4-15-2021

Representational Enactivism

Zhexi Zhang
University of Missouri-St. Louis, zhexi.zhang2018@gmail.com

Representational Enactivism

Zhexi Zhang

Bachelor of Management, Accounting, Lanzhou University, 2017
M.A. Philosophy, King's College London, 2018

A Thesis Submitted to The Graduate School at the University of Missouri-St. Louis
in partial fulfillment of the requirements for the degree
Master of Arts in Philosophy

May
2021

Advisory Committee

Committee Chairperson, Gualtiero Piccinini, Ph.D.
Chairperson

Committee Faculty Member 2, William Dunaway, Ph.D.

Committee Faculty Member 3, Lauren Olin, Ph.D.

Representational Enactivism

Zhexi Zhang

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.

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.

In this essay, I will argue that there is a version of representationalism enactivists can accept: content pragmatism (Egan 2010, 2014, forthcoming; 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. Egan admits the indispensability of representations and content in fruitful cognitive science, but she rejects the idea that representational content is real in cognitive systems. For her, content is dependent on pragmatic concerns about how the specific content of a representation can offer a helpful explanation. It appears to me, then, that enactivists can accept content pragmatism, and appeal to representations for their explanatory value when representational terms are used in the best scientific theories. But they do not need to accept robust realism about representational content.

When it comes to representationalism, there are at least three 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. Egan's content pragmatism is a third option. Unlike eliminativists, content pragmatists hold that representations are ineliminable parts of cognitive systems. Unlike robust realists, they argue that the notion of content does

not capture anything real in cognitive systems. I will argue that both representationalists and enactivists should consider this third 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 examine how Hutto and Myin's radical enactivism is problematic: it seems to contradict 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.

The representational enactivism defended in this essay has two main benefits. First, it allows enactivists to draw on successful representational explanations in the areas where enactive theories have almost nothing to say (e.g., about the neural mechanisms underlying human propositional thought and language). Second, when no neural evidence is available in brainless organisms like bacteria, representationalists can appeal to the dynamical resources emphasized by enactivists to explain their patterns of cognition and behavior. This is possible as long as enactivism does not

imply eliminativism about representations.

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. I intend to show that even radical enactivism can be made compatible with representationalism.

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.[1]

For radical enactivists, basic cognition (e.g., that of simple organisms) is intentional, but not representational. They hold that there is a form of intentionality which cannot be explained by representationalist approaches. In representationalist frameworks, intentionality, which is the ability to be about something, can be

[1] Views on the nature of representations differ. For example, representations are symbols for Fodor (1975), while they are subsymbolic activation patterns for connectionists. But many theorists may agree with Carey (2009, p. 5) that representations are "states of the nervous system that have content, that refer to concrete or abstract (or even fictional) entities, properties and events" (*cf.* Schlicht and Starzak 2019).

explained in terms of representation. For example, when I see a glass of water, some perceptual state of mine is representing it; when I am thinking about the usefulness of the glass of water, some computational states are representing how I will use it. Radical enactivists just deny that such relations between intentionality and representation always hold.

Radical enactivists try to "disentangle" *directedness* and *aboutness* (Hutto and Myin 2017) in order to make this point. Directedness and aboutness are important as they are two aspects of intentionality which can be traced back to Brentano (1874/1995). "*Directedness*" means the direction toward an object in external perception or the direction toward humans' own mental phenomena in inner perception (Brentano 1874/1995, p. 22). "*Aboutness*" means the intentional reference to a content in some propositional attitude.

To argue for the existence of intentionality without representational content, 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. 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 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. However, especially in light of Schlicht and Starzak (2019)'s discussion on the dilemma proponents of radical enactivism face, there are reasons to think that their focus on so-called contentless cognition is problematic.

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. The notion of content as they interpret it can be encapsulated in the following: there is content if and only if there are specified satisfaction conditions for truth or accuracy (Hutto and Myin

2013). The idea is that in order to have representational content, an agent should be sensible to norms and capable of exhibiting public and intersubjective behavior that conforms to those norms. If not, the agent and other agents around her would have no clue about whether the content of the agent's mental state is true or accurate. It is possible to have contentful states only after the "construction of sociocultural cognitive niches in the human lineage" (Hutto and Myin 2017, p. 134). Therefore, according to this line of thought, it seems that only humans can have contentful cognition and mental representations. By contrast, basic and contentless forms of cognition are prevalent in animals. It is wrong, therefore, to conclude that all forms of cognition are representational just because human thought is representational. This is one of the most important arguments put forward by radical enactivists.

Schlicht and Starzak (2019) argue against Hutto and Myin's view by questioning whether perceiving Gibsonian affordances can happen without any conditions on truth or accuracy. They point out that many theorists do believe that experiencing affordances presuppose accuracy conditions of some kind.[2] For example, before I react appropriately to some affordances of the mug, my perceptual experience must present the mug as having certain properties, such as being in a certain shape and size and being able to be picked up or thrown away (Martens and Schlicht 2018). Affordances can also be misperceived, when nonvisual information is at odds with the visual information. A bench may appear to someone sittable enough for an adult to sit on, but in fact it is severely rotten (Palmer 1999, p. 412). When the experience

[2] The philosophers they refer to include Siegel (2014) and Palmer (1999).

becomes inaccurate, the organism may fail to respond to the affordances in an ideal way. The concern about truth or accuracy motivates some philosophers to advance ways in which the perception of affordances can be integrated into a representationalist framework (e.g., Bruner 1964, *cf*. Schlicht and Starzak 2019). There is content whenever there are satisfaction conditions for truth or accuracy, but radical enactivism rejects content completely on the level of "basic phenomena" including perceptions. The difficulty with specifying satisfaction conditions without using any representationalist framework constitutes part of the reason why I want to bridge enactivism and representionalism.

As I mentioned earlier, radical enactivists maintain that basic cognitive phenomena such as perception are non-representational. No representation is needed because the perception of affordances is direct. Schlicht and Starzak (2019) break this direct perception of affordances into two parts: the no-preceding-processing claim and the no-subsequent-processing claim. For the sake of time, we will only see how the former is refuted: neuroscientific evidence demonstrates that we cannot perceive affordances directly.

The no-preceding-processing claim is that perceiving affordances is direct: it does not require any preceding processing before the generation of responses to what the environment affords. But, as Schlicht and Starzak (2019) argue, the best neuroscientific evidence regarding perception shows that the perception of affordances is not direct. There are prior stages before perceiving affordances, like for instance the categorization of objects. Perceiving objects' affordances is the final

stage. For instance, 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). Perceiving affordances then involves some prior processing related to features like shape and color. Representationalists call the internal states about shape and color sensory representations.

At this point, I need to say more about why 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 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, pp. 195-196). 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.

Sensory representations give an organism a fallible but reliable access to the environment and can be used to guide behavior. Thomson and Piccinini (2018) argue that neuroscientists have discovered sensory representations at multiple levels of organization in the nervous system. There are low-level sensory representations in the

retina. Higher-level visual regions also contain representations of visual motion; for example, area MT (Movshon and Newsome 1996) represents visual system.

Besides sensory presentations, philosophers are also interested in representations that are uncoupled from current sensory stimulation. Uncoupled representations are activation patterns in the nervous system that carry information about the past state of the environment. The past state is a part of the broader mapping between internal and external states that guides action (Thomson and Piccinini 2018, p. 205). An example is birdsong learning. The two learning stages are *sensory learning stage* and *sensorimotor learning stage* (Mooney 2009). In the former stage, young birds listen to a tutor sing a song and acquire a memory of that song. In the latter stage, young birds reproduce the song sang by the tutor. The song was stored in long-term memory so that young birds can receive some error signals if there are differences between the memory and the song they produce.

Motor representations are commands sent from the brain to the body to move. Neuroscientists use motor maps to show which regions control which kinds of movement. [3] Researchers have done studies on how the primary motor cortex (M1) act together to control movement in the primate brain. There are two theories about how M1 represents movement (Fetz 1992; *cf*. Thomson and Piccinini 2018, pp. 214-215). Neuronal populations in primate animals either implicitly or explicitly represents some feature by extracting information about that feature. The representation is implicit when using it requires a great number of computational

[3] An example is motor homunculus, which shows where muscles twitch when a person is electrically stimulated (Penfield and Boldrey 1937; *cf*. Thomson and Piccinini 2018, p.211).

steps, while the representation is explicit when not many computational steps are needed. No matter which specific representational theory of M1 function is taken, goal-directed commands are representational.

Sensory representations and motor representations are not always isolated from one another. Piccinini holds that there are mixed representations (Piccinini 2020; Thomson and Piccinini 2018), which are constructed in processes of sensorimotor transformations. This happens when sensory representations are converted into motor representations. Neural populations then produce sensory representations and motor representations simultaneously.

The above are examples of different forms of representations that are supported by neuroscience. Piccinini (2018, 2020) concludes that the complex control function in neural systems necessarily involve *structural representations*. Structural representations are surrogates for what they represent. The structure of system A is a representation of structure B if A is homomorphic with B. Importantly, the informational *content* of representational system A is relevant to the control *functions* of system B (Morgan and Piccinini 2018). [4]

A representation has two parts: a semantic content and a functional role (Thomson and Piccinini 2018). 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

[4] Structural representations must have four elements: "(i) a homomorphism (partial isomorphism) between a system of internal states and their target, (ii) a causal connection from the target to the internal states, (iii) the possibility for
the internal states to be decoupled from their target, and (iv) a role in action control" (Piccinini 2018, p.3).

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, pp. 194-195). This relation between function and content, however, is not held in every representational theory (e.g., Egan's content pragmatism, which will be assessed in later sections). Since we have seen the importance of representations in neuroscience, I will stop here for now and return to enactivism.

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, p. 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.

Following Gallagher, it seems open for enactivists to limit their understanding of direct perception to the phenomenological level, and to concede that the phenomenology of direct perception is compatible with a representational account of the machinery of perception, which in turn is supposed to provide neuroscientific details for explaining how specific neural processing contributes to the appearance that perception is direct. The neural details are not part of the phenomenology of seeing how something in the environment affords certain actions.

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): “Ecological psychologists argue that cognition is primarily explained in terms of dynamical variables characterizing the interaction between agents and environments. According to them, uncovering inner mechanisms is unnecessary” (p. 2).

Enactivists seem to be at a crossroads. On one hand they are silent about what the specific neural mechanisms behind perception are. The neuroscientifc evidence about perception seems to suggest that the processing of representations is real. But some enactivists, on the other hand, claim that there is simply no representational content involved in basic forms of cognition. Either enactivists can maintain that basic

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. But then it looks as though radical enactivism is ruled out.

One way out would be implied by an account of representation that is still compatible with an enactivist view of cognition. One option for enactivists is to appeal to content pragmatism, which is developed by Egan (2010, 2014, forthcoming): 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 this? 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, (at least) *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, forthcoming; Coelho Mollo 2020). Egan (forthcoming) 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* (p. 22). 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.

To see how content pragmatism can be defended, here is a question about how a frog's internal state represents a bug. Is it a fly, frog food, or a small dark moving thing (Egan forthcoming, pp. 10-11)? Egan argues that this question cannot be

answered if the specific explanatory concern is unknown. The representational content is more likely to be a fly if the goal of the theory is to explain the frog's behavior within its environmental niche. On the other hand, the content being assigned is more likely to be a small dark moving thing if the explanatory concern is about the frog's visual mechanisms. The specific goals or the pragmatic considerations are directly contrary to naturalistic considerations about representational content. For example, tracking theorists would claim that there is a causal relation between internal states and external objects.[5]

Egan does not believe that there is a naturalistic content-determining relation that will rule out the need to appeal to pragmatic considerations. However, it is noteworthy that Egan does not intend to compromise the naturalistic credentials of neuroscience (ibid., p. 14). This is because the "theory proper" is about representational vehicles rather than about representational content. For Egan, only representational content is independent of naturalistic considerations. There are still naturalistic constraints on representational vehicles.

For organisms which are much simpler than human beings, I take content pragmatism to be the best representationalist theory. Egan's observation that human beings' mental representations, such as thoughts and feelings, are different from the representations in mindless cognitive systems is noteworthy in this sense. For example, although plants do not obviously have mental representations like humans

[5] For example, Tye (1995) sees intentionality as a causal-informational relation between the internal state and the distal entity, so that the representational content of the internal state is determined by the function of tracking the distal entity.

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., p. 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).[6] 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.

[6] 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.

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-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.

Based on my earlier discussion, it might be wrong for us to conclude that representation is unnecessary for cognition. Everyone in this debate agrees that simple organisms do not have nervous systems capable of supporting the formation of representations like human beings. Radical enactivists are happy with this observation. However, we should choose an alternative route offered by content pragmatists: cognition in simple organisms is still representational, but the

representational content is dependent on specific explanatory purposes. Without neuroscientific evidence, it seems impossible for us to identify representational content in a reliable and purely naturalistic way.

Thompson (2007) develops his 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 (p. 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 (p. 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"[7] (p. 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 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.

[7] Dennett 1995, p. 205.

4.2 Anti-realism about Representational Content

Interestingly, content pragmatism and enactivism have a similar attitude towards representational content. I will examine the reason why representationalism based on structural similarities is rejected by some enactivists. As we have briefly seen in Section 2, structural similarity, or isomorphism, is a relation between a model or a map and what it represents. Structural similarity holds when an internal state and an external condition are structurally similar. Defenders of structural representationalism claim that structural similarity ground the representational content. As O'Brien (2016) says, representational vehicles are "contentful in virtue of resembling their represented objects" (p. 9). But Egan (forthcoming, p. 7) argues that structural similarity is not sufficient to underwrite determinate contents. Enactivists raise a similar objection.

The first reason offered by Segundo-Ortin and Hutto (2019) is that if the representational content of a model or map is completely determined by what it structurally mirrors, then what it represents is indeterminate (pp. 7-8). They give an example: suppose that there is a map which is said to contentfully represent Sydney. The metrical relations among the constituent elements of the map mirror those of Sydney. But they imagine that the same map also mirrors the spatial layout of New York City (or another city with a similar layout), though to a different degree. Then it seems that the map represents both Sydney and New York City, and so the representational content is indeterminate.

Segundo-Ortin and Hutto (2019) have taken two responses into consideration. One is that the representational content of the map is fixed not only by what it is structurally similar with, but also by what the map is "in fact used to deal with" (Godfrey-Smith 2006, p. 58). Another response is that representational content is also fixed by the casual relations that bring the representational vehicle into existence, in other words, by what they are selected for in the cognitive process (Ramsey 2016; *cf*. Segundo-Ortin and Hutto 2019). However, some enactivists respond that whether the fact that one item structurally mirrors another item does not *suffice for* one to contentfully represent the other. Hutto and Segundo-Ortin contends that structural similarities are not *inherently* contentful (2019, p. 8). For example, we can use variations in the height of liquid column of a mercury thermometer to accurately make an inference about the changes of temperature, but it is a separate question whether the structural similarities (here the variations in the height) in themselves contentfully represent things (here the changes in the temperature).

Following this reasoning, we can make some truth-relevant inference about a fact in the environment, by relying on a device with a similar structure. This practice can be very successful, but it does not entail that the structural similarity suffices for the representation to have a definite content. At least, the content seems not always be inherently fixed by the similarity. We can presuppose that structural similarities are contentful, but to explain *how* and *why* they are contentful is a difficult task. This problem is labeled as "Hard Problem of Content" by radical enactivists (Hutto and Myin 2013). Since I do not know how to solve this problem, I would for now contend

that, at least in some cases, people endow representations with certain contents out of their pragmatic concerns—not because contents are always fixed by certain structural similarities.

Segundo-Ortin and Hutto (2019) argues that "if structural similarities do not suffice for or entail content, then *a fortiori* they do not get their cognitive work done in virtue of possessing content" (p. 9). Content pragmatists can happily accept this contention, because for them, there is no natural condition to constrain representational content. What gets the cognitive work done are representational vehicles. However, radical enactivists seem to go too far when they claim that cognition does not always involve representations. If enactivists stop at the claim that cognition does not always involve representational content, then enactivism can be compatible with content pragmatism. It appears to be acceptable to say that structural similarity is insufficient for representational content (instead of representations). Being anti-realists about representational content but realists about representations seems to be an option which radical enactivists have ignored.

4.3 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, p. 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. For example, if you are interested in interpreting a frog's behavior in terms of its ability to self-organize, it seems more appropriate to assign the representation in the frog's brain the content of nutritious food, not a small dark moving thing.

It seems that integrating content pragmatism and autopoietic enactivism benefits both views: enactivists can take a pragmatic stance towards representational content without committing to robust realism in cases where representational terms feature in our best scientific explanations. At the same time, enactivism can allow that pragmatic considerations on the biological level sometimes help to identify the appropriate representational contents.

I will use a bacterium's cognitive ability as an example. Thompson (2007, p. 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 already noted in Section 2 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 (forthcoming, p. 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 she 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 (p. 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[8] 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 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

[8] 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.

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, 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

Brandom, R. B. (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford: Oxford 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.

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Casper, M. & Artese, G. F. (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, D. (2020). Content Pragmatism Defended. *Topoi* 39 (1):103-113.

Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.

Egan, F. (2010). Computational models: a modest role for content. *Studies in History and Philosophy of Science Part A* 41 (3):253-259.

Egan, F. (2014). How to think about mental content. *Philosophical Studies*, 170(1), 115–135.

Egan, F. (forthcoming). A Deflationary Account of Mental Representation. In Joulia Smortchkova, Krzysztof Dolega & Tobias Schlicht (eds.), *Mental Representations*. New York, USA: Oxford University Press.

Fetz, E. E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences*, 15, 679–690.

Fodor, J. (1975). *The language of thought*. New York: Harvard University Press.

Gallagher, S. (2008). Direct perception in the intersubjective context. *Consciousness and Cognition*, 17, 535–543.

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

Godfrey-Smith, P. (2006). Mental representation, naturalism, and teleosemantics. In D. Papineau & G. MacDonald (Eds.), *Teleosemantics: New philosophical essays* (pp. 42–68). Oxford: Oxford University Press.

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

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

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

Martens, J., & Schlicht, T. (2018). Individualism versus interactionism about social understanding. *Phenomenology and the Cognitive Sciences* 17 (2):245-266.

Menary, R. (2016). Pragmatism and the pragmatic turn in cognitive science. In Karl Friston, Andreas Andreas & Danika Kragic (eds.), *Pragmatism and the Pragmatic Turn in Cognitive Science*. Cambridge MA: M.I.T. Press. pp. 219-236.

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

Mooney, R. (2009). Neural mechanisms for learned birdsong. *Learning & Memory*, 16(11), 655–669.

Morgan, A. & Piccinini, G. (2018). Towards a Cognitive Neuroscience of Intentionality. *Minds and Machines* 28 (1):119-139.

Movshon, J. A. & Newsome, W. T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. *Journal of Neuroscience*, 16(23), 7733–7741.

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

O'Brien, G. (2016). How does mind matter? Solving the content causation problem. In Thomas Metzinger (ed.), *Open MIND Philosophy and the Mind Sciences in the 21st Century. Volume 2,*. Cambridge, USA: MIT Press. pp. 1137-1150.

Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. MIT Press.

Piccinini, G. (2008). Computation without representation. *Philosophical Studies*, 137, 205–241.

Piccinini, G. (2018). Computation and Representation in Cognitive Neuroscience. *Minds and Machines* 28 (1):1-6.

Piccinini, G. (2020). *Neurocognitive Mechanisms: Explaining Biological Cognition.* Oxford University Press.

Ramsey, W. M. (2016). Untangling two questions about mental representation. *New Ideas in Psychology*, 40, 3–12.

Schlicht, T. (2018). Does Separating Intentionality From Mental Representation Imply Radical Enactivism? *Frontiers in Psychology* 9.

Schlicht, T. & Starzak, T. (2019). Prospects of enactivist approaches to intentionality and cognition. *Synthese*:1-25.

Segundo-Ortin, M. & Hutto, D. D. (2019). Similarity-based cognition: radical enactivism meets cognitive neuroscience. *Synthese* 198 (Suppl 1):1-19.

Siegel, S. (2014). Affordances and the Contents of Perception. In Berit Brogaard (ed.), *Does Perception Have Content?* Oxford University Press. pp. 39-76.

Thompson, E. (2007). *Mind in life.* Cambridge, MA: Harvard University Press.

Thomson, E. & Piccinini, G. (2018). Neural Representations Observed. *Minds and Machines* 28 (1):191-235.

Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.

Varela, F. & Maturana, H. (1980 [1973]). *Autopoiesis and Cognition: the Realization of the Living*. D. Reidel Publishing Company.

Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.