Understanding Structural Representations

Marc Artiga

The notion of structural representation has recently been suggested to play two key dialectical roles: on the one hand, against eliminativist challenges it is supposed to capture some cognitive processes in which a representionalist explanation can earn its keep. On the other, against liberal views it vindicates a concept of representation that avoids the problem of trivialization. A first goal of this article is to show that the notion of structural representation can be understood in different ways. Second, I will argue that on any interpretation it is unclear this concept can play the theoretical role attributed to it in the recent literature.

---

1. Introduction

Since Ramsey ([1931]) famously claimed that beliefs are maps by which we steer, many have explored the analogies between cognitive states and map-like representations. A particularly interesting version of this idea has been developed around the notion of structural representation, which has been brought to bear on one of the central questions of cognitive science: the representational status of cognitive processes.

Roughly, a structural representation can be defined as follows:

Structural Representation: A mechanism in which a structural correspondence between a set of vehicles and a target domain is used to accomplish certain task.

As a first approximation, structural correspondence refers to some kind of structural similarity between a set of entities and a target domain. Underground maps or nautical charts, for example, are paradigmatic examples of structural representation.

It has recently been proposed that structural representations play at least two key theoretical roles. First, eliminativist approaches have put some pressure on the concept of representation, which has forced philosophers who posit them to revise their position and seek a better justification for their postulation (for example, Beer [1995]; van Gelder [1995]; Clark [1996]; Chemero [2009]; Hutto and Myin [2013], [2017]). In the context of this debate, the concept of structural representation is meant to identify a set of processes that cannot easily be explained without an appeal to representations. Accordingly, the claim that some cognitive mechanism qualifies as a structural representation is supposed to vindicate representationalist explanations.

Second, structural representations have been used to oppose liberal representationalist views. The debate between liberal and non-liberal accounts concerns the scope of attributions: according to the former, the class of representations extends well beyond the classical examples of beliefs, desires, perceptual states, linguistic knowledge and the like and include, for instance, firing patterns in photoreceptors or ganglion cells. More generally, according to liberal views, detectors (that is, stimulus-specific mechanisms that are typically triggered by a very specific stimuli and cause a particular effect downstream) can qualify as genuine representations. Many philosophers reject liberal representationalism for various reasons and seek to provide stringent criteria for representationality that are not met by mere detectors, but which are satisfied by more sophisticated mechanisms (Ramsey [2007]; Gładziejewski and Miłkowski [2017]; see also Schulte [2015]; Butlin [2020]).[1]

Of course, these two debates intermingle in various ways. A key point of connection is the objection of trivialization: anti-representationalists typically accuse theories of representation of being too liberal (Hutto and Myin [2017]). Thus, by appealing to structural representations, representationalists aim to vindicate some representational explanations in cognitive science and, at the same time, block the objection that these theories are excessively liberal (Ramsey [2007]; Gładziejewski and Miłkowski [2017]; cf. Artiga [2016]; Rupert [2018]; Ganson [2020]).

As a result, a common core among recent proposals that appeal to this view is that this notion 'serves as an important, distinct, and explanatorily valuable posit of classical computational accounts of cognition' (Ramsey [2007], p. 79). Since a main goal of the article is to discuss this view, let us provide a more precise characterization of what we might call 'structuralism':

Structuralism: Structural representations constitute a distinctive and paradigmatically representational kind in virtue of the fact that a structural correspondence is used or exploited to accomplish a certain task.

Here 'distinctive' means that it is a mechanism that significantly differs from detectors. Crucially, note that structuralism not only holds that representational mechanisms use structural correspondences, but also they are paradigmaticaly representational in virtue of using them. The claim is that exploiting such correspondences is one of the key features that turn a mechanism into a clear and explanatory central instance of representation. Different versions of this view have been defended, for instance, by Bartels ([2006]), Ramsey ([2007]), Churchland ([unpublished]), Głaziejewski and Miłkowski ([2017]), Opie and O'Brien ([2004]), Plebe and de la Cruz ([2018]), Williams and Colling ([2018]), Shea ([2018]), Swoyer ([1991]), or Lee ([2019]). Piccinini ([2018], p. 3) describes the state of art as 'an emerging consensus that the best way to understand representation in the context of cognitive explanation is structural'.

---

[1] A third motivation is to provide a theory of representational content (O'Brien and Opie [2004]; Bartels [2006]; Neander [2017]; Plebe and De La Cruz [2018]). This issue will not be discussed in this article (see Nirshberg and Shapiro [2020]; Segundo-ortin and Hutto [2021]).).

As suggested above, a key motivation for structural representation is the analogy with maps. The reasoning can be succinctly expressed as follows: Maps are well-known and paradigmatically representational artefacts, so if some mental mechanisms are employed in an analogous way, their representational status can also be secured (Cummins [1996]; Ramsey [2016], p. 7). More precisely, a map user manages to navigate a certain environment by exploiting a structural correspondence that it bears on the target domain. Thus, if a mental mechanism possesses a similar structural correspondence and the brain exploits it to successfully perform some task, their representational nature can rest on firm ground. Furthermore, since mere detectors significantly differ from maps, this position can vindicate the representational status of certain mental mechanisms without trivializing this notion.

In this article, I would like to assess this strategy in two steps. The first one is a project of conceptual clarification: I consider there are significantly different ways of understanding the notion of structural representation that do not primarily derive from distinct interpretations of structural correspondence (as one may have thought), but from various ways of defining the concepts of exploitation and structure. Indeed, I will suggest that these notions have been understood in various ways in the recent literature, which has caused certain misunderstandings (see sec.2).

The second goal is more critical: I will argue that it is unclear that any interpretation of the concept of structural representation can vindicate structuralism. At least, plausible ways of understanding this concept do not seem to support structuralism. One the one hand, if one embraces a broad definition, then this concept is much more liberal than structuralists envisage, and it might fail to capture the set of states that theorists are willing to categorize as distinctive, explanatorily valuable and paradigmatically representational. On the other hand, if structuralists provide a more stringent analysis of structural representations relying on different ways of specifying 'structure' and 'exploitation' in more detail, one might worry that the real explanatory work is being done by these additional features; furthermore, the resulting analysis might fail to identify a robust scientific kind. Consequently, any interpretation of structural representation designed to support structuralism faces important challenges.

The article is organized as follows: In section 2, I will provide an initial definition of structural representations that is compatible with the claims of some authors and I will argue that it is too liberal to vindicate structuralism. Although others have argued for that claim before, I will show that a lack of precision has led to certain misunderstandings in the debate. This initial result will be used to motivate the need for clearer and more precise analyses of this concept and to assess if they can vindicate structuralism. In section 3, I will consider additional conditions for restricting the notion of structural representation and I will present particular cognitive mechanisms that meet them: using relations between vehicles (sec. 3.1), relation-sensitivity (sec. 3.2), richness (sec. 3.3) and off-line use (sec. 3.4). The goal is not to provide an exhaustive analysis, but to illustrate that 'structure use' is too cheap a requirement, and that more interesting kinds are identified by focusing on different structures and the particular ways of exploiting them. In section 4, I will argue that any interpretation that seeks to vindicate structuralism faces important challenges. In the final section (sec. 5), I will summarize the main results and point to the different options available to the structuralist.

2. Structural Representations: First Pass

What is a structural representation? A central idea of the standard conception claims that a structural representation preserves or, at least, possesses a structure that is similar to its target. This relation is often cashed out as a homomorphism between representations and target systems. A homomorphism is a function from one set of entities to another set (in this case, from the set of vehicles to a set of world entities). Let us use h_i for the world entities, H for a relation between them, v_i for the putative vehicles, and V for the relationship between them. There is a homomorphism between the set $\{v_i, V\}$ and the set $\{h_i, H\}$ if and only if there is a function f that maps each v_i to a h_i and, crucially, if the relation V obtains between two vehicles v_1 and v_2, then relation H obtains between $f(v_1)$ and $f(v_2)$. More precisely (see Shea [2018], p. 117):

Structural Correspondence: There is a structural correspondence between $\{v_m, V\}$ and $\{h_n, H\}$ if and only if there is a function, f, that maps the v_m onto the h_n and $\forall_{i,j}(V(v_i, v_j) \leftrightarrow H(f(v_i), f(v_j)))$.

Of course, the mere existence of a structural correspondence does not suffice for there to be a structural representation, since it is not difficult to find homomorphisms between sets of entities that are not representational at all (Goodman [1976]; cf. Cummins [1996]). To qualify as a structural representation, a mechanism needs to be used, employed or exploited by a system in order to deal satisfactorily with a target system because of the structural correspondence it bears to it (Opie and O'Brien [2004]; Bartels [2006]). Thus, it is in virtue of exploiting a structural correspondence between representations and a target domain that the system can successfully perform a task.

Although the appeal to structure use certainly helps narrow down the set of mechanisms that qualify as structural representations, some philosophers have recently objected that this proposal is still too liberal. In particular, they argue that according to this analysis '[detectors] just are structural representations' (Morgan [2014]; see also Nirshberg and Shapiro [2020]; Facchin [2021]). Obviously, if detectors (such as photoreceptors or themostats) actually qualify as structural representations, that would seriously compromise the theoretical role that the notion of structural representation is supposed to play according to structuralists, as one of the main goals of developing this concept is precisely to oppose liberal views.

So why do these philosophers think detectors meet the conditions for being a structural representation? Consider the following example: Many bats emit ultrasonic sounds to echolocate prey and, as a response, some insects such as the nocturnal cricket (*Gryllus bimaculatus*) have evolved auditory organs sensitive to bat ultrasounds (Miller et al. [2001]; Marsat and Pollack

[2012]). Suppose a cell in the cricket auditory organ can be in two states, on and off, which respectively correspond to the presence and absence of ultrasound waves. It is relatively trivial to find some relation between these two vehicles (on and off) and some relation between the two relevant world events (that is, presence and absence of ultrasonic sounds), such that the former maps onto the latter; for example, consider the relation '50 times higher frequency' that holds between the firing rate of the cell being on and its being off, and the relation '20,000 times higher frequency' that holds between a low frequency sound (regular environmental noise) and ultrasound waves. This an instance of homomorphism between, on the one hand, the states on and off and their relation and, on the other, the two relevant stimuli (ultrasound present and absent) and their relation. Consequently, $\{v_i, v_j, 50\ \textit{times higher frequency}\}$ structurally corresponds to $\{h_i, h_j, 20,000\ \textit{times higher frequency}\}$. Furthermore, since crickets exploit the fact that a change in cell activity corresponds to a change in the presence of ultrasounds to behave successfully (for example, to escape from bats), it seems that this simple mechanism meets all conditions for qualifying as a structural representation.

Discussing this debate between structuralists and critics is relevant not only to illustrate the need for conceptual clarification, but also because it will help us formulate a first precise way of understanding the notion of structural representation. As a matter of fact, I would suggest that structuralists, on the one hand, and Morgan, Nirshberg, Shapiro and Facchin, on the other, are probably talking at cross-purposes. This unfortunate situation derives from not being sufficiently clear about how the notion of structural representation should be understood. Let me elaborate.

We saw that structural representations involve a mechanism that bears a structural correspondence with a target domain, and which the system uses to accomplish certain tasks. Here is a particular way of making this idea more precise:

DEF1: In system S, $\{v_m, V\}$ is a structural representation if and only if

(1) there is a structural correspondence between $\{v_m, V\}$ and $\{h_n, H\}$;

(2) condition 1 provides an explanation for how system S uses v_m to successfully perform some task F.

Crucially, note that according to this interpretation the system uses a set of vehicles to behave successfully. Nonetheless, the homomorphism is essential to explain success because, given the way the system is built, each vehicle manages to be triggered in the right circumstances thanks to the existence of a homomorphism between neuronal states and their images.

I would say that Morgan ([2014]), Nirshberg and Shapiro ([2020]) and Facchin ([2021]) are probably right that if structural representations are understood as DEF1, detectors typically qualify as such: the system employs each particular state of an on–off cell to detect the presence or absence of an ultrasound, and the mechanism works properly when a structural correspondence between levels of activity and wave frequency holds. Thanks to the homomorphism each vehicle is tokened at the right time. Accordingly, in a broad sense, the mechanism exploits or uses the homomorphism to behave successfully. However, this is but one way of understanding the concept of structural representation. As I will argue in the next section, there are other interpretations that also fit the claims that structuralists typically make and which fail to include detectors within this category (whether this account is completely successful will be discussed below). In other words: when Morgan and others argue that detectors meet the conditions for being a structural representation, they are assuming one possible understanding of 'structure use'. However, I will present alternative ways of developing this notion according to which these examples are not classified as such. Thus, some conceptual clarification is needed for the debate to move forward.

With respect to our main issue, structuralism, this example shows two things. First, that according to some interpretations (for example, DEF1) structural representations include detectors, such as the simple on–off cells presented above. This indicates that some definitions of structural representations might fail to capture the set of states that proponents want to classify as distinct and paradigmatically representational. Accordingly, if structural representations are understood as DEF1, they fail to support structuralism. Second, this short discussion suggests that other interpretations are possible; DEF1 is only one way of understanding structural representations, and other (perhaps more charitable) readings might be in a better position to support structuralism. As a result, we still need to assess whether a slight modification of DEF1 can provide a concept that could eventually play the two main theoretical roles that would vindicate structuralism. This is the task of the remainder of this article.

3. Structural Representations: Second Pass

Assuming that DEF1 is too weak to vindicate structuralism, can a more stringent analysis of structural representations identify a mechanism that is distinctive and paradigmatically representational in virtue of the fact that a structural correspondence is used or exploited to accomplish certain tasks? In this section four strategies will be discussed: appealing to exploiting relations, relation-sensitivity, richness and off-line use. For ease of exposition, in this section I will only present the different definitions and some mechanisms that satisfy them; whether they can vindicate structuralism will be discussed in the next section.

3.1. Exploiting relations

Let us think again about the key example inspiring the notion of structural representation: maps. If DEF1 were an accurate description, in using a map of Europe one would only exploit the fact that each dot corresponds to a different city, yet this is a poor characterization of our use of cartographic representations. We not only exploit the fact that dots correspond to places, but also (and crucially) we gain knowledge about the relations between cities by paying attention to the relations between dots. Only by exploiting the connections and distances between dots can we ascertain the connections and distances between cities. This is how we actually manage to travel from Paris to Berlin using a map of Europe.

Based on this idea, structural representations could be defined as follows:

DEF2: In system S, $\{v_m, V\}$ is a structural representation if and only if

(1) there is a structural correspondence between $\{v_m, V\}$ and $\{h_n, H\}$;

(2) condition 1 provides an explanation for how system S uses $\{v_m, V\}$ to successfully perform some task F.

It should be understood that the explanation mentioned in condition 2 appeals to v_m corresponding to h_n and V corresponding to H. In other words, in this analysis the system not only exploits the fact that each vehicle corresponds to a particular entity, but also relies on relations between vehicles corresponding to relations between their images. As a result, and in contrast to DEF1, success here crucially depends on the correct mapping of both vehicles and relations.

Now, can DEF2 vindicate structuralism? Does it capture a set of states that that are distinctive, explanatorily valuable and paradigmatically representational in virtue of using structure to perform certain tasks? Note that not just any brain state would count as a structural representation according to DEF2. Let me describe three different mechanisms that fail to satisfy DEF2, some of which have been claimed to be structural representations by some critics (see, for instance, Morgan [2014]; Nirshberg and Shapiro [2020]).

Let us consider again the simple mechanism discussed in the previous section. Following Morgan ([2014]) and others, I argued that simple on–off cells that are supposed to track ultrasounds meet DEF1. In contrast, I think this mechanism fails to satisfy DEF2. The system certainly uses the individual states on and off because they carry information about ultrasounds, but there is no sense in which the system employs the relation between these states to deal better with the relation between the presence or absence of ultrasounds. The system cares about the different stimuli not about the relationship between them. Of course, it is actually built in such a way that in normal conditions each vehicle only correctly informs about the world when the relation between vehicles corresponds to a relation between their images, but that does not mean that the system exploits internal relations to learn about external ones.

To provide some support for that claim, we can employ the following test (suggested for different reasons by Gładziejewski and Miłkowski [2017]; Shea [2018], p. 142):

Test: If strengthening the correspondence between relation V on vehicles v_m and relation H on entities h_n increases the likelihood of the system achieving its function F (and weakening the correspondence decreases this likelihood), then we have some evidence that $\{v_m, V\}$ and the structural correspondence between $\{v_m, V\}$ and $\{h_n, H\}$ are used or exploited to perform task F.

Suppose we inject a substance into crickets that modifies the relationship between the firing patterns of a particular cell being on and its being off (for example, on has now 100 times higher frequency than off, rather than the usual 50), but make sure that, nonetheless, on and off states still correspond to the presence and absence of ultrasounds as effectively as before. If the system were using the correspondence between relations, this manipulation would alter it and lead the system to failure. However, it should be obvious that intervening in this cell in this way would not affect the success rate of crickets. Hence, it fails to pass the test above, which indicates that the mapping between relations is irrelevant in the case of simple detectors in crickets.[2]

Consider now an analogue detector, such as a photoreceptor that modifies its firing rate in response to light changes. There is a homomorphism between different levels of activity and different light intensities but, again, it is not obvious that the correspondence between the relations between firing rates and relations between light intensities is being used by the system at all (cf. Morgan [2014]). What seems essential is that every light intensity corresponds to one and only one level of cell activity. Here is one way of presenting the worry[3]: this system could be modelled as a list of representational vehicles v_m, their corresponding light intensities, and their effects downstream. We can explain perfectly well how the system manages to perform a task by referring to this list, which does not consider the relations between vehicles or between elements in the target domain at all. Indeed, the relationship between vehicles v_m could be changed at will and, insofar as each vehicle still corresponds to one light intensity, the mechanism would still work perfectly well. Thus, it seems that weakening the correspondence between relations V and H need not result in the system achieving its function less well, and thus not not passing the test. Therefore, this mechanism fails to satisfy DEF2 and, as a result, does not qualify as a structural representation in this sense.

Interestingly, even a topographic structure might not qualify as a structural representation according to DEF2 (cf. Nirshberg and Shapiro [2020]). V1 contains a topographic representation of the visual field, meaning that the distance between two cells in V1 roughly corresponds to the distance of their respective receptive fields. Neighbouring points in the cortex correspond to nearby points in parameter space. One possible explanation for that organization is that this structure is more efficient in terms of energy costs; cells with nearby receptive fields are likely to develop strong connections, which are energetically costly, so by placing cells with strong connections close to each other, costs are minimized (Durbin and Mitchison [1990]). If this hypothesis were on the right track, V1 would not qualify as a structural representation according to DEF2 (at least, not in virtue of there being a topographic mapping). Despite the existence of an homomorphism in which spatial relations in V1 map onto spatial relations in the visual field (so, in that respect, just like a map, which is the paradigmatic case inspiring the

---

[2] To be clear, I am arguing here that some detectors don't satisfy DEF2; others, however, might satisfy it (see below).

[3] I would like to thank Nicholas Shea for helping me think through these examples.

idea of structural representation), these relations are not exploited. Again, we can run our test: if the only reason for having a topographic representation in the primary visual cortex is cost minimization, the mechanism would achieve is function as well as it actually does even if we changed the position of some cells (without modifying their receptive fields or efferents).

I think the same reasoning applies to some other suggestions made in the literature. Garzón ([2009]), Churchland ([2012]), or Shagrir ([2012]), for example, have argued that neural networks are homomorphic to the target domain and that they are extremely useful to perform certain tasks. What remains to be shown, however, is that the system employs the fact that relations among vehicles map onto relations among represented entities in such a way that, for instance, weakening the correspondence between relations while maintaining the rest of correspondences causes the network to perform certain tasks less successfully. Perhaps this can be shown; my point is simply that, assuming DEF2, this is the work to be done.

Consequently, DEF2 is not as liberal as DEF1, since it can correctly exclude some counterexamples suggested in the literature that the previous analysis could not rule out. Does that mean that DEF2 can vindicate structuralism? Not quite; there are still very simple mechanisms that do satisfy the two conditions of DEF2, so this definition is probably too weak to capture the distinct, explanatorily valuable and paradigmatically representational kind that structuralists are after. Associative learning provides such an example.

Associative learning is a learning process by means of which animals learn about predictive relationships between events and behave appropriately as a result (Shettleworth [2010], p. 107). Bees, for instance, learn to extend their proboscis (an elongated appendage that works like a tongue) as a response to a particular odour (conditioned stimulus, or CS), if it is regularly presented with sucrose (unconditioned stimulus, or US) (Takeda [1961]; Bitterman et al. [1983]). Interestingly, associations between some kinds of stimuli turn out to be easier to learn than others. For instance, bees will more easily learn a connection between sucrose and colours or odours, than between sucrose and other kinds of stimuli (Menzel and Müller [1996]). This is also true of instrumental conditioning: pigeons learn more easily to associate a light than a tone with food, and a tone than a light with danger (Shettleworth [2010], p. 113). A plausible explanation seems to be that associative learning is a mechanism for learning causally related events,[4] so if one event is more likely to have a causal connection with another, then it should take less evidence for the animal to learn it. This feature is what has been called ‘relevance’ (Dickinson [1980]) or ‘predispositions’ (Escobar and Miller [2004]), among other expressions (Shettleworth [2010], p. 114).

This mechanism satisfies DEF2 (cf. Shea [2014], p. 127). In the case of associative learning just described, two vehicles (CS and US) map onto two properties (yellow and sucrose,

---

[4] As a reviewer pointed out, the evidence mentioned here does not allow us to conclude that the strength of the neural connection corresponds to a causal connection rather than, say, a strong correlation. In any case, that does not affect my main point: the success of the mechanism requires that the relation between vehicles latches onto some objective relation, be it causal, correlational or of some other sort.

respectively). Furthermore, the relationship between these vehicles is supposed to correspond to some causal relationship between the entities (as the existence of 'predispositions' shows), so part of the explanation of why a learned association leads to successful action is that the connection between vehicles maps onto a causally underpinned connection between the entities represented by them. In a natural setting, a bee's association of yellowness and nectar is useful precisely because of the causal mechanism underpinning this connection.

This mechanism also passes the test. Imagine that some playful child paints all flowers of her father's garden pink and as a result, bees from a nearby nest learn to associate pinkness with nectar. Compare it to a different group of bees whose nest lies near a sunflower field and learn to associate yellowness with food. The correspondence between the strength of the neuronal connection (relation V) and the robust causal link between properties (relation H) is much worse in the first group than in the second and, as predicted by the test, this results in worse performance. In particular, the second group possess a much more resilient mechanism for identifying food, thanks to the fact that strong connections actually correspond to a causally underpinned link between properties. In contrast, when the paint disappears, the first group of bees will lack a reliable mechanism for finding nectar. Consequently, associative learning satisfies DEF2 and, as result, should count as a structural representation according to it.

Now, can DEF2 lend support to structuralism? We will address this question in section 4. Let us now consider a different way of developing the concept of structural representation: relation-sensitivity.

3.2. Relation-sensitivity

Shea ([2014], [2018]) provides a still more restrictive analysis of structural representations. According to him, in order to exploit a structural correspondence, some component needs to be sensitive to the relation between vehicles, such that this sensitivity is part of the explanation of how the mechanism manages to perform its functions. 'Sensitivity' is here a causal notion. As he argues, 'a necessary condition for an isomorphism to be exploited is that relations between putative representations should make a difference to downstream computations or behavior' (Shea [2018], p. 127). Thus, let us define the following property:

Relation-Sensitive: A mechanism, M, is relation-sensitive if and only if some internal component, C, of M is sensitive to the relation V between vehicles.

With this notion in mind, we can define a more stringent analysis of structural representations (see Shea [2018], chap. 5):

DEF3: In system S, $\{v_m, V\}$ is a structural representation if and only if

(1) there is a structural correspondence between $\{v_m, V\}$ and $\{h_n, H\}$;

(2) relation-sensitive: some internal component C is sensitive to relation V;

(3) conditions 1 and 2 provide an explanation for how system S uses $\{v_m, V\}$ to successfully perform some task F.

In this sense, for a mechanism to qualify as a structural representation it is not enough for a system to use the relations between vehicles to learn about relations in the target domain; in addition to this, the system must exploit the structural correspondence partly by including a relation-sensitive component. The mechanism underlying associative learning that I described above, for instance, probably lacks such a component; in general, associative learning does not require a mechanism to be sensitive to the relationship between the vehicles CS and US. To meet the relation-sensitive requirement, there should be a third mechanism that modifies its response in reaction to changes on the strength of the connection between CS and US, for example.

DEF3 is certainly more restrictive than DEF2. One mechanism that satisfies this analysis is that of binding. The so-called binding problem is a classic difficulty in understanding how the brain processes perceptual information. We know that complex stimuli (such as a yellow circle and a green square) are broken down during sensory processing and processed in different parts of the brain. Nonetheless, somehow the brain manages to keep track of which properties belong to the same object and which do not (for instance, that yellowness and roundness are instantiated in the same object, whereas greenness goes with squareness). How does our perceptual system solve the problem of binding the different properties that belong to the same object after splitting them and processing them in different brain areas? A candidate mechanism for accomplishing this task involves the use of neural oscillations.

For a long time it has been known that neuronal activity follows a rhythmic pattern, although the role of these neural oscillations remains a matter of speculation. One hypothesis suggests that they play an important role in solving the binding problem: In short, the idea is that if the firing patterns of two neurons are synchronized (which means they fire at roughly the same phase of the oscillatory pattern), then the brain interprets that the properties represented by these two cells are co-instantiated (von der Malsburg and Schneider [1986]; Singer and Gray [1995]; Engel and Singer [2001]). Rhythmic co-firing indicates that the properties represented by the synchronic neural states belong to the same entity. For instance, if a neural state v_1 carries information about roundness and v_2 stands for yellowness, the coordinated tokening of v_1 and v_2 indicates that yellowness and roundness are instantiated by the same object. A classical experiment supporting this hypothesis is presented in figure 1.

This example satisfies all conditions included in DEF3 (Martínez and Artiga [2023]). Consider the (highly simplified) example involving a yellow circle and a green square: if v_1 represents roundness and v_2 represents yellowness, relation V (synchrony) between vehicles v_1 and v_2 is supposed to correspond to the relationship of co-instantiation between $f(v_1)$ and $f(v_2)$. Accordingly, a relation of structural correspondence holds between $\{v_1, v_2, V\}$ and {yellowness, roundness, co-instantiation}. Second, note that this mechanism is relation-sensitive, since some

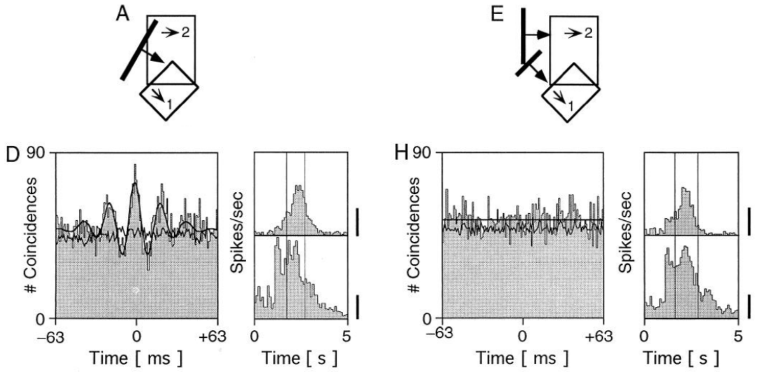

Figure 1. The experiment considers the synchronic firing patterns of two cells in the middle temporal area of an awake macaque monkey when (A) a single bar crosses the receptive fields of two neurons (1 and 2), and (E) when two bars cross the same receptive fields. (D) shows that when the two stimuli belong the same object, the firing patterns of the two cells become synchronized, whereas (H) shows no temporal correlation when the bars do not belong to the same object (from Kreiter and Singer [1996]; copyright 1996 Society for Neuroscience)

component downstream takes into account the synchronic pattern. More precisely, it has been suggested that certain neurons work as 'coincidence detectors' (Abeles [1982]; König et al. [1996]). As a result, this mechanism meets all conditions included in DEF3.

So far, we have discussed three different ways of understanding structural representations, DEF1, DEF2, and DEF3. However, these do not exhaust the plausible interpretations.

3.3. Rich structures

A different strategy to develop a more sophisticated analysis of structural representations focuses on one of the limitations of neural oscillations as a solution to the binding problem. Notice, for example, that in the mechanism described in the previous section, the relation between vehicles always has the same content: it indicates co-instantiation. So although the structural correspondence is used, it is always exploited to latch onto one and the same property, namely, the fact that two features are instantiated by the same object. Perhaps a more interesting notion can be developed by considering mechanisms that perform more diverse operations on vehicles.

Gallistel ([1990]) provides one way of developing this idea. Among what he labels 'functioning isomorphisms' (which roughly correspond to what others call 'structural representations'), he makes a distinction between rich and impoverished structures. What distinguishes them is the diversity and complexity of the operations that the brain performs on vehicles. For instance, the most impoverished form of numerical representation ('nominal representations') only accepts the 'equal' or 'identity' being operated (Gallistel [1990], p. 26). The mapping from players on a team to jersey number is a nominal representation: the player that scored a goal and the player that fouled an opponent are the same player if and only if the number of the first player and the number of the second player are the same. An example he provides of rich functional homomorphism is path integration, in which organisms need to combine different sources of information (coming, for instance, from head direction, optic flow, polarized light and proprioception; see Stone et al. [2017]) and perform certain operations in order to find their way.

Relying on this interesting suggestion, we can define the following property:

Rich: The use of a structural correspondence is rich if and only if it performs different mathematical operations (addition, multiplication, and so on) on vehicles.

If this condition is added to DEF3, we obtain a more stringent analysis, that can be labelled DEF4. The solution to the binding problem involving neural oscillations, for instance, probably fails to satisfy DEF4. If neural oscillations roughly work in the way I explained above, they actually constitute a very impoverished sort of functioning isomorphism, in the sense that the only operation of this system involves a co-instantation sign: if neuron A and neuron B fire at the same time, they represent properties of the same object. This system does not include other kinds of mathematical operations (sum, multiplication, orderings, and so on). Again, to illustrate this analysis, let us discuss a mechanism that seems to satisfy DEF4: template matching.

'Template matching' refers to a mechanism animals use to locate a food source (Cartwright and Collett [1983]; Shettleworth [2010]). In brief, the idea is that if an animal finds food near some landmark at time t_1, it can take a snapshot of how the world looks from that place and store it in its memory, which will be compared with its perceptual representation at a subsequent time t_{1+n} to move to the very same location. More precisely, when at t_{1+n} it spots the very same landmark from a distance, it can compare the current perceptual scene at t_{1+n} with the stored snapshot taken at t_1 and the difference between them will indicate the distance to the food source. By moving in the direction that minimizes the difference between the two pictures, the animal will approach the food source. Template matching is employed by a variety of organisms, such as insects (Collet and Baron [1994]), chickens (Dawkins and Woodington [2000]), rodents (Leonard and McNaughton [1990]), and fish (Schuster and Amtsfelt [2002]).

Although nowadays the existence of template matching is widely accepted, there is still an ongoing dispute as to how bees actually compute the difference between the snapshot and the current perceptual state. Different models appeal to operations such as subtracting the average landmark vector of the current scene from the average landmark vector of the snapshot stored in memory or using the ratio of the luminance-amplitude and luminance-background within a 3 × 3-pixel window, which is then weighted by the inverse of the distance, among others (Muüller and Wehner [2010]; Egelhaaf et al. [2014]; Doussot et al. [2020]). It should be obvious that any of these models would vindicate the idea that this mechanism involves not only a functioning isomorphism, but also a rich one.

This mechanism seems to satisfy DEF4. First, the stored snapshot v_1 represents the position h_1 where food was found; the current perceptual state v_2 represents the current position h_2, and the difference V between vehicles represents the distance H from h_1 to h_2. Thus, there is a structural correspondence between (v_1, v_2, V) and (h_1, h_2, H) and relation V is used to learn about H. Furthermore, it passes our test: the mechanism relies on the correspondence between the difference value V and the distance between positions H, in such a way that if there were a mismatch between these two values, they would do less well in foraging. The relation-sensitive condition is also met, since some component is sensitive to V and, finally, it is also rich: the calculation of the discrepancy between snapshot and perception involves different mathematical operations (vector subtraction, averaging, and so on).

To conclude this discussion of different interpretations of the concept of structural representation and before addressing the plausibility of structuralism, let us finally consider the classic examples that structrualists typically: cognitive maps in the rat's hippocampus.

3.4. Offline

Structuralists' most central example of structural representation is the cognitive map in rats (Shea [2014], [2018]; Gładziejewski and Miłkowski [2017]; Ramsey [2016]; Williams and Colling [2018]; Lee [2019]). For this reason, I want to devote this last subsection to briefly discussing this example and showing how it fits into the reasoning of this article.

Our best theories of how rats develop and employ cognitive maps are rooted in the line of research started by O'Keefe and Dostrovsky ([1971]), who discovered that some cells in the hippocampus preferentially fire when the rat is in a particular area of the environment. These neurons were aptly called 'place cells', and boosted a research programme that led to an intense study of their properties and the discovery of similar cells with different features (grid cells, head-direction cells, boundary cells, and so on) that also seem to play a role in the formation and deployment of cognitive maps (Hartley et al. [2014]; Colgin [2020]). Despite the extremely interesting details of this research (see Bechtel [2016]), given the lack of space I will focus here on those features that are relevant to our purposes.

A first remarkable aspect is that place cells do not form a topographic map (or, at most, only possess a very limited topographic organization restricted to small, distributed clusters of neurons; see Redish et al. [2001]; Nakamura et al. [2010]). In other words, two neighbouring cells are not more likely to represent nearby places than any other two random cells. Consequently, in the same way that I argued that having a topographic structure is not sufficient for qualifying as a structural representation (see sec. **??**), this shows that neither is it necessary. In this sense at least, cognitive maps are rather unlike maps. In any case, although distances between places do not correspond to distances between cells, there is actually a relation between cells that maps onto a relation between places: neurons that represent nearby places tend to co-activate each other. Thus, a relation of structural correspondence holds in virtue of an abstract mapping relation between co-activation patterns and distances.

Interestingly, some evidence suggests that rats use their place cell configuration in order to plan their future actions. For instance, it has been shown that before traversing a path the very

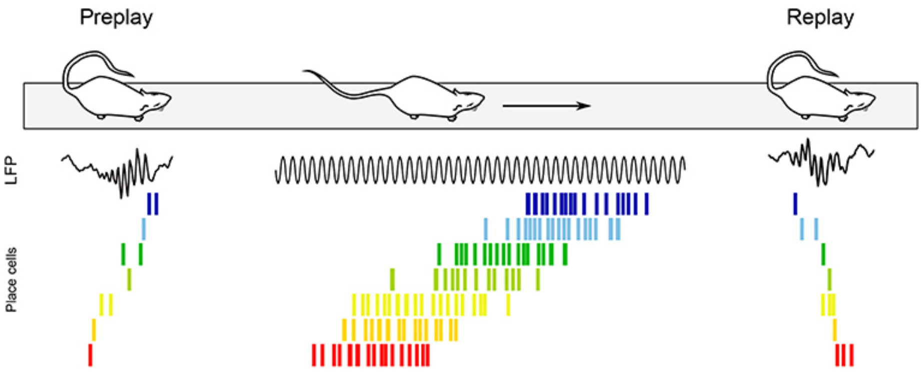

Figure 2. Spike trains of place cells are shown before, during, and after a single traversal. Sequences that occur during the track are reactivated before (forward preplay) and after (reverse replay) the run (from Drieu and Zugaro [2019], p. 13).

same cells that would fire in the traverse are activated in the very same temporal order (see fig. 2). This forward preplay indicates that the rat is anticipating or planning the future route (Diba and Buzsaki [2007]; Dragoi and Tonegawa [2011]; Drieu and Zugaro [2019]). Following Shea ([2018], p. 115), let us suppose that preplay help rats choose the shortest routes by selecting the shortest sequence of cell firing. This seems to be an interesting example of offline reasoning: roughly, some process takes place offline (as opposed to online) when it is not directly driven by current environmental input. In this case, the rat's capacity to reason offline relies on the existence of a structural correspondence between the sequence of place cells being activated and the path formed by the cells' represented locations.

Note that this example meets all conditions stated in DEF4: there are certain vehicles v_i and relation between vehicles V, which are supposed to map onto a certain spatial arrangement. It also contains some component that is sensitive to the relation (since we are supposing that the rat is able to compare different routes by sensing the sequences of place cell activations) and it has a rich structure that probably involves performing different operations on vehicles. Nonetheless, cognitive maps possess an extra feature that distinguishes them from previous mechanisms: offline use. So, as a matter of fact, cognitive maps fit a more stringent notion of structural representation that we could label DEF5, which adds this aspect to the list. Indeed, structuralists tend to stress this feature; as Morgan ([2014], p. 237) suggests, 'my impression is that there's a tendency in the literature to assume that for a representation to be used as a surrogate *just is* for it to be manipulated offline'.

Now, can DEF5 or any of the previous definitions vindicate structuralism? This is the main question to be addressed in the next section.

4. Discussion: Structuralism and Structural Representations

At the beginning of the article I put forward two goals for this article: to elaborate in some detail different ways of understanding the notion of structural representation and to critically assess structuralism, that is, the claim that structural representations constitute a distinctive and paradigmatically representational kind in virtue of the fact that a structural correspondence is used or exploited to accomplish certain tasks. In this section I aim to discuss how the different definitions and mechanisms explored in previous sections bear on these two goals.

4.1. Conceptual clarification

With respect to the project of conceptual clarification, the previous discussion suggests a wide variety of interpretations of the concept of structural representation. All definitions I have provided can be understood as specifications of the idea that a structural representation is a mechanism in which a structural correspondence between a set of vehicles and a target domain is used to accomplish a certain task. More precisely, alternative readings differ on how to understand using or exploiting a structural correspondence (using relations, relation-sensitivity, off-line use) and some properties of the structure (rich or impoverished isomorphism). Obviously, this brief survey of different features is not supposed to be exhaustive (other properties like 'decoupling' or 'error detection' could also be mentioned; see Gładziejewski [2015]; Piccinini [2020], p. 261; Arnellos and Moreno [2021]), but I think it suffices to show that when someone describes a mechanism as using or exploiting a structural correspondence, care need to be taken in disentangling different senses of this expression.

Indeed, I also argued that structural representations have actually been understood in different ways. Critics of structuralism have often embraced a relatively liberal interpretation of structural representations along the lines of DEF1, whereas a (perhaps more charitable) alternative reading suggests that structuralists probably have some other more sophisticated approach in mind. Unfortunately, they have typically failed to spell out in sufficient detail which specific interpretation they favour, thus a more careful analysis of this aspect will help the debate move forward.

4.2. Structuralism

The second goal of the article has been to argue that once these different interpretations are specified in detail, it is unclear that any of them can actually support structuralism. It is time to unpack and develop this claim.

First, we saw that the initial definitions of structural representation, DEF1 and DEF2, are too liberal to vindicate structuralism. DEF1 is met by simple detectors and DEF2 by associative learning, which might simply be composed of two connected detectors. Thus, these definitions do not allow us to draw a clear distinction between detectors and structural representations, as structuralism intends. Indeed, on this analysis any mechanism in which transitions between vehicles are supposed to map onto some important relation between the represented entities would count as a structural representation; as a result, perhaps any computational structure would qualify as such (Shea [2015]). The mere existence of a homomorphism between two systems is too cheap a requirement, and adding to it the idea that vehicles or relations between them are used to learn about world entities or their relations slightly constrains the set of mechanisms that satisfy this analysis, but not sufficiently. This shows that if the requirement of 'structure use' is left unconstrained, it actually provides an extremely weak condition. Understood in this way, the concept of structural representation fails to capture a distinctive representational kind that clearly differs from mere detectors.

There are various reasons why structuralists probably want to avoid this result (and lead to their rejection of DEF1 or DEF2). First, as mentioned above, one of the key reasons to develop a concept of structural representation is to oppose liberal representationalism and, in this way, block the objection of trivialization raised by anti-representationalists. If detectors are just structural representations, then the concept of structural representation cannot play one of the key dialectical roles that it is intended for. Second, even those who accept that detectors are minimally or marginally representational, would balk at the idea of classifying them as paradigmatic representational mechanisms; however, structural representation is said to be paradigmatically representational, so accepting detectors as structural representations implies giving up one of these entrenched assumptions. In other words, there is a tension between the following claims: (1) structural representations are paradigmatically representational, (2) detectors are not paradigmatically representational, and (3) detectors are structural representations. Third, many of the claims and descriptions put forward by structuralists would need to be revised; for instance, conceiving of structural representations as internal models or the link between structural representations and off-line use or surrogative reasoning (see Swoyer [1991]). Finally, some theorists explicitly claim that one of their key goals is to distinguish structural representations from detectors (for example, Ramsey [2007]; Gładziejewski and Miłkowski [2017]). Thus, to this extent at least, adopting this option would have certain theoretical costs.[5]

Structuralists, however, can appeal to alternative interpretations of the concept of structural representation that significantly differ from DEF1 and DEF2 and which might ultimately support structuralism. To assess this strategy on behalf of the structuralist, I explored alternative ways of understanding the notion of structural representation, by specifying different interpretations of 'use' and 'structure': relation-sensitivity, rich structures and offline use. Adding more constraints (for example, DEF3, DEF4, and DEF5) certainly brings us closer to the set of cases that the notion of structural representation is supposed to capture, since it selects a set of mechanisms that differ from mere detectors and play a clear representational role. Yet can these interpretations lend support to structuralism? The previous discussion suggests two key challenges that structuralists will need to meet to vindicate this strategy.

A first concern takes issue with the idea that structural representations are paradigmatically

---

[5] Let me stress that some people might happy with this result. For instance, if one is not really interested in drawing a line between detectors and 'genuine' representations and is not moved by trivializing objections raised by anti-representationalists, then one could adopt DEF1 or DEF2. However (and this is the crucial point), one would be abandoning structuralism, which I think is a widespread (standard, perhaps) approach in the area. Furthermore, some of the arguments of this section will also have consequences for this view, such as the tension highlighted in the main text.

representational in virtue of using a structural correspondence. The worry is that as we bring more complex analyses in, the notion of structural correspondence progressively loses its explanatory bite. Richness, offline use or relation-sensitivity add an important level of sophistication to a system and could by themselves explain why mechanisms having them are regarded as paradigmatically representational. In other words, my concern is that in more stringent definitions the bulk of the work is being done by these more fine-grained concepts. Let me provide two reasons to think this is the case.

First, it is certainly true that a structural correspondence probably constitutes a necessary condition for many different uses of a mechanism. The existence of a rich structure, for instance, presupposes a mapping between relations, and offline reasoning probably requires some form of structural correspondence. In this respect, perhaps a structural similarity is 'poised' (Gładziejewski and Miłkowski [2017]) or has a potential to be used in surrogative reasoning (Swoyer [1991]). However, I have argued structural correspondence is also necessary for a neural state to work as a simple detector: for a bipolar cell to match its possible states onto two environmental features, there has to be some structural correspondence that is being exploited (in some broad sense of the term). The argument, then, is that structural correspondences are being used in very simple and, according to many people, marginal instances of representation, so although they are also used in more interesting cases, it is reasonable to question if their explanatory value as representations is linked to properties shared with uninteresting cases, rather than the specific kind of exploitation.

This point is reinforced by the second reason for questioning the idea that structural representations are paradigmatically representational in virtue of using a structural correspondence: we have independent reasons for thinking that many of the properties added in DEF3–DEF5 account for important representational features. Being relation-sensitive, for instance, is far from a trivial property. According to some teleological theories this might a necessary condition for the relation between vehicles V to qualify as a genuine representation, since some of these theories require a receiver or consumer (for example, Millikan [1984]; Stegmann [2009]). It has also been argued that relation-sensitivity might distinguish explicit from implicit representations (Shea [2015]). Similarly, the distinction between rich and impoverished structures marks an important contrast; according to Gallistel, for instance, it distinguishes genuine representations from the rest. This supports the worry that the extent to which any of the mechanisms satisfying DEF3–DEF5 look like clear instances of representational mechanism might not depend on the fact that a structural correspondence is exploited as such, but on the specific kind of use or the structure richness.

Offline use probably provides the clearest illustration of this worry. One of the general and distinctive properties of representations is that A can represent B even if B is absent or even non-existent. Relying on this idea, Clark and Toribio ([1994], p. 419) suggest that 'reasoning about absent, non-existing or counterfactual states of affairs' is an instance of a 'representation-hungry' task, which they describe as a class of problems that are specially difficult to account in non-representational terms. Now, since offline use implies reasoning about absent or non-existing states of affairs, there are good reasons for thinking that any mechanism with this capacity constitutes a clear case of representation. Consequently, although the rat's cognitive map is a remarkable example of representational system, it is not obvious that it supports structuralism: what makes this mechanism interesting, explanatorily valuable and hard to accommodate by eliminativists crucially depends on the specific way of exploiting the structural correspondence (that is, offline) and its complexity, rather than the mere fact that a structural correspondence is used.

Let me make the structure of this argument as clear as possible. Imagine two mechanisms, M_1 and M_2. M_1 is paradigmatically representational and has properties F and G, while M_2 is not paradigmatically representational and has properties F and H. Additionally, suppose we have independent reasons for thinking that G accounts for some of the explanatory role that makes mechanisms clearly representational. Since F is possessed by both M_1 and M_2 and we know that G is important for qualifying as a representation, it would be natural to suppose that representational role is linked to G, rather than F or H. Of course, this is an inference to the best explanation, and it could go wrong in different ways. Nonetheless, I think it raises a challenge for structuralists.

To avoid possible misunderstandings, let me clarify that I am not denying that structural correspondences play some explanatory role. For instance, Gładziejewski and Miłkowski ([2017]) plausibly argue that they are causally relevant for behavioural success. Indeed, I argued above that detectors rely on structural correspondences to work successfully. My worry does not concern their explanatory role in general, but whether they make an explanatory contribution that is relevant to qualifying as a representation. In other words, these difficulties question the connection between their alleged explanatory contribution and representational status. Structuralism, recall, claims that structural representations are paradigmatically representational in virtue of using a structural correspondence; this is the link being questioned here.

So far I have developed the first challenge for structuralists, who are willing to adopt an analysis of structural representations along the lines of DEF3–DEF5. Yet assuming one of the complex definitions to vindicate structuralism faces a second difficulty. I assume a central goal of theorizing is to develop concepts that somehow carve up nature at its joints.[6] Now, if the notion of structural representation captures a distinct and paradigmatically representational kind as structuralism contends, it should presumably specify a cluster of properties that tend to co-occur, even if locally or imperfectly (Shea [2018]). However, it is unclear that the features I appealed in order to develop the different interpretations form anything like a natural cluster. More generally, the idea of 'using' or 'exploiting' a structural correspondence might actually be a hodgepodge of very different ways that structural correspondences can contribute to a

---

[6] I'm assuming foes and critics of structuralism alike regard representation as a scientific concept (cf. Rupert [2018]); the disagreement concerns the theoretical role that this concept can play.

system's performance. Consider DEF5: if the meaning of structural representation is captured by a definition along these lines, it probably picks out a gerrymandered set of properties that do not naturally cluster together and can easily come apart. For instance, detectors can be used offline (as admitted by Gładziejewski and Miłkowski [2017], p. 350); circadian rhythms might provide an example of time detector that can work correctly for some days without having direct input for the environment (Bechtel [2009]; for other examples along similar lines, see Facchin [2021]). Likewise, the question of whether a computational process is rich or impoverished seems to be independent of whether relations between vehicles are used to learn about relations between their objects and the existence of a relation-sensitive component is not obviously related to the existence of an offline use or a rich homomorphism. Hence, if one of the more stringent definitions of structural representation I developed above is adopted, the concept might fail to capture a cluster of features that tend to co-occur. As a consequence, the resulting concept might fail to latch onto a robust explanatory kind.

This worry concerns the concept of structural representation, not structural correspondence. The latter captures a clear and well-defined mapping relationship and it might well capture a scientific kind. In contrast, a definition of structural representation along the lines of DEF5, for example, includes some features that, *prima facie*, do not seem to cluster together. The reasoning here is similar to Griffiths's attack on the notion of innateness: according to Griffiths ([1997], p. 72), the concept of 'innate' bundles together three different ideas (developmental fixity, species nature and intended outcome) and since 'the three are empirically dissociated, a theoretical construct that conflates them is undesirable'. Similarly, relation-sensitivity, rich structure and offline use seem to be empirically dissociated; if this is true, the a concept that includes them might be misleading. This worry can be understood as a challenge for structuralists who want to accept a complex definition along the lines of DEF3–DEF5: they need to show that the properties included in the analysis tend to cluster together (or, alternatively, argue that it is acceptable to keep a concept that bundles together an unrelated set of phenomena).

Summing up, I think the results of this article put some pressure on the structuralist position. First of all, since multiple interpretations of structural representations are possible, they need to specify and spell out in detail which understanding they favour. Second, they have to develop a notion that avoids the Scylla of being too liberal, and the Carybdis of mentioning certain properties that do the important explanatory work, which, in addition, might fail to constitute a scientific kind.

5. Conclusion

Time to wrap up. In recent years, the notion of structural representation has increasingly been used to identify a set of mechanisms that are supposed to play a distinctive role associated with paradigmatic instances of representations. The main aim of this article is to contribute to a better understanding of this concept and to assess the theoretical role that it can play. First, I argued there are significantly different ways of understanding this notion and that failure to see this plurality of interpretations may have contributed to certain misunderstandings. In this regard, more clarity would help the debate move forward. Second, while I think structural correspondences are important and I discussed different ways that their exploitation gives rise to powerful representational mechanisms, I argued that it is not obvious any interpretation of the concept of structural representation can support structuralism, that is, the view that structural representations constitute a distinctive and paradigmatically representational kind in virtue of the fact that a structural correspondence is exploited to accomplish a certain task.

If the arguments are on the right track, there are broadly two ways to go. One the one hand, one could accept an analysis along the lines of DEF1–DEF2 and, consequently, include detectors within the category of structural representations. Among other things, that would require revising the idea that structural representations are paradigmatically representational: either one rejects this claim, or one also includes detectors within the category of the paradigmatically representational. Likewise, one would need to find a different way to address the trivialization argument raised by anti-representationalists.

An alternative strategy is to embrace a more sophisticated analysis along the lines of DEF3–DEF5. In this case, the challenges are to vindicate the idea that structure exploitation (rather than, say, offline use) accounts for the explanatory role associated to representations and also to justify (or renounce) the idea that the concept of structural representation captures a scientific kind. Unless these challenges are met, structuralism should be regarded with suspicion.

Acknowledgements

I would like to thank Dimitri Coelho Mollo, Mariela Destéfano, Sérgio Farias de Souza Filho, Javier González de Prado Salas, Marta Jorba, Manolo Martínez, Gömöri Márton, Marcin Miłkowski, Jesús Navarro, Jake Quilty-Dunn, Miguel Ángel Sebastián, Peter Schulte, Nicholas Shea, Joulia Smortchkova, Laszlo Szabo, and Agustín Vicente. I am specially indebted to Nicholas Shea for long, deep, and extremely helpful discussions during a three-month research stay at the University of Oxford. Earlier versions of this article were presented at IX Conference of the Spanish Society of Analytic Philosophy at the University of Valencia (2019), at the workshop 'Information and Representation' at the University of Barcelona (2020), at the Logic and Philosophy of Science Seminar at the Eötvös Loránd University (2020) and at the Workshop on Representation in Cognitive Science at the University of Barcelona (2022).

Financial support was provided by the MICINN projects 'Varieties of Information' (PGC2018-101425-B-I00), 'The Representational Penumbra' (PID2021-127046NA-I00) and EFE-VMN (PID2019-106420GA-I00) and the project CIGE/2021/160 awared by the Generalitat Valenciana.

Marc Artiga
Department de Filosofia

Universitat de València
València, Spain
marc.artiga@uv.es

References

Abeles, M. [1982]: ‘Role of the Cortical Neuron: Integrator or Coincidence Detector?’, *Israel Journal of Medical Sciences*, **18**, pp. 83–92.

Arnellos, A. and Moreno, A. [2021]: ‘Visual Perception and the Emergence of Minimal Representation’, *Frontiers in Psychology*, **12**, available at <doi.org/10.3389/fpsyg.2021.660807>.

Artiga, M. [2016]: ‘Liberal Representationalism: A Deflationist Defense’, *Dialectica*, **70**, pp. 407–30.

Bartels, A. [2006]: ‘Defending the Structural Concept of Representation’, *Theoria*, **21**, pp. 7–19.

Bechtel, W. [2009]: ‘Generalization and Discovery by Assuming Conserved Mechanisms: Cross-Species Research on Circadian Oscillators’, *Philosophy of Science*, **76**, pp. 762–73.

Bechtel, W. [2016]: ‘Investigating Neural Representations: The Tale of Place Cells’, *Synthese*, **193**, pp. 1287–321.

Beer, R. D. [1995]: ‘A Dynamical Systems Perspective on Agent–Environment Interaction’, *Artificial Intelligence*, **72**, pp. 173–215.

Bitterman, M. E., Menzel, R., Fietz, A. and Schäfer, S. [1983]: ‘Classical Conditioning of Proboscis Extension in Honeybees (*Apis mellifera*)’, *Journal of Comparative Psychology*, **97**, pp. 107–19.

Butlin, P. [2020]: ‘Representation and the Active Consumer’, *Synthese*, **197**, pp. 4533–50.

Cartwright, B. A. and Collett, T. S. [1983]: ‘Landmark Learning in Bees’, *Journal of Comparative Physiology*, **151**, pp. 521–43.

Chemero, A. [2009]: *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press.

Churchland, P. M. [2012]: *Plato’s Camera: How the Physical Brain Captures a Landscape of Abstract Universals*, Cambridge, MA: MIT Press.

Clark, A. [1996]: *Being There: Putting Brain, Body, and World Together Again*, Cambridge, MA: MIT Press.

Clark, A. and Toribio, J. [1994]: ‘Doing without Representing?’, *Synthese*, **101**, pp. 401–31.

Colgin, L. L. [2020]: ‘Five Decades of Hippocampal Place Cells and EEG Rhythms in Behaving Rats’, *Journal of Neuroscience*, **40**, pp. 54–60.

Collett, T. S. and Baron, J. [1994]: ‘Biological Compasses and the Coordinate Frame of Landmark Memories in Honeybees’, *Nature*, **368**, pp. 137–40.

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

Dawkins, M. S. and Woodington, A. [2000]: ‘Pattern Recognition and Active Vision in Chickens’, *Nature*, **403**, pp. 652–55.

Diba, K. and Buzsáki, G. [2007]: 'Forward and Reverse Hippocampal Place-Cell Sequences during Ripples', *Nature Neuroscience*, **10**, pp. 1241–42.

Dickinson, A. [1980]: *Contemporary Animal Learning Theory*, Cambridge: Cambridge University Press.

Doussot, C., Bertrand, O. J. N. and Egelhaaf, M. [2020]: 'Visually Guided Homing of Bumblebees in Ambiguous Situations: A Behavioural and Modelling Study', *PLOS Computational Biology*, **16**, available at <doi.org/10.1371/journal.pcbi.1008272>.

Dragoi, G. and Tonegawa, S. [2011]: 'Preplay of Future Place Cell Sequences by Hippocampal Cellular Assemblies', *Nature*, **469**, pp. 397–401.

Drieu, C. and Zugaro, M. [2019]: 'Hippocampal Sequences during Exploration: Mechanisms and Functions', *Frontiers in Cellular Neuroscience*, **13**, p. 232.

Durbin, R. and Mitchison, G. [1990]: 'A Dimension Reduction Framework for Understanding Cortical Maps', *Nature*, **343**, pp. 644–47.

Egelhaaf, M., Kern, R. and Lindemann, J. P. [2014]: 'Motion as a Source of Environmental Information: A Fresh View on Biological Motion Computation by Insect Brains', *Frontiers in Neural Circuits*, **8**, available at <doi.org/10.3389/fncir.2014.00127>.

Engel, A. and Singer, W. [2001]: 'Temporal Binding and the Neural Correlates of Awareness', *Trends in Cognitive Sciences*, **5**, pp. 16–25.

Escobar, M. and Miller, R. [2004]: 'A Review of the Empirical Laws of Basic Learning in Pavlovian Conditioning', *International Journal of Comparative Psychology*, **17**, pp. 279–303.

Facchin, M. [2021]: 'Structural Representations Do Not Meet the Job Description Challenge', *Synthese*, **199**, pp. 5479–508.

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

Ganson, T. [2020]: 'A Role for Representations in Inflexible Behavior', *Biology and Philosophy*, **35**, pp. 1–18.

Garzón, F. C. and Rodríguez, F. C. [2009]: 'Where Is Cognitive Science Heading?', *Minds and Machines*, **19**, pp. 301–18.

Gładziejewski, P. [2015]: 'Explaining Cognitive Phenomena with Internal Representations: A Mechanistic Perspective', *Studies in Logic, Grammar and Rhetoric*, **40**, pp. 63–90.

Gładziejewski, P. and Miłkowski, M. [2017]: 'Structural Representations: Causally Relevant and Different from Detectors', *Biology and Philosophy*, **32**, pp. 337–55.

Goodman, N. [1976]: *Languages of Art: An Approach to a Theory of Symbols*, Indianapolis, IN: Hackett.

Griffiths, P. E. [1997]: *What Emotions Really Are: The Problem of Psychological Categories*, Chicago: University of Chicago Press.

Hartley, T., Lever, C., Burgess, N. and O'Keefe, J. [2014]: 'Space in the Brain: How the Hippocampal Formation Supports Spatial Cognition', *Philosophical Transactions of the Royal Society of London B*, **369**, available at <doi.org/10.1098/rstb.2012.0510>.

Hutto, D. D. and Myin, E. [2013]: *Radicalizing Enactivism: Basic Minds without Content*, Cambridge, MA: MIT Press.

Hutto, D. D. and Myin, E. [2017]: *Evolving Enactivism: Basic Minds Meet Content*, Cambridge, MA: MIT Press.

König, P., Engel, A. K. and Singer, W. [1996]: ‘Integrator or Coincidence Detector? The Role of the Cortical Neuron Revisited’, *Trends in Neurosciences*, **19**, pp. 130–37.

Kreiter, A. K. and Singer, W. [1996]: ‘Stimulus-Dependent Synchronization of Neuronal Responses in the Visual Cortex of the Awake Macaque Monkey’, *Journal of Neuroscience*, **16**, pp. 2381–96.

Lee, J. [2019]: ‘Structural Representation and the Two Problems of Content’, *Mind and Language*, **34**, pp. 606–26.

Leonard, B. and McNaughton, B. L. [1990]: ‘Spatial Representation in the Rat: Conceptual, Behavioral, and Neurophysiological Perspectives’, in *Neurobiology of Comparative Cognition, Comparative Cognition, and Neuroscience*, Hillsdale, NJ: Lawrence Erlbaum, pp. 363–422.

Marsat, G. and Pollack, G. S. [2012]: ‘Bursting Neurons and Ultrasound Avoidance in Crickets’, *Frontiers in Neuroscience*, **6**, available <doi.org/10.3389/fnins.2012.00095>.

Martínez, M. and Artiga, M. [2023]: ‘Neural Oscillations as Representations’, *British Journal for the Philosophy of Science*, **74**, pp. 619–48.

Menzel, R. and Müller, U. [1996]: ‘Learning and Memory in Honeybees: From Behavior to Neural Substrates’, *Annual Review of Neuroscience*, **19**, pp. 379–404.

Miller, L. A. and Surlykke, A. [2001]: ‘How Some Insects Detect and Avoid Being Eaten by Bats: Tactics and Countertactics of Prey and Predator’, *BioScience*, **51**, pp. 570–81.

Millikan, R. G. [1984]: *Language, Thought, and Other Biological Categories*, Cambridge, MA: MIT Press.

Müller, M. and Wehner, R. [2010]: ‘Path Integration Provides a Scaffold for Landmark Learning in Desert Ants’, *Current Biology*, **20**, pp. 1368–71.

Morgan, A. [2014]: ‘Representations Gone Mental’, *Synthese*, **191**, pp. 213–44.

Nakamura, N. H., Fukunaga, M., Akama, K. T., Soga, T., Ogawa, S. and Pavlides, C. [2010]: ‘Hippocampal Cells Encode Places by Forming Small Anatomical Clusters’, *Neuroscience*, **166**, pp. 994–1007.

Neander, K. [2017]: *A Mark of the Mental: A Defence of Informational Teleosemantics*, Cambridge, MA: MIT Press.

Nirshberg, G. and Shapiro, L. [2020]: ‘Structural and Indicator Representations: A Difference in Degree, Not Kind’, *Synthese*, **198**, pp. 7647–64.

O’Keefe, J. and Dostrovsky, J. [1971]: ‘The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely-Moving Rat’, *Brain Research*, **34**, pp. 171–75.

Opie, J. and O’Brien, G. [2004]: ‘Notes toward a Structuralist Theory of Mental Representation’, in H. Clapin, P. Staines and P. Slezak (*eds*), *Representation in Mind: New Approaches* *to Mental Representation*, Amsterdam: Elsevier, pp. 1–20.

Piccinini, G. [2018]: 'Computation and Representation in Cognitive Neuroscience', *Minds and Machines*, **28**, pp. 1–6.

Piccinini, G. [2020]: *Neurocognitive Mechanisms: Explaining Biological Cognition*, Oxford: Oxford University Press.

Plebe, A. and De La Cruz, V. M. [2018]: 'Neural Representations beyond "Plus X"', *Minds and Machines*, **28**, pp. 93–117.

Ramsey, F. [1931]: *The Foundations of Mathematics*, London: Kegan Paul.

Ramsey, W. [2016]: 'Untangling Two Questions about Mental Representation', *New Ideas in Psychology*, **40**, pp. 3–12.

Ramsey, W. M. [2007]: *Representation Reconsidered*, Cambridge: Cambridge University Press.

Redish, A. D., Battaglia, F. P., Chawla, M. K., Ekstrom, A. D., Gerrard, J. L., Lipa, P., Rosenzweig, E. S., Worley, P. F., Guzowski, J. F., McNaughton, B. L. and Barnes, C. A. [2001]: 'Independence of Firing Correlates of Anatomically Proximate Hippocampal Pyramidal Cells', *The Journal of Neuroscience*, **21**, available at <doi.org/10.1523/jneurosci.21-05-j0004.2001>.

Rupert, R. D. [2018]: 'Representation and Mental Representation', *Philosophical Explorations*, **21**, pp. 204–25.

Schulte, P. [2015]: 'Perceptual Representations: A Teleosemantic Answer to the Breadth-of-Application Problem', *Biology and Philosophy*, **30**, pp. 119–36.

Schuster, S. and Amtsfeld, S. [2002]: 'Template-Matching Describes Visual Pattern-Recognition Tasks in the Weakly Electric Fish *Gnathonemus petersii*', *The Journal of Experimental Biology*, **205**, pp. 549–57.

Segundo-Ortin, M. and Hutto, D. D. [2021]: 'Similarity-Based Cognition: Radical Enactivism Meets Cognitive Neuroscience', *Synthese*, **198**, pp. 5–23.

Shagrir, O. [2012]: 'Structural Representations and the Brain', *British Journal for the Philosophy of Science*, **63**, pp. 519–45.

Shea, N. [2014]: 'Exploitable Isomorphism and Structural Representation', *Proceedings of the Aristotelian Society*, **114**, pp. 123–44.

Shea, N. [2015]: 'Distinguishing Top-Down from Bottom-Up Effects', in D. Stokes, M. Matthen and S. Biggs (*eds*), *Perception and Its Modalities*, Oxford: Oxford University Press, pp. 73–91.

Shea, N. [2018]: *Representation in Cognitive Science*, Oxford: Oxford University Press.

Shea, N., Godfrey-Smith, P. and Cao, R. [2018]: 'Content in Simple Signalling Systems', *British Journal for the Philosophy of Science*, **69**, pp. 1009–35.

Shettleworth, S. J. [2010]: *Cognition, Evolution, and Behavior*, Oxford: Oxford University Press.

Singer, W. and Gray, C. M. [1995]: 'Visual Feature Integration and the Temporal Correlation

Hypothesis', *Annual Review of Neuroscience*, **18**, pp. 555–86.

Stegmann, U. E. [2009]: 'A Consumer-Based Teleosemantics for Animal Signals', *Philosophy of Science*, **76**, pp. 864–75.

Stone, T., Webb, B., Adden, A., Weddig, N. B., Honkanen, A., Templin, R., Wcislo, W., Scimeca, L., Warrant, E. and Heinze, S. [2017]: 'An Anatomically Constrained Model for Path Integration in the Bee Brain', *Current Biology*, **27**, pp. 3069–85.

Swoyer, C. [1991]: 'Structural Representation and Surrogative Reasoning', *Synthese*, **87**, pp. 449–508.

Takeda, K. [1961]: 'Classical Conditioned Response in the Honey Bee', *Journal of Insect Physiology*, **6**, pp. 168–79.

Van Gelder, T. [1995]: 'What Might Cognition Be, If Not Computation?', *The Journal of Philosophy*, **92**, pp. 345–81.

von der Malsburg, C. and Schneider, W. [1986]: 'A Neural Cocktail-Party Processor', *Biological Cybernetics*, **54**, pp. 29–40.

Williams, D. and Colling, L. [2018]: 'From Symbols to Icons: The Return of Resemblance in the Cognitive Neuroscience Revolution', *Synthese*, **195**, pp. 1941–67.