Forthcoming in *Life and Mind: Theoretical and Applied Issues in Contemporary Philosophy of Biology and Cognitive Sciences*, edited by Mariano Martín-Villuendas, Juan Gefaell, and Ana Cuevas-Badallo. Please cite the published version.

Where There is Life There is Mind… And Free Energy Minimisation?

Juan Diego Bogotá

Department of Social and Political Sciences, Philosophy, and Anthropology, University of Exeter, Exeter, United Kingdom
Egenis – The Centre for the Study of Life Sciences, University of Exeter, Exeter, United Kingdom
Grupo de investigación Filosofía y cognición, Universidad Nacional de Colombia, Bogotá, Colombia
e-mail: jb1244@exeter.ac.uk

Abstract This chapter explores the possibility of integrating the enactive and the Free Energy Principle's (FEP) approaches to life and mind. Both frameworks have been linked to the life-mind continuity thesis, but recent debates challenge their potential integration. Critics argue that the enactive approach, rooted in autopoiesis theory, has an internalist view of life and a contentful view of cognition, making it challenging to account for adaptive behaviour and minimal cognition. Similarly, some find the FEP's stationary view of life biologically implausible. Here, I address recent challenges in integrating the FEP and enactivism, focusing on the life-mind continuity thesis. I suggest that the FEP, without explicitly defining life and mind, can be used to model the autopoietic dynamics of organisms. Additionally, I argue that the enactive conception of cognition as sense-making overcomes issues associated with contentful views of cognition. Furthermore, I refute the misinterpretation of the FEP's assertion of stationary organisms, allowing for the modeling of enactive adaptive behaviour through free energy minimisation. Ultimately, I offer a constructive and interactionist approach to life and mind, transcending internalist and externalist perspectives.

Keywords Adaptivity, Autopoiesis, Enactivism, Free Energy Principle, Life and mind continuity thesis, Sense-making

1 Introduction

Artificial life (A-Life) is a discipline that makes use of computational models and simulations to study life, not only as we know it, but as it could be. With the advent of A-Life and its close links to artificial intelligence methods and ideas, theorists have noticed the possibility of using it to shed some light on our understanding of the mind. As Michael Wheeler has put it, "A-Life has the potential to be the intellectual engine of a biological cognitive science" (1997, p. 10). There are at least two assumptions behind this claim. The first assumption is that research on artificial models can tell us something about the nature of non-artificial phenomena such as life and mind. The second one is that learning about life will tell us something about the mind, which is to say that cognition is in some non-trivial sense a biological phenomenon. The latter assumption is known in the literature as the life-mind continuity thesis (LMCT).

To say that life and mind are continuous broadly means that, while they cannot be reduced to one another, there are not breaks or gaps between them, implying that mind arises from life. This understanding of 'continuity' can be traced back to John Dewey's (1938/2008) works. Note, however, that there are different ways of interpreting LMCT. Roughly, we can distinguish between ontological and epistemological interpretations of LMCT. According to the former, life and mind *are* continuous. In this chapter, I mostly focus on a strong reading of an ontological interpretation of LMCT. From that perspective, life and mind share a common set of functional characteristics of which those that are characteristic of cognition are enriched versions of those that are characteristic of life (Godfrey-Smith 1996a; 1996b). An epistemological reading, in contrast, is not concerned with what life and mind *are*, but with the scientific tools and concept needed for properly understanding those phenomena (see, e.g., Clark 2001, p. 118). So, if we were anti-realists concerning scientific theories, this interpretation of LMCT would suggest that even if life and mind were not alike or continuous, our best scientific theories of them would suggest they are insofar as they rely on the same set of core concepts.

In the current cognitive science landscape, two prominent approaches advocate for a strong ontological interpretation of LMCT.[1] The first one, the enactive approach (Varela 1997; Thompson 2007), has its roots in the theory of autopoiesis (Maturana and Varela 1980), which itself subscribes to a form of LMCT. For the enactivist, the organisation that defines an organism is sufficient for cognition. Therefore, the biological study of life should be continuous with the scientific study of the mind.

More recently, the Free Energy Principle (FEP) has been linked to LMCT (Kirchhoff and Froese 2017; Bruineberg et al. 2018). Initially, the FEP was introduced as a theory of brain functioning (Friston 2005; 2010). Under this guise, the FEP is meant to explain cognition, perception, and action under a single unified brain mechanism, i.e., free energy minimisation (known as 'prediction error minimisation' in the predictive processing literature, see Hohwy 2013; Clark 2016). During the last decade, however, the FEP has been generalised as a theory of biological systems (Friston 2013), as well as a "theory of every 'thing'" (Friston 2019, p. 4; see also Friston et al. 2021; Friston et al. 2022a), where 'thing' is given a technical definition within the framework. In the context of biological systems, the idea is that, for an organism to maintain its integrity, it must engage in inferential behaviour that involves acting adaptively within its environment. Such adaptive behaviour is taken to be cognitive.

Ever since the scope of the FEP was broadened to include biological systems, it has been suggested that it can be integrated with enactivism to form a single unified approach to life and mind (e.g., Kirchhoff and Froese 2017; Allen and Friston 2018; Ramstead et al. 2018b; van Es and Kirchhoff 2021; Wiese and Friston 2021; Korbak 2021).[2] Recently, however, the idea of such integration has come under attack from both camps. On the one hand, Michael Kirchhoff (2018) has argued that in contrast to the FEP, the theory of autopoiesis is too internalist in its explanatory focus, making it hard for it to accommodate cognition in its framework—and even if the concept of autopoiesis was expanded as it is in later enactive literature (e.g., Thompson 2007), the enactive conception of cognition as 'sense-making' is at odds with the FEP's contentless view of cognition. On the other hand, Ezequiel Di Paolo, Evan Thompson, and Randall Beer (2022) have suggested that links between enactivism and the FEP tend to rest on

[1] From now on, and unless stated otherwise, I use 'LMCT' to refer to the strong ontological interpretation of the continuity between life and mind.

[2] There is further literature that has attempted to interpret the FEP through an enactive lens without focusing specifically on LMCT (see Bruineberg and Rietveld 2014; Bruineberg et al. 2018; Kiverstein 2020; Ramstead et al. 2020b; Ramstead et al. 2021).

misunderstandings of core enactive concepts, and that, on a closer look, both approaches conceive of life in fundamentally different ways, making the integration between the two approaches implausible.

My purpose is to evaluate to what extent the enactive and the FEP approaches to life and mind are related and whether they can be integrated into a unified approach. In this regard, in contrast to some of the literature that has tried to integrate these two perspectives, I focus exclusively on how enactivists and FEP theorists address the continuity between life and mind. Additionally, furthering the literature that has focused on LMCT, the FEP, and enactivism, I address recent challenges that have been raised vis-à-vis the possibility of integrating these frameworks. To do so, in section 2, I first introduce the enactive and the FEP approaches to life and mind.[3] In sections 3 and 4, I respectively present some of the current attempts to integrate the FEP and enactivism, as well as both Kirchhoff's and Di Paolo et al.'s arguments against such views. Finally, in section 5, I propose a way in which enactivism and the FEP can be taken to be complementary, circumventing the arguments put forth by Kirchhoff and Di Paolo et al.

2 Life and mind from the enactive and Free Energy Principle perspectives

2.1 Autopoiesis, enaction, and sense-making

As made clear by Thompson's (2007) *Mind in Life,* the idea that life and mind are continuous is one of the core ideas of the enactive approach. As he states at the very beginning of the book, "[w]here there is life there is mind, and mind in its most articulated forms belongs to life" (2007, p. ix). This idea is inherited from the theory of autopoiesis, from which the enactive approach first arose.[4] In this sub-section, I briefly present the theory of autopoiesis and how it has been expanded upon within the enactive approach to life and mind.

Enactivists define life as "sense-making in precarious conditions" (Thompson, 2011a, p. 114, emphasis omitted). Rooted in the theory of autopoiesis, which conceptualises living systems as self-producing and self-distinguishing (Maturana and Varela, 1980), the enactive approach focuses on the organisational properties that characterise living systems. One of the core ideas in the theory of autopoiesis is that life can be defined by referring to the organisation (as distinct from the material structure) of living systems. It is at that level that we can see that living systems are *operationally closed*, which means that they are constituted by a set of recurrent processes that have their own topology and dynamics. An operationally closed system is an autonomous system, and an autonomous system in the molecular domain is defined as autopoietic. Thus, in the molecular domain, operational closure implies self-production and self-distinction.

---

[3] If the reader is already acquainted with the enactive conception of life and mind and the recent literature on the FEP, I suggest them to skip to section 3.

[4] Thompson's enactive take on LMCT goes beyond the strong continuity between life and mind we can find in the theory of autopoiesis. According to Thompson, the continuity between life and mind is 'deep', which means that life and mind not only share organisational properties but also phenomenological ones: "certain existential structures of human life or phenomenological structures of human experience […] are applicable to life itself" (Thompson 2011b, p. 216; see also Thompson 2007, pp. 129, 157). Addressing the significance of Thompson's claim for a potential integration of the FEP and enactivism is beyond the scope of this chapter. Suffice it to say that I believe that the constructivist position I endorse in section 5.3 is consistent with the phenomenological commitments of the enactive approach. I am grateful with Julian Kiverstein for directing my attention to this point.

In the enactive framework, living systems are not only considered autopoietic (and thus, autonomous) but also adaptive and existing in precarious conditions (Di Paolo, 2005). Because it is precarious, and to keep its constituent processes running, the organism must interact with its environment adaptively, i.e., regulating its states whenever it is approaching its limits of viability, resulting in the transformation of those states into tendencies toward the prevention of reaching viability limits in the future. Additionally, according to Di Paolo, "[o]nly of the subset of autopoietic systems that are […] adaptive can we say that they posses [sic] operational mechanisms to potentially distinguish the different virtual implications of otherwise equally viable paths of encounters with the environment. This differential operation is […] sense-making" (2009, p. 14). Given the focus on paths of encounters in its definition, sense-making must be understood partly in terms of sensorimotor interactions between organism and environment.

Importantly, sense-making is intentional because it implies an organismic perspective from which meaning is brought forth. Because of the organism's precariousness, certain interactions with the environment are either positive or negative from the perspective of the organism. Consider a bacterium swimming up a gradient of glucose (Varela 1997). The glucose is meaningful for the bacterium because it provides nutrients that are necessary for maintaining its metabolic processes (i.e., its autopoiesis). The meaning of the glucose gradient only makes sense from the perspective of the bacterium. Take away the bacterium from the situation and the glucose stops having its meaning. What brings forth such meaning is the organism's adaptive and autopoietic organisation. Such bringing forth of meaning is what enactivists call sense-making.

Enactivists thus claim that organisms bring forth or enact significance via sense-making, implying that sense-making is the basic form of cognition. Put this way, life and cognition come together, entailing a continuity between life and mind.

2.2 A free energy principle for biological systems

Since its introduction, the Free Energy Principle (FEP) has been presented as a formal approach to understand brain functioning, biological systems, and self-organising systems more generally (Friston 2005; 2010; 2012; 2013; 2019; Ramstead et al. 2023). It is partly because of its applicability to both cognitive and biological systems that the FEP is sometimes taken to be able to address LMCT from a mathematical perspective. I now turn to draw the connection between the FEP and LMCT. Here, I provide an informal and heuristic presentation of the FEP framework. For an in-depth mathematical treatment, see Friston (2019).

Simply put, the FEP is a mathematical framework that allows us to describe the dynamics of self-organising systems using the tools of several fields including statistical physics, random dynamical systems, and Bayesian probability theory. The upshot of that description is that it may allow for the formulation of mechanical theories that explain the dynamics of self-organising systems and how they are coupled with their surroundings (Ramstead et al. 2023).

In a nutshell, the FEP states that the dynamics of a self-organising system can be described as if they encoded a probabilistic model of their environment, and that, given that description, the system must minimise an information-theoretic quantity called *free energy* to retain its integrity. Describing a self-organising system in that way requires formally distinguishing the dynamics of the target system from the dynamics of the system's environment. This can be done by using the Markov blanket formalism (Pearl 1988).

A Markov blanket is the statistical interface between two sets of random variables that are conditionally independent of each other. The FEP framework makes use of Markov blankets to model the dynamical coupling between two partitions of a unified dynamical system. Under the FEP, a set of variables is interpreted as the *internal states* of a self-organising system, another set as the *external states* (i.e., those of the environment), and the statistical interface between them as the *blanket states*. Blanket states are further specified as either *sensory* (if they influence internal states) or *active states* (if they influence external states). Given the use of the Markov blanket formalism, the overall dynamical system is partitioned between *particular states* (which include both the internal and the blanket states) and external states (Figure 1). The particular states are meant to represent the dynamics of the target self-organising system, often referred to as a *particle*.

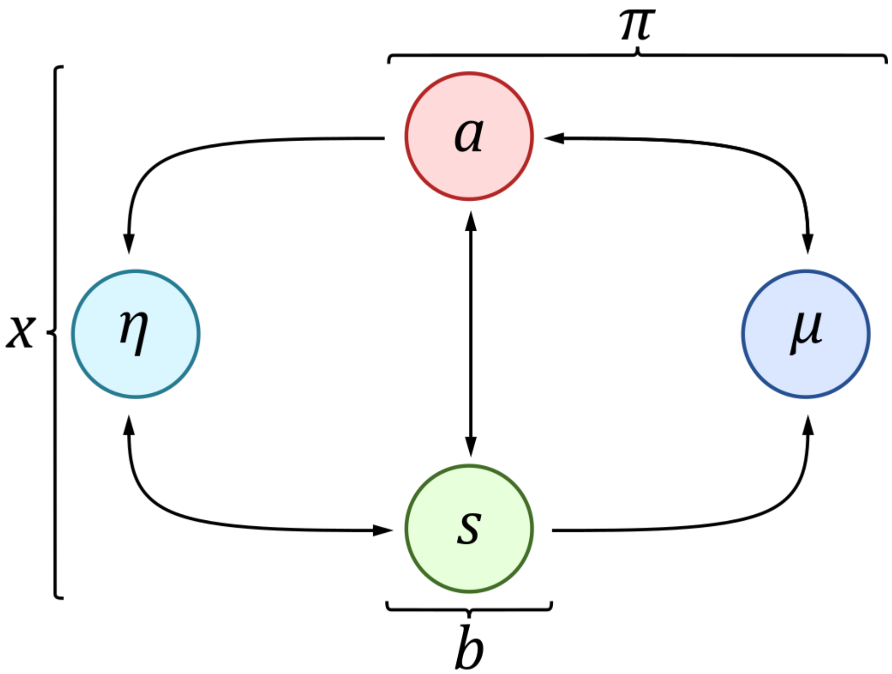

Fig 1. A Markov blanketed system under active inference (adapted from Parr et al. 2022, p. 44). This diagram represents the influences between the different variables in a simple four-state system x. The system is partitioned into external (cyan, η), sensory (green, s), active (red, a), and internal states (blue, μ). The combination of sensory and active states comprises the blanket states ($b = \{s, a\}$), and the combination of the blanket and the internal states comprises the particular states ($\pi = \{b, \mu\}$).

The main move of the FEP is to give the partitioned system a probabilistic interpretation. To do so, one can note that for a self-organising system to retain its integrity over time, it must avoid states in which it would no longer be viable. For instance, a person with an internal temperature of 45ºC would probably die. Therefore, in terms of probability theory, an internal temperature of 45ºC in a human being has a high *surprisal*. Technically speaking, surprisal is an information-theoretic quantity that equates to the negative log-probability of finding a particular system in a given state, meaning that the increase of surprisal can be cast in terms of an increase of uncertainty or entropy. Therefore, to retain its integrity over time, the particle must minimise surprisal.

Under the FEP, surprisal can be minimised in two different ways which lead to two different formulations of the FEP formalism (Ramstead et al. 2023). The first and most common formulation focuses on determining how to minimise surprisal over states (call it *state-based formulation*; see Friston 2013 for an often-cited example of such formulation). The second formulation focuses on determining how to minimise surprisal over paths or trajectories in state space (call it *path-based formulation*; see Friston et al. 2022a; Ramstead et al. 2023). Both formulations have in common that surprisal is minimised by minimising another information-theoretic quantity known as variational free energy. Variational free energy is mathematically an upper bound on surprisal, meaning that minimising it implicitly minimises surprisal. The reason why free energy is needed is because surprisal is intractable, whereas free energy is not (Parr et al. 2022, p. 35).

Variational free energy can be cast as the divergence between a probabilistic model of the external states and the observations acquired via the sensory states of the Markov blanket. Such a probabilistic model is said to be encoded by the internal states of the particle, hence the FEP's claim that, for a self-organising system to retain its structure, it must encode a probabilistic model of its environment and thus minimise variational free energy.

How can a particle minimise free energy? Under active inference, a corollary of the FEP, free energy minimisation is cast in terms of inference: How should the probabilistic model be for it to correctly infer the external states? Such inference can be put in terms of approximate Bayesian inference. There are two ways of minimising free energy: Perception, which consists in fitting the model with the sensory evidence acquired via the sensory states of the Markov blanket; and action, which consists in making the external states fit with the probabilistic model via the active states of the Markov blanket. Both perception and action work in tandem (Parr et al. 2022, p. 27), implying a sensorimotor loop: the particle modifies its environment via its active states, indirectly causing new sensory observations which in turn call for an updating of its probabilistic model and new actions.

Active inference moves from variational free energy to expected free energy, which roughly corresponds to future free energy (Parr et al. 2022, p. 81). Simply put, action requires anticipatory dynamics, given that the particle must infer how an action undertaken in the future will minimise free energy. Mathematically, expected free energy depends upon a complex probabilistic model that also includes (and therefore is biased toward) a set of preferences (Tschantz et al. 2020). There are actions that are preferred over others precisely because, given the internal dynamics of the particle, they will probably minimise more free energy.

The main difference between the state- and the path-based formulations of the FEP lies in how they formally capture the prior preferences encoded by the probabilistic model of the particle under active inference. In the state-based formulation, these preferences are cashed out in terms of a non-equilibrium steady state (NESS) density. It is assumed that, with enough time, the particle will converge into a set of attracting states (i.e., a random global attractor; see Friston 2013). In other words, the particle's actions tend toward such a set of states. In contrast, in the path-based formulation, there is no need for the assumption of a NESS density (Friston et al. 2022b; Ramstead et al. 2023). In this formulation, the problem that the FEP is addressing is cast as one concerning paths of least action (roughly, expected free energy) given both the internal and the external dynamics. Under this interpretation, internal and active paths (i.e., the flow of internal and active states) look as if they tracked external paths via predictive dynamics by naturally evolving through paths of least action, thus minimising expected free energy without the need for the NESS density.

What does all this have to do with LMCT? When addressing life and mind, FEP theorists often claim that the formalism implies that living systems are intentional (e.g., Ramstead et al., 2019; Ramstead et al., 2020a; Wiese & Friston, 2021) and/or sentient (Friston et al., 2023; Parr et al., 2022; Ramstead et al., 2023). When doing so, they do not mean that the target system is phenomenally conscious or self-aware, but rather that it is "responsive to sensory impressions" (Friston et al., 2020, p. 3). Such a responsiveness is taken to be a consequence of the Markov blanket formalism and its partition into sensory and active states. A system equipped with a Markov blanket is taken to be 'sentient' because it will statistically be affected by the external states through its sensory states.

Additionally, biological systems are described as intentional under the FEP given the description of the internal states of the particle as encoding probabilistic beliefs about the external states. Even if these beliefs should not be understood as propositional having propositional content (as the concept of 'belief' is usually understood in the philosophy of mind), they are still *about* the environment, meaning that they are intentional (Ramstead et al., 2019). In other words, from this perspective, living systems are taken as cognitive, implying LMCT.

3 A unified path?

When looking at both the enactive and FEP views on mind and life, it is undeniable that there are some similarities. Both approaches put emphasis on self-organisation, and on adaptive encounters with the environment via sensorimotor interactions. Some authors have built upon these similarities to propose integrations of the enactive and FEP approaches, paving the way towards a unified approach to life and mind. In this section, I present the main ideas that these proposals have in common.

One of the main bases for proposals of a unified approach is Friston's suggestion that, under the FEP, one of the main attributes of biological self-organisation is that organisms are autopoietic, which for him means that

> because active states change—but are not changed by—hidden [i.e., external] states […], they will appear to place an upper (free energy) bound on the dispersion (entropy) of biological states. This homoeostasis is informed by internal states, which means that active states will appear to maintain the structural and functional integrity of biological states. (2013, p. 5)

For Friston, the link between the FEP and autopoiesis is straightforward given Maturana and Varela's (1980, p. 79) description of autopoietic systems as homeostatic. Under the FEP, homeostasis is understood as the maintenance of the system against entropic decay, which is formally expressed as the persistence its Markov blanket. Consequently, although Friston does not put it explicitly so, it seems plausible to say that he takes autopoiesis to be achieved within the FEP by means of the maintenance of the Markov blanket. This maintenance is achieved via active inference. From this perspective (exemplified by, e.g., Allen and Friston 2018; Wiese and Friston 2021), the FEP would entail autopoietic organisation and adaptive behaviour (but see section 4.2), which, within the enactive approach, would be sufficient for sense-making.

It has been also suggested that the link between the FEP and enactivism can be drawn by casting the operational boundaries of a cognitive system (i.e., the boundaries of the topology that arises via operational closure) in terms of Markov blankets.[5] For instance, Thomas van Es and Kirchhoff claim that

---

[5] The emphasis in this argument is somewhat different from the one focused on autopoiesis because enactivists distinguish between autonomous and autopoietic organisation (Thompson 2007). Simply put, operational closure implies autonomy, which

> any operationally closed system is inherently composed of multiple individually distinguishable component processes. […] If we take each component process to have a Markov blanket, and the larger, operationally closed network to have a Markov blanket too, the generation of a self-enabling and self-individuating process network can be cast as the hierarchical self-organization of a Markov blanketed *ensemble* of Markov blankets. (2021, p. 6637)

In this proposal, each Markov blanketed process is taken to be tightly coupled to the other, entailing a set of recurrent processes that, as a whole, constitute an operationally closed network. Operational closure is thus cast as a blanketed set of Markov blankets and tight coupling between inner and blanket states. As Maxwell Ramstead and colleagues suggest, "it is fairly straightforward to establish that the Markov blanket formalism provides a statistical formulation of operational closure" (2021, p. 55). Given that surprisal can be formulated as entropy (see section 2.2), it can be said that the whole blanketed set of Markov blankets will naturally tend toward entropy (making it a precarious system), unless it minimises its free energy—such minimisation, under active inference, can be described as adaptive behaviour. So, from this perspective, casting operational closure in terms of Markov blankets entails that adaptive operationally closed systems in precarious conditions minimise their free energy.

One potential motivation to unify the FEP and enactivism has to do with the so-called pebble challenge, i.e., the claim that the FEP's formalism may equally apply to both organisms and pebbles, entailing that it does not say anything interesting about mind or life. To avoid this challenge, van Es and Kirchhoff (2021) have suggested using enactive ideas to constrain the applicability of the FEP formalism to specifically biotic systems. Another motivation has to do with the explanatory potential that the FEP framework may provide by linking self-organising dynamics to ideas from statistical physics and computational modelling. Some enactivists seem to consider that potential as a suitable way to formally address the dynamics that life and mind may share. The FEP framework, nevertheless, may lead to cognitivist and representationalist views that are at odds with the enactive approach. Thus, enactivists often reconceptualize FEP notions such as 'inference' and 'model' in less cognitivist ways (e.g., Kirchhoff and Froese, 2017; Bruineberg et al., 2018). By casting free energy minimisation in enactive terms, one may describe active inference as enactive inference (Ramstead et al., 2020b).

4 A forking path

The above proposals notwithstanding, arguments coming from both the FEP and enactive camps have recently suggested that the integration between the two approaches is unlikely. In this section, I present the arguments given by Kirchhoff (2018) and Di Paolo et al. (2022).

4.1 Internalist vs. Externalist explanations

The main target of Kirchhoff's (2018) argument is not precisely the enactive approach, but rather Maturana and Varela's original theory of autopoiesis which, as shown in section 2.1., has been complemented in the enactive literature with the concepts of adaptivity, precariousness, and sense-making. However, despite asserting that it is not his intention to "claim that

in the context of molecular systems equates to autopoiesis. Consequently, all autopoietic systems are autonomous, but not the other way around. Something characteristic of the enactive approach (as opposed to the classical theory of autopoiesis) is the focus on autonomous systems rather than simply autopoietic ones (see Varela et al. 1991/2016). Here, I have focused on autopoiesis because of its importance for LMCT.

enactivism and the FEP cannot complement one another and so jointly combat orthodox 'cognitivist' view in cognitive science" (Kirchhoff 2018, p. 2521), Kirchhoff has also some reservations concerning the possible integration between FEP and enactivism given the latter's conception of cognition as sense-making.

Kirchhoff's main argument is that the FEP should be preferred over the theory of autopoiesis when it comes to LMCT. His argument relies on Godfrey-Smith's (1996a, chap. 2) distinction between internalist and externalist causal explanations in biology: "[T]he term 'externalist' will be used for all explanations of properties of organic systems in terms of properties of their environments. Explanations of one set of organic properties in terms of other internal or intrinsic properties of the organic system will be called 'internalist'" (1996a, p. 30). Godfrey-Smith's distinction has to do with the epistemological question about the nature of our explanations of biological properties, not with the ontological question about the boundaries of the mind. It is possible to argue that a given framework about life and mind is externalist (in Godfrey-Smith's epistemological and biological sense), as well as internalist (in the ontological sense about the nature of the mind). One may, for instance, suggest that the biological dynamics of a cognitive system operate in the way they do to cope with environmental complexities while maintaining that that same cognitive system's cognitive capacities are realised exclusively by the brain.

Borrowing Godfrey-Smith's distinction, Kirchhoff notes the autopoietic focus on self-production maps to an internalist explanation, whereas the FEP's (via active inference) focus on self-preservation maps to an externalist explanation. The idea is that the theory of autopoiesis explains the organisation of living systems by referring almost exclusively to the internal dynamics self-producing of the organism, thus arguably reducing the epistemic importance of the environment, making it an internalist explanation. As Godfrey-Smith puts it when discussing autopoiesis: "If environmental features matter, then these too are internal to the system, not explanatorily important elements outside of it: all feedback is internal" (1996a, p. 49).

The FEP is interpreted as externalist because of two reasons. First, as explained in section 2.2., the FEP describes the internal dynamics of an organism is as if they encoded a probabilistic model of the environment. That entails that an explanation of the internal dynamics of an organism would necessarily have to refer to the specific environment it is embedded in. Therefore, the inner is explained in terms of the outer.

Second, Kirchhoff emphasises that, under active inference,

> [g]iven that the environment is non-stationary, organisms can optimize their [inferences] through embodied activity in the environment […]. It is because organisms can act on the environment that they can occupy a domain far from (terminal) phase-boundaries. Indeed, organisms strive to achieve a maximal fit between their probabilistic models and environmental niche via embodied activity. (2018, p. 2529)

Action arises from self-preservation. As mentioned earlier (section 2.2.), perception is not enough for the minimisation of surprisal; action is also required to minimise surprisal. Here the emphasis is on how biological systems must interact *adaptively* with their environment, not on their inner dynamics alone as autopoiesis arguably does. Thus, the FEP gives an externalist explanation of life.

The crux of Kirchhoff's argument lies in the link he draws between cognition to adaptivity: "cognition is usually taken to involve a future-oriented regulation […]. It follows cognition is *prima facie* a relational feature, enabling living systems to improve conditions of self-maintenance through *adaptive* behavior in their environmental niche" (2018, p. 2530). Via active inference, the FEP accommodates adaptive behaviour within its (externalist) explanatory framework and, thus, for Kirchhoff, it is equipped to approach LMCT. In contrast, he argues that autopoiesis is ill-suited for that. Simply put, Maturana and Varela's classical formulation of autopoiesis downplays the role of the environment, which in turn entails that it is not clear whether autopoiesis involves adaptive behaviour. If autopoiesis is not necessary for adaptive behaviour, then the theory of autopoiesis cannot satisfactorily address LMCT. Thus, Kirchhoff concludes, the FEP should be preferred over the theory of autopoiesis.

At this point, one may refer to the enactive complementation of autopoiesis with adaptivity and the introduction of sense-making (i.e., cognition) into the autopoietic picture. As Kirchhoff himself acknowledges, doing so "shoehorns nicely with the FEP […]. [O]n both accounts, life and mind share the same basic property: adaptivity—a capacity for future-oriented regulation of functional and structural integrity" (2018, p. 2535). There is, however, a caveat: "[T]he idea that sense-making is the right way to go in cognitive science is not obvious. […]. [I]f the notion of sense-making entails the existence of meaningful content, then it is not at all clear that the FEP and autopoietic enactivism can be brought into a truly fruitful interplay" (2018, pp. 2536-2537). Relying on so-called 'radical enactivist' arguments against basic cognition involving content (Hutto and Myin 2013), i.e., the specification of conditions of satisfaction, Kirchhoff suggests that there is some tension between the FEP and enactivism insofar as sense-making is conceptualised as an *evaluation* of the environment given the metabolic concerns of the organism (e.g., Di Paolo 2009, p. 14; Colombetti 2014, p. 17). In contrast, it is argued that inference within the FEP framework does not involve content given that it is better understood as a mathematical (hence, purely formal) process (Ramstead et al. 2020a). Unless enactivists take a deflationary stance on sense-making, an integration between their views and the FEP seems unlikely. However, enactivists would insist that sense-making, understood as involving a meaningful (and hence, 'radical enactivists' would argue, contentful) perspective over the world, necessarily follows from autopoiesis and adaptivity.

In sum, Kirchhoff apparently suggests that the problems with autopoiesis and enactivism have to do with whether these approaches can satisfactorily address basic cognition in relation to LMCT. On the one hand, the internalist explanation given by the theory of autopoiesis seems to leave aside any explanatory role of the environment, which in turn arguably implies the impossibility of accounting for cognition within the theory. On the other hand, even if autopoiesis is complemented with adaptivity—as it is done within enactive literature—, it remains unclear to what extent sense-making is the way to go. For Kirchhoff, the FEP avoids these issues by accommodating adaptivity and the explanatory role of the environment via active inference, which in turn must be conceived of contentless.

4.2 Autopoietic Markov blanketed systems at a non-equilibrium steady state?

Aside from Kirchhoff's arguments, which explicitly draw something of a gap between enactive and FEP ideas by favouring the latter, Di Paolo et al. (2022) elaborate on what they consider to be deep tensions between the two approaches. In line with recent critical evaluations of the FEP framework (see, e.g., Aguilera et al. 2021; Colombo and Palacios 2021; Raja et al. 2021), Di Paolo et al. question some of the assumptions of the FEP, as well as its applicability to biological systems, suggesting that enactivism is better posed to address life (and hence, under LMCT, mind).

According to Di Paolo et al., FEP theorists tend to misread the concept of autopoiesis. Di Paolo et al. stress the distinction between organisation and structure in relation to the idea that organisms retain their integrity over time. As mentioned in section 3, some FEP theorists claim that, under the FEP, biological systems are autopoietic given how they maintain their integrity against entropic decay. However, as Di Paolo et al. clarify, these claims rely on a misunderstanding of what autopoiesis is. Autopoiesis does imply that organisms maintain their integrity and persist over time—but at the level of their *organisation*. Given constant metabolic changes, as well as material and energetic exchanges with the environment, the *structure* of an organism rarely persists over time. In the FEP literature, in contrast, it is often unclear what it is precisely that self-organising systems preserve under the FEP. It is usually claimed that they preserve their integrity. But should we interpret such integrity as organisational or as structural? Sometimes it is claimed that the Markov blanket is preserved at the level of self-organisation (Friston 2019, p. 50). Other times it is claimed that what is preserved is the structural and functional integrity (Friston 2013, p. 9). Thus, Di Paolo et al. assert, "[t]here is no obvious organization-level equivalent in Friston's systems to the network of processes that through transformations *realize* the conditions of its own production" (2022, p. 13).

A second misreading Di Paolo et al. identify refers to the concept of boundary. In the enactive literature, one way of referring to the boundary of a living system is with the notion of operational closure, which refers to the topology of a recursive network of processes that make up a given autopoietic system at the organisation-level. Given that such a topology is closed, it seems right to say that it is bounded. The operational boundaries of a system may not coincide with what we would usually identify with the boundaries of its structure. Enactivists argue that organisms can extend themselves by incorporating material structures from the environment into their organisation (Di Paolo 2009; Colombetti 2017). As mentioned in section 3, some FEP theorists suggest that the FEP accommodates the notion of operational closure using the Markov blanket formalism. However, Markov blankets are used in several, sometimes ambiguous and even contradictory ways within the FEP literature, as if they captured the structural, cognitive, operational, and several kinds of boundaries.[6]

Di Paolo et al. do not see the need for mediation between an operationally closed system and its environment like the one established by blanket states between internal and external states. For them, "Markov blanket conditions are too strong and too undiscriminating to account for the complexity of the organism-environment relation, and this is a point of tension" (2022, p. 30). They take Markov blanket conditions to be too strong because some biological processes (e.g., sweating) do not involve or instantiate a clear structural boundary involving conditional independency between the inner and the outer, and too undiscriminating because of the wide applicability of Markov blankets without a clear principled way of defining when they should be applied. In contrast, operational closure involves a (flexibly) bounded topology at the organisation-level defined in a principled way via the self-distinction condition of autopoiesis, which is correlated with the condition of self-production (see section 2.1). Under the FEP, the definition of boundaries rests on "the simplifying assumption that over a suitable time scale, blanket states are well defined—as a subset of attracting states" (Friston 2019, p. 50). So, at best, the mapping of boundaries via Markov blankets rests on nothing but an assumption. At worst, as it has been argued by Vicente Raja and colleagues (2021, pp. 59-60), the applicability of the Markov blanket formalism within the FEP literature is arbitrary. Using Kirchhoff and colleagues' (2018) active inference model of two coupled pendulums connected by a beam as an example, Raja et al. question the decision to interpret the beam as a Markov blanket that separates each pendulum. Why cannot we interpret the atoms on the surface of the pendulums as their Markov blankets? There is nothing about the beam that prescribes the need for partitioning the system in the way in which Kirchhoff et al. do. Because of this reason, Raja et al.

[6] See Menary and Gillet (2020), Raja et al. (2021) and Bruineberg et al. (2022) for discussions on this issue.

suggest that, within the FEP literature, the Markov blanket partition responds to what is convenient for the purposes of the model, making it arbitrary.

Given the reasons above, Di Paolo et al. (2022) claim that not only there is no straightforward relation between operational closure and Markov blankets, but also that there is some tension between how the FEP and enactivism conceive of boundaries.

Perhaps the deepest tension between enactivism and the FEP revolves around the non-equilibrium steady state (NESS) assumption in the latter. Recall that one way of determining the surprisal of the states of a given system under the FEP involves assuming that the target system will tend toward a set of attracting states. It is from this assumption that several points within the FEP framework follow, including the descriptions of the target self-organising system as encoding a probabilistic model (Andrews 2021, p. 13), and as behaving as if it performed approximate Bayesian inference (see Parr et al. 2022, pp. 47-50). Ramstead et al. clarify that the NESS assumption allows "us to say *more informative things about the flow* of target systems" (2022, pp. 17-18). However, as has been suggested by Matteo Colombo and Patricia Palacios (2021), this assumption places costly constraints on the study of biological systems since one of the most interesting characteristics they have is precisely the fact that their dynamics may change radically, thus giving rise to a completely new attractor landscape. This argument puts into question the idea of biological systems converging into a NESS. Simply put, the NESS assumption precludes significant change over time.

Di Paolo et al. compare the NESS assumption in the FEP with the historicity of sense-making. Via the enactive conception of adaptivity, enactivists conceive of life and mind as historical (i.e., non-stationary) in the sense that the space of possibilities for a given organism changes qualitatively over time. As they suggest, "adaptivity entails that adaptive interventions necessarily modulate the dynamical landscape, for instance, by changing parametric relations to the environment, (un)coupling to other systems, or altering constraints or boundary conditions" (2022, p. 21). So, contrary to Kirchhoff's (2018) claims, Di Paolo et al. argue that (the state-based formulation of) the FEP cannot truly accommodate adaptivity. As long as the FEP retains its NESS assumption, it is argued that it is incompatible with the enactive conception of life and mind as adaptive and historical.

Summing up, Di Paolo et al.'s arguments mainly target the translation of enactive ideas into the FEP ideas—most notably, that of operational closure into Markov blankets—, which for them is usually misguided, as well as the biological plausibility of the FEP given the NESS assumption.

5 Laying down a converging path (of sorts)

Are then all attempts at integration between enactivism and the FEP futile? Is this really an either/or situation in which we are left to choose between either potential issues concerning basic cognition in enactivism and autopoiesis (as Kirchhoff suggests), or potential issues concerning life in the FEP (as Di Paolo et al. suggest)? I believe not. In this final section, I argue for a potential way of complementing these two approaches to address LMCT. As it will be clear, my aim is not to show that enactivism and the FEP *ought* to be integrated. My claim is more modest. I simply argue that they can indeed be seen as potentially working in tandem to address LMCT, and that if FEP theorists wish to address life and mind, they should probably use the enactive approach as a guide.

I first discuss the concept of sense-making to address one of Kirchhoff's arguments, showing not only that sense-making is contentful in a non-problematic way, but also that cognition without content seems implausible, which in turn puts the FEP conception of cognition into question. I then address the different concepts of 'boundary' at play in the discussion, suggesting that by neatly distinguishing between operational, statistical, and structural boundaries, there could be no issue with mapping operational closure onto Markov blankets. I thus pave a converging path between the FEP and enactivism in which the former is better seen as a way of possibly modelling the cognitive dynamics explained by the latter. Finally, I suggest that this converging path is neither internalist nor externalist, but rather constructivist, which is a third kind of explanation discussed by Godfrey-Smith that Kirchhoff leaves aside.

5.1 Cognition and content

As mentioned in section 4.1, the biggest strength that Kirchhoff sees in the FEP over enactivism is a contentless conception of basic cognition via active inference. Given this view, one might wonder what is so problematic about content. The idea of contentful basic cognition is seen as problematic from a naturalistic perspective: meaningful content is nowhere to be found in the physical world (all there is is information as co-variance) and, therefore, cognitive processes cannot be understood as content-bearing from the bottom-up—at least if one wants to maintain a naturalist view. This claim, introduced by Daniel D. Hutto and Erik Myin (2013), is known as the 'hard problem of content'. If there is indeed a hard problem of content, then basic cognition cannot be contentful. Hutto and Myin arrive at this conclusion because, assuming that intentionality is the mark of cognition, and that intentionality must be specified in terms of semantic content, then it follows that there is an unbridgeable gap between basic cognition and the natural origins of intentional content. Thus, either we accept that basic cognition is contentless, or we introduce an inflated notion of information beyond covariance that is at odds with naturalism.

Recently, Thompson (2018) has called attention to the fact that Hutto and Myin seem to equate content (i.e., conditions of satisfaction) with *representational* content. If cognition were representational from the bottom-up it would certainly involve meaning and conditions of satisfaction. Representational content can be right or wrong concerning what it represents. Part of what defines representational content is the fact that it carries information about a state of affairs x while it being possible for it to get x wrong. It is then a necessary condition for a representation to be capable of misrepresenting what it is meant to represent (Rowlands 2009). From this perspective, representational content does indeed fit Hutto and Myin's characterisation of content.

As opposed to the conception of content as representational, Thompson (2018) briefly points out that phenomenologists (e.g., Maurice Merleau-Ponty) and embodied cognitive scientists (e.g., Walter Freeman) have argued for basic cognition as involving content, but not representational.[7] Somewhat surprisingly, Thompson does not mention sense-making. However, I believe that one of the main claims of the enactive approach is precisely that sense-making is contentful without it being representational.

Consider Adrian Cussins' definition of content: "The term 'content' […] refers […] to the way in which some aspect of the world is presented to a subject; the way in which an object or property or state of affairs is given in, or presented to, experience or thought" (2003, p. 133).

[7] To be fair, Hutto and Myin do briefly discuss non-representational views of intentional content, but they dismiss them as being "lax and liberal" (2017, p. 12). Perhaps such views are *too radical* for them.

This definition is consistent with Hutto and Myin's conception of content as involving conditions of satisfaction. Arguably, x can be presented to a subject as having property P, but it may still be possible for x to not have property P regardless of how it is presented to the subject. Cussins's definition, however, is broad enough to accommodate both representational and non-representational content. Sense-making, I suggest, also Cussin's definition without implying representational content. Let me elaborate on this point.

Aside from the fact that representations carry information about what they represent, and that they may misrepresent it, it is often claimed that representations are decouplable from what they represent (Haugeland 1991; Rowlands 2009). The decouplability criterion is used to capture the fact that a representation can serve as a proxy or a stand-in for what it is meant to represent. If a representation is decouplable from what it represents, it then can guide behaviour even in the absence of its target. Decouplability thus helps when accounting for cases of offline cognition. From this perspective, decouplability and the capability of misrepresentation are two distinct characteristics of representational states. I believe, however, that these two characteristics must be seen as deeply linked to one another. Consider a case in which an information-bearing state of affairs is not decouplable from what it provides information about: the number of rings in a tree trunk. As it is well known, that number covaries with the age of the tree. The older the tree, the more rings it will have in its trunk. The number of rings and the age are two coupled variables. Can the number of rings 'misrepresent' the age of the tree? Not exactly. If we know the rate at which the two variables covary, we should be able to reliably infer the age of a given tree by counting the rings in its trunk. I suspect that the same logic applies to other cases in which two variables or states of affairs are not decouplable: even if there is a sense in which one state of affairs carries information about the other, there can be no misrepresentation. How could something misrepresent something else if the information-bearing state is not decouplable from its target? I thus suggest that a representation R can only misrepresent x if R and x can be decoupled.[8]

Now, recall the example of the bacterium swimming up the glucose gradient (sec. 2.1). This behaviour is described as cognitive by the enactivist because we must treat the glucose as meaningful from the perspective of the bacterium. By saying that an organism makes sense of its environment, it follows that the latter is *presented* in a certain way. Following Jakob von Uexküll (1957), via sense-making, one can say that an organism's environment becomes an *Umwelt* (i.e., the world as it is lived by the organism). For the enactivist, the world cannot be separated from the organism, nor the latter from the former. Therefore, cognition understood as sense-making cannot be decoupled from the environment, entailing that it is a non-representational form of cognition.

The correlation between the organism and its world implies that sense-making cannot involve the possibility of decouplability. For the enactivist, cognition is not internal to the cognitive system. If it were, it would be possible for it to be decouplable from the world—even if the organism as a whole could not be decoupled from the environment. Sense-making occurs when an organism interacts adaptively with the environment to prevent reaching viability limits. Consequently, the significance of the world for the organism is neither in the world nor inside the organism. Such significance arises in the relation between the organism and its environment, meaning that it is a relational form of cognition that cuts across the internal-external distinction (Di Paolo 2009).

---

[8] I thank an anonymous reviewer for pressing me into elaborating better the connection I see between decouplability and misrepresentation.

But does sense-making have conditions of satisfaction? If it does, it would seemingly fit into Hutto and Myin's conception of content. The answer to this question depends on what is meant by 'conditions of satisfaction'. To the best of my knowledge, Hutto and Myin do not provide a concrete definition. They do, nevertheless, state the following: "At its simplest, there is content wherever there are specified conditions of satisfaction. And there is true or accurate content wherever the conditions specified are, in fact, instantiated" (2013, p. x). It then seems possible to say that, for Hutto and Myin, there are conditions of satisfaction wherever one can speak of truth or accuracy. I take this to be a *narrow sense of 'conditions of satisfaction'*. This view fits perfectly with a representational view of content. If representationalism were true, then a mental state R could *inaccurately* or *falsely* represent a worldly state of affairs x. As argued above, R can inaccurately (or falsely) represent x because R and x are decouplable. Thus, although Hutto and Myin do not state it explicitly, their definition of content presupposes decouplability. From this perspective, sense-making does not involve conditions of satisfaction because it precludes decouplability.

There is, however, a *broad sense of 'conditions of satisfaction'*. Under this broader definition, to say that something has conditions of satisfaction means that it is normative. For instance, a statement P may be true or accurate because it complies with the norms of truth or accuracy. But normativity cannot be reduced to the norms of truth and accuracy.

Because sense-making emerges from both the autopoietic and adaptive dynamics of the organism, it is also normative. This normativity, however, is not that of validity or truth, which are all-or-nothing concepts. Rather, the normativity inherent to sense-making is that of maintaining the organism within bounds of viability. This norm entails that a given interaction should be taken as positive if it contributes to the organism's self-production and negative if it turns up to be harmful. Whatever affects the organism "acquire[s] a valence which is dual at its basis: attraction or rejection, approach or escape" (Weber and Varela 2002, p. 117). This valence is inherent to the normative dimension of sense-making. Hence, cognition (understood as sense-making) presents the world to the organism while being normative, meaning that it is a non-representational but still contentful form of cognition.[9]

Having shown that sense-making is an example of a form of cognition that is contentful but not representational, I now address whether it is still affected by the hard problem of content. Recall that that problem states that content does not exist in nature and, therefore, it cannot be used as a basis for a naturalist understanding of cognition. That problem only arises if 'content' is conflated with 'representational content'. For Hutto and Myin (2013), the only kind of information that can be found in physical nature is covariance. For instance, the number of rings in a tree trunk covaries with the age of the tree, allowing them to provide information about it. As

[9] An anonymous reviewer notes that Hutto and Myin (2017, chap. 5) introduce the notion of 'Ur-intentionality' to account of cases of basic cognition that, while contentless (in their sense), nevertheless exhibit a biological form of normativity. Is not sense-making simply what Hutto and Myin call Ur-intentionality? I do not think so. A clue regarding the difference between sense-making and Ur-intentionality is Thompson's (2018) claim that phenomenologists have argued for a non-representational form of intentionality. That phenomenological conception of intentionality informs the enactive approach to mind and life (Thompson, 2007). Fully fleshing out the connection between phenomenology and sense-making is beyond the scope of this chapter. Let me note, however, that one of the main features of the phenomenological conception of intentionality is that we should distinguish between an object-directed form of intentionality, and a broader and more basic form of intentionality that can be thought of as a general openness to the world. For a subject to intend a specific object, she must already be open to the world. Sense-making is closer to the latter, more basic form of intentionality than to the object-directed form. In contrast, Hutto and Myin's Ur-intentionality, given that it is 'target-focused' (2017, p. 104), is closer to the object-directed form of intentionality. In sense-making, a meaningful Umwelt is presented to the organism. It is in that meaningful Umwelt that things can become targets of the organism's Ur-intentionality.

already mentioned, those rings do not *represent* the age of the tree; they are not decouplable from the age of the tree, and they do not have conditions of satisfaction (in the narrow sense). Therefore, covariance does not constitute content (Hutto & Myin 2013, p. xv). Since what is lacking is both coupling and conditions of satisfaction (in the narrow sense), the content that is not constituted by covariance is representational content. Sense-making, however, does not involve decouplability, which is, in the first place, the source of the hard problem of content. Sense-making nevertheless involves conditions of satisfaction (in the broad sense) because it is normative. The normativity of sense-making is consistent with naturalism because it is a *biological* normativity. The same idea can be put in LMCT terms. If mind (i.e., sense-making) is continuous with life, and life is a natural phenomenon, then cognition must be a natural phenomenon as well. Notice that this continuity goes both ways: not only mind must be a natural phenomenon, but nature must be thought of as allowing for the rise of (subjective) perspectives such as that of the organism.

If cognition lacked content, it would be unclear whether it would be truly cognition. If the world were not presented at all to a cognitive system, how could it behave adaptively in it? There would be nothing, from its perspective, to which it would have to adapt. Cognition requires content, which, for the enactivist, implies that it involves a perspective over a meaningful world.

The above point puts into question the FEP approach to cognition. As mentioned in section 4.1, the FEP is sometimes presented as a contentless conception of cognition given that it relies on formal constructs that describe organisms as if they encoded probabilistic beliefs. These beliefs must be distinguished from how 'belief' is commonly used in the philosophy of mind and cognitive science (Ramstead et al. 2022, p. 3 n. 2), that is, as involving semantic content. The probabilistic beliefs mentioned in the FEP literature are sometimes presented as a form of generalised synchrony as mutual Shannon information between internal and external states (Kirchhoff and Robertson 2018; Friston 2019). This means that the concept of Bayesian inference as optimisation of probabilistic beliefs refers to the coupling of two covarying sets of variables which, under the FEP, can be construed as if one (internal states) were inferring the other (external states). The tighter the coupling between these set of variables, the less divergence between model and world. However, as Hutto and Myin rightly claim, covariance does not constitute content—not even non-representational content. So, even if there is covariance between the inner dynamics of a cognitive system and its environment, such covariance would not constitute (i.e., would not be sufficient for) cognition. If the FEP conceptualises cognition as inference (i.e., a form of covariance), then it is unclear whether the FEP truly addresses cognition.[10]

---

[10] This point has been seemingly conceded in the most recent FEP literature: "[A]nything that is an organised system generated via emergent dynamics does indeed satisfy some core properties of cognition. […] [E]verything which could be modelled as performing some kind of inference can indeed be understood as performing an elemental sort of inference (as a species of generalised synchrony), without the metaphysical baggage of statements about 'mind' and 'cognition,' are restorative. […] Indeed, the commitment to a principled distinction between cognitive and non-cognitive systems, or living and non-living ones, commits to a sort of *élan vital*, wherein the substance and laws of learning, perception, and action should not be grounded in the same laws of physics as a stone, as though they provide a different, more implacable sort of organisation or coherence of states" (Ramstead et al. 2022, pp. 39-40). From Ramstead and colleagues' perspective, inference is not sufficient for cognition. There is, however, a problem with the consequence they draw from that claim. I believe it is uncontroversial to assume a continuity (in Dewey's sense) between non-life, life, and mind. Just as mind arises from life, life arises from non-life. If vitalists were right and life arose from *élan vital*, then there would be a break between life and non-life, implying a lack of continuity. Dewey's postulate of continuity, however, does not only preclude breaks, but also reduction. By trying to avoid a cognitive analogous to the *élan vital* of vitalism (which would imply a break), the position taken by Ramstead et al. (2022)

5.2 Boundaries and models

I suggest that not everything is lost for the FEP theorist who wants to address cognition, or more generally, LMCT. What must be done, however, is to keep a neat distinction between what the FEP does and what enactivism provides. More specifically, a neat distinction between different senses of 'boundary' must be kept, while clarifying how these senses may relate to one another.

As I have argued elsewhere (Bogotá 2022), we must distinguish between the enactive concept of operational closure, the structural boundaries of an organism, and the Markov blanket formalism used in the FEP literature. Di Paolo et al. (2022, p. 15) rightly claim that the use of Markov blankets in the FEP literature is perplexing and arbitrary. They nevertheless fall victim to the generalised confusion on the use of Markov blankets in the FEP literature by taking them to be the structural boundaries of biological systems. That is the reason why, while attempting to draw a gap between enactivism and the FEP, they claim that some biological processes do not involve a clear structural boundary that would imply conditional independence between internal and external states. Markov blankets, however, are not structural boundaries. They are better understood as statistical boundaries of a system (or particle) *within a statistical model* thereof. A modeller may draw a Markov blanket where there is conditional independence. But that independence is not a property of the world neither at a structural nor an organisation level. There is conditional independence whenever, *in the statistical context of Bayesian networks and probabilistic inferences*, acquiring information about a set of variables will not give *us* any new information about another set of variables. Conditional independence arises from the perspective of an observer that is modelling a set of causal processes using the tools of probabilistic inference. In contrast, the structural boundaries of, say, a living cell are a physical thing (i.e., the cell membrane)—they do not have anything to do with the perspective of an external observer. Similarly, the operational boundaries of the same organism, even if they are conceptualised as an abstract topology of constitutive processes, do not have anything to do with a given model of the organism. In other words, statistical boundaries (i.e., Markov blankets) are boundaries *within a model*, whereas structural and operational boundaries are not.

The reason why the application of Markov blankets within the FEP literature is so ambiguous and perplexing is the lack of a principled way of applying them. The application of a Markov blanket depends upon the possibility of identifying conditional independencies between sets of variables within a probabilistic model of a system. Those independencies are nevertheless relatively easy to draw in cases where there is room to introduce the Markov blanket. Consider the simple system represented in Figure 1. There is conditional independence between μ and η because of the presence of a and s (i.e., the blanket states). If there was a further node that directly influenced, say, η, one could draw a new Markov blanket involving η because there now would be conditional independence between the new node and the set of a, s, and μ, effectively constituting a boundary within the model.

A clear distinction between operational, structural, and statistical boundaries provides the possibility of applying Markov blankets in a principled way. Depending on the interests of the modeller, Markov blankets can be used to represent different kinds of boundaries within a

seems to double-down on the opposite way, suggesting that there might be mind almost everywhere ("anything that is an organised system generated via emergent dynamics does indeed satisfy some core properties of cognition"), implying reduction. As pointed out to me by Julian Kiverstein (a comment for which I am grateful), the enactive approach does have a principled way of distinguishing life from non-life without implying an *élan vital*: adaptive autonomy. Adaptive autonomous dynamics are a non-mysterious form of self-organisation that nevertheless distinguishes life from non-life.

model. Before building a model, however, the relevant (either structural or operational) boundaries of the organism must be identified. If the interest of the modeller is life and mind, then the relevant kind of boundary is the one marked by the operational closure of the organism. It is only after identifying the operational boundaries of the target system that the Markov blanket must be drawn within the model, effectively mapping operational closure within a statistical context that is meant to represent the self-organising dynamics of an autopoietic system (Figure 2). By doing so, it is possible to model the autopoietic and, ultimately, sense-making dynamics of living systems with the FEP formalism. From this perspective, the FEP does not provide an explanation of what either life or cognition are. It simply gives a probabilistic description of their dynamics—a description that would not make sense in isolation from the enactive conception of life and mind. In other words, the FEP is simply a way of modelling or describing some self-organising systems, but not necessarily a proper theory thereof. This point explains the constant use of the 'as if' clause in the FEP literature. It could be said that, from the perspective I am proposing, autopoietic systems can be described *as if* they engaged in active inference while *literally* engaging in sense-making.

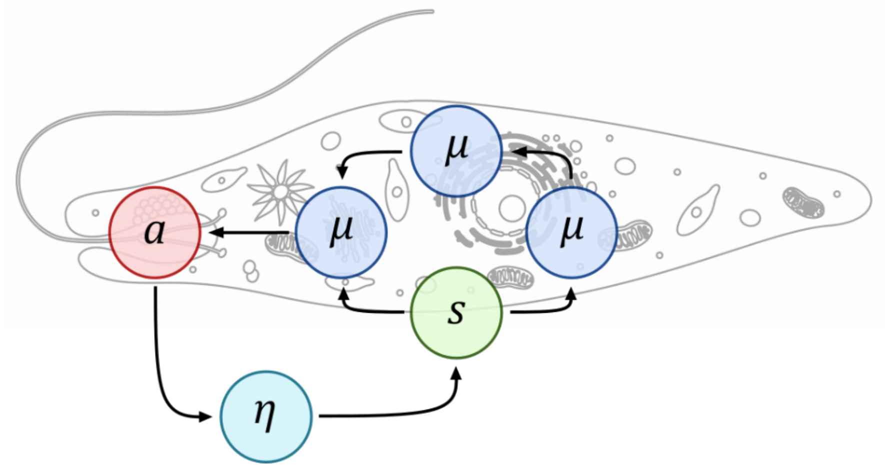

Fig 2. A Markov blanket model of a living cell of the genus Euglena. In this toy example, the recurrent influences between internal (blue, μ), sensory (green, s), and active states (red, a) are meant to represent the operational closure that defines the living cell. Here, the sensory state would represent the surface of the pellicle (i.e., the outer membrane) and the active state would represent the flagella of the cell. Note that the note representing external states (cyan, η) is an oversimplification of what would be several external states, most of which would not directly influence sensory states, meaning that they would not be part of the operational network of the living cell.

As opposed to what van Es and Kirchhoff (2021) suggest, I am not suggesting that the autopoietic dynamics of living systems "can be cast as the hierarchical self-organization of a Markov blanketed *ensemble* of Markov blankets" (p. 6637). My view has a subtle, but important difference. When studying life and mind, Markov blanketed ensembles do not literally self-organise because they are nothing but models of physical systems that self-organise autopoietically. If Markov blanketed ensembles self-organised in the way that van Es and Kirchhoff suggest, it would mean that organisms somehow self-produced Markov blankets. But, as modelling tools, Markov blankets are not produced by organisms; they simply are drawn within a model by an external observer. From this perspective, it would be a category mistake to claim that adaptive behaviour and sense-making are a matter of inference and encoding a probabilistic model. Claiming that we can model autopoietic systems *as if* they engaged in inferential behaviour and encoded a probabilistic model of their environment does not entail that they *literally* do so. Maintaining a clear conceptual difference between the statistical, operative, and structural levels avoids representational and cognitivist readings of life and mind that rely on the FEP (e.g., Wiese and Friston 2021).

One clarification must be made at this point. When making use of a modelling framework like the Markov blanket formalism or the FEP framework, one must be aware of its limitations. By their very nature, models distort and intentionally misrepresent their target systems for the sake of simplicity (Toon 2012; Potochnik 2017). Some phenomena are infinitely complex and having a completely detailed scientific description of them is impossible, as well as undesirable—imagine describing to somebody in every single detail how to go from, say, Exeter to Bristol instead of simply showing them a map which, as a model, oversimplifies the territory it represents. As Angela Potochnik clarifies, "the ultimate epistemic aim of science is not truth but understanding" (2017, p. 71). Depending on what it is that one wants to understand, some models may be more useful than others. In the case at hand, I believe, the FEP does provide a useful model to understand the continuity between life and mind. By modelling life and mind under the FEP, one gets a description of an organism that treats it as if it encoded a probabilistic model of its environment. This perspective does not tell us what life and mind *are*, but it does tell us how life and mind can be *understood* in a scientifically rigorous way so that both are linked from the bottom up. In other words, when guided by an enactive conception of life and mind, which is itself committed to a strong ontological reading of LMCT, the FEP gives us a description of life and mind as continuous which is better understood as a form of the epistemological reading of LMCT (see section 2). Given the roots and connections between the FEP and other approaches in physics (see Friston et al. 2022a), the FEP may even be able to provide an epistemological link between life, mind, and non-living physical systems. This project, however, is yet to be developed (see Levin 2022 for a related project).

The understanding of life and mind that the FEP may provide is nevertheless fundamentally limited. As mentioned in section 4.2, the NESS assumption in the states-focused formulation of the FEP may lack biological plausibility. Organisms behave in such a way that their dynamics can be hardly described as stationary. An interesting challenge to the NESS assumption is that of metamorphosis (Di Paolo et al. 2022, p. 20 n. 9; Clark 2017). Both the kind of behaviour and viability conditions that apply at one life-stage may change so drastically when changing to another that assuming that a metamorphic organism remains at NESS its whole life is an oversimplification whose epistemic usefulness is far from clear. The more complex the organism, the less likely it is to remain at a NESS. As argued by Di Paolo et al. (2022), historicity is a core characteristic of adaptivity and, thus, of organisms.

To address the historicity challenge within the FEP framework one must note that the NESS assumption is only needed in the state-based formulation of the FEP, but not in the path-based one. In their critical assessment of the possibility of integrating enactivism and the FEP, Di Paolo et al. (2022) exclusively focus on the state-based formulation—and understandably so because of the focus that the main advocates of the FEP framework have given to that formulation. In the path-based formulation, the evolution of the dynamics of an active particle can be described as 'Red Queen dynamics' via their tendency toward a strange attractor (Friston et al. 2022b). Red Queen dynamics are those that, despite constantly moving forward and changing (thus allowing for historicity), end up in the same place (thus addressing the fact that biological systems retain their structure). In other words, under the path-based formulation of the FEP, organismic dynamics are described as relatively stable structures that are nevertheless in constant flux, thus providing a link between the FEP formalism and processual philosophy of biology (e.g., Dupré and Nicholson 2018; Dupré 2021), where change is the norm rather than the exception. Put this way, biological systems under the FEP can be described as displaying the kind of historicity that adaptivity requires.

Interestingly, given that the path-based and the state-based formulations are formulations of the same mathematical principle, the former allows for an interpretation of the latter that may allow for historicity despite the NESS assumption. Here, I follow Julian Kiverstein and colleagues' (2022) interpretation of the random global attractor underwritten by the NESS density as the extended phenotype of the organism. This interpretation is justified in the case of biological systems because the ideal set of states in which such systems remain viable cannot only be limited to their internal dynamics, but must also include the external dynamics and how the former are coupled with the latter. Thus, Kiverstein et al. claim: "The states belonging to the system's attracting set will therefore be the subset of all possible states the system can occupy that are highly probable given the system's phenotype and the niche it inhabits" (2022, p. 6). Put this way, the minimisation of free energy implies tending toward both ideal internal and external conditions. But such a tendency may not be achieved. The NESS density implies that, with enough time, the particle will converge in its random global attractor. But there is no need for assuming that such a stationary point will be reached. The random fluctuations of the environment entail that the dynamics of the particle will be constantly changing, even if they are trying to converge in its attracting set, and depending on the strength of such fluctuations, the changes in dynamics needed may be of a greater or a lesser extent. For instance, heuristically, one may find oneself in unfavourable environmental situations such as heat or cold waves that call for immediate action to avoid death. But, importantly, even if one acts accordingly and avoids dying because of the external temperature, that does not mean that one has achieved perfectly ideal states (both internal and external). There will always be room for adaptive behaviour and change. What the NESS density allows is capturing the kind of teleology that characterises organismic behaviour (Ramstead et al. 2023, p. 10). Such teleology is also found at the core of sense-making (Weber and Varela 2002).

Even if the NESS assumption proved to be too restrictive for modelling complex biological systems, there is a degree to which the FEP may still be usefully integrated with the enactive framework to address LMCT. The set of attracting states implied by the NESS density can be thought of as less complex when we are treating simple organisms instead of complex organisms. It may be useful for a scientist (be it a theoretical biologist or a cognitive scientist) to treat a simple life form (e.g., a bacterium) as if it remained at NESS to show that, under these simplifying assumptions, that same system can be modelled as showing adaptive (and hence cognitive) behaviour. In other words, by modelling simple living organisms under the FEP, and using the enactive conception of life and mind as a guide, one may address the continuity of life and mind in the simplest life forms. This kind of approach is characteristic of A-Life approaches to minimal autopoiesis (i.e., the cell, either physical or simulated, whose organisation is the minimally required one for autopoiesis; for discussion, see Thompson 2007, pp. 107-118). Integrating the FEP and enactivism in the proposed way may prove to be useful to address LMCT both epistemologically and ontologically at the level of minimal autopoiesis.

5.3 Interactionist constructivism

To conclude this chapter, I briefly address Kirchhoff's (2018) discussion concerning internalist and externalist explanations of the properties of organic systems.

Recall that Kirchhoff borrows Godfrey-Smith's distinction between internalist and externalist explanations to argue that the theory of autopoiesis is internalist and the FEP is externalist (section 4.1). At first glance, it might seem as if my proposal was externalist because of the emphasis given to adaptivity. This externalist reading would be supported by the complementation of the FEP with the enactive approach. However, it would be a mistake to identify my proposal as either externalist or internalist.

Apart from internalism and externalism, Godfrey-Smith introduces a third kind of explanation that Kirchhoff omits: constructivism. According to Godfrey-Smith, “*constructive* explanations explain environmental properties in terms of organic properties” (1996b, p. 131). The idea is that the explanation of some properties of an organism in terms of environmental properties does not preclude explaining those same environmental properties in terms of organismic properties. Therefore, an externalist explanation can be integrated into a constructive one. Such integration would entail what Godfrey-Smith (1996b, p. 132) dubs an ‘interactionist view’. In interactionist views, the focus is the reciprocal interactions between the organism and its environment.

Rather than externalist, enactivism (both by itself and as complemented with the FEP) is constructivist, specifically an interactionist form of constructivism. Via sense-making, an autopoietic system enacts its world, which means that a meaningful domain arises from its perspective. This meaningful domain is related to the inner dynamics of the system given that its meaning is relative to the autopoietic organisation of the organism. For the enactivist, the properties of what is external to the organism are explained in terms of its organismic self-producing dynamics. Correlatively, the continuous self-producing internal dynamics of the organism only make sense as coupled with the environmental dynamics via sensorimotor interactions with which the organism acquires the nutrients needed for it to preserve its autopoietic operations. As enactivists constantly clarify, operational closure does not entail thermodynamic or material closure (e.g., Thompson 2007, pp. 44-45). Thus, within the enactive framework, both the internal and the external are understood in relation to one another.

The importance of this interactionist constructivist view for LMCT lies in the fact that it implies that life and mind must be understood as a relation between a living system and its environment. Notice, however, that there is an asymmetry in the way in which environment and organismic dynamics relate to one another. Whereas the link from environment to organismic dynamics is causal, the link from organismic dynamics to environment is both causal and phenomenological. It is not only that the organism causally modifies its environment by acting adaptively on it, but also that the environment is disclosed as a meaningful domain from the perspective of the organism. Whereas the organism has a perspective over the world, the world lacks a perspective over the organism. It is precisely the fact that we find this organismic perspective that allows for the introduction of cognition and sense-making into the enactive framework. Thus, the enactivist may argue, the continuity between life and mind is not only operational or functional but also phenomenological (Thompson 2007, chap. 6). This perspective implies that, to properly address LMCT, a theoretical-biological and mathematical perspective on life and mind like the one proposed here is not sufficient given that it must be complemented by a phenomenological one. Such a work, however, is beyond the scope of this paper.

6 Conclusion

I have presented two different ways in which recent philosophers and cognitive scientists have tried to address the so-called life and mind continuity thesis. On the one hand, enactivists have emphasised how adaptive autopoietic dynamics may not only explain life, but also cognition (understood as sense-making). On the other hand, advocates of the Free Energy Principle have suggested that what characterises biological systems is that they can be described as displaying a form of inferential behaviour that allows them to retain their biological integrity. Such behaviour is cast as both sentient (given the sensitivity to sensory information) and intentional (given the fact that the inferences are about the environment).

Despite attempts at integrating both frameworks, recently authors from both camps have suggested that integration seems unlikely. Whereas Kirchhoff (2018) has argued that autopoietic and enactive views of life and mind subscribe to an internalist and contentful conception of biological and cognitive dynamics that is at odds with adaptive behaviour, Di Paolo et al. (2022) have argued that the Markov blanket formalism does not correspond to the operational boundaries of living systems as they are understood within the enactive approach, and that the assumption that biological systems tend toward a non-equilibrium steady state precludes the FEP from properly addressing the historicity that is characteristic of adaptive autopoietic systems.

By noting the difference between operational, structural, and statistical boundaries, I have argued that it is indeed possible to apply the FEP formalism to model the autopoietic and adaptive dynamics of organisms (as they are conceived of within the enactive framework) in terms of active inference. Such integration is desirable for the advocates of the FEP framework because, as I have argued, what allows us to properly talk about cognition is the fact that there is meaningful content. Sense-making, I suggested, is a non-problematic form of contentful cognition that ought to be distinguished from representational views of cognition. In other words, it is possible to conceive of cognitive content as non-representational.

Additionally, I have suggested that the NESS assumption does not need to be seen as problematic as Di Paolo et al. put it. On the one hand, such an assumption is unnecessary in the path-based formulation of the FEP. On the other hand, even if one wanted to keep the NESS assumption, it is possible to make sense of historicity within the FEP by noting that the tendency toward a set of attracting states does not entail reaching a stationary point, but rather displaying teleological behaviour like the one enactivists identify in adaptive autopoietic systems.

Finally, against the idea that enactivism subscribes to an internalist explanatory view on life and mind that would preclude adaptivity, I have argued that it is better conceived of as interactionist and constructivist, meaning that both the environment and the organism must be understood in relation to one another. Importantly, the environment must be understood in relation to the sense-making dynamics of the organism, implying the need for a phenomenological view of life and mind that would complement the theoretical-biological and mathematical approaches that characterise enactivism and the FEP as presented here.

Acknowledgements
I would like to thank Giovanna Colombetti, Sam Wilkinson, Julian Kiverstein, Tom Roberts, Adam Toon, Bartosz Radomski, Lucía Guerrero, and an anonymous reviewer for their thoughtful comments on earlier drafts of this paper. I am also grateful for the feedback I received while presenting talks based on these ideas at the Fourth Bochum Early Career Researchers Workshop in Philosophy of Mind and Cognitive Science, and at the XI Research Workshop on Philosophy of Biology and Cognitive Sciences in Salamanca, Spain. This research was supported by the Ministerio de Ciencia, Tecnología e Innovación (Minciencias) of Colombia (grant 885 of 2020).

References

Aguilera, M., B. Millidge, A. Tschantz, and C.L. Buckley. 2021. How Particular are the Physics of the Free Energy Principle? *Physics of Life* 40: 24-50. https://doi.org/10.1016/j.plrev.2021.11.001.

Allen, M., and K.J. Friston. 2018. From Cognitivism to Autopoiesis: Towards a Computational Framework for the Embodied Mind. *Synthese* 195: 2459-2482. https://doi.org/10.1007/s11229-016-1288-5.

Andrews, M. 2021. The Math is Not the Territory: Navigating the Free Energy Principle. *Biology & Philosophy* 23:30. https://doi.org/10.1007/s10539-021-09807-0.

Bogotá, Juan Diego. 2022. Why not Both (but also, Neither)? Markov Blankets and the Idea of Enactive-Extended Cognition. *Constructivist Foundations* 17: 233-235.

Bruineberg, J., and E. Rietveld. 2014. Self-Organization, Free Energy Minimization, and Optimal Grip on a Field of Affordances. *Frontiers in Human Neuroscience* 8: 1-14. https://doi.org/10.3389/fnhum.2014.00599

Bruineberg, J., J. Kiverstein, and E. Rietveld. 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.

Bruineberg, J., K. Dolega, J. Dewhurst, and M. Baltieri. 2022. The Emperor's New Markov Blankets. *Behavioral and Brain Sciences* 45: E183. https://doi.org/10.1017/S0140525X21002351.

Clark, Andy. 2001. *Mindware. An Introduction to the Philosophy of Cognitive Science.* Oxford: Oxford University Press.

Clark, Andy. 2016. *Surfing Uncertainty. Prediction, Action and the Embodied Mind*. Oxford: Oxford University Press.

Clark, Andy. 2017. How to Knit Your Own Markov Blanket: Resisting the Second Law with Metamorphic Minds. In *Philosophy and Predictive Processing*, ed. Thomas Metzinger and Wanja Wiese, 3, Frankfurt am Main: MIND group. https://doi.org/10.15502/9783958573031.

Colombetti, Giovanna. 2014. *The Feeling Body. Affective Science Meets the Enactive Mind.* Cambridge, MA: MIT Press.

Colombetti, G. 2017. Enactive Affectivity, Extended. *Topoi* 36: 445-455. https://doi.org/10.1007/s11245-015-9335-2.

Colombo, M., and Palacios, P. (2021). Non-Equilibrium Thermodynamics and the Free Energy Principle in Biology. *Biology & Philosophy* 36: 41. https://doi.org/10.1007/s10539-021-09818-x.

Cussins, Adrian. 1993. Content, Conceptual Content, and Nonconceptual Content. In *Essays on Nonconceptual Content*, ed. York H. Gunther, 133-163. Cambridge, MA: MIT Press.

Dewey, John. 1932/2008. *Logic: The Theory of Inquiry*. Carbondale: Southern Illinois University Press.

Di Paolo, E. 2005. Autopoiesis, Adaptivity, Teleology, Agency. *Phenomenology and the Cognitive Sciences* 4: 429-452. https://doi.org/10.1007/s11097-005-9002-y.

Di Paolo, E. 2009. Extended Life. *Topoi* 28: 9-21. https://doi.org/10.1007/s11245-008-9042-3.

Di Paolo, E. A., E. Thompson, and R. Beer. 2022. Laying Down a Forking Path: Tensions Between Enaction and the Free Energy Principle. *Philosophy and the Mind Sciences* 3: 1-39. https://doi.org/10.33735/phimisci.2022.9187.

Dupré, John. 2021. *The Metaphysics of Biology*. Cambridge: Cambridge University Press.

Dupré, John, and Daniel J. Nicholson. 2018. A Manifesto for a Processual Philosophy of Biology. In *Everything Flows. Towards a Processual Philosophy of Biology*, ed. Daniel J. Nicholson and John Dupré, 3-45. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198779636.003.0001.

Friston, K. 2005. A Theory of Cortical Responses. *Philosophical Transactions of the Royal Society B* 360: 815-836. https://doi.org/10.1098/rstb.2005.1622
Friston, K. 2010. The Free-Energy Principle: A Unified Brain Theory? *Nature Reviews Neuroscience* 11: 127-138. https://doi.org/10.1038/nrn2787.
Friston, K. 2012. A Free Energy Principle for Biological Systems. *Entropy* 12: 2100-2121. https://doi.org/10.3390/e14112100.
Friston, K. 2013. Life as We Know It. *Journal of the Royal Society Interface* 10: 20130475. https://doi.org/10.1098/rsif.2013.0475.
Friston, Karl. 2019. A Free Energy Principle for a Particular Physic. arXiv. https://doi.org/10.48550/arXiv.1906.10184. Accessed 26 May 2021.
Friston, K., L. Da Costa, and T. Parr. 2021. Some Interesting Observations on the Free Energy Principle. *Entropy*, *23*(8), 1076. https://doi.org/10.3390/e23081076.
Friston, Karl, Lancelot Da Costa, Noor Sajid, Conor Heins, Kai Ueltzhöffer, Grigorios A. Pavliotis, and Thomas Parr. 2022a. The Free Energy Principle Made Simpler but Not Too Simple. arXiv. https://doi.org/10.48550/arXiv.2201.06387. Accessed 23 September 2022.
Friston, Karl, Lancelot Da Costa, Dalton A. R. Sakthivadivel, Conor Heins, Grigorios A. Pavliotis, Maxwell Ramstead, and Thomas Parr. 2022b. Path Integrals, Particular Kinds, and Strange Things. https://doi.org/10.48550/arXiv.2210.12761. Accessed 9 May 2023.
Friston, K., F, Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald, and G. Pezzulo. 2015. Active Inference and Epistemic Value. *Cognitive Neuroscience* 6: 187-214. https://doi.org/10.1080/17588928.2015.1020053.
Friston, K., J., Rosch, R., Parr, T., Price, C., and Bowman, H. 2017. Deep Temporal Models and Active Inference. *Neuroscience and Biobehavioral Reviews* 77: 388-402. https://doi.org/10.1016/j.neubiorev.2017.04.009.
Godfrey-Smith, Peter. 1996a. *Complexity and the Function of Mind in Nature.* Cambridge: Cambridge University Press.
Godfrey-Smith, Peter. 1996b. Spencer and Dewey on Life and Mind. In *The Philosophy of Artificial Life*, ed. Margaret A. Boden, 314-331. Oxford: Oxford University Press.
Haugeland, John. 1991. Representational Genera. In *Philosophy and Connectionist Theory*, ed. William Ramsey, Stephen P. Stich, and David E. Rumelhart, 61-89. New York: Psychology Press.
Hohwy, Jakob. 2013. *The Predictive Mind*. Oxford: Oxford University Press.
Hohwy, J. 2016. The Self-Evidencing Brain. *Noûs* 50: 259-285. https://doi.org/10.1111/nous.12062
Hutto, Daniel D., and Erik Myin. 2013. *Radicalizing Enactivism. Basic Minds Without Content.* Cambridge, MA: MIT Press.
Hutto, Daniel D., and Erik Myin. 2017. *Evolving Enactivism. Basic Minds Meet Content.* Cambridge, MA: MIT Press.
Kirchhoff, M.D. 2018. Autopoiesis, Free Energy, and the Life-Mind Continuity Thesis. *Synthese* 195: 2519-2540. https://doi.org/10.1007/s11229-016-1100-6.
Kirchhoff, M.D., and T. Froese. 2017. Where There is Life There is Mind: In Support of a Strong Life-Mind Continuity Thesis. *Entropy* 19: 1-18. https://doi.org/10.3390/e19040169.
Kirchhoff, M.D., and I. Robertson. 2018. Enactivism and Predictive Processing: A Non-Representational View. *Philosophical Explorations* 21: 264-281. https://doi.org/10.1080/13869795.2018.1477983.
Kirchhoff, M., T. Parr, E. Palacios, K. Friston, and Kiverstein, J. 2018. The Markov Blankets of Life: Autonomy, Active Inference and the Free Energy Principle. *Journal of The Royal Society Interface* 15: 20170792. https://doi.org/10.1098/rsif.2017.0792
Kiverstein, J. 2020. Free Energy and the Self: An Ecological-Enactive Interpretation. *Topoi* 39: 559-574. https://doi.org/10.1007/s11245-018-9561-5

Kiverstein, J., M.D. Kirchhoff, and T. Froese. 2022. The Problem of Meaning: The Free Energy Principle and Artificial Agency. *Frontiers in Neurorobotics* 16: 844773. https://doi.org/10.3389/fnbot.2022.844773
Korbak, T. 2021. Computational enactivism under the free energy principle. *Synthese* 198: 2743–2763. https://doi.org/10.1007/s11229-019-02243-4
Levin, M. 2022. Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds. *Frontiers in Systems Neuroscience* 16: 1-43. https://doi.org/10.3389/fnsys.2022.768201.
Maturana, Humberto R., and Francisco J. Varela. 1980. *Autopoiesis and Cognition. The Realization of the Living.* Dordrecht: D. Reidel Publishing Company.
Menary, Richard, and Alexander James Gillett. 2020. Are Markov Blankets Real and Does It Matter? In *The Philosophy and Science of Predictive Processing*, ed. Dina Mendonça, Manuel Curado, and Steven S. Gouveia, 39–58. New York: Bloomsbury Academic. https://doi.org/10.5040/9781350099784.ch-003.
Millidge, B., A. Tschantz, and C.L. Buckley. 2021. Whence the Expected Free Energy? *Neural Computation* 33:447-482. https://doi.org/10.1162/neco_a_01354.
Parr, Thomas, Giovanni Pezzulo, G, and Karl J. Friston. 2022. *Active Inference. The Free Energy Principle in Mind, Brain, and Behavior*. Cambridge, MA: MIT Press.
Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco: Morgan Kaufmann.
Potochnik, Angela. 2017. *Idealization and the Aims of Science*. Chicago: Chicago University Press.
Raja, V., D. Valluri, E. Baggs, E., A. Chemero, and M.L. Anderson. 2021. The Markov Blanket Trick: On the Scope of the Free Energy Principle and Active Inference. *Physics of Life Reviews, 39*, 49-72. https://doi.org/10.1016/j.plrev.2021.09.001.
Ramstead, M. J. D., P.B. Badcock, and K.J. Friston. 2018a. Answering Schrödinger's Question: A Free-Energy Formulation. *Physics of Life Reviews* 24:1-16. https://doi.org/10.1016/j.plrev.2017.09.001.
Ramstead, M. J. D., P.B. Badcock, and K.J. Friston. 2018b. Variational Neuroethology: Answering Further Questions. *Physics of Life Reviews* 24:59-66. https://doi.org/10.1016/j.plrev.2018.01.003.
Ramstead, M. J. D., A. Constant, P.B. Badcock, and K.J. Friston. 2019. Variational Ecology and the Physics of Sentient Systems. *Physics of Life Reviews* 31: 188-205. https://doi.org/10.1016/j.plrev.2018.12.002.
Ramstead, M. J. D., K.J. Friston, and I. Hipólito. 2020a. Is the Free-Energy Principle a Formal Theory of Semantics? From Variational Density Dynamics to Neural and Phenotypic Representations. *Entropy* 22: 889. https://doi.org/10.3390/e22080889.
Ramstead, M. J. D., M.D. Kirchhoff, and K.J. Friston. 2020b. A Tale of Two Densities: Active Inference is Enactive Inference. *Adaptive Behavior* 28: 225-239. https://doi.org/10.1177/1059712319862774.
Ramstead, M. J. D., M.D. Kirchhoff, A. Constant, and K.J. Friston. 2021. Multiscale Integration: Beyond Internalism and Externalism. *Synthese* 198: 41-70. https://doi.org/10.1007/s11229-019-02115-x.
Ramstead, M. J. D., D.A.R. Sakthivadivel, C. Heins, M. Koudahl, B. Millidge, L. Da Costa, B. Klein, and K.J. Friston. 2023. On Bayesian Mechanics: A Physics of and by Beliefs. *Interface Focus* 13:1-27. https://doi.org/10.1098/rsfs.2022.0029.
Rowlands, Mark. 2009. Situated Representation. In *The Cambridge Handbook of Situated Cognition*, ed. Philip Robbins and Murat Aydede, 117-133. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511816826.007.

Smith, R., K.J. Friston, and C.J. Whyte. 2022. A Step-By-Step Tutorial on Active Inference and Its Application to Empirical Data. *Journal of Mathematical Psychology* 107: 102632. https://doi.org/10.1016/j.jmp.2021.102632.

Thompson, Evan. 2007. *Mind in Life. Biology, Phenomenology, and the Sciences of Mind.* Cambridge, MA: Belknap Press.

Thompson, E. 2011a. Living Ways of Sense-Making. *Philosophy Today* 55: 114-123. https://doi.org/10.5840/philtoday201155Supplement14

Thompson, E. 2011b. Reply to Commentaries. *Journal of Consciousness Studies* 18: 176-223.

Thompson, Evan. 2018. Review of Daniel D. Hutto and Erik Myin, Evolving Eenactivism: Basic Minds Meet Content. Notre Dame Philosophical Reviews. https://ndpr.nd.edu/reviews/evolving-enactivism-basic-minds-meet-content/. Accessed 1 March 2021.

Thompson, E. 2022. Could All Life Be Sentient? *Journal of Consciousness Studies* 29: 229-265. https://doi.org/10.53765/20512201.29.3.229.

Tschantz, A., A.K. Seth, and C.L. Buckley. 2020. Learning Action-Oriented Models Through Active Inference. *PLOS Computational Biology* 16: e1007805. https://doi.org/10.1371/journal.pcbi.1007805.

Toon, Adam. 2012. *Models as Make-Believe: Imagination, Fiction and Scientific Representation*. New York: Palgrave Macmillan.

van Es, T., and M.D. Kirchhoff. 2021. Between Pebbles and Organisms: Weaving Autonomy into the Markov Blanket. *Synthese* 199: 6623-6644. https://doi.org/10.1007/s11229-021-03084-w.

Varela, Francisco J. 1979. *Principles of Biological Autonomy*. New York: Elsevier North Holland.

Varela, F. J. 1997. Patterns of Life: Intertwining Identity and Cognition. *Brain and Cognition* 34: 72-87. https://doi.org/10.1006/brcg.1997.0907.

Varela, Francisco J., Evan Thompson, and Eleanor Rosch. 1991/2016. *The Embodied Mind. Cognitive Science and Human Experience (Revised Edition)*. Cambridge, MA: MIT Press.

von Uexküll, J. 1957. A Stroll Through the Worlds of Animals and Men. A Picture Book of Invisible Worlds. In *Instinctive Behavior. The Development of a Modern Concept*, ed. Claire H. Schiller, 5-80. New York: International University Press.

Weber, A., and F.J. Varela. 2002. Life After Kant: Natural Purposes and the Autopoietic Foundations of Biological Individuality. *Phenomenology and the Cognitive Sciences* 1: 97-125. https://doi.org/10.1023/A:1020368120174.

Wiese, W., and K.J. Friston. 2021. Examining the Continuity between Life and Mind: Is There a Continuity between Autopoietic Intentionality and Representationality? *Philosophies* 6: 1-17. https://doi.org/10.3390/philosophies6010018.

Wheeler, Michael. 1997. Cognition's Coming Home: The Reunion of Life and Mind. In *Proceedings of the Fourth European Conference on Artificial Life*, ed. Phil Husbands and Inman Harvey, 10-19. Cambridge, MA: MIT Press.