<u>ACCEPTED IN *ERKENNTNIS* - PLEASE REFER TO THE PUBLISHED VERSION</u>

Maps, simulations, spaces and dynamics: on distinguishing types of structural representations

Marco Facchin, Ph.D. [corresponding author]
Research Fellow
Antwerp University
Centre for Philosophical Psychology
Prinstraat 13, Antwerp 2000
marco.facchin.marco.facchin@gmail.com
https://orcid.org/0000-0001-5753-987

Abstract:

Structural representations are likely the most talked about representational posits in the contemporary debate over cognitive representations. Indeed, the debate surrounding them is so vast virtually every claim about them has been made. Some, for instance, claimed structural representations are different from indicators. Others argued they are the same. Some claimed structural representations mesh perfectly with mechanistic explanations, others argued they can't in principle mash. Some claimed structural representations are central to predictive processing accounts of cognition, others rebuked predictive processing networks are blissfully structural representation free. And so forth. Here, I suggest this confusing state of affairs is due to the fact that the term "structural representations" is applied to a number of distinct conceptions of representations. In this paper, I distinguish four such conceptions, argue that these four conceptions are actually distinct, and then show that such a fourfold distinction can be used to clarify some of the most pressing questions concerning structural representations and their role in cognitive theorizing, making these questions more easily answerable.

Keywords: Structural representations; Indicator representations; Structural similarity; Simulations

Declarations:

Founding: This research was founded by the FWO grant "Towards a globally non-representational theory of the mind" (grant number 1202824N)

Conflict of Interests/competing interests: The author declares no conflict of interests.

Availability of data and Material: Not applicable.

Code availability: Not applicable.

Authors' contribution: Marco Facchin is the sole author of the paper.

Acknowledgements:

Many thanks to Peter Schulte for having made the observation that led me to write this paper. Thanks also to the anonymous referees of the journal *Synthese* who persuaded me that my article titled *Neural Representations Unobserved* was getting too long: cutting it down finally gave me the chance to write this paper . Also, many thanks to the anonymous referees of the journal *Erkenntnis* for their insightful feedback - the paper has really improved since the first, lazily written draft. Thanks to (in random order) Matej Kohar, Beate Krickel, Nir Fresco, Marcin Miłkowski and the other participants of the MeREx workshop in Berlin for the engaging discussion.

1 - Introduction

The field of cognitive representations - that is, the sub-personal explanation working as *explanantia* in cognitive scientific theories - blooms with structural representations. But what are these? My younger self would have promptly responded that "they are representational vehicles that *stand-in-for* their targets *by reproducing* the targets' inner structure - so that each resembles the target it represents in its inner relational structure. Go read Gładziejewski (2015, 2016)."

As it often happens, I now look back at my younger self and cringe at his naive answer. The debate over structural representations is *way* more articulate than that simplistic answer indicates. Consider, for example, what sorts of cognitive architectures host structural representations. "Classical" rule-and-representations-based architectures surely do - as Ramsey (2007) and Cummins (1989) persuasively argued... unless "classical" architectures represent symbolically; that is, without using structural representations at all (cf. Williams & Colling 2017). Connectionists and post-connectionist architectures rely on structural representations too (Cummins 1989; Churchland 2012)... unless such systems are simply "wired up" to respond to specific stimuli in specific ways (cf. Ramsey 2007, Orlandi 2014). Structural representations can be observed taking up a mechanistic perspective on representational systems (Gładziejewski 2015)... but also by taking up a dynamical perspective on them (cf. Shagrir 2012); even if, in the eyes of many, such a dynamical perspective should be the *opposite* of a mechanistic one (cf. Silberstein & Chemero 2013) and should more or less naturally lead to a form of anti-representationalism (Van Gelder 1995; Chemero 2009). Structural representations are the linchpin of the cognitive neuroscience revolution: an ambitious view that aims to unify cognitive (neuro)science under the banner of multilevel mechanistic explanations (Boone & Piccinini 2016; Williams & Colling 2018)... even if structural representations and mechanistic explanation may fail to coherently merge (Kohar 2023). Structural representations have apparently been observed in our brains (Thompson & Piccinini 2018; Piccinini 2020a); but that may have been just a trick of the light (Facchin 2024). Structural representations are central in "predictive processing" accounts of cognition (Gładziejewski 2016; Kiefer & Hohwy 2018), except that perhaps such accounts actually rely on no inner representations at all (Downey 2018; Facchin 2021a). Structural representations, however, are definitely distinct from indicator representations (Ramsey 2007; Gładziejewski & Miłkowski 2017) - unless the two are one and the same (Morgan 2014, Nirshberg & Shapiro 2021, Facchin 2021b).

Structural representations are thus at the center of a dense web of philosophical debates. And a close analysis of these debates reveals that the term "structural representation" is actually used to refer to a variety of different posits, which have not yet been sufficiently teased apart.[1] As a consequence, sometimes remarkably different conceptions of structural representations are used

[1] As this paper was undergoing its first round of review, the situation changed somewhat. Facchin (2024, *appendix*) noticed that two different notions of structural representations are used interchangeably in the literature - and indeed, if I interpret him correctly, he basically distinguishes, in a very rough manner, what in **§2** I will dub "structural maps" and "structural spaces". Artiga (2023) too proposed a fivefold distinction of structural representations. However, unlike the distinction sketched by Facchin, his distinctions seem to consist in the progressive "fine graining" of what in **§2** I will call "structural spaces". So, if I understand him correctly, whereas my distinctions operate "horizontally", telling apart various different types of structural representations, Artiga's distinctions operate "vertically", within a single kind. Be as it may, it seems clear that there is a need to distinguish various forms of structural representations, and my paper offers an answer to that need.

within the same debate by different authors, in a way that prevents them from reaching a consensus and making progress in the issues they tackle. Hence the rather minimal aim of this paper: I wish to point out that there is an unnoticed four-fold distinction to be made in the current literature on structural representations, and that making such a distinction is beneficial to a number of debates.

Here's how I will proceed. **§2** surveys the current literature to disentangle four different conceptions of structural representations. **§3** argues that these four different conceptions are actually distinct, and do not reduce to each other. **§4** puts these distinctions to use showing how they can be fruitfully applied in some recent debates over structural representations. **§5** concludes the paper.

2 - Four types of structural representations

I distinguish four different conceptions of structural representations, presenting each in a dedicated subsection. My aim is *descriptive* and *clarificatory*: I want to show that, in the present literature, the phrase "structural representation" has been used to name - and think about - very different representational posits.

Does my fourfold distinction I'm making provide a *complete* enumeration of structural representations types? Maybe not - I can't exclude that there's a fifth (or sixth, or seventh, and so forth) type of structural representation. But whether my enumeration is complete or not is irrelevant here. My aim is not to chart *all possible* types of structural representation that may exist. I aim to make order in the current literature, and my distinction does it well enough. Does my fourfold distinction identify *genuinely* different types of structural representations? Yes - as I will argue in **§4**. Does my distinction identify metaphysically different sub-kinds[2] of structural representations? I think so, even if no actual argument backs this assertion up. However, whether the distinction I'm introducing is a distinction *between sub-kinds* is irrelevant here. Not all relevant and useful distinctions need to be distinctions in kind!

Some introductory words on structural representations. In general, philosophers think of representations as concrete items, called *vehicles*, standing-in-for other things, which I will call *targets*. This definition is purposefully liberal, and it only needs two restrictions. One concerns what may count as a vehicle: vehicles must be *concrete particulars* that are *tokenizable* in the system within which they function as representations. The other concerns the fact that, in the case of structural representations, a vehicle (or a system of vehicles) must be *structurally similar* to the target (or target system). Sure, current definitions of structural representations are typically more demanding than this (cf. Gładziejewski 2015; 2016, Pezzulo 2008). But these definitions are *not* universally accepted, and they would prevent me from identifying current conceptions of structural representations *as conceptions of structural representations*. So, these more demanding conceptions are here best ignored.

What's the *structural similarity* that I mentioned above? Simplifying to the extreme by "structural similarity" I have in mind a relation of one-to-one correspondence between (at least) some of the

[2] I owe this observation to Marcin Miłkowski (private communication), who correctly observed that the four *genera* of structural representations I will identify here are sub-kinds of a single superordinate kind "structural representation" identified by at least some necessary conditions, such as, for example, the presence of some relevant structural similarity.

elements of two distinct domains, such that a same *abstract pattern* of relations holds between the elements of the two domains involved in the correspondence.[3] In the case at hand, the two domains are representational vehicles on the one hand and targets on the other,

Following a common practice in the literature (e.g. Ramsey 2007, Godfrey-Smith 2009, Gładziejewski 2015; 2016, Facchin 2021b) I will also take *cognitive* representations as inner analogs of *public* representations - and since we are all familiar with the latter, I will extensively use them for clarificatory purposes.

Lastly, following many (e.g. Ryder 2009; Artiga 2016; Neander 2017; Schulte 2023) I will distinguish two different questions an account of representation may be attempting to answer. An account of representations may answer what I call the *content question*: why does this type of representation represent *this* target rather than *that* target? An account of representation may also attempt to answer what I will call the *status question*: why does a state or item have content at all and qualifies as a representational vehicle, rather than having no content at all and remaining a simple non-representational state or item? Notice how the status question asks why something is a representation in the first place, whereas the content question *assumes* that something is a representation to then inquire into its content. The two questions are importantly connected, and answers to the status question must logically precede answers to the content question - after all, in order for a vehicle to represent *this* target rather than *that other* target it seems that the vehicle must be a representation in the first place, and so answer to the status question shape and constrain the space of possible answers to the content question. Still, despite this tight link, the two questions are conceptually distinct, and it is possible for their answer to diverge: it is at least in principle possible that the set of factors in virtue of which an item or a state is a representational vehicle are distinct from the set of factors in virtue of which that vehicle has the specific target it has.

One last word on the status question. Following Ramsey (2007; 2016; 2023) this question is at times rephrased in functional terms, assuming that representations form a functional kind and asking what sort of (long-armed) functional role[4] must a state or item play to qualify as a representation within the system in which it is tokened. Now, whilst this understanding of the status question is extremely popular, I don't want to assume from the onset that the status question must be answered functionally. It is, however, important to notice that it is increasingly receiving such an answer. For this reason, I will take claims concerning the functional role of representations as answers to the status question.

Time, now, to introduce my fortfold distinction.

[3] See O'Brien and Opie (2004) for a more rigorous definition. See also Kohar (2023, ch. 4) for a convincing argument to the effect that this is the *only* relevant definition of structural similarity in the current debate.

[4] As a reviewer rightfully noticed, taking this functionalist reading of the status question presupposes a long-armed (externalist, world-involving) reading of functional roles (on this point, see also Shea 2018, pp. 31-36). It would be indeed very hard to specify a representational functional role in terms of short-armed, internalistic terms!

2.1 - Structural maps

According to a very prominent view, structural representations are individual map-like vehicles whose inner relational structure mirrors the inner relational structure of the target they represent. This is what I'll call the *structural map* conception of structural representations.

Unsurprisingly, structural maps are often characterized by in reference to four features of ordinary cartographic maps (cf. Gładziejewski 2015; 2016). Consider, for example, a map of Europe. **(1)** The map is a single object - a single *representational vehicle* - depicting Europe. The vehicle itself is made of various *constituents* - say, various blobs of color that stand for the various countries, blue blobs representing seas, oceans and lakes, dots representing various important cities and so forth - and these constituents stand in a variety of spatial relations with each other. The represented target itself is a single object - a single *represented target* - with various constituents - the various countries, cities, seas etc. - which stand in a variety of spatial relation to each other. Now, in a well-formed, functioning map, there is a *one-to-one* correspondence between the constituents of the vehicle and the constituents of the target. Each blue blob (standing for a sea or a lake) *corresponds* only to one sea or lake, each dot *corresponds* only to one city, and so forth. Moreover, this one-to-one correspondence preserves *the same pattern of relation* on both sides of the correspondence. If two vehicle constituents a and b are in a relation R, then the corresponding target constituents a* and b* are in a corresponding relation R^*.[5] Thus, the map is **(1) structurally similar** to its representational target. **(2)** That structural similarity, moreover, is *essential* to the map's functioning. The better the map resembles the target, the better of a representation it will be.[6] If the map of Europe closely resembles Europe, then it can be used to reliably draw inferences about Europe, and to plan one's road through Europe. Conversely, if the map does not resemble Europe, it won't allow us to reliably draw inferences or plan one's road. Thus, the structural similarity must **(2) be exploitable** by the user of the representation.[7] Maps can also **(3) be decoupled** from their targets - they can be used to represent it even when no causal connection ties them together. Lastly, maps allow for **(4) error detection**. If a map reliably leads us astray, we conclude that it is a wrong map that misrepresents the terrain - and find a different means to navigate it. The structural map conception conceives cognitive structural representations (i.e. structural maps) as single vehicles satisfying **(1)**-**(4)**.

Who defends this view? Whilst Cummins (1996) immediately springs to mind as a precursor it is Gładziejewski (2015; 2016) who explicitly defined structural maps in terms of **(1)**-**(4)**. Since the literature on predictive processing (cf Williams 2017; 2018a,b; Kiefer and Hohwy 2018; 2019; Wiese2017, 2018; Facchin 2021a) and the "cognitive neuroscience revolution" (Williams & Colling 2017; Piccinini 2020a,b; Thomson and Piccinini 2018; Lee 2019, 2021) adopt Gładziejewski's

[5] One unfortunate fact about the cartographic map example is that the pattern of relations preserved on both sides of the mapping holds amongst relations of the same kind (that is, *spatial* relations). But that is by no means necessary. In oceanographic maps, for example, various *depths* map onto a gradient of *colors*. And so, while the relations holding amongst the constituents are different (i.e. "deeper than" for the target constituents and "darker than" for the vehicle constituents) the *pattern* of these relations remains identical.

[6] Given, of course, certain limits of utilizability and intelligibility: a useful map need not be (and indeed *can't* be) a 1:1 replica of the mapped territory.

[7] For a less casual definition of exploitability, see (Shea 2018; Gładziejewski & Miłkowski 2017)

conception of structural representations, they appear to be committed to structural maps too.[8] Structural maps are thus quite popular. But, contrary to what my younger self believed, they're not the only type of structural representation around.

What does the structural map account offer? Does it answer the content question or the status question? It is not hard to answer this question: Gładziejewski (2015; 2016) explicitly presents structural maps as a functional, Ramsey (2007) inspired, answer to the status question. Thus, **(1)**-**(4)** should be conceived as establishing what "makes" something into a representation by providing a functional profile detailing how structural maps *function as representations* within the systems in which they are tokenized.[9] Followers of Gładziejewski from both the "cognitive neuroscience revolution" and the predictive processing movement agree on this point. Thus, for example, both Piccinini (2020a,b) and Wiese (2017; 2018) try to supplement structural maps with an independent theory of content so as to enable them to face the content question too. Of course, such an independent account might not be necessary, as the structural map conception might also implicitly offer an account of content (Lee 2019). Be as it might, the fact that *if* structural maps answer the content question they do so *implicitly* indicates that structural maps are *explicitly* thought of as an answer to a different question, namely the status question.[10]

To recap: structural maps are individual vehicles that represent their targets by having features **(1)**-**(4)**, where such features are typically intended not as an account of what gives the vehicle its content, but rather as an account providing a (functional) specification of what in virtue of which certain states *are* representations in the first place.

Not all occurrences of "structural representations" refer to structural maps, however.

2.2- Structural simulations

Structural maps provide an answer to the status question. Such an answer is inspired by Ramsey's (2007) analysis of theories of representations, and the functionalist reading of the status question Ramsey suggested. And indeed, Ramsey himself argued that structural representations offer an answer to the status question (functionally interpreted). Yet, when he claimed so, he was not talking about structural maps. Rather, he was talking about what I shall call *structural simulations*. The name honors Cummins's "simulation representation" (1989) account, which, on Ramsey's (2007) view, is the original notion of structural representation.

---

[8] Even if perhaps this literature should start to deploy a different conception, see **§4.3** below.

[9] One may wonder how these features make up a *functional* profile, given that presumably resemblance relations are not causal/functional, and thus that **(1)** can't plausibly be taken to spell out a functional feature. It's worth noticing, however, that according to Gładziejewski structural similarities *have* a causal role, given an interventionist understanding of causality (see Gładziejewski & Miłkowski 2017). Here, I'll concede this point for the sake of discussion.

[10] Cummins (1996) is the only clear cut exception to this general trend, as he proposes structural maps as an account of content. Kiefer and Hohwy (2018; see also Kiefer 2023) *may* also consider structural maps as providing an answer to the content question. Yet their position on this matter is unclear, as they *also* seem to endorse a form of functional role semantics, which arguably fails to merge with structural maps, regardless of whether they are taken to answer the status or the content questions (see Facchin 2021c, pp. 295-297 for discussion of this point).

Cummins (1989) conceived of representational vehicles as inner states of computational systems. These computational states bear certain computational relations to each other, which specify computational state transitions. The general form of these computational state transitions is $o = f(i,s)$, meaning that, when coupled to an input *i*, an inner computational state *s* yields a second computational state *o* as outcome; where *i* may be an inner computational state in its own right (e.g. the computational state my compute transition when I press the button "r"). Importantly, these computational states can be put into a one-to-one correspondence to external targets in a way such that their computational state transitions mirror certain relations (typically state transitions) holding amongst the targets. When this happens, the computational process involving vehicles *i*, *s* and *o* mirrors - that is, is structurally similar to - the relational structure holding amongst the targets. Hence the former can be used to *simulate* the latter, and can thus be called a *structural representation* of them. For example, if *i* maps onto "the merch I sold", *s* maps onto "the merch I have" and *o* maps onto "the merch after the selling", and *f* is roughly the subtraction function, then the computational state transition $o = f(i,s)$ can used to *simulate* the storage state of my boutique as I sell items, thereby coming to *structurally represent* it.

Who defends structural simulations? Quite obviously, Cummins (1989) and Ramsey (2007).[11] But also Shagrir (2012, 2018), who holds that various neurocomputational models involve structural simulations. Grush (2004) may be read as making a similar claim, though with the scope narrowed down to a single class of neurocomputational models. And, in spite of her deflationary attitude towards representations, Egan (2014; 2020) can be rightfully said to this position too - even if she imposes some very rigid restrictions on the relevant vehicle-target correspondences.[12] Following Egan, Wiese (2017; 2018) Wiese accepts it too - perhaps surprisingly, given that he also endorses a *structural map* view of structural representations.

Is Wiese's position coherent? That's a bit hard to say. Surely Wiese (2017; 2018) *thinks* it is coherent, as he takes the structural simulation account to provide an answer to the content question, to be merged with a structural map account of representational status. Egan (2014) also seemingly takes the structural simulation account as an account of representational content. However, she recently changed her mind, taking it as a functional account of representational status (Egan 2020). Ramsey (2007) agrees. Cummuns (1989), moreover, took it as an account of both the *status and content* of representations. Needless to say, if structural simulations were answering (either also or only) the status question, then Wiese's sophisticated conception of structural representations would be in trouble. For one thing, Wiese's account would suddenly lack an answer to the content question. For another thing, it is not clear that structural maps and structural simulations can coherently merge in a single account of representational status. If a structural simulation account of function is accepted as an account of representational status, there seems to be no intelligible reason as to why the relevant computational states must *also* have the features identified by the structural map account in order to qualify as representations (and *vice*

[11] Ramsey (2007) seems also to dissociate structural simulation to what he calls "input-output" representations, that is, representations of the inputs and outputs of a computational process. As clarified by Sprevak (2011), however, the distinction is spurious: structural simulations *are* "input-output" representations.

[12] Basically, in Egan's view, the only actual vehicle-target correspondences doing *real* representational works map the vehicle on some argument or value of the function computed by the device, so that the inner dynamics of the computational mechanism *simulates* the function computed. For a critical analysis of this view, see (Facchin 2022).

versa). Moreover, a combined account of representational status according to which structural representations must satisfy both the structural map *and* the structural simulation conception risks being unduly restrictive, failing to recognize genuine representational posits *as such*. I won't, however, try to ascertain the tenability of such an account any further here. I merely want to flag out that, whereas structural maps are (almost) unanimously thought of as providing an answer to the status question, structural simulation accounts have a less clear status. It is not clear whether they provide an account of representational content, an account of representational status, or an account answering both questions at once.

2.3 - Structural spaces

Structural representations are often presented as radically distinct from indicators (cf. Ramsey 2007; Cummins 1996; Gładziejewski & Miłkowski 2017). This makes intuitive sense: indicators are tied to their representational targets by causal/informational links rather than any structural similarity. Moreover, indicator states are often compared to the representational states generated by measurement tools such as thermometers, or barometers, rather than public representations such as maps. Yet, a small, but noisy, group of philosophers (Morgan 2014; Nirshberg & Shapiro 2021, Facchin 2021b) have argued that indicators actually qualify as structural representations. Actually, looking closely at the literature, there are *two* distinct senses in which indicators can be said to be structural representations. Here, I will examine one, which I dub the *structural space* conception.

The structural space conception applies only to graded, non-categorical indicators; that is, indicators that can tokenize a range of indicator states in response to a variety of targets, most typically the amount or magnitude of what they indicate. Think, for example, about a thermometer indicating a range of different temperatures - roughly, the *amount* of mean kinetic energy in a determined region of space. In the case of graded indicators there is always an indicator-specific relation (in the case of thermometers, the *higher than* relation) and an indicated target-specific relation (in the case of thermometers, the *hotter then* relation[13]) that impose a strict total ordering amongst the indicator states and the indicated states respectively. And, indicator states maps onto indicated states in a way such that the strict total order is preserved on both sides of the mapping - basically for each pair of indicator states and indicated states, if indicator state A is in a the indicator-specific relation with indicator state B, then the state indicated by A is in the target-specific relation with the state indicated by B. As such, the state space (that is, the set of all *state types* of the indicator) is structurally similar to the state space of the target.[14] For some researchers, this is enough to turn an indicator into a structural representation of its target.

Who holds this view? (Morgan 2014; Nirshberg & Shapiro 2021, Facchin 2021b) of course, but also defenders of structural representations such as O'Brien and Opie (2004) and Churchland (2012). For, in both of their account, the relevant structural similarity holds between the state space of an entire artificial neural network and the targets the network was tasked to categorize; and artificial neural networks can be considered just as a large number of indicator like units, or even as a single complex indicator (Gosche & Koppelberg 1991; Ramsey 2007).

[13] Or "having more meak kinetic energy than" relation.

[14] At least, within the range of sensitivity of the indicator.

Should we understand the structural space conception of structural representations as providing an answer to the content question or the status question? Arguably, as an answer to the content question. After all, Nirshberg and Shapiro, O'Brian and Opie, Churchland, and O'Brien (2015) all deal with structural spaces when dealing with the *content* of representations. So, it seems charitable to interpret structural spaces as an answer to the content question. Structural spaces, however, could also read as providing an answer to the status question. In this view, an indicator state would (paradoxically) acquire the status of representation not because it *indicates* its target, but because it partakes in the structural similarity that connects the various states of the indicator and the target. The discussion of indicators in O'Brien and Opie (2004) and Williams and Colling (2017) *may* be read as advancing this point. I am, however, hesitant in asserting that that is the intended reading of both papers. In the case of O'Brien and Opie, whilst it is true that they distinguish (in footnote 9) theories of content from theories of *representations* - which can naturally be interpreted as answers to the content question and the status question respectively - their work seems to focus on finding an account of content allowing content to be causally efficacious. As such, they are most naturally read as providing an answer to the content question. Williams and Colling endorse a structural space-based answer to the function question, as they all explicitly endorse Gładziejewski's structural map account, which *already* provides an answer to the status question. So, at present, I just flag that whereas the structural space conception is often advanced as an answer to the content question, it may also be "recycled" as an answer to the status question.

2.4 - Structural dynamics

There is, however, also a second structural similarity that ties together indicators to their targets - and that structural similarity holds for *all* indicators, including categorical ones - that is, indicators having just two states, that simply indicate the presence/absence of a target (think, for example, at the infrared sensor of a garage door). This structural similarity, however, does not hold between the state space of the indicator and its target. Rather, it holds between the *dynamics* of the two. Roughly put, the idea is that the way in which indicator states are tokenized through time is structurally similar to the way in which indicated states are tokenized in time (cf. Facchin 2021b). For example, if the infrared sensor of a garage door tokenizes the state that registers the presence of something at time t, and then stops registering any presence from $t+1$ to $t+n$, to then re-tokenize the presence state at time $t+m$, then (on the assumption that the indicator is working correctly) a target is present at t, missing from $t+1$ to $t+n$, and present again at $t+m$. Thus, token indicator states map one-to-one onto some token target states, and their *temporal relations are the exact same* on both sides of the mapping. In this way, the way in which the state of the indicator *dynamically* changes, tracking the indicated target, comes to be structurally similar to the way in which the target dynamically evolves. For this reason, the former may be said to be a structural representation of the latter. Let me dub this the *structural dynamics* conception of structural representations.

Who holds this view? That's a tricky question to answer. Officially, no one, as the structural space and structural dynamics conception are often run together in the literature (see especially Morgan 2014; Facchin 2021b). However, Facchin (2021b) explicitly notes that the relata involved in the structural similarity underpinning structural spaces *differ* from the relata underpinning the structural similarity in structural dynamics cases. So, the structural dynamics view can be ascribed to him. Morgan (2014) also may hold it, as he too claims that binary indicators are structural representations of their targets. However, if Morgan holds the structural dynamics view, he does not state it explicitly. Lee and Calder (2023) may endorse this view too, at least insofar they attempt to answer Facchin's challenge to structural representation, which is based on the structural dynamics conception of structural representations.[15]

What does the structural dynamics view of structural representations offer? It seems safe to say that it is supposed to address the status question. Indeed, both Morgan (2014) and Facchin (2021b) resort to it to retort against certain claims in Ramsey (2007), claims that concern a specific (functional) interpretation of the status question. Thus, the structural dynamic conception seems naturally interpreted as explaining how indiators qualify as (structural) representations of their targets: they qualify as representations because the way in which they token their states through time "mirrors" the temporal evolution of a given target.

2.5 - A quick summary

I have surveyed the literature on structural representations, distinguishing four different conceptions of structural representation. The structural map conception provides an account of why certain items qualify as representations, by spelling out how they can *function as* (structural) representations within the systems in which they are tokenized. On this view, items qualify as representations because they bear four relevant functional features that make them function *as* maps.

The structural simulation conception is less clear in its aims, as it has been alternatively interpreted as an account of representational contents, an account of representational status, or both. Whatever its aim, however, the structural simulation conception sees (structural) representation as computational states, whose computational state transitions allow them to simulate a certain target.

The structural space conception applies exclusively to non-categorical indicators, and it is mainly intended to yield a theory of content for indicator states. On this view, indicator states carry the content they carry because the state space of the indicator structurally resembles the state space of the indicated target.

Lastly, the structural dynamics view seems to provide a functional account of the representational status of indicators, according to which indicators function as representations of their targets by tokening indicator states so as to mirror the temporal dynamics of the target.

My main claim here is simple: the four conceptions of structural representations I have distinguished above are indeed four *different* conceptions, and the fact that in the current literature they have not been sufficiently distinguished - that is, the fact that "structural representations" and similar terms are currently used to refer to any of these conception, without specifying which is the intended one - is preventing philosophers if not from reaching a consensus in a number of debates, at least from making genuine progress on them.

---

[15] Notice, however, how the fact that no one clearly and explicitly defends the structural dynamics conception backs up one of main claims of this paper - namely, that different sub-kinds of structural representations have been conflated in the literature.

So, I have two tasks ahead of me. First, I need to show that the four conceptions I identified are indeed *different* conceptions. I'll do so in **§3**. Then, I have to show that distinguishing them is indeed useful in the current philosophical landscape. I'll do so in **§4**.

*3 - Different types of structural representations are actually *different* types*

The claim that structural maps, structural simulations, structural spaces and structural dynamics are actually *distinct* notions of structural representation can be backed up by a number of different arguments.

A first argument points out that these conceptions have different *aims*. They're made to address different questions. As highlighted above, the structural map and the structural dynamics conception aim at tackling the representational status question from a functional point of view. The structural space conception, on the contrary, aims at tackling the content question. And the structural simulation conception seemingly aims at answering both types of questions. So, these four conceptions of structural representations have quite different aims.

The above suggests that they are actually *different*. Sadly, however, it does not *establish* it. For, as I noticed above, in many cases it is a bit unclear what each conception aims to do. Moreover, it is possible that, whilst a conception is aimed to answer one question, it actually ends up answering the other (or *also* the other). Compare: just like a person aiming to craft a screwdriver may inadvertently end up crafting an ice pick, a philosopher intending to answer the content question may inadvertently end up answering (either also or exclusively) the status question. Thus, whilst the observation that the different conceptions have different aims supports and provides circumstantial evidence in favor of them being distinct, it does not yet establish that they are.

A second - more convincing - argument is thus needed. The argument is based on the fact that the four different conceptions I identified are *not* extensionally equivalent. They have different extensions, and so they apply to *different* things. Hence they are conceptions *of different things*. Consider, first, the structural simulation conception. It applies exclusively to *computational states and processes*. But none of the other conceptions I identified is so restricted: they can all apply to non-computational states and processes. So, the extension of the structural simulation notion differs from the extension of all the other notions. The extension of the structural map notion also differs from the extension of the structural space and structural dynamics notion. This has actually been persuasively argued for by Facchin (2021b) who showed that *some*, *but not all*, indicators have all the features of structural maps. In particular, Facchin discussed a number of (real and imaginary) systems relying on indicators to orchestrate their behaviors, showing that, whilst *all* indicators are (and *must be*) structurally similar to their targets (minimally in the structural dynamics sense) and that such a structural similarity is essential to their functioning and so it is (and *must be*) an exploited structural similarity, it is not the case that *all* indicators are decouplable from their targets, or allow their user to identify any representational error (cf. Facchin 2021b, pp. 5487-5497; see also Lee & Calder 2023 for discussion). Thus, whereas all indicators satisfy conditions **(1)** and **(2)** of structural maps, only some satisfy **(3)** and **(4)**. However, all indicators (whether categorical or graded) satisfy the structural dynamics conception; indeed, as Facchin (2021b, p. 5488-5492) convincingly argued, each indicator *must* satisfy this conception; else, it would simply not be an indicator of its target.[16] Moreover, all graded indicators also satisfy the structural spaces conception. It follows that these three conceptions - structural maps, structural spaces and structural dynamics - have different extensions. One the assumption that a concept's extension is at least partially constitutive of its identity, it follows that the four conceptions I have identified are actually four distinct conceptions.

A third argument establishing that these four conceptions are different is that the structural similarities they revolve around take quite different *relata* on the "vehicle side" of the relation. In the case of the structural map conception, the structural similarity holds between an *individual vehicle* and a target. So, the single vehicle reproduces the target's structure with its *inner* structure, just like a map reproduces the spatial layout of the environment with its spatial layout. In the case of structural spaces, structural dynamics and structural simulations, however, the relevant *relatum* on the vehicle side of the relation is not an individual vehicle. According to these conceptions, the structural similarity holds amongst *many* vehicles (computational state types and indicator states) and a target. So, whereas the structural map conception prescribes that each individual vehicle satisfying it must be structurally similar to the target it represents, all other three conceptions do not prescribe anything of the sort.[17]

The structural space, structural dynamics and structural simulation conception differ too for their *relatum* on the vehicle side of the relation. The structural similarity involved in the structural dynamics conception involves *token* indicators states and their temporal relations. The structural space conception, however, invokes a structural similarity holding between *types* of indicator states and types of target states. Moreover, according to the structural space conception, the relevant pattern of relation holding among indicator state types need not involve *temporal* relations holding amongst such states. On the structural dynamics conception, however, the relevant pattern of relation holding amongst the indicator state tokens is always a pattern of *temporal* relations. Thus, the structural space and structural dynamics conception clearly involve different *relata* on the vehicle side of the relation. Moreover, both the structural space and the structural dynamics conception invoke a structural similarity that is different from the one invoked in the structural simulation conception: in the latter case, the structural similarity holds between computational state types. But computational state types need not indicate anything, and indicators states need not be involved in any computation. Hence, the relevant structural similarity invoked in these three notions holds among different types of relata.

A potential wrinkle to the verdict above is that although computational state types *need not* indicate anything, they *might* be used as indicators. Cerebral states, such as, for example, certain

[16] Summarizing it roughly, Facchin's argument is that if an "indicator" does not satisfy what I've called the structural dynamics conception, then its states do not covary in time with the states of its target. But if so, then there is no sense in which the states of the "indicator" indicate the states of the targets, and thus the former is *not* an indicator of the latter.

[17] As I will argue in (**§4.1**), this poses a problem for arguments to the effect that indicators are structural representations, such as the one offered in (Facchin 2021b). However, my second argument to establish that the four conceptions of structural representations I have presented are different relied heavily on Facchin's argument. So, haven't I just shot myself in the foot? No, I haven't. For, if Facchin's argument is wrongheaded (as it *is*), then he's wrong in claiming that indicators satisfy conditions **(1)** and **(2)** of the structural maps conception. Hence the extensions of structural maps, structural spaces and structural dynamics would be even more different than the argument just provided shows.

patterns of activation in our sensory cortices, are presumably indicators. And, if mainstream contemporary cognitive science is on the right track, they are also computational states. Do states such as these cause the collapse of the distinction between the structural simulation and structural space conception?

The answer is negative, as there remains an important difference between the structural similarity involved in structural spaces and the structural similarity involved in structural simulations. The structural similarity involved in the structural space conception give raise to an *analog magnitude representation* (cf. Beck 2015). In these cases, the representational vehicle (that is, the indicator state) indicates a quantity in a way such that increases and decreases of the indicated quantity correspond to increases or decreases of the indicator state. For example, as the temperature *grows hotter*, the height of the thermometer mercury bar *grows taller*. Notice that this fact is at the heart of the functioning of non-binary indicators, and that such a fact *constitutes* the relevant structural similarity the structural space conception revolves around. Computational state types, however, may, but *need not*, map onto quantities and magnitudes so as to form an analog magnitude representation. The strings of digits "10101" and "101" may for example respectively map onto quantities *a* and *a+b*, even if the "bigger" string of digits represents the smallest quantity. This does not mean that computational state types may *never* map onto magnitudes so as to give rise to an analog magnitude representation. Indeed, according to some philosophers, that is precisely what singles out *analog computation* as a special and interesting case of computation (cf. O'Brien and Opie 2006; Maley 2021a,b). So, if these philosophers are right, the structural space and the structural simulation conception partially overlap. But notice that this partial overlap does not pose a threat to my claim that the structural space and the structural simulation conception are different. For the overlap is, in fact, *partial*, and limited to a very special set of cases. Were the two conceptions identical, their overlap would have been *total*.

The arguments I provided above should also have clarified that the four conceptions I have identified are not hierarchically related by superordination or subordination. None of the notions I have previously laid out is a more restricted version of another notion. Hence, no notion qualifies as a *subordinate* of any other notion; that is, as a more restricted and specific variant of a more general notion. And, clearly, since no notion is a subordinate of any other notion, no notion is a *superordinate* (that is, a more generic, less restricted version) of any other notion either. One quick argument for this conclusion is the following: if a notion *n* is a subordinate of notion *N*, then the extension of *n* is a subset of the extension of *N* (consider, for example, how the extension of "mammal" is a subset of the extension of "chordate"). But, as argued above, none of the notions I have distinguished has an extension that is a subset of the extension of another notion - at best, there is a partial overlap in the case of structural simulation and structural spaces. Thus, no notion is a subordinate of another. For the same reason, no notion is a superordinate of any other notion either. A similar line of thought also establishes that none of these for notions can be reduced to any of the other three notions.

One may try to disconfirm the verdict above arguing that structural maps are a "limit case" of structural simulations, in which the simulation is static (so to speak) and realized in a single vehicle, rather than through the tokening of a series of vehicles. So, if this view were correct, structural maps would be a subclass of structural simulations.

However, that view is not correct, and structural maps cannot be conceived as structural simulations. For one thing, the relevant relations holding amongst vehicle constituents in the structural map case *need not* be computational relations. But, in the case of structural simulations, the relevant relations are computational (on the view Cummins 1989 proposes, they are so even by definitional point). As a result, there are some structural maps that lack the relations needed to qualify as structural simulations, and so the former cannot be a subclass of the latter. Notice further that, even when computational relations are at play, in the case of structural maps such relations should hold amongst vehicle *constituents*, rather than vehicles. But, at least according to a very widespread view of computation according to which computational processes consist in the rule based manipulation of representations (eg. Fodor 1975, p.34; 1981), computational processes hold amongst vehicles, rather than vehicle constituents.[18] Thus, the relevant computational relations the structural simulation notion hinges upon hold amongst vehicles, rather than vehicle constituents. As such, these relations cannot underpin structural maps, and so, *a fortiori*, structural maps can't be a subclass of structural simulations. Of course, one *could* theoretically revise the very notion of computation to define it as holding amongst both vehicles and vehicle constituents. But what would justify such a revision? This question, as far as I can see, has yet to receive an answer.

One may also try to disconfirm my verdict by claiming that the structural dynamics conception picks up a subclass of structural simulations. Morgan (2014), for example, can be read as claiming that even very simple binary categorical indicators are (very simple, binary) *simulations* of what they indicate, in a way that suggests that they are a special case of structural simulations. Similarly, one could argue that structural simulations are a subclass of structural dynamics by saying that, in general, simulations try to capture the *dynamics* of the simulated thing, in a way that makes them a special case of structural dynamics

Again, I don't think any of the above views is correct. Consider first whether structural dynamics may be a subclass of structural simulations. The relevant relations in the structural dynamics case are temporal relations, which must be the exact same temporal relations holding amongst the indicated states (cf. above **§2.4**; Facchin 2021b). In the structural simulation case, however, the relevant relations are computational relations. Now, since in general we don't take temporal relation to be computational, nor we do take indicators such as thermometers, hygrometers and barometers to be computers, it follows that, in most cases, items that satisfy the structural dynamics conception of structural representations *do not* satisfy the structural simulation conception - hence the former can't be a subclass of the latter. And vice versa: the state of structural simulations don't need to be temporally related in a way such that their temporal relations are the exact same temporal relation holding between portions of the simulated target. The computational simulation of an evolutionary process does not need to unfold at an evolutionary, extremely long timescale - and it is indeed useful and informative precisely because of that. Hence at least some structural simulations do not qualify as structural dynamics, and so the former can't be a subclass of the latter.

[18] The same is typically true even for "non-semantic" accounts of computation (eg. Piccinini 2015, 2020a). These accounts do not *define* computation in representational terms (and do not take computational implementation to be a matter of representation), but still they hold that, in most (if not all) normal cases, computational state transitions happen between representations and representational states.

So, the four notions of structural representation I have individuated appear to be actually *different* notions, each picking up a distinct type of structural representation.

It's now time to put them to use and show how they can make the debate over structural representations clearer, contributing to making certain pressing questions more answerable.

4 - Putting the distinction to use

In **§1** I quickly sketched the current, rather messy state of the current debate on structural representations. Here, I show that my fourfold distinction may help clarify certain thorny and contested issues in that debate, in a way that may help reach a philosophical consensus on these issues. I will do so by quickly applying my distinction to three active areas of debate, concerning (a) whether indicators and structural representations are one and the same, (b) whether structural representation can mesh with mechanistic explanations and the "cognitive neuroscience revolution" and (c) whether structural representations are present in connectionist and post-connectionist systems.

4.1 - Are indicators structural representations?

Consider, first, the debate on whether indicators are a type of structural representation of not (cf. Ramsey 2007; Gładziejewski & Miłkowski 2017; Morgan 2014; Nirshberg & Shapiro 2021; Facchin 2021b). Thus far, the debate has been dominated by a functionalist reading of the status question, trying to ascertain whether the *functional profile* of indicators is the same functional profile of structural representations. My fourfold distinction allows this question to receive a straightforward answer: yes (in some senses), but no (in others). All indicators are structural representations according to the structural dynamics conception. So, if the question is whether all indicators are structural representations according to that conception, the answer is positive. Non-categorical, graded indicators are also structural representations according to the structural spaces conception. So, if that is the relevant notion of structural representation deployed, then the answer is that a particular class of indicators (non-categorical, graded ones) are structural representations. However, in light of the discussion above, it should be clear that indicators *are not* structural maps nor structural simulations. They are not structural maps because no individual indicator state (that is, no individual vehicle) is structurally similar to its target, which is what the structural map conception requires.[19] And indicators are not structural simulations, for the relevant relations holding amongst indicator states whereby indicator states come to structurally represent their targets are not computational state transitions, contrary to what the structural simulation conception requires.

So, my fourfold distinction allows us to easily, clearly and decisively answer the question whether indicators are structural representations. It also allows us to reconstruct the dialectic exchange that led to the present confusion on whether indicators and structural representations are of the same kind. The dialectic is roughly the following: Ramsey (2007) *correctly* denied that structural simulations are indicators, Morgan (2014) *incorrectly* replied that structural spaces (or perhaps even structural dynamics) are structural simulations, to which Gładziejewski and Miłkowski (2017) replied, challenging the "structural representation" status of some *structural spaces* (i.e. graded

[19] And so, to be clear, I was wrong in my (Facchin 2021b) to claim that indicators were structural maps.

indicators) and lastly Facchin (2021b) replied - *incorrectly and again changing the subject matter* - that structural dynamics are structural maps. If this succinct reconstruction is correct, it can be seen that the focus of the debate shifted at least twice, with the papers of Gładziejewski and Miłkowski and Facchin - if not *thrice*, if Morgan's paper is best read as dealing with structural dynamics rather than structural spaces. Such unsignaled, and often unnoticed, changes of topic surely have contributed to the present confusion enshrouding indicators and structural representations, and they may even be the prime responsible for it.

4.2 - Structural representations, mechanisms and neuroscience

Another open debate concerns the compatibility of structural representations and mechanistic explanations. Some, especially defenders of the "cognitive neuroscience revolution" take structural representations to be essential to mechanistic neuroscience (Boone & Piccinini 2016; Williams & Colling 2017). Others claim the two can't coherently mesh (Kohar 2023). Can my distinction help settle this debate?

Maybe it can't *settle* that debate, but it can surely help advance it. For, Kohar's (2023) attack focuses on the *content* of structural spaces, providing a long and sustained argument to the effect that such content cannot coherently mesh with mechanistic explanations. Kohar's argument, however, seems to be based on a structural space (or at least structural space-like) conception of structural representations. So, even if they were correct, as I think they are, these arguments would only establish that the "cognitive neuroscience revolution" can't rely on structural spaces. It may still rely on other types of structural representations.

Importantly, Kohar himself argues that motor emulators - as conceptualized by Grush (2004) - *are* compatible with the "cognitive neuroscience revolution". Since motor emulators are structural simulations (at least, if one conceives of them as Grush does, see **§2.2**), then it might be the case that the "cognitive neuroscience revolution" is best understood and defended by committing to structural simulations rather than other types of structural representations. And indeed, some defenders of the cognitive neuroscience revolution are *already* moving towards a structural simulation-like understanding of representations, according to which a representation's content is at least in part dependent on the overall computational role the vehicle plays (cf. Piccinini 2022).

Such a shift towards structural simulations might also help defenders of the "cognitive neuroscience revolution" fend off another objection, recently articulated by Facchin (2024). Roughly put, the objection is that according to the cognitive neuroscience revolution, the vehicles of structural representations should be publicly observable (see Facchin 2024, see also Thomson & Piccinini 2018). Yet, we have not observed them yet, and so we should not endorse the "cognitive neuroscience revolution". Crucially, however, Facchin's argument is based on a structural map conception of structural representations. And whilst Facchin does argue somewhat convincingly that oftentimes defenders of the cognitive neuroscience revolution *do* point to that conception as the relevant one, defenders of the cognitive neuroscience revolution are still free to revise their view so as to base it on a different conception of structural representations. In the light of the above, the structural simulation view recommends itself as the one that most likely allows the cognitive neuroscience revolution to reach its epistemic aims.

Such a revision would likely be somewhat painful. After all, the structural simulation conception is grounded in *classical* cognitive science, and the cognitive neuroscience revolution is supposedly a revolution against classicism (cf. Boone & Piccinini 2016; Williams & Colling 2018).[20] Moreover, notice that, if the structural simulation conception is adopted as the relevant meaning of "structural representation", then the syntagm "(structural) representations" *cannot* refer to the physical computational states involved in computational processes, as the term refers to the *computational processes themselves*. Hence, if this reading is accepted, it can't literally be true cognition consists of computational processes operating over structural representations, as the cognitive neuroscience revolutions seemingly claims. So, the revision would be somewhat painful - but which revision isn't?

At any rate, what matters for the purpose of this paper is not whether the cognitive neuroscience revolution is ultimately correct or not. What matters is whether the distinctions I made in **§2** are helpful in advancing philosophical debates, making them inch closer to a satisfactory conclusion. The above, I submit, shows that the distinctions I made are indeed helpful.

4.3 - Structural representations, connectionism and predictive processing

My fourfold distinction also helps clarify the representational credentials of connectionist systems, including the neural networks involved in "predictive processing" accounts of cognition (cf. Gładziejewski 2016). Famously, Ramsey (2007) denied that connectionist systems host any genuine representation - structural or otherwise. Shagrir (2012) and Morgan (2014) replied by noticing that (at least some) artificial neural networks satisfy a structural simulation conception of structural representations: the computational state transitions by means of which they pair inputs and outputs are such that they *structurally simulate* some process. Who is right? In a way, both. Shagrir and Morgan are right in pointing out that *certain* neural networks satisfy a structural simulation conception of structural representation. Yet, in all fairness to Ramsey, his account focused on a rather different class of connectionist systems (namely 3 layer perceptrons used for classification). And these connectionist systems do not seem to fit the structural simulation conception.[21] Their computational dynamics simply maps input vectors onto output labels. It does not *simulate* any external system. It simply "projects" inputs onto labels. So, Ramsey was right in saying that these specific artificial neural networks do not tokenize structural representations (meaning structural simulations). He was, however, wrong in the *scope* of his claim. Some neural networks *do* seem to tokenize structural simulations.

Moreover, Ramsey's verdict is *wrong* under other conceptions of structural representation. For, the inner *state space* of three layer classifiers has repeatedly been shown to be structurally similar to the target domain the classifies has been trained to operate upon (see O'Brien and Opie 2004; Churchland 2012). So, a *different* notion of structural representation appears to be best suited to defend the representational credentials of connectionist systems.

---

[20] I'm conceding that the cognitive neuroscience revolution *is* a revolution for the sake of argument. Petrovic & Viola, however, have recently argued that the cognitive neuroscience revolution is *not* a revolution, at least not in the Khunian sense of the term (cf. Petrovic & Viola 2022). The fact that such a "revolution" seems to be in need to rely on a type of representation deeply embedded in *classical* cognitive science is in line with Petrovic & Viola's verdict on the matter.

[21] *Pace* (Cummins 1989).

What, lastly, about a particular family of post-connectionist models, namely the artificial neural networks involved in "predictive processing" theories of cognition? Their representational credentials are hotly debated (e.g. Gładziejewski 2016; Kiefer & Hohwy 2018, Downey 2018; Facchin 2021a), in particular when it comes to ascertain whether predictive processing systems tokenize structural representations. As I noticed in (§2.1), the whole debate seems largely based on the structural map notion, according to which individual vehicles need to be structurally similar to their targets. As Facchin (2021a) argued, however, it is hard to find such vehicles in connectionist systems, especially predictive processing ones. This suggests that, to better defend their position - representationalists about predictive processing should change the relevant notion of structural representation at play.

Importantly, such a change may be less traumatic than it may *prima facie* appear. As I noticed in **§2.2**, Wiese (2017, 2018) is *already* in part resorting to a structural simulation account of content. Perhaps, then, he can simply drop the structural map response to the representational status question (which, as I argued in **§2.2** might not combine well with a structural simulation account of content and function) and embrace a full structural simulation account of representations in the predictive mind. In a somewhat similar vein, Kiefer and Hohwy (2018) and Kiefer (2023) do already conceive structural representations in terms of what I've called structural simulations, even if the picture of structural simulations they propose is complicated by an appeal to functional role semantics to determine the contents of these simulations.

Again, what matters for my purposes here is not to provide (or even sketch) an account of representations that "fits" connectionists and post-connectionist systems. What matters is only that the distinction I have sketched in **§2** allows us to advance philosophical debates in a productive way. That, in my mind, is enough to adopt and deploy the quadripartite distinction I am proposing.

5 - Conclusion

Have structural representations been observed in our brains? Are they compatible with mechanistic explanations? Are they actually posited in predictive processing accounts of cognition? Do they pass the "job description" challenge? Do they have well-determined, non-disjunctive contents? In this paper, I have argued that such questions are ill-posed. For, in the current literature, the generic term "structural representation" is used to designate a large variety of different and distinct types of representations. Recognizing - and carefully separating - them is necessary in order to pose questions such as the ones that open this section in a more meaningful, and hopefully answerable, manner.[22]

References

Artiga, M. (2016). "Liberal Representationalism: A Deflationist Defense."dialectica 70 (3): 407–430

[22] Also, since I'm an anti-representationalist, I can't resist the urge of writing *somewhere* that the answer to all these questions is negative. But to argue for that is definitely beside the point of this paper.

Artiga, M. (2023). Understanding Structural Representations.The British Journal of Philosophy of Science. https://doi.org/10.1086/728714

Beck, J. (2015). Analogue magnitude representations: A philosophical introduction. The British Journal for the Philosophy of Science.

Boone, W., & Piccinini, G. (2016). The cognitive neuroscience revolution. Synthese, 193, 1509-1534.

Chemero, A. (2009). Radical embodied cognitive science. The MIT Press.

Churchland, P. (2012). Plato's Camera. The MIT Press.

Cummins, R. (1989). meaning and mental representation. The MIT Press.

Cummins, R. (1996). Representation, targets and attitudes. The MIT Press.

Downey, A. (2018). Predictive processing and the representation wars: A victory for the eliminativist (via fictionalism). Synthese, 195, 5115-5139.

Dretske, F. (1988). Explaining behavior. The MIT Press.

Egan, F. (2014). How to think about mental content. Philosophical Studies, 170, 115-135.

Egan, F. (2020). A deflationary account of mental representations. In J. Smortchkova, K. Dolega, & T. Schlicht (Eds.), What are Mental Representations? (pp. 26–54). Oxford University Press

Facchin, M. (2021a). Predictive processing and anti-representationalism. Synthese, 199(3-4), 11609-11642.

Facchin, M. (2021b). Structural representations do not meet the job description challenge. Synthese, 199(3), 5479-5508.

Facchin, M. (2021c). Are generative models structural representations?. Minds and Machines, 31(2), 277-303.

Facchin, M. (2022). Troubles with mathematical contents. Philosophical psychology. https://doi.org/10.1080/09515089.2022.2119952

Facchin, M. (2024) Neural representations unobserved—or: a dilemma for the cognitive neuroscience revolution. Synthese 203, 7 https://doi.org/10.1007/s11229-023-04418-6

Fodor, J. A. (1975). The Language of thought. Harvard University Press.

Fodor, J. A. (1981). "The Mind-Body Problem." Scientific American 244 (January 1981). Reprinted in J. Heil, (Ed.) (2004a), Philosophy of Mind: A Guide and Anthology (168–82). Oxford: Oxford University Press

Gładziejewski, P. (2015). Explaining cognitive phenomena with internal representations: a mechanistic perspective, *Studies in Logic, Grammar and Rhetoric*, 40(1), 63-90.

Gładziejewski, P. (2016). Predictive coding and representationalism, Synthese, 193(2), 559-582.

Gładziejewski, P., & Miłkowski, M. (2017). Structural representations: causally relevant and distinct from detectors. Biology and Philosophy, 32(3), 337-355

Godfrey-Smith, P (2009). Mental representations, naturalism and teleosemantics. In Teleosemantics, (pp. 42-68), New York, Oxford University Press

Gosche, T., & Koppelberg, D. (1991). The concept of representation and therepresentation of concepts in connectionist models, in W. Ramsey, S. P. Stich, D. E.Rumelhart (eds.), Philosophy and Connectionist Theory (pp. 129-163). New York,Rutledge.

Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. Behavioral and brain sciences, 27(3), 377-396.

Kiefer, A. B. (2023). What, precisely, is a Bayesian belief?: Comment on “Path integrals, particular kinds, and strange things” by Friston et al. Physics of Life Reviews.

Kiefer, A., & Hohwy, J. (2018). Content and misrepresentation in hierarchical generative models. Synthese, 195(6), 2387-2415.

Kiefer, A., & Hohwy, J. (2019). Representation in the prediction error minimization framework. In S. Robins, J. Symons, P. Calvo (Eds.), The Routledge Companion to Philosophy of Psychology (2nd Ed.) (pp. 384-410). New York: Routledge.

Kohar, M. (2023). Neural Machines: a defense of non-representationalism in cognitive neuroscience. Springer.

Lee, J. (2019). Structural representations and the two problems of content. *Mind & Language*, *34*(5), 606-626.

Lee, J. (2021). Rise of the swamp creatures. Philosophical Psychology, 34(6), 805-828.

Lee, J., & Calder, D. (2023). The many problems with S-representation (and how to solve them). Philosophy and the Mind Sciences, 4.

Maley, C. (2021a). Analog computation and representation. The British Journal of Philosophy of Science. https://doi.org/10.1086/715031

Maley, C. J. (2021b). The physicality of representation. Synthese, 199(5-6), 14725-14750.

McLendon, H. J. (1955). Uses of similarity of structure in contemporary philosophy. Mind, 64(253), 79–95.

Morgan, A. (2014). Representations gone mental. Synthese, 191(2), 213-244.

Neander, K., (2017). A mark of the mental. The MIT Press.

Nirshberg, G. (2023). Structural Resemblance and the Causal Role of Content. Erkenntnis, 1-20.

Nirshberg, G., & Shapiro, L. (2021). Structural and Indicator representations: a difference in degree, not in kind. Synthese, https://doi.org/10.1007/s11229-020-02537-y

O'Brien, G. (2015). How does the mind matter? Solving the content-causationproblem. In T. Metzinger, J. M. Windt (Eds.). Pen MIND: 28(T). Frankfurt am Main: The MIND Group. https://doi.org/10.15502/9783958570146

O'Brein, G., & Opie, J. (2004). Notes towards a structuralist theory of mental representations, in H. Clapin; P. Staines & P. Slezak (eds.), Representation in Mind: New Approaches to Mental Representaion (pp. 1-20). Oxford: Elsevier.

O'Brien, G., & Opie, J. (2006). How do connectionist networks compute?. Cognitive Processing, 7, 30-41.

Orlandi, N. (2014). The Innocent Eye. Oxford University Press.

Petrovich, E., & Viola, M. (2022). The "cognitive neuroscience revolution" is not a (Kuhnian) revolution. Evidence from scientometrics. Rivista internazionale di Filosofia e Psicologia, 13(2), 142-156.

Pezzulo, G. (2008). Coordinating with the future: the anticipatory nature of representation. Minds and Machines, 18, 179-225.

Piccinini, G. (2015). Physical computation. A mechanistic account. Oxford University Press.

Piccinini, G. (2020a). *Neurocognitive Mechanisms*. New York: Oxford University press

Piccinini, G. (2020b). Nonnatural mental representations. In G. Smortchkova, K. Dolega, & T. Schlicht (Eds.), What Are Mental Representations? Oxford University Press.

Piccinini, G. (2022). Situated neural representations: Solving the problems of content. Frontiers in Neurorobotics, 16, 846979.

Ramsey, W. (2007). Representation reconsidered. Cambridge University Press.

Ramsey, W. (2016). Untangling two questions about mental representation. New Ideas in Psychology, 40, 3-12

Ramsey, W. M. (2023). The Hard Problem of Content is Neither. Review of Philosophy and Psychology, 1-22.

Ryder, D. (2009a). Problems of representation I: nature and role. In S. Robins, J. SImons, P. Calvo (Eds.), The Routledge Companion to Philosophy and Psychology (2nd ed.) (pp. 233-250). New York: Routledge.

Schulte, P. (2023). Mental Content. Cambridge University Press.

Shagrir, O. (2012). Structural representations and the brain. The British Journal for the Philosophy of Science.

Shagrir, O. (2018). The brain as an input–output model of the world. Minds and Machines, 28, 53-75.

Shea, N. (2018). Representations in cognitive science. OUP

Silberstein, M., & Chemero, A. (2013). Constraints on localization and decomposition as explanatory strategies in the biological sciences. Philosophy of Science, 80(5), 958-970.

Sprevak, M. (2011). Review of Representation Reconsidered. The British Journal of Philosophy of Science. 62, 669-675.

Thomson, E., & Piccinini, G. (2018). Neural representations observed. Minds and Machines, 28, 191-235.

Van Gelder, T. (1995). What might cognition be, if not computation?. The Journal of Philosophy, 92(7), 345-381.

Wiese, W. (2017). What are the contents of the representations in predictive processing?. Phenomenology and the Cognitive Sciences, 16(4), 715-736.

Wiese, W. (2018). Experienced Wholeness. Cambridge, MA.: The MIT Press.

Williams, D. (2017). Predictive Processing and the Representation Wars. Minds And Machines, 28(1), 141-172.

Williams, D. (2018a). The Mind as a Predictive Modeling Engine: Generative Models, Structural Similarity, and Mental Representation. Ph.D. Disseration, University of Cambridge, UK. Accessed athttps://www.repository.cam.ac.uk/bitstream/handle/1810/286067/Daniel%20Williams%20PhD%20Thesis.pdf?s equence=1. Last accessed 17/09/2020.

Williams, D. (2018b). Predictive minds and small scale models: Kennet Craick's contribution to cognitive science. Philosophical Explorations, 21(2), 245-263.

Williams D., & Colling L. (2017). From symbols to icons: the return of resemblance in the cognitive science revolution, Synthese, 195(5), 1941-1967.