

ORIGINAL ARTICLE

Why Imagining Requires Content: A Reply to a Reply to an Objection to Radical Enactive Cognition

Luke Roelofs

Institut für Philosophie II, Ruhr-Universität Bochum

'Radical enactivism' (Hutto and Myin 2013, 2017) eschews representational content for all 'basic' mental activities. Critics have argued that this view cannot make sense of the workings of the imagination. In their recent book (2017), Hutto and Myin respond to these critics, arguing that some imaginings can be understood without attributing them any representational content. Their response relies on the claim that a system can exploit a structural isomorphism between two things without either of those things being a semantically evaluable representation of the other. I argue that even if this claim is granted, there remains a problem for radically enactive accounts of imagining, namely that the active establishing and maintenance of a structural isomorphism seems to require representational content even if the exploitation of such an isomorphism, when established, does not.

DOI:10.1002/tht3.393

1 The aims of REC

Radically enactive cognition, or 'REC', claims that all the activities of "basic" minds go on without representational content, and should instead be understood by reference to features of the environment and organism, the way they interact, and the way those interactions develop over time (Hutto and Myin 2013, Hutto and Myin 2017). A basic mind is roughly one that does not use a full human language, or more broadly is not embedded in human sociocultural practices of truth-telling. These minds (including all nonhuman animal minds) are, according to REC, cognitive but not contentful: they may be very good at dealing with their environment in a sensitive and goal-directed way, but nothing they do is strictly capable of truth or falsity, accuracy or inaccuracy. In a situation where we might intuitively attribute false beliefs to such a mind (e.g. when a dog who has seen some food placed somewhere continues to seek it there even though the food has, "unbeknownst to it", been moved somewhere else), REC claims that we should in strictness only say that its activities are pragmatically mis-aligned, insensitive to certain features of the environment in a way that causes their efforts to fail. They can go *wrong* but cannot *get it wrong*.

Moreover, REC maintains that the cognitive activities characteristic of basic minds remain largely contentless even in mature humans who have mastered the sociocultural

Correspondence to: E-mail: luke.roelofs@utoronto.ca

practices which bestow content. That is, our minds, though not “basic”, do lots of basic things, and do them in largely the same contentless way that animals do. Those activities are taken to include perception, action, emotion, some forms of memory, and, most importantly, “basic imagining”. It is regarding this last item that some critics have specifically objected. Even if representation turns out to be unnecessary for processes like perception and motor control (what we might call “online” processes, dealing with the actually present environment), how on earth, these critics ask, can we make sense of “offline” imagining, dealing with objects that may not be present, and may not actually exist, without talking of representation?

The next section examines this objection in more detail, and the section after that looks at Hutto and Myin’s reply. In this section I will first comment on the two key terms of this debate: what do ‘basic imagining’ and ‘representational content’ mean?

The closest we get to a definition of the sort of imagining under discussion is: “REC denies that the most basic and primitive varieties of imagination have representational content.” (Hutto and Myin 2017, p. 183). Beyond this, the discussion has tended to revolve around a few examples, notably:

- Mental rotation tasks, as studied by Shepard and Metzler (1971): subjects must judge whether two shapes are congruent, and seem to do this by visualising one shape rotating until it matches, or clearly does not match, the other.
- Efforts to solve the Tower of Hanoi puzzle (mentioned by Shapiro 2014): a set of discs must be moved from one pole to another subject to certain rules, and subjects must consider in advance how the game set will look after each prospective move.
- The construction of stone tools by early hominins (as discussed in Malafouris 2013): this requires finding the right stones for breaking other stones, and it is presumed that imagining the needed stones and the ways they might be employed was important here.¹

In each of these cases, the subject imagines tangible objects moving in relation to each other, in order to inform and guide a manual task involving them. These are basic² at least in that there is no obvious reason to think that language or linguistic thought is involved: rather, they seem to be essentially a matter of reusing perceptual and motor abilities, taking them, as is sometimes said, off-line. It seems plausible that nonhuman basic minds could and do employ imagining in this sense—indeed, Hutto and Myin take it that the early hominins discussed by Malafouris had basic minds and nevertheless were able to employ basic imagining. Since I will confine my discussion to these examples, I will use ‘basic’ simply to mean ‘relevantly like these’, without attempting to spell out which other cases qualify and why.

Next, what does “representational content” mean here? Hutto and Myin specify that what they deny to basic minds is “any notion of content that assumes the existence of some kind of specified correctness condition. To be in a contentful state of mind is to take (“represent,” “claim,” “say,” “assert”) things to be a certain way such that they might not be so.” (Hutto and Myin 2017, p. 10). Content in this sense

need not be linguistically structured, or have any other specific format, nor need it be compositional in the manner of a language. It just requires conditions of some kind of “correctness”, which include conditions of truth, accuracy, veridicality, and so on. One thing to note is that this distinguishes content (which Hutto and Myin deny to basic minds) from intentional directedness onto an object (which they ascribe to basic minds). Basic minds can be directed onto objects, but according to REC they do not “say anything” about how that object is, and so cannot misdescribe it or otherwise represent it falsely.

It is worth emphasising that, in the context of the broader “representation wars”, the debate I am commenting on is a debate between moderates. Hutto and Myin do not reject representational content entirely, and they admit that *some* cognitive processes can only be made sense of by positing representational content. Likewise, the critics I will be considering do not insist on the omnipresence of representation in cognition, and accept that *some* cognitive processes may be nonrepresentational. Consequently, I am setting aside definitions of “representation” which would imply either that cognitive processes are never representational (e.g. Tonneau 2011), or that all brain processes, or even all biological processes, automatically qualify as representation. The question is whether, *given* some understanding of representation that allows for it to be involved in some but not all cognitive processes, Hutto and Myin are right to draw the boundary between the processes distinctive to enculturated humans and those present in nonhuman animals. The critics discussed below argue, and I agree, that if such a boundary is to be drawn, it should lie (roughly) between on-line and off-line processes, that is, between different prelinguistic cognitive processes displayed by at least some nonhuman animals.

2 The challenge from models

The criticism of REC offered by Foglia and Grush (2011) is the most focused on imagining, and receives the most explicit discussion by Hutto and Myin, so I will focus on it; very similar claims about the difficulty of accounting for basic imagining without content are made by Shapiro (2014) and Clowes and Mendonça (2016). Foglia and Grush discuss mental-rotation tasks, and their argument has two steps:

Premise 1: Success in completing a mental-rotation task is, empirically, best explained as involving a mental image of the shape being itself rotated;

Premise 2: A mental image of a shape is a representation of the shape;

Conclusion: *Therefore*, the best empirical account of how subjects solve mental-rotation tasks involves mental representations.

Premise 1 is established by various pieces of evidence (e.g. Foglia and Grush 2011, pp. 37–41; Shepard and Metzler 1971) that are taken to support what Foglia and Grush call an ‘emulation’ account (where the brain simulates both the subject’s sensorimotor activities and the external object), against both ‘simulation’ accounts (where the brain simulates only the subject’s sensorimotor activities) and purely propositional/linguistic accounts (on which there is no role for a ‘mental image’).

The heart of the debate lies in premise 2. Foglia and Grush are careful not to treat mental images as strictly analogous to either sentences or pictures: instead, they describe them as ‘models’:

The crucial idea might be loosely described by saying that something, *M*, is a model of *X* (for some agent *A*) if *A* can use *M* as a stand-in for *X*. It might also be glossed as: *A* can interact with *M* in a manner analogous to how *A* interacts with *X*. (Foglia and Grush 2011, p. 42)

Foglia and Grush do not say very much more than this about how to understand this notion of a model, or why it should be understood as having representational content. Fortunately, Gładziejewski (2016), in the course of a similar but distinct attack on REC, employs what seems to be a substantively similar notion of a ‘model’, said to represent an object by being a ‘stand-in’ for it, and unpacks it much more fully. Gładziejewski is not appealing to imagining specifically, but to the role of ‘generative models’ in Predictive Coding Theory (PCT), which has been primarily developed as an account of perception (though one of its attractions is how readily the models it posits to explain perception could be reused for imagining). Nevertheless it will be useful to examine his account of models, to see how Foglia and Grush’s critique of REC could be spelled out and defended.

Gładziejewski writes that “Generative models guide action by constituting map-like internal stand-ins for worldly states of affairs.” (2016, p. 576) But what is it to constitute a map-like internal stand-in? We are told that:

[R]epresentations as PCT sees them earn their status because the functional roles they play non-trivially resemble the roles played by cartographic maps ... [which] have four functional properties that are constitutive of them playing a role of representations. Maps are:

- (1) structural representations [i.e. ones which structurally resemble some target] that
- (2) guide the actions of their users [with respect to that target],
- (3) do so in a detachable way and
- (4) allow their users to detect representational errors.” (Gładziejewski 2016, p. 566; cf. O’Brien and Opie 2015; Gładziejewski and Miłkowski 2017; Ramsey 2007)

These four criteria apply readily to the visual images rotated in the Shepard and Metzler studies: they structurally resemble the perceived objects in their shape and size; they can be interacted with like those objects, for example, by rotation or comparison to other objects, and guide their users decisions about congruence; they can be used without being able to see the objects; and (most importantly for the argument of this paper) when they diverge from the objects, this can be registered as an error.

3 Models without content?

Hutto and Myin respond by granting, “for the sake of argument” (2017, p. 198) that internal models may play a crucial role in basic imagining, but denying that models must be representations. Thus they deny premise 2, rather than premise 1, in Foglia and Grush’s argument. This commits them to the initially surprising claim that subjects might solve a mental-rotation task by using a mental model of a shape, without that model representing the shape. How does this work? We are told that:

It is plausible that the use of such surrogates makes it possible to respond intelligently even in the absence of modeled items. Even if all of this forms part of the best explanation of how models might enable the completion of certain cognitive tasks, it is important to recognize that such success would still only depend on: certain correspondences holding between the model [and] the modeled, and the imaginer’s systematically engaging with the model appropriately because such correspondences hold. What does not seem to be required, and no argument has been supplied to suggest otherwise, is that imaginers or their subparts need to “take” such correspondences to hold... (Hutto and Myin 2017, p. 200)

It is useful to compare this to Hutto and Myin’s very similar response to Gładziejewski: “even when structural resemblances are exploited for tasks such as mapping a domain and guiding behavior with respect to it ... Additional, independent reasons are needed to explain any content stand-ins might bear.” (2017, p. 158). This claim that ‘stand-ins’ need not have content is developed in two more specific ways: first that meeting Gładziejewski’s criteria for being a stand-in is *insufficient* for having content, and second that attributing content to such stand-ins is explanatorily *superfluous*.

The claim of insufficiency is motivated by a scenario where something qualifies as a stand-in without qualifying as a representation:

Consider a mundane case. Someone, Y, acts as a surrogate for someone else, X, by delegating for X and playing X’s role at some function or event. Clearly, even though we might say that Y is X’s representative, it is clear such surrogacy is achieved without [Y] actually representing [X] in any contentful sense. We can also assume that in order to play the surrogacy role on some occasion, it is further required that Y structurally resembles X. Perhaps Y needs to structurally resemble X in several respects at once. Allow that some of these respects are quite abstract, such as sharing X’s height or X’s capacity for witty repartee. Even in such cases Y does not represent X in any contentful sense. As it is with one person acting as a surrogate or stand-in for another, so it is with a cognitive state or process that acts as a surrogate or stand-in for another. (Hutto and Myin 2017, pp. 157–158)

In the next section I will suggest that this argument fails through not incorporating all four of Gładziejewski’s criteria into the imagined scenario; for now let us move on to the second claim, of explanatory superfluity:

even if we allow that mental models in some sense of that notion are important for understanding basic imaginings, there seems no compelling reason to suppose that assuming such models have representational contents and correctness conditions will help us to understand how basic sensory imaginings execute their important cognitive offices. (Hutto and Myin 2017, p. 201)³

There is room to dispute Hutto and Myin's claim that using one thing as a stand-in for another—exploiting an isomorphism—is not sufficient for content, and gains no explanatory force from positing content. But this focus specifically on exploitation neglects another important feature of the conception of models as representations: error detection. This neglect allows for the attack on REC to be sustained even if all the positive claims made by Hutto and Myin are true.

4 Why this response fails

Suppose we grant Hutto and Myin their central claim: that we can understand a subject *using* a model ('exploiting an isomorphism') without appealing to the model having representational content. I believe their response would still fail, because it considers only the use of an existing model, and not the equally vital process of *shaping* a model. By "shaping" a model I here mean both the initial construction or deployment of a mental model for a task, but also the on-going process of maintaining it, correcting errors that become apparent if they are relevant to its function, and updating it when the thing modelled is found to have changed in relevant ways. After all, someone who mentally rotates a shape *could*, quite easily, imagine a different shape rotating, that does not match the one in front of them; they could start imagining the perceived shape but then deform it or shrink it while it rotates, or they could start imagining it, notice that their image does not actually match the perceived shape, but not bother to correct it. Insofar as subjects avoid these pitfalls we need to consider the processes, whatever they are, that ensure this: these processes are what I am calling the "shaping" of the model.

The shaping process seems to have a very clear role for representational content, understood in Hutto and Myin's own terms as conditions of correctness. In order to have a good model, it must be made to match the thing modelled in the relevant respects; such matching cannot generally just be assumed, or happen by accident, but must be actively ensured. Ensuring it means monitoring the model for relevant sorts of mismatch with the thing modelled, which can trivially be redescribed as monitoring the model for violations of a set of conditions specified by the thing modelled. And when the model does violate those conditions (i.e. when it is found to diverge from the thing modelled in relevant respects) that prompts a specific sort of response, namely adjusting it to meet them. All of this is simply to articulate what must happen for a model to nonaccidentally match the thing modelled.

Are the 'conditions' enforced on a model by the shaping process 'correctness conditions'? They certainly seem like they fit the bill. They are set by the external object which

is being modelled, and when they are violated the adjustments this prompts would naturally be called ‘corrections’. Moreover, there is a clear basis for distinguishing different representations of the same object under different ‘modes of presentation’, namely those which are shaped to resemble the object in different ways—that is, different structural properties might be ‘relevant’ for different models and occasions.⁴

If the foregoing reasoning is sound, it allows for direct answers to Hutto and Myin’s insufficiency argument. Recall their scenario:

Y [is] playing X’s role at some function or event ... in order to play the surrogacy role on some occasion, it is further required that Y structurally resembles X ... [e.g.] sharing X’s height or X’s capacity for witty repartee. Even in such cases Y does not represent X in any contentful sense. (Hutto and Myin 2017, pp. 157–158)

This scenario seems to me to ignore Gładziejewski’s fourth criterion, that a model should allow for “detect[ing] representational errors.” Equivalently, it ignores the necessity of error-detection and error-correction in the shaping of models. To remedy this, suppose we add that Y is assisted by an attendant (possibly Y themselves) who dresses them with an eye to matching X, and during the event checks constantly whether they resemble X in the right ways, intervening when they don’t to correct them. It’s still not that obvious whether we should think of Y as representing X in any intuitive sense (in part because, after all, they are a person and not a state), but what does seem obvious is that what’s going on in this scenario involves correctness/accuracy conditions. After all, Y might spend the whole evening being corrected for their inaccurate portrayal of X.⁵

We can also answer Hutto and Myin’s claim of explanatory superfluity. To say what a model is representing, and which features the model represents it as having, is explanatory in a way that goes beyond simply saying that an isomorphism holds and that this isomorphism is exploited. It explains *why* that isomorphism holds, namely because the system’s resources are devoted to actively constructing and maintaining it; moreover it explains which differences between the model and the thing modelled are failures—inaccuracies—and which are not, which in turn explains why some will tend to be corrected when detected and others will not.

5 Concluding remarks

I cautiously conclude that basic imagining does require representational content. At least, if there is a division to be drawn between cognition that does, and that does not, involve representational content, then imagining, even in its basic forms, should probably be put on the ‘representational’ side. This might support the position of Clowes and Mendonça (2016), who argue that cognitive science should treat some areas (e.g. perception and action, ‘on-line’ processes) in enactive terms, and other areas (e.g. imagination, memory, ‘off-line’ processes) in representationalist terms. However, as Hutto and Myin recognise (2017, p. 180), the problem may not be so easily contained. It is often suggested that imagination somehow plays an essential role in everyday perception—that perception is ‘infused’ with imagination (Clark 2013; Grush 2004; Nanay 2010; Strawson 1974). If that

is true, then representationalism about imagining might support representationalism about perception, or at least about human perception. It would still be left open that there might be nonrepresentational forms of cognition, but human perception would not be one.

Notes

- 1 A fourth potential example mentioned elsewhere in the book is rat dreams: recordings of hippocampal activity seem to show that rats, when resting, simulate taking certain paths through mazes they have seen but not run through (Ólafsdóttir et al. 2015, cf. Bechtel 2016).
- 2 At times Hutto and Myin use the word “basic” specifically to mean “not involving content,” but this would turn the claim under dispute (“basic imagining does not require content”) into a tautology. I prefer to use ‘basic imagining’ for the sort of imagining whose contentfulness is being debated (i.e. the sort of imagining which Hutto and Myin claim is basic in the aforementioned alternative sense).
- 3 An even more succinct statement, by Myin in conversation: “If you have the isomorphism you can give the explanation just in terms of the isomorphism without appealing to content”; cf. Tonneau 2011, p. 342: “Whenever behavior is explained by appealing to an internal isomorph of the environment, what does the explanatory job is the notion of isomorphism and not that of representation.”
- 4 A supporter of REC might respond that the shaping of models still does not involve genuine content, because, as Hutto and Myin put it, we have no reason to say “that imaginers or their subparts need to “take” such correspondences to hold” (2017, p. 200). But this is a red herring, if ‘take’ here has the same meaning as when they say that “To be in a contentful state of mind is to take ... things to be a certain way” (2017, p. 10). For then demanding that the imaginer takes the model to correspond to the external object amounts to demanding that the imaginer represents their model as corresponding to the object. But of course critics of REC are not suggesting that imaginers represent their own models; they are suggesting that imaginers represent external objects, by means of their models.
- 5 This also bears on the question whether teleosemantics is an adequate general theory of content, something which Hutto and Myin repeatedly deny (2017, pp. 43–45, ff, cf. Miłkowski 2015). Teleosemantics is normally thought of in terms of the way an evolutionary process has shaped a mechanism to match certain objects, but can also be thought of in terms of ongoing shaping of particular states by the cognitive system, with personal history taking over from evolutionary history (see esp. Gärtner and Clowes 2017, pp. 66–68).

References

- Bechtel, W. “Investigating Neural Representations: The Tale of Place Cells.” *Synthese* 193 (2016): 1287–1321.
- Clark, A. “Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science.” *Behavioral and Brain Sciences* 36 (2013): 181–253.
- Clowes, R. and D. Mendonça. “Representation Redux: Is there Still a Useful Role for Representation to Play in the Context of Embodied, Dynamicist and Situated Theories of Mind?” *New Ideas in Psychology* 40.A (2016): 26–47.

- Foglia, L. and R. Grush. "The Limitations of a Purely Enactive (Non-representational) Account of Imagery." *Journal of Consciousness Studies* 18.5–6 (2011): 35–43.
- Gärtner, K. and R. Clowes. "Enactivism, Radical Enactivism and Predictive Processing: What Is Radical in Cognitive Science? Kairos." *Journal of Philosophy & Science* 18.1 (2017): 54–83.
- Gładziejewski, P. "Predictive Coding and Representationalism." *Synthese* 193.2 (2016): 559–582.
- Gładziejewski, P. and M. Miłkowski. "Structural Representations: Causally Relevant and Different from Detectors." *Biology and Philosophy* 32.3 (2017): 337–355.
- Grush, R. "The Emulation Theory of Representation: Motor Control, Imagery, and Perception." *Behavioral and Brain Sciences* 27 (2004): 377–442.
- Hutto, D. and E. Myin. *Radicalizing Enactivism: Basic Minds without Content*. Cambridge: The MIT Press, 2013.
- Hutto, D. and E. Myin. *Evolving Enactivism: Basic Minds Meet Content*. Cambridge: The MIT Press, 2017.
- Malafouris, M. *How Things Shape the Mind: A Theory of Material Engagement*. Cambridge: The MIT Press, 2013.
- Miłkowski, M. "The Hard Problem of Content: Solved (Long Ago)." *Studies in Logic, Grammar and Rhetoric* 41.1 (2015): 73–88.
- Nanay, Bence. "Perception and Imagination: Amodal Perception as Mental Imagery." *Philosophical Studies* 50 (2010): 239–254.
- O'Brien, G. and J. Opie. "Intentionality Lite or Analog Content?" *Philosophia* 43.3 (2015): 723–730.
- Ólafsdóttir, F., C. Barry, A. Saleem, D. Hassabis, and H. Spiers. "Hippocampal Place Cells Construct Reward Related Sequences through Unexplored Space." *eLife* 4 (2015): e06063.
- Ramsey, William. *Representation Reconsidered*. Cambridge: Cambridge University Press, 2007.
- Shapiro, L. "Radicalizing Enactivism: Basic Minds without Content, by Daniel D. Hutto and Erik Myin." *Mind* 123.489 (2014): 213–220.
- Shepard, R. N. and J. Metzler. "Mental Rotation of Three-Dimensional Objects." *Science* 171 (1971): 701–703.
- Strawson, P. F. "Imagination and Perception," in *Freedom and Resentment and Other Essays*. London: Methuen, 1974, 45–65. Originally published 1970, in *Experience and Theory*, ed. L. Foster and J.W. Swanson. Amherst: University of Massachusetts Press.
- Tonneau, François. "Metaphor and Truth: A Review of 'Representation Reconsidered' by W. M. Ramsey." *Behavior and Philosophy* 39/40 (2011): 331–343.