

THE HARD PROBLEM OF CONTENT AND WHY WE SHOULD SOLVE IT

SASKIA JANINA NEUMANN¹

Abstract: The world is its own best model (Brooks, 1989)! Or at least it is for Brooks' behavior-based robots. Robots such as Herbert, whose task consisted in stealing empty soda cans from offices and Baxter, whose task was to interact with human beings in a factory environment, prove that complex behavior can be achieved without any inner representation of the world. If we believe Hutto and Myin (2013), we human beings are not too different from Herbert or Baxter. Our behavior can be explained without assuming that we mentally represent our world. In this paper, I want to show evidence against this claim based on empirical data. Although certain aspects of our behavior can be understood without assuming internal representation of the world, most of our actions and how we store interactions with the world strongly suggest internal representation. This claim will be based on the following empirical evidence: non-linguistically based spreading activation (Barr et al., 2014), object recognition (Banich & Compton, 2018), agnosia (Banich & Compton, 2018) and visual reconstruction (UC Berkeley, 2011). I will, furthermore, extend this theory and argue for human beings being different in knowledge acquisition to behavior-based robots due to the nature of our memory. While behavior-based robots do not need to represent the world, we human beings have to represent the world in a specific manner to be able to interact with the world. The knowledge we acquire is mainly based on our inner representation of the world.

¹ Saskia Janina Neumann is currently pursuing a Doctor of Philosophy (PhD) focused on the Epistemology of Memory at Eötvös Loránd University, Budapest.

Keywords: AI, Epistemology, Epistemology of Memory, Technology, Philosophy of Technology, Hard problem of content, Radical Enactivism, Enactivism

1. Introduction

When I remember Aunt Abby's birthday, I remember quite a lot of what I have experienced as something. For instance, I remember her birthday as her (57th) birthday or the cake we got as red velvet cake. My ability to represent my experiences as something is indicative of our memory traces not only being a form of reactivation (and nothing else on top of that) but as saving content from the experiences we had and transmitting this content up until the point, we recall our experience. While this is a plausible explanation of how we can remember experiences as something, the claim that memory traces have content has its opponents. David Hutto (2023) claims that memory traces do not need content, that reactivation is simple reactivation and nothing more and together with Erik Myin (2013) that cognition in general can be explained without assuming that content exists in the world or our cognition. In what follows, I will counter their claim, asserting that there is a compelling reason to assume cognition relies on content—specifically, our incapacity to engage with the world in a sophisticated and continuous manner without content.

After introducing Hutto and Myin's theory, namely, radical enactivism, I will provide arguments for the assumption that our cognition relies on content. However, I will also show that while our cognition needs content, there is cognition which is possible without content. I will introduce Rodney Brooks' behavior-based robots to give an example of cognition without content. While I will admit that their interaction with the world is quite sophisticated² and that cognition without content is possible, I will also explain that this form of cognition

²I will not define the word 'sophisticated' in this paper. While there may be important distinctions between sophisticated and non-sophisticated behavior, I will assume that our common understanding of sophisticated behavior is captured by the behavior I will introduce in the form of Brook's behavior-based robots.

is limited. I will suggest that these limits can only be surpassed by contentful cognition, and that radical enactivism is wrong for this reason.

2. Radical enactivism

Imagine you use your phone to tell your heating system at home to heat your home up to 20 degrees before you even arrive at home, to open the curtains and to start the coffee machine. With current smart home technology, this is already a possibility. You can control your home remotely. But what if you would not need your phone to do so? What if you were able to control your home with your brain directly? How much of that connection would still count as your cognition or your mind, and how much is still part of the external world? An enactivist would gladly admit that there is not any clear distinction between 'inner' and 'outer' here. They would claim that cognition can only be understood if we take our interactions with the environment into consideration. If we were not to interact with our environment, our body would have nothing to interact with and with this, our mind would not have anything to engage with. However, as we constantly interact with the environment, the environment significantly influences our mind and the non-neural body as well. Our mind and body, however, also constantly influence our environment. Due to the fact that our interactions with the environment are continuous and highly recurrent, we cannot distinguish between a linear input to the mind and an output to the world. Our interaction with the environment is loopy, not linear, and therefore, there is no clear distinction between 'inner' and 'outer' when it comes to cognition (Hutto & Myin, 2013, 6ff.).

Usually, we would assume that something is still transported in this loop, and this something is usually thought to be information. For instance, when we see an apple, we assume that the information of there being an apple in front of us is carried to our brain, processed, stored there and maybe retrieved again once we want to eat an apple and grasp it. An enactivist would agree with this claim. Our cognition carries content.

According to Hutto and Myin (2013, 79), this also implies the following: Cognition consists in representing our interactions with our world as something. For instance, when we see an apple, we represent this apple as an apple in front of our mind. This is achieved by our representational states of representing an apple as an apple having semantic content. For instance, when we see an apple and it is already brown, we represent the apple as rotten. This representation has the semantic content of “The apple is rotten”. Through the semantic content, the content we have has truth-values. Either it is true that the apple is rotten, or it is false that the apple is rotten.

While Hutto and Myin identify themselves as enactivists, they explain that enactivism needs to be radicalized when it comes to content. We need to dispose of the current idea that there is informational content in the world which is acquired by us. Informational content is not the raw material of mental consumption and content also does not enter the game at a later point (Hutto & Myin, 2013, 73) as there is no reason to assume that there is content, truth-values and semantic content present in most of our interactions with the world. We can explain the behavior of beings without assuming that there is content (Hutto & Myin, 2013, 2). Beings are able to deal with aspects of the environment in quite remarkable and sophisticated ways even if they do not have the capacity for content-involving deliberation (Hutto and Myin, 2013, 14). Consider how crickets deal with their environment without any representational content. Female crickets can find their male counterpart by attending the acoustic signals of male songs since the sound they follow has a species-specific characteristic tone and rhythm—one that uniquely matches the particular makeup of the female’s auditory system. By locating their male partner, female crickets ‘[...] engage in a continuous interactive process of engagement with the environment. [...] a beautiful cooperation arises because of the way the cricket’s body and features of the wider environment enable successful navigational activity—activity that involves nothing more than a series of dynamic and regular embodied interactions (Hutto & Myin, 2013, 42). Female crickets do not have representational internal knowledge about their environment. All they do

is to be guided by continuous, temporary, extended interactions with their environment (Hutto & Myin, 2013, 42).

Hutto and Myin (2014, 42) explain that crickets can act intelligently without creating and relying on internal representations and models and that their cognition might even be the only real possibility of basic cognition *per se*. Our cognition does not need to represent or be semantically contentful for us to be able to interact with the world (Hutto & Myin 2013, 82). There is only one possible scenario in which Hutto and Myin grant interaction with the world content. Language is without any doubt semantically contentful (Hutto & Myin 2013, 82). Yet, language is to be seen as the tip of the cognitive iceberg (Hutto & Myin, 2013, 46). Everything 'below' can make do without content. World-directed, action-guiding cognition exhibits intentional directedness that is not contentful (Hutto & Myin, 2003, 82).

Apart from being able to explain cognition without content, Hutto and Myin also introduce another reason to not argue for contentful cognition in most cases. It is the so-called hard problem of content. The problem of telling where content comes from if it cannot be picked up from the world. This problem is subject to what Hutto and Myin (2013, 66) call the Muggle constraint. 'One's explanation of some phenomenon meets the Muggle constraint when it appeals only to entities, states and processes that are wholly nonmagical in character (Hutto & Myin, 2013, 66)'. Or, in other words, there needs to be a clear explanation for the question of where content comes from apart from the world just magically conveying it. To solve the hard problem, we would need to discover how basic physical properties such as a brown apple relate to contentful properties such as 'This apple is rotten' (Hutto & Myin, 2013, 69).

3. Why we should face the hard problem of content

In what follows, I will argue against radical enactivism. I agree with Hutto and Myin when it comes to the claim that there is no content to be picked up from our environment. Our environment does not represent itself to

us as something. For instance, if I am about to be hit by a car, the car itself does not present itself as a (dangerous) car to us. There is no content to be picked up here. However, I disagree with Hutto and Myin's claim that we can interact with the world in quite sophisticated ways without content. Even though we might be able to interact with the world in quite sophisticated ways, I will argue for the view that we do not do so without content. This means, that I have to solve the hard problem of content. But with my suggestion, it is an easy solve. The source content comes from is by the processing powers of our brain itself. I will support this claim and the claim that we might be able to interact with the world in quite sophisticated ways but do not do so based on empirical evidence. The argument will be the same for all four cases. Our brain sorts the input it gets into specific categories. If it sorts the input it gets into specific categories, it needs to represent the input it gets as something. Yet, if it represents the input as something, we have internal mental (probably unconscious) representation and with this content. We do not have cognition without content. Moreover, I will also shortly explain why content needs to be more than linguistic content if Hutto and Myin's argument should be able to bite in the first place.

I'll begin by presenting the argument against content solely relying on linguistic aspects initially. If content was only based linguistically, only beings with the ability to structure their representations semantically would be able to have cognition to start with. This restricts cognition to a very limited field. Therefore, I claim that we necessarily have to extend content to more than semantic content. If we do not do so, Hutto and Myin have it very easy to argue that there is nearly no (except for language) cognitively important content as there is no possibility for content to have a broader base than language. Thus, in what follows, I will assume that content extends beyond the limits of linguistic content. I will show four different cases of empirical evidence which are suggestive of our brain working with content and to assume that cognition cannot make do without it.

The first case is the so-called spreading activation theory by Loftus and Collins (1975). Based on a theory by Quillian, they propose that we store verbal input related to semantic distance/relatedness. For instance, we would store the word 'fox' physically/neuronally closer to the word 'cat' than the word 'book' as both belong to the category 'animal'. To be able to store words depending on their semantic relatedness, our brain needs to interpret those words as what these words mean. It needs to store the acoustic content it gets according to what this content means. Otherwise, there would not be any storing according to semantic relatedness. While this shows that there is the need for content for our brain when it comes to storing, this example only relates to the very tip of the cognitive iceberg so far according to Huttenlocher as it refers to semantic content. Yet, spreading activation goes beyond semantic content. While research around spreading activation is usually done with verbal material (Barr et al., 2014, 550), spreading activation can also be found in other areas than verbal content, for instance, in visuospatial memory networks. Foster et al. (2017), for example, have done studies which suggest that verbal and visuospatial memory networks are dissociated but still show spreading activation in both cases. However, an even more convincing case for spreading activation being present apart from semantic input can be found in preverbal infants. If spreading activation can also be found in preverbal infants, it is not only connected to language and content goes beyond language for the simple reason that preverbal infants do not have language yet. In what follows, I will cite empirical evidence which is indicative of preverbal infants showing spreading activation.

One way to show spreading activation with nonverbal features is to pair a task which would usually be inaccessible after a short period of time with another task which is remembered for a longer period³. After a timely delay which is usually longer than the short-lived task would be accessible, infants are tested on whether the longer-lived task helps to

³ This can be seen as a nonverbal analogue of the verbal priming task.

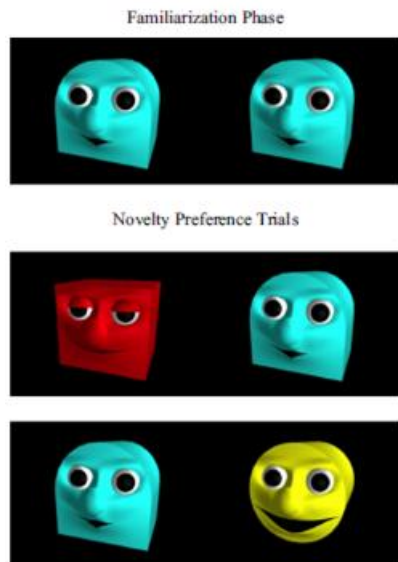


Figure 1

retrieve the shorter-lived task which would usually have been forgotten at that point in time (Barr et al., 2014, 550). Studies like these have been done multiple times (e.g., Barr et al., 2001, 2002; Barr et al., 2011; Greco et al., 1990; Hayne et al., 1993; Timmons, 1994). To understand the individual steps, I will introduce a study by Barr et al. (2014) more closely. Barr et al. (2014) have given infants of different ages (6, 12 and 18 months) a non-verbal analogue of a spreading activation task and have found that spreading activation can also be found with preverbal infants.

They have based their study on two different tasks. The first one is the shorter-lived task and makes use of the visual recognition memory paradigm. The VRM paradigm exploits the fact that young children pay more attention to newer stimuli than to older ones (Barr et al., 2014, 551). During an initial familiarization phase, infants were presented with the same stimuli on two different screens and with two different stimuli during the test phase (stimuli are depicted in figure 1). The already known stimulus from the familiarization phase also changed screens within the test phase and during the second trial of the test phase, a new stimulus was used to exclude the possibility of the infant having become familiar with the novel stimulus from the first trial. Infants tended to show a so-called novelty preference. They preferred to look at the novel stimulus. 10 seconds of familiarization was sufficient for all age groups to show a novelty preference if they were presented with the novel stimulus immediately. 18 months old infants also showed a novelty preference when 24 hours had passed but did not show a novelty preference after

one week anymore. 12- and 6-months old infants did not show a novelty preference if they were not shown the novel stimulus immediately.

The second task consisted of a pairing of the VRM task and a deferred imitation task and was designed to see whether the deferred imitation task would help to prolong retention of the first task. In the deferred imitation task, infants were shown a hand puppet which wore a removable felt mitten on its right hand. During the demonstration session, the experimenter removed the mitten, shook the puppet three times, and placed the mitten on the puppet's hand again. Infants were not allowed to touch the puppet. During the test session, infants were allowed to touch the puppet and monitored to see whether they would imitate the interaction with the puppet of the experimenter. To see whether the deferred interaction task would protract the retention of the VRM task, the demonstration phase of the deferred imitation task was directly followed by the VRM task. After 24 hours, so longer than just an immediate confrontation with the novel stimulus (6- and 12-months old infants), or one week (18 months old infants), as 18 months infants had still shown a novelty preference after 24 hours, infants were tested on whether they would still show a novelty stimulus. All three age groups have shown retention in the DI task. After some changes in the experiment design, the novelty preference could also be found in all age groups. This is indicative of the pairing of the deferred stimulus task with the VRM task prolonging preverbal infant's retention interval.

The given study suggests that similar to spreading-activation in semantically related content, there is spreading-activation in the case of preverbal infants when it comes to contextual or situation dependent information as the ability of preverbal infants to retrieve short lived input increased when it was paired with long lived input. This is indicative of preverbal infants storing this input according to its contextual or situational relatedness. If infants indeed store input situationally or contextually related, they, similar to semantic relatedness, need to store the given input as something. Yet, if they store the given input as something, we have content. Thus, we have the first good reason to assume that there is content in memory traces.

The second case, which is indicative of memory traces having content and our brain not being able to make do without content, is the case of object recognition. If we would not mentally represent an object as that object, there would be no way for us to recognize that object as that object at a later point in time since we would not have any means to compare the former object with the latter object without an inner mental representation. However, we recognize objects. People can recognize objects as the same even if an object has changed locations, configuration or orientation (Banich & Compton, 2018, 539 ff.). For example, if we observe a rabbit at time t_1 eating dandelions on our lawn under a tree and lying on its back, and then at a later time t_2 , we see the same rabbit inside our house on our table, standing on its hind legs, we can recognize that the rabbit in our house is the same one we observed on our lawn at t_1 . Moreover, if we could not recognize objects as the same at a later point in time, we would not be able to change our behavior towards it. But this is what we constantly do. For instance, when we see a rabbit on our lawn at t_1 and see it again at a later point in time at t_2 , our reaction to it should be the same at t_1 and at t_2 as our interaction with the world should trigger the same response if there is no internal representation of that rabbit available. Yet, we react differently to that rabbit at t_2 . For instance, we might feed the rabbit a dandelion at t_2 because we saw it eating dandelion at t_1 . If we were to just blindly react to our world, this would not be possible. We are only able to change our behavior once we represent the rabbit at t_2 as the same rabbit we have met at t_1 .

Because we can change our behavior towards the same object and since we are able to recognize an object as the same object even when it changes configuration, location or orientation, I claim that we are able to represent an object as the same object. Yet, if we are able to represent an object as the same object, we represent objects as something. If we represent an object as something, we have content. Thus, the second reason to assume that our cognition needs content is object recognition.

The third case, which is indicative of memory traces having content and our brain not being able to make do without content, is the case of agnosia. Technically, it is an extension of the second argument. While I

have explained that we are able to recognize an object as the same object and thus have to assume that there is content present, this argument does not only apply to sameness as content but to us usually linking meaning to the input we get and with this giving our input content.

People with agnosia have some physical impairment in their brain. This impairment makes it hard or even impossible for them to give meaning to specific input they get (Banich & Compton, 2018, 500 ff.). For instance, while someone with agnosia might see an apple, they might not be able to tell that they are seeing an apple. For them, the visual information of seeing an apple is meaningless. It does not come with any content. People with a non-impaired brain, in contrast, do not have this problem. They see an apple as an apple. Therefore, I suggest that giving the input we get content is our brain's default reply to input. On top of that, there is further evidence which is indicative of us representing input as something and cognition having content. It is the fact that agnosia comes in different forms. To name only a few, there is prosopagnosia, the inability to recognize faces as such, the inability to recognize human voices as such and to distinguish between them, the inability to recognize printed words as such, the inability to identify a sound as such (for instance, recognizing the honking of a car as honking of a car) and category specific agnosia. Category specific agnosia only applies to specific categories. Someone with category specific agnosia might, for example, not have any problems with recognizing animals as such but might have severe problems with recognizing fruit as such (Banich & Compton, 2018, 500 ff.). If people with agnosia have a deficit in recognizing input to such a specific extent, the input they got must have been stored in the lesioned area. To store input in a specific area in the brain, however, there must have been some sorting of the input into a specific category such as 'noise' or 'fruit'. If we sort the input we get into specific categories, we need to represent the input we have as something. Otherwise, we would not be able to sort. If we see the input we get as something, we need content again. This is why the third reason why memory traces should be seen as having content is the case of agnosia.

The fourth case, which is indicative of memory traces having content and our brain not being able to make do without content, is the case of vision reconstruction. Technically, it is an extension of the third case. In the third case, I have considered what lesioned brain areas can tell us about the question of whether cognition needs content. In the fourth case, I will consider what a non-lesioned brain tells us about content.

Vision reconstruction was first done by UC Berkely (2011) researchers. They have successfully generated videos of what people have seen on a screen by carefully monitoring brain activation while being shown a video and reconstructing this activation into a video based on the very brain activation being present while watching the video. For further reference, feel free to refer to UC Berkeley (2011).

Specific different areas of the brain were, for example, active when it came to specific local features such as edges, color, motion, and texture. This means that we store the input we get to such a reliable extent that others can reconstruct what we have seen by reading our individual brain patterns. If this is possible, our brain sorts the input we get into specific categories to store it at specific brain areas. Yet, if it does so, it needs to represent the input we get as something (for instance as that color or as that edgy form) and with this, we have content again.

I have provided four examples which support the claim that our brain usually represents the input it gets as something and that it thus works with content. Spreading activation, agnosia, object recognition, and visual reconstruction demonstrate the brain's typical reliance on content, although it may demonstrate some functionality without it; its operation is predominantly content-oriented. However, as it works with content, I claim that radical enactivism is wrong in the claim that cognition does not need content. It may not need content in a strictly logical necessary way, but we usually work with content generated by our brain sorting the input we get. This also helps us to solve the hard problem of content. Content does not come from our environment but from our brain.

How human beings are different

In the last section, I have given a possible solution to the hard problem of content. Moreover, I have claimed that even though cognition might make do without content, this is not how we work. While we might not need content in a logical necessary way, we still work with content. Yet, beings such as crickets seem to be able to interact with the world in a sophisticated way without having content. In what follows, I will introduce cases of mechanisms being able to interact with the world in quite sophisticated ways even though they do not work with content. Subsequently, I will show how they differ from us and give a possible explanation for why they are able to interact with the world without content while our cognition continues to need content. While we can adapt our behavior, to learn from the interaction with our environment and to avoid making mistakes in the future, contentless organisms cannot do this. This is why they can make do without content. As long as we want organisms which are able to adapt to their goals though, we need content. Therefore, I claim that radical enactivism does not apply to human beings.

Apart from the already introduced example of crickets, there are further examples of organisms interacting with the world in quite sophisticated ways without content involving cognition. These organisms are Brooks' (1990) behavior-based robots. They are based on a bottom-up approach. Behavior-based robots rely on simple stimulus-response mechanisms without one central executive. They interact with their environment on the basis of these simple mechanisms and in quite sophisticated ways that is (Rodney Brooks and Bottom-Up Robots, 2011). These robots seem to show goal-directed behavior even though their interaction with the world stems from non-goal directed behaviors.




Figure 2



Figure 3

For instance, robot Herbert (depicted in Figure 2⁴) runs around in offices and collects empty soda cans. For this, it scans its environment for hinderances. Once it senses a soda can, it extends its arm, equipped with simple sensors and grasps the soda can (Brooks, 1990). For an impressive showcase of its abilities, feel free to refer to Mechanimal (2010).

Herbert is able to avoid obstacles, follow a wall, recognize a soda can, move its arm into the right orientation to grasp a soda can, search for the soda can, locate it and pick it up (Brooks, 1999).

Brooks' Packbot (depicted in figure 3⁵) is even able to interact with its environment in a more impressive way. It is used as a surveillance device in war areas in which it would be too unsafe to send a soldier into. Packbot can explore its environment autonomously, flip itself over. If necessary, climb stairs and go into communication range in an uncontrolled environment. That is, the Packbot does not need someone to clean up before it enters an area. It can interact with its environment and climb possible hinderances itself even if the terrain should be highly

⁴ Figure 2: Robot Herbert, from: Copeland, 2014.

⁵ Figure 3: Robot Packbot, Defense update, 2007.




Figure 4




Figura 5

inaccessible (TED, 2009). A video demonstration of its behavior can be found at TED (2009). The robot Kismet (depicted in figure 4⁶) can interact with a human being via emotional, facial expressions and words. Its facial features include eyebrows, eyelids, and a mouth. It can mimic expressions analogous to anger, fatigue, fear, disgust, excitement, happiness, interest, sadness, and surprise (Brooks et al., 1999). Kismet's interaction with a human being is surprisingly life-like. When Kismet is told that he has done something wrong, he reacts convincingly upset. Kismet has also been found to have a 25-minute-long conversation with a researcher in the lab. For a video demonstration of this conversation and its emotional responses, feel free to refer to TED (2009).

The two-armed collaborative robot with an animated face Baxter (depicted in figure 5⁷) shows even more surprising behavior. Baxter is able to sense human beings and to interact with them via its animated face. As it senses human beings, it is safe to work around Baxter without getting into accidents with it as Baxter will stop what it is doing once it senses that it might collide with a human being. Baxter's animated face makes it

⁶ Figure 4: Robot Kismet. From: Geo, 2023.

⁷ Figure 5: Robot Baxter and Rodney Brooks Create, 2018.

easy to predict its behavior as it gazes into the direction it will move to. Baxter's two arms enable it to sense its environment and to move objects from one place to another. It can also be trained to make specific movements with its hands by people manipulating its arms (CNBC International TV, 2015). For a peak into how Baxter interacts with its environment, feel free to access CNBC International TV (2014).

Herbert, Packbot, Kismet, Baxter and many more examples show that even though their cognition seems to be contentless as their behavior is only controlled by sensors and certain behaviors as reactions to the world pre-programmed within them, their behavior can be interpreted as quite sophisticated. Yet, I have claimed that we human beings might be able to interact with the world in quite sophisticated ways but simply do not do so. This leaves us with the question of where the difference between us and them might be found.

Herbert might be able to interact with its environment in quite sophisticated ways, but only once following one specific goal, the collection of soda cans in its case remains a feasible goal. A possible world in which it is asked to collect glass bottles will stop it from interacting with the world in a quite sophisticated manner. This even applies to Baxter who is able to acquire new hand movements. If its task was to now find soda cans, it could not do it. The only way for them to be able to achieve these new goals would be re-programming. Fortunately, human beings can do that for them. However, human beings cannot be that easily reprogrammed. If their goal changes, they need to adapt to their changed environment and change behaviors without external help. Human beings are able to do so. Human beings are able to learn from the interaction with their world, change their goals and to also change their behavior towards their environment in a more sophisticated manner than behavior-based robots. The difference between them and behavior-based robots is that they use content to interact with the world. Thus, I suggest that radical enactivism only applies to organisms which cannot adapt their goals. If organisms are able to adapt their goals, their input needs content. Radical enactivism does not apply to human beings.

Conclusions

In this paper, I have presented arguments against radical enactivism, asserting the necessity for extending content beyond linguistic boundaries to avoid an overly restricted definition of cognition. Additionally, I have demonstrated that while humans, akin to other beings, may have the capacity to interact with their world without content, in practice, they seldom do so. Supporting this assertion, I have referenced spreading activation, object recognition, agnosia, and visual reconstruction. Furthermore, I have provided examples of beings capable of sophisticated interaction without content, highlighting their dependency on unchanging goals. In contrast, human beings can maintain sophisticated interactions even when their goals change, suggesting a connection to the use of content in cognition. This implies that radical enactivism is applicable only to beings with unchanging goals, a scenario seldom observed in our dynamically evolving world.

References

- Banich M. T. & Compton, R. J. (2018). *Cognitive Neuroscience*, Cambridge University Press.
- Barr, R., Vieira, A., & Rovee-Collier, C. (2001). Mediated imitation in 6-month-olds: Remembering by association, *Journal of Experimental Child Psychology*, 79, 229–252.
- Barr, R., Vieira, A., & Rovee-Collier, C. (2002). Bidirectional priming in infants, *Memory & Cognition*, 30, 246–255.
- Barr, R., Rovee-Collier, C., & Learmonth, A. E. (2011). Potentiation in young infants: The origin of the prior knowledge effect?, *Memory and Cognition*, 39, 625–636.
- Barr, R., Walker, J. Gross, J. & Hayne, H. (2014). Age-related changes in spreading activation during infancy, *Child Development* 85(2), 549-563.

- Brooks, R.A. (1990). Elephants don't play chess, *Robotics and Autonomous Systems* 6, 3-15.
- Brooks, R.A., Breazeal, C., Marjanovic, M., Scassellati, B., Williamson, M. (1999). The Cog Project: Building a Humanoid Robot, *Computation for Metaphors, Analogy, and Agents*. C. Nehaniv (ed), *Lecture Notes in Artificial Intelligence* 1562, Springer, 52-87.
- CNBC International TV (2015). *Baxter, the bionic robot*, <https://www.youtube.com/watch?v=JWBqXLHlqjE>, last access: 15.07.2023.
- Collins, A.M. & Loftus, E. F. (1975). A Spreading-Activation Theory of Semantic Processing, *Psychological Review* 82(6), 407-428.
- Copeland, B.J. (2014). Nouvelle artificial Intelligence, *Britannica*, <https://www.britannica.com/technology/nouvelle-artificial-intelligence>, last access: 15.07.2023.
- Create (2018). *Rethink Robotics revisited: A look back at the collaborative robots pioneer*, <https://createdigital.org.au/rethink-robotics-collaborative-robots-pioneer/>, last accessed: 15.07.2023.
- 'Deedlydeedee' (2011). *Rodney Brooks and Bottom-Up Robots*, https://www.youtube.com/watch?v=9u0CIQ8P_qk (last accessed 14.07.2023).
- Defense update (2007). *PackBot Tactical Robot*, https://defense-update.com/20070120_packbot.html, last access: 15.07.2023, Figure 3.
- Foster, P. S., Wakefield, C., Pryjmak, S., Roosa, K. M., Branch, K. K., Drago, V., Harrison, D. W. & Ruff, R. (2017), Spreading activation in nonverbal memory networks, *Brain Informatics* 4, 187-199.
- Geo (2023). *What are Social Robots?* https://faculty.cc.gatech.edu/~athomaz/classes/CS8803-HRISpr08/Geo/pv_social_robots.html, last access: 15.07.2023. (Figure 2)
- Greco, C., Hayne, H., & Rovee-Collier, C. (1990). Roles of function, reminding, and variability in categorization by 3-month-old infants, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 617-633.

- Hayne, H., Greco-Vigorito, C., & Rovee-Collier, C. (1993). Forming contextual categories in infancy, *Cognitive Development*, 8, 63–82.
- Hutto, D. D. (2023). Remembering without a trace? Moving beyond trace minimalism, in: *Current controversies in Philosophy of Memory*, ed: Sant'Anna, A., McCarroll, C. J., Michaelian, K., 59-82.
- Hutto, D. D. & Myin, E. (2013). Radicalizing enactivism, *Basic minds without content*, MIT Press.
- Hutto, D. D. (2005). Knowing what? Radical vs conservative enactivism, *Phenomenology and the Cognitive Sciences* 4, 389–405.
- Meade, M. L., Watson, J. M., Balota D. A. & Roediger, H.L. (2007). The roles of spreading activation and retrieval mode in producing false recognition in the DRM paradigm, *Journal of Memory and Language* 56, 305-320.
- Roediger, H., Balota, D. A., Watson, J. M. (2001). Spreading activation and the arousal of false memories, *The nature of remembering: Essays in honor of Robert G. Crowder*, (ed) Roediger, H.L., Nairne, J.S., Surprenant, N. A. M., American Psychological Press, 95-115.
- TED (2009). Rodney Brooks: *How robots will invade our lives*, <https://www.youtube.com/watch?v=UdyRmdv-KiY> (00:05:01-00:07:23), last access: 15.07.2023.
- Temple C.M. (1995). The kangaroo's a fox, *Broken Memories Case Studies in Memory Impairment*, (ed) Campbell R. & Conway M.A., 383-397.
- Timmons, C. R. (1994). Associative links between discrete memories in early infancy, *Infant Behavior and Development*, 17, 431–455.
- UC Berkeley (2011). *Vision reconstruction*, <https://www.youtube.com/watch?v=6FsH7RK1S2E>, last access: 14.05.2023.