

Prospects of enactivist approaches to intentionality and cognition

Tobias Schlicht & Tobias Starzak


Synthese

An International Journal for
Epistemology, Methodology and
Philosophy of Science

ISSN 0039-7857

Synthese

DOI 10.1007/s11229-019-02361-z



Your article is protected by copyright and all rights are held exclusively by Springer Nature B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".



Prospects of enactivist approaches to intentionality and cognition

Tobias Schlicht¹ · Tobias Starzak¹

Received: 1 September 2018 / Accepted: 9 August 2019
 © Springer Nature B.V. 2019

Abstract

We discuss various implications of some radical anti-representationalist views of cognition and what they have to offer with regard to the naturalization of intentionality and the explanation of cognitive phenomena. Our focus is on recent arguments from proponents of enactive views of cognition to the effect that basic cognition is intentional but not representational and that cognition is co-extensive with life. We focus on lower rather than higher forms of cognition, namely the question regarding the intentional and representational nature of cognition found in organisms simpler than human beings, because enactivists do not deny that more sophisticated cognitive phenomena are representational and involve content. After introducing the debate on the naturalization of intentionality (Sect. 2), we briefly review different varieties of enactivism and introduce their central claims (Sect. 3). In Sect. 4 we turn to radical enactivism in order to focus on the arguments for a thoroughly non-representational, enactive account of perception and basic cognition. In particular, we discuss three major issues: First, what is supposed to replace the representational analysis of perception in a radical-enactive explanation of perception? How does the enactive explanation of perception compare to the best scientific work on the neuroscience of perception? Second, what is—on an enactive account—the function of neural processing in the brain for the generation of perception if not to produce representations? This question is especially pressing since one implication of autopoietic enactivism (accepted by radical enactivists) is that even the simplest organisms, i.e. single-celled organisms, have cognitive capacities (Sect. 5). Since they lack brains and nervous systems, enactivists must specify the (possibly) unique contribution of the brain and nervous system in those animals who have them. In Sect. 5, we evaluate the advantages of an autopoietic–enactive approach to the naturalization of intentionality and end with a suggestion how cognition may relate to intentionality and representation.

Keywords Intentionality · Enactivism · Autopoiesis · Cognition · Representation

✉ Tobias Schlicht
 tobias.schlicht@rub.de

Extended author information available on the last page of the article

1 Introduction

In recent philosophical debates, we are facing a clash between representationalist and anti-representationalist views of cognition. In this paper, we discuss various implications of some radical anti-representationalist views, and what they have to offer with regard to the naturalization of intentionality and the explanation of cognitive phenomena. Our focus is on recent arguments from proponents of enactive views of cognition according to which basic cognition is intentional but not representational and that cognition is co-extensive with life. We focus on lower rather than higher forms of cognition because enactivists do not deny that more sophisticated cognitive phenomena are representational and involve content. After setting the stage by introducing the debate on the naturalization of intentionality (Sect. 2), we briefly review different varieties of enactivism and introduce their central claims (Sect. 3). In Sect. 4 we turn to radical enactivism and focus on the case for a thoroughly non-representational, enactive account of basic cognition. In particular, we discuss the following major issues. First, rather than aiming at a defense of representationalist views, we focus on the question what is supposed to replace the representational analysis in a radical-enactive explanation of perception. Moreover, we ask how an enactive explanation of perception compares to the best scientific work on the neuroscience of perception. In this context, we address Gibson's claim, central for enactivism, that we can directly perceive affordances. Second, since enactivists emphasize the importance of body and world in explanations of basic cognition, what is—on their accounts—the function of neural processing in the brain for the generation of perception, if not to produce representations? This question is especially pressing since one implication of autopoietic enactivism (accepted by radical enactivists) is that even the simplest organisms, i.e. single cells, have cognitive capacities (Sect. 5). Since they lack brains and nervous systems, enactivists must specify the (possibly unique) contribution of the brain and nervous system in those animals possessing them. Therefore, in Sect. 5, we consider a really radical account of cognition starting from the enactivist's emphasis on autopoiesis with respect to the naturalization of intentionality.

2 Naturalizing intentionality

Philosophers introduced “Intentionality” as a technical term for that feature of mental states (like perceptions, beliefs, and desires, i.e. paradigmatically cognitive phenomena), in virtue of which they are *about* or *directed at* something. When I look at the coffee mug in front of me, I am intentionally directed at the mug. The same applies to other mental states like beliefs, imaginations and desires: I can imagine or desire a coffee mug by being intentionally directed at the mug in different ways. But I can also imagine unicorns or desire to meet Santa Claus, both of which don't exist. Although this seems trivial, it raises the question

of how I can bear any relation to something that does not exist. Typical physical relations like kicking or sitting in are unlike intentionality since they presuppose the existence of the things that they relate, which makes intentionality a peculiar phenomenon. Brentano (1874/1995) famously claimed that intentionality is the mark of the mental: all and only mental phenomena are intentional, no physical phenomenon exhibits intentionality. And yet, many philosophers take intentional attitudes to play a vital role in explaining behavior: I pick up the coffee mug *because I believe* there to be coffee in it and *because I desire or intend* to drink coffee. The fact that reference to intentional states enables successful causal explanations and predictions of behavior makes intentionality a good *prima facie* candidate for being a natural phenomenon, what Sterelny (2003) calls the *argument from success*. However, as the ongoing discussions concerning mental causation bear witness to, what seems so plausible at first sight turns out to be a delicate matter once we try to spell out the details. The puzzle of intentionality is thus how it fits into the natural order given its peculiarity.

Finding a place for intentionality in the natural world has occupied philosophers of mind for decades, yet without having reached a clear consensus. Quine (1960, p. 202) forced us to choose between either defending the “indispensability of intentional idioms”—and thus that naturalism is false—or accepting “the baselessness of intentional idioms”. Quine opted for the elimination of intentionality. Fodor, by contrast, defends the indispensability of intentionality and argues that once physicists will have completed the “catalogue of the ultimate and irreducible properties of things”, then “the likes of *spin*, *charm*, and *charge* will perhaps appear on their list. But... intentionality... doesn’t go that deep... If aboutness is real, it must be really something else” (Fodor 1987, p. 97). The challenge for naturalistically minded philosophers has since been to spell out how to be realists about intentionality by explaining it in naturalistically acceptable terms, i.e. “in nonintentional, nonsemantic, nonteleological, and in general, non-question-begging vocabulary” (Fodor 1987, p. 126).

The debate about intentionality is still very much alive, with an abundance of approaches to its naturalization still on the table, including causal (Fodor 1987), structuralist/isomorphic (Cummins 1997; Shea 2014; O’Brien and Opie 2004), teleological (Dretske 1986; Millikan 1989; Neander 2016), instrumentalist (Dennett 1987), fictionalist (Sprevak 2013), measurement-theoretic (Matthews 2007), function-theoretic (Egan 2014), eliminativist (Churchland 1981), and pragmatist (Brandom 1994) positions and no consensus is in sight.¹ To complicate matters, the renaissance of phenomenological philosophy and progress in the empirical mind sciences have led to a proliferation of notions of intentionality, including e.g. affective intentionality (Slaby et al. 2011), phenomenal intentionality (Kriegel 2013), motor intentionality (Kelly 2002; Jacob and Jeannerod 2003; Sinigaglia 2008), enactive intentionality (Gallagher 2017), skilled intentionality (Van Dijk and Rietveld 2017), biological intentionality (Thompson 2007), neural intentionality (Damasio 2011), nano-intentionality (Fitch 2008), and even Ur-intentionality (Hutto and Myin

¹ For recent updates see Montague (2010), Smortchkova et al. (2019) and Dolega et al. (2018).

2017). An adequate theory that illuminates the conceptual relations among these notions and that demonstrates their natural origins is still an important desideratum at the heart of the Philosophy of Mind and Cognitive Science where multiple projects (either implicitly or explicitly) presuppose there being such an account.

3 Representation, cognition, and the enactivist challenge

For a long time, research in cognitive science proceeded under the assumption that cognitive capacities must be explained in terms of computational processes operating on representations. Although views about representations differ, many would subscribe to Carey's (2009, p. 5) view that representations are "states of the nervous system that have content, that refer to concrete or abstract (or even fictional) entities, properties and events". With regard to visual perception, for example, Palmer (1999, p. 77) holds that "a representation refers to a state of the visual system that stands for an environmental property, object or event: it is a *model* of what it represents". Cognitivists like Fodor (1975) conceived of representations as symbols, whereas connectionists characterize them as subsymbolic activation patterns (e.g. Smolensky 1988). All these frameworks work under the assumption that intentionality and representation are more or less equivalent notions. Thus, to say that I am *intentionally directed* at the coffee mug is—on these views—supposed to be equivalent to saying that my perceptual state *represents* the coffee mug: When I see the mug, it features in the representational content of my perceptual state. This equivalence assumption can be traced back at least to Brentano's claim that *all* mental phenomena can be characterized by intentionality, which he described in various terms as the "intentional... inexistence of the object", a "relation to a content, direction upon an object... or immanent objectivity" (Brentano 1874/1995, p. 88). While Brentano himself admits that all these descriptions are "not quite unambiguous", Crane (2003, p. 31) holds that "things are simpler here than they might initially seem" since these phrases, "despite superficial differences between them, are all different ways of expressing the same idea: that mental phenomena involve representation or presentation of the world."

This equivalence assumption is not only widespread in philosophy of mind but also in cognitive science. In their discussion of the bounds of cognition, Adams and Aizawa (2008, p. 31) hold that "cognition involves non-derived representations, representations that mean what they do independently of other representational or intentional capacities" and that "cognition is to be individuated by specific kinds of information-processing mechanisms". This view of intentionality implies that while there may be other representational entities and processes, like linguistic, pictorial and conventional ones like traffic signs, their being representations at all depends on the non-derived representations characteristic of cognition.

Although not uncommon in cognitive science, Ramsey is right in criticizing this way of conceiving of cognition because it *defines* cognitive phenomena as involving representations. Construing the relation between representation and cognition as a conceptual one, Ramsey (2017) argues, is problematic for a variety of reasons, even if one subscribes to a representational theory of mind. *First*, defining cognition

in representational terms unnecessarily constrains psychological theorizing. The scope of a new branch of science is usually defined via its explanatory target, not via theoretical constructs that figure in explanations of those targets. The explanatory target, however, is determined by paradigmatic cases of cognition, like perception and memory (Allen 2017; Rupert 2018), while representations should be seen as theoretical constructs introduced to explain these phenomena. The history of science has shown repeatedly that scientific theories tend to turn out very different from our starting assumptions, and we should be open to this possibility for cognitive science as well.

Secondly, and relatedly, defining cognition in terms of representation undermines their explanatory power, making representations figure in both explanans and explanandum. As a consequence, defining cognition in terms of representation shifts the scope of application of representational theories from explanatory questions (what is cognition? how does it work?) to merely distributive questions (which processes are cognitive, i.e. representational?). The strength of this argument primarily derives from the assumption that defining cognition in terms of representation cannot even be in the interest of advocates of the Representational Theory of Mind, given that they pursue explanatory goals.

Finally, Ramsey argues that we may be able to interpret in representational terms whatever we find, but only if our notion of representation is flexible enough such that it denotes anything that is being causally activated by some input and plays a mediating role in subsequent processing (Ramsey 2007). Adhering to a necessary connection between representation and cognition thus potentially (and following Ramsey (2007) *in fact*) deflates the notion of representation. As Chemero (2009, p. 77) argues, this kind of representational gloss “does not predict anything about the system’s behavior that could not be predicted by dynamical explanations alone”. But if representations do not bring anything to the table, why would we need them in our explanatory enterprise?

With Ramsey’s critique in mind we can derive criteria of adequacy for representational theories of cognition. First, they should be formulated as empirical hypotheses that could turn out to be wrong, which implies that we should not rule out a priori that cognition is best explained in non-representational terms. Secondly, we should evaluate representationalist hypotheses according to how conceptually sound and specific they spell out the notion of representation. For this they must demonstrate how mental representations acquire their determinate content, for example by developing a teleosemantic theory of information (Neander 2016). Furthermore, they must *justify* why a given mechanism or structure in nature should be conceived of as a representation in the first place. This comprises an account of the function of representations, how representations can be identified within a representational system, and how to specify the tasks they perform (Kirchhoff 2011). Thirdly, their success depends on how explanatory these theories are in the light of empirical findings, e.g. how well they explain paradigmatic cases of cognition (see also Rupert 2018). And finally, we have to ask how these theories fare in comparison to alternative, non-representational views of cognition.

This is exactly one entry point for radical enactive accounts of cognition. Varela et al. (1991, p. 173) characterized the enactive approach in general as the claims that

“(1) perception consists in perceptually guided action and (2) cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided.” These claims are characteristic of all enactive proposals, including the *sensorimotor* version (O’Regan and Noë 2001), the *autopoietic* version (Thompson 2007), and the *radical* version (Hutto and Myin 2013). All of them are supportive of the autopoietic aspect of enactivism to some extent (cf. Noë 2009, p. 42; Hutto and Myin 2013, p. 33), to which we return at length in Sect. 5.

The claims that (a) basic cognition should be conceived of as representational and (b) that positing representational content with truth or accuracy conditions adds anything to the explanation of basic forms of cognition like perception and action has been rejected most strongly by radical enactivists like Hutto and Myin (2013, 2017). They do not doubt that—to the extent that human thought involves concepts, linguistic capacities and public symbol systems—human thinking involves intensional content (in the sense of propositions having truth conditions) (Hutto and Myin 2013, p. x). However, they take representational cognition to be a very special case that only arises with the mastery of socio-cultural practices and that is the result of internalized public representations (Vygotsky 1997). Thus, they criticize traditional views as advocating a “narrow vision of intentionality” (Hutto and Myin 2017, p. 98) that takes for granted that intentionality in all its forms must be modeled on the most sophisticated form of intentionality, namely human thought. Thus, Hutto and Myin (2017, p. 95) propose a “radical enactivism” (REC), according to which we should “think of the most primitive form of intentionality... in noncontentful, nonrepresentational ways” but still as an “attitude directed toward an object”. Intentionality is then no longer a feature of contentful mental states that represent but “an attitude of the whole organism expressed in their behavior” (Hutto 2008, p. 57). The idea is to “disentangle” *directedness* and *aboutness* that figure in Brentano’s characterization of intentionality as two sides of the same coin, and make the latter—contentful representations with satisfaction conditions—dependent on linguistic capacities. This bold move leads Hutto and Myin to the radical claim that there are forms of all basic cognitive capacities including perceiving, imagining and remembering that can be exhaustively explained without positing mental representations:

REC’s signature view is that such basic forms of cognition do not involve the picking up and processing of information that is used, reused, stored, and represented in the brain. The usual form of what REC calls basic, contentless cognition is nothing short of organisms actively engaging with selective aspects of their environment in informationally sensitive, spatiotemporally extended ways. (Hutto and Myin 2017, p. xiv)

Furthermore, enactivists also reject the methodological move, implicated by cognitivism and connectionism, to focus on the brain as the “locus” of cognitive processing. While brains do play a role for cognition, focusing too narrowly on brains cannot account for basic cognition, or so they argue. The problem they see is that of continuous reciprocal causation—whenever “the causal contributions made by components of a system partially determine and are partially determined by causal contributions of other systemic components” which makes it “impossible to assign a *specific* subtask to an *identifiable* subsystem within a larger system” (Kirchhoff

2011, p. 2). Enactivists typically argue that explanations of cognitive phenomena like perception need to take into account the complex dynamics between brain, body and environment, since the perceiver—as embodied cognitive agent—is always necessarily endowed with a characteristic set of sensorimotor capacities and coupled to a specific environmental niche, which enables and constrains the agent’s possibilities for perception and action. Our knowledge concerning the relation between brain and cognitive processing would then be much overstated. For instance, neuroscientists like Engel (2010) and Fuchs (2018) support the enactivist claim that cognitive brain processes must always be conceived in their dynamical relation to the body and environment.² Engel (2010, p. 226) recognizes, for example, that the ‘pragmatic turn’ in cognitive science requires us to see the function of neurons “with proper reference to other subsystems and the actions of the whole cognitive system”. But rather than merely downplaying the role of brain processes in the dynamics of brain, body and world, enactivists should also focus on providing a positive vision of what brain processes do contribute to cognition in those organisms who have brains and nervous systems.

The claim shared by many enactivists that *intentionality is a feature of whole embodied agents, i.e. organisms*, not of representational mental states that can be localized in the brain (Hutto 2008, p. 57; Thompson 2007, pp. 13, 159–160) has important and wide-ranging implications for various philosophical discussions about kinds of minds and the nature of cognition as we will see. In what follows, we do not intend to dwell on objections to representational accounts or defend representationism against such objections. Rather, we would like to discuss what enactivists offer as a positive alternative with respect to the following questions: What are the natural sources of intentionality? How should we explain basic cognitive phenomena if not in representational terms? In the following section we first turn to the implications of radical enactivism with respect to the explanation of perception.

4 Radical enactivism, affordances, and cognitive processing

Enactivists stress the importance of explaining basic cognition in terms of the dynamics of brain, body and world. Embodied agents are viewed as being coupled to an environmental niche, which affords certain actions. The notions of *coupling* and of *affordance* are important conceptual tools in enactivist accounts of perception. O’Regan and Noë (2001) make extensive use of them, Hutto and Myin (2013, ix, xvi, 8, 16) prefer to speak of “worldly offerings” rather than of affordances. Although the notion of coupling is rarely defined rigorously, the general idea is quite clear: If we want to explain a given creature’s cognitive abilities, we have to pay close attention to its immediate environment because cognition is always situated in the sense of taking place in the context of an embodied agent being embedded in a given niche. The features of this environment and the features of the agent mutually

² Further work in this area includes Anderson (2014) and the collection of papers edited by Wilson et al. at <https://www.frontiersin.org/research-topics/1713/radical-embodied-cognitive-neuroscience>.

constrain each other such that agent and world constitute one complex *coupled* system. Yet, while the idea may be intuitively clear, the notion of coupling is never defined, neither in Varela's work nor in subsequent writings by enactivists. De Jaegher and Di Paolo (2007) make an attempt to spell out some conditions for coupling to emphasize the strength of the relationship between agent and world, but they do not inquire into the preconditions of coupling on a sophisticated level. They make clear that taking the same subway to work and exchanging body heat is not sufficient for two agents to count as one coupled system. But stronger forms of interaction—either between agent and world or between two agents—may be in need of specific conditions to be in place or may be in need to be actively achieved by (at least one of) the agents such that enactive uses of this notion are underspecified (see Martens and Schlicht 2017). Because of this underspecification we focus on the correlative notion of affordance and its role in enactive accounts of basic cognition. In this section we would like to lay bare and discuss some of the implications we take to be important for an evaluation of the viability of the enactive approach.

4.1 Perceiving and acting on affordances

By way of a positive story, radical enactivists claim that “experiencing organisms are set up to be set off by certain worldly offerings—that they respond to such offerings in distinctive sensorimotor ways that exhibit a certain minimal kind of directedness and phenomenality” (Hutto and Myin 2013, p. 19). The features of the environment which are relevant here are often identified with what Gibson (1986, p. 127) called affordances, i.e. “what it [the environment] *offers* the animal, what it *provides* or *furnishes*, either for good or ill.” In fact, enactivists rely heavily on Gibson's ecological approach to perception.³ Gibson focused on the informational basis of perception in the environment and claimed that the structure of light in the environment is specific enough for the perceiver to have immediate access to—and pick up directly—how things are (cf. Palmer 1999, p. 53). Furthermore, Gibson claimed that perceiving affordances is not only direct, affordances are perceived *first*. Chemero observes with respect to Gibson's theory:

If perception is direct, no information is added in the mind; if perception also guides behavior, the environment must contain sufficient information for the animal to guide its behavior. That is, the environment must contain information that specifies opportunities for behavior. In other words, the environment must contain information that specifies affordances. (Chemero 2009, p. 106)

The claim is that perceiving an apple's or a coffee mug's affordances is immediate in the sense that it does not require *any preceding* processing. Nor is any *subsequent* processing needed for an agent to generate appropriate reactions to the possibilities provided by the environment. In this dynamic process of being sensitive

³ It is often stressed that Gibson shared with enactivists the hostility to mental representations in the explanation of perception. It should be said, however, that his theory did not address the mechanistic basis of perception in the brain at all.

and reacting to the world's affordances, the agent's brain does not need to produce representations of its environment.

Hutto and Myin (2017, p. xiv) claim that rather than representing the environment, the brain is engaged in “anticipating, influencing, and coordinating responses in a strong, silent manner”. In a similar vein, Gallagher says:

In contrast to the standard conception of the brain making inferences (as found in classical computational accounts and the more internalist predictive processing accounts) the enactivist view is that the brain, as part of the body-environment system (not only regulating body, but regulated by the body and its affective processes) is, as Jesse Prinz puts it, ‘set up to be set off’ (2004, p. 55) by prior experience and plastic changes. The brain works as an integral part of the organism which, as a whole, responds dynamically to environmental changes. It's not clear that this is equivalent to the notion of ‘active inference’ in predictive processing accounts, but from the enactivist perspective, it may be the best way to think of how the brain works. (Gallagher 2017, p. 24)

These passages illustrate that enactivists mostly talk about what the brain is *not doing*. They are typically neither very specific nor very committal with respect to the brain's contribution to perception and other basic cognitive phenomena. Of course, assigning the brain an important role for cognitive processes does not imply that interpreting brain processing in representationalist terms is adequate. However, enactivists should not remain silent about the details, but at least attempt to account for what and how the brain contributes to cognitive processes. Merely stating that the brain is set-up to be set-off by the environment is insufficient. In the following sections, we will have a closer look at both, the *no-preceding-processing* and the *no-subsequent-processing* claim, and argue that radical enactivism is either less radical than often assumed, or not supported by the evidence from behavioral—and neuroscience.

4.2 The *no-preceding-processing* claim and vision science

Consider seeing a coffee mug in front of you. Its design affords certain simple actions like grasping and picking it up for creatures like us with arms and hands of the right size and shape. When claiming that we can *directly* see this or that, enactivists typically argue on a *phenomenological* level, i.e. on the level of conscious experience: it certainly *feels* like I can directly see the mug and directly exploit its affording to be grasped and picked up by me. Similarly, it certainly feels like you can see your red car directly instead of merely seeing a big red mass of steel and it certainly feels like you can directly see your daughter's joy in her facial expression, her gestures and her posture (Gallagher 2008). Furthermore, some versions of enactivism hold that you perceive the coffee mug as a three-dimensional object with a backside and an interior although these features are currently out of sight. They argue that you do not need to *infer* them but that you can *perceive* these features “as absent”, as Noë claims: “The sense of their presence that I enjoy is manifestly *visual*. We have a visual sense of the presence of the hidden parts of the things we see.” (Noë 2014,

p. 95). Again, *enjoying* it suggests that what Noë is talking about is the level of conscious experience. However, saying how something phenomenally feels to me leaves open the nature of the processing in my brain, before and during my grasp. Sometimes it seems that enactivists take the phenomenological level as the only level of explanation that is philosophically relevant. As Gallagher writes:

Of course I am not denying that when I see my car all kinds of complex processing is going on in my brain. The visual cortex is processing information about shape and about color, and so forth, in a distributed fashion. Neuroscientifically I may in fact be able to carve up this early processing in the visual system where shape and color are processed in V1 and intermediate visual areas... More processing in the inferior temporal cortex, and “top-down” processes may focus and integrate the sensory information. Perhaps the neuroscientist even thinks there is a binding problem and is motivated to ask how all this distributed processing gets glued together to form a coherent perception of my car. But that is a problem for the neuroscientist; not for the perceiver. (Gallagher 2008, p. 537)

It may not be a problem for the perceiver, but it surely is a problem for the philosopher aiming to explain perception. Thus, if enactivists restrict their claim about direct perception (*no-preceding-processing*) to the phenomenological level, as the Gallagher-quote above suggests, then their theory is not in conflict with classical representational accounts of perception, which aim to explain perception in the light of what (neural) processing contributes to perception. Moreover, restricting our theory of perception to how it feels to perceive something is unsatisfying for everyone who is interested in the underlying mechanics of perception. Enactivism understood this way does not appear to be particularly radical.

If, however, the *no-preceding-processing* claim is to be understood as a radical alternative to representationalist accounts of perception, the thesis should be evaluated in the light of the best neuroscientific evidence concerning perception. The question then is whether affordances are perceptually basic or whether perceiving them presupposes prior processing of some kind. Following Stephen Palmer, who in his book *Vision Science* (1999) defends an information-processing account of visual perception, it seems pretty clear that affordances are *not* perceptually basic. Rather, perceiving affordances presupposes perceiving much more basic features that have to be “glued” together to yield a percept of a coffee mug that can be picked up, or of my car that I can drive. In Palmer’s view, “the *final* stage of processing must be concerned with recovering the functional properties of objects: what they afford the organism...” (Palmer 1999, p. 91, our italics), suggesting that there are prior stages. He concedes to Gibson that in some cases it is possible for us to directly perceive (i.e. to directly ‘pick up’ from information provided by the optic array) an affordance without first categorizing the object (as a mug, chair, or car, say), but he argues that this only works for affordances of objects that have to do with the objects’ physical structure or surface properties. Most functions of the objects around us, however, are not available from visible information alone (Palmer 1999, p. 411). In many countries, a mailbox and a litter box share many physical properties although only one of them affords mailing letters. So it makes an important difference in which slot

I put my envelope. Grasping the affordance of the mailbox depends on knowledge about the mailing system, about the mailman having to remove letters from the box etc. Likewise, grasping the affordances of a litter box presupposes knowledge of the larger context related to litter being picked up etc. In such cases, grasping an affordance is not independent of prior categorization. Thus, perceiving affordances that are somewhat removed from the visible properties of an object requires prior categorization, and even if perception of some functional features of objects may be possible before categorization, this still presupposes earlier image-based, surface-based and object-based processing stages (Palmer 1999, pp. 85–92). And if this is right, then perceiving affordances presupposes perceptual representation of more basic features like shape, color, and so on. In order for the brain to “support” vision by “enabling exercise and mastery of sensorimotor contingencies” (Engel 2010, p. 225), i.e. of the ways in which bodily movement determines what we see, it must first process information about the object with which one is interacting right now.

Moreover, Palmer submits that in order to be able to “see”, visually enjoy, absent properties of an object, we first have to process the visible ones:

After perceiving the structure of visible surfaces, perceivers may make further inferences... but anything beyond visible surfaces is at least one step further removed from the retinally available information. (Palmer 1999, p. 201)

That is, by Palmer’s lights, Noë’s assumed perceptual experience of absent features of an object presupposes information processing of the visible, non-absent features that apply to the current profile of the mug from where I stand and look at it. Leaving aside the dispute between Palmer and Noë whether experiencing the absent features is a result of perceptual processing or whether it requires inference,⁴ we can ask whether perceiving affordances is more like seeing absent features than perceiving present features. That is, does perceiving affordances presuppose information processing of features that are simpler or on lower levels? The answer, we suggest, is yes. The culminating evidence from neuroscience and neuropsychology (especially from pathological cases) concerning the complexity of the information processing involved in vision alone is perplexing.

Gallagher reviews and acknowledges this complexity in the quote cited above but thinks it does not affect a *philosophical* theory of perception. In order to yield a perception of my car as a red drivable object, brain processes must at least put together shape, color, (possible) movement and other features of my car that are processed in a distributed and parallel fashion, since these features trigger different neural mechanisms. Pathological cases show that people can be impaired with respect to a very specific aspect of vision, e.g. color or movement, while remaining unaffected in other aspects of perception. Neither of these separately processed aspects by themselves affords sitting in or driving or much else. Note that enactivists cannot simply buy into the story told by the opposition since this involves information-processing

⁴ According to the recently popular predictive processing accounts, perception itself is a highly inferential process (Clark 2016; Hohwy 2013). On such accounts, it makes no difference whether we called the process inferential or visual, of course, because this comes down to the same thing. See Gallagher (2017, ch. 6), for critical discussion.

of representations. But in the absence of any positive enactivist account of what the brain must do in order to yield a percept of my car based on the processing of these features, the most informative account at hand (so far) holds that the brain must represent these features and perform computations on these representational units which, once combined, result in a full-blown representation (see Treisman 1988). Thus, the *no-prior-processing* claim poses a dilemma for radical enactivists: either their view is about the phenomenological level of perception only, in which case it is much less radical than advertised; or it involves a claim about what brains actually do during perception. In this case their theory is in fact radically different than representational theories of perception—but it is not supported by neuroscientific evidence.

4.3 The *no-subsequent-processing* claim and behavioral flexibility

In the last section we argued that enactivism is in trouble if it uncritically adopts Gibson's claim that affordances are perceived directly without *preceding* processing. It is insufficient to insist that perception emerges only as a personal-level phenomenon if it is at the same time conceded that the brain contributes to the dynamics of brain, body and world in a significant way. In this section we focus on a further, maybe even more radical, claim of radical enactivism, the claim that *no further* processing takes place to generate appropriate reactions to what is being perceived either. The brain is simply *set-up to be set-off*. Taken together, these claims are supposed to draw a picture of cognition according to which agents are tuned to their environment in a way that action-selection, i.e. the generation of appropriate responses to what a given environment affords, is not dependent on further processing. A problem with this view is that it suggests a rigidity of stimulus–response that is particularly *uncharacteristic* of human and lots of animal behavior.^{5, 6} Recall that cognitive scientists have introduced mental representations precisely to account for behavioral flexibility and because certain behaviors exhibit aspects that *seem to depend* on how the organism in question *represents* the world. Examples include explanations of spatial navigation in terms of cognitive maps, mental time-travel like episodic memory (Dally et al. 2006) and future-oriented behavior in many food-caching birds like scrub jays (Raby et al. 2007) and in chimpanzees (Bourjade et al. 2012), or spontaneous problem solving under novel circumstances (Mendes et al. 2007; Hanus et al. 2011), to mention just some examples. Even the behavior of simple animals, like desert ants which, after a nonsystematic search for food, return to their nest in a straight line (Wehner 1999), suggests the existence of some internal mechanism representing the location of the nest. In “representation-hungry” (Clark and Toribio 1994) cognitive tasks like these,

⁵ This is true even if we take into account ‘soliciting affordances’ (Siegel 2014) which, so to speak, ‘pull’ actions out of us automatically.

⁶ Alternatively, if we do not understand the phrase in terms of rigidity, it does not seem to be at odds with classical cognitivism: Of course, in a very general sense, the brain is at any given time set up to be set off in some way. Understood in this loose way, however, classical cognitivism seems to be explanatorily stronger (given that explanatory power is understood as how good an explanation is under the assumption that it is true (see Ylikowski and Kuorikoski 2008), since it offers an account of what the brain is doing once it is set off: it computes representations which carry information about how things stand in the world.

involving amongst others memory, future-oriented behavior, or spatial navigation, the animal is directed at something in its absence: hidden nuts, a predicted craving for a specific kind of food, an out-of-sight nest. However, if these things are out of sight, they cannot set-off any brain, no matter how it is set-up. Positing internal representations (understood as decoupled states that carry information about, or *stand-in* for some environmental feature) can account for this flexibility. Pointing to the (residual or substantial) problems of representational accounts surely is important, and representationalists need to address these problems. But likewise, alternative accounts must match their explanatory power such that ongoing debates can come to a verdict on which direction is better suited to the task at hand. In the absence of a satisfying positive enactivist account, the problems that may beset representationalism do not yet justify abandoning this serious contender in the light of its explanatory power with respect to flexible and stimulus-independent behavior and the role brains may play in this context. While traditional accounts may downplay the role of other factors (like the non-neural body and the environment) that are important for cognitive processing, they do have a much more well-defined account of the brain's contribution. Enactivists still owe us a determinate and committal characterization of what the brain is doing in cognitive respects.

It is striking that not even neuroscientists who support enactivist claims have provided sufficient characterizations of brain processes in this context. For instance, Engel (2010), who rejects a representationalist characterization of neural processing, sees the need for developing a new interpretation of neural processing in the context of his proclaimed “pragmatic turn” in cognitive science. But he remains vague when he claims that “the brain enables us to see” (2010, p. 225) and that neural states “support the capacity of structuring situations through action” (2010, p. 226). He stresses that neural activity is in the service of guiding action and that his notion of a “directive”—replacing the “representational”—mind does not simply refer to an internal state of the brain; but he does hold that “directives correspond to functional roles of neural states” and that “neural activity patterns support and partially implement directives” (2010, p. 230). Again, this does not address the question *how* neural processes contribute to enable action. As far as the job description of neural processes is concerned, Engel's focus is on the temporal synchronization of neural firing in response to perception of aspects of one and the same object. He emphasizes that a significant body of evidence suggests, “that synchrony may provide a dynamic binding principle for structuring and selecting sensorimotor couplings” (2010, p. 235) and that neural firing patterns should be interpreted in an “action-oriented” sense as implementing “procedural knowledge of sensorimotor contingencies” (2010, p. 237). But this leaves unaddressed the worry that the system may first have to “identify” or even “categorize” the object in question. As long as the job description of the neural information processing resulting from encountering a given object remains vague, the representational view has the advantage of providing such a job description, namely, to use the stand-in (representation) to enable action and other kinds of further use. This can be sustained even if perception is seen in the service of action (Clark 1997), since pushmi-pullyu-representations (Millikan 1995) have the dual character of describing the relevant aspect of the environment and of prescribing particular actions given the setup of the agent in question. Moreover,

such stand-ins can then explain cognitive activities in the absence of the stimulus in question that is represented in the brain.

In any case, in order for enactivists to provide a full viable alternative to representationalism, their characterization must find a way out of the dilemma formulated in the last sections. It should avoid both the *no-preceding*- and the *no-subsequent*-processing claims, because the cognitivist gloss on the behavioral—and neuroscientific evidence so far seems superior and explanatorily more powerful. The task at hand is thus for the enactivists to provide an equally strong non-cognitivist interpretation of these data, while at the same time not restricting the scope of their theory to the phenomenological level.

4.4 Affordances and truth- or accuracy-conditions

So the amount of computations brains perform to generate percepts and act on them so far supports a representational view of perception. Enactivists might reply at this point that they do not deny that brain processing is required to perceive what the world affords us and to act on what we perceive. What they do deny is that we must interpret this neural processing of action possibilities in terms of representations. We have argued that perceiving affordances requires prior processing. We now turn to the question whether perceiving affordances involves content. Radical enactivists consider themselves radical because they reject any content on the level of basic cognitive phenomena, including perception and action. In their interpretation of Brentano's passage, Hutto and Myin (2017) defend a liberal notion of intentionality. However, by rejecting Clark's moderate notion of action-oriented representations (Hutto and Myin 2013, pp. 51–56) and the notion of non-conceptual (or non-propositional) representation in their discussion of Crane's approach to perception (Hutto and Myin 2017, pp. 100–101) they defend strong constraints on the notion of representational content, most plausibly construed as propositional content: "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." (Hutto and Myin 2013, p. x) In their view, an explanation of content requires "saying how basic minds came to master the relevant public practices that made it possible to fix the right kinds of standards" (2017, p. 120) and add that "such intersubjective practices and sensitivity to the relevant norms comes with the mastery of the use of public symbol systems". From this they conclude that "as it happens, this appears only to have occurred in full form with construction of sociocultural cognitive niches in the human lineage" (2017, p. 134). Linking content to truth conditions of linguistic utterances is reminiscent of Chisholm's (1957) claim that intentional content is propositional, i.e. that it can be expressed in a sentence using a "mental" verb and a that-clause (see also Quine 1960, p. 201; Davidson 1970, p. 210). But the question that matters for us is whether the constraints on what counts as representational content could not also be more liberal than Hutto and Myin suppose. It may be true that mastery of public, intersubjective practices and sensitivity to norms only comes with the mastery of the use of public symbol systems, but tying the possibility of having contentful states to this ability means placing very high demands indeed.

Against such a view a number of philosophers like Siegel (2014) and earlier Millikan (1995) and also Palmer (1999) provided arguments according to which experiencing Gibsonian affordances in fact presuppose accuracy (if not truth-) conditions of some kind: the mug does not only look as *having property F* (being full or empty say), but it also looks as *to be phi'd* (picked up, grasped or thrown away say etc.). In order for me to be able to react appropriately to such soliciting affordances of the mug my perceptual experience must present the mug to me as having certain properties (cf. Martens and Schlicht 2017). And, importantly, even considering perception of affordances to be direct does not imply that it is immune to error. As Palmer points out:

[A]ffordances can be misperceived whenever nonvisual information carries information that is at odds with the visual information. A wooden bench or log might *look* eminently sittable-upon, but if it is sufficiently rotten, it will not afford sitting to a normal adult. Similarly, a baseball might look as though it affords throwing, but if it has been super-glued to the table, it will not, in fact, be throwable. (Palmer 1999, p. 412)

More generally, if my experience is in some way illusory or too inaccurate, my actions will fail, or to put it in enactive terms, my responses to the world's offerings will miss their mark. Thus, affordances have a normative dimension, too. This, however, amounts to such perceptions having content, even by the lights of Hutto and Myin (2013, p. x).⁷ The relevant content of affordance-perceptions is geared towards possible actions but given the need for accuracy conditions, it is content nonetheless. To capture both their descriptive and prescriptive aspect, Millikan (1995) proposed to conceive of Gibsonian affordances in representational terms. Her notion of a “pushmi-pullyu-representation” exhibits both directions of fit (Anscombe 1957) of informational and motivational states (like beliefs and desires) and is considered the most basic and widespread kind of representation, which we share with many animals. This is echoed in Clark's (1997) and Wheeler's (2008) attempt to integrate the insights from embodied approaches into a more traditional framework by conceiving of representations as being “action-oriented”, “ego-centric”, and “context-sensitive”.⁸

⁷ One way to avoid this conclusion is to accept accuracy conditions that are not yet truth conditions. One obvious move could be to interpret these accuracy conditions in terms of non-conceptual content. Yet, given Hutto and Myin's rejection of non-conceptual representations, it is not easy to see just how these accuracy conditions should be characterized within their framework.

⁸ Bruner (1964) introduced the term “enactive representation” (which today has an ironic ring to it) and distinguished such representations from “iconic” and “linguistic” representations. He conceived of enactive representations as “appropriate skills necessary for sensorimotor acts, for organizing percepts, and for organizing our thoughts” (1964, p. 1) and argued that all our cognitive acts “depend upon techniques rather than upon wired-in arrangements in our nervous system”. By an enactive mode of representing Bruner denotes a kind of representation that cannot be decoupled from an appropriate motor act directed towards some object. Several segments of our environment, like bicycle riding, tying knots, aspects of driving etc., Bruner argues, “get represented in our muscles” (1964, p. 2); these are cases of “representa-

This sketch of ways in which perception of affordances can be integrated into a representationalist framework may suffice here. Indeed, the important points resulting from this discussion are (a) that perception of affordances requires prior information processing, (b) that acting on affordances requires further information processing, and (c) that experiencing affordances involves accuracy conditions. This section was intended not so much as a rebuttal, but as a critical evaluation of the prospects of radical enactivist accounts of perception and the challenges they face. In particular, radical enactivists need a much more specified account of how brain processing contributes to perception, and in doing so avoid falling back into representational accounts. This shortcoming gives representational accounts, which are worked out in much more detail, an explanatory advantage and we thus conclude that the radical enactive account of perception in its current state is not persuasive.⁹ But in addition to our criticisms of enactivist's shortcomings, we now want to close by adding a rather positive evaluation of enactivism with respect to its contribution to the naturalization of intentionality. In this context, we focus on the idea of autopoiesis.

5 Towards a really radical view of cognition

Rather than siding with either representationalism or radical enactivism, we emphasize that it is an empirical question which of the frameworks (or yet another one) will ultimately be adequate to explain perception and basic cognition. Instead we would like to consider a few steps towards a truly radical view of cognition that is radical concerning the scope of cognitive abilities in nature.

As a starting point, we can take Lyon's (2006) proposal to distinguish between an *anthropogenic* and a *biogenic* approach to cognition. The former is widespread in philosophy, taking human cognition as the paradigm test case for the presence of cognitive capacities in nonhuman animals and organisms in nature. The latter, by contrast, starts in a bottom-up fashion to extract organizational features of basic cognitive phenomena: it takes the "principles of biological organization and their links to fitness [...] to be the most productive means by which we can understand what cognition is, what it does, and how and why it evolved" (Barrett 2018). From this perspective, given its demands on content, the radical enactive approach proposed by Hutto and Myin seems to have ties to the familiar group of anthropogenic approaches. By contrast, other enactivist approaches to cognition that focus on autopoiesis as central

Footnote 8 (continued)

tion by action alone" (1964, p. 3). A more recent example from the literature is Milner and Goodale's (1995) patient D.F. whose posting action of a letter can be very well explained in terms of such action-oriented or "enactive" representations.

⁹ A problem we ignore here is that with respect to the metaphysics of perception (Drayson 2018), Hutto and Myin more or less face the choice between naive realism and the sense datum theory once they reject the intentional theory. Although they do not elaborate on this, passages indicating that perception is "world-involving" (Hutto and Myin 2017, pp. 60, 92) suggest they opt for naive realism. This would require them to provide a theory of illusory or non-veridical experiences, which is a notorious difficulty for proponents of naive realism. But such a discussion is beyond the scope of this paper.

for intentionality and cognition are better characterized as biogenic.¹⁰ In this section we argue that a biogenic approach, e.g. in the guise of autopoietic enactivism, can make a persuasive case for a naturalization of intentionality while leaving systematically open the possibility that even the most basic kinds of intentionality are contentful and require an explanation in terms of (some sort of) representation. In closing, we make a suggestion how intentionality, cognition, and representation relate to one another, and formulate some open questions for future research.

5.1 Organisms as self-organizing systems

Philosophers from the autopoietic enactivist tradition have suggested that tracing back the natural origins of intentionality may not lead us *all* the way down to the level of physical particles, but at least to the level of biological self-organization. Organisms are defined by the feature of self-organization or “autopoiesis”, as Maturana and Varela (1980) call it. Central assumptions of this view about organisms yield, according to Thompson’s (2007, p. 159) autopoietic enactivism, “an explicit hypothesis about the natural roots of intentionality: Intentionality arises from the operational closure and interactive dynamics of autopoiesis”, and is thus grounded in the structural organization of living organisms, namely self-organization. According to this, intentionality is a basic feature of embodied acts, not of mental states (ibid., p. 25). This view can be seen as a descendant, indeed a naturalized version, of Kant’s definition of an organism as a “natural purpose”, developed in his *Critique of the power of Judgment* (Kant 1790/1998, §65–66). A natural purpose is a system in which the parts of the system (1) are only possible through their relation to the system as a whole, and moreover, in which (2) the parts of the system are “combined into a whole by being reciprocally the cause and effect of their form”. That is, “in such a product of nature each part is conceived of as if it exists only *through all the others*, thus as if existing *for the sake of the others* and *on account of the whole*, i.e. as an instrument (an organ) [...]” (Kant 1790/1998, p. 373). Unlike an artifact, such as a watch, an organism is—simply in virtue of being a natural purpose—not caused by any external rational agent, such as a watchmaker, but by its own formative powers. An example alluded to by Kant that illustrates self-organization is an organism’s ability to repair itself in response to damage to the body. Kant seems to have been aware of Abraham Trembley’s discovery that after cutting hydra (multicellular organisms found in unpolluted fresh waters) into two halves, they regenerate by developing two complete organisms. As Fitch (2008) observes, many organisms like Salamanders can regrow entire body parts like lost limbs. This astonishing fact raises the question how it is possible for these animals (or their parts anyway) to “know” what they should grow, i.e. what the overall animal is supposed to be like in order to supplement what’s still left with what’s missing. Thompson, elaborating on Kant’s theory, dubs this power “circular causality”, a combination of *local-to-global*- and *global-to-local*-determination: emergent structures and properties on the macro-level are generated and sustained by the behavior of the components on the

¹⁰ See Thompson (2018) for a brief discussion of the differences between his thoroughly autopoietic enactivism and the radical enactivism defended by Hutto and Myin.

micro-level, while at the same time global structures and processes constrain local interactions” (2007, p. 62). He discusses this in the context of neurodynamics:

Coherent and ordered global behaviors, which are described by collective variables or order parameters, constrain or govern the behavior of the individual components, entraining them so that they no longer have the same behavioral alternatives open to them as they would if they were not interdependently woven into the coherent and ordered global pattern. At the same time, the behavior of the components generates and sustains the global order. This two-sided or double determination is known as circular causality [...] (Thompson 2007, p. 62)

Understood in this way, teleology or circular causality is not opposed to causality but introduces a differentiation into the notion of causation in terms of a two-sided dependency. Kant’s notion of a natural purpose delineates the group of those entities which are at the same time products of nature and which necessarily have to be understood teleologically, as being intrinsically directed toward some purpose or goal. But while Kant held teleological descriptions as providing merely indispensable heuristics rather than objective explanations (Kant 1790/1998, p. 389), Varela and Thompson both argue that a modern empirical theory of life (based on the theory of autopoiesis and dynamical systems theory) can be seen as a naturalized version of Kant’s notion of a natural purpose, providing us with a “non-reductionist yet ‘hard’ explanation of the living” (Weber and Varela 2002, p. 102). According to this modern understanding, the organism is conceived of as a “creator of ‘real teleology’: [...] *organisms are subjects having purposes according to values encountered in the making of their living*” (ibid.). In the present context of the elaboration of the basic biological intentionality of organisms, it is crucial that such biosystems are in an important sense autonomous, i.e. “a cell or multicellular organism is not merely self-maintaining, like a candle flame; it is also self-producing and thus produces its own self-maintaining processes, including an active topological boundary that demarcates inside from outside and actively regulates interaction with the environment” (Thompson 2007, p. 64, cf. p. 103 for the defining elements of autopoietic systems). This can be illustrated by looking at the simplest organism, the living cell, out of which all complex organisms are ultimately composed. In a single cell, a biochemical network “produces the metabolites that constitute both the network itself and the membrane that permits the network’s bounded dynamics” (Thompson 2007, p.65). In other words, a cell qua self-organized and self-producing system generates a simple biological self-world distinction. However, in order to survive and maintain its identity, the cell must continually exchange matter and energy with its environment. Some molecules are imported through the membrane and participate in processes inside the cell, whereas other molecules are excreted as waste. In this way, the cell produces its own components including its boundary, which in turn produces and maintains it as a unified system, in an ongoing process, i.e. autopoiesis (Thompson 2007, p. 97ff). In order to sustain itself, the cell must realize biological purposes. With respect to the cell’s biological needs, its formerly neutral physicochemical environment thus obtains a certain *value*, i.e. features of the physicochemical environment “become” nutrition; but only in relation to the cell’s metabolism do they acquire the status of food. As the environment becomes “a place of valence,

of attraction and repulsion, approach or escape” (Thompson 2007, p. 158), living itself is a way of bringing forth value and significance: by being semi-autonomous, the living cell is both *directed* towards and dependent on its environment. This is, according to Thompson, the source of intentionality in the (biological) domain.

Thus, since “the principles constraining biogenic theorizing are empirically solid, rest on firm theoretical foundations, and the key concepts involved enjoy a high degree of interdisciplinary agreement” (Lyon 2006, p. 25), autopoietic enactivism as an exemplification of a biogenic approach is a promising candidate for a naturalized theory of intentionality. Working from this hypothesis about the natural origins of intentionality, various further questions arise from looking at simple biological systems: (1) should we consider all forms of intentionality down to single-celled organisms as cognitive, and (2) if so, are these basic cognitive phenomena best explained in terms of representations?

5.2 Biological intentionality, cognition, and representation

How does intentionality relate to cognition? The traditional understanding of intentionality as both directedness *and* aboutness suggests an immediate link to cognition, which has often been understood as computations over intentional states, understood as states with non-derived content, see Adams and Aizawa 2008, p. 31). Similarly, Fitch (2008) uses *intentionality* and *aboutness* interchangeably, suggesting an equivalence. He regards basic biological intentionality as a causal power of organisms and dubs it—provisionally—

nano-intentionality – a microscopic form of aboutness, inherent in individual eukaryotic cells, that make up a goal-directed capacity to respond adaptively to novel circumstances. The core causal power underlying nano-intentionality is the cell’s ability to arrange and rearrange its own molecules in a locally functional-manner, thus preserving and extending its individual existence... (Fitch 2008, p. 158f)

Thompson (2007, p. 159) also writes that “intentionality first emerges in nature in the form of autopoiesis and sense-making”. To call what the cell is doing here, namely its intentional directedness towards its environment based on its biological needs, “sense-making”, suggests that it already exhibits a basic kind of meaningful cognition. Thompson (2007, p. 127) does not want to take a stand in this matter, but his suggestive way of describing it is criticized by Hutto and Myin (2013, pp. 32–36), and indeed if intentionality is used equivalently to representation, then this seems to imply that representational cognition can be traced back all the way down to single-celled organisms. However, it is not obvious that directedness and aboutness are two sides of the same coin. Taking intentional directedness as a feature of whole embodied agents, rather than states, suggests that there can be intentionality without representation. Moreover, this does not preclude that non-representational intentionality is nevertheless cognitive: as we have seen above, we should not *define* cognition in representational terms.

The debate concerning representations is a tricky one, and its solution is beyond the scope of this paper. Representationalists need to take Ramsey’s *job description*

challenge (Ramsey 2007) seriously to justify positing representations (see Artiga 2016 for an unrestricted representationalism all the way down). Anti-representation-ists, on the other hand, need to spell out in detail how they can account for cognitive processing in non-representational terms without falling back into behaviorism. Until then, merely shifting the burden of proof to the other side does not amount to a justification of either position.

So putting the question of representation aside for the moment, the first question to address is whether single-celled organisms should be considered as cognitive systems in the first place. For some philosophers and scientists, this is clearly the case, since many single-celled organisms control what they do and how they respond to their surroundings. Because of their “eyespots”, Eukaryote cells, which make up animals, are sensitive to light and thus are capable of using visual information for controlling their behavior. *E. coli* are thought to use memory to track favorable nutrient conditions. As long as nutrient conditions improve, they keep swimming in the same direction; if conditions get worse, they tumble randomly until conditions improve once again (Allen 2017). In this context, Godfrey-Smith (2016, p.16) argues that parts of them must be sensing and others must be active in order for the organism to behave like this: while one mechanism “registers” the outside conditions right now, another mechanism “records” the state of affairs moments ago, allowing the bacterium to determine via comparison whether conditions are improving or deteriorating. Some cells are sensitive not only to the presence of other cells but also to the chemicals they produce, since these are “chemicals that are made because they’ll be perceived and responded to by others.” (Godfrey-Smith 2016, p. 19) This can be interpreted as passing the threshold to signaling and communication. Furthermore, some bacteria coordinate their behavior with that of other bacteria: they make their behavior dependent on what others around them do, while at the same time influencing the behavior of those around: the amount of chemicals of some sort they produce depends on how much of it is being produced by others around them; the more they sense, the more they produce. In concert, this so-called *quorum sensing* brings about macro-behaviors like illumination in Hawaiian squid. A point can be made that this marks the beginnings of social behavior (Godfrey-Smith 2016).¹¹ However, as Barrett puts it, the problem with these examples is that

there is very little consensus on what entitles something to be called a cognitive process; often [...] we simply point toward the set of abilities we consider to be cognitive, like reasoning and language, and leave it at that, without specifying exactly why such processes count and others don’t. (Barrett 2018, p. 720)

As a result, it seems that whether we count them as cases of cognition depends on our theory of what cognition is. However, this should not be a question of taste. Ramsey’s argument against a priori representationalism can be expanded to other a priori definitions of what cognition is. The best way to really turn the question concerning the nature of cognition into an empirical one is to tie it to the sciences that investigate cognition. Doing this, we should be open to the possibility that there is no single answer to that question, i.e. that cognition eludes our attempts to define

¹¹ For a defense that also plants have cognitive abilities, see Maher (2017).

it: if cognitive phenomena are those which cognitive scientists investigate, there may be nothing that all and only phenomena considered cognitive have in common (Allen 2017). Furthermore, the question may not be all too important in actual scientific discussions, as Allen argues:

Detailed scientific investigation of the adaptive information processing capacities of a variety of systems [...] is typically conducted at the level of concepts such as learning, memory, problem solving, and decision making rather than at the more abstract level of “cognition”. [...] Although the term “cognition” is used in scientific contexts, in my view it serves as a kind of umbrella term under which more specific capacities are grouped, whereas [...] all serious modeling efforts and investigation of mechanisms pertain to the specific capacities. (Allen 2017, p. 6)

Thus, it seems unproblematic to understand single-celled organisms as cognitive, given the examples of memory, perception, signaling etc. presented above. This translates to other, more complex organisms like plants, which are often excluded from the group of cognitive systems (cf. Dennett 1995). Calvo Garzón and Keijzer (2011) and Calvo and Friston (2017) consider the different timescales that are relevant for the behavior of plants and make a persuasive case for including plants, like bacteria, in the domain of cognitive systems. Moreover, while it has been argued that *genuine cognition* requires *genuine intentionality*, and thus talk of bacterial memory or perception is not to be taken literally, this does not pose a problem once we take intentionality to be a basic feature of life, as proposed by enactivists.

Taking cognition as an umbrella term that refers to whatever cognitive scientists investigate, the relation between intentionality and cognition is hard to determine. Since cognitive science could make basic intentionality its subject, it could turn out that the domains of intentionality and cognition coincide. However, given that autopoietic enactivism takes intentionality to arise from basic metabolic functions of living organisms, this would also mean that cognitive science coincides with biology (see Schlicht 2018).

A promising alternative is Moreno et al.'s (1997) idea that cognition requires *meta-metabolic processes*, i.e. their function being “to identify and optimize the conditions needed to maintain metabolic processes within acceptable limits” (Barrett 2018, p. 725). In the present context, this suggestion is interesting for various reasons. First, it can be characterized as a biogenic approach: it does not start with high-level cognitive abilities, but rather assigns a biological function to cognition whose development we can investigate in a bottom-up fashion from an evolutionary perspective. Secondly, it gives us a useful criterion to demarcate cognitive science from general biology. Thirdly, it is formulated in a way that leaves room for scientific development: bacterial memory matches this description and thus counts as cognitive, although the idea of understanding cognition as meta-metabolic processes was originally intended to defend the peculiar role of nervous systems for cognition. Finding these cognitive processes in organisms without nervous systems (like bacteria) forces us to rethink the role of nervous systems for cognition.

Understanding cognition in terms of meta-metabolic functions allows us to specify the relation of autopoietic intentionality and cognition: Cognition is in the service of intentionality, its function is to allow continuous autopoiesis under non-optimal

conditions. If intentionality arises from autopoiesis, non-cognitive intentionality is concerned with metabolic processes. If cognitive processes arise on top of that, they can enrich and transform basic biological intentionality. To the extent that some organisms have no need to “identify and optimize the conditions needed to maintain metabolic processes within acceptable limits”, and metabolic processes are all we find, then there is intentionality without cognition. And if there is intentionality without cognition, then basic intentionality is best understood as directedness without aboutness (i.e. representation), in line with the enactivist proposal. However, the reverse does not hold: even if it turns out that wherever there is intentionality, there is also cognition, that does not entail that there is also representation. This is so because it is a further question whether basic cognition is representational or not. Importantly, however, even if cognition in simple organisms like bacteria were best characterized as non-representational, this would not easily generalize to more complex organisms and thus would not support the radical enactivist claim that memory, perception etc. are non-representational in all non-human animals. While the view considered here supports the claim that cognition can be found in organisms as simple as bacteria, it is open with respect to representations and gives rise to a set of empirical questions that need to be addressed in future research: Is cognition always (even in bacteria) representational? If not: Is basic cognition (such as perception, memory etc.) always non-representational, as claimed by radical enactivists? What is the role of brains and nervous systems for cognitive processing given that we find cognition in brainless organisms? How can organisms be at the same time autonomous cognitive systems and part of other, more complex systems (like *E. coli*, which have memory, but are constitutive parts of humans and other animals, for example)? And finally: How does this notion of biological cognition relate to the notion of cognition applied in artificial intelligence research?

6 Conclusion

In this paper, we evaluated advantages and problems of enactivist approaches to intentionality and cognition. On the one hand, we found radical enactivist approaches still wanting because they have provided neither a robust conception of the role of brains and nervous systems for cognition nor a clear alternative explanation in terms of affordances and coupling to a representationalist account of perception and basic cognition. On the other hand, we argued that in applying a biogenic approach that takes as a starting point basic biological processes rather than high-level human cognitive abilities to extract important features of intentionality, autopoietic enactivism yields a compelling hypothesis about the biological roots of intentionality without taking a definite stand on the question of representation. Starting from this notion of intentionality by applying the biogenic approach to cognition as well, we defended the claim that cognition evolved to perform the function of supporting the basic metabolic processes that give rise to intentional directedness by identifying optimizing conditions needed for metabolic processing. This functional characterization is specific enough to demarcate the domain of cognitive science from the domain of biology, but open enough to accommodate new empirical developments.

Acknowledgements This article was supported by the project “The structure and development of Understanding actions and reasons”, funded by Deutsche Forschungsgemeinschaft (SCHL 588/3-1) and by the Volkswagen Foundation’s funding for the project “Situating Cognition. Perceiving the world and understanding other minds” (Az. 87 105). We are grateful for this support and want to thank two anonymous reviewers for valuable comments on an earlier version.

References

- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Oxford: Oxford University Press.
- Allen, C. (2017). On (not) defining cognition. *Synthese*. <https://doi.org/10.1007/s11229-017-1454-4>.
- Anderson, M. (2014). *After phenology. Neural reuse and the interactive brain*. Cambridge, MA: MIT Press.
- Anscombe, E. (1957). *Intention*. Cambridge: Harvard University Press.
- Artiga, M. (2016). Liberal representationalism: A deflationist defense. *Dialectica*, 70(3), 407–430.
- Barrett, L. (2018). The evolution of cognition. A 4E perspective. In A. Newen, L. de Bruin, & S. Gallagher (Eds.), *Oxford Handbook of 4E cognition*. Oxford: Oxford University Press.
- Bourjade, M., Thierry, B., Call, J., & Dufour, V. (2012). Are monkey able to plan for future exchange? *Animal Cognition*, 15(5), 783–795.
- Brandom, R. (1994). *Making it explicit, reasoning, representing and discursive commitment*. Cambridge, MA: Harvard University Press.
- Brentano, F. (1874/1995). *Psychology from an Empirical Standpoint* (A. C. Rancurello, D. B. Terrell & L. L. McAlister, Trans.). London: Routledge.
- Bruner, J. (1964). The course of cognitive growth. *American Psychologist*, 19(1), 1–15.
- Calvo, P., & Friston, K. (2017). Predicting green: Really radical (plant) predictive processing. *Journal of the Royal Society Interface*, 14, 2070096.
- Calvo Garzón, P., & Keijzer, F. A. (2011). Plants: Adaptive behavior, root brains, and minimal cognition. *Adaptive behavior*, 19, 155.
- Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Chisholm, R. (1957). *Perceiving: A philosophical study*. Ithaca: Cornell University Press.
- Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78, 67–90.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (2016). *Surfing uncertainty*. Oxford: Oxford University Press.
- Clark, A., & Toribio, J. (1994). Doing without representation. *Synthese*, 101, 401–431.
- Crane, T. (2003). The intentional structure of consciousness. In A. Jokic & Q. Smith (Eds.), *Consciousness: New philosophical perspectives*. Oxford: Oxford University Press.
- Cummins, R. (1997). *Representations, targets, and attitudes*. Cambridge, MA: MIT Press.
- Dally, J. M., Emery, N. J., & Clayton, N. S. (2006). Food-caching western scrub-jays keep track of who was watching when. *Science*, 312, 1662–1665.
- Damasio, A. (2011). *Self comes to mind*. New York: Pantheon.
- Davidson, D. (1970). *Essays on actions and events*. Oxford: Oxford University Press.
- De Jaegher, H., & Di Paolo, E. (2007). Participatory Sense-Making: An enactive approach to social cognition. *Phenomenology and the Cognitive Sciences*, 6(4), 485–507.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1995). *Kinds of minds*. New York: Basic Books.
- Dolega, K., Roelofs, L., & Schlicht, T. (Eds.). (2018). Enactivism, representationalism, and predictive processing. Special Issue of *Philosophical Explorations*, 21(2).
- Drayson, Z. (2018). Direct perception and the predictive mind. *Philosophical Studies*, 175(12), 3145–3164.
- Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, content, and function*. Oxford: Oxford University Press.
- Egan, F. (2014). How to think about mental content. *Philosophical Studies*, 170(1), 115–135.
- Engel, A. (2010). Directive minds: How dynamics shapes cognition. In J. Stewart, O. Gapenne, & E. Di Paolo (Eds.), *Enaction. Towards a new paradigm for cognitive science* (pp. 219–244). Cambridge, MA: MIT Press.

- Fitch, W. T. (2008). Nano-intentionality: A defense of intrinsic intentionality. *Biology and Philosophy*, 23, 157–177.
- Fodor, J. (1975). *The language of thought*. New York: Harvard University Press.
- Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- Fuchs, T. (2018). *Ecology of the brain*. Oxford: Oxford University Press.
- Gallagher, S. (2008). Direct perception in the intersubjective context. *Consciousness and Cognition*, 17, 535–543.
- Gallagher, S. (2017). *Enactive interventions*. Oxford: Oxford University Press.
- Gibson, J. J. (1986). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Godfrey-Smith, P. (2016). *Other minds. The octopus, the sea, and the deep origins of consciousness*. New York: Farrar, Straus and Giroux.
- Hanus, D., Mendes, N., Tennie, C., & Call, J. (2011). Comparing the performances of apes (Gorilla gorilla, Pan troglodytes, Pongo pygmaeus) and human children (Homo sapiens) in the floating peanut task. *PLoS ONE*, 6(6), e19555. <https://doi.org/10.1371/journal.pone.0019555>.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Hutto, D. D. (2008). *Folk-psychological narratives*. Cambridge, MA: MIT Press.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing enactivism*. Cambridge, MA: MIT Press.
- Hutto, D. D., & Myin, E. (2017). *Evolving enactivism*. Cambridge, MA: MIT Press.
- Jacob, P., & Jeannerod, M. (2003). *Ways of seeing. The scope and limits of visual cognition*. New York: Oxford University Press.
- Kant, I. (1790/1998). In Guyer (Ed.) *Critique of the power of judgment*, 2002. Cambridge: Cambridge University Press.
- Kelly, S. D. (2002). Merleau-Ponty on the body: The logic of motor intentionality. *Ratio*, XV(new series), 376–391.
- Kirchhoff, M. (2011). Anti-representationalism: Not a well-founded theory of cognition. *Res Cogitans*, 2, 1–34.
- Kriegel, U. (Ed.). (2013). *Phenomenal intentionality*. New York: Oxford University Press.
- Lyon, P. (2006). The biogenic approach to cognition. *Cognitive Processing*, 7(1), 11–29.
- Maher, C. (2017). *Plant minds. A philosophical defense*. London: Routledge.
- Martens, J., & Schlicht, T. (2017). Individualism versus interactionism about social cognition. *Phenomenology and the cognitive sciences*. <https://doi.org/10.1007/s11097-017-9499-x>.
- Matthews, R. (2007). *The measure of mind*. Oxford: Oxford University Press.
- Maturana, H. R., & Varela, F. G. (1980). *Autopoiesis and cognition: The Realization of the living. Boston studies in the philosophy of science* (Vol. 42). Dordrecht: D. Reidel.
- Mendes, J., Hanus, D., & Call, J. (2007). Raising the level: Orangutans use water as a tool. *Biology Letters*, 3, 453–455.
- Millikan, R. G. (1989). Biosemantics. *Journal of Philosophy*, 86, 281–297.
- Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, 9, 185–200. **AI, connectionism, and philosophical psychology**.
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Montague, M. (2010). Recent work on intentionality. *Analysis*, 70, 765–782.
- Moreno, A., Umerez, J., & Ibañez, J. (1997). Cognition and life: The autonomy of cognition. *Brain and Cognition*, 34, 107–129.
- Neander, K. (2016). *A mark of the mental*. Oxford: Oxford University Press.
- Noë, A. (2009). *Out of our heads*. Newyork: Hill and Wang.
- Noë, A. (2014). *The varieties of presence*. Cambridge, MA: Harvard University Press.
- O'Brien, G., & Opie, J. (2004). Notes toward a structuralist theory of mental representation. In H. Clapin, P. Staines, & P. Slezak (Eds.), *Representation in mind: Approaches to mental representation* (pp. 1–20). Oxford: Elsevier.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–1031.
- Palmer, S. (1999). *Vision. Photons to phenomenology*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Raby, C. R., Alexis, D. M., Dickinson, A., & Clayton, N. S. (2007). Planning for the future in western scrub jays. *Nature*, 445, 919–921.
- Ramsey, W. (2007). *Representation reconsidered*. Cambridge: Cambridge University Press.
- Ramsey, W. (2017). Must cognition be representational? *Synthese*, 194(11), 4197–4214.
- Rupert, R. (2018). Representation and mental representation. *Philosophical Explorations*, 22(2), 204–225.

- Schlicht, T. (2018). Does separating intentionality from representation imply radical enactivism? *Frontiers in Psychology*, 9, 1497. <https://doi.org/10.3389/fpsyg.2018.01497>.
- Shea, N. (2014). Exploitable isomorphism and structural representation. *Proceedings of the Aristotelian Society*, 64(2), 123–144.
- Siegel, S. (2014). Affordances and the content of perception. In B. Brogaard (Ed.), *Does perception have content?*. Oxford: Oxford University Press.
- Sinigaglia, C. (2008). Enactive understanding and motor intentionality. In F. Morganti, A. Carassa, & G. Riva (Eds.), *Enacting intersubjectivity: A cognitive and social perspective on the study of interactions* (pp. 17–32). Amsterdam: IOS Press.
- Slaby, J., Stephan, A., & Walter, H. (Eds.). (2011). *Affektive Intentionalität*. Paderborn: Mentis.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1–74.
- Smortchkova, J., Dolega, K., & Schlicht, T. (Eds.). (2019). *What are mental representations?*. New York: Oxford University Press.
- Sprevak, M. (2013). Fictionalism about neural representations. *The Monist*, 96, 539–560.
- Sterelny, K. (2003). *Thought in a hostile world. The evolution of human cognition*. Hoboken: Blackwell Publishing.
- Thompson, E. (2007). *Mind in life*. Cambridge, MA: Harvard University Press.
- Thompson, E. (2018). Review of Daniel D. Hutto and Erik Myin, *evolving enactivism: Basic minds meet content*. *Notre Dame Philosophical Review*, January 11, 2018.
- Treisman, A. (1988). Features and objects: The fourteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, 40A, 201–236.
- Van Dijk, L., & Rietveld, E. (2017). Foregrounding sociomaterial practice in our understanding of affordances: The skilled intentionality framework. *Frontiers in Psychology*, 7, 1969. <https://doi.org/10.3389/fpsyg.2016.01969>.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Vygotsky, L. V. (1997). *The collected works of L.S. Vygotsky: Problems of the theory and history of psychology* (Vol. 3). Berlin: Springer.
- Weber, A., & Varela, F. (2002). Life after Kant. Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences*, 1, 97–125.
- Wehner, R. (1999). Large-scale navigation. In C. Freksa (Ed.), *Spatial information theory*. Berlin: D. Mark.
- Wheeler, M. (2008). Minimal representing: A response to Gallagher. *International Journal of Philosophical Studies*, 16(3), 371–376.
- Wilson, A., Golonka, S., & Barrett, I. (Eds.). Radical embodied cognitive neuroscience. *Frontiers in Psychology*. <https://www.frontiersin.org/research-topics/1713/radical-embodied-cognitive-neuroscience>. Accessed 6 Sept 2019.
- Ylikowski, P., & Kuorikoski, J. (2008). Dissecting explanatory power. *Philosophical Studies*. <https://doi.org/10.1007/s11098-008-9324-z>.

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

Affiliations

Tobias Schlicht¹  · Tobias Starzak¹

Tobias Starzak
tobias.starzak@rub.de

¹ Institute of Philosophy II, Ruhr-Universität Bochum, Universitätsstr. 150, GA3/29, 44780 Bochum, Germany