

Let The Explanation Fit The Theorist - Enactive Explanatory Pluralism And The Representation Debate

Simon M^cGregor

University of Sussex
s.mcgregor@sussex.ac.uk

Abstract

Discourse in the representation debate within cognitive science employs a problematic model of explanations, which needs to be challenged. What makes a good explanation, even on a philosophically realist interpretation, depends not only on the agreed facts regarding a phenomenon, but also intrinsically on the purposes of the theorist. While one might expect traditional cognitivists to be oblivious to these considerations, it is ironic that outspoken proponents of enactivism ignore them.

It is also open for debate whether there is any objective fact of the matter regarding whether a given system makes use of representational content or not. Intentional systems theory is compatible with anti-realism regarding representations, as are philosophies of science such as constructive empiricism.

I argue that, due to the diversity of the discipline, different theorists within cognitive science have legitimately different explanatory needs, and that this merits both non-representational and representational explanations of the very same system.

Introduction

It is now becoming widely acknowledged that cognition is not a sterile process of information transfer and processing, but rather one that is given meaning and structure by purposeful, embodied, sensorimotor interaction. The mind is neither philosophically nor empirically independent of bodily striving in an agent's *umwelt* (environment-for-the-agent). However, the extent to which traditional approaches to cognition need to be rethought, in order to address embodied, enactive concerns, is a matter of heated debate.

This paper focuses on a specific question: whether the concept of representation always plays a legitimate role in cognitive science explanations, or whether there are cognitive systems for which the concept of representations should not be used. Simulated agents play an important part in this debate, since there is nothing hidden or unknown in their mechanism; hence, the question of whether or not they merit representational explanation is a particularly stark one. For instance, does the notion of representation have a role to play in the analysis of Beer's 'minimally cognitive agents'? (Beer and Williams, 2015; Mirolli, 2012)

I contend that these questions cannot properly be answered without understanding, as van Fraassen (1980) puts it, that "The discussion of explanation went wrong at the very beginning when explanation was conceived of as a relation like description: a relation between a theory and a fact. Really, it is a three-term relation between theory, fact, and context." (van Fraassen, 1980, p.156).

Focusing on simple systems - the Watt governor and bacterial metabolism - I will argue that legitimately differing scientific purposes can require the same system to be explained in representational or non-representational terms, depending, for instance, on whether one is primarily interested in predicting Watt governors' behaviour or, in contrast, in understanding representation-like phenomena. To my mind, this consideration of context and purpose aligns naturally with core enactivist concerns, while an insistence that explanations of a single type should be uniquely privileged for certain types of system seems more in keeping with the spirit of traditional cognitivism.

Structure Of The Article

The term 'enactivist' can be used with a variety of different connotations, and I describe what I mean by it in the next section, recapping Hutto & Myin's distinction between 'radical' and 'conservative' enactivism, in relation to the use of representational explanations of systems. The section following that discusses the relation between models of explanation and philosophical realism. Subsequently, I present the notion of explanatory pluralism and motivate it from (pragmatic) philosophical considerations as well as specific everyday and mathematical examples.

The article proceeds with a discussion of representational explanations in Watt governors, bacteria and Beer's minimally cognitive agents, concluding in each case that the 'correct' explanation depends on the theorist's purposes. A brief section at the end relates explanatory pluralism to reflective enactivism and participation in cultural politics.

Enactivism: Plain, Conservative And Radical

I will follow Hutto and Myin (2013) in using the term “enactivist” very broadly; I mean to include any approach that places purpose, embodiment and real-time interaction with an environment at the centre of the theoretical stage in cognitive science. Enactivism in this sense does not imply any beliefs about the relationship between life and mind (of the sort that autopoietic enactivists such as Thompson and Di Paolo propound) or even that there is a meaningful division between systems with minds and those without. ‘Plain’ enactivism, so construed, involves few philosophical commitments, and its primary significance lies in its import for methodology within the cognitive sciences.

One very significant question regarding scientific methodology is the following: what systems, if any, call for the concept of representation in their explanations? One might distinguish two extreme poles here: traditional cognitivists who insist that *every* cognitive system mandates representational explanation; and hardcore anti-representationalists who insist that representational explanations are *never* valid.

Hutto and Myin (2013) have advocated a relatively mild, but nonetheless significant form of enactivist anti-representationalism that they call *radical* enactive / embodied cognitive science (REC), which asserts that for *some* cognitive systems (Hutto and Myin call them ‘basic minds’), representation / information-processing talk is strictly forbidden. Representational talk is, on this view, appropriate only for particular and special sorts of cognition, ‘emerging late in phylogeny and ontogeny’ (Hutto and Myin, 2013, p.13).

They contrast radical enactivism with *conservative* enactive / embodied cognitive science (CEC), which seeks to absorb enactive concerns entirely within a traditional framework of representation, information processing, and possibly even computation. See, e.g. Clowes and Mendonça (2016) for an overview of how representational talk might fit into embodied, dynamicist and situated theories of mind.

REC is *radical* because it asserts that a revolution is necessary, and that we must start from the ground up with entirely new explanatory concepts: typically, bio-centric ones such as autopoiesis (Thompson, 2007); or perhaps the explanatory tools of dynamical systems theory (Chemero, 2009).

Realism and Pluralism

Theorists like Hutto and Myin and Clark (2016) write as though they are philosophical *realists* about representational content; they assume that there is a factual distinction between representational *systems* and non-representational ones, and that this distinction determines whether a representational or non-representational *explanation* is appropriate for the system. Consider: “our most elementary ways of engaging with the world... *are* non-representational” (Hutto and Myin, 2013, p.13, italics added); “bacteria that follow

chemical gradients... [*do not* deploy] internally represented models” (Clark, 2016, p.84, italics added).

This assumption begs an important question: *anti-realist* and *instrumentalist* theories, including intentional systems theory (Dennett, 1987, 2009), will object to it. (Even assuming realism, it is worth mentioning that if *panpsychism* is possible, it does not logically follow that there must be non-representational systems.)

Where the arguments relate to scientific *practice*, we should bear in mind that the professional obligations of scientists do not include realism regarding unobservable theoretical entities; at most, they require a realist mode of discourse regarding measurable observations. As van Fraassen (1980) argues, explanations *qua scientific activity* need only ‘save the appearances’, i.e. be empirically adequate.

Hence, some writers refrain from having their arguments depend on the realist assumption. For instance, after arguing at length that “identifying representations is an important aspect of neuroscientists’ quest to explain mental phenomena”, Bechtel (2016) avoids the question of “whether [neuroscientists] are correct in their assumption that there are representations in the brain”, providing instead the sociological observation that neuroscientists will continue to employ a representation-focused methodology as long as they feel they are making progress with it.

However, even writers who explicitly make only “epistemic” (or sometimes even merely “methodological”) claims often presuppose that there is a unique objectively best explanation: “our *best* explanations of cognitive systems will not involve representations” (Chemero, 2009, p.67, italics added).

The next section will show that multiple distinct explanations can be valid even if one believes there are objective facts of the matter describing the phenomenon to be explained, let alone if one doubts this premise.

Explanatory Pluralism

This section will advance arguments for explanatory pluralism in everyday discourse and in mathematics. I will distinguish *pluralism*, which accepts *multiple* distinct responses as legitimate answers to a question, from two other positions: the first position, which I will call *strawman relativism*, accepts *any* response as a legitimate answer; the second, which I will call *unique realism*, accepts at most one response as a legitimate answer. Hence, explanatory pluralism accepts multiple distinct narratives as legitimate explanations, but rejects others.

We will begin with a brief overview of pragmatist theories of explanation, following this with illustrations from causal explanation in everyday discourse and mathematical explanation. In the examples given, explanatory value is demonstrably theorist-relative even if the facts of the matter are known.

Pragmatist Philosophy of Explanation

Pragmatist philosophers of explanation such as van Fraassen (1980) and Achinstein (2010) have long emphasised the significance of context for explanation. Achinstein (2010) insists that a *good* explanation must be more than just correct; it must meet the idiosyncratic needs of the audience, in the context of the explanation's momentary use. This can be evaluated only by the exercise of human social expertise, and not according to any "objective, nonpragmatic, universal criteria" (Achinstein, 2010, p.167). Similarly, van Fraassen requires that explanations be evaluated according to whether they are *acceptable*: a criterion which includes such contextual factors.

However, a realist notion of correctness may beg the question in regard to representations, and van Fraassen's philosophy of constructive empiricism provides a metaphysically weaker criterion: whether an explanation is *empirically adequate*, i.e. whether there exists a (model-theoretic) model of the theory that contains structures isomorphic to (past and future) observable data. According to van Fraassen (1980), empirical adequacy does the same job as truth (as far as scientific practice is concerned), but requires fewer philosophical commitments.

It is worth noting that *both* of these pragmatist philosophers of explanation offer a view of scientific explanation that is pluralist, despite one being a realist regarding theoretical entities and one an anti-realist regarding them. Achinstein is realist without being a unique realist regarding explanations, because he considers that multiple different explanations can be objectively correct; while van Fraassen is anti-realist without being a strawman relativist, because he rejects scientific explanations that are not empirically adequate.

Explanatory Pluralism In Everyday Discourse

In general, it is fairly obvious that the same phenomenon can admit multiple, distinct, explanatory narratives that are all legitimate. van Fraassen (1980, p.125) quotes Norwood Russell Hanson here on causal narratives: "Consider how the cause of death might have been set out by a physician as 'multiple haemorrhage', by the barrister as 'negligence on the part of the driver', ... by a civic planner as 'the presence of tall shrubbery at that turning'."

The novelist Michael Frayn gives a similar example, regarding the explanation for the cruise liner Titanic's sinking: "[T]he most salient cause' is a concept that depends entirely upon the point of view of the proposer, and the use to which it is being put... For the commission of enquiry charged with establishing responsibility, ... the salient cause might appear to be the captain's decision... For the business analyst, it would be the competitive pressures of the North Atlantic passenger trade..."

The 'right' explanation is underdetermined by the facts; there are additional contextual factors, related to the expla-

nation's *purpose*, that lead the physician to employ one explanation and the civic planner another. To demonstrate that the identity of the explainer is relevant only insofar as it affects the purpose of the explanation, imagine that the physician, barrister and civic planner were one and the same person: in their *role* as physician, the correct explanation would still be multiple haemorrhage, while in their role as barrister, the correct explanation would still be negligence.

Explanatory Pluralism In Mathematical Proofs

In case the reader supposes this explanatory multiplicity to be an artefact of complex multicausal webs, we will consider its operation in the mathematical realm, where I will take a mathematical proof to be an explanation of why a particular result holds. A paper by Wagon (1994) gives fourteen proofs of a result due to de Bruijn (e.g. using integration or induction), arguing that the proofs are all different since they lead to different generalisations (e.g. integration allows the theorem to be proved for any dimension, while induction allows it to be proved for toroidal or cylindrical topologies).

Which of these proofs is the best explanation of why the theorem holds? It goes without saying that the best explanation must be a valid proof in a suitable formal system, and some (arguably) theorist-independent factors such as the length of the proof, or the number of axioms assumed by it, may also be relevant.

But surely the best explanation must also depend on the theorist's familiarity with the branch of mathematics in question, and whether the proof readily permits the result to be generalised in directions that the theorist cares about. Note that this latter criterion can in principle be given a rigorous and explicitly purpose-relative meaning: say, the number of additional steps that are required to prove theorems in a particular target set T , given the original proof.

On this view, mathematical (and, by extension, scientific) explanations need to be seen as tools in an active cognitive process, where the best explanation is relative to the current explanatory needs of the theorist.

Explanatory Pluralism And Representations The Watt Governor

Chemero appeals to a pragmatic principle when he asserts that the dynamical systems story of the Watt governor tells us everything that the representational story does. But I have argued that pragmatic considerations are relative; and I contend that Chemero has underestimated the diversity of possible explanatory needs a theorist might have.

Firstly, the dynamical systems story of the Watt governor only tells us everything that we might want to know *about the behaviour of the Watt governor*. We may want a representational explanation of how the Watt governor works for other legitimate cognitive science purposes: for instance, for understanding how representational stories relate to dynamical systems ones.

Clowes and Mendonça (2016, sec 3.4) make the same point with respect to so-called basic minds: “One great advantage of specifying a basic mode of representation is that it allows us to make sense of the various extended modes of representation without postulating a strong discontinuity.”

The second objection to Chemero’s claim (that the dynamical systems story of the Watt governor tells us everything that the representational story does) is that there are problems which will *in practice* be easier to solve using a representational story than a non-representational one. For instance, suppose we want to put in-principle upper bounds on the average performance that the governor can exhibit when various component couplings are subject to noise. This is a non-trivial mathematical problem when faced with the raw dynamical equations, and it is the sort of thing that can often be solved more easily when a system is interpreted as dealing in task-relevant information about its environment: for instance, using the framework in Touchette and Lloyd (2004).

Note that it is precisely the fact of a theorist’s physical embodiment that imposes these convenience considerations. If theorists did not care about the time spent in coming to a conclusion, or the metabolic energy expended in doing so, there would be no pragmatic reason to prefer one form of notation over another. There is a parallel with Nozick’s famous Martian, to whom humans are almost as simple as thermostats (Dennett, 2009, p.25): the Martian has no incentive to use the intentional stance (or, correspondingly employ representational explanations), since it, like a Laplacian God, can effortlessly use the more detailed physical stance to make all the same predictions.

Anti-representationalists are not the only ones to privilege one type of explanation for the Watt governor: for instance, Bechtel (1998) argues that the spindle arm angle does represent the speed of the flywheel, since it serves the mechanical purpose of standing in for the flywheel’s speed: “It is because the spindle arms rise and fall in response to the speed of the flywheel that the angle of the arms can be used by the linkage mechanism to open and shut the valve.” The same paragraph concludes by invoking explanations: “[I]f someone does not understand how the governor works, the first thing one would draw attention to is how the spindle arm angle registers the speed of the flywheel.”

But as Bechtel points out in the same article, Chemero’s dynamical systems equations conform to a different model of explanation: a ‘covering law’ that describes regularities in the governor’s behaviour, rather than a ‘mechanistic’ explanation identifying how linkages between separable parts achieve a purpose; Bechtel’s argument can at best establish that representational talk is called for when the theorist wants to understand *how the Watt governor works*, not when they merely wish to *predict its behaviour*.

Bacterial Cognition

Even a theorist like Clark (2016), who endorses Friston’s (representational) Bayesian hierarchical predictive coding account of the human brain (e.g. Friston et al., 2010), agrees with Hutto & Myin that bacterial chemotaxis is non-representational. But this line, drawn somewhere between humans and bacteria, is unnecessary. Cellular and systems biologists are very interested in information-processing within bacteria: for instance, Libby et al. (2007) argue that the transcription rate of the *lac* operon in *e. coli* should be interpreted as encoding a Bayesian posterior over what sort of environment the bacterium is in, given the concentrations of certain intracellular chemicals.

It should be emphasised here that a Bayesian posterior is far more than a mere correlation. A posterior consists of a conditional probability distribution; according to the predominant perspective in Bayesian statistics, probability distributions represent quantifications of subjective uncertainty, and conditionalisation represents an ideally rational process of belief modification based on observable evidence.

In other words, Libby et al. are explaining the transcription rate of the *lac* operon as approximately encoding an *ideal inference* about survival-relevant aspects of the environment, based on noisy intracellular data. If they are right, then key elements of bacterial *metabolism* (never mind chemotaxis) can be interpreted representationally within the same fundamentally Bayesian paradigm that Clark favours for the brain, and Clark is potentially inconsistent in asserting that bacterial cognition is non-representational.

The question of whether or not *lac* operon transcription *intrinsically* mandates a representational explanation arguably lies outside the scope of science. Philosophers will no doubt continue to debate it, but as far as scientific practice goes, I suggest that a representational explanation of *lac* operon transcription appears to be appropriate for the explanatory purposes of Libby et al., and that this does not logically imply that it is the right type of explanation for everyone.

Beer’s Minimally Cognitive Agents

Traditional ‘good old-fashioned’ approaches to artificial intelligence made not only the cognitivist assumption that representations were essential to (or constitutive of) cognition; they went further in assuming that the best ways to build cognitive systems were by explicitly modelling the processing of representational content, labelled as such in a program’s source code.

This view came under criticism not only from ‘nouvelle AI’ roboticists such as Brooks (1991), but from models of simple cognition in Artificial Life. A famous early example involves Beer’s (2003) simulated minimal agents, which perform discriminatory locomotion tasks involving falling objects, using limited sensory information. The controllers for these agents are evolved neural networks, and their dy-

namics can be analysed using the mathematics of continuous dynamical systems, whose lexicon does not include the notion of representation.

Are representational explanations of these agents useful? In fact, both Mirolli (2012) and Beer and Williams (2015) themselves conclude that, in a suitable (action-oriented, teleologically grounded) sense, it is appropriate to invoke representations in the evolved agents (at least for tasks involving memory). For them, a representational state “must have the purpose or function of carrying the information it does”, where information carried is to be understood within a technical framework of Shannon information theory.

Interestingly, Beer and Williams stop short of claiming that these representations are an objective fact of the matter: “Because our model was evolved for performing a relational categorization task, such a functional story can be told about its operation.” Moreover, Beaton and Aleksander (2012) have argued that Shannon information itself should be understood as theorist-relative, with variables providing subjective information content for the theorist.

In the same paper, Beer and Williams apply both information-theoretic and dynamical systems theoretic analysis to one of the evolved agents, and conclude that “information processing and dynamical explanations of our model agent are both consistent and complementary.” For instance, “although [information theory] can describe the rise of relative size information in N1 as an important event in the agent’s categorization, [dynamical systems theory] is required to understand the bifurcation that underlies it.”

These authors, who operate firmly within the enactive and dynamical tradition, take a pluralist perspective that explicitly recognises the context-relativity of good explanations: “Each lens provides a unique explanatory perspective, distinctive language, and a complementary set of tools for framing and answering certain kinds of questions.”

Reflection, Enactivism and Realism

Scientific theorising is surely an exemplar of cognitive activity. Hence, a general scientific theory of cognition should be able to account for the theorist themselves, and their engagement with that very theory.

Traditional cognitivist conceptions of mind have attempted to characterise cognition in terms of the processing of abstract mental representations that qualify as such by reflecting the world, with no reference to any teleological context. Within this paradigm, reflective practice is straightforward: scientific theories are representations and reflect the world, which is a unique objective entity.

However, enactivist conceptions of cognition emphasise the meaning of the interaction for the agent, which necessarily differs from individual to individual. This has different implications for reflective practice: von Uexküll (1934) emphasises that the *umbegung* of a biological system (its ‘objective’ physical surroundings) is nothing more than the

umwelt (the subjectively meaningful world) of the human theorist as a biological system.

Similarly, second-order cyberneticists have long insisted on acknowledging the theorist’s perspective in the framing of cognitive science: ‘Everything said is said by someone’ (Maturana and Varela, 1987); ‘From the constructivist perspective, science cannot transcend the domain of experience’ (von Glasersfeld, 2001).

There is, at the very least, a notable irony in Hutto and Myin opposing even action-oriented representations for basic minds (which are held to be mental by virtue of purposeful action), and at the same time employing a realist mode of discourse, where scientific explanations apparently reflect the external world in the objective, purpose-independent manner of a traditional cognitivist representation.

Pluralism and Cultural Politics

I have advocated an explanatory pluralism, motivated by a number of quintessentially enactivist concerns: a reflexive interest in the theorist as, simultaneously, object and active user of the theory; a focus on fitness for subjective purpose (as opposed to correspondence with objective reality); and an appreciation of how the theorist’s embodiment affects the uses to which they can put a particular explanatory narrative.

This perspective is also, as I hope has been evident, motivated by political concerns regarding diversity and power relations. By reflecting, self-consciously, on the ways in which scientific practice involves our own agency, we can allow that other theorists may have explanatory purposes that are not identical with our own. We must then decide, as a matter of cultural politics, which purposes we are prepared to recognise as legitimately scientific ones (i.e. having a place in the cooperatively constituted endeavour of scientific discourse).

Cognitive science is an extraordinarily diverse discipline. We should bear in mind, when discussing ‘our’ best explanations, as Chemero does, that there is not guaranteed to be a relevant ‘we’: the best explanation for Chemero’s purposes may genuinely not be the best explanation for Beer and Williams’s purposes, or for Bechtel’s.

Under this view, attempts to privilege representational (or non-representational) explanations constitute a worrying power play: one that seeks to dismiss the genuine diversity of explanatory needs within the multidisciplinary field of cognitive science. Is there really a one-size-fits-all explanatory stance (even for ‘basic minds’) that will serve the purposes of all roboticists, psychologists, bacteriologists, machine learning researchers, philosophers, neuroscientists, non-equilibrium physicists and ethologists?

For my own part, a reflective participation in scientific sense-making inclines me to take an anti-realist perspective, in the sense of Dummett (1978): there is *no fact of the matter* about whether a system uses representational content, any more than there is a fact of the matter about whether a spa-

tial trajectory is in polar or rectangular coordinates. This view resembles intentional systems theory (Dennett, 2009), but emphasises a fine-grained relativism, and does not assume that one type of explanation need necessarily be more detailed or abstract than the other.

Nonetheless, there would be a certain hypocrisy were I to assert that realism regarding representations (and explanations) was objectively incorrect. As Stewart (2001) comments, “constructivism actually rehabilitates objectivism as a viable, pragmatic possibility in the field of cognitive science”. I can only observe that I do not, personally, find the realist view helpful; hence, I (and others like me) can be expected to oppose realist attempts to purge representational language from (or impose it upon) our explanatory vocabulary.

Summary

Discourse around the concept of representation in cognitive science explanations tends to default to a unique realist model, in which there is a single correct explanation for real-world phenomena. I have argued, contrarily, that both representational and non-representational explanations of systems simple enough to qualify as basic minds (including bacteria, Watt governors and Beer agents) can be scientifically legitimate, depending on the theorist’s individual explanatory needs.

There are non-realist perspectives available to cognitive scientists: for instance, according to van Fraassen’s constructive empiricism, theoretical entities do not have to be seen as real; von Glasersfeld’s radical constructivism says that there is a real world but that we can say nothing about it. However, even if (for the sake of argument) we grant a realist metaphysics, it is fallacious to suppose this entails a uniquely best explanation, since the right explanation depends on a theorist’s purposes, as well as their available cognitive resources.

There is space in cognitive science for more than one way of understanding cognitive systems. I am with Gärtner and Clowes (2017) when they conclude, “we prefer to let a hundred flowers bloom”.

Acknowledgments

I would like to thank several anonymous reviewers whose detailed feedback helped me to improve the paper.

References

Achinstein, P. (2010). *Evidence, explanation, and realism: Essays in philosophy of science*. Oxford University Press.

Beaton, M. and Aleksander, I. (2012). World-related integrated information: enactivist and phenomenal perspectives. *International Journal of Machine Consciousness*, 4(02):439–455.

Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist’s challenge in cognitive science. *Cognitive Science*, 22(3):295–318.

Bechtel, W. (2016). Investigating neural representations: the tale of place cells. *Synthese*, 193(5):1287–1321.

Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243.

Beer, R. D. and Williams, P. L. (2015). Information processing and dynamics in minimally cognitive agents. *Cognitive science*, 39(1):1–38.

Brooks, R. A. (1991). Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159.

Chemero, A. (2009). *Radical embodied cognitive science*. MIT press.

Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

Clowes, R. W. and Mendonça, D. (2016). 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*, 40:26–47.

Dennett, D. (1987). *The Intentional Stance*. MIT Press, Cambridge, MA.

Dennett, D. (2009). Intentional systems theory. *The Oxford handbook of philosophy of mind*, pages 339–350.

Dummett, M. (1978). *Truth and other enigmas*. Harvard University Press.

Friston, K. J., Daunizeau, J., Kilner, J., and Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biol Cybern*, 102(3):227–260.

Gärtner, F. and Clowes, R. W. (2017). Enactivism, radical enactivism and predictive processing: What is radical in cognitive science? *KAIROS*, forthcoming.

Hutto, D. D. and Myin, E. (2013). *Radicalizing enactivism: Basic minds without content*. Mit Press.

Libby, E., Perkins, T. J., and Swain, P. S. (2007). Noisy information processing through transcriptional regulation. *Proceedings of the National Academy of Sciences*, 104(17):7151–7156.

Maturana, H. R. and Varela, F. J. (1987). *The Tree of Knowledge: Biological Roots of Human Understanding*. Shambhala, Boston, MA.

- Mirolli, M. (2012). Representations in dynamical embodied agents: Re-analyzing a minimally cognitive model agent. *Cognitive Science*, 36(5):870–895.
- Stewart, J. (2001). Radical constructivism in biology and cognitive science. *Foundations of Science*, 6(1–3):99–124.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Touchette, H. and Lloyd, S. (2004). Information-theoretic approach to the study of control systems. *Physica A: Statistical Mechanics and its Applications*, 331(1):140–172.
- van Fraassen, B. C. (1980). *The scientific image*. Oxford University Press.
- von Glasersfeld, E. (2001). The radical constructivist view of science. *Foundations of Science*, 6(1):31–43.
- von Uexküll, J. (1934). *Streifzüge durch die Umwelten von Tieren und Menschen*. Springer. Read in English trans. J.D. O’Neill, “A Foray Into The Worlds Of Animals And Men”, 2010.
- Wagon, S. (1994). Fourteen proofs of a result about tiling a rectangle. In *The Lighter Side of Mathematics: Proceedings of the Eugène Strens Memorial Conference on Recreational Mathematics and Its History*, page 113. Cambridge University Press.