive map, the content is *not* entirely grasped via

concepts. Thus, the mental state in question is non-conceptual (see Sects. 3, 4.3, and Camp (2018)). The argument from robustness is inspired by some remarks of Evans’.

[A] fundamental identification of a place would identify it by simultaneous reference to its relations to each of the objects constituting the frame of reference. A place would be thought about in this way if it was identified on a map which represented, simultaneously, the spatial relations of the objects constituting the frame of reference. This identification has a holistic character: a place is not identified by reference to just one or two objects, and so the identification can be effective even if a few objects move or are destroyed. (Evans, 1982: 151).

The point we want to take from Evans is that if a cognitive map can represent a place effectively even when some of the objects constituting the frame of reference have been removed, a natural explanation would be that the cognitive map represents places holistically. Let us call the property of effectively identifying places even when relevant landmarks are removed, *robustness*. It turns out that cognitive maps are typically robust. O’Keefe and Nadel (1978) discuss some of the relevant experiments in, *The Hippocampus as a Cognitive Map*, a book that Evans cites. So, it is possible that Evans had in mind the very experiment we are about to discuss.

O’Keefe and Conway (1976) took single-unit recordings from place cells of rats moving around a cross-shaped maze. The researchers initially wanted to see what sensory cues anchored the rats’ place fields. The maze was surrounded by black curtains, on which were mounted sensory cues in different modalities. There was a lightbulb, a fan, a place-card, and a buzzer. The researchers kept the sensory cues in constant spatial relation to one another, but rotated the locations of the cues relative to the arms of the maze, and alternated which arm the rats started from. The researchers found that the place fields were fixed relative to the sensory cues, and not relative to the physical arm of the maze, or the body turns the rats made before or after entering the place fields. Having determined that place fields are identified relative to the sensory cues, the researchers sought to determine whether any one sensory cue was essential, as O’Keefe and Nadel explain:

For some of these place units it has proved possible to ascertain whether the place field was dependent on any specific stimulus within the environment by removing each of the four stimuli during different probe trials. In some units two stimuli at a time were removed... Almost all animals, and place units, behave normally after the removal of any single stimulus. This is the theoretically crucial finding, and demonstrates that the units are not responding to the position of any one of the stimuli. (O’Keefe & Nadel, 1978: 208).

The robustness of cognitive maps suggests that each place is identified relative to several others. No one landmark location is primary. The point we take from Heck’s discussion is that when spatial relationships are represented in holistic format, relevant revisions to representations of spatial relationships between entities do not have to be produced individually. Rather, if a landmark is removed from a holistic representation, the relations it formerly stood in to other landmarks are revised automati-

cally. We acknowledge that if a cognitive map were a set of propositionally-structured strings, it would be possible *in principle* to produce all of the appropriate revisions by deductive procedures (Cf. Sect. 4.3). But, and this is related to Heck's point, the revision processes required would be much more computationally costly than if the system were holistic (Cf. Camp, 2018). Moreover, if revisions to the representations of relations between entities have to be deduced individually, there is opportunity for error at every step of the inferences involved, while appropriate revisions of a holistic representation would be achieved automatically and without opportunity for error upon removal of the representation of the entity.

That place cells do not miss a beat, as it were, when anchoring landmarks are removed from a space, shows that revisions are made to cognitive maps with great accuracy and efficiency. These features of the operative revision process are better explained by the hypothesis that cognitive maps are holistic than by the hypothesis that they are sets of propositionally-structured strings. At this stage, we have no clear sense of how to flesh-out the latter hypothesis—no clear sense of what the relevant update principles could be, or of how they could produce appropriate revisions with the observed reliability and speed. We have a comparatively clearer sense of how to explain the reliability and efficiency of cognitive map revision on the hypothesis that cognitive maps are holistic. This should give some support to the latter hypothesis, even if this support is modest and provisional. If cognitive maps are indeed holistic, and standing in a propositional attitude requires the subject to grasp the relevant proposition wholly by way of discrete concepts, then cognitive maps are non-conceptual (Sect. 3). In the next section, we will develop more conclusive considerations that cognitive maps are indeed non-conceptual.

5.1 The argument from relational coding

Cognitive maps do not fit within the framework of conceptual mental states. We ascribe conceptual mental states to explain/predict what subjects will believe and do as the conclusions of chains of reasoning. Conceptual mental states play double-duty as both causal and rational links in these chains of thought. To play such a rational role, to serve as a premise in reasoning, the content of the state must be grasped by the subject entirely with distinct and repeatable concepts. The strategy of the argument from relational coding is to show that cognitive maps are not conceptually structured in this way.

Cognitive maps encode up-to-the-millisecond information about the changing spatial environment and the organism's changing positions within it, by mirroring spatial relationships with representational relationships. Cognitive maps' content is not grasped entirely via discrete concepts. This is why cognitive maps are not steps in a subject's deliberate deductive reasoning, and we are not able to predict or explain a subject's conclusions or behaviors within the model of deductive rationality by ascribing cognitive maps to her (cf. Sect. 4.3). The structure of cognitive maps is importantly different from that of conceptual mental states. This is best seen with examples.

In O'Keefe and Recce (1993) discovered a central representational strategy of cognitive maps, called *phase precession*. O'Keefe and Destrovsky (1971) had dis-

covered the existence of place cells, which emit bursts of action potentials when the organism moves through a specific location in the environment and are relatively silent otherwise (see Sect. 2). During movement, other cells in the surrounding population were known to fire synchronously, in a largely constant rhythm which EEG records as theta waves. Using single-unit recording, O'Keefe and Recce (1993) compared the timing of place cell bursts with the phase of the background theta rhythm in rats running back and forth along a linear track. They found evidence of a systematic relationship between place cell bursts and the background theta rhythm. As the rat enters a place cell's place field,

The first burst of firing consistently occurs at a particular phase of the reference theta [just after the peak of a wave], but each successive firing burst moves to a point earlier in the theta cycle. This precession of the phase correlates often continues throughout the field so that by the time the rat exits the field the phase shift may have moved through an entire cycle of the theta wave. (O'Keefe & Recce, 318).

In other words, a place cell's spike-bursts begin near the peak of a theta wave form, but as the organism moves through the center of the cell's place field, the bursts occur earlier and earlier relative to the phase of the theta wave in which they commenced. There is a roughly linear inverse correlation between the timing of a place cell's bursts and the phase of the theta wave in which they onset, as the organism moves through the cell's place field. The theta phase of place cells' bursts encodes much more detailed information about the organism's location than place cells' spike rates considered alone (O'Keefe & Recce, 1993), as evidenced by the fact that these phase correlates can be used to reconstruct an organism's location in the environment much more accurately than can rate information alone (Jensen & Lisman, 2000).

Phase precession is observed not only in place cells, but also in grid cells (Hafting et al., 2008) which have multiple, regularly spaced firing fields across the entirety of a space.¹⁴ To get a more precise sense of what theta phase-precession might represent in spatial navigation, Jeewajee et al. (2014) measured the phase correlates of place cell and grid cell activity in rats during navigation of an open arena. Summarizing their results, they explain that,

[T]he theta phase of firing of place cells and grid cells appears to represent the location of the animal within the currently occupied firing field. The way in which this location is represented during foraging in open environments is best captured by the distance of the field peak [i.e. the point at which the cell fires maximally] ahead of or behind the animal along its current direction of motion... [T]his representation is directly useful for navigation. (Jeewajee et al., 2014: 8).

¹⁴Phase precession is also described in other, non-spatially tuned neurons, indicating a potential role for it in encoding non-spatial sequences (Bush & Burgess, 2020; Qasim et al., 2021).

In other words, place cells and grid cells represent the spatial *relation* (distance) between the organism and the center of the relevant cells' firing fields continuing along the current direction of travel by way of the temporal *relation* between these cells' bursts and the phase of the background theta rhythm in which the bursts commenced.

Phase precession is not limited to rodents but has recently been reported in humans (Qasim et al., 2021), bats (Eliav et al., 2018), and marmosets (Courellis et al., 2019). Qasim et al., (2021) took single-unit recordings from four human epileptic patients navigating a virtual reality environment. Summarizing their results, they write,

Our findings show that neurons in the human brain spike in rhythm with local network oscillations to represent [relative] spatial position... We provide evidence for rodent-like spatial phase precession in human hippocampus and entorhinal cortex, in which spatially tuned neurons spike at earlier phases of theta (2-10 Hz) LFP oscillations as subjects moved through the putative place field center. (Qasim et al., 2021: 3250).

There is a broad consensus that cognitive maps represent extremely subtle changes in relative spatial position via subtle changes in the relationships between burst timing and theta phase.

Head direction (HD) cells exhibit similar relational coding of fine-grained changes in relative orientation via changes in representational relations. HD cells are cells that tend to fire only when the animal's head is pointed in a particular range of allocentric directions (usually composed of about 90° of angular rotation) (Sect. 2). But HD cells' representational strategy is more subtle than this. When the animal's head is facing outside an HD's cell's preferred range, the cell's spike rate is close to zero. But once the head faces a direction within the cell's preferred range, the cell's spike rate increases roughly linearly as the head's direction approaches the center of the range (Taube, 2007). Thus, an HD cell's tuning curve on an x-y graph plotting head direction against spike rate will be triangular, with peak firing in the center of the range (Ibid.). Subtle changes in an HD cell's spike rate relative to max seem to represent subtle changes in head direction relative to the center of the cell's firing range.

Subjects do not grasp the contents of cognitive maps entirely via concepts in the way subjects must grasp the contents of conceptual mental states such as belief. The reason for this is not merely because cognitive map representations draw more specific, subtle, or fine-grained distinctions between contents than do concepts. Rather, as Shea (2018) has argued in an important discussion, cognitive maps seem to be S-representations. Shea argues that in route planning, for example during virtual trial and error (Sect. 4.2), the brain uses co-activation relations between place cells to represent spatial relations between firing fields that make up possible routes to a goal.

The animal picks an efficient route to a goal by picking the sequence that takes the shortest time to unfold during preplay. It then follows that sequence. That algorithm has been stabilized by learning in part because of a structural correspondence between co-activation on the place cells and spatial proximity on locations, relied on to calculate the route. The structural correspondence... is

exploited... [Thus] it is content-constituting: co-activation of place cells represents spatial proximity of locations. (Shea, 2018: 116).

The precise algorithms used by the brain during route planning are still debated, but it is immensely plausible that they exploit relations of place-cell activations as proxies for, i.e. to represent, spatial relations between places. Thus, route planning plausibly depends on the fact that spatial relations in the world are mirrored by relations between activations of different place cells.

Moreover, our preceding discussion shows that even outside the context of route planning, cognitive maps are S-representations of relative self-location and relative self-orientation. In the activity of place cells and grid cells, cognitive maps appear to represent where one is relative to place and grid field centers by the timing of cell bursts relative to a consistent backbeat LFP. HD cells appear to represent one's facing direction relative to the center of their firing fields by the relationship between cells' current spike rates and their max spike rates. Thus, the reason subjects do not grasp cognitive map contents entirely via concepts in the way subjects must grasp conceptual contents is that cognitive maps represent worldly relations differently than conceptual mental states do. Cognitive maps represent (some) worldly relations not by means of distinct representation elements, such as concepts, but rather by means of structural relations *between* representation elements (see Sect. 3).

6 A taxonomy of functional information

If there are different kinds of mental states, non-conceptual mental states as well as conceptual mental states, then questions about their respects of similarity and difference become pressing. This section does not presume to answer these questions, some of which may not *be* answerable until empirical research progresses. Rather, we aim to sketch a set of hypotheses suggested by the taxonomy of functional information (Fresco et al., 2020), primarily to clarify the problem-space and invite future discussion.

Lewis (1969) proposed a game-theoretic approach to social conventions, according to which they are solutions to coordination problems, (i.e., Nash equilibria). This approach has since been fruitfully applied to animal signaling and communication (see, e.g., Skyrms, 2010). Another tradition, beginning with Dretske (1981), appeals to Shannon's mathematical theory of communication to understand how states of organisms can register and extract information from sensory impingement. The taxonomy of functional information proposed by Fresco et al. (2020) draws from both traditions to produce a taxonomy of functional information.

The notion of functional information is that of a difference in the (internal or external) milieu of a system which has made a systematic difference to the system's goal-directed behavior (Fresco et al., 553). The notion of information is thus receiver-dependent. The taxonomy of different types of functional information has a hierarchical structure. A *datum* is only functionally informational in a potential sense. It is a difference in the environment of or within a system that could potentially be exploited by the system to make a reliable difference to its goal-directed behavior.

A *sign* is a datum that a receiver has evolved through selection (phylogenetic, ontogenetic, or cultural) to exploit, or to be disposed to exploit, in a way that makes a reliable difference to the system's goal-directed behavior. A *signal* is a sign, which may have a learned component, that a sender has evolved through some process of selection to send to particular receivers. And lastly, a *symbol* is a signal that is part of a rule-governed and self-referential system. But symbols are rule-governed in a specific way. The rules governing symbolic systems makes it generally the case that, "truth conditions can be assigned to symbolic messages" (Fresco et al., 2020: 562). Non-symbolic messages are typically evaluated by the standard of functional accuracy, rather than truth.

Within this taxonomy, states of the cognitive map might be best understood as signals. Experimentally confirmed commonalities in the structure and function of cognitive maps, for example, phase precession in place and grid cells, across phylogenetically distant mammalian species, [rats, bats humans, etc. (Qasim et al., 2021)], strongly suggest that this neural machinery has been evolved to encode and transmit location information. The hippocampus receives input from the entorhinal cortex housing grid cells, and the subiculum housing HD cells. Characterization of the receiver of the signals that the cognitive map sends will need to await more mature computational models. Cognitive maps may also, in a sense, be self-reading (Godfrey-Smith, 2013). In the present taxonomy of functional information, conceptual mental states are seen as symbols. The idea that conceptual mental states are governed by syntactic rules has been argued plausibly by Fodor in many places (Fodor, 1987; Fodor & Pylyshyn, 1987). The contents of conceptual mental states, propositions, are truth-evaluable. By contrast, cognitive maps are not governed by rules in a way that makes them straightforwardly truth-evaluable.¹⁵ Cognitive maps represent (in part) by exploiting natural correspondences between worldly relations and representational relations. Accuracy in a cognitive map is determined by the degree of structural analogy between represented worldly relations and representational relations. The structural analogies by which cognitive maps represent worldly relations determine *scales* of accuracy by which to evaluate cognitive maps. The representational structures of symbolic systems do not provide scalar standards of accuracy in this way. The literature on rule-following is vast, and much more would have to be said than we can say here to secure the claim that cognitive maps are not rule-governed in the way that symbolic systems are. Our hypothesis is that symbolic systems are governed by rules that are algorithmic in something like the mathematical sense. Their steps are definite and effective means for producing specific outputs, given specific inputs (Papayannopoulos, 2023). Non-conceptual mental states such as cognitive maps represent (in part) by S-representation, and thus by analogy. If so, non-conceptual mental states such as cognitive maps would be better evaluated for accuracy, (according to the closeness of the representational analogy they produce), than simply for truth/falsity. In that case, the principles governing non-conceptual mental states such as cognitive maps may be a matter of more-or-less rather than all-

¹⁵Cognitive maps are also not self-referential in the way that language is, for example. Words sometimes refer to other words within the lexical system, rather than to any worldly object or event. Place cells, grid cells, and HD cells never refer to other such cells—their reference is always to worldly entities or relations.

or-nothing. That is, it might consist in mapping from a given input to any of a range of possible outputs, some, perhaps, more correct than others.

Whether or not the taxonomy of functional information correctly describes the relationships between cognitive maps and propositional attitudes, it may provide a useful framework for future discussion. Any theory of the relationships in question must, as the taxonomy of information purports to, show what unifies conceptual and non-conceptual mental states. The conceptual framework of functional information suggests that mental states are linked in being selected differences that make a reliable difference to the system's goals. It presents conceptual mental states as higher than non-conceptual mental states, (e.g., cognitive maps), in a constitutional hierarchy of informational states.

7 Conclusion

Eavesdropping on the mind with recording and imaging technologies has led to remarkable progress in our understanding of mammalian navigation. Cognitive maps deserve more philosophical attention. If cognitive maps are indeed non-conceptual mental states, as we have argued, this would have far-reaching implications. First, it would open the possibility of commonalities between the navigational mental states of creatures with different conceptual repertoires. This is desirable given the extensive overlap in neurology relevant to navigation across mammals. Second, if cognitive maps are non-conceptual, this would clear the way for an account of the formation of certain concepts in part from the raw materials of cognitive maps. This is a project that Evans would approve, and indeed one that he began in *The Varieties of Reference*. Last, the existence of non-conceptual content in cognitive maps would raise the question of how cognitive maps can epistemically license or otherwise the formation of certain beliefs. It would make for simpler philosophy if the whole mental domain were populated exclusively by conceptual mental states. But it makes more sense that there would be states of mind, more primitive than belief, that allow an animal to get around in the world, and whose contents need not be fully grasped through the use of concepts.

Acknowledgements Many people helped make this article better. For invaluable comments on earlier drafts, we wish to thank to Ori Beck, Ali Boyle, Matteo Colombo, Zoe Drayson, Jonathan Najenson, Will Sharp, Rotem Yahel, and three anonymous referees for this journal. We are also grateful to helpful audiences at the 26th Meeting of the Israel Philosophical Association, and at the Grenoble-Bochum Memory Colloquium (May, 2023).

Funding A.S. gratefully acknowledges the support of the Israel Academy of Sciences and Humanities & Council for Higher Education's "Excellence Fellowship Program for International Postdoctoral Researchers", as well as support from United Kingdom Research and Innovation through the Future Leaders Fellowship Grant to Ali Boyle, "Episodic Memory: Uniquely Human?" Grant Reference: MR/W00741X/1. This research was also supported by an Israel Science Foundation Grant to N.F. Grant Reference: 386/20.

Data availability There are no data associated with this article.

Declarations

Conflict of interest The authors declare no conflicts of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bechtel, W. (2016). Investigating neural representations: The tale of place cells. *Synthese*, 193(5), 1287–1321. <https://doi.org/10.1007/s11229-014-0480-8>
- Bermúdez, J. (2003). *Thinking without words*. Oxford UP.
- Bermúdez, J. L. (2007). What is at stake In the debate on nonconceptual content? *Philosophical Perspectives*, 21, 55–72.
- Brewer, B. (2005). Perceptual Experience Has Conceptual Content. In S. Matthias, E. Sosa (Eds.), *Contemporary debates in epistemology*. Blackwell Pub, 217–230.
- Burge, T. (2018). Iconic representation: Maps, pictures, and perception. In S. Wuppuluri & F. A. Doria (Eds.), *The map and the territory: Exploring the foundations of science, thought and reality* (pp. 79–100). Springer. https://doi.org/10.1007/978-3-319-72478-2_5
- Bush, D., & Burgess, N. (2020). Advantages and detection of phase coding in the absence of rhythmicity. *Hippocampus*, 30(7), 745–762. <https://doi.org/10.1002/hipo.23199>
- Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, 21(1), 145–182. <https://doi.org/10.1111/j.1520-8583.2007.00124.x>
- Camp, E. (2018). Why maps are not propositional. In A. Grzankowski & M. Montague (Eds.), *Non-propositional intentionality*. (Vol. 1). Oxford University Press. <https://doi.org/10.1093/oso/9780198732570.001.0001>
- Clarke, S. (2022). Mapping the visual icon. *The Philosophical Quarterly*, 72(3), 552–577. <https://doi.org/10.1093/pq/pqab054>
- Clarke, S. (2023). Compositionality and constituent structure in the analogue mind. *Philosophical Perspectives*, 37(1), 90–118. <https://doi.org/10.1111/phpe.12182>
- Courellis, H. S., Nummela, S. U., Metke, M., Diehl, G. W., Bussell, R., Cauwenberghs, G., & Miller, C. T. (2019). Spatial encoding in primate hippocampus during free navigation. *PLOS Biology*, 17(12), e3000546. <https://doi.org/10.1371/journal.pbio.3000546>
- Crane, T. (1992). The nonconceptual content of experience. In T. Crane (Ed.), *The contents of experience*. Cambridge University Press.
- Crane, T. (2012). The given. In J. Schear (Ed.), *Mind, reason, and being-in-the-world: The McDowell-Dreyfus debate* (pp. 229–249). Routledge.
- Cummins, R. (1996). Representations, targets, and attitudes. *The MIT Press*. <https://doi.org/10.7551/mitpress/5887.001.0001>
- Cussins, A. (1990). The connectionist construction of concepts. In M. Boden (Ed.), *The philosophy of artificial intelligence* (pp. 368–440). Oxford University Press.
- Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469, 397–401.
- Dreske, F. (1969). *Seeing and knowing*. Routledge and Kegan Paul.
- Dreske, F. (1981). *Knowledge and the flow of information*. MIT Press, Bradford Books.
- Dreske, F. (1995). *Naturalizing the mind*. MIT Press.
- Ekstrom, A., Kahana, M., Caplan, M., Fields, T., Isham, E., Newman, E., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, 425(6954), 184–188.

- Ekstrom, A., Spiers, H., Bohbot, V., & Rosenbaum, R. (2018). *Human spatial navigation*. Princeton UP.
- Eliav, T. E., Geva-Sagiv, M., Yartsev, M. M., Finkelstein, A., Rubin, A., Las, L., & Ulanovsky, N. (2018). Nonoscillatory phase coding and synchronization in the bat hippocampal formation. *Cell*, 175(4), 1119–1130.e15. <https://doi.org/10.1016/j.cell.2018.09.017>
- Evans, G. (1982). *The varieties of reference*. Oxford UP.
- Fodor, J. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. MIT Press.
- Fodor, J. (2007). The revenge of the given. In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 105–116). Wiley-Blackwell.
- Fodor, J., & Pylyshyn, Z. (1987). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3–71.
- Foster, D. (2017). Replay comes of age. *Annual Review of Neuroscience*, 40, 581–602.
- Fresco, N., Gisburg, H., & Jablonka, E. (2020). Functional information: A graded taxonomy. *Review of Philosophy and Psychology*, 11, 547–567.
- Godfrey-Smith, P. (2013). Signals, icons, and beliefs. In: Millikan and her critics. John Wiley and Sons Inc., p. 41–62.
- Grice, H. P. (1961). The causal theory of perception. *Proceedings of the Aristotelian Society*, 35, 121–152.
- Hafting, T., Fyhn, M., Bonnevie, T., Moser, M.-B., & Moser, E. I. (2008). Hippocampus-independent phase precession in entorhinal grid cells. *Nature*, 453(7199), 1248–1252. <https://doi.org/10.1038/nature06957>
- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801–806.
- Hassabis, D., Chu, C., Rees, G., Weiskopf, N., Molyneux, P., & Maguire, E. (2009). Decoding neuronal ensembles in the human hippocampus. *Current Biology*, 19(7), 546–554.
- Heck, R. K. (2000). Nonconceptual content and the ‘Space of Reasons.’ *The Philosophical Review*, 109(4), 483–452.
- Heck, R. K. (2007). Are there different kinds of content? In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 117–138). Wiley-Blackwell.
- Jacobs, J., Weidemann, C., Miller, J., Solway, A., Burke, J., Wei, X., Suthana, N., Sperling, M., Sharan, A., Fried, I., & Kahana, M. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, 16(9), 1188–1190.
- Jeewajee, A., Barry, C., Douchamps, V., Manson, D., Lever, C., & Burgess, N. (2014). Theta phase precession of grid and place cell firing in open environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635), 20120532. <https://doi.org/10.1098/rstb.2012.0532>
- Jensen, O., & Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: Contribution of theta phase coding. *Journal of Neurophysiology*, 83(5), 2602–2609. <https://doi.org/10.1152/jn.2000.83.5.2602>
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Martin, M. (1992). Perception, concepts, and memory. *The Philosophical Review*, 101, 745–763.
- McDowell, J. (1994). *Mind and world*. Harvard University Press.
- Muller, R., & Kubie, J. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7(7), 1951–1968.
- Nadel, L. (2013). Cognitive maps. In L. Nadel & D. Waller (Eds.), *Handbook of spatial cognition* (pp. 155–172). American Psychological Association.
- Neander, K. (2017). *A mark of the mental*. MIT Press.
- O’Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171–175.
- O’Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Clarendon Press.
- O’Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3), 317–330.
- O’Keefe, J., & Conway, D. (1976). Hippocampal place units in the freely moving rat—why they fire where they fire. *Experimental Brain Research*, 31, 573–590.
- Papayannopoulos, P. (2023). On algorithms, effective procedures, and their definitions. *Philosophia Mathematica*, 31(3), 291–329. <https://doi.org/10.1093/philmat/nkad011>
- Peacocke, C. (1992). *A study of concepts*. MIT Press.
- Qasim, S. E., Fried, I., & Jacobs, J. (2021). Phase precession in the human hippocampus and entorhinal cortex. *Cell*, 184(12), 3242–3255.e10. <https://doi.org/10.1016/j.cell.2021.04.017>
- Ramsey, W. M. (2007). *Representation reconsidered* (1st ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511597954>

- Ranck Jr., J. B. (1984). Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. *Society for Neuroscience Abstracts*.
- Rescorla, M. (2009). Cognitive maps and the language of thought. *British Journal of the Philosophy of Science*, 60, 377–407.
- Rescorla, M. (2017). Maps in the head? In K. Andrews & J. Beck (Eds.), *The Routledge handbook of animal minds* (pp. 34–45). Routledge.
- Roskies, A. L. (2008). A new argument for nonconceptual content. *Philosophy and Phenomenological Research*, 76(3), 633–659.
- Rowland, D. C., & Moser, M. B. (2014). From cortical modules to memories. *Current Opinion in Neurobiology*, 24(1), 22–27.
- Shea, N. (2018). *Representations in cognitive science*. Oxford University Press.
- Shine, J., Valdes-Herrera, J., Hegarty, M., & Wolbers, T. (2016). The human retrosplenial cortex and thalamus code head direction in a global reference frame. *Journal of Neuroscience*, 36, 6371–6381.
- Skyrms, B. (2010). *Signals: Evolution, learning, & information*. Oxford University Press.
- Stensola, H., Stensola, T., Solstad, T., Froland, K., Moser, M.-B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492, 72–78.
- Sugar, J., & Moser, M.-B. (2019). Episodic memory: Neuronal codes for what, where, and when. *Hippocampus*, 29, 1190–1205.
- Taube, J. (2007). The head direction signal: Origins and sensory-motor integration. *Annual Review of Neuroscience*, 30, 181–207.
- Tolman, E. (1948). Cognitive maps in rats and men. *The Psychological Review*, 55(4), 189–208.
- Tolman, E. C., Ritchie, B. F., & Kalish, D. (1946). Studies in spatial learning. I. Orientation and the short-cut. *Journal of Experimental Psychology*, 36(1), 13–24.
- Yartsev, M., Witter, M., & Ulanovsky, N. (2011). Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature*, 479, 103–107.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.