Reconceiving Rationality - Situating Rationality into Radically Enactive Cognition

Giovanni Rolla

This is a postprint version of the paper published at *Synthese*. Please quote from the published version: https://link.springer.com/article/10.1007%2Fs11229-019-02362-y
For open access, see: https://rdcu.be/bQJ3y

Abstract

Rational beliefs and actions are typically evaluated against certain benchmarks, e.g., those of classical logic or probability theory. Rationality therefore is traditionally taken to involve some sort of reasoning, which in turn implies contentful cognition. Radically Enactive (and Embodied) views of Cognition (REC), on the other hand, claim that not all cognition is contentful. In order to show that rationality does not need to lie outside of REC's scope of radicalizing cognition, I develop a Radically Enactive notion of Rationality (RER), according to which rationality is embodied, situated and contentless. For RER, an organism acts rationally insofar as it sustains a proficient interaction with its environment, which in turn requires the coordination of cognitive abilities in accordance with environmental constraints. Rationality is thus distinguished from reasoning, for reasoning is understood as a capacity to coordinate representational cognitive abilities.

Keywords

Radically Enactive Cognition; Contentless Rationality; Adaptive Behavior;

1. Contentless cognition and rationality

A traditional view on rationality takes it to be indissociable from reasoning. Thus, a person's beliefs are rational if their grounding meets certain standards, which are expressed in the reasoning that leads to them. When it comes to practical rationality, a person's actions are rational if there are good reasons she can adduce in their favor. Here is Nozick with the clearest statement of that view:

> Two themes permeate the philosophical literature. First, that rationality is a matter of reasons. A belief's rationality depends upon the reasons for holding that belief. These will be reasons for thinking the belief is true (or perhaps for thinking that it has some other desirable cognitive virtue, such as explanatory power). Second, that rationality is a matter of reliability. Does the process or procedure that produces (and maintains) the belief lead

> to a high percentage of true beliefs? A rational belief is one that arises through some process that reliably produces beliefs that are true (or that have some other desirable cognitive virtue). (Nozick 1993, p. 64).

That is not to say that there is widespread agreement about the standards for rational thought and action. Whereas the predominant position in philosophy identified such standards with classical logic, alternative views emerged recently promoting coherence and probability as more suitable benchmarks (see Schurz & Hertwig 2019). Regardless of that dispute, the assumption that rationality necessarily involves reasoning of some sort has not raised much suspicion. Even the more unorthodox view known as Ecological Rationality, proposed for instance by Todd and Gigerenzer (2012), holds that a theory of rationality should describe the *heuristic* reasonings used by real agents, where heuristics involve following certain environmental cues and ignoring excessive information - which is a matter of reasoning nonetheless.

Now, if rationality requires reasoning, and if reasoning involves the evaluation of beliefs against certain benchmarks (whichever they are), then rationality implies contentful cognition. A cognitive act involves content if it exhibits accuracy conditions - such is the role of representational content postulated by cognitive scientist adherent to some sort of cognitivism. Because beliefs are essentially representational (and arguably conceptual as well), the traditional view makes no room for contentless rationality.

The standard take on rationality provides an immediate fit with old-school cognitivism, for the latter claims that the mind's underlying structure exhibits syntactical and semantical properties, hence its commitment with widespread contentful cognition. Recently, however, the Pragmatic Turn in the cognitive sciences (see Engel et al. 2013) has challenged the postulate that all cognition is representational. Within that movement, the cluster of theories known as Radically Enactive (and Embodied) views of Cognition (henceforth, REC) converges towards the idea that it is possible to explain a basic level of cognition without appealing to representational structures and content more generally - see Hutto and Myin (2013); Hutto and Myin (2017); Chemero (2009); Di Paolo, Burhmann and Barandaran (2017); Gallagher (2017). Thus, Hutto and Myin write: "the biologically basic modes of organismic responding don't involve content, where content is understood in terms of either reference, truth, or accuracy" (Hutto & Myin 2013, p. 78). And:

> A truly radical enactivism - REC - holds that it is possible to explain a creature's capacity to perceive, keep track of, and act appropriately with respect to some object or property

> without positing internal structures that function to represent, refer to, or stand for the object or property in question. Our basic ways of responding to worldly offerings are not semantically contentful. (Hutto & Myin 2013, p. 82).

REC might seem to lack a positive characterization of basic cognition, but this issue can be countered with the notion of *sensorimotor engagement*. When properly exercised, in accordance to environmental solicitations, sensorimotor abilities attune the organism to an environment through sensation and movement.[1] Perceptual states and other basic (contentless) cognitive states thus arise when an organism actively explores its immediate environment through sensorimotor engagement. Crucially, RECers do not (or, at least, should not) claim that *knowledge* of sensorimotor abilities is necessary for basic cognition, for that would smuggle content back into the picture in the form of propositional knowledge, thus rendering REC a variant of Conservative Enactive Cognitive Science - a charge raised against Alva Noë's (2004) sensorimotor enactivism by Hutto and Myin (2013) and Rowlands (2012).[2] The slide into conservative enactivism can be avoided insofar as RECers make a clear stance on *exercises* of sensorimotor abilities as being necessary and sufficient for contentless cognition. As we shall see, this notion plays an important role in the view of rationality presented here.

If REC is correct, representational states are not needed for explaining at least some cases of cognitive activity, which can be done through sensorimotor engagement. The epistemological point made by RECers, in conjunction with a principle of parsimony, allows us to conclude that not all cognition is representational, thus not all cognition is contentful.[3] Although REC acknowledges contentful

[1] Sensorimotor attunement can be understood as the transmission of motor information through perception and perceptual information through movement. Importantly, there is no agreement on whether that notion of *information* is compatible with REC. Hutto and Myin, for instance, straightforwardly reject that basic cognition requires information *processing*, whereas the ecological approach favored by Chemero centers on a Gibsonian notion of information. Hereby I rely only on the deflationary view of information as *information for action*, according to which certain environmental structures inform what actions an organism can perform, given its goals, bodily morphology and abilities. This view does not imply that cognition requires *semantically laden information*, the more contentious notion that comes along with representational content. Information conveyance, thus understood, is not inimical to REC, for it implies only *covariation*: actions undertaken by an organism with specific bodily morphology and abilities covary reliably with the achievement of certain goals in specific environments. See van Dijk et al. (2015) for an argument in favor of an ecological notion of information for enactivism.

[2] One possibility available to REC is to construe the knowledge of sensorimotor abilities as kind of know-how that is irreducible to propositional knowledge.

[3] Could REC go even further and deny the very existence of contentful cognition and content-involving cognitive capacities? Note that REC does not offer a definite scope of contentless or basic cognition, so it could in principle end up including other cognitive activities that so far seem to require representational states. I do not explore this view here, but if a complete radicalization of cognition turns out to be viable, REC has to accept a radically enactive view of rationality as I develop below.

cognition, it is impossible to situate rationality in what REC conceives of as contentless cognition insofar as we subscribe to a picture of rationality as a content-involving capacity. At this juncture, RECers have two options: they can either maintain the traditional view on rationality and argue that rational states emerge once an organism has developed the capacities to engage in contentful cognitive acts, or they can revise the traditional view and offer a novel account that befits the contentless nature of basic cognition.

The first option is endorsed for instance by Engel, for he considers enactivism to be the view that "cognition [...] is grounded in a *prerational* understanding of the world that is based on sensorimotor acquisition of real-life situations." (Engel 2014, p. 219, emphasis added). A few pages further, Engel indirectly explains what he means by "prerational understanding" when he writes that, for Merleau-Ponty and Heidegger, "the relation to the world can be only one rooting in practice, in acting, and practice, in turn, is mediated through the body" (ibid. p. 223). These passages reveal that Engel assumes a traditional view of rationality - for, unlike acting and practice, rationality is not bona fide embodied. Rationality would be a content-involving capacity clearly distinguished from basic, embodied cognition. Besides depending on an explanation for the emergence of contentful cognition, this line of reasoning also has to explain how the emergent features of contentful cognition exert downward causation on basic (contentless) states - for it seems that rational states, which according to this view are contentful ones, must in some way inform and coordinate perception and action, which REC claims to be fundamentally contentless. Thus, it remains to be seen how the task of informing and coordinating contentless processes is to be performed by the contents of rational states, unless something other than their content is doing the work of qualifying or enhancing the behavior of rational creatures. That is an open possibility which I will not explore here.

The second option is in some aspects more economical, but also more revisionist. It is more economical because it does not depend on explaining the emergence of contentful cognition (although, as we will see, it does benefit from it), nor does it have to explain how rational states intervene upon contentless ones.[4] On the other hand it requires rejecting the canonical assumption regarding rationality and reasoning. Susan Hurley, for instance, presents the outlines of such a view. She argues that, if action and perception are mutually constitutive, we should not conceive our

[4] It may seem that the problem of explaining the interplay between contentful and contentless cognition is simply pushed further down the line. But this is not the case, at least when it comes to rational cognition, for the notion of *ability* is central for the explanation of the emergence of contentful *and* contentless states. As I argue in the following section, reasoning is typically a kind of rational contentful cognition, which means it is a matter of exercising content-involving abilities in a proficient way. It is the fact that both contentful and contentless cognition consist in proficient engagements that provides the link between them.

cognitive architecture as primarily vertical in structure, with a dedicated layer for all cognitive matters starkly situated between crude perceptual stimuli and goal-oriented actions (what she famously called the "sandwich conception of the mind"). A better model of cognitive architecture has a horizontal structure, where each layer or module of cognition is dedicated to a specific task and "loops dynamically through internal sensory and motor processes well as through the environment" (Hurley 2001, p. 7), thus ensuing a continuous interplay between perception and action from which cognition emerges. Since rationality is a cognitive capacity, she speculates that:

> Rationality might emerge from a complex system of decentralized, higher-order relations of inhibition, facilitation, and coordination among different horizontal layers, each of which is dynamic and environmentally situated [...] Rationality reconceived in horizontally modular terms is substantively related to the environment. It does not depend only on internal procedures that mediate between input and output, either for the organism as a whole or for a vertically bounded central cognitive module. Rather, it depends on complex relationships between dedicated, world-involving layers that monitor and respond to specific aspects of the natural and social environment and of the neural network, and register feedback from responses. (Hurley 2001, p. 10).

This alternative account of rationality - hitherto largely unexplored by enactivists, let alone cognitivists - changes the explanatory focus from content-involving abilities (having or assessing reasons) to other cognitive abilities, which must be embodied and situated. To be clear, I do not claim that upholding the classical view on rationality is unsustainable within the embodied and enactive framework, for REC does not *entail* a contentless view on rationality. However, divorcing rationality from content is in tune with the main theme of radical enactivism, and developing a contentless view on rationality broadens the reach of the program. Here I choose the latter option and follow Hurley's lead in order to develop in more details what I call RADICALLY ENACTIVE RATIONALITY (RER), which may further the enactivist understanding of the mind. I thereby distinguish reason from rationality. To reason is to exercise content-involving abilities according to certain rules, whereas an organism acts rationally insofar as it maintains a proficient engagement with its environment. So conceived, rationality does not need to deal with contentful cognition, but also does not exclude it. In section 2, I present RER in its details. In section 3, I conclude by considering some objections.

2. Radically Enactive Rationality

What do we call rational?

Ascriptions of rationality provide a useful, yet defeatable, guide to understand the different conceptions of rationality, thus allowing us to better delineate the different primary targets these conceptions have. In order to see how rationality does not always require reasoning, we can begin by distinguishing at least three senses of 'rational'.

The first, more general sense denotes the property of a person of being open to a special kind of normative assessment, either positive or negative. When a person is said to be rational in this sense, the implicit contrast is not with irrational, but with *arational*. Normal adults are rational, whereas small children and non-human animals are deprived of rationality (according to a certain received view), for they simply are not under that kind of normative assessment, be it positive or not. There are different ways to explore whatever makes this kind of normative assessment special, one of them is through the ideas of responsibility and agency: normal adults are rational because they are responsible for what they say and do, which in turn implies they are agents. Small children and some animals, on the other hand, supposedly do not answer for what they do. Even when someone's actions or beliefs violate specific norms of rationality, that is, when we evaluate her negatively in a certain framework; we still consider her rational in this general sense. This sense is important to us here to the extent that the class of things that can be normatively assessed is dependent upon specific conceptions of rationality. Because RER represents a liberal take on rationality, it implies that more things are rational in this first sense comparing to more conservative views.

Alternatively, rationality may be ascribed to a belief held by a person.[5] Generally, the explanation for ascriptions of rationality in this sense is that the person has an ability to perform some sort of reasoning, or at least is sensitive to the relevant reasoning (even if she cannot consciously access it). Human beings are typically capable of reasoning, or of updating their beliefs according to some reasoning, but most other animals (supposedly) are not, at least not to the same extent that we are. Thus, this second sense ties rationality with what I will call here the capacity to reason, viz., the capacity to engage in contentful cognitive acts in accordance to certain rules. The primary bearers of rationality in this sense are beliefs, and secondary bearers of rationality are the individuals that hold these beliefs. A weak requirement for a belief to be rational is that there are reasons that can be adduced to in its favor, whereas a stronger requirement is that a belief is rational insofar as the person can become aware of these reasons (an even stronger requirement would be that the person holds a rational belief at time *t* only if she has access to the reasons for her belief at *t*). These are different

[5] Or by a group of persons, but that does not affect the point made above.

ways to specify the normative realm implicit in the first kind of ascription of rationality, ways that are more demanding than the appeal to the ideas of responsibility and agency we saw above. Something like this underlies McDowell's (1994) idea that the realm of normative assessment is "the logical space of reasons". According to such a view, a person who does not meet those requirements manifests some level of irrationality, despite still being rational in the first, more general sense.

Now, one can also describe *actions* as rational, thus using 'rationality' in a sense that is intimately related to goal-directedness and efficiency, but not explicitly related to the epistemic character of the beliefs (if any) held by an individual at the time or before she acts. To say that an action is rational in this way is to say that it is the most efficient way to achieve a specific goal, so it is evaluated positively in a normative framework. To say that a person acted rationally in a given situation means that her actions were efficient in those circumstances, i.e., the actions were both causally relevant for the achievement of the relevant goal and their performance was less costly than it could otherwise have been.[6] This third kind of ascription is therefore neutral regarding the basis upon which an action is evaluated as rational, for it refers only to the relation between means and ends, and to the efficiency with which the relevant ends are achieved by the available means. Thus, if rationality ascriptions are of any use in understanding the ontology of rationality, it is plausible that in at least some cases a person can behave rationally without backing up her actions with any reasoning whatsoever.

The latter use of 'rational' is also what Hurley has in mind when she writes that rationality "might emerge from a complex system of decentralised, higher-order relations of inhibition, facilitation, and coordination" (Hurley 2001, p. 10). Acting rationally increases the efficiency of independent cognitive abilities by attuning them to environmental solicitations, thus enhancing the organism's engagement with its immediate environment. This is a radically enactive view of rationality (RER), because it emphasises action without calling for content-involving abilities, which is how reasoning is typically understood. RER can be summarised as follows: *an organism acts rationally insofar as it maintains a proficient interaction with its environment, coordinating its cognitive abilities according to environmental constraints.*

Hurley subsequently developed an account of *islands of practical rationality* (Hurley 2003), according to which correctly ascribing reasons for action to an agent requires holism and normativity, for relations between stimuli and responses in goal-oriented actions depend on a network of cognitive

[6] This kind of ascription also reflects the idea that rationality is instrumental, that is, it selects efficient means to an end. But whether the end itself is rational may be harder to fit in this view. One could say that an end is rational insofar as it enables and produces further rational engagements, so that rationality is self-constitutive in the sense that it is an instrument for its on maintenance.

abilities and exhibit a flexible structure. Importantly, Hurley argues that reasons for acting may be context-bound - they may fail to satisfy the generality constraint typically implied by having conceptual abilities. Although her account is more progressive than most views on reason and rationality, rationality is still indissociable from reasoning on that view: rationality may not involve conceptual abilities, but it does require the exercise of (context-bound) content-involving abilities. The radically enactive view on rationality that I explore in the rest of this section is even more progressive than Hurley's latter account, for it allows for the possibility that an agent acts rationally even in the lack of reasons for acting.

Agency and Proficient Interaction

Rationality implies *agency*, for in the relevant sense only an act, something an organism *does*, can be rational. Something that merely *happens with it* does not meet the minimum requirements for behavior.[7] Although we are able to adopt an intentional stance and maybe reliably discern agents from automata, criteria for agency cannot be merely attributive if they have to capture its nature. The theory of autopoiesis, as originally put forth by Maturana and Varela (1980), provides an initial insight into that matter. According to Autopoietic Enactivism (AE), a living organism is *autopoietic* or self-producing. That is, a living system is a precarious network of processes whose organization and physical boundaries are produced by the system itself, in a way that it continually determines its own viability conditions. The production of its viability conditions is known as "sense-making", the development by the system of an internal perspective that evaluates meaningful relations with its environment. An agent makes sense of its relations with the outside world in a bivalent way, it avoids encounters that are detrimental to its self-production and it approaches encounters that enable or facilitate further self-production processes.

AE's core ideas are prima facie compatible with REC and RER, so they can provide a sound basis for agency.[8] Take for instance the case of a Roomba. It is a complex system designed to track and respond

[7] This intuition holds for traditional and revisionist views on rationality alike. One could not easily accept in the traditional view that automata are rational. Consider a hypothetical belief-producing machine that undergoes a series of truth-conducting steps hardwired in its architecture and whose outcome is a reasonable, yet inevitable belief. Intuitively, the hypothetical machine is not rational - plausibly because in that example it is not described as an agent, so its belief-producing processes do not amount to reasoning, even if they are positively evaluated against an accepted benchmark of rationality.

[8] Note, however, that a contentful interpretation of sense-making would render REC and AE incompatible. AE claims that there is a strong continuity between life and mind - for "mind is life-like and life is mind-like" (Thompson 2007, p. 128). Thus, the same organizational structures that explain the emergence of life would also explain the emergence of cognition, which goes all the way up to complex cognitive activity. This would allow for a unified theory of life and mind - which admittedly is one of the most promising features of AE.

dynamically to specific environmental circumstances. Despite its complexity, the Roomba is not an agent because it does not exhibit authorship of the norms that govern its patterns of stimuli-responses - patterns that consequently fall short of fully-fledged agency. Agents, on the other hand, develop what we may call *habits*, that is, self-sustaining patterns of sensorimotor coordination (Barandiaran 2017) that select and improve their own norms of behavior. If AE is correct, a machine that does not produce and improve its own viability conditions is not an agent, a fortiori, it does not act rationally. RER is willing to accept this result - for, despite its liberal stance on rationally, RER would be easily trivialized if everything that exhibits a complex pattern of stimuli-responses turns out to be rational. For that reason, a notion of agency is needed, and RER can inherit it from AE.

As it is usually understood, however, autopoiesis is insufficient to ground sense-making and agency. As Di Paolo (2005) shows, autopoiesis is an all-or-nothing process: it is essentially conservative of the system's organizational structure and has no saying on whether specific relations between the organism and its medium are better or worse for its survival. So conceived, autopoiesis only evaluates interactions as either preserving or destroying the system's structure, thus failing to provide an explanation of how a system is able to *improve* its viability conditions by choosing more proficient interactions with its environment. Agency, on the other hand,

> is an 'asymmetrical concept': there is the actor and that which is acted upon [...] Only when a process is established that is able to regulate this exchange [between the unity or system and its medium] so that in general the result is an improved condition of viability, only then it is possible to speak of a true asymmetry. [...] Behaviour defined not as physical coupling, but as its regulation, is always asymmetrical, has an intentional structure, and can be said to either succeed or fail. It is only at this stage, when the organism behaves, that we may speak of an *agent*. (Di Paolo, 2005, p. 443, emphasis in the original).

Di Paolo complements the traditional theory of autopoiesis with the notion of *adaptivity*, which he defines as a capacity of an organism to improve its conditions of self-production (Di Paolo, 2005, 438). Unlike autopoiesis, which merely describes the structural coupling between the organism and

Insofar as our phenomenology is taken to be the ground for the strong life-mind continuity thesis, and insofar as the phenomenology of higher cognition seems to be contentful, it follows that all cognition (and consequently all life) is contentful (de Jesus 2016). In order to avoid this issue, RECers can either deny the strong life-mind continuity thesis, or they can construe "sense-making" in a deflationary manner, such as the contentless attitude of specifying environmental solicitations and prohibitions that serve as the basis for exploratory actions from which basic cognition emerges.

its medium, adaptivity is an asymmetrical concept intended to capture the capacity of an organism to monitor and to regulate internal changes in order to select the best available course of action. Thus, in this revised autopoietic framework, which provides the scaffolding for RER's distinction between rational action and mere reaction, agency is realized through the exercise of abilities that promote changes within the organism in accordance with environmental constrains. Usually the performance of an action aims the achievement of optimal relations with environmental solicitations, which in turn require feedback sensitivity of its internal states as well as openness to external ones. An optimal relation of this kind is hereby understood as one where internal and external states exhibit a perfect attunement, such as the required pressure in squeezing a sponge in order to feel its softness (Di Paolo et al. 2017). Given that sensorimotor abilities convey motor and sensorial information, and given they are developed, acquired or selected due to their success in a history of interaction, and not merely designed externally as in the Roomba case; the exercise of sensorimotor abilities is a display of agency which is necessary for the achievement of an optimal relation between organism and environment. Coordinating these abilities in accordance to environmental solicitations and constraints enables a complex net of nuanced actions, which cannot be reduced to mere stimuli-responses patterns. Moreover, to the extent that an agent's sensorimotor engagement is bound by its bodily morphology, it follows that RER takes rationality to be *embodied*, i.e., it cannot be sufficiently understood as detached from the organism's phylogeny and ontogeny.

A proficient interaction between an organism and its medium unfolds dynamically in closed loops: what the organism does affects its environmental layout, that is, the way the environment is structured, thus opening way for further actions and so on. Over short timescales, the interplay between organism and environment depends upon anticipation, action, response and adjustments, which in turn lead to further anticipations and so on (over broad timescales, interactions of that kind cause phylogenetic adaptations that allow the organism to thrive more efficiently in its niche).[9] The interruption of a proficient interaction leads to the increase of discrepancy between internal and external states, which eventually leads to systemic death (Gallagher & Allen, 2016). When all interactions are terminated, the organism ceases to exist as an organized system able to self-sustainment, and so it tends to dismantle. Rationality, therefore, is an important self-maintenance tool because it enhances the continuous exchange between organism and environment, which is crucial for survival. It follows that rationality is a capacity that is not only embodied, but also *situated*, for it depends on the attunement between organism and its immediate environment.

[9] To understand the notion of *anticipation* in a radically enactive framework is notoriously tricky because it requires explaining *prediction* without assuming it represents an absent target. In the next session, I present the outlines of an enactive view on prediction.

In order to see these points more clearly, imagine a mammal, *O*, that finds itself in an excessively hot environment, E_1. *O* has the goal of cooling itself, *G*. Achieving *G* is crucial for thermoregulation, which in turn enables other proficient interactions within *O*'s niche. Therefore, accomplishing *G* is the rational thing for *O* to do at that point. Let us imagine that achieving *G* can be done by submerging in a nearby body of water which has a lower temperature than the average environmental temperature. Doing so in E_1 requires the coordination of a subset of *O*'s abilities, say, $\{A_1\}$. Exercising these abilities is causally relevant for *G*, which is preferable than more costly alternatives. But imagine an alternative scenario E_2 where the environmental layout is sufficiently different. In E_2 the water is boiling, so submerging in it may cause permanent damage. Exercising $\{A_1\}$ in E_2 may be detrimental to *O*'s survival, so accomplishing *G* in the circumstances of E_2 may require other actions, such as seeking for shade, which in turn may require the exercise of other subset of *O*'s abilities, say, $\{A_2\}$. So whether *O* achieves *G* depends on whether it exercises $\{A_1\}$ or $\{A_2\}$ according to environmental constraints. The coordination of its abilities turns out to be the key for an adaptive behavior, which enables the organism to cope with a dynamic environment. The more rational an organism is, the more adaptive its behavior and more nuanced the coordination of its abilities according to the environment it inhabits.

Because RER takes rationality to be situated, the unfolding relations between a rational organism and its embedding environment in principle can be operationalized with dynamical systems theory. When the proficient interaction between organism and environment is sustained, there is said to be a causal coupling between the two. A rational organism *O* is causally coupled with an environment *E* in a way that the mapping of the state space of *O* in *E* predicts what variations of *E* would affect the responses by *O* without terminating the interactions between *O* and *E* (in our example above, submerging in a body of water could enable *O*'s interactions in *E*). In dynamical systems theory, these relations are typically described through stochastic differential equations in which the constants that describe *O*'s behavior are variables in equations that describe changes in *E*, and the constants that describe changes in *E* are variables in the equations that describe *O*'s behavior. The unbounded increase of discrepancy or mismatch between *O* and *E*, at a limit, would cause the systemic death of the organism, thus terminating its otherwise proficient interaction with its environment. A maximally rational organism would therefore act as to maintain its causal coupling with *E* in all possible circumstances, thus preventing its systemic death in all possible situations, whereas a minimally rational organism would

be critically vulnerable given any changes in *E*.[10] Consider again our example above: as we would expect, *O* would behave more rationally if it were able to coordinate its abilities in order to achieve *G* in E_1 as well as in E_2, and less rationally if it could achieve *G* only in one of those environments - for, in the latter case, there would be more environmental situations that would lead to a mismatch between *O* and *E*, thus leading to its systemic death.

The Ontological and Conceptual Status

Some further points need clarification. For RER, if an organism coordinates its cognitive abilities in a proficient manner, it is rational, even if minimally so. This may seem profoundly counterintuitive, and I will address this objection in the next section. For now, we have two more pressing questions: first, what is the ontological status of rationality? We began with the assessment of attributions of rationality, but our claim is that certain organisms *act* rationally, not only that we say they do. Secondly, if adaptive behavior is rational, what is the conceptual difference between adaptivity and rationality?

In order to answer the first question, one could say that RER conceives of rationality as a *second-order capacity* because it involves the coordination of other cognitive abilities. But there is a caveat: Acting rationally is not hereby conceived as exercising a separate cognitive ability that works apart from more basic ones. Instead, acting rationally *is* the coordination (selection, adjustment, exercise, withholding) of the organism's abilities according to its current environmental constraints. RER is committed only to an ontologically deflationary view on rationality, for it is not a second-order cognitive state whose job is to oversee lower-level cognitive abilities. RER does not compel us to reify rationality or to postulate a special faculty wherein it operates. Instead, the ontological status of rationality is *adverbial*, that is, it transforms the organism's engagement with its environment, promoting successful interactions which can vary gradually. Thus we can understand how rationality influences behavior without assuming that rationality is contained in a special layer of cognition. Another way to express this idea, which is in tune with Hurley's rejection of a vertically modular cognitive architecture, is that rationality is operative in all layers of cognition. Rationality thus is already present in the basic sensorimotor engagement one holds within its environment. This means that rationality is not a dedicated layer that operates additionally to other, more basic ones, i.e., that it acts upon stimuli and actions that would be, by themselves, rationally neutral. On the contrary,

[10] Because the mapping of state spaces of two or more coupled systems requires the observation of previous states of each system, and given the immense variety of nuanced actions and reactions that each organism is able to realize in an environmental setting, it follows that a systematization of this kind is, despite theoretically possible, frequently intractable, at least for sufficiently complex cases.

rationality is *transformative* of other cognitive operations, because it enhances their success in varying circumstances.

Now, what is the conceptual distinction between adaptivity and rationality? If whenever an organism adapts its behavior - that is, when it improves its viability conditions - it behaves rationally, what is the difference between these two concepts? In order to answer that, consider that an organism behaves adaptively if it avoids encounters that are detrimental to its self-production and approaches encounters that are beneficial to it, and it does so in a way that recognises better and worse courses of action. Establishing better and worse courses of action is a matter of developing its own norms of engagement. The conceptual difference between adaptivity and rationality is that to act rationally is a matter of *how well* the organism conforms to its self-produced norms of engagement, that is, of how efficiently it achieves the ends that constitute a proficient interaction in its environment. Rationality, therefore, has a prescriptive dimension, fot it tells what an organism must do. But because for RER rationality is embodied and situated, it is skeptical about describing specific rules of engagement that serve all purposes. The only one it is able to provide is a consequentialist lemma: achieve the desired ends more efficiently.

Reasoning and Rationality

We began this section by examining how we can describe an action as rational without assuming it is based on reasoning. RER attempts to go beyond merely ascriptive considerations, thus offering an account of the nature of contentless rationality. But even if we distinguish rationality from reasoning, it remains plausible that a person that reasons in accordance with certain standards displays some sort of rationality. Therefore, if RER is to provide a more complete explanation of rationality, it has to specify how contentful cognition, particularly reasoning, can be a rational endeavor. In order to do so, let us take a look at how RECers solve the Scale-up Problem. The Scale-up Problem asks how REC can explain higher levels of cognitive activity, levels that prima facie involve representational content. The challenge is offering such an explanation without compromising REC's tenets of contentless, basic cognition. Showing how contentful cognition enters the picture within the enactivist program is important not only because reasoning implies content, but also because we are able to perform many cognitive tasks in the absence of a target domain. These are known as representational-hungry tasks, i.e., tasks that do not refer to the immediate environment through action, for there are no sensorimotor structures to be aptly explored. Typical cases are acts of remembering what you ate the day before, or of planning what you will do the next weekend if the weather is nice, or envisioning how to write a paper and so on. Clearly, the fact that representational content is at least weakly

decoupable offers a good prima facie explanation of how we can perform that kind of task.[11] So representations may still have a role to play in at least some cognitive levels, but the scope of content manipulation is far more restricted than in classical cognitivism. Acknowledging so prevents REC from becoming an excessively radical view that would reject a priori any possible role for contentful cognition.

A compelling answer to the Scale-up Problem is that contentful cognition emerges once an organism with the appropriate cognitive abilities enters a sociocultural domain with public symbols and rules for their usage (Hutto & Satne 2015; Hutto & Myin 2017). Accordingly, being capable of engaging in contentful cognition - and, more to the point, being capable of reasoning - ultimately depends upon being socioculturally embedded, which is of course plausible independently of REC. Moreover, the fact that psychology of reasoning is able to map cross-cultural differences in reasoning capacities shows that reasoning depends upon one's sociocultural environment (see Nisbett et al. 2008; Choi & Nisbett 2000; Burnett and Medin 2008; Medin et al. 2006). But it also lends indirect support for the dissociation between reason and rationality as advocated by RER, for individuals equally proficient in coping with their immediate environments may have significantly different reasoning capacities. Take for instance the studies conducted by Medin and his colleagues regarding folk-biological taxonomies. The studies involved two groups of fishermen living in North Wisconsin, the European-Americans and the Menominee Indians. Participants were asked to group fishes that "go together by nature" by looking at sets of cards with images of fishes. One relevant finding by Medin and his colleagues is that Menominee Indians favor ecological taxonomies, such as "bottom feeder" and "river fish", whereas European-Americans prefer morphological taxonomies, such as "bass", instead of ecological ones. Now, because both groups live in the same area, catch adult fish with the same methods and have the same goal, which is recreational, nothing indicates a notable difference in their capacity to coordinate the relevant fish-catching abilities in an efficient manner. However, because they do employ conceptual capacities differently in categorizing the fish they interact with, the experiment shows that members of each groups tend to vary substantially when performing inductive reasoning about them. That was shown by asking the subjects about what kinds of fish are more liable to develop a disease (which is, unbeknownst to them, fictional). Plausibly, differences of that kind are due to contingencies and peculiarities of the social systems and institutions within which contentful cognition emerges, which in turn affect the capacity to perform contentful cognitive acts, for instance, arriving at reasonable beliefs.

[11] For an alternative approach, which combines ecological psychology with enactivism, see Kiverstein & Rietveld (2018).

Thus, to say that rationality is manifested in contentless cognitive processes is not to say that rationality is *exclusively* contentless. Accordingly, RER does not exclude that rationality involves the coordination of content-involving abilities, and REC's answer to the Scale-up Problem allows us to see such a continuity in cognitive capacities. The core idea of RER is that acting rationally is the proficient behavior of an organism in its environment, which depends on the coordination and exercise of sensorimotor abilities. If contentful cognition arises in a sociocultural scaffolding, one would expect that a contentful cognitive state is rational if it displays a proficient behavior of a *socioculturally embedded organism* in its *sociocultural environment*, which in turn depends on the coordination and exercise of *content-involving* abilities. This is the case, for instance, of having reasonable beliefs. In normal social environments, an agent that holds reasonable beliefs exhibits a proficient sociocultural engagement, because her beliefs may be aptly deferred to by others, for she can meet conversational challenges regarding their epistemic grounds. These kinds of engagement require the coordination of cognitive abilities associated with reasoning, analogously to the way that a proficient engagement with one's immediate environment requires coordination of sensorimotor abilities in goal-oriented actions. That is why holding a reasonable belief is typically a display of rationality - which does not imply that one can only act rationally if one's actions are grounded on reasons.[12]

In this section, I presented RER as an alternative to the canonical view that rationality entails reasoning. RER aims to characterize rational activity in the same way that REC characterizes basic, contentless cognitive activity, namely, through the exercise of sensorimotor abilities that respond to environmental solicitations and enable adaptive behavior. Accordingly, an organism is rational if it

[12] If reasoning necessarily involves abilities that deal with content, whereas rationality does not, then reason and rationality can come apart even more radically than in the example above. This of course is not to be expected in normal cases of bona fide cognition, where reasoning is a rational display of abilities. In deviant cases, however, things could be different. Consider for instance the case of a brain in a vat (BIV), where a disembodied brain is fed misleading sensorial stimuli about its surroundings, thus forming false perceptual beliefs that lead to largely false mental states. Despite wildly far-fetched, the BIV case supposedly shows that we are willing to ascribe a positive epistemic status to the envatted brain, as long as it exhibits good inferential performances, does not jump to conclusions, avoids fallacious reasoning and so on. But there is also something clearly amiss with the BIV's overall epistemic situation, given that it cannot engage successfully with its environment: all its goal-oriented actions concerning external objects necessarily fail. If the BIV aims to pick up and savor a cup of hot coffee, and if, from its first-person perspective, the BIV seems to accomplish that task, it is only because it relies on misleading impressions of doing so. Now, if we consider more mundane cases of deviant perceptual states, ranging from visual hallucinations in cases of Charles Bonnet syndrome to paranoid schizophrenia, it seems that relying on misleading sensorial information and failing to cope with one's immediate environment do undermine one's overall rational status, even if the subjects are able to keep their ability to reason relatively intact. The BIV case is of course an extrapolation, and its legitimacy can be questioned from a naturalistic standpoint, but I suggest it can be better understood as showing that rationality and reasoning could in principle come apart, and not as simply eliciting internalist intuitions regarding the BIV's epistemic situation, as it is typically construed.

maintains proficient interactions with the environment, which do not call for contentful cognition. Reasoning, on the other hand, can be understood as exercise of content-involving abilities, which, given REC's answer to the Scale-up Problem, only come into play through sociocultural interactions.

3. Objections

RER is too liberal

I now turn to objections to RER. The first one consists in eliciting an apparently undesirable consequence, thus motivating a reductio: if a rational action is the coordination of cognitive abilities, it follows that rationality is gradual, for an organism could be more or less proficient in its environmental interactions. Additionally, by rejecting contentful constraints on rationality, RER becomes vulnerable to the charge of excessive liberalism. This is indeed the case, for RER clearly implies a broadening of the extension of the concept of rationality (considered in the first sense mentioned in §2), given that many organisms achieve and maintain an adaptive behavior in their niches with different levels of success.

Nevertheless, is that truly a problem? First, RER does preserve the intuitive idea that well-functioning adults are cognitively superior to simpler animals such as snails, for instance. Since the snail has a restricted set of abilities, its rational capacity is critically limited, it is only capable of responding and adapting to a small set of environmental changes. Well-functioning human beings, on the other hand, have a wide array of abilities and their combinatory result is potentially infinite, thus enabling innumerable ways of coping with dynamical environments, including ones never encountered before. Therefore, RER matches our intuition about the difference between humans and cognitively simpler creatures: we are *more* rational than other beings - but that is not to say that rationality is exclusively human. Secondly, proponents of a received view on rationality might resist the idea that rationality is non-binary. Given a naturalistic assumption that cognitive capacities should be understood by reference to naturalistically well-behaved structures, by parity of reasoning, proponents of the received view should also resist the idea that there are many non-binary capabilities shared unevenly among members of different species (e.g. being able to jump, to run, to communicate, etc.). Otherwise, they carry the burden of showing why rationality is an entirely different kind of capacity, maybe one that cannot be fully naturalized. Thus, a resolute rejection of a continuum of rational capacities across species is incompatible with a naturalist perspective on rational cognition.

Predictive Processing and Rationality

How does RER fare against the Predictive Processing paradigm (henceforth, PP), the new heavyweight contender of the cognitive sciences? PP takes the brain to be a homeostatic machine that runs subpersonal statistical inferences. Cognition is thus understood as a process of error minimization that occurs in bidirectional neural hierarchies (Clark 2012; Clark 2013; Friston 2009; Hohwy 2013; Hohwy 2016). Thus the brain anticipates the distal causes of sensorial stimuli by generating internal models for the activity of its own layers and minimizes discrepancy between prediction and current stimulus signal. The standard reading of PP claims that anticipation and action consist in the realization of bidirectional *inferences* that broadly conform to Bayesian probability theory. The inferences are bidirectional because they happen top-down through the generation of internal models and bottom-up through the propagation of error signals. In higher layers of neural processing hierarchy, the brain generates models of the environment in order to explain the causal structure underlying the received stimuli. These models function as prior probabilities for the prediction of the activity of inferior layers in the hierarchy, a process that iterates downwards to the most basic level, viz., the input of sensorial stimuli. If there is a discrepancy between prediction and reception, the input layers receive the stimuli and transfer an error signal to the upper layer. The upward processes of transmitting error signal lead to the adjustments of the probabilities determined by the higher layers through an update of the hypotheses that predict the causal structure of the external world. Good predictions increase posterior probability, viz., the probability of a distal cause of sensorial stimulus given the set of prior beliefs and the evidence (the actual signals).

PP's claim to success lies partly in its novelty, for it conceives the brain not as a passive receiver of external stimuli, but as essentially active, for cognition consists in the ongoing generation of hypotheses and the explaining away of error signal. Thus, the upward flow of information only conveys error, whereas the downward flow only conveys prediction for the activity of inferior layers. Indirect evidence for PP is abductive: the paradigm offers the best explanation of well-known cognitive phenomena. Take for instance binocular rivalry.[13] Binocular rivalry can be stimulated in a controlled environment where each subject's eye receives a different image (such as a house and a person's face). Although the perceptual signal remains constant, subjects report seeing a recurrent alternation between the two images, instead of an amalgamation of those images. PP offers an account

[13] Other cases of indirect evidence assessed by Clark are receptive fields effects, end-stopping effects and repetition suppression (see Clark 2013, pp. 191-192). Direct evidence for the neural structures postulated by this cognitive architecture, however, is still inconclusive, largely because no study has been able to identify the neural units responsible for representing the activity of inferior layers (representational neurons) and the neural units responsible for encoding predictive errors (error units) (see Clark 2013, p. 192).

of how their perceptual consciousness alternates from one interpretation of the signal to another in a bi-stable manner:

> [...] the driving (bottom-up) signals contain information that suggests two distinct, and incompatible, states of the visually presented world, for exemple, face at time *t* at location *x* and house at time *t* at location *x*. When one of these is selected as the best overall hypothesis, it will account for all and only those elements of the driving input that the hypothesis predicts. As a result, prediction error for that hypothesis decreases. But prediction error associated with the elements of the driving signal suggestive of the alternative hypothesis is not thereby suppressed, so it is now propagated up the hierarchy. To suppress *those* prediction erros, the system needs to find another hypothesis [...] In Bayesian terms [...] no unique and stable hypothesis combines high prior and high likelihood. (Clark 2016, pp. 35-6).

Whether PP's promise of revolutionary cognitive science is fulfilled remains to be seen. Even if that turns out not to be the case, at this juncture it is important to consider whether it is compatible with RER - for, if PP and RER are indeed incompatible, RER is under threat, given that PP enjoys growing support and may be in the right track to become the main paradigm in cognitive sciences. At first sight, PP does side with most enactivists views when it comes to the rejection of the old-school cognitivist claim that the brain is a passive recipient of crude stimuli. However, in tune with old-school cognitivism, PP conceives of inferential abilities as central to all cognition, therefore it seems to entail pervasive contentful cognition. Commenting that PP implies blurring the limits between cognition and perception, Clark says that: "The process of perception is thus inseparable from *rational (broadly Bayesian) processes* of belief fixation, and context (top-down) effects are felt at every intermediate level of processing" (Clark 2013, p. 190, emphasis added). That idea recalls Helmholtz's (1962) view that perception involves subpersonal statistical inferences, and that is what makes it a rational process in the standard interpretation of PP.[14] Thus, it seems to follow that PP and RER are indeed incompatible.

Nevertheless, there is a non-Helmholtzian reading of PP that allows us to distinguish its putative commitments to pervasive contentful cognition from its precious insights (see Kirchhoff 2018;

[14] Although Clark (2015) does acknowledge that a more radical reading of PP is possible, that reading would still preserve the role of representational content. Representations, however, are not "action neutral" as in the Representational Theory of Mind, but function as codes that make predictions about the world and revise these predictions through the organism's interactions, thus supposedly meeting the enactivist claim about the indispensable role of action in perception and cognition (cf. Clark 2015, p. 22).

Bruineberg & Rietveld 2014; Gallagher & Allen 2016). Kirchhoff argues that a strictly neural hierarchy is not sufficient to explain the prior probability distributions at the onset of cognitive processing, given that the sensory barrage is underdetermined by a potentially infinite set of explanatory hypotheses. Thus we should consider the organism itself a model that is phylogenetically selected to limit the scope of hypotheses that represent the prior probabilities at the beginning of cognitive processing. Thus: "minimization of prediction error is not restricted to the brain alone but involves the entire organism (morphology, action capacities, and so on) and its embedding environment" (Kirchhoff 2018, p. 762). Therefore:

> By 'model' it does not follow that an organism has an internal, representational model of its niche and that it is this model that does all the cognitive work (if you like). Instead, an organism is a model, viz., the causal and statistical regularities reflected in some environment are reflected in some phenotype, i.e., model. (ibid.).

Kirchhoff's argument shows that PP cannot do without the considerations about the embodied nature of cognition, thus restricting the scope of contentful cognition. Another argument against Helmholtzian PP consists in the idea that there is a deflationary sense of 'prediction', according to which the values of two or more variables covary or vary reliably (Anderson & Chemero 2013), thus allowing for an external observer to predict the value of one based on another. So conceived, prediction is not an internal relation holding between two or more variables, such as internal states and external stimuli, but a performance by an external observer based on the available data.

The notion of "active inference", which plays a central role in PP's integration of action and perception under the same set of rules (Friston et al. 2010), can also be understood in a deflationary, contentless way. An agent performs an active inference when he or she "moves its sensors in ways that amount to actively seeking or generating the sensory consequences that they (or rather, their brains) expect" (Clark 2013, p. 186). However, the way this idea is fleshed out according to the Helmholtzian view of PP seems to inadvertently merge two senses of inference (Bruineberg et al. 2018). Traditionally, inference is taken to be a process that relates a set of premises to a set of conclusions according to some rules. Content is clearly indispensable in this traditional sense. But PP conceives of inference as a subpersonal process, usually one that cannot be consciously reached or intervened upon by the individual. There is no compelling reason, other than covert representationalism, to construe processes of this kind as contentful, for once again *content is in the eye of the beholder*: she who interprets how the selection of sensorial input approximates internal predictions (in a deflationary sense) can describe it in inferential terms, modelling it for instance with

the aid of Bayes' law. Thus, active inference can be more naturally construed as a norm for the efficiency of goal-oriented behavior: the agent has to act in certain ways in order for his or her states come to fruition more efficiently (less costly). An external observer could describe the match between input selection and internal states by invoking accuracy conditions, conceiving of prediction in a robust sense that necessarily involves content, but this does not imply the more contentious claim that the generative models themselves display and perform predictive inferences (involving content). Accordingly, input selection allows for a better covariation between internal states and external causes (i.e., a lower discrepancy between agent and environment) which does not imply that the former predicts the latter. Thus, one could say that to act rationally means to achieve a proficient engagement with the environment through prediction and error minimization. What Clark calls "rational (and broadly Bayesian) processes" need not imply contentful cognition, but the selection of input and the coordination of cognitive abilities in world-involving manners, which is exactly the main claim of RER.[15] Therefore, RER is not only compatible with PP, but can prove to be a powerful ally in explaining rational processes within the paradigm.

4. Conclusion

My aim in this paper is to show that we can offer a plausible account of radically enactive rationality, that is, an account whose explanatory focus lies in the proficient interaction held by an organism in its medium, rather than in reasoning as it is in the traditional view. This new account meets our attributions of rational action in everyday discourse, but goes beyond merely attributive considerations and intends to capture the nature of rationality without smuggling content in. The contrast between radically enactive rationality and traditional (contentful) views is not merely a verbal dispute, for denying that content-involving capacities are a requisite for rationality affects what kinds of creatures and behaviors qualify as rational. If we can make sense of contentless rationality and put it to work in the broader enactivist framework, there is no reason to keep the canonical view in place. Thus, if REC truly wants to go radical, why not radicalize rationality as well? There is no need to set rationality outside of REC's program of radicalizing cognition - so rationality can be properly REConceived.

Acknowledgements

[15] Proponents of a Helmholtzian (or cognitivist) PP can still argue that, since it is possible to show that predictive processes may diverge from the distal causes of received stimuli (construed as probability distributions), then these processes may *misrepresent*. See Kirchhoff and Robertson (2018) for an argument that statistical divergence is not enough for representation, so that the information encoded in generative models does not need to be construed as representational.

I am indebted to two anonymous reviewers whose constructive criticisms and suggestions have greatly improved this manuscript. I also thank the audience at the Brazilian Society for Analytic Philosophy for their challenging questions to a previous version of this paper.

References

Anderson, M. L., & Chemero, T. (2013). The problem with brain GUTs: Conflation of different senses of "prediction" threatens metaphysical disaster. *Behavioral and Brain Sciences*, 36(03), 204–205. doi:10.1017/S0140525X1200221X

Barandiaran, X. E. (2017). Autonomy and Enactivism: Towards a Theory of Sensorimotor Autonomous Agency. *Topoi*, 36(3), 409–430. doi:10.1007/s11245-016-9365-4

Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese*, 195(6), 2417–2444. doi:10.1007/s11229-016-1239-1

Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8. doi:10.3389/fnhum.2014.00599

Burnett, R., & Medin, D. (2008). Reasoning across Cultures. In J. Adler & L. Rips (Eds.), *Reasoning: Studies of Human Interfaces and Its Foundations* (pp. 935–955). Cambridge: Cambridge University Press.

Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, Massachusetts: The MIT Press.

Choi, I., & Nisbett, R. E. (2000). The cultural psychology of surprise: Holistic theories and recognition of contradiction. *Journal of Personality and Social Psychology*, 79, 890–905.

Clark, A. (2012). Dreaming the Whole Cat: Generative Models, Predictive Processing, and the Enactivist Conception of Perceptual Experience. *Mind*, 121(483), 753–771.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(03), 181–204. doi:10.1017/S0140525X12000477

Clark, A. (2015). Radical Predictive Processing. *The Southern Journal of Philosophy*, 53, 3–27. doi:10.1111/sjp.12120

Clark, A. (2016). *Surfing Uncertainty*. Oxford: Oxford University Press.

de Jesus, P. (2016). Autopoietic enactivism, phenomenology and the deep continuity between life and mind. *Phenomenolgy and the Cognitive Sciences*, 15(2), 265–289.

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

Di Paolo, E., Burhmann, T., & Barandiaram, X. (2017). *Sensorimotor Life: an Enactive Proposal*. Oxford, New York: Oxford University Press.

Engel, A. K. (2014). Directive Minds: How Dynamics Shapes Cognition. In J. Stewart, O. Gapenne, & E. Di Paolo (Eds.), *Enaction: toward a new paradigm for cognitive science* (pp. 219–243). Cambridge: MIT Press.

Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences*, 17(5), 202–209. doi:10.1016/j.tics.2013.03.006

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301. doi:10.1016/j.tics.2009.04.005

Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biological Cybernetics*, 102(3), 227–260. doi:10.1007/s00422-010-0364-z

Gallagher, S. (2017). *Enactivist Interventions: Rethiking the Mind* (Vol. 1). Oxford: Oxford University Press. doi:10.1093/oso/9780198794325.001.0001

Gallagher, S., & Allen, M. (2018). Active inference, enactivism and the hermeneutics of social cognition. *Synthese*, 195(6), 2627–2648. https://doi.org/10.1007/s11229-016-1269-8

Helmholtz, H. (1962). *Treatise on Physiological Optics*. (J. Southall, Ed.). Mineola, New York: Dover Publications Inc.

Hohwy, J. (2016). The Self-Evidencing Brain. *Noûs*, 50(2), 259–285. doi:10.1111/nous.12062

Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.

Hurley, S. (2001). Perception and Action: Alternative Views. *Synthese*, 129, 3–40.

Hurley, S. (2003). Animal Actions in the Space of Reasons. *Mind & Language*, 18(3), 231–256.

Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism: Basic Minds without Content*. Cambridge, Massachusetts: MIT Press.

Hutto, D. D., & Myin, E. (2017). *Evolving Enactivism: Basic Minds Meet Content*. Cambridge, Massachusetts: The MIT Press.

Hutto, D. D., & Satne, G. (2015). The Natural Origins of Content. *Philosophia*, 43(3), 521–536.

Kirchhoff, M. D. (2018). Predictive processing, perceiving and imagining: Is to perceive to imagine, or something close to it? *Philosophical Studies*, 175(3), 751–767. doi:10.1007/s11098-017-0891-8

Kirchhoff, M. D., & Robertson, I. (2018). Enactivism and Predictive Processing: A Non-Representational View. *Philosophical Explorations*, 21.

Kiverstein, J. D., & Rietveld, E. (2018). Reconceiving representation-hungry cognition: an ecological-enactive proposal. *Adaptive Behavior*, 26(4), 147–163. https://doi.org/10.1177/1059712318772778

Maturana, H., & Varela, F. (1980). *Autopoiesis and Cognition: the realization of the living*. Dordrecht: D. Reidel Publishing Company.

McDowell, J. (1994). *Mind and World*. Cambridge: Harvard University Press.

Medin, D., Ross, N., Atran, S., Cox, D., Coley, J., Proffitt, J., & Blok, S. (2006). Folkbiology of freshwater fish. *Cognition*, 99, 237–273.

Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2008). Culture and Systems of Thought: Holistic versus Analytic Cognition. In J. Adler & L. Rips (Eds.), *Reasoning: Studies of Human Interfaces and Its Foundations* (pp. 956–985). Cambridge: Cambridge University Press.

Noë, A. (2004). *Action in Perception*. Cambridge, Massachusetts: MIT Press.

Nozick, R. (1993). *The Nature of Rationality* (1st ed.). Princeton: Princeton University Press.

Rowlands, M. (2010). T*he New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Massachusetts: The MIT Press.

Schurz, G., & Hertwig, R. (2019). Cognitive Success: A Consequentialist Account of Rationality in Cognition. *Topics in Cognitive Science*, 11(1), 7–36. doi:10.1111/tops.12410

Thompson, E. (2007). *Mind in Life: Biology, Phenomenology and the Sciences of the Mind*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

Todd, P. M., & Gigerenzer, G. (2012). What is Ecological Rationality? In P. M. Todd & G. Gigerenzer (Eds.), *Ecological Rationality*. Oxford: Oxford University Press.

van Dijk, L., Wihtagen, R., & Bongers, R. M. (2015). Information without content: A Gibsonian reply to enactivists' worries. *Cognition*, 134, 2010–2014.