



A match made in heaven: predictive approaches to (an unorthodox) sensorimotor enactivism

María Jimena Clavel Vázquez^{1,2} 

Published online: 20 December 2019
© The Author(s) 2019

Abstract

It has been pointed out that Sensorimotor Enactivism, a theory that claims that perception is enacted and brought about by movement, says very little about the neural mechanisms that enable perception. For the proponents of the predictive approach to Sensorimotor Enactivism, this is a challenge that can be met by introducing predictive processing into the picture. However, the compatibility between these theories is not straightforward. Firstly, because they seem to differ in their stand towards representations: while Sensorimotor Enactivism is said to belong to the non-representational wing of cognitive science, predictive processing has a representational profile. And secondly, because they exhibit different explanatory strategies: while Sensorimotor Enactivism prioritizes the interactions of the embodied agent, predictive processing has internalist commitments. The aim of this paper is to address these concerns and show that a predictive approach to Sensorimotor Enactivism is viable. More specifically, I focus on the Free-Energy approach, a theory that falls within the ballpark of predictive processing. In this paper I argue for the following claims. I argue that (a) both Sensorimotor Enactivism and the Free-Energy approach may be understood for some systems in representational terms. The non-representational reading of Sensorimotor Enactivism is not mandatory and neither is the representational reading of the Free-Energy approach. (b) Sensorimotor Enactivism is, in this respect, compatible with both representational and non-representational interpretations of the FEA. So, the position towards representations of these frameworks should not stand in the way of a predictive approach to Sensorimotor Enactivism. I also show that (c) the Free-Energy approach allows for an account that prioritizes the interaction of the embodied agent with the environment. This is the explanatory strategy followed by Sensorimotor Enactivism. To justify this strategy and following other proponents of Sensorimotor Enactivism, I argue that by referring to the interactions of the embodied agent a better account of the phenomena in question is provided. On this basis, I claim that (d) Sensorimotor Enactivism and the Free-Energy approach are compatible in what concerns their explanatory strategy as well. Thus, making the case for the viability of the predictive approach to Sensorimotor Enactivism.

✉ María Jimena Clavel Vázquez
menaclavazq@gmail.com

Extended author information available on the last page of the article

Keywords Sensorimotor Enactivism · Predictive processing · Free-energy approach · Perception · Perceptual experience

1 Introduction

From the perspective of sensorimotor enactivism (henceforth SMEn), perception is “something we do” mediated by something we know (O’Regan and Noë 2001, p. 970).¹ This theory joins other enactivist proposals in understanding cognition, action, and perception as active processes that are deeply intertwined, and that should be explained by referring primarily to the agent’s bodily constitution and her interactions with the environment. Perception is, for SMEn, enacted and brought about by movement. It is an interactive phenomenon that is assimilated to an exploratory activity through which the agent probes the environment as it offers possibilities for action. In that sense, perception is understood, at least in part, as the process of accessing sensory information from the world as an external memory, thus opposing other views of perception, according to which this would consist in the process of constructing a model of the world. According to this theory, perception requires the possession and skilful execution of sensorimotor knowledge: an implicit and practical knowledge of the lawful ways in which sensory input changes given certain interactions between the perceiver and the object perceived. Consider, for instance, how one spontaneously squints the eyes when an image looks blurry (Noë 2004, p. 1). We do this in order to make the image appear clearer, perhaps to get a better focus. Squinting our eyes in this case suggests that we know how the image would change after doing this. SMEn claims that perception is constituted by our handle on the way sensory input systematically changes after an interaction.

SMEn is a theory of perception that aims at explaining perception from the point of view of the agent and the agent’s experience. Nonetheless, it is a naturalistic theory that is taken to be continuous with a causal account of perception. For SMEn, a theory of perception must straddle both levels, the agential and the causal (Noë 2004, pp. 31–32). The causal story that is favored by SMEn is one that accounts for perception in terms of the interactions of an embodied agent with her environment.² This might explain something that has been pointed out in the literature: SMEn says very little about the neural mechanisms that are involved in perception (Chemero 2009; Clark 2009; Seth 2014, 2015; Downey 2017). It is important to consider that this lack of neuro-talk is not seen as a challenge within the framework. After all, SMEn starts from the claim that looking inside the head does not take us too far in our understanding of perceptual experience. SMEn emphatically rejects that an account of the neural mechanisms of perception is sufficient for perceptual experience. The latter is constituted by the interactions of the agent with her surroundings.

¹ See Hurley (1998), O’Regan and Noë (2001), Noë (2004), O’Regan (2011).

² The agential and the causal are continuous in the sense that the former depends on the latter: features exhibited at the agential level emerge from the complex sensorimotor interactions between the agent and the environment. Hurley (1998) illustrates this continuity with the case of a circus performer whose complex movements are both a source of and dependent upon complex dynamic feedback loops between herself and the environment. As Hurley puts it: the content of the performer’s intentions and perceptions depends on this complex causal flow (Hurley 1998, p. 2).

Nonetheless, the theory does not deny that the brain is involved in perception. See, for instance, the following claim by Noë: “The brain has a job to do (...) a careful examination of the way experience and the brain’s activity depend on each other makes plausible the idea that the brain’s job is, in effect, to coordinate our dealings with the environment. It is thus only in the context of an animal’s embodied existence, situated in an environment, dynamically interacting with objects and situations, that the function of the brain can be understood” (Noë 2009, p. 65). Neural mechanisms, although not sufficient, are nevertheless necessary for perceptual experience. Furthermore, these should be understood within the wider context of the interactions of the embodied agent.

So, does SMEn owe us an account of the way the brain is involved in perception? One thought might be that it does not. Surely the brain plays a role in supporting perceptual experience, but a story at the neural level does not place a strong constraint on our account of perception because perception is explained at the level of the whole organism and its interactions with the environment. Nonetheless, by claiming that the causal and the agential level are continuous, SMEn establishes a tighter link between the two levels. The causal story goes hand in hand with the story at the level of the agent’s interactions. And part of that causal story is about the neural structures that underpin perceptual experience. The causal story is not complete until some account of the role of the brain is provided.

The latter seems to be the thought behind the proposals for a predictive approach to SMEn. Predictive processing³ offers an account of the neural mechanisms involved in perceptual experience that might allow SMEn to explain how the brain contributes to perceptual experience. For instance, Adrian Downey (2017) appeals to predictive processing in order to respond to the objection advanced to sensorimotor approaches to perception that claims that they are empirically vacuous, while Anil Seth (2014) appeals to predictive processing to provide, among other things, a causal story that accounts for sensorimotor knowledge.

At the outset, a predictive approach to SMEn is appealing because both theories share a view of perception as an anticipatory phenomenon.⁴ While for SMEn, one squints the eyes expecting an image to appear more focused, predictive processing claims that the brain advances predictions about the object of perception. That said, there are at least two good reasons to think that these views are incompatible and that, in consequence, they do not belong together. Firstly, predictive processing has a representational profile that seems incompatible with the non-representational profile that is traditionally associated with SMEn (see e.g. O’Regan and Degenaar 2014; Di Paolo et al. 2017). Secondly, SMEn’s explanatory strategy might clash with predictive processing. While predictive processing is a theory that is usually interpreted as centred around the brain, SMEn prioritizes the interactions of the embodied agent over the specification of the neural mechanisms that enable perception.

³ For a philosophical overview of predictive processing, see Clark (2016) and Hohwy (2013). In section 3 of this paper, I offer a review of the framework.

⁴ They are both what Bruineberg et al. (2018) call future-oriented frameworks.

Importantly, as has been shown by Clark (2016, 2017b), a description of predictive processing as a neuro-centric approach is not mandatory. For him, there are some valuable lessons to be learned from developments in embodied cognition. For instance, that problem-solving can be distributed between brain, body and world (Clark 2016, p. 246). And there is a sense in which these lessons are already reflected in predictive processing. Neural dynamics, as they are modelled by the framework, should be considered to be in the service of “world-engaging action” (Clark 2016, p. 250). In line with this idea, Clark (2017b) argues that neural processing can be understood as enabling efficient and reliable interactions with the world.

In this paper, I focus on the Free-Energy approach (henceforth FEA), a theory that, although independent from predictive processing, has been developed hand in hand with it because predictive processing is a possible computational implementation of the FEA. The FEA starts from the perspective of the organism to understand neural dynamics. According to the FEA, in order to reach viability, biological systems seek to avoid surprising dealings with their surroundings.⁵ In section 2, I will say more about the relation between surprise and free-energy. At this point, it will suffice to say that, according to the FEA, the agent seeks to avoid states that would challenge its viability. Surprise, in this context, is not to be confused with the personal-level phenomenon of feeling surprised by an unexpected event. Rather, the term refers to the improbability of an agent visiting certain states “on average and over-time” (Kirchhoff and Kiverstein 2019, p. 57). Free-energy is taken as a principle of surprise minimization that can be implemented by predictive processing. As I intend to show, the FEA is a good prospect when it comes to providing an account of the neural mechanisms involved in perception that is compatible with SMEn. Clark (2017a), for instance, shows that, if cognition is understood in terms of surprise minimization, we should not limit our explanations to neural activity given that non-neural elements of the body and things external to the agent can also be considered part of these processes of surprise minimization. The FEA is promising precisely because it understands the role of the brain within the wider context of the interactions of the organism, interactions that are for SMEn constitutive of perception.

There are two tasks that must be undertaken by the proponents of a predictive approach to SMEn. Firstly, it is necessary to justify why predictive processing is the right causal story for SMEn in what concerns the role of the brain in perception. However, before doing that, there is a second task that must be tackled: it is necessary to show that such an approach is viable at all. To do that, it is necessary to address the concerns introduced above, namely: (i) that a predictive approach to SMEn is not viable in virtue of the respective stand of these frameworks towards representations; and (ii) that a predictive approach to SMEn is not viable in virtue of the commitment of the FEA to an internalist account of perception. The aim of this paper is to address these concerns and show that, in principle, a predictive approach to SMEn is viable. More

⁵ It is under discussion whether the FEA falls within the wider enactivist framework that accounts for perception by referring to the wider perspective of the organism (e.g. Seth 2015; Bruineberg et al. 2018; Kiverstein et al. 2019). In that line, Clark (2016, 2017a), has argued that the FEA is not in tension with embedded-enactive-extended approaches to cognition. On the other hand, the FEA has also been linked to ecological approaches of perception (e.g. Orlandi 2016; Kirchhoff 2018).

work would then be required to tackle the first task and show that predictive processing is the best or the only causal story for SMEn.⁶ In this paper I argue for the following claims: (a) for some systems, both SMEn and the FEA may be understood in representational terms; (b) SMEn is, in this respect, compatible with both a representational and a non-representational interpretation of the FEA; (c) the FEA allows for an account that prioritizes the interaction of the embodied agent with the environment; and (d) given that the latter is the explanatory strategy of SMEn, SMEn and the FEA are compatible in this respect as well.

The structure of the paper is the following. In section 2, I begin by briefly presenting an overview of the most relevant aspects of SMEn for my purposes: SMEn's position on representations and its explanatory strategy.

In section 3, I present the general features of predictive processing and of the FEA. Although the FEA has been interpreted as a non-representational framework (see e.g. Bruineberg et al. 2018; Downey 2017, 2018; Hutto and Myin 2017), I focus on the representational interpretation of the framework because, more often than not, both predictive processing and the FEA are read in this way. Moreover, establishing a marriage between SMEn and predictive processing when the FEA is read in its more common representational key is much less straightforward than when the FEA is read in its less common nonrepresentational key, since SMEn is often understood as having a non-representational profile. It is worth clarifying that I am not committed to a representational reading of the FEA. Given that this interpretation of the FEA represents the greater challenge in arguing for the compatibility with SMEn, in the paper I discuss at large the representational commitments of the FEA and reconstruct the strongest version of the position to motivate it. Nonetheless, the aim of this discussion is not to provide knock out arguments in favor of (or against it, for that matter) the representational reading of the FEA. Given the purpose of this paper, showing the viability of the marriage between the representational FEA and SMEn is the harder case, but the one which reflects what justifiably qualifies as the default understanding of the framework in the wider cognitive science literature.

To argue for the viability of this joint approach, in section 4, I advance an account of SMEn's stand towards representations, an account which, although slightly unorthodox, respects the foundational elements of the framework. More precisely, I argue that given SMEn's notion of sensorimotor knowledge, perception can be interpreted as a domain that admits the positing of representations, i.e. that sensorimotor knowledge can sometimes be enabled by representations (although it need not to be). In this

⁶ See Flament-Fultot (2016) for a case against the predictive approach to SMEn. Flament-Fultot emphasizes two shortcomings of predictive processing. Firstly, he questions the plausibility of the model posited by predictive processing. The theory states that the architecture of the brain must be hierarchical, but for him it is not clear that hierarchical models can be implemented given the anatomy of the brain. Furthermore, given the complexity of the models posited by predictive processing, he worries that these might be computationally intractable for the brain. On the other hand, he joins others in their concern for the representational profile of predictive processing. Flament-Fultot is concerned with the first question about whether predictive processing is the right causal story to account for sensorimotor knowledge. Given that I do not tackle this question, I will not address Flament-Fultot's worries directly. In section 3, however, I discuss the representational profile of predictive processing. Here I would like to note that the first worry expressed by Flament-Fultot is an open empirical question. As things stand, this should not prevent us yet from taking the predictive approach as a competing account of sensorimotor knowledge.

section, I also argue that the representations posited by the FEA are not problematic for SMEn. This addresses the first of the concerns raised above, namely that a representational reading of the FEA obstructs the marriage of the two frameworks.

In section 5, I briefly present the proposal of Anil Seth for a predictive approach to SMEn and one important objection advanced by Ezequiel Di Paolo to the approach. Seth's proposal allows me to bring forward the second concern that arises from a marriage between the FEA and SMEn, namely that an internalist reading of the FEA according to which the brain is necessary and sufficient to account for perceptual experience stands in the way of the joint approach.

Finally, in section 6, in light of Di Paolo's objection, I revisit the explanatory strategy advanced by SMEn. Along with other proponents of SMEn, I argue that the strategy of referring to the interactions of the embodied agent to account for perception is justified because this strategy provides a better account of the phenomena in question. Following Hurley and Noë, I claim that this explanatory strategy should not rest on an ontological claim about the substrates of perception. On the basis of this analysis, I finish this section by discussing briefly how the FEA is compatible with the proposed strategy, and so address the second of our concerns. I thus show that the FEA is compatible with SMEn.

2 Sensorimotor enactivism

SMEn claims that perception is mediated by sensorimotor knowledge. This is an implicit and practical kind of knowledge of the lawful set of Sensorimotor Contingencies (henceforth SMCs): covariations between movement either of the perceiver or of the object perceived, and its sensory outcome (Buhrmann et al. 2013, p. 2; see also O'Regan and Noë 2001, p. 943; Noë 2004, p. 63). For SMEn, perception is in part constituted by the possession and skillful execution of sensorimotor knowledge which consists in our capacity or ability to act with sensitivity to SMCs, or to the way sensory input covaries with movements of the object or the agent (see Silverman 2016, 2018). Sensorimotor knowledge can be glossed counterfactually as the knowledge of how sensory information changes after an interaction between the perceiver and an object.

By positing sensorimotor knowledge, SMEn is able to account for the qualitative features of perceptual experience. Sensorimotor knowledge is advanced to account, for instance, for phenomena such as amodal perception or perceptual presence. This refers to a puzzling aspect of our perceptual experience, namely that even when we are not attending to a feature of a scene or object perceived or even when an object is partially occluded, this feature appears in our experience as being absent or as being out of sight (Noë 2004, p. 59). Think of the experience we have when seeing e.g. a tomato: despite there being one sense in which we see only the side of the tomato facing us, Noë points out, the tomato appears in our experience as a three-dimensional, round object. The side of the tomato that is hidden also appears in our experience (Noë 2004). But it does not appear in the same way as the visible side of the tomato, rather it appears as absent. For SMEn, this aspect of perceptual experience is explained by our possession of sensorimotor knowledge in the following way.

Perceptual experience, according to the theory, requires the satisfaction of two conditions (see O'Regan and Noë 2001). Firstly, it requires sensorimotor interactions

to be ruled by SMCs: the interactions of the perceiver are governed by the lawful ways in which sensory input covaries with the interactions of the agent given its embodied structure and the structure of the object. Secondly, perceptual experience requires the mastery of sensorimotor laws. It is not only that the interactions of the perceiver correspond to or fit with a pattern of SMCs, but that the perceiver actively interacts masterfully with the object in a way that is governed by these laws. As O'Regan and Noë put it, the perceiver “must be actively exercising its mastery of these laws” (O'Regan and Noë 2001, p. 943).⁷ Consider, for instance, the well-known case of a cat behind a fence (Hurley and Noë 2003; Noë 2004). Perceiving the cat is governed by the practical knowledge of the way the hidden bits of the cat would appear if I were to move myself. It is not only that this can be described under e.g. certain laws of occlusion. Rather, for SMEn, the agent's perceptual experience of the cat already manifests that she has a handle on these laws of occlusion.

With the notion of sensorimotor knowledge, SMEn explains that perception is a skillful activity. The agent abstracts certain sensorimotor lawlike regularities that are available for her and the knowledge of these is “latent, potentially available for recall” (O'Regan and Noë 2001, p. 945). In this way, sensorimotor knowledge explains the agent's mastery of SMCs: the agent possesses knowledge of the rules that govern sensorimotor interactions. The phenomenal qualities of perceptual experience are explained by the agent's ability or capacity to access e.g. the hidden side of the tomato or the occluded bits of the cat. Sensorimotor knowledge fulfils two explanatory roles. On the one hand, it is meant to explain that perceptual experience consists in the *execution* of embodied skills. And, on the other hand, that some of the phenomenal features of perceptual experience are determined by the systematic sensory consequences of *non-actualized* interactions (Silverman 2018, p. 164).

Now, O'Regan and Noë (2001) explicitly reject the thought that sensorimotor knowledge consists of propositional knowledge. So, in order to make sense of sensorimotor knowledge, many argue in favor of its practical nature: it is an implicit practical knowledge, a kind of know-how or embodied skill that is executed by the agent (see e.g. Flament-Fultot 2016; Loughlin 2018a; Silverman 2016, 2018). However, another plausible interpretation of sensorimotor knowledge consists in taking at face value the counterfactual formulation introduced at the beginning of this section. From this perspective, sensorimotor knowledge is knowledge about the way input would change given certain interactions, i.e. it is implicit propositional knowledge of counterfactuals involved in perceptual experience.

Neither of these readings comes without problems. On the one hand, the more orthodox practical reading of sensorimotor knowledge faces the challenge of fulfilling the two explanatory roles mentioned above. While sensorimotor knowledge understood in these terms can easily explain that perception consists in the *execution* of an embodied skill, it might struggle to explain the qualitative features of perceptual experience that depend on *non-actualized* interactions, e.g. perceptual presence. When we see the tomato, we see it as a round object even when we do not interact with it in a way that would make it obvious that it is a round object: we do not need to hold the tomato, nor do

⁷ Importantly, these interactions should be considered within the context of the agent's goals. O'Regan and Noë add as well that, to properly talk about perceptual awareness, the agent must be able to use the mastery of SMCs for thought and planning (O'Regan and Noë 2001, p. 944).

we need to be disposed to turn it around to perceive it as a round object. The qualitative features of perceptual experience rely on the mastery of sensorimotor patterns of interactions that are not currently being actualized. To account for this, the proponent of SMEn must posit a second-order practical knowledge (see Flament-Fultot 2016; Silverman 2018).

On the other hand, although the (unorthodox) counterfactual reading allows for a more straightforward account of how SMEn fulfils the two explanatory roles mentioned above, it relies on the positing of internal representations to account for this (see Seth 2014). On this reading, perception is understood as the “deployment” of internal representations of SMCs (Silverman 2018, p. 158). Thus, this reading inherits the challenges that come with internal representations, such as: i. struggling to account for their content (see e.g. Hutto and Myin 2012, 2017); and ii. arguing convincingly that the functional role attributed to putative representational states is, in fact, of a representational nature (see e.g. Ramsey 2007).⁸ In addition, the sense in which sensorimotor knowledge is practical knowledge becomes less straightforward.

Now, the concerns that arise from positing representations can be deflated if we consider that these challenges must be met by the causal account plugged into SMEn. So, whether this approach is able to meet the challenges advanced to representationalism will depend in part on whether the causal account to which it is paired with can provide a sound answer to them. I address some of these worries in connection to the FEA in section 2.1 of this paper.

As Silverman (2018) notes, SMEn seems to face a dilemma: either it takes sensorimotor knowledge as practical and faces the drawbacks mentioned above, or it takes sensorimotor knowledge to be counterfactual and risks losing the practical aspect of it. To resolve this dilemma, Silverman (2018) advances a proposal in which sensorimotor knowledge fulfils the two explanatory roles that it is meant to fulfil, while maintaining its practical character.⁹ His proposal is to take perception as consisting of embodied skills that are, in turn, grounded on sensorimotor knowledge. Silverman defends the view that sensorimotor knowledge is an ability that grounds other embodied abilities. In that sense, it is a condition of possibility or a requirement to execute other embodied skills. In perception, agents manifest the possession of sensorimotor knowledge by exhibiting sensitivity to non-current states of affairs. This knowledge grounds the execution of the embodied skills that constitute perception. Importantly, this is compatible with accepting that there are multiple ways in which sensorimotor knowledge might be *causally enabled*, including representational states. As I will discuss briefly in

⁸ I am following Williams’ (2018) summary of the challenges to representational accounts of cognition. He notes that the literature advances a third challenge to representationalism, namely the cognitive functional challenge. This challenge consists of two steps: (i) accepting that cognition should be subsumed under action or understood as an active phenomenon; and (ii) in consonance with (i), the positing of representations should be avoided partially or altogether (see Williams 2018, pp. 143–149). I do not address this point here because I do not consider it to be a challenge to representational accounts as such, but rather an explanatory strategy. Moved by the idea that cognition is an active phenomenon, some opt not to prioritize explanations that rely on internal representations. This strategy is similar to the explanatory strategy followed by SMEn. In section 1.1, I present this strategy.

⁹ In addition, Silverman’s position reconciles the practical and propositional readings of sensorimotor knowledge. For him, sensorimotor knowledge can be understood as attunement and can be construed as propositional knowledge depending on the complexity of the agent in question (Silverman 2018, p. 166).

section 2.2 and in section 4.1, this does not imply that perception is *constituted* by representational states.

This position allows us to incorporate the notion of sensorimotor knowledge that is at play in predictive approaches to SMEn. Whatever the brain is doing, it should support the agent's sensitivity to non-current sensorimotor interactions. As we will see in section 4, Seth is committed to a counterfactual reading of sensorimotor knowledge. With the aim of specifying the neural implementation of sensorimotor knowledge, Seth exploits the aspects of predictive processing that are relevant for SMEn. Specifically, Seth is interested in the way predictive processing can accommodate the counterfactual aspect of sensorimotor knowledge.

In what follows, I discuss two aspects of SMEn that are relevant to the viability of a predictive approach to this theory: (1) SMEn's explanatory strategy; and (2) its position on representations.

2.1 The explanatory priority of the interactions of an embodied agent

As mentioned before, SMEn holds a commitment to a kind of naturalism according to which both levels within a theory of perception, the agential and the causal, should be continuous. This continuity is guaranteed by a common explanatory strategy. If we wish to argue that a predictive approach to SMEn is viable, we must make this strategy explicit.

SMEn's explanatory strategy is made explicit in the theory of SMCs advanced by Di Paolo et al. (2017). Let us take, for instance, one of their cases: a case of tactile discrimination between two different shapes, e.g. discriminating between two light switches of different size but similar shapes, using only the tip of the finger (Di Paolo et al. 2017, p. 62).

Di Paolo et al. advance a model of this behavior starting with a minimal cognition model (in this case, a network comprised of only two neurons). These models use agents that are very simplistic in what concerns their sensory and motor structures. Instead, the coupling with the environment becomes central to understanding the performance of the agent in a specific task (Di Paolo et al. 2017, p. 62).¹⁰ This explanatory strategy does not intend to deny the presence and influence of the elements that are internal to an agent, but instead to explain the behavior of the agent by emphasizing the determining role of her interactions.

The explanatory strategy comprises two claims: (a) on the one hand, it argues against the sufficiency of the brain to explain the behavior of an agent; (b) on the other hand, it argues that the interactions of the embodied agent enjoy explanatory priority over any reference to the internal organization of the system.¹¹ In other words, as a strategy to understand and explain behavior, the

¹⁰ The authors study the behavior of minimal agents in discriminating narrow from wide shapes. The agents are meant to move away from the wider shape and approach the more narrow one. The task cannot be undertaken by relying on direct sensory information because the shapes are randomly positioned at each trial. After several trials, agents were selected that performed better, i.e. that distinguished wide from narrow shapes across different positions. The authors, then, model the agents' sensorimotor patterns in order to identify and predict the trajectories visited by the agents, their long-term tendencies and regularities, and the differences between behaviors. The analysis also models the agent's internal structure. See Di Paolo et al. (2017, pp. 62-69).

¹¹ This strategy is implicit also in e.g. Di Paolo (2014) and Loughlin (2018a).

interactions with the environment through the embodied situation of the agent should precede any reference to its inner structure. In that sense, neural dynamics are subordinated to, and should be understood within, the context of the embodied agent's interactions.

This same strategy is present in the work of Noë and of O'Regan. See, for instance, this claim by O'Regan and Noë:

But the brain's activation does not in itself constitute the seeing. In partner dancing, specifying the bodily configuration or brain state of the dancer is not sufficient to specify the dance (because we need additionally to know how the partner is currently interacting). Likewise, in seeing, specifying the brain state is not sufficient to determine the sensory experience, because we need to know how the visual apparatus and the environment are currently interacting. (O'Regan and Noë 2001, p. 966).

And the following by Noë:

On the enactive approach, brain, body, and world work together to make consciousness happen (Thompson and Varela 2001). (...) Experience is not caused by and realized in the brain, although it depends causally on the brain. Experience is realized in the active life of the skilful animal. (Noë 2004, p. 227).

As can be seen in these quotes, the claim for the explanatory priority of the interactions of the embodied agent is accompanied by remarks about the ontological nature of perception. So, the claim goes: an account of the interactions of the embodied agent is required to account for perception because perception is not realized in the brain. Rather, the realizing base of perception extends beyond the inner (neural) structure of the agent. This ontological commitment is also formulated in terms of the causal/constitutive distinction: an account of the interactions of the embodied agent is required because perception is, constitutively, an interaction between the agent and the world. The brain either enjoys a 'mere' causal relevance or it is only partly constitutive. More needs to be said to explain the relation between the ontological and the explanatory claim, and to make the case for the ontological commitments at play. I discuss this in section 5. For now, it is sufficient to note that the explanatory claim relies on the ontological claim, that is, the claim about the realizing base of perceptual experience. For SMEn, given that perception is constituted by the interactions of the embodied agent (or given that the agent-environment system realizes perception), accounting for perception by referring to the agent's interactions is required.

I now move on to the second aspect relevant for the discussion about the viability of the predictive approach to SMEn, namely its position towards representations.

2.2 Sensorimotor Enactivism's position on representations

Although traditionally associated with the non-representational wing of cognitive science, whether the proponents of SMEn are willing to accept at least some kind of

representations remains unclear.¹² However, *if* the theory were to posit representations, what kind of representations would these be? This is the question I am concerned with in this subsection. Those who claim that SMEn does not reject representations per se, claim instead that the theory places constraints on the positing of representations. On the one hand, it seems that the theory constrains the *kind* of representations that can be posited (see e.g. Noë 2004; O'Regan 1992, 2011). On the other, it seems that SMEn also constrains what the *content* of these representations should be.

I begin by going through the first constraint, the constraint placed to the kind of representations that can be posited. See, for instance, the following claim by Noë:

The claim is not that there are no representations in vision. (...) The claim rather is that the role of representations in perceptual theory needs to be reconsidered. (...) It is a mistake to suppose that vision just is a process whereby an internal world-model is built up, and that the task-level characterization of vision (...) should treat vision as a process whereby a unified internal model of the world is generated. This is compatible with there being all sorts of representations in the brain, and indeed, with the presence of such representations being necessary for perception. (Noë 2004, p. 22).

This position with respect to representations involves, firstly, the rejection of a view that takes perception to be in the business of building an internal model of the scene perceived. For SMEn, a model or a representation of the world should not be the end-result of perception. Perceptual processing might be accounted for in representational terms, but perceptual experience does not consist of representations. The position expressed here by Noë is consistent with a view according to which perceptual experience is *causally enabled* by representations but is not *constitutively* representational. I will further discuss the distinction between enabling and constitutive representationalism in section 3.1.

But that is not the only way in which at least some proponents of SMEn constrain the positing of representations. SMEn also constrains the content these representations might have. The theory rejects an account of perception that involves the generation of a mental model of the scene perceived that is neutral to the actions and constitution of the perceiver. For instance, SMEn rejects internal allocentric representations of space, e.g. representations that preserve metric relations between the objects of the scene perceived. Consider the following claims by O'Regan: "Viewers simply take the incoming information as it comes, and do not attempt to integrate it into a precise, metric-preserving internal representation, but only into a kind of non-metric, schematic

¹² Noë (2004), O'Regan and Noë (2001), O'Regan (2011), Silverman (2018), for instance, can be taken to argue for a weaker perspective on the positing of representations. However, Chemero (2009), Hutto and Myin (2012, 2017), Noë (2012), O'Regan and Degenaar (2014), Myin and Loughlin (2018), among others, offer a more radical perspective on their rejection of representationalism. The arguments against representationalism within the enactivism literature in general, and within SMEn in particular, are diverse. Chemero (2009), for instance, advances his radical perspective as a competing theory to representational alternatives. For him, the radical, non-representational, stance can better explain the same phenomena explained by positing representations. Hutto and Myin (2012, 2017), on the other hand, argue against the possibility of providing a naturalised account of content. Given that representations are contentful internal states, we should avoid positing them. Given the purposes of this paper, I will not discuss further how the representational profile of SMEn stands with respect to these arguments.

mental framework, or structural description” (O’Regan 1992, p. 471). And: “the passive sensation we get from the retina or some iconic derivative of the information upon it (...) is being used to supplement a mental schema we have about the results of the possible actions that we can undertake with our eyes (or heads or bodies)” (O’Regan 1992, p. 472).

So, if SMEn is to posit some kind of representations, besides these not being the end-result of perception, these should meet the following requirements: they should be representations of the scene perceived that are indicative of further possible contextual interactions and that are agent-specific, thus, referring to the agent’s constitution and history of interactions. This kind of representations can be assimilated to “action-oriented representations” (see Clark 1997; Wheeler 2005; Cappuccio and Wheeler 2012 for discussion).

In this section I have introduced the aspects of SMEn that will become relevant in the discussion of the viability of the predictive approach to SMEn. As will be seen in the next section, predictive processing in general, and the FEA in particular, are often considered frameworks with a representational profile. For that reason, it is important to bear in mind the kind of representations that would be acceptable if SMEn were to posit representations. With respect to SMEn’s explanatory strategy, in section 5, I discuss the potential conflict between SMEn and the FEA. This conflict will become evident in section 4 where I discuss Anil Seth’s proposal for a predictive approach to SMEn.

3 The free-energy approach

The FEA and predictive processing are independent theses, although often discussed hand in hand. While the FEA is a theory about the adaptive behavior of biological systems, predictive processing is a functional approach to the brain with the potential of providing a unifying strategy to understand neural activity. Predictive processing is a possible implementation of the FEA. In that sense, within the context of the FEA, predictive processing can be taken as a possible account about the workings of the central nervous system and its contribution to the adaptive behavior of biological systems.¹³ For this reason, I start by saying a bit more on predictive processing before going through the FEA.

Predictive processing takes the brain to be driven by top-down processing, emphasizing that neural activity is hierarchically structured. Moshe Bar, for instance, characterizes the brain as proactive because, according to predictive processing, it advances a prediction of the objects of perception instead of merely reacting to input (Bar 2009). An example that can be useful to illustrate this position relates to visual tracking (adapted from Friston et al. 2012; Seth 2014). How are we able to follow the trajectory

¹³ It is worth noting that nothing in the FEA’s view of the adaptive behavior of biological systems mandates that surprise minimization should be implemented in hierarchical models such as those posited by predictive processing (see Clark 2017a). Nonetheless, it is also important to note that others in the predictive processing literature indicate that the link between the FEA and predictive processing is tighter. Hohwy, for instance, does not take these frameworks to be fundamentally different or independent (Hohwy 2013, p. 4). The FEA expresses principles that are essential to predictive processing, so buying into predictive processing already implies the acceptance of some core principles of the FEA. Nonetheless, in this paper, I take the theses to be independent.

of an object in the sky, for instance, of a bird, if we have no clue where it is heading? According to predictive processing, our capacity to follow a visual cue can be explained by taking the brain to predict, based on previous knowledge, the future position of the bird in the sky. Processing coming from a higher level of the neural hierarchy advances a hypothesis about stimuli received at a lower level. Going back to our example, we could say that based on past information the brain advances a prediction about the position of the bird. The prediction is met by sensory information coming from a lower level of the hierarchy. The brain processes an error signal that encodes the difference between the prediction advanced and sensory input, providing thus feedback about the accuracy of the prediction. The process allows the system to advance better predictions over time.

For predictive processing, neural activity relies on prior knowledge that constitutes the system's best model of the world, a generative model from which predictions are generated. Neural activity is interpreted as inferring the most likely causes of stimuli based on said model. The generative model is constantly updated on the basis of the system's interactions with the aim of achieving better predictions over time. *Error* feedback is provided by input coming from a lower level. In addition to this, the brain also processes the reliability and salience of error feedback, namely its precision.¹⁴

Most approaches use an approximation to Bayesian inference to model neural activity a computational method that allows the modelling of the process just described. Bayesian inference, however, has a downside: it is computationally very complex and, in many cases, intractable. Approximate Bayesian computations posit a likelihood function that limits the space of probability from which the prediction is derived. This makes them computationally tractable. The FEA is the story behind the choice of free-energy as the function that limits or bounds the space of probability (Friston and Stephan 2007, p. 425). It is with this approach that I show SMEn to be compatible.

To understand the FEA, we should begin with its notion of a biological system, an embodied agent whose boundaries are defined in virtue of its interactions with the environment. This notion has a heuristic function and should be taken as a guide for the formulation of mathematical models that explain and predict neural processing.

According to the FEA, biological systems are open systems that are not in equilibrium and seek to minimize dispersion, thus reaching viability. Adaptive behavior, in this approach, consists in the minimization of surprising dealings with the environment. As mentioned before, surprise does not refer to a personal-level phenomenon. Rather, surprise is a measure of the probability of states in which an agent might find itself given its model. Surprising states are those that have a low probability of being visited by an agent on average and over time. These are states that would jeopardize an agent's viability. Unsurprising states, on the other hand, are states that have a high probability of occurring (see Kiverstein 2018, pp. 5–6).

Now, the free-energy principle states that, in order to minimize disorder or dispersion, biological systems “make implicit inferences about their surroundings” (Friston and Stephan 2007, p. 420). One way to put this is to say that, to ensure their viability, biological systems minimize surprise by inferring which states they are expected to visit

¹⁴ For an overview of the notion of precision, see Clark (2016, chapter 2) and Kirchoff and Kiverstein (2019, chapter 5).

over time, i.e. which states have higher probability or are less surprising. However, the agent does not have access to the measure of probability of the states it is expected to visit over time because the agent would need to evaluate all the possible states it could visit. This is the information provided by the notion of free-energy (Kiverstein 2018, p. 6).

The free-energy principle originated in statistical thermodynamics. In that context, free-energy measures the discrepancy between the energy of a system and its entropy, i.e. its tendency towards disorder. In our context, the notion of free-energy that is at play is *informational* or statistical free-energy. Here, free-energy is a function of probability distributions (Friston and Stephan 2007, p. 420). Free-energy is an upper-bound on surprise, given a generative model. It measures the probability of states given the agent's model. For Friston, free-energy can be evaluated because, simply put, it is a function of sensory states and the current state of the brain. Given that free-energy is always greater than surprise, by minimizing free-energy, the agent minimizes surprise (Friston 2009, p. 294).

Neural activity is, thus, understood in terms of surprise minimization. That is, as an instance of the adaptive behavior of biological agents. A system deploys two strategies to reduce the probability of error in its predictions over time (and, thus, to minimize surprise): (a) changing its internal states or parameters (i.e. perceptual inference), and (b) acting upon the environment (i.e. active inference). These two strategies are deeply intertwined and must be understood as belonging to the same cycle of surprise minimization that defines the interactions of a system (Friston 2010, p. 129).¹⁵ For the FEA, interaction acquires the primary role over perception. The system seeks its preservation in its engagements with the environment; however, for these engagements to be unsurprising it requires updated beliefs about the world. What traditional predictive processing developed as perceptual inference is thus subordinated to the active engagements of the system. Just as in perceptual inference, what is predicted is the most likely state of the world, in active inference, what is predicted is the state of the world that would serve the preservation of the system through a motor command that will bring about this state. In the visual tracking example mentioned above, the result of the inferential process is the movement of our head directed towards the predicted future position of the bird.

An important claim within the FEA is that the generative model from which the system generates its predictions is the biological system itself. Furthermore, the biological system as a whole is taken to model or represent the environment, as can be seen in the following claim by Friston and Stephan: “adaptive systems (...) represent the state and causal architecture of the environment in which they are immersed. Conversely, this means that causal regularities in the environment are transcribed into the system's configuration” (Friston and Stephan 2007, p. 426). The proponents of the FEA emphasize that a system represents the causal architecture of the environment and that this architecture is embedded in the structure of the system. In line with this, Clark

¹⁵ In this context, the subordination of perceptual inference to active inference is such that some have proposed to take the two strategies of surprise minimization as constitutive of active inference (see e.g. Friston et al. 2012).

claims that: “(...) the free-energy minimizing ‘model’ that does the mirroring is in fact the whole embodied, active organism” (Clark 2017a, p. 6; see also Clark 2017b).

Friston, for instance, characterizes the generative model as hierarchical, nonlinear and dynamic features that it inherits from the world and that are reflected in the model according to the timescale of the phenomenon represented (Friston 2012, p. S 174). The generative model codes external changes based on their different timescales, distinguishing thus between: (i) rapid environmental changes caused by “structural instabilities or other organisms”; (ii) changes that happen over seconds such as illumination or “slowly varying fields”, i.e., contextual changes that are state dependent, in that they are expected given previous states or cues in the environment; and, finally, (iii) long-term changes encoded in the generative model, that might correspond to “an invariant aspect of [the environment’s] causal architecture” which corresponds to “physical laws and structural regularities” (Friston and Stephan 2007, p. 429).

It is important to note that the generative model posited by the FEA is not agent independent. As Williams (2018) shows, the generative model does not track the world in objective terms. Rather, it tracks the environment as it appears to the agent, in agent-dependent terms. Williams claims that the generative model is a recollection of “the environment as it *matters* to the organism and its physical integrity” (Williams 2018, p. 158). Given that the generative model serves the agent’s ultimate goal which is surprise minimization, both agent- and interest-dependent. The agent is attuned to relational information that is useful to guide the interactions of the agent (Williams 2018, pp. 166–167). This point is also emphasised by Bruineberg et al. (2018) who note that the FEA’s description of the biological agent is not that of a system that advances perfect predictions or predictions that correspond to an unbiased and objective perspective of the environment. I come back to a more detailed description of their perspective in the following section. For now, it is important to keep in mind that, for the FEA, the generative model is an agent specific model of the environment. Furthermore, what is recapitulated in the generative model is not an objective environment, but the environment as it is significant for the survival of the biological system, the system’s *Umwelt* (Williams 2018; see also Clark 2013, 2017b).

3.1 The representational profile of the free-energy approach

Representational language is pervasive in the FEA’s literature. If the FEA were to be characterized as a representational framework, then the marriage with SMEn becomes considerably more difficult to justify. In this section, my aim is to show why the FEA might be understood in representational terms to show in the following sections that, even if the FEA were to be interpreted in these terms, it can still be used to explain sensorimotor knowledge. With this in mind, I build the strongest case possible for the representational reading of the FEA to show that the two frameworks are compatible. So, my aim here is not to defend the representationalist interpretation of the FEA. For that reason, the arguments presented here should not be taken to be knock-out arguments in favor of this reading.

So, why does the FEA deserve the ‘representational’ label, in the first place? After all, the relation between environmental changes that are encoded in the generative model and the generative model could be merely a causal correlation. If this is the case, some might find it worrisome. Ramsey, for instance, considers that if the term refers

only to the activation of an internal state of a system given an input that “plays a mediating role in the processing” (Ramsey 2015, p. 10), then an extremely deflationary notion of representation is on the table and, in consequence, it becomes trivial. In order to identify cognitive structures that are truly representational, Ramsey (2007) advances what he calls the job description challenge. A theory is expected to provide an account of the essential features in virtue of which a cognitive structure can be considered a representation. It is necessary to provide “a job description that tells us what it is for something to function as a representation in a physical system” *qua* representation (Ramsey 2007, p. 27). Meeting the job description challenge is sufficient for a structure to be described as a representation.

So, is the FEA’s use of the term representation ‘trivial’? Gładziejewski (2016) offers a good case for the representational profile of predictive processing, including the FEA.¹⁶ He claims that predictive processing meets the job description challenge advanced by Ramsey. To make his case, Gładziejewski compares the functional features of the generative models posited by predictive processing with the functional features of cartographical maps, a paradigmatic case of representations.¹⁷

Gładziejewski proceeds as follows: if generative models share with cartographical maps the functional features that allow the latter to fulfil a paradigmatic representational role, then the representational profile of predictive processing is justified. These features are i. action-guidance, ii. (partial) detachability, and iii. error detection. Gładziejewski’s claim is not that generative models display some functional features of a paradigmatic case of representation. Rather, his claim is that the generative models display all three functional features (Gładziejewski 2016, p. 570). For this reason, Gładziejewski takes his position to be provisionally immune to anti-representational arguments that seek to trivialize the use of representations in the framework.

Just like cartographical maps, generative models track certain structural features of the environment. Generative models can be characterised as structural representations (S-representations) because, as shown above, they track the causal structure of what is represented in agent-dependent terms.¹⁸ Furthermore, it is in virtue of sharing the structural features of an agent-dependent environment that generative models can fulfil the role they do. The relation of similarity between the generative model and the world

¹⁶ In these sections, I refer mostly to the work of Gładziejewski (2016) and Williams (2018). Their arguments refer to predictive processing in general and consider the FEA within it. For that reason, the distinction between predictive processing and the FEA becomes blurry in this section. Nonetheless, the arguments are presented in light of the FEA as a theory about the adaptive behavior of biological systems and predictive processing as a possible implementation of the approach.

¹⁷ This is what Gładziejewski calls the ‘compare-to-prototype’ strategy—the argumentative strategy recommended by Ramsey (2007). The idea is that generative models posited by predictive processing and cartographical maps share some functional features. These structures, however, do not need to feature the exact same characteristics. For instance, cartographical maps are allocentric representations of space, and they do not guide an agent by providing possibilities of action, among other differences (Gładziejewski 2016, p. 569). Importantly, Gładziejewski notes, the paradigmatic representation and the posited representation might exemplify these features in a different manner (Gładziejewski 2016, p. 564). These differences should not break the functional analogy that is established in this argument. Featuring the distinctive functional characteristics presented by the cartographical map warrants the label of representation for the generative model.

¹⁸ See also Williams and Colling (2018) for a discussion on iconic representations. The authors define these as representations that fulfil their functional role “through their resemblance to their target” (Williams and Colling 2018, p. 1951). As Williams (2018) notes, in the case of predictive processing, the idea is roughly that the representation and its target can be described by the same model.

is, for Gładziejewski, an exploitable relation (Gładziejewski 2016, pp. 566–567). The functional features attributed to S-representations are due to the structural similarity between the representation and the represented.

Let us start with i. action-guidance, the first functional feature of S-representations. This kind of representations serves to guide the actions of an agent. Maps, for instance, allow an agent to move from one point to the other. An agent knows how to find her destination by following a map (Gładziejewski 2016, p. 567). Generative models also exhibit this functional feature. On the basis of predictions advanced by the model, the agent minimizes surprising dealings with the environment and achieves viability. With the notion of active inference, the FEA emphasizes the active character of predictions: the agent advances, as a prediction, a command for acting in one way or the other. At this point, it is important to recall that perceptual inference is subordinated to active inference (Gładziejewski 2016, pp. 574–575).

Secondly, S-representations are ii. detachable. They can fulfil their role (i.e. guiding actions) without the represented being present or available. Detachability can be either complete or partial. Think back again to the case of the cartographical map. The agent guides her actions by following the map. Although her destination is not present, the information about the location of her destination is made available by the map (Gładziejewski 2016, pp. 567–568). The agent decides on the path to take based on the map. It is in this sense that these are detachable. In a similar way, the agent's interaction is mandated on the basis of the generative model. Predictions advanced as a result of the predictive process are anticipations: non-actualized states of affairs. However, detachability, in both cases, is not complete but partial. Although a route is decided based on the map, it is verified in the world: either one gets to the intended destination or not. Similarly, the prediction advanced on the basis of the generative model meets actual sensory feedback and is advanced as an anticipation of what will be encountered in the world. Given that predictions are verified by the interactions of the agent, generative models are attuned to the environment and never fully detached (Gładziejewski 2016, p. 577).¹⁹

Finally, S-representations allow iii. error-detection, i.e. the detection of misrepresentations. It is possible for the agent to detect an error in the representation when she encounters a practical failure, e.g. when a map leads to unsuccessful actions (Gładziejewski 2016, p. 568). Generative models exhibit this functional feature as well because predictions advanced by the model are tested constantly. When the model delivers an inaccurate prediction, the interactions of the agent will be unsuccessful. Furthermore, error prediction is an essential feature of the prediction process. The generative model is updated on the basis of error feedback. Error detection plays a fundamental role: the model is updated to achieve more accurate predictions and reduce the possibility of error (and of surprising dealings) over time (Gładziejewski 2016, p. 579).

¹⁹ According to Gładziejewski, generative models are characterized by partial detachability, in part, because they have been mainly used to explain cases of online cognition. So, it is an open question whether they can be fully detached from the environment, that is, whether they can be used to explain cases of offline cognition. For Gładziejewski, one step in this direction is precisely the model advanced by Friston and by Seth, according to which generative models include coding of counterfactuals. These models involve “counterfactual sampling” of non-actualized actions and could be thought of, in that specific sense, as completely detached. (Gładziejewski 2016, pp. 576–577).

The representational interpretation of the FEA (and of predictive processing) has been contested in the literature.²⁰ And it is possible to provide a non-representational alternative reading to the functional features presented by Gładziejewski. Bruineberg et al. (2018), for instance, offer an account of error detection in terms of dis-attunement. To make their case for a non-representational interpretation of the FEA, they start by emphasising that the FEA should not be conflated with predictive processing. While predictive processing is a strategy advanced to model neural activity, the Free-Energy principle is advanced to understand the behavior of a biological system. The same features that are attributed to the biological system are taken to explain cognition, broadly construed. For Bruineberg et al., the FEA offers a formalization of self-maintenance, an essential feature of biological systems from the perspective of enactivism, according to which biological systems bring forward their own boundaries in their interactions with the environment. These interactions are ordered to the system's viability (Bruineberg et al. 2018, pp. 2420–2423).²¹

The FEA approach mandates a view of the biological system as one that is attuned with its environment and that, in order to maintain viability, requires a generative model that “embodies in its structure and organisation longer-term regularities between action, environment and the state of the organism” (Bruineberg et al. 2018, p. 2425). The environment is reflected in the generative model as it is affectively significant for the biological system, i.e. in terms of what the biological agent cares about. Furthermore, if we stick to the FEA's perspective, it is the biological system (and not the brain in isolation) that should be considered as the system in question. This notion of the generative model is advanced in opposition to a perspective on the generative model that takes it to be an objective and agent-independent representation of the environment. The latter is, for Bruineberg et al., the notion of the generative model that is in place in Bayesian models of the brain.

Bruineberg et al. (2018) take issue with representational readings of the FEA and, particularly, with the link between the representational interpretation of the FEA and the Bayesian brain. The reason is that the latter recommends a view of a brain that behaves as “an exemplary scientist”, inferring from representational states the hidden causes of sensations, and not as a system that is biased in its recollection of the environment and its relation to it. Bruineberg et al. (2018) thus offer an alternative non-representational reading of the framework.

With the presentation of Gładziejewski's arguments I do not aim to show that a representational reading of the FEA is mandatory. However, if the concern of the non-representationalist were to relate more specifically to the positing of representations that are action-neutral, the position sketched has the tools to address the worry. As Williams (2018) has shown, the representational interpretation of the FEA can address this worry

²⁰ Hutto and Myin (2017) argue that although interpreting the framework in this way comes naturally, the proponents of the FEA must face the gaps left by taking such stance. If the FEA posits contentful representations, it must explain where the contents of such representations come from (Hutto and Myin 2017, pp. 154–155). I am not addressing this discussion in more detail because Hutto and Myin's is an objection that is general to representational frameworks, and not specific to the predictive processing and FEA debate. For an account of the contents of representations within predictive processing, see e.g. Wiese (2017) and Kiefer and Hohwy (2018).

²¹ For an overview on the FEA's notion of autopoiesis as self-maintenance and how it compares to enactivism's notion of self-production as developed mainly by Humberto Maturana and Francisco Varela, see Kirchhoff (2018).

because the commitment to representationalism does not imply a commitment to the positing of objective representations that stand for an agent-independent environment. This is also recognised by Gładziejewski: generative models are not an objective mirror of the causal structure of the environment. They rather recapitulate the causal relations as they are relevant for the biological system (Gładziejewski 2016, p. 571; see also Wiese 2017). These structures are ultimately at the service of the system's coping with the environment (Williams 2018).

Nico Orlandi (2014, 2016) has also contested the representational interpretation of predictive processing, in general.²² In particular, she objects to an understanding of vision as an inferential process. Her point is that visual processing does not consist in “transitions between representational states” (Orlandi 2014, p. 18). To show this, she argues that the states that are involved in the inferential process that constitutes the Bayesian view of vision should not be understood as representations. For Orlandi, the only key concept within predictive processing that can be described representationally is that of perceptual hypotheses, i.e. the result of the process, and not the process itself. Hypotheses concern distal or absent conditions, misrepresent and can be described as performance-guiding states.²³

Orlandi notes firstly that when it comes to early stages of visual processing and to intermediate levels of the cascade of prediction that lead to perceptual hypotheses, characterizing these states as representations becomes difficult. After all, these states (i.e. error- and prediction-signals) do not concern distal or absent causes at all, but the immediate lower level or what is immediately present to the senses (Orlandi 2016, p. 344). On the other hand, in what concerns the generative model, i.e. priors and likelihoods, Orlandi claims that similar considerations apply. Priors and likelihoods are not about distal conditions such as the causal structure of the environment. Rather, they are about other states at lower levels of the generative model.

There are a couple of things to consider in connection to Orlandi's objection. Firstly, her objections are directed at the Bayesian considerations of the brain rather than at the FEA as a story about the adaptive behavior of biological systems. One of the reasons she gives to inscribe predictive processing within ecological approaches to visual perception is that Bayesian approaches do not provide a story of how the space of probability from which predictions are advanced is bounded. Importantly, this is precisely the story offered by the FEA.

Secondly, the FEA's claim about the representational nature of generative models is that the structure of the generative model represents the causal structure of the environment. The generative model is constituted by variables that represent not only likelihoods and prior probabilities, but also the dynamic relations between variables at different levels. The proposal of the FEA is that the relation between the different levels, i.e. the structure of the generative model, can be described by the same model that describes the system's environment (see Wiese 2017, pp. 725–726). So, Orlandi might be right in that the process is not inferential because it does not involve transitions between

²² It is worth noting that Orlandi advances a proposal that opposes the view defended throughout the paper. For her, while early visual processing should not be characterized as representational, the end-result of visual perception should be accounted for in representational terms (see Orlandi 2014).

²³ For a critical overview of Orlandi's position, see Downey (2018).

representational states. However, that argument does not address the point made by the proponents of the representational reading of the FEA that the generative model, as a whole, recapitulates the structure of the system's environment.

The arguments just sketched show how the representational account of the FEA can respond to the worries that drive anti-representational interpretations of the framework. To that extent they support the plausibility of the representational reading of the FEA. Nonetheless, it is important to note that these arguments are not conclusive since they do not show that the representational interpretation is mandatory. More would need to be said to make a decisive case for either a representational or a non-representational reading of the FEA. This has not been the purpose of this section. Again, for the purposes of arguing for the viability of a predictive approach to SMEn, it is important to take the plausibility of the representational profile of the FEA seriously, given that this represents an important obstacle for the joint account. If we can claim that SMEn is compatible with such an account, arguing for the viability of a non-representational predictive approach to SMEn comes easily.

4 A representation-friendly version of sensorimotor enactivism

Given the features of the FEA, including its representational profile, is a predictive approach to SMEn viable? With respect to the representational profile of the FEA, there are two questions that need to be addressed. Firstly, is there room for representations in SMEn? And secondly, if there is, are the representations posited by the FEA acceptable for SMEn? In the following section, I address the first question.

4.1 Perception: A representation-friendly domain

There is one good reason to think that SMEn's model of perception posits representations. Consider sensorimotor knowledge. As mentioned earlier, it is practical knowledge defined precisely as our "capacity to respond (...) with sensitivity" to the systematic ways sensory input changes with movement (Silverman 2016, p. 283; see also Silverman 2018). The agent invokes non-current states of affairs that are not readily available in the environment to make sense of or to successfully engage with the object of perception. Furthermore, this explains the qualitative aspects of perceptual experience.

The description of sensorimotor knowledge as a capacity to respond to absent features can be assimilated to Clark and Toribio's (1994) prominent description of representation-hungry domains, domains that require the positing of representations in order to be explained. Representation-hungry domains are those in which: "[t]he problem requires the agent to be selectively sensitive to parameters whose ambient physical manifestations are complex and unruly (for example, open-endedly disjunctive)". Clark and Toribio add that a domain is representation-hungry if it requires sensitivity to "states of affairs whose manifestation in the sensory inputs is (...) attenuated" (Clark and Toribio 1994, p. 419).

Now, given that sensorimotor knowledge invokes absent aspects of the object of perception, it becomes plausible to describe object perception as a representation-

hungry domain and, in consequence, as a domain that *must* be explained representationally.²⁴ However, this description of sensorimotor knowledge commits SMEn to what Silverman (2018) calls constitutive representationalism, the view that perception is constituted by representations. Silverman relies on John McDowell's distinction between *constitutive* and *enabling* features of perception (see McDowell 1994). While the latter are features that "play a causal role" in perception, the former are defining features of perception: these are just what perception is. This kind of representationalism claims that cognition (broadly construed) "could not occur unless there was internal representation" (Silverman 2018, p. 168). The problem, as Silverman sees it, is that constitutive representationalism goes one step too far. The description of perception as a representation-hungry domain becomes problematic because it risks losing the practical aspect of sensorimotor knowledge and mandates the positing of representations.

Silverman's (2018) suggestion is to consider these domains as *prima facie* representation-hungry. That way, SMEn is not committed to constitutive representationalism but to an enabling kind of representationalism that does not *require* representations. *Prima facie* representation-hungry domains are those in which some "theoretical work (...) could be done by internal representation" (Silverman 2018, p. 168). According to this view, then, representations could be posited to account for perceptual processing, without these constituting perceptual experiences.

Silverman's suggestion becomes more plausible when we consider that the idea that representation-hungry domains *require* representations has been challenged in the literature. Wheeler (2005), for instance, shows that domains that are thus described can be explained using strategies that do not posit representations. The notion of representation-hungry domains has also been challenged from the trenches of enactivism. Kiverstein and Rietveld (2018) show that the notion of representation-hunger relies on the idea that cognitive systems are, at times, decoupled from the environment. Against this idea, they claim that what allows the agent to deal with the absent and the abstract is the agent's skillful coordination with the environment and the affordances it offers. Finally, Degenaar and Myin (2014) show that neither cognitive activities in domains that involve something absent, nor domains that involve something abstract *necessitate* representations.

These arguments show that cognitive domains that are described as representation-hungry domains do not *require* representations. Rather, this indicates the plausibility of a representational explanation to account for the domain in question. *Prima facie* representation-hungry domains can be explained following both a representational and a non-representational strategy. Positing representations is just one of the strategies to account for problem domains that exhibit the features previously described. So, given the description of sensorimotor knowledge advanced by SMEn, the theory can, in principle, be described as representation-friendly. One possible strategy to explain that sensorimotor knowledge involves invoking non-current states of affairs is of a

²⁴ Downey (2017) also points out that SMEn's description of notions like attention and sensorimotor knowledge is such that it seems like they should be explained in representational terms (Downey 2017, pp. 2–3). Myin and Loughlin (2018), on the other hand, consider that although SMEn is committed to the rejection of representations for explaining perceptual experience, the view might posit representations for other theoretical purposes (Myin and Loughlin 2018, p. 212).

representational nature. The notion of sensorimotor knowledge allows representational strategies. However, these are not mandatory.²⁵

In the next subsection, I show that the representations posited by the FEA comply with the conditions advanced by SMEn. Firstly, I argue that representations are not, for the FEA, the end-result of perception. I then show that the generative model complies with the conditions advanced by SMEn. The FEA posits representations that are not neutral representations of the scene perceived, that are indicative of further possible contextual interactions, and that are agent-specific.

4.2 Are the representations posited by the free-energy approach adequate for sensorimotor enactivism?

It is important to begin by noting that although the cycle of surprise minimization in the FEA's model does imply feeding and adjusting a generative model, the creation and subsequent adjustment of said model is not the end-result of perception. Firstly, this is because perceptual inference (i.e. the adjustment of the model) belongs to a cycle of surprise minimization that has as its ultimate goal delivering successful interactions between the agent and the environment. The adapted biological system is itself a model of the world that enables its engagement with its her surroundings. Perception contributes to the adequate engagement of the agent with the environment which allows for its viability. So, it is difficult to claim that the end-result of perception, from this perspective, is the updating of the model.

At this point, a distinction advanced by Kirchhoff and Kiverstein (2019) is relevant, namely the distinction between the *generative model* and the *generative process*. As mentioned earlier, for the FEA it is the whole agent that comprises the generative model. As Kirchhoff and Kiverstein claim, the generative model is not something that can be abstracted away from the biological agent as a whole (Kirchhoff and Kiverstein 2019, p. 57). In active inference, the agent generates its own constitution and maintains itself as a self-model. While the generative model “selects (...) ways of acting”, the *generative process* brings forth or enacts those actions. The generative process corresponds to the dynamic coupling with the environment that “enables an agent to reduce uncertainty” over time (Kirchhoff and Kiverstein 2019, p. 57). It is through the generative process that the agent brings about the states that serve its preservation.

So, is the generative model acceptable for SMEn as a representation? In the case of the generative model, talk of representations should not be problematic. The generative model taken as a representation complies with the conditions advanced by SMEn, as identified earlier. Firstly, consider that part of the information coded in the generative model concerns the sensory stimuli that is expected to follow an interaction. The generative model is thus indicative of further interactions as it presents the agent with possibilities of action. The generative model is agent specific in that it is in part the embodied agent itself that constitutes said model and in that the neural implementation of the generative model recollects the environment as it is significant for the system.

²⁵ It is worth noting that this strategy of distinguishing *enabling* from *constitutive* representationalism still faces the challenges generally advanced against representationalism in connection to content ascription (Hutto and Myin 2012, 2017).

Furthermore, the generative model is not neutral to the agent's history of interactions, on the contrary, it is enriched by learning.

Finally, in connection with the description of SMEn's model of perception as a *prima facie* representation-hungry domain, it is important to consider the following. The FEA admits both representational and non-representational strategies to explain these domains. The sensitivity to the environment that is necessary to account for perceptual experience relies on the generative model, which is implemented both in the embodied structure of the agent and in neural activity. In consequence, the distinction between representational and non-representational structures does not correspond to a distinction between the brain and the body. Moreover, from the perspective of the FEA, these two explanatory strategies (i.e. relying on the embodied constitution of the agent and her interactions, and positing representations) are not incompatible.

As can be seen, there is nothing about the FEA's take on representations that should concern SMEn and, in consequence, that would indicate that a predictive approach to SMEn is not viable. In the following section, I present Seth's predictive approach to SMEn to show how these frameworks fit together. This discussion will allow me to address the second worry raised to a predictive approach to SMEn, i.e. the concern that the explanatory strategy recommended by SMEn conflicts with the FEA.

5 The predictive approach to SMEn

Anil Seth proposes to account for SMCs, and the mastery thereof, by means of the generative model. It is important to bear in mind that, according to Seth, for predictive processing, neural mechanisms are necessary and sufficient to enable perceptual experience because the content of perceptual experience is determined by the generative models (Seth 2014, p. 98).

Seth's proposal is to integrate SMCs into the generative models, thus providing the causal story that SMEn is missing (Seth 2014, p. 98). It is important to recall that SMCs are the lawful relations that result from the interaction between the agent and the object of perception. Perception, for SMEn, is grounded in the practical knowledge of SMCs. Incorporating the theory of SMCs comes naturally for the FEA because the generative model includes the specification of stimuli that follow a given movement (Friston et al. 2012). It is in that sense that Seth takes SMCs to be coded in the generative model. He claims that the generative model encodes counterfactual "perception-action couplings" that result in active inference. According to this approach, then, the agent deploys a set of SMCs as a result of top-down inferential processing. At this point, it will be useful to recall the visual tracking example. In this case, the movement of the eyes and head as the trajectory of the bird is followed results from the prior belief about the position of the bird and generates a posterior belief from which a new inference—hence, a new interaction—results (adapted from Friston et al. 2012, p. 15; Seth 2014, p. 104).

The *mastery* of SMCs, for Seth, relates to the "conditional aspect" of perception, i.e. to the counterfactual knowledge of, for instance, how the appearance of a tomato would change if I were to turn it around. This mastery is attributed to the hierarchical structure of the brain. Seth claims that this can be reflected in the layered structure of generative models: predictions are generated from a cascade of top-down processes that are associated with more abstract predictions (Seth 2014, pp. 102–103).

Seth's proposal provides SMEn with a story about the neural mechanisms involved in perception. In this way, his account fills in the gap left by the question of the brain's role in perceptual experience. So, is there anything within Seth's view that should make us question the viability of the predictive approach to SMEn? At this point, it is important to recall that, according to Seth, the neural implementation of the generative model suffices to account for perceptual experience. And, in consequence, in Seth's model we find a shift in the explanatory strategy. He abandons SMEn's explanatory strategy of prioritizing the interactions of the embodied agent over the inner structure of the agent. Seth's account leads to the second aspect of conceptual tension between SMEn and the FEA and, thus, falls prey to the second concern regarding a predictive approach to SMEn, i.e. that this joint account is not viable in virtue of the internalist commitments of the FEA according to which the brain is necessary and sufficient to account for perception.

In the following subsection, I put pressure on Seth's assumption that the generative models are necessary and sufficient to account for perceptual experience.

5.1 Against the sufficiency of the brain to explain perception

It is worth recalling that, from the perspective of predictive processing, while the updating of the generative model is a result of the agent's current interactions with the world, the shape of the generative model is due to the agent's history of interactions. For Seth, the coding of SMCs in the generative model is sufficient to explain perceptual experience. Di Paolo (2014) takes issue with this. According to him, there are also cases of interactions that constitute SMCs, despite neither being the current interactions of the agent, nor appearing in the agent's repertoire of possible interactions that conform to the neurally implemented generative model.

To understand Di Paolo's objection, we must note firstly that he distinguishes between a virtual structure of interactions from actual interactions. So, for instance, going back to the first example about our experience when observing a tomato, we must distinguish the interaction currently taking place from other possible interactions that are implied in the current interaction: from the way the tomato looks when we hold it in one hand, certain interactions are implied, for instance, that we could turn the tomato around and see its other side. Di Paolo refers to the former as the actual structure of interactions and to the latter as the virtual structure of interactions.²⁶

For Seth, the virtual structure of interactions is implemented by the generative model and this should suffice to account for the mastery of SMCs and, therefore, to account for perceptual experience. For Di Paolo, on the other hand, generative models require regular feedback from the world because they require regular updating to maintain a certain level of accuracy. For him, the brain is insufficient because the world is necessarily involved and already constitutes, in part, the virtual structure of interactions.

²⁶ From the perspective of enactivism, the environment is constituted by the virtual interactions that are available to the agent and that are actualized as the agent interacts with its surrounding. Kiverstein and Rietveld (2018), for instance, take the agent's environment to be virtual precisely because it is enacted by the interactions of the agent. The virtual structure of interactions are the states a system could occupy given its tendencies and current interactions. For them, the virtual conditions that constitute a system's environment are "the set of states of the agent-environment system that 'surround a currently actualized trajectory of states' (Di Paolo et al. 2017)" (Kiverstein and Rietveld 2018, p. 152).

The virtual structure of interactions is not exhausted by the generative model. In the interactions of the agent with the world, there is already a structure of possible states that can be visited by her. Some possibilities of action are available to the agent not in virtue of the counterfactual structure of the generative model, rather they are implied in the agent's engagement itself. So, even if the generative model can explain the agent's *mastery* of SMCs, SMCs cannot be exhaustively accounted for by referring to the generative model. The possibilities of action offered by the generative model are constrained by the agent's current interaction and the virtual structure of possibilities that emerges from it.

To illustrate this, Di Paolo takes a case of dynamical operationalizations of SMCs: "if I walk on a slope there is a strong downward tendency for my movement, even when I walk uphill. This is real and does not depend on my having enough sensitivity to detect it. Most "nearby" trajectory options are implied in the enacted movement" (Di Paolo 2014, p. 2). In this case, Di Paolo claims, the virtual structure is entailed by the movement itself. The movement of walking uphill on a slope entails a virtual structure that constrains the possibilities of actions to which we are sensitive. The neural generative model does not suffice to explain why *these* possibilities of action are available and not others.²⁷

Di Paolo objects to Seth's claim that the neural generative model is sufficient to account for perceptual experience. For him, the generative model is insufficient to account for the virtual structure of interactions to which the agent is sensitive. Moreover, the generative model does not suffice to explain *how* the world contributes when updating the generative model: the interactions of the agent already constrain the virtual structure of interactions. Di Paolo's claim is that the available sensorimotor interactions exceed those that appear in the generative model.

This objection should not lead us to the rejection of Seth's model. And, importantly, nothing in Seth's model prevents him from accepting Di Paolo's proposal of constraining the available sensorimotor interactions by way of the agent's embodiment and interactions.²⁸ Di Paolo's point is compatible with accepting that the mastery of SMCs, i.e. the capacity of the agent to respond with sensitivity to SMCs, might still be supported by the generative model.

Rather, this shows that the conflict that arises between the frameworks is due to their conflicting explanatory strategies. While Seth seems to think that referring to the inner structure of the agent suffices to explain perception, Di Paolo claims that it does not because the virtual structure of sensorimotor interactions to which the agent is sensitive cannot be thus explained. So, in the face of these two conflicting strategies, what arguments does SMEn offer to support its explanatory strategy? And, importantly, can the FEA follow the same explanatory strategy? These are the issues I address in the last section.

²⁷ Flament-Fultot calls this change in the sensorimotor patterns available a change in the "landscape of *possibilia*" (Flament-Fultot 2016, p. 170, italics in original).

²⁸ It might be objected that Di Paolo and Seth rely on incompatible views of sensorimotor knowledge. While Di Paolo is committed to a practical view of sensorimotor knowledge, Seth is committed to a counterfactual reading. However, as I tried to show in the first section, sensorimotor knowledge should be understood as grounding perception, which is an embodied skill. Sensorimotor knowledge is the agent's capacity to respond to possible sensorimotor interactions, and this capacity can be enabled in different ways: by the agent's embodiment or by internal representations.

6 Revisiting the explanatory strategy of sensorimotor enactivism

At this point it is important to recall that, for SMEn, the interactions of the embodied agent enjoy explanatory priority over the inner structure of the agent. This priority should be understood as an explanatory strategy that comprises two claims: (a) the claim for the insufficiency of the neural mechanisms in an account of perception; and (b) the claim that the neural mechanisms of the agent should be understood in virtue of the interactions of the embodied agent. As discussed earlier, this strategy is accompanied by an ontological claim about the vehicles of perceptual experience. The question I am concerned with here is: how does SMEn justify its explanatory strategy?

If SMEn were to rely simply on an ontological claim according to which perceptual experience constitutively consists in the interactions between the embodied agent and the environment, the theory would be vulnerable to an objection from the internalist. The claim falls prey to the causal-constitution fallacy objection advanced by Adams and Aizawa (2009, 2010), according to which SMEn would be conflating factors that are only causally relevant, with factors that are constitutive to the phenomenon in question. By relying on an ontological claim, Clark argues, SMEn confuses evidence for the role of these interactions in training and tuning the system with evidence for their role as vehicles of perceptual experience (Clark 2009, p. 970).²⁹

SMEn justifies its strategy of referring primarily to the agent's embodiment and environment on the basis of an argument for the best explanation. Its explanatory strategy is justified because the best way to account for perceptual experience is by referring to the interactions of the agent. Furthermore, this is the best way to make sense of the agent's inner structure. So, the claim is not only that external elements are constitutive of perceptual experience, but that this provides the best explanation (Clark 2009, pp. 970–971).

This is what Noë (2007) calls explanatory externalism. He claims: “the nervous system modifies experience by *keeping track of our relation to things around us*” (Noë 2007, p. 465, italics in original). To make sense of the nervous system and of neural activity, it is necessary to pay attention to the interactions of the agent. For Noë (2007), the ontological claim that accompanies SMEn's explanatory strategy is justified on the basis of the explanatory strategy, and not the other way around. Perception is constituted by the interactions of the agent (or, as he puts it, by the world) because these are the explanatory substrates, i.e. the realizers that are required to explain perception. He adds: “[n]eural systems are essential for experience, but they are not all that it is essential, and so we cannot understand experience in neural terms alone” (Noë 2007, p. 467).

Susan Hurley (2010) recommends a similar strategy. Hurley advances an argument in favor of externalism about perceptual experience, the view according to which the realizing base of perceptual experience encompasses, in at least some cases, the agent's dynamic interactions with the world. To make her case, Hurley examines cases of internal simulations that mimic sensory events and distinguishes between offline and online cases of simulation. While the former are cases where the system mimics sensory events but is not currently interacting with the environment, the latter, are cases where the system is currently interacting with the environment (Hurley 2010, pp. 139–141). The FEA's model of perceptual processing can be

²⁹ For an overview of the debate, see Kirchhoff and Kiverstein (2019, chapters 2 and 3).

assimilated to online simulation cases because the brain advances a prediction that mimics input coming from a lower level. From the perspective of the FEA, this is done within a cycle of interactions with the environment.

Although in both cases it is undeniable that internal simulations are necessary to account for the simulated sensory event, it is not obvious that the agent's interactions are necessary. Hurley supports her claim that interactions are indeed necessary to account for these cases by arguing for the derivative character of simulations. In online simulations, the simulations' parameters are continuously updated by means of the agent's interactions. The interactions of the embodied agent provide "ongoing tuning and maintenance" (Hurley 2010, p. 142). In the case of offline simulations, on the other hand, the role of the feedback loops that are mimicked results from the agent's developmental history and history of interactions. In consequence, to make sense of both cases of simulations, it is necessary to make sense of the agent's (history of) interactions. Similar arguments are raised by Kirchhoff and Kiverstein (2019). They claim that, when the relevance of the generative process in grounding perceptual experience is recognised, the dependence of simulations on the agent's interactions and the primacy of the latter become evident (Kirchhoff and Kiverstein 2019, p. 60).

There are a couple of arguments that are worth considering at this point, both advanced by Clark (2009). If the claim is that it is necessary to refer to the interactions of the embodied agent with the environment to make sense of the contributions of the brain to perceptual experience, Clark worries that the very same strategy might be advanced by the internalist. Standard internalist views might argue that a mental state is neurally supported in many different ways and that, to distinguish neural contributions, considerations about external elements might be relevant. There is nothing about the argument that supports externalism (Clark 2009, p. 971).

Clark himself notes that the argument of SMEn is slightly different: SMEn provides an account that offers an additional explanatory advantage. It is not only that referring to external elements allows for a distinction between neural structures that are involved in perceptual experience, distinction that would otherwise be ignored. It is also that SMEn provides an answer as to *why* this pattern of interactions supports e.g. visual rather than auditory experience (Clark 2009, p. 972). But why should we think that the internalist does not have the resources to explain this and do so without conceding that the vehicles of perceptual experience extend beyond the boundaries of the brain? For Clark, we should not underestimate the internalist because she might be able to provide an answer to that question.

Clark might be right in that: just like SMEn, some internalist views might be able to provide an answer to this question. Nonetheless, it becomes a matter of providing the best possible explanation. It is no longer a matter that is solved by referring to an ontological claim. SMEn's claim is that in domains in which the contributions of internal and external factors are entangled, the best possible explanation will be provided by an account that encompasses both. In connection with this, Hurley claims that in dynamical explanations, the causal-constitutive distinction is not explanatorily transparent (Hurley 2010, p. 126).³⁰ Furthermore, when we think of this issue in the

³⁰ In a similar vein, Kirchhoff and Kiverstein (2019) challenge the standard causal-constitutive distinction. They argue for a notion of constitution as a diachronic relation. In this way, the distinction between the causal and the constitutive becomes blurry (see Kirchhoff and Kiverstein 2019, chapter 6).

context of the FEA, it becomes difficult to ignore the contributions made by the agent's interactions. This is ultimately the point made by Di Paolo against Seth: for SMEn, to understand the pattern of sensorimotor interactions available to the agent that support perceptual experience, it is not sufficient to look at the neurally implemented generative model. The best explanation available refers to the interactions of the embodied agent.

At this point, the ontological claim has almost become otiose: distinguishing between the causal and the constitutive is not explanatorily informative. It is perhaps for this reason that Hurley briefly entertains the possibility of dropping the distinction altogether. She claims that: "it isn't clear how causal-constitutive talk can be mapped onto complex dynamical explanation, or even what work a criterion of the constitutive is supposed to do in this context. We don't have such a criterion here, and it isn't clear that the cognitive sciences need one in order to provide good explanations" (Hurley 2010, p. 126).³¹

In the context of SMEn, the worry might be that dropping the distinction misses the point of SMEn. To see this, it might be useful to think back to Di Paolo's objection to Seth. It is not only that paying attention to the contributions of the world allows for a better account of perceptual experience. The claim is stronger. The claim is that the virtual structure of interactions determines the content and quality of perceptual experience. That virtual SMCs in part constitute and determine perceptual experience becomes evident when my explanation cannot do without referring to the virtual structure of possible interactions. We come to know about this virtual structure by means of the role it plays in our explanation, because otherwise we would not be able to account for perceptual experience.

So, where does this leave the predictive approach to SMEn? Hurley's take on simulations allows us to understand the priority of the interactions of the embodied agent. Neural dynamics are subordinated to and should be understood within the context of the agent's interactions. Deploying this strategy is consistent not only with SMEn, but also with the FEA. Although the latter is an approach that aims at accounting for neural activity involved in perception, this does not necessarily mean that it takes the brain to be explanatorily sufficient for perception. The FEA is compatible with this reading of the explanatory priority of the interactions of the embodied agent since it offers an account of the biological system that underlies its theory of perceptual processing. It is important to note that claim (b), namely that the neural mechanisms of the agent should be understood in virtue of the interactions of the embodied agent, qualifies how the necessity should be understood: referring to the interactions of the embodied agent is necessary to explain perception in the sense that, even when referring to the inner structure of the agent, the interactions of the embodied agent should throw some light on our understanding of the agent's inner (neural) structure. The FEA is compatible with this reading of the explanatory priority of the interactions of the embodied agent since it offers an account of the biological system that underlies its theory of perceptual processing.

The consequence of this explanatory strategy for the predictive approach to SMEn is the following. The predictive approach to SMEn should be concerned with the role of the predictive brain in perception. However, an account of sensorimotor knowledge and

³¹ For another approach to the causal-constitutive distinction that is in the vicinity, see Loughlin and Zahidi's (2017) and Loughlin's (2018b) wide-view of mentality.

of perceptual experience cannot be exhausted by an appeal to predictive processing and a neurally implemented generative model.

7 Concluding remarks

I have argued that a predictive approach to SMEn is viable. Such an argument is necessary to support the views that bring predictive processing and SMEn together because the viability of a joint account is not straightforward. We have reasons to think that they are incompatible, even though they both share a view of perception as an active and anticipatory phenomenon. Firstly, they seem to belong to different sides within the representation-wars: predictive processing and the FEA are standardly thought to have a representational profile, while SMEn is often said to belong to the non-representational wing of cognitive science. Secondly, they exhibit different explanatory strategies to account for perception: while SMEn prioritizes the interactions of the embodied agent, predictive processing is a framework that aims at understanding neural activity involved in sensory processing.

To argue for the marriage of the FEA and SMEn, I began by discussing their respective position towards representations. I discussed the representational profile of the FEA at large. The aim was not to make a case for this interpretation of the FEA, but to show that, even when the FEA is read in this way, a marriage with SMEn is justified. In order to show the compatibility of SMEn with the more challenging interpretation of the FEA for SMEn, I departed from the usual reading of SMEn as a view that is non-representational. I defended an unorthodox account of SMEn that is representation-friendly and that does not rely on an ontological claim to justify the way it proceeds. The FEA can be made compatible with the features of this account. Firstly, because if we take the generative model as representing the world, it turns out to be a representation that complies with the constraints fixed by SMEn: it is a non-neutral representation, it is indicative of further interactions, and it is specific to the agent. Furthermore, updating this model is not the end-result of perception.

Now, the claim that referring to the interactions of the embodied agent is necessary to account for perception also posed a challenge for the compatibility of these theories. The reason for this was the assumption reflected in Seth's model that, for predictive processing in general, the brain is the realizing base of perception, thus resulting in an explanatory strategy that considers that an account of the neural mechanisms involved in perception is necessary and sufficient to account for this phenomenon.

For SMEn it is necessary to refer to the interactions of the embodied agent to account for perception because it is in virtue of these interactions that the inner structure of the agent can be best understood. Here I follow Hurley and Noë's explanatory externalism in claiming that SMEn justifies this explanatory strategy by means of an argument to the best explanation. Moreover, I show that the FEA is compatible with such an approach. For the FEA, the dynamics instantiated in the brain are part of a wider cycle of surprise minimization that aims at the preservation and viability of the agent. The FEA is compatible with a view according to which the brain is not explanatorily sufficient to account for perceptual experience.

This paper leaves the door open to further research, in particular, on the question about the incorporation of the predictive brain into a general theory of SMCs. In this

paper I concluded that a predictive approach to SMEn is viable. However, more work is needed to claim that the FEA is the best or the only causal account for SMEn.

Acknowledgments I am grateful to Michael Wheeler, Alisa Mandrigin and two anonymous referees for their very helpful comments on the various drafts of this paper. I would also like to thank audiences at the Open Minds XIII conference at the University of Manchester, the Models of the Mind conference at the University of Edinburgh and the Philosophy of Mind and Phenomenology Summer School at the University of Copenhagen for their helpful feedback.

Funding information This work was supported by Mexico's National Council of Science and Technology (Consejo Nacional de Ciencia y Tecnología, CONACYT) under PhD scholarship 295962/440722.

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

References

- Adams, F., & Aizawa, K. (2009). Why the mind is still in the head. In A. Murat & P. Robbins (Eds.), *The Cambridge handbook of situated cognition* (pp. 78–95). Cambridge: Cambridge University Press.
- Adams, F., & Aizawa, K. (2010). *The bounds of cognition*. Sussex: Wiley-Blackwell.
- Bar, M. (2009). The proactive brain: Memory for predictions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1235–1243. <https://doi.org/10.1098/rstb.2008.0310>.
- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese*, 195, 2417–2444. <https://doi.org/10.1007/s11229-016-1239-1>.
- Buhrmann, T., Di Paolo, E., & Barandiaran, X. (2013). A dynamical systems account of SMCs. *Frontiers in Psychology*, 4, 1–19.
- Cappuccio, M., & Wheeler, M. (2012). Ground-level intelligence: Action-oriented representation and the dynamics of the background. In Z. Radman (Ed.), *Knowing without thinking. mind, action, cognition, and the phenomenon of the background* (pp. 13–36). Basingstoke: Palgrave Macmillan.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge: The MIT Press. Bradford Books.
- Clark, A. (2009). Spreading the joy? Why the machinery of consciousness is (Probably) still in the head. *Mind*, 118(472), 963–993. <https://doi.org/10.1093/mind/fzp110>.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–253. <https://doi.org/10.1017/S0140525X12000477>.
- Clark, A. (2016). *Surfing uncertainty*. Oxford: Oxford University Press.
- Clark, A. (2017a). How to knit your own Markov blanket. Resisting the second law with metamorphic minds. In *Philosophy and predictive coding* (Vol. 3). Frankfurt am Main: MIND Group. <https://doi.org/10.15502/9783958573031>.
- Clark, A. (2017b). Bursting out: Predictive brains, embodied minds, and the puzzle of the evidentiary veil. *Noûs*, 51(4), 727–753. <https://doi.org/10.1111/nous.12140>.
- Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese*, 101 (Connectionism and the frontiers of artificial intelligence) (pp. 401–431).
- Clark, A. (1997). *Being There. Putting Brain, Body, and World Together Again*. Cambridge, MA: The MIT Press. Bradford Books.
- Degenaar, J., & Myin, E. (2014). Representation-hunger reconsidered. *Synthese*, 191, 3639–3648.
- Di Paolo, E. (2014). The worldly constituents of perceptual presence. *Frontiers in Psychology*, 5, 450. <https://doi.org/10.3389/fpsyg.2014.00450>.

- Di Paolo, E., Buhmann, T., & Barandiaran, X. (2017). *Sensorimotor life: An enactive proposal*. Oxford: Oxford University Press.
- Downey, A. (2017). Radical sensorimotor enactivism & predictive processing. Providing a conceptual framework for the scientific study of conscious perception. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing* (Vol. 11). Frankfurt am Main: MIND Group. <https://doi.org/10.15502/9783958573123>.
- Downey, A. (2018). Predictive processing and the representation wars: A victory for the eliminativist (via fictionalism). *Synthese*, 195, 5115–5139.
- Flament-Fultot, M. (2016). Counterfactuals versus constraints. Towards an implementation theory of sensorimotor mastery. *Journal of Consciousness Studies*, 23(5–6), 153–176.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>.
- Friston, K. (2009) The free-energy principle: a rough guide to the brain?. *Trends in Cognitive Sciences* 13 (7), 293–301
- Friston, K. (2012). Embodied inference and spatial cognition. *Cognitive Processing*, 12(Suppl 1), S171–S177. <https://doi.org/10.1007/s10339-012-0519-z>.
- Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3), 417–458. <https://doi.org/10.1007/s11229-007-9237-y>.
- Friston, K., Adams, R. A., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: Saccades as experiments. *Frontiers in Psychology*, 3, 1–20. <https://doi.org/10.3389/fpsyg.2012.00151>.
- Gładziejewski, P. (2016). Predictive coding and representationalism. *Synthese*, 193(2), 559–582. <https://doi.org/10.1007/s11229-015-0762-9>.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Hurley, S. (1998). *Consciousness in action*. Cambridge: Harvard University Press.
- Hurley, S. & Noë, A. (2003) Neural Plasticity and Consciousness. *Biology & Philosophy*, 18, 131–168.
- Hurley, S. (2010). The varieties of externalism. In R. Menary (Ed.), *The extended mind* (pp. 101–154). Cambridge: The MIT Press. Bradford Books. <https://doi.org/10.7551/mitpress/9780262014038.001.0001>.
- Hutto, D., & Myin, E. (2012). *Radicalizing enactivism. Basic minds without content*. Cambridge: The MIT Press.
- Hutto, D., & Myin, E. (2017). *Evolving enactivism. Basic minds meet content*. Cambridge: The MIT Press.
- Kiefer, A., & Hohwy, J. (2018). Content and misrepresentation in hierarchical generative models. *Synthese*, 195(6), 2387–2415. <https://doi.org/10.1007/s11229-017-1435-7>.
- Kirchhoff, M. (2018). Autopoiesis, free energy, and the life-mind continuity thesis. *Synthese*, 195(6), 2519–2540. <https://doi.org/10.1007/s11229-016-1100-6>.
- Kirchhoff, M., & Kiverstein, J. (2019). *Extended consciousness and predictive processing. A third-wave view*. London: Routledge.
- Kiverstein, J. (2018). Free energy and the self: An ecological-enactive interpretation. *Topoi*, <https://doi.org/10.1007/s11245-018-9561-5>.
- Kiverstein, J., & Rietveld, E. (2018). Reconciling representation-hungry cognition: An ecological-enactive proposal. *Adaptive Behaviour*, 26(4), 147–163.
- Kiverstein, J., Miller, M. & Rietveld, E. (2019). The feeling of grip: novelty, error dynamics, and the predictive brain. *Synthese*, 196 (7), 2847–2869.
- Loughlin, V. (2018a). Sensorimotor theory, cognitive access and the ‘absolute’ explanatory gap. *Phenomenology and the Cognitive Sciences*, 17, 611–627.
- Loughlin, V. (2018b). Going wide: Extended mind and Wittgenstein. *Adaptive Behaviour*, 26(6), 275–283.
- Loughlin, V., & Zahidi, K. (2017). What is left of the active externalism debate? *European Journal of Philosophy*, 25(4), 1614–1639.
- McDowell, J. (1994). The content of perceptual experience. *The Philosophical Quarterly*, 44(175), 190–205.
- Myin, E., & Loughlin, V. (2018). Sensorimotor enactivism and enactive approaches to consciousness. In R. Gennaro (Ed.), *The Routledge handbook of consciousness* (pp. 202–215). New York: Routledge.
- Noë, A. (2004). *Action in perception*. Cambridge: The MIT Press.
- Noë, A. (2007). Magic realism and the limits of intelligibility: What makes us conscious. *Philosophical Perspectives*, 21(1), 457–474.
- Noë, A. (2009). *Out of our heads. Why you are not your brain, and other lessons from the biology of consciousness*. New York: Hill and Wang.
- Noë, A. (2012). *Varieties in presence*. Cambridge: Harvard University Press.
- O’Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46(3), 461–488.

- O'Regan, J. K. (2011). *Why red doesn't sound like a bell: Understanding the feel of consciousness*. Oxford: Oxford University Press.
- O'Regan, J. K., & Dehaene, S. (2000). Predictive processing, perceptual presence, and sensorimotor theory. *Cognitive Neuroscience*, 5(2), 130–131.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–1031.
- Orlandi, N. (2014). *The innocent eye. Why vision is not a cognitive process*. Oxford: Oxford University Press.
- Orlandi, N. (2016). Bayesian perception is ecological perception. *Philosophical Topics*, 44(2), 327–351.
- Ramsey, W. (2007). *Representation reconsidered*. Cambridge: Cambridge University Press.
- Ramsey, W. (2015). Must cognition be representational? *Synthese*, 194(11), 4197–4214. <https://doi.org/10.1007/s11229-014-0644-6>.
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2), 97–118. <https://doi.org/10.1080/17588928.2013.877880>.
- Seth, A. K. (2015). The cybernetic bayesian brain. From interoceptive inference to sensorimotor contingencies. In T. Metzinger & J. M. Windt (Eds.), *Open MIND*, 35. <https://doi.org/10.15502/9783958570108>.
- Silverman, D. (2016). Representationalism and the sensorimotor theory. *Constructivist Foundations*, 11(2), 282–284.
- Silverman, D. (2018). Bodily skill and internal representation in sensorimotor perception. *Phenomenology and the Cognitive Sciences*, 17, 157–173.
- Thompson, E. and Varela, F. J. (2001). Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Sciences*, 5 (10), 418–425.
- Wheeler, M. (2005). *Reconstructing the cognitive world. The next step*. Cambridge, MA/London: The MIT Press. Bradford Books.
- Wiese, W. (2017). What are the contents of representations in predictive processing? *Phenomenology and the Cognitive Sciences*, 16, 715–736.
- Williams, D. (2018). Predictive processing and the representation wars. *Minds & Machines*, 28, 141–172.
- Williams, D., & Colling, L. (2018). From symbols to icons: The return of resemblance in the cognitive neuroscience revolution. *Synthese*, 195(5), 1941–1967. <https://doi.org/10.1007/s11229-017-1578-6>.

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

Affiliations

María Jimena Clavel Vázquez^{1,2}

¹ Division of Law and Philosophy, University of Stirling, Stirling FK9 4LA, UK

² Department of Philosophy, University of St Andrews, St Andrews KY16 9AR, UK