ARTICLE IN PRESS

Studies in History and Philosophy of Science xxx (xxxx) xxx

Contents lists available at ScienceDirect

Studies in History and Philosophy of Science

journal homepage: www.elsevier.com/locate/shpsa

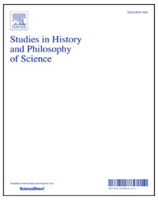

Book Review

Representation in Cognitive Function by Nicholas Shea: Organization and Structure in the Service of Systematicity and Productivity☆

This is the book I've been seeking, not knowing whether it existed (such are the wonders of intentionality). Like Shea, I favor a broadly pragmatic approach to mental representation: methodologically, we posit representations to explain complex patterns of behavior; substantively, representing is a capacity to track information in the service of action. But like him, I want my pragmatism to be realist: mental representations are physically implemented and causally efficacious, and can be false even if useful. Teleosemantics has long been a leading candidate to fill this niche, but it has been beset by accusations of circularity and wishful thinking. Shea places a broadly teleosemantic approach on a firm, if unapologetically non-reductive, footing. I especially appreciate his insistence that realism requires explaining in detail *how* a system tracks information, which leads him to grapple with multiple representational formats and forces. Here, as in his analysis of functions, he is admirably ecumenical and specific, covering cases of varying complexity in a way that meshes elegantly with actual practice in cognitive neuroscience.

I want to focus on how Shea's realism interacts with his format ecumenicalism, and more specifically on how different representational systems distribute the burden of exploiting structural correlations to track information. Representation requires systematicity; but different formats underwrite systematicity in different ways and degrees. This matters functionally, by affecting their representational capacities and vulnerabilities. And it matters theoretically, by affecting where we posit and how we test for representational mechanisms. The upshot is that all of us have more fine-grained work to do in establishing what and how a system represents.

1. Organizational and structural systematicity

There is a sense in which any differential response to stimuli is systematic: an organism, like magnetotactic bacteria, discriminates and reacts to instances of a certain kind in a certain way, which it does not to others. But representation in a substantive sense requires a more robustly systematic contingency between discrimination and action. Vervet alarm calls provide an elegant illustration: vervets' behaviors are best explained by abstracting away from indefinitely many variations among heard calls and responses and positing that the remaining pattern has stabilized because there are three types of calls, each of which is sufficiently systematically correlated with the presence of a certain type of predator (*eagle, snake, leopard*) that producing that type of behavior (*look up, look down, climb)* on hearing that call (C_e, C_s, C_l) is sufficiently helpful for survival. Vervets' overall behavior displays a common structure, from a predator-type P_x to an action-type A_x, mediated by a representation-type R_x. Each token representation (e.g. r_e) is an element within that system: it would not be a representation, or represent what it does, were it not an instance of the type (R_e), where that type is constituted as the type it is by being embedded within a causal structure which also subsumes the other two types.

But while there is a system, and systematicity, here, it is still quite thin. Each representational type is an independent unit: the interpretation function from each of R_x to P_x is *pointwise* (119), with the three types related only by their parallel structure, which could just as well have one (or thirteen) instances, each implemented by a fully distinct mechanism. Other systems' representational types are more systematically interconnected. Thus, analogue magnitude representations (AMRs) exploit a unidimensional correlation between the magnitudes of an external property *P* and an internal register *R* (98). Because this function is analogue, it generates indefinitely (potentially continuum-) many representational types ($R_1, R_2, R_3, \ldots$). Because the function generating those types is structural rather than pointwise, the types are systematically ordered with respect to each other. And because the structure among the resulting types mirrors the structure of the contents they represent, AMRs can be used not just to track indefinitely many distinct contents in an efficient and reliable way, but also to compare those contents to one another. Further, many AMR systems track not just one but multiple properties – for instance, duration, distance, and numerosity; in that case, *R* generates multiple classes of degree-types ($R_{ai}, R_{aj}, \ldots; R_{bi}, R_{bj}, \ldots; R_{ci}, R_{cj}, \ldots$), which can potentially be used to compare or aggregate multiple quantities of multiple properties.

Following Godfrey-Smith (2017), Shea describes systems with systematically related representational types as “organized” (128). He distinguishes these from genuinely *structural* systems, which generate representational types that have internal parts that are related in a systematic, representationally significant way. A representation's having parts governed by interpretation functions that systematically exploit structural correspondences does not suffice to make the representation itself structural. Thus, bee dances have two parts (waggle duration and orientation), each exploiting an (analogue) dimension of structural covariance with a represented magnitude (distance and direction). But those two dimensions are representationally independent (163): the functions ($\mathit{duration}_m \rightarrow \mathit{distance}_m$) and ($\mathit{orientation}_n \rightarrow \mathit{direction}_n$) operate separately, with instantiation in a common dance being a mere implementational convenience.

Shea offers cognitive maps in rats as a paradigm case of what more is needed for a genuinely structural representation. Neurons in rats' hippocampus function as ‘place cells’ representing locations. A set of place cells does not yet constitute a structural representation – even if they instantiate a topological configuration that mirrors their denoted locations, which rats' place cells don't. Rather, Shea claims that the rat hippocampus implements a map because relations of cellular co-activation systematically co-vary with relations of distance and direction among

---

☆ Thanks to Nick Shea and Eliot Michaelson for discussion.

https://doi.org/10.1016/j.shpsa.2021.05.015

Please cite this article as: Camp, E., *Representation in Cognitive Function* by Nicholas Shea: Organization and Structure in the Service of Systematicity and Productivity, Studies in History and Philosophy of Science, https://doi.org/10.1016/j.shpsa.2021.05.015

denoted locations, and rats use this covariance for efficient navigation (116).

But when does systematic exploitation of a structural – more specifically: spatial, geometrical, or topological – correspondence make a map? A spatially implemented map, like a paper seating chart or road atlas, is constituted as a single map because its parts (e.g. names, blue dots, black lines) actually stand in spatial relations, which are themselves representationally significant (Camp, 2007). It belongs to an organized system because it is generated via principles that also govern other actual and possible maps, representing other domains and other configurations of properties. And it is a structural representation because its overall content depends not just on which marks it contains, but on how they are related.

More specifically, maps – unlike sentences – exhibit *holistic* representational structure. Because all of the marks in a spatially implemented map actually stand in all of the spatial relations they represent, the map directly, explicitly, and simultaneously represents all of those spatial relations among those denoted locations, and in turn among all of the properties that are represented as being at those locations. Its parts comprise a functionally integrated whole in a way an informationally equivalent list of conjoined sentences does not. Interpretively, information about spatial relations among represented properties comes along as a "free ride," in virtue of representing them at their locations (Shimojima, 1996). And implementationally, any alteration to the placement of a property-representing mark automatically alters all of the spatial relations in which that property is represented as standing, so that merely partial changes to those represented relations are not just inconsistent, but impossible (Camp, 2018).

The shift from a concrete spatial representation to an abstract functional one buys the rats implementational flexibility. But it also weakens theorists' grip on the idea of format, and thereby on the difference between organizational and structural systematicity. And this in turn makes pressing the question *why* we should treat co-activation relations among rats' place cells as "proxies" for represented spatial relations in a way that would justify classifying them as structural representations, and specifically as maps. Shea appeals here to the fact that rats' offline activation of sequences of place cells systematically correspond to routes through the represented domain (116). I agree that this demonstrates the representational significance of a structural correspondence between the co-activation of place cells and the geometry of their denoted locations. But by itself, it does not yet establish that this structural correspondence is being exploited within "a single representation with representational parts," and thus as a structural representation, as opposed to "a series of different representations" (128).

Contrast three implementations of rats' navigation system, each built on a set of location-denoting individual place cells (L_1, L_2, L_3, …). *System 1* constructs potential routes from source to goal separately (in succession or parallel) by scanning a list of recorded pairwise transitions between cells ($L_1 \rightarrow L_2$, $L_2 \rightarrow L_4$, $L_3 \rightarrow L_4$, …), compiling a list of sequences of pairs that share at least one cell, and then scanning that list of sequences for any that contain the cells denoting the source and the goal. It represents route lengths with numerical labels corresponding to the number of pairs in a sequence, and selects the sequence with the numeral denoting the smallest number. *System 2* constructs potential routes separately (in succession or parallel) by "re-playing" sequences of cellular activation (115) of previously travelled routes ($L_1 \rightarrow L_2 \rightarrow L_4 \rightarrow L_6$, …), beginning with the source-denoting cell, segmenting sequences at cells where an attempted movement was blocked, compiling segments with common cellular endpoints, and halting when a sequence including both source-denoting and goal-denoting cells has been compiled. It tracks route lengths with an analogue 'duration' register, and selects the sequence with the smallest magnitude. *System 3* constructs potential routes by progressive activation from the source-denoting cell to every cell that has previously been co-activated with it, and then to every cell that has been co-activated with them. It halts as soon as the goal-denoting cell is activated, then progressively de-activates every cell not connected to at least two active cells, and initiates action to the source-neighboring cell that remains activated after pruning is complete (selecting at random if multiple routes remain active).

There is a sense in which all three systems are computationally equivalent; a sense in which all are functional rather than concrete 'maps'; and a sense in which all exploit structural correspondences between external circumstances and inner states and operations. But they exploit those structural correspondences via quite different mechanisms, in ways that make them vulnerable to different types of malfunction and impose different implementational constraints. Among other things, Systems 1 and 2 require "little chains of reasoning" (194) to compile and extract information, and memory to store it. This means they can fail to perform a comprehensive search, or introduce transcription errors, in a way that the more robustly map-like System 3 can't. But they also produce stable representations of distinct routes where System 3 doesn't, which might then be available for further informational exploitation.

These more fine-grained differences in how a system encodes, compiles, and extracts information make a practical difference. They also make an evidential and theoretical difference. In particular, if rats' generation and comparison of potential routes does occur via parallel diffusion over the entire array of place cells (115, fn. 4), then this supports positing a more holistic representation, à la System 3. The theoretical relevance of such still-tentative empirical details suggests that Shea's parade case for a system of structural representations – and with it the contrast between organized and structural representational systems – is not yet conclusive. But it also thereby illustrates his more basic point: that earning representational realism requires a full causal explanation of how a system deploys stable information-exploiting mechanisms to perform tasks in an unstable environment.

2. Systematicity, productivity, and inference

How systematic a representational system is, and how it exploits structure in the service of systematicity, doesn't just affect how it handles a given body of information, but also the range of information it is capable of representing. The vervet call system's pointwise, parallel structure entails that it can only represent three predators; it has the potential to generate new representational types only in the counterfactual sense that its overarching causal structure might be reconstructed to do so. By contrast, the analogue organizational structure of ARMs entails that, as constituted, they produce indefinitely many representational types, despite being internally unstructured. Languages also produce indefinitely many types, but by means of compositional structure: by exploiting a highly abstract correspondence between the asymmetrical, pairwise structures of predication and metaphysical instantiation, implemented via recursive application over a pointwise base (Camp, 2015).

In what ways does the systematicity of rats' navigational system support productivity? Because the worldly locations that their place cells denote cannot themselves be permuted, it does not make sense to ask whether the system can recombine a given set of cells to form multiple maps, as languages do. However, independently of whether the combinatorial structure supporting rats' navigational system turns out to be more like System 1, 2, or 3, presumably it does have the capacity to permute which property-types (e.g. *food*, *obstacle*, *nest*) it represents as being at those locations, thereby generating a finite but factorially-large number of distinct representational types. In this sense, the structural systematicity of their system supports productivity where vervets' alarm calls do not.

The rats' system also appears to differ from the vervets' in having a stable mechanism for recruiting new place cells. If so, this constitutes a form of *system-level productivity* – much as languages achieve productivity not just structurally, by recombining terms from a fixed lexicon, but systemically, by forming new terms. This mechanism might itself be more or less productive and organized. Thus, it might select place cells randomly and assign their reference pointwise, à la names. It might select cells in a systematic, organized way but assign reference pointwise, à la indexical terms (e.g. ‘here$_i$’, ‘here$_j$’ …, ‘there$_i$’, ‘there$_j$’ …). Or it might exploit structure in order to both select and interpret cells, à la cartographic maps (125).

What location- and property-representing mechanisms rats' maps employ will again make a functional difference – *inter alia*, to how many locations and properties they are capable of representing, to whether they can represent entirely disconnected domains, and to what information about relations among properties can be extracted from relations among place cells. For instance, introducing a new, name-like place cell into System 1 will presumably require establishing co-activation relations piecemeal, while a new place cell introduced into System 3 via a structure-exploiting mechanism will automatically inherit a full suite of co-activation relations in virtue of that mechanism's structure.

Thus far, we have retained the assumption that rats employ a structural interpretive function overlaying a stable, discrete base. A cognitive map system might take the systematic exploitation of structure a step further, swapping that discrete base for AMRs representing distance and direction from an egocentric origin. Indeed, it might go fully analogue, employing AMRs to represent not just locations but also gradable properties, like danger and hunger-satisfiability, at those locations, in a manner akin to Peacocke’s (1992) scenario contents. Depending on whether and how the system exploited these AMRs, such a system might count as representing higher-order relational contents – say, the relative distances of routes, or relative trade-offs in distance, risk, and reward – directly and at a systemic level, without the explicit local encoding and inferential extraction that would be employed by a more thoroughly pointwise system. (Shea allows for system-level representational types (219).) Such a system would be highly holistic. Its implementation could be massively distributed and functionally abstract. But it would still be strongly systematic and productive. And it would achieve systematic productivity by being compositional and computational in fairly robust, familiar senses of those terms.

None of these conclusions undermine Shea's core claims about functions or structure. Rather, they demonstrate the power of his analysis to clarify the range of ways that representational systems can exploit stable covariations among inner and outer states to track, compile, and deploy information. Philosophers have long assumed a fundamental dichotomy between perceptual, stimulus-dependent, unstructured, iconic representations and conceptual, stimulus-independent, systematic, propositional ones. But we have seen that systematic, stimulus-independent representation can be achieved by exploiting structure in a wide range of ways, varying along multiple dimensions: either concretely or abstractly, locally or systemically, in analogue or digital form, by piecemeal or holistic combination. Shea scrupulously abjures extending his analysis to conceptual, personal-level representations like belief and desire. But if systematic, stimulus-independent representation suffices for conceptual thought, as I believe it does (Camp, 2009), then we should expect even conceptual thought to take a multitude of forms as well.

References

Camp, E. (2007). Thinking with maps. *Philosophical Perspectives, 21*(1), 145–182.
Camp, E. (2009). Putting thoughts to work: Concepts, systematicity, and stimulus-independence. *Philosophy and Phenomenological Research, 78*(2), 275–311.
Camp, E. (2015). Logical concepts and associative characterizations. In E. Margolis, & S. Laurence (Eds.), *The conceptual mind: New Directions in the Study of concepts* (p. 591621). Cambridge, MA: MIT Press).
Camp, E. (2018). Why maps are not propositional. In A. Grzankowski, & M. Montague (Eds.), *Non-Propositional Intentionality* (pp. 19–45). Oxford: Oxford University Press.
Peacocke, C. (1992). Scenarios, concepts, and perception. In T. Crane (Ed.), *The Contents of experience* (pp. 105–135). Cambridge: Cambridge University Press.
Shimojima, A. (1996). Operational constraints in diagrammatic reasoning. In J. Barwise, & G. Allwein (Eds.), *Logical Reasoning with diagrams* (pp. 27–48). Oxford: Oxford University Press.

Elisabeth Camp
Rutgers University, Department of Philosophy, New Brunswick, USA
E-mail address: elisabeth.camp@rutgers.edu.