Representational Enactivism

Abstract: In the literature on enactive approaches to cognition, representationalism is often seen as a rival theory. In this paper, I argue that enactivism can be fruitfully combined with representationalism by adopting Frances Egan's content pragmatism. This representational enactivism avoids some of the problems faced by anti-representational versions of enactivism. Most significantly, representational enactivism accommodates empirical evidence that neural systems manipulate representations. In addition, representational enactivism provides a valuable insight into how to identify representational content, especially in brainless organisms: we can identify representational content by investigating autopoietic processes.

Word count: 4547

1. Introduction

Representationalists hold that cognition involves internal states that represent things in the external world. Enactive approaches to cognition are usually understood as rejecting representations. Instead, they begin their accounts of cognition with *action*, and concentrate on how an organism dynamically interacts with the external world. For example, autopoietic enactivists contend that the origin of cognition should be found at the level of biological self-organization. That is, where actions such as growth, repair and regeneration occur (Varela and Maturana 1980; Thompson 2007). Radical enactivists explicitly argue that basic forms of cognition, such as perception, are completely non-representational (Hutto and Myin 2013, 2017).

Significant problems emerge, however, by rejecting representations. One problem is that some higher forms of cognition like thought and language clearly do involve

representations. This is something even radical enactivists concede (Hutto and Myin 2013, 2017). In response, radical enactivists mostly focus on basic cognition. So, at its best, radical enactivism seems to be an incomplete theory of cognition, even if it ends up being true. A second problem is that dismissing the explanatory value of representations makes it difficult for enactivism to explain why representational terms are so widely used in cognitive science.

When it comes to representationalism, there are more than two options. First, robust realists believe that both representations and content are real. Second, eliminativists such as, radical enactivists, hold that both representations and content should be eliminated. There are a range of options lying between robust realism and eliminativism (Ramsey 2021). Egan's content pragmatism is one of those options (Egan 2010, 2014, 2020; Coelho Mollo 2020). Content pragmatism acknowledges the utility of representations but does not commit to the ontological reality of representational content. It treats representational content as an explanatory gloss that depends on pragmatic considerations. I argue that both representationalists and enactivists should consider this option when it comes to non-human cognition: the representational content of non-human cognition is not a real property of the cognitive system but can still be attributed to the system based on pragmatic concerns.

The primary aim of this essay is to defend a representational version of enactivism. Here is how I will proceed. In Section 2, I introduce how Hutto and Myin's radical

enactivism contradicts our best neuroscientific evidence. Motivated by that concern, I argue for representational enactivism by bridging enactivism and content pragmatism. In Section 3, I explain how it is possible to understand content pragmatism in ways that make it look compatible with enactivism. In Section 4, I take a step further and explore how the idea of autopoietic enactivism can be incorporated into mainstream representationalist approaches: representational enactivism helps to identify the content of representations by analyzing organisms' exercise of skillful know-how in the environment.

2. Radical Enactivism and Representationalism

This section will introduce radical enactivism and explain why it is at odds with representationalism. Referring to radical enactivism is important since, of all the versions of enactivism, radical enactivism opposes representationalism most strongly.

Radical enactivists do not believe that basic cognitive phenomena are representational, including cognitive phenomena like perception. They do admit that certain sophisticated cognitive phenomena are representational, such as human thoughts involving linguistic or otherwise symbolic content (Hutto and Myin 2013, 2017). Radical enactivism is most distinctive for its commitments regarding the non-representational structure of "basic cognition". This is despite the fact that it is widely assumed in cognitive science that cognition should be explained in terms of computations that

manipulate representations.

For radical enactivists, some cognition is intentional, but not representational. They hold that directedness and aboutness are two aspects of intentionality.[1] Radical enactivists emphasize that aboutness is contentful and may depend on complex linguistic capacities, whereas directedness is manifested in basic, contentless cognition. Our understanding of the former—the contentful and sophisticated form of intentionality—cannot be generalized to basic intentionality (Hutto and Myin 2017). In this way, radical enactivists reject the idea that sophisticated human thought is the most typical or paradigmatic form of intentionality. The two types of intentionality should not be treated the same, and the most primitive form of intentionality lacks content.[2] This basic intentionality without content or representations is sometimes called "Ur-intentionality" (Hutto 2008). By using this terminology, radical enactivists intend to capture how the attitude of the whole organism is directed toward an object, without turning to contentful mental states of the organism.

In this sense, the radical enactivist view is clearly related to Gibson's ecological approach to perceptions. Gibson (1979) puts stress on what the environment *offers* the

[1] This can be traced back to Brentano (1874/1995).

[2] Hutto and Myin (2013, 2017) reject content on the level of basic cognitive phenomena because of the connection they make between the content and the conditions for the truth, or accuracy, of linguistic utterances. An agent should be sensible to certain norms and capable of exhibiting public and intersubjective behavior that conforms to those norms, before having representational content. In their words, it is possible to have contentful states only after the "construction of sociocultural cognitive niches in the human lineage" (Hutto and Myin 2017, 134).

organisms: organisms *directly* perceive what is accessible to them in the environment, including both good things and bad things. In other words, the perceiver *immediately* gets a bunch of information about the *affordances* in the environment; affordances are operationalized in terms of informational availability. Hutto and Myin (2013) use the term "worldly offerings", which has a similar meaning to Gibson's affordances, to demonstrate the enactivist idea: an organism' immediate response to certain worldly offerings does not require its brain to produce representations of the environment. This form of the agent's dynamic response to the stimulus, as it moves around in the environment, is contentless. This is because the minimal kind of intentionality that it requires to explain those responses, *directedness*, is contentless.

In a word, radical enactivists maintain that basic cognitive phenomena such as perception are non-representational. No representation is needed because the perception of affordances is direct.[3] At this point, I need to describe a completely opposite view held by robust representational realists, that is, neuroscience suggests that representations exist.

Thomson and Piccinini (2018) use empirical evidence from neuroscience to argue for the existence of neural representations. The three kinds of representations discussed in that article are sensory representations, representations uncoupled from current sensory

[3] Schlicht and Starzak (2019) argue against Hutto and Myin's view by pointing out that the best neuroscientific evidence regarding perception shows that the perception of affordances is not direct. For example, grasping the possible affordances of a mailbox cannot solely rely on the perception of the mailbox—one should already have some knowledge about the mailing system in order to identify it as such (Palmer 1999).

stimulation, and motor representations. On their view, neuroscientists do not merely posit neural representations, rather, they observe and manipulate them. Neuroscientists do this by establishing that some neural signals fit the criteria of representations (Thomson and Piccinini 2018, 195-96). The criteria are (1) the signals carry information about a current state or a future state of the environment, (2) there is a systematic mapping between the signals and a current or future state of the environment, and (3) the system can use them to guide future behavior, or the signals actually cause the future state of the environment. In addition, Thomson and Piccinini note that representations have two parts: a semantic content and a functional role.

Semantic content can be indicative or imperative (Millikan 1984). Indicative content is about how the world is like, and imperative content is about how the world will be. On the account of structural representationalism, sensory representations have indicative content, and their functions are to track the actual states of the world. Motor representations have imperative content, and their functions are to produce behaviors and bring about new states of the world. Thus, there is a match between representational content and function. Taking sensory representation as an example, according to *informational teleosemantics* (Neander 2017), teleosemantics assigns semantic content to sensory representations based on natural semantic information in the environment. At the same time, the function of the sensory system is to carry the same natural semantic information, and then the neural signals can transmit that information and guide the future

behavior (*cf*. Thomson and Piccinini 2018, 194-95).[4] Representational theories are claimed to be well supported by neuroscientific evidence.

Enactivists may be correct to claim that the *phenomenology* of experience supports the idea that we directly experience the ways that something in the world is useful. However, the details about neural processing leave open the possibility that a representationalist account is correct, phenomenology aside. As Gallagher (2008, 537) comments, problems like how the inferior temporal cortex works are for neuroscientists to solve empirically, not for ordinary perceivers to decide based on the structure of their experience. Some enactivists, such as, Gallagher, seem content to concentrate on phenomenology. All enactivists should accept neuroscientific theories about neural mechanisms.

If enactivism about human perception is mostly confined to phenomenology, there appears to be no genuine conflict between representational and enactive accounts of perception. This is because they are targeting different explananda: neural mechanisms and phenomenology respectively. Because enactivism lacks a positive account of the underlying mechanisms, enactivism and representationalism seem to be concerned about different things. This is echoed in Piccinini (2018, 2): "Ecological psychologists argue that cognition is primarily explained in terms of dynamical variables characterizing the

[4] This relation between function and content, however, is not held in every representational theory (e.g., Egan's content pragmatism, which is assessed in later sections).

interaction between agents and environments. According to them, uncovering inner mechanisms is unnecessary".

Enactivists seem to be at a crossroads. Either enactivists can maintain that some cognition is non-representational, or they should keep the notion of representation because it is very useful in explaining the fruitful results in empirical research. If the latter option is preferred, it looks as though radical enactivism is ruled out.

One way out would be implied by an account of representation that is compatible with an enactivist view of cognition. One option for enactivists is to appeal to content pragmatism, which is developed by Egan (2010, 2014, 2020): accept the *explanatory* value of representation when representational contents seem unavoidable in cognitive science, but deny the ontological reality of representations at the same time.

3. Content Pragmatism

Is it plausible to suppose that enactivists would accept content pragmatism? Radical enactivists argue that it is incorrect to define cognition in representational terms because only sophisticated cognition has representational content. Radical enactivists are also eliminativists who claim that representational content should be eliminated. For radical enactivists, positing representational content with truth or accuracy conditions does not add anything to the explanation of basic forms of cognition like perception and action (Hutto and Myin 2013). Consider now robust realists. Robust realists hold that

representations and representational content are real features of cognitive systems. Eliminativists, on the other hand, believe that someday the best science will get rid of representations and content and instead explain cognition in purely functional or neurophysiological terms (*cf*. Coelho Mollo 2020). Radical enactivists focus on basic forms of cognition, while autopoietic enactivists often base their theory on simple organisms such as biological cells and bacteria (as we will see in Section 4). Then, *prima facie*, it seems that in both cases it would be difficult to persuade enactivists to accept robust realism about representations.

Content pragmatism, however, does not involve this commitment to robust realism. The core of this theory is that representational content is part of an explanatory gloss informed by pragmatic considerations (Egan 2010, 2014, 2020; Coelho Mollo 2020). Egan (2020, 22) emphasizes the distinction between representational *vehicle* and representational *content*. The account of mental representation defended by Egan couples a realist account of representational *vehicles* with a pragmatic account of representational *content*. The realization function specifies the physically realized vehicles of representation, which are structures or states of some sort. The interpretation function specifies the content of representation. Therefore, like the states or structures posited in all well-confirmed scientific explanations, the representational vehicles pinpointed by neuroscientists are real. By contrast, representational content serves heuristic purposes but is not part of what Egan calls the "theory proper" (Egan 2014). She means that

representational content is not the target phenomenon for computational theories to explain, and it should be seen an explanatory gloss dependent on pragmatic considerations.

Egan's observation that human beings' mental representations, such as thoughts and feelings, are different from the representations in mindless cognitive systems is noteworthy. For example, although plants do not obviously have mental representations like humans do, plants are believed to represent temporal properties because they have circadian clocks. Egan points out that a naturalized account of representational content cannot explain what is so special about mental representations and that "from a detached, naturalistic perspective there may not be any *distinctively* mental representation" (ibid., 27).

This idea is coherent with the motivation behind my defense of representational enactivism. In Section 2 I showed how radical enactivists draw a distinction between basic cognition and sophisticated cognition. Disregarding radical enactivists' disputable viewpoints about human perception, there does appear to be a real distinction between representations in mindless systems and human thoughts.

Without claiming that basic cognition is non-representational, it does seem right that explaining some forms of cognition in much simpler organisms requires theories besides the computational theories that function to explain human cognition well. So, enactivists might also see content pragmatism as an attractive representational account because it

allows them to treat basic cognition and sophisticated cognition differently.

It might be a separate question whether Egan is correct to claim that (all) representational content is purely an explanatory gloss, since much neuroscientific evidence suggests that neural representations are real (Thomson & Piccinini 2018).[5] But tackling that issue is not the aim of this essay. In the next section, we will have a closer look at how a form of enactivism can be compatible with representationalism.

4. Enactivism and Representationalism

4.1 Autopoietic Enactivism

In earlier sections, I discussed why radical enactivism is problematic and explained the idea of content pragmatism. Now I will focus on how autopoietic enactivism and representationalism can be complementary. Compared with radical enactivism, autopoietic enactivism is a better candidate for being integrated into a representationalist framework because it does not "radically" preclude the existence of representations. The representational enactivism I defend in what follows here is an integration of autopoietic enactivism and content pragmatism.

The concept of autopoiesis, or self-organization, was introduced by Maturana and Varela. In their usage, a closed system that has autonomy, self-reference and self-

[5] I think that content pragmatism is more likely to be true when it comes to simpler organisms whose representational content is more difficult to determine.

construction is an autopoietic system (1980). Autopoiesis is said to be the nature of living systems, and based on this theory, they define cognition as a biological phenomenon. Cognition is present in all forms of organisms, simple as well as complex. No matter how simple an organism is, as long as it shows the behavior of self-maintenance, it is cognitive by definition.

Thompson (2007) develops autopoietic enactivism based on the autopoietic organization of biological life. On his view, mental life should be seen as a part of bodily life such that it cannot be reduced to processing in the brain. Instead, it should be understood through its role in the world (2007, ix). A cell or a multicellular organism is not merely self-maintaining. We should also see an actively topological boundary demarcating the inside of an organism from its outside and actively regulating the organism's interaction with the environment (2007, 64). Accordingly, Thompson encourages us to find intentionality in organisms' operational closure and dynamic interaction with the environment. In this way, Thompson agrees with Dennett's statement that "intentionality doesn't come from on high; it percolates up from below"[6] (2007, 160).

Proponents of autopoietic enactivism employ the idea of autopoiesis in order to discover the origin of cognition. This so-called "bottom-up" approach insists that finding the principles of biological organization is the most productive way to understand what cognition is, what it does, and how it evolved (Barrett 2018; *cf*. Schlicht and Starzak

[6] Dennett 1995, 205.

2019). This idea opposes that of taking human cognition as a paradigm and generalizing it to simpler forms of cognition. I am not taking a side between these two approaches; there is probably something valuable in both.

4.2 Representational Enactivism

Autopoietic enactivism seems convincing mostly in the case of very basic forms of cognition; perhaps also for the kinds of cognition characteristic of simple organisms like bacteria. However, as Schlicht and Starzak comment (2019, 23), even if we should understand cognition in simple organisms as non-representational, non-representational cognition would not easily generalize to more complex organisms. The existence of the most primitive form of cognition does not, therefore, support the radical enactivists' claim that perception in all non-human animals is non-representational.

A relevant worry is whether enactivism can make any contribution to fruitful representationalist theories. In the opinion of content pragmatists, the identification of the representational content is always dependent on pragmatic considerations. But if the pragmatic considerations are to explain how an organism actively interacts with the environment for its self-organization, autopoietic enactivism can help to identify the content of representations.

I will use a bacterium's cognitive ability as an example. Thompson (2007, 103) lists three criteria for characterizing life in terms of autopoiesis: (1) Semipermeable Boundary:

the system is defined by a semipermeable boundary made up of molecular components. This boundary can be used to discriminate between the inside and outside of the system. (2) Reaction Network: the components are being produced by a network of reactions that take place within the boundary. (3) Interdependency: (1) and (2) are interdependent. Bacteria are autopoietic because they satisfy all three criteria. Therefore, according to Varela and Maturana (1980)'s definition of cognition, cognition is present in bacteria, even though they are relatively simple lifeforms.

There is a bacterium having a cognitive ability to sense the concentration of sucrose in the environment and to move accordingly. We can interpret that cognitive ability as an instance of fulfilling self-maintenance done by a closed and automatic system, which satisfies the definition of autopoiesis. I have noted in Section 2 that Egan does not think there are any naturalistic conditions for content (of representations in mindless systems), otherwise there is nothing special about human's mental representations (2020, 27). According to content pragmatism, the job of connecting the naturalistic theory with the target phenomenon is left for a gloss. Thus, when it comes to bacteria, we can find a gloss "acquiring nutrition for self-maintenance", which is inspired by biological evidence, so that we can make sense of that bacterium's cognitive process described above. In this case, "acquiring sucrose because it is nutritional for self-maintenance" seems to be an appropriate gloss for the representational content.

I have so far emphasized the ways that representational enactivism claims the

advantages of both autopoietic enactivism and content pragmatism. But a representationalist need not believe that representational content is always dependent on pragmatic considerations to benefit from what autopoietic enactivists say about the origin of cognition. A robust realist can have very good neuroscientific reasons to believe that human cognition should be defined by representational mental states. But that realist can also embrace the idea of autopoietic enactivism when other forms of cognition are of concern, and whenever neuroscientific evidence is not available.

Enactivists can also benefit from this understanding of representational content, but need to commit to its relevance across the board. Notice that although Thompson (2007) focuses on embodied and enactive acts, he does not give up representations altogether. Instead, in several places of his book, he seems to challenge some interpretations of representations and promotes an understanding of representation that fits well with autopoietic enactivism. Thompson's understanding of representations can be found in his comparison between code and DNA (2007, 182): a representational system of code is composed of arbitrary referential relations between the symbols and what they stand for. However, DNA is not representational in this way because it is a component of the autopoietic process and the relation between DNA and the information it contains is not arbitrary. It is unacceptable to say that DNA "contains the information for phenotypic design" because this piece of information is contained in the autopoietic network as a whole instead of any component.

The representational version of enactivism I defend has the same feature: the representational content of any part of the organism should not be considered independently of other parts of the whole autopoietic network. The autopoietic process determines the content of representations by specifying the characteristics of each part of the organism. The representational content, therefore, should be understood by reference to how a sub-function contributes to a global autopoietic process. This idea can be useful when understanding the representational content of a non-human animal or plant (or of its part): the content should be identified based on the autopoietic network because the autopoietic system in its entirety specifies the function and the semantics.

One difficulty of integrating enactivism and computationalism[7] pointed out by Casper and Artese (2020) is that different heuristics are implicit in the two views. While enactivists explain cognitive processes in integrative terms, computationalists often attempt to localize cognitive processes, and specify how each local part of a functional mechanism processes input and produces output. For this reason, they claim that radical embodied views, including enactivism, are incompatible with decompositional and localizational strategies. It is not clear, however, why enactivists must reject localization.

Non-radical enactivists like Thompson, for instance, would have no problem appealing to representations found in any part of the organism, such as representations

[7] Piccinini (2008) separates representationalism from computationalism. The form of computationalism that is difficult to be integrated with enactivism, according to Casper and Artese (2020), is representational computationalism.

within DNA. But they may be more interested in finding out how a certain sub-function contribute to the global function. And according to pragmatic concerns, the *content* of local representations can be determined by the global autopoietic process.

I am not saying that the enactive approach should be taken on every form of cognition. However, for the study of simpler organisms where neurological evidence is not available—e.g., if techniques like fMRI are not suitable—a thorough investigation into the autopoietic system might helpfully inform efforts to identify the content of the relevant representations. The representational enactivism defended in this essay can in these ways supplement mainstream representationalist theories.

5. Conclusion

Finally, I want to emphasize that the enactive point of view need not undermine representational approaches. For example, it seems that autopoietic enactivism can be even integrated into an informational teleosemantic framework. Advocates of informational teleosemantics say that what it takes to possess and process representations is to process information for a control function (Neander 2017; cf. Piccinini 2020). In an enactive context, the content of a teleosemantic representation should match the specific function played by the parts of the autopoietic system. Consider an organ of the autopoietic system. It processes representations to identify and optimize the conditions needed to maintain itself. The representational content is determined by that specific

function.

What I described is representational enactivism based on informational teleosemantics. For radical enactivists, a view like this would be much more difficult to accept. Informational teleosemantics is a naturalistic theory, and it claims that each token of representation is caused by a certain function (Neander 2017). The representational content is fixed once the function is confirmed. Representational enactivism based on content pragmatism, on the other hand, would deny that there is any naturalistic basis for content.

Since radical enactivists are eliminativists about representations, they are not willing to admit that representational content can be settled by some biological functions. Nor that any representational content is real. But they might be willing to say that the use of representational content is done only out of pragmatic concerns. But non-radical enactivists, like Thompson, could accept the idea that representational content is fixed by certain function as long as that the function is an autopoietic function.

This essay has argued that an integration between enactivism and representationalism is possible and deserves serious consideration. Embracing representationalism would allow enactivists to give a more convincing, neuroscientifically informed account of cognition, and embracing enactivism can help representationalists identify the representational contents of basic and non-human forms of cognition. For these reasons, representational enactivism may provide a fuller picture of cognition.

References

Barrett, Louise. 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.

Casper, Mark-Oliver, and Giuseppe Flavio Artese. 2020. “Maintaining coherence in the situated cognition debate: what computationalism cannot offer to a future post-cognitivist science.” *Adaptive Behavior*, OnlineFirst, 1-15.

Coelho Mollo, Dimitri. 2020. “Content pragmatism defended.” *Topoi* 39 (1):103-13.

Dennett, Daniel Clement. 1995. *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.

Egan, Frances. 2010. “Computational models: a modest role for content.” *Studies in History and Philosophy of Science Part A* 41 (3):253-59.

Egan, Frances. 2014. “How to think about mental content.” *Philosophical Studies*, 170(1), 115–35.

Egan, Frances. 2020. “A deflationary account of mental representation.” In Joulia Smortchkova, Krzysztof Dolega & Tobias Schlicht (eds.), *Mental Representations*. New York, USA: Oxford University Press.

Gallagher, Shaun. 2008. “Direct perception in the intersubjective context.” *Consciousness and Cognition*, 17, 535–43.

Gibson, James. J. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Hutto, Daniel. D. 2008. *Folk-psychological narratives*. Cambridge, MA: MIT Press.

Hutto, Daniel. D., and Erik Myin. 2013. *Radicalizing Enactivism: Basic Minds Without Content*. Cambridge, MA: MIT Press.

Hutto, Daniel. D., and Erik Myin. 2017. *Evolving Enactivism: Basic Minds Meet Content*. Cambridge, MA, USA: MIT Press.

Millikan, Ruth. G. 1984. *Language, thought, and other biological categories: New foundations for realism*. Cambridge, MA: MIT Press.

Neander, Karen. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge, MA: MIT Press.

Palmer, Stephen. 1999. *Vision Science: Photons to Phenomenology*. MIT Press.

Piccinini, Gualtiero. 2008. "Computation without representation." *Philosophical Studies*, 137, 205–41.

Piccinini, Gualtiero. 2018. "Computation and representation in cognitive neuroscience." *Minds and Machines* 28 (1):1-6.

Piccinini, Gualtiero. 2020. *Neurocognitive Mechanisms: Explaining Biological Cognition.* Oxford University Press.

Ramsey, William. 2021. "Defending representation realism." *What are Mental Representations?* 54-78. Oxford University Press.

Schlicht, Tobias, and Tobias Starzak. 2019. "Prospects of enactivist approaches to intentionality and cognition." *Synthese*: 198 (Suppl 1):89-113.

Thompson, Evan. 2007. *Mind in life.* Cambridge, MA: Harvard University Press.

Thomson, Eric, and Gualtiero Piccinini. 2018. "Neural representations observed." *Minds and Machines* 28 (1):191-235.

Varela, Francisco J., and Humberto R. Maturana. 1980 [1973]. *Autopoiesis and Cognition: the Realization of the Living*. D. Reidel Publishing Company.