



A continuum of intentionality: linking the biogenic and anthropogenic approaches to cognition

Matthew Sims¹

Received: 28 May 2020 / Accepted: 22 October 2021 / Published online: 3 November 2021
© The Author(s) 2021

Abstract

Biogenic approaches investigate cognition from the standpoint of evolutionary function, asking what cognition does for a living system and then looking for common principles and exhibitions of cognitive strategies in a vast array of living systems—non-neural to neural. One worry which arises for the biogenic approach is that it is overly permissive in terms of what it construes as cognition. In this paper I critically engage with a recent instance of this way of criticising biogenic approaches in order to clarify their theoretical commitments and prospects. In his critique of the biogenic approach, Fred Adams (Stud Hist Philos Sci 68:20–30, 10.1016/j.shpsa.2017.11.007, 2018) uses the presence of intentional states with conceptual content as a criterion to demarcate cognition-driven behaviour from mere sensory response. In this paper I agree with Adams that intentionality is the mark of the cognitive, but simultaneously reject his overly restrictive conception of intentionality. I argue that understanding intentionality *simpliciter* as the mark of the mental is compatible with endorsing the biogenic approach. I argue that because cognitive science is not exclusively interested in behaviour driven by intentional states with the kind of content Adams demands, the biogenic approach's status as an approach to cognition is not called into question. I then go on to propose a novel view of intentionality whereby it is seen to exist along a continuum which increases in the degree of representational complexity: how far into the future representational content can be directed and drive anticipatory behaviour. Understanding intentionality as existing along a continuum allows biogenic approaches and anthropogenic approaches to investigate the same overarching capacity of cognition as expressed in its different forms positioned along the continuum of intentionality. Even if all organisms engage in some behaviour that is driven by weak intentional dynamics, this does not suggest that every behaviour of all organisms is so driven. As such, the worry that the biogenic approach is overly permissive can be avoided.

Keywords Intentionality · Biogenic approach · Anticipation · Allostatic control · Basal cognition · Representational complexity · Evolution of cognition

Introduction

One kind of starting point for investigating cognition begins by asking questions pitched at the kind of ‘high-level’ capacities paradigmatic of human cognition (e.g., beliefs, desires, concept formation, language, reasoning, consciousness, etc.). Some examples of such questions are: Is rationality a necessary property of intentional systems? (Dennett 1983)? Do all thoughts have a language-like, propositional structure (Fodor 1975)? What kind of neural processing does consciousness involve (Carruthers 2019)? This starting point then attempts to account for possible instances of such ‘high level’ capacities within an explanatory framework tailored specifically to them. Philosophers taking this starting point tend to conceptualize cognition in terms of different mental states which are taken to be individuated—like linguistic performances—by their truth conditions (Cf. Fodor 1987). Although often empirically-informed, biological considerations regarding cognition when used are non-essential and generally serve as a means to theorize about features typical of human cognition. From this *anthropogenic* starting point, human cognition is often viewed as the standard against which all other forms of behaviour are recognized as a proper topic for cognitive scientific inquiry.

Another kind of starting point, the *biogenic approach* (Lyon 2006) sees biology as essential to understanding cognition. This biologically-grounded approach (or strictly speaking, family of approaches) starts from an evolutionary continuity assumption that ‘higher’ cognition in humans is different in degree but not in kind from cognition in simpler organisms (Darwin 1871). It holds that the most fruitful route to understanding cognition in its various forms is to begin with considerations regarding adaptivity, fitness, and principles of biological organization. As such the biogenic approach (BA) starts with biological facts—as typically exemplified by the adaptive behaviour of basal organisms—and attempts to work its way up to human cognition. Crucially, by widening the scope of model organisms under consideration to include organisms such as bacteria and plants as cognizers, proponents of BA suggest that it is well placed to identify possible organizational principles and conserved mechanisms underlying a full range of cognitive capacities, from simple to complex (see Ben Jacob et al. 2006; van Duijn et al. 2006; Baluška and Mancuso 2009; Bechtel 2014; Baluška and Levin 2016; Levin 2019; Ginsburg and Jablonka 2019; Barrett 2019; Calvo et al. 2020; Bechtel and Bich 2021).

One worry which arises for BA is that it is overly permissive in terms of what it construes as *bona fide* cognition. The behaviour of bacteria and plants is something that philosophers have been prone to consider as an inflexible, reflexive response to sensory stimuli (Dennett 1984; Sterelny 2001; Godfrey-Smith 2002; Schlicht 2018). If such behaviour under BA qualifies as an expression of cognition, then it seems—one might worry—that the very notion of cognition is rendered too unconstrained to be explanatorily useful. This kind of worry is voiced by Fred Adams (2018), who in a recent critique of BA responds by providing a strict demarcation criterion that allows for drawing a sharp line between

the cognitive and the non-cognitive. This classification methodology, which is informed by traditional philosophy of cognitive science, uses the presence of intentional states with conceptual-level content as a manner of picking out behaviour that is underwritten by cognitive processing. Armed with this conception of the mark of the cognitive, Adams mounts an argument against BA: because the various cases of behaviour that biogenic theorists have classified as cognitive are not underwritten by intentional states with conceptual-level content, such cases of behaviour fail to be expressions of cognition; they are more accurately classified as hardwired responses to sensory information processing. What I call Adams' 'argument from intentional content' goes on to conclude that this conflation is harmful to the endeavours of cognitive science and that BA should be abandoned as an approach to cognition.

The general worry that BA is overly permissive (or even incoherent) relies on an implicit appeal to some mark of the cognitive. Adams is thus useful to engage with because he makes a rare attempt to explicitly spell out what this mark of the cognitive is in a way that purports to be informed by the nature and methodology of cognitive science. Identifying the infelicities of Adams' critique gives us a better view of the explanatory tools that BA might have at its disposal, and the manner in which anthropogenic and biogenic approaches relate to each other. Is there, however, another classification methodology available that is compatible with the range of cognitive behaviours recognized by biogenic theorists, yet which nonetheless offers a manner of distinguishing the cognitive from the non-cognitive? In what follows, I will argue that this question may be answered affirmatively. For from the fact that many cognitive scientists have been interested in some specific mode of functional organization such as intentional states with conceptual-level content, "one cannot infer the difference between exhibiting that mode of organization and not exhibiting it must in some way play a decisive role in marking off the cognitive from the noncognitive" (Wheeler 2010, pp. 256–267). The difference between the presence of that mode of functional organization and its absence may simply fall within the domain of the cognitive and hence within the domain of cognitive science.

For the purposes of this paper, I propose to grant Adams' contention that intentionality is the mark of the cognitive. However, unlike Adams, the position that will be taken here is that it is intentionality *simpliciter* rather than intentionality with conceptual-level content that is the mark of the cognitive. This view of intentionality, I shall argue, is compatible with the manner in which biogenic theorists approach cognition. I do so by proposing and arguing for the *intentionality continuum thesis*, a thesis that is both biogenic and yet informed by looking at the range of projects cognitive scientists actually spend their time doing. On one end of the taxonomy I present, there is *weak intentionality* of the internal dynamics that drive anticipatory sensorimotor and biochemical behaviour. Meaningful content at the level of weak intentionality is, we shall see, phenotype-relative. It is directed at yet-to-be encountered environmental/sensory states that are to be pursued (or avoided) given the conditions which an agent must remain in in order to preserve its organization over the long run. On the other end there is *strong intentionality* which involves mental states with conceptual content that is deployable offline. Intentionality is a necessary feature of cognition, but where a particular behaviour/state falls upon the continuum of

intentionality is of no relevance to that behaviour's being (or not being) an expression of cognition. I shall argue Adams' anthropogenic-based hallmark of the cognitive looks arbitrary once we appreciate that the intentionality that is of interest to cognitive scientists is not exhausted by intentional states with truth evaluable content. Given that contemporary cognitive science does indeed investigate and recognize putative cognitive phenomena and mechanisms that fall all along the continuum of intentionality, the manner in which Adams delineates cognitive from non-cognitive processes is undermined by the range of interests and aims of researchers in the very discipline he attempts to safeguard. Understanding intentionality as existing along a continuum allows biogenic approaches and anthropogenic approaches to investigate the same overarching capacity of cognition as expressed in its different forms positioned along the continuum. Even if all organisms engage in some behaviour that is driven by weak intentionality, this does not suggest that every behaviour of all organisms is so driven. As such, the general worry that BA is overly permissive can be avoided. The continuum thesis is promising not only as grounds for a reply to cognitive chauvinists who use an overly sophisticated notion of cognition, but because it illustrates that both biogenic and anthropogenic approaches are valid starting points for investigating a common notion of cognition.

This paper will proceed as follows: Sect. 2, provides an overview of BA. In Sect. 3, I present Adams' argument from intentional content against BA. In Sect. 4, after quickly surveying some of the practices and projects of contemporary cognitive science, I argue that because cognitive science is not exclusively interested in behaviour driven by intentional states with truth evaluable representational content, BA's status as an approach to cognition is not called into question in the way that Adams suggests. I then propose a novel view of intentionality whereby it is seen to exist along a continuum that ranges in the degree of representational complexity. In this view complexity is a measure of how far into the future representational content can be directed and drive anticipatory behaviour. This section concludes by returning to the worry that BA is overly permissive. I argue that even if cognition turns out to be a ubiquitous feature of living organisms, it is simply not the case that everything that an organism does would qualify as cognition or that the cognition that various organisms engage in could all be accounted for in the same way. As such, the worry that cognition becomes a trivial notion under BA is misplaced.

The biogenic approach

Biogenic approaches constitute a loose family of approaches that are unified by themes and motivations rather than strict commitments.¹ In her seminal paper 'The biogenic approach to cognition' Pamela Lyon (2006) identifies a particular investigatory strategy which she claims has been central to two distinct subfamilies of biogenic cognitive explanation: *self-organizing complex systems theories* and *autopoietic theory*. The first of these kinds of explanation, examples of which may

¹ I am grateful to an anonymous reviewer for pushing me to clarify this important point.

be seen in the work of Goodson (2003), Rosen (1985/2012), Piaget (1970), Popper (1965/1972), Vertosick (2002), Christensen (2004), Bickhard (2009), and Deacon (2012), among others, places emphasis upon the relation between cognition and the second law of thermodynamics. More precisely, it sees cognition as a process that underwrites an organism's ability to remain in thermodynamically improbable non-equilibrium steady states despite the tendency for all systems to move towards thermodynamic equilibrium. Homeostasis, the preservation of steady-states (e.g., physiological states such as oxygenation, core temperature, and metabolic energy levels), is an ongoing result of a dissipative system's ability to modulate its parameters (i.e., internal processes and external behaviour) so as to avoid persistence-compromising situations (e.g., starvation or predation).² Self-organizing complex systems theorists thus see cognition as the host of capacities, varying in their complexity, that guide a system's parameter modulation so as to preserve its homeostatic balance and keep it at thermodynamic disequilibrium despite environmental perturbation.

In contrast, rather than understanding cognition in terms of thermodynamic regulation, autopoietic theory (Maturana 1970/1980; Maturana and Varela, 1980; Thompson 2007; Di Paolo 2005) understands cognition in terms of a biological system's ability to continuously produce itself and distinguish itself from its milieu. *Self-production* is the internally driven process with which a biological system continuously constructs and realizes its own network of processes (Maturana and Varela 1980).³ *Self-distinction*, on the other hand, is the process of generating and preserving systemic boundaries (e.g., a cell membrane) by which an organism partitions itself from its non-systemic environment. Self-production and self-distinction are construed as necessary and sufficient conditions on life. According to autopoietic theory, cognition is entailed by life, expressing a deep continuity between life and mind. In other words, "living systems are cognitive systems, and living as a process is a process of cognition" (Maturana and Varela 1980, p. 13).⁴ The overlap between self-organizing complex systems theory and autopoietic theory should be clear: both kinds of theory construe cognition in terms of the functional contribution that it makes towards an organism's adapting to its environment so as to remain alive. Self-organizing complex systems theory emphasizes capacities to navigate energetically dynamic environments so as to self-preserve; autopoietic theory emphasizes the capacities leading to self-production of the organismic boundary that defines the autonomous organism in relation to its external environment.

This overlap is illustrative of the core explanatory strategy of BA; BA starts from the notion that cognition is a functional capacity that, being driven by various selection pressures, has been gradually selected for—perhaps many times—over

² A dissipative system is a thermodynamically open system that, which in virtue of exchanging matter and energy with the environment, remains far from thermodynamic equilibrium.

³ There are important disagreements between 'autopoietic theory' as propounded by Maturana and the enactive approaches developed by Varela, Thompson and Di Paolo. See Villalobos and Ward (2015).

⁴ Although autopoietic theory is an example of a biogenic explanatory approach, it is not required that a biogenic approach be committed to the kind of entailment thesis (i.e., that life entails mind and mind entails life) that is central to autopoietic theory. The entailment thesis, for example, is not held by self-organizing complex systems theories.

evolutionary timescales. Thus, BA starts with the facts of biology in attempting to understand *what cognition does for biological systems*. This is clear when Lyon writes:

An investigator adopting a biogenic approach assumes that the principles of biological organization and the requirements of survival and reproduction present the most productive route to a general understanding of the principles of cognition. Cognition, whatever it may be in the future, is naturally a biological process and a biological function (2006, p. 12).

The explanatory agenda of BA is determined by biological considerations. We saw above that BA may be contrasted with the anthropogenic approach (AA) (Lyon 2006). Generally, cognition in AA is conceptualized with no essential reference to biological function. AA like BA is a loose family of approaches, which rather than sharing any particular commitments, are unified by common explanatory foci and motivations. The cognitive is often demarcated from the non-cognitive by using distinctively human capacities (e.g., believing, thinking, planning, decision making, reasoning, etc.) as the non-negotiable standard of cognition. The notion of cognition that these kinds of capacities and their putative mechanisms suggest is then used comparatively to determine whether or not the behaviour of non-human systems should be considered as expressive of cognition. It is in this sense that AA works from the human case “downwards” to possible cases of non-human cognition (Lyon 2006). Because of its human-centred demarcation criteria, AA generally assumes that the possession of a central nervous system, the use of representational states, and rationality are all central features of cognition.

To take an example from the philosophy of perception literature to illustrate the different approaches. Tyler Burge (2010) in his opus ‘Origins of Objectivity’ considers what makes perception distinct from mere sensory registration. Although his empirically-informed analytical method is biologically thoroughgoing, his *desiderata* for what constitutes *bona fide* perception are thoroughly anthropogenic. For instance, one of Burge’s criteria for perception is that it must involve the possibility of representational success and failure. In other words, perceptual states must be either veridical or non-veridical. Crucially, the veridicality conditions that are used to individual perceptual states, according to Burge, are not the same as—nor do they answer to—the conditions under which an organism would or would not survive; a state’s being veridical is not the same as a state’s functioning in a manner that reflects its selection history. Considering the biological function of perception is inessential to his analysis. He motivates his use of this veridicality criterion with the fact that perceptual psychology has been successful in accounting for the problem of perception *and* that the veridical/non-veridical distinction has been explanatorily central in accounting for the problem.⁵ Given that perceptual psychology and the problem of perception has been largely grounded in human perception, this criterion for perception is an instance in which putative features of human cognition are used

⁵ This is the problem of how the perception of particulars and properties is possible given impoverished and/or underdetermining proximal stimuli.

to carve nature at its joints. As a result of his analysis, Burge concludes, speculatively, that *bona fide* perception can be found as far down the phylogenetic tree as arachnids.

Compare Burge's approach to perception with that of Paco Calvo et al.'s (2020) approach to perception, the latter of which is easily interpretable as an instance of researchers taking a biogenic approach. Calvo and colleagues analyse plant perception (and cognition more generally) within an ecological context. This context, a real-world environment as opposed to a laboratory environment, is one that houses a variety of selection pressures (e.g., predation, resource acquisition, etc.) that any organism must behaviourally adapt to in order to bring about its continued survival. Cognition is seen as any adaptive behaviour that improves an individual's chances for survival over the course of its lifetime. Against this background, Calvo et al. (2020) analyse perception as part of the more encompassing process of cognition from within the framework of complex systems theory; perception both feeds into learning/memory and assessment systems and receives feedback from behavioural response systems. As such, perception and behaviour are understood as a cyclical process, which, *pace* Burge, is not in the game of constructing veridical representations but rather of controlling adaptive behaviour that allows a system to persist in the face of environmental stressors. The notion of veridicality is bypassed altogether.⁶ Importantly, to investigate how and what a particular kind of organism perceives Calvo et al. stress that one must look to the selection pressures that such an organism must cope with in its niche. On this account, whether or not (and what) plants perceive is not a question that can be addressed by using standards derived from a perceptual psychology that has been tailored to human capacities. On the other hand, one may start with identifying basic cognitive organizational principles as exhibited by basal organisms and plants and attempt to understand more complex forms of cognition in relation to them (i.e., in terms of the complexification of cognitive architectures, internal models, and functional and structural organization).

By understanding cognitive capacities in terms of their biological function to guide adaptive behaviour *from the very start*, biogenic theorists are not constrained by human-centred demarcation criteria; the proponent of BA may choose to investigate how the same biological function may be expressed in humans as well as in non-neural organisms (e.g., bacteria, slime moulds, plants etc.) despite the fact that such living systems differ vastly in terms of their organismal complexity. This being said, there is an interesting overlap between BA and AA. It is exposed by the fact that the distinction between BA and AA "relates to starting assumptions about the problem space, not to how 'biological' a proposal is" (Lyon 2006, p. 16). For example, proponents of BA needn't reject that some forms of cognition may involve the kinds of features which are taken by AA to be central to cognition. For example, Goodson (2003), a biogenic theorist through and through, provides an account of cognition that helps itself to representational states. His use of representations in explanation is consistent with the fact that his investigation of cognition is premised

⁶ For an argument as to how Burge's veridicality constraint on perception can be satisfied by plants when veridicality is construed as accurate enough for the guidance of adaptive behaviour see (Sims 2019).

upon the assumption that cognition is a biological process *essentially*, and that this assumption determines how his investigation of cognition proceeds.

Similarly, an AA theorist may take biological considerations into account when formulating an anthropogenic theory. For example, AA theorists may attempt to construct a general biologically informed account of how representational content is fixed by looking to candidate representational mechanisms in simple systems (Cf. Dretske 1981). A theorist can do this whilst still holding that propositional attitudes, like beliefs, in humans are the gold standard for *bona fide* cognition. More generally, although AA proponents may investigate cognition in simple systems (e.g., nematode brains), such investigation will ultimately be led by a conception of what cognition is that premised upon human capacities. For this reason, BA's proponents argue that the biogenic perspective allows for a more complete picture of cognition and how various cognitive phenomena are related despite being selected for by different evolutionary pressures. Lyon writes:

It is important to stress that the two approaches are not mutually exclusive, indeed, both are necessary for a complete picture. We must understand what cognition is and what it does as a natural phenomenon, but we also have to understand human psychological experience (2006, p. 26).

Here, Lyon makes it explicit that the two approaches to investigating cognition are linked in the sense that investigating what cognition does for biological systems will ultimately include investigating how it does what it does in the human case. That being said, there is a tension between AA's demarcation criterion and BA's biological starting point. For BA, many genuine instances of cognition may take a very different form than that of human cognition. It is due to the AA theorist's more restrictive conception of cognition that she may attempt to dismiss the broader set of non-human-like capacities recognized as cognitive by BA theorists as something else. As will be seen in the next section, it is just this kind of charge that Fred Adams foists against BA.

The argument from intentional content

Adams (2018), on behalf of AA, has recently raised an objection against BA. Appealing to the interests and practices of contemporary cognitive science, he argues that cognitive scientists use the term "cognition" to refer to something very different than what BA theorists use the term to refer to. Hence, thinks Adams, attending to cognitive scientific practice shows that we should prefer AA to BA. In line with a venerable trend in philosophy of cognitive science (Fodor 1990; Dretske 1981), Adams's general anthropogenic classification methodology distinguishes cognition from non-cognition by holding that the former is a kind of processing that necessarily involves internal representational states with intentionality. Intentionality, the property of "aboutness" that mental states have to be directed at the world was originally proposed by Brentano (1874/1995) as the mark of the mental. In other words, "all and only mental phenomena exhibit intentionality" (Schlicht 2018, p. 8) (see also Crane 2016). These intentional states, Adams contends, take the format of

representations with conceptual-level content, the meaning of which is determined by states of affairs in the world that they are about. Adams writes:

[...] cognition requires specific kinds of representational formats. And these formats are built on top of information processes. The symbols that are cognitive are built from information exchanged between system and environment. But the format of these representations is *semantic* at a level above that of the information itself (Adams 2018, p. 24; my emphasis).

It is only when intentional states with conceptual-level content operate on sensory representations (i.e., sensory information states) that the resulting behaviour can truly be said to be cognitively driven.⁷ Adams's clarifies this format distinction between meaning-bearing representations with conceptual-level content and sensory representations when, referring to his own work and that of Dretske (1981, 1998) he writes:

"But conceptual representations are different kinds of representations than sensory ones. The change from information to meaning (where the content of a concept is its meaning) involves changes in the format of representation" (Adams 2018, 28).

And again

"I am suggesting that cognition is this kind of information processing which alters the representational format to a different level—to the level of meaning and not just information" (Adams 2018, 28).

To get an idea of the kind of representational format that he argues underwrites genuine cognition, Adams offers an example. If a person senses that the temperature of a room is increasing, whether or not she is aware of it, she also senses the increasing mean molecular kinetic energy (mmke) of the air surrounding her. However, if she thinks that the room's temperature is increasing, she needn't think that the mmke of the surrounding air is also increasing. This example demonstrates that although thoughts, beliefs, desires, and other representations with intentional content may be about the same referent (in this case a target event), thinking one kind of thought (distinguished by its meaning) about a referent does not entail thinking all corollary thoughts about the same referent.⁸ Although sensory information processing may involve representations, such representations fail to have the right format, on Adams' construal, and thus they never to rise to the level of meaning.⁹

⁷ The mere deployment of representations with intentional content without resulting behaviour (e.g., thinking, reasoning, etc.) or in the absence of sensory information states would also qualify as cognitive according to Adams.

⁸ This is an instance of Frege's problem (i.e., referential opacity).

⁹ The distinction between what Adams calls sensory representations and representations with intentionality may be interpreted as being analogous to Grice's (1957) distinction between "natural meaning" and "non-natural meaning". Natural meaning is indication-based. For example, smoke indicates fire, but smoke is not *about* fire in the sense that one cannot intelligibly speak about the presence of smoke in the absence of fire as something that is 'false'. Non-natural meaning, on the other hand, rises to the level of semantics. For example, that the word "water" means "tasteless, odourless, colourless liquid" makes it the case that the proposition "water is black and tastes like tar" is false.

In contrast to the deployment of intentional states with conceptual-level content, Adams claims that what BA calls ‘cognition’ is mere information processing of sensory states which results in hardwired, adaptive behaviour. Although cognitive and non-cognitive processes both involve information processing and adaptive behaviour, the kind of adaptive behaviour that is driven by mere sensory information (in a sensory format) alone does not rise to the status of being cognitive behaviour. It is because what biogenic theorists refer to as cognition is actually “information-driven behaviour that is adjusted in response to variable environmental conditions” (Adams 2018, p. 27), that biogenic plant scientists and biologists erroneously categorize the adaptive behaviour of organisms such as plants and bacteria as cognition-driven processes. When biogenic theorists claim that plants and bacteria exhibit cognitive capacities what they are really doing is changing the subject altogether from cognition to something else. For in such cases “there are no internal states that mean (in the sense of having truth values) things outside the system” (Adams 2018, p. 29). Adams concludes that because intentionality is a distinctive characteristic of the cognitive and this is something which BA by ascribing cognition to plants and bacteria fails to respect, biogenic researchers are in the practice of investigating something very different than what cognitive science recognizes as cognition.

I will now respond to this argument by providing another classification methodology for distinguishing the cognitive from the non-cognitive, one which is compatible with the range of cognitive behaviours that BA recognizes.

Responding to the argument from intentional content

It may be assumed that Adams’s primary reason for locating cognition at the level of intentional states with conceptual content is based upon his view of what cognitive science recognizes as cognition.¹⁰ If biogenic theorists when ascribing cognition to bacteria and plants are using “cognition” to mean something other than processing that reaches the level of intentional states with conceptual-level content “then no one in current cognitive science would be alarmed” (Adams 2018, p. 28). However, if such theorists are not ascribing cognitive capacities to plants and bacteria instrumentally and actually “intend the cognitive ascriptions to be true, then it is not harmless” (Adams 2018, p. 29).¹¹

¹⁰ To be clear, Adams is in agreement with Lyon, who when setting out what she identifies as the core anthropogenic principles (i.e., the core principles of contemporary cognitive science), claims that intentionality is a distinctive feature of cognition according to AA. BA principles as put forth by Lyon (2006) do not include intentionality. Interestingly, if the continuum that I propose below is correct then Lyon has overlooked intentionality as a necessary feature of cognition that is common to BA and AA.

¹¹ Although BA theorists are not ascribing cognition instrumentally, one might question whether the plant scientists and microbiologists whose research biogenic theorists help themselves to are engaging in instrumental talk. For a sustained argument as to why there is no non-question-begging reason to dismiss the use of psychological predicates as merely instrumental or metaphorical when biologists apply them to basal organisms see Figdor (2017, 2018).

I will now argue that proponents of BA can both agree with Adams that intentionality is a distinct feature of the cognitive and yet consistently reject his claim that intentionality must be concept-based in order for processing to qualify as cognition. More specifically, using a classification methodology that is compatible with what cognitive scientists actually do, I shall propose a view of intentionality *simpliciter* that conceives of it as a continuum, ranging from weak to strong intentionality. Because cognitive science investigates and recognizes cognitive phenomena all along the continuum and not just at the level of strong intentionality where Adams restrictively locates cognition, his argument against BA turns out to be ill-motivated. In locating cognition exclusively at the level of strong intentionality, Adams disregards the range of research that cognitive scientists have engaged in over the course of the last 40 years. I begin with an initial response that considers some of the current practices of cognitive science. This response challenges the motivation underwriting Adams' demarcation criterion. I shall then introduce the intentionality continuum thesis as a primary response to Adams' argument from intentionality and more generally as a ground for rejecting the kind of cognitive chauvinism that Adams' argument is representative of.

The practices of current cognitive science

Does cognitive science recognize and investigate only those phenomena that are underwritten by intentional states in the form of beliefs and desires with concepts as their constituents? While there is plenty of cognitive science that is aimed at investigating phenomena at the level of Adams' strong intentionality, it appears that there is plenty that is not. A wide range of cognitive scientific research programmes such as action-blindsight research (Kentridge et al. 1999), two visual streams research (Milner and Goodale 1992), 4E cognition (i.e., enactive, embedded, embodied and situated cognition) (Clark 1997/2008; Chemero 2009; Beer 2000; Thompson 2007; O'Regan 1992; Lakoff and Johnson 1980; Kirsh 2010), action-oriented cognition (Jeannerod 2006; Cisek and Kalaska 2010; Engel et al. 2013; Pezzulo 2011; Tversky 2019), and dynamic systems accounts (Thelen and Smith 1994; Beer 2000; Kelso 2016; Newell and Liu 2012)—to name a few—investigate and explain psychological processes in ways that do not essentially implicate states with conceptual-level contents. This is not to deny that much of cognitive science is focused upon phenomena at the level of strong intentionality. What I am arguing however is that cognitive science does not *exclusively* investigate phenomena at the level of strong intentionality. The existence of each of these many programmes demonstrates that cognitive science recognizes and investigates processes that do not exclusively involve conceptual representations. To take an example, the core of embodied cognition is based upon the claim that a complete understanding of cognition requires understanding the various constitutive roles that the sensorimotor systems play in not only conceptualization (see Lakoff and Johnson 1980) but also how we respond to environmental changes and solve problems in the here and now (Clark 2008). A classic example of this is the ecological solution to the baseball outfielder problem in which by moving so as to keep a flyball's position steady in one's field of view—staying coupled

to it via sensorimotor engagement—one is able to adjust one’s movement’s in light of the ball’s perceptible trajectory and eventually catch it (McBeath et al. 1995).¹² The kind of conceptual-level representation that Adams argues is definitive of intentionality is rendered superfluous by this simple solution to the problem of how to catch a fly ball.

More generally, the kind of cognitive science which has emerged in the last 40 years has recognized a problem space that calls for the use of adaptive strategies that are not limited to detached planning, abstract thought, and conceptualization. The agreed problem space of cognitive science has been widened to include problems that arise on faster timescales. Attempting to address these problems has motivated many researchers to abandon the kind of traditional “sandwich” model of cognition that Adams uses (i.e., where cognition is conceived of as a process occurring sandwiched between perceiving and acting).¹³ Many of the problems in this wider problem space may be efficiently solved by eliciting frugal and fast solutions of environmentally coupled perception and action loops (Clark 2008). This is not to say (radically) that all solutions required within this widened problem space involve the coupling of an organism to the environment via action and perception, but rather that there are various kinds of cognitive strategies that may be deployed given the nature of the problems which organisms are faced with.

One counterargument that Adams might potentially raise at this point is that cognitive science is merely a misnomer and hence what many cognitive scientists study actually falls outside of the domain of cognition proper. If this is the case, then identifying what cognitive scientists actually do is no argument against Adams’ demarcation criterion. This kind of objection can be resisted, however, for the following reason: assuming for the sake of argument that cognitive science studies both cognition and pre and post-cognitive processing, then Adams’ motivation for claiming that cognition is restricted to intentional states with conceptual-level content cannot be that of honouring the practices of cognitive science; his demarcation criterion for cognition returns a much narrower range of capacities than those studied by cognitive scientists. If his demarcation criterion is motivated by a subset of cognitive scientific practices, then he must supply a way to delineate that subset independently of his own criterion. He cannot say, circularly, that his criterion for cognition is picked out by a narrow range of capacities that cognitive science studies and that the narrow range of *bona fide* cognitive capacities that cognitive science studies is picked out by his demarcation criterion. Hence, even if cognitive science were a misnomer, then the practices of cognitive science alone cannot motivate Adam’s demarcation criterion.¹⁴

To sum up: Adams’ argument from intentionality would like to conclude that BA is untenable and is harmful to cognitive science because it conflates actual cognition with mere sensory information processing. However, if one takes into account

¹² For different theoretical perspectives regarding this process see McLeod et al. (2006) on optical acceleration cancelation and see McBeath et al. (1995) on linear optical trajectory.

¹³ Susan Hurley (1998) is responsible for coining the term “sandwich model”.

¹⁴ I am grateful to an anonymous reviewer for raising this objection.

the range of research programmes that have been developing over the last 40 years (i.e., what cognitive scientists have actually been investigating and how the field is developing), his claim about what “cognition” means to cognitive science is simply not true (anymore). Adams hence fails to provide a valid reason to reject BA as a *cognitive* explanatory framework.

BA however is not off the hook. Even if Adams’ demarcation criterion fails to be motivated in the manner that he thinks it is, the onus is nonetheless placed upon the BA theorist to make explicit the shared core aspect of cognition which links BA and AA explanations. For without such a common core, Adams’ contention that BA and AA are investigating different processes remains a real threat to BA. Far from serving its intended purpose, I would like to suggest that the argument from intentionality provides a valuable opportunity to demonstrate just how BA and AA are linked. Disagreeing with Adams that the mark of the cognitive must take the specific form of intentionality with conceptual-level content does not entail rejecting the idea that intentionality *simpliciter* is a distinctive feature of cognition. In the next subsection I will develop the intentionality continuum thesis and argue that when taking it into account, demarcating the cognitive from non-cognitive by way of an intentionality criterion fails to be a consideration against BA. If the intentionality continuum thesis is correct, BA and AA merely lie upon two poles on the continuum, separated by the degree of intentionality involved in the kinds of cognitive phenomena they investigate.

The intentionality continuum thesis

One the basic BA assumptions is that cognition is a set of evolved traits that has been selected for in virtue of its functional contribution towards organismic adaptivity. Cognition as such is seen as graded, moving from less to more complex cognitive capacities. If intentionality is the mark of the cognitive, and assuming that intentionality also underwrites the differences in degrees of cognitive complexity, then one might reason that intentionality is also a graded feature. The intentionality continuum thesis provides one plausible way of understanding the graded nature of intentionality. If correct, it offers a biologically plausible manner of theorizing about the possible evolution of intentionality and the evolution of the various cognitive capacities that have arisen in tandem with the complexification of representational content.

The intentionality continuum thesis proposes that intentionality varies along an axis of representational complexity, moving from less complex to more complex representational content and hence from simpler to more sophisticated forms of cognition.¹⁵ *Future-oriented-ness* will be used in what follows as a general measure of representational complexity; it describes how far into the future content allows an organism to behaviourally adapt to yet-to-be-encountered environments/sensory

¹⁵ The intentionality continuum thesis assumes that representation is an essential aspect of intentionality; “aboutness” or “directedness” is necessarily a matter of representing some target state of affairs. This is consistent with what has called “the entailment thesis” (see Schlicht 2018).

states. Future-oriented content that operates over the timescales of minutes is more representationally complex than content that operates over seconds. Moreover, more complex representational content (i.e., strong intentional content) that operates over longer/slower timescales comes about as the result of the combined activity of less complex representational content (i.e., weak intentional content) at shorter/faster timescales.

Understanding representational complexity in terms of content's future-orientedness finds support in various programmes including associative learning research (Ginsburg and Jablonka 2019); active inference theories of action planning (Pezzulo et al. 2015; Corcoran et al. 2020), and computational models of cognitive boundaries in theoretical biology (Levin 2019) to name a few. The idea common to all of these research programmes is that organisms embody internal models of their environments and their behavioural consequences upon the environment (Cf. Bechtel and Bich 2021). Evolutionarily, as organisms became able to incorporate increasingly temporally distal consequences of both their current behaviour and environmental causes into their behavioural regime, their internal models became more complex. Model complexity is scored by a model's *temporal depth* (Pezzulo et al. 2015; Friston et al 2017; Ginsburg and Jablonka 2019; Levin 2019; Calvo, et al. 2020). An increase in temporal depth requires that an internal model must have more interacting levels of control mechanisms. Because faster interactions tend to occur between spatially adjacent levels, organisms with shallow models are limited to faster timescale interactions, the kinds that are involved in tightly coupled organism-environment causal traffic. As such, the preferences of less complex organisms are *largely* restricted to what is occurring at the present moment in their immediate environment (Levin 2019). Under the assumption that models serve anticipatory behaviour (see below), the mechanisms of shallow models that anticipate sensory/environmental input needn't anticipate far into the future, and hence, they have little temporal depth. The interaction between distant mechanisms at different levels in a deep model (i.e., with increased levels of organization) not only means that there is a delay across increasingly distant levels, but also that in order to anticipate the activity at the spatially distant sensorium, higher levels must be able to anticipate further into the future. Thus, it is the increased hierarchical organization of temporally deep models that allows for the increased decoupling of such models from the direct influence of the environment and immediate behaviour.¹⁶ Since there are plausible reasons to interpret internal models in representational terms (see Rosen 1985/2012; Sims and Pezzulo 2021), increased representational complexity can be understood as increased model complexity (i.e., increased temporal depth). It is in the notion of temporal depth that the measure of future-orientedness is grounded.

Although there have been previous philosophical accounts that have attempted to understand representational complexity as existing along a continuum (see Millikan 1989, 2004; Sterelny 2003; Schulz 2018), the intentionality continuum thesis

¹⁶ Strictly speaking, in the general context of biological systems this levelled organization may be more accurately understood in terms of heterarchical arrangement as opposed to hierarchical arrangement (see Bechtel and Bich 2021).

breaks with these proposals in two crucial ways. Firstly, the proposed continuum thesis takes a biogenic approach to investigating intentionality; it starts from the biological facts and works its way up to the human case. Importantly, the motivation for working up to the human case is not to stress differences between human cognition and non-human cognition, as is typical of AA, but rather to expose a possible commonality (weak intentionality) underwriting all forms of cognition. The notion of meaning and goal-directedness that weak intentionality is grounded in will be approached using the conceptual tools of control theory (Wiener 1961; Conant and Ashby 1970). Accordingly, organisms will be cast as self-organizing dynamic systems, the adaptive behaviour of which is finessed by sensory feedback and driven by control mechanisms that maintain setpoints (i.e., homeostatic norms) at different nested levels of biological organization. These setpoints must be kept within a limited range in order for a system to remain alive. A system's integrated dynamics are thus directed towards the maintenance of homeostatic norms and, importantly, what biological representations mean to a given system is relative to that system's homeostatic norms.

Secondly, the intentionality continuum thesis differs from other continuum approaches in taking anticipation to be an essential property of cognitive systems. For example, Millikan (2004) in providing an account of the complexification of representations, starts from what she calls “pushmi-pullyu” representations. These are primitive representations that have undifferentiated indicative and imperative content. According to Millikan, a splash of a beaver's tail may indicate something like “there is currently a predator flying above”. The same tail splash may simultaneously be a sign for other beavers to “initiate escape behaviour now!”. The distinction between the kind of primitive content (i.e., weak intentional content) that the intentionality continuum locates at one pole of the continuum and pushmi-pullyu content is this: Whereas pushmi-pullyu representations describe how the world *is now*, states with weak intentional content describe how the world *will be* (given how the world is now). Moreover, whereas pushmi-pullyu representations specify what an organism should do *in response* to how the world is, weak intentional content specifies what an organism should do *in preparation* for how the world will be. If cognition is fundamentally anticipatory, then this difference seems to be an important one. For any account of primitive intentionality must be able to account for the nature of anticipatory behaviour. This is something that the kind of weak intentionality that is central to the intentionality continuum is in the position to do.

One reason to interpret basic biological representations as describing how the world will be may be arrived at by asking the following question: In what circumstances would representations be needed by the simplest biological systems? Thus, reasoning from biological need to future-oriented representational content. On one understanding of representations, they function as stand-ins for absent targets (van Gelder 1995; Pezzulo 2008; Sims and Pezzulo 2021). Accordingly, accounting for simple reflex-like responses do not require representational explanation because the mechanisms with which reflexes are produced are purely reactive and sustained by external stimuli; reflex profiles are caused by specific aspects of detected stimuli (i.e., stimulus kind, strength, and/or duration) in the presence of that stimuli. However, under the assumption that being limited to purely reactive compensation—like

reflexes—would be detrimental to the ability of an organism to adapt to the selective pressures in its environment (Rosen 1985/2012), it is reasonable to think that an organism's ability to successfully adapt requires a means for it to extrapolate from present stimuli to yet-to-be encountered states of affairs and guide its behaviour accordingly.¹⁷ We thus arrive at a reason to interpret weak intentional content as playing the role of standing in for yet-to-be encountered sensory/ environmental states via the route of acknowledging a basic biological need for proactive compensatory behaviour.

The assumption that cognition is fundamentally anticipatory finds support with Lyon (2006), who in describing one of the biogenic principles of cognition writes: “cognition is intrinsically future-oriented (what happens next?) and thus predictive” (p. 23). The idea that cognition is fundamentally anticipatory also follows a long tradition in psychology and cognitive science (Bartlett 1932; Craik 1943; Piaget 1970; Neisser 1976; Drescher 1991; Arbib 1992; Riegler 2001; Grush 2004; Castelfranchi 2005; Bar 2007). Despite the fact that cognitive science has primarily focused on the anticipatory nature of human cognition, it is important to stress that anticipatory behaviour is not unique to neuronal organisms. For instance, Tagkopoulos et al. (2008) have shown that there is anticipatory aerobic respiration repression in bacteria *E. coli*. that occurs prior to oxygen level depletion. Saigusa et al. (2008) have provided striking evidence that slime moulds (*P. polycephalum*) anticipate periodic change in dryness conditions. Research by Rodaki et al. (2009) have shown that fungi too respond to yet-to-be-encountered stress in an anticipatory manner. Novoplansky et al., (1990) have provided evidence that the common purslane (*P. oleracea*) anticipatorily grows away from other plants that will compete with it for light in the future. It is these kinds of anticipatory behaviours that are overlooked and left unaccounted for by previous accounts of primitive representation and thus warrant the place of the intentionality continuum thesis. Since, according to this thesis, all anticipatory behaviour is driven by weak intentional content (see below), future-oriented-ness is both fundamental to understanding intentionality and how this mark of the cognitive exists along a continuum of increasing complexity.

To be sure, because there are additional measures that are needed for a complete characterization of increasing representational complexity (e.g., format, compositionality, recursion, etc.) future-oriented-ness will serve us here as one kind of heuristic for an initial investigation into intentional strength. I will assume that such additional complexity metrics are largely independent dimensions of a

¹⁷ Imagine, for instance, of the difference between two hypothetical bird populations, *Fasts* and *Slows*, that are identical other than the following phenotypical difference: *Fasts* can behave according to acquired colour cues. They have learnt that every time Mrs. Miller's bright red coat is present, shortly thereafter they will receive breadcrumbs. As such *Fasts* have a tendency to swoop down and flock around Mrs. Miller upon her arrival. *Slows* on the other hand, do not have this capacity to learn colour cues and, as such, they merely respond to *Fasts* eating bread by flocking around *Fasts*. As such, by the end of Mrs. Miller's daily feeding routine, *Slows* have secured—if anything at all!—the crumbs that *Fasts* have left behind. Now, imagine that the only food source for both *Fasts* and *Slows* regularly came from Mrs. Miller. *Slows*' reactive behaviour would place them at a significant fitness disadvantage relative to *Fasts*, who due to their ability to use cues to anticipatorily behave would most likely outcompete the population of *Slows*.

multidimensional space. For simplicity of exposition and owing to space limitations, any additional measures and how they interact with future-oriented-ness will not be included in the following treatment of the continuum thesis.

With these preliminary remarks to hand, let us now move on to the details of the intentionality continuum.

Weak intentionality

On one end of the continuum is what I call *weak intentionality*; I shall define this as the aboutness of internal dynamics that is directed at yet-to-be encountered environmental state targets and that causally underpins adaptive behaviour. Such dynamics are meaningful insofar as they bring about environmental conditions that are relevant to an organism's long-term homeostatic maintenance by proactively guiding behaviour. Weak intentionality is phenotype-relative in the sense that not every environmental condition has the same homeostatic relevance to every kind of organism (e.g., being an air environment is important to a human's continued homeostasis but not to a fish's). Much like what autopoietic theorists have called "basic intentionality" (Thompson 2007; Schlicht 2018), weak intentionality is a feature of the entire organism and its behaviour rather than a feature of mental states or underlying mechanisms. It is a property of an organism's integrated internal dynamics that is causatively manifested in sensorimotor and biochemical behaviour.

Not just any adaptive behavioural response counts as weak-intentionality-driven. In order to qualify as weak-intentionality-driven, behaviour must be directed at having *higher-order homeostatic influence*; an organism behaviourally invests its current short-term homeostatic stability (i.e., metabolic resources) for the long-term homeostatic stability brought about as a result of that behavioural investment. In some organisms, this will manifest as overt behaviour, exerting a direct influence upon its surrounds (e.g., a plant's spending metabolic resources now in directing its growth away from the location that light competing plants are likely to eventually occupy (Novoplansky et al. 1999)). In other organisms, this higher-order influence may take the form of metabolically expensive investments in gene transcription in preparation for yet-to-be-encountered stress conditions (see example below).

To be sure, although behaviour that is driven by weak intentionality allows homeostatic maintenance, it is not reducible to homeostatic reflex response. The distinction between homeostatic control and higher-order homeostatic influence may be cast as the difference between a purely reactive form of self-regulation (see Cannon 1929) and *allostasis* or proactive "stability through change" (Sterling and Eyer 1988).¹⁸ Whereas reflexes involve direct responses to stimuli in virtue of sensory deviations from genetically encoded homeostatic setpoints (Pezzulo et al. 2015; Stephan et al. 2016), behaviour driven by weak intentionality is anticipatory. Its anticipatory nature implies that such behaviour has a degree of independence from the

¹⁸ Although allostatic theory was originally applied to describe the behaviour of neuronal organism, it has more recently been expanded to include the behaviour of non-neuronal organism (Alfieri 2008; Corcoran et al. 2020; Kiverstein and Sims 2021).

influences of the environment; it is freed from being fully determined by the current states of the encountered environment. As a result, the same environmental stimuli may not always result in the same behaviour. This is significant for two reasons: firstly, it suggests that weak intentionality driven behaviour exhibits more flexibility than mere reflex-like response¹⁹; secondly, it opens up a space for a kind of primitive agency; rather than being “pushed around” by sensory stimuli, an agent’s behaviour is normatively constrained by maintaining certain long-term homeostatic conditions that its persistence depends upon.²⁰

In order to exhibit weak intentionality a living system’s inner dynamics requires a minimal degree of future-oriented-ness. We may roughly understand future-oriented-ness to operate according to the following kind of rule:

Given an environmental state type A that is regularly followed by environmental state type B and that this A-B correlation has been acquired (over phylogenetically or ontogenetically timescales) by an organism O; anytime O detects an instance of state A, its weak intentional content will be directed at an instance of state B as a function of O’s homeostasis. This will be the case despite the fact that B is yet-to-be-encountered.

The temporal distance between A’s occurrence and B’s occurrence will typically determine the measure of future-oriented-ness. If B regularly occurs in a matter of

¹⁹ Following Schulz (2018) I would argue that even reflexes exhibit some degree of flexibility. Whether or not a reflex response occurs not only depends upon the detection of an appropriate stimulus but also often depends on an organism’s current physiological conditions. For example, if an organism is fatigued, it may not respond to a stimulus that it tends to respond to stereotypically.

²⁰ Walsh (2015) in arguing for a continuum of “ecological agency” similarly recognizes a kind of agency that even simple systems such as bacteria may be endowed with. Walsh analyses ecological agency in terms of three related notions of “goal-directedness”, “affordances”, and “behavioural repertoires”. The idea, roughly, is that being a goal-directed system is a requirement for experiencing environmental conditions as opportunities for action (i.e., “affordances”) and experiencing conditions as affordances requires having a behavioural repertoire. Behavioural repertoires come in degrees; some biological systems having restricted repertoires and others having a larger range of behaviours in which they may engage. As such, agency also comes in degrees. For example, according to Walsh bacteria qualify as having a low degree of agency as a result of their restricted behavioural repertoires. Even though bacteria are ecological agents, Walsh—engaging in a kind of anthropogenic cognitive chauvinism reminiscent of Adams’ own—argues that bacteria’s behaviour falls short of being driven by full-blow cognition because they fail to entertain propositional attitudes or engage in reasoning. Thus, agency is prior to cognition according to him. Weak intentionality, like Walsh’s notion of ecological agency, emphasizes behaviour that facilitates an organism’s remaining near homeostatic equilibrium. Unlike Walsh, however, I specify *cognitively driven behaviour* in terms of *anticipatory* biochemical and sensorimotor behaviour; behaviour that is biased in virtue of future-oriented representational content. This is to say that Walsh’s account of agency does not—as far as I can tell—distinguish between purely reactive behaviour that is driven by homeostatic goals and proactive higher-order homeostatic control. Moreover, weak intentionality, as I am arguing here is the mark of the cognitive. That a primitive form of agency may arise in tandem with weak intentionality is a contingent feature of the continuum thesis; the correctness of the intentionality continuum thesis does not hinge upon the correctness of the claim that a form of agency arises with weak intentionality. More generally, whilst Walsh argues for a continuum of ecological *agency*, I am arguing for a continuum of *intentionality* and hence there is a clear distinction between Walsh’s *explanandum* and mine. While it is evident that he takes a biologically-informed approach to agency, he does not take a biogenic approach to cognition.

seconds after A, then future-oriented content will operate at the timescale of seconds. Similarly, if B happens minutes after A, then future-oriented content will operate at the timescale of minutes. Because weak intentional content drives anticipatory behaviour, the degree of future-oriented-ness must allow behaviour to occur prior to B and not subsequent to B, the latter being a case of mere reflex-like response.

This correlational rule is implicit in the internal dynamics (i.e., biophysical patterns of activity) of a living system and it explains what it is about weak intentional content that elicits anticipatory behaviour. Such rules are adaptively acquired over phylogenetic timescales in virtue of the fact that organisms which are or become more responsive to the correlations have more selective fitness than those that do not. This results in correlation-sensitive populations. In other words, organisms become models of the correlational regularities within their regulated environments (Cf. Conant and Ashby 1970). Over ontogenetic timescales, classical conditioning may be one example of how such rules are learnt. Despite the differences in timescales at which such rules are learnt, they always bias behaviour that is relative to an organism's homeostatic norms. For example, if environmental state B is beneficial to organism's long-term homeostasis, then that organism's behaviour will be biased so as to prepare it for encountering and interacting with B. If state B is incompatible with an organism's long-term homeostasis, then that organism's behaviour will be biased so as to avoid an encounter with state B. It is in this manner that future-oriented content via anticipatory behaviour mandates the temporary alteration of homeostatic equilibrium *now* in order to avoid irrecoverable deviation from long-term homeostatic equilibrium *later*. See Fig. 1.

Before looking at an example of a system that exhibits future-oriented weak intentional dynamics, let us look at a few cases of systems that do not. Doing so will be telling of the fact that future-oriented-ness is not itself an overly liberal notion. A moon, when orbiting a planet decelerates as it approaches its apoapsis and accelerates afterwards. Why shouldn't one describe this deceleration as a moon's anticipating its approach to apoapsis and so exhibiting future-oriented content?²¹ Firstly, because behaviour driven by future-oriented-ness is grounded in *homeostatic norms*, and because a moon is not a biotic system, it has no need for future-oriented content; abiotic entities do not have homeostatic setpoints to defend. A moon's deceleration can be explained simply in terms of Newtonian mechanistic causes (i.e., in terms of gravitational force, mass, and distance between two bodies) given initial conditions. Secondly, there is no learning (or adaptation) involved in the moon case, and hence, no correlational rules that a moon's slowing down reflects. To put it differently, a moon's deceleration is not the result of a learnt model of its own homeostatic relevant behaviour.

What then about a simple embodied AI agent (e.g., a robot cell) that *can learn* environmental correlations and use them to anticipate and adaptively behave with respect to temporally distant targets? Its behaviour too, given the account above, would fail to exhibit future-oriented weak intentionality. The reason, again, being

²¹ Many thanks to an anonymous reviewer for helpfully raising this particular objection and for pushing me to respond to it.

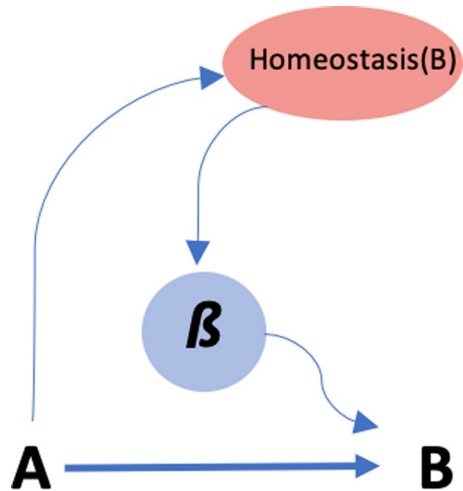
that such an agent's behaviour is not relative to homeostatic norms. Weak intentionality cannot be reduced to mere learning-driven anticipatory behaviour. There is of course the possibility that an embodied AI might be programmed to behaviourally defend a particular setpoint range that is relevant to the continued functioning of one or another of its parts (e.g., the cell might be programmed to swim away from areas of its environment that are too hot or too cold given its effectors only function in a particular temperature range). Whatever the programmed rules by which some of the AI's parts continued functioning occurs are, they are distinct from the kind of homeostatic norms that future-oriented dynamics answer to. Not only are homeostatic norms intrinsic to biological systems, but they are *systemwide*; every component part of a biological system is subject to metabolic dysfunction given that all such parts are the product of metabolisms. Moreover, the functional status of one part is often connected to the continued functioning of another. Disrupt the flagellar motor system of an *E. Coli*, and one also disrupts its ability to pursue local nutrients and engage in the metabolic activities that the continued functioning of each of its parts are dependent upon. Since long-term dyshomeostasis brings about irreversible metabolic disfunction, it is not only the continued functioning of any one part (or selection of parts) but the continued functioning of the entire organism that is at stake when long-term homeostasis is threatened. In contrast to this, there is often (at least currently) a high degree of functional modularity when it comes to embodied AI (Bongard and Levin 2021); the functional disruption of one part of an AI (e.g., its effector system) needn't affect or be affected by the functional stability of any other of its parts.

It is the interdependence of systematic *metabolic function* and *homeostasis* that allows systems exhibiting weak intentionality to temporarily leverage their current homeostatic stability to behave in ways that bring about anticipated effects across multiple connected homeostatic norms. As such, future-oriented-ness answers to the long-term stability of homeostatic norms across all integrated level of a system's nested organization. Without systemwide homeostatic norms, an AI's "concern" with a particular functionally relevant setpoint (e.g., its local environment's temperature) is only superficial; it fails to be a goal or a norm for the entire system.²² This is evident by the fact that the AI's continuing to swim so that it maintains a particular temperature is compatible with its "starving to death" (i.e., its failing to seek out sources of energy (e.g., chemical, electric, solar, etc.) that would support its continued its activity).

Now, to take an example of a system that does exhibit weak intentional dynamics: after detecting decreasing levels of its preferred nutrient (glucose) in its environmental medium, some strains of wild yeast, *Saccharomyces cerevisiae*, begin the metabolically expensive process of galactose utilization pathway induction, allowing them to catabolize galactose in preparation for *eventual* glucose depletion (Wang

²² I suspect that for an embodied AI to have systemwide homeostatic norms and hence to exhibit weak intentional dynamics, such agents may very well need to implement something akin to metabolism. Supporting this claim, however, must be left as a project of future research.

Fig. 1 Future-oriented weak intentional dynamics. The correlational rule (environmental state *A* will be followed by environmental state *B*) is internalized by an organism such that when *A* is sensed, future-oriented dynamics (red oval) are elicited and directed at *B* as a function of the organism's homeostasis. Anticipatory behaviour, β , despite its metabolic cost (or other short-term homeostasis destabilizing effects), is initiated prior to the coming about of *B* to ensure long-term homeostasis



et al. 2015).²³ Importantly, *S. cerevisiae* prepare for galactose consumption *while still in the presence of glucose* and preparation can occur up to a striking 3.8 h before total glucose depletion! One manner of understanding this phenomenon is that the internal dynamics that facilitate galactose utilization pathway induction are modified in part by the yeast's sensitivity to a tendency in the global structure of its environmental medium. Decreases in glucose concentration are regularly correlated with, and hence specify, the eventual exhaustion of glucose. This modification biases the system to behave anticipatorily. *S. cerevisiae* acts not solely with respect to its short-term homeostatic maintenance and the sensory/environmental states that it is currently detecting, but with respect to the yet-to-be-encountered absence of glucose. The induction of the galactose utilization pathway proactively counters the long-term homeostatic instability that would occur systematically across all levels of *S. cerevisiae*'s organization as a result of local glucose depletion. The activity of *S. cerevisiae*'s gene regulatory network exhibits future-oriented-ness that manifests itself in higher-order homeostatic influence.

That weak intentional content is future-oriented and allows for anticipatory behaviour at all is contingent on the fact that content exhibits some degree of past-oriented-ness. Again, content can be future-oriented only insofar as it is the product of acquired environmental correlations between a cue and a cued state. Having past-oriented character thus implies the presence of a simple form of memory that is deployed to guide future behaviour.²⁴ One might at this point balk at the claim

²³ Wang et al. (2015), write: "We now have shown that low or decreasing levels of a preferred nutrient can serve as a predictive cue for eventual depletion. Since this is inevitable when cells deplete a mixture of nutrients at unequal rates, and mixed-nutrient environments are ubiquitous in nature, environmental anticipation may be a more widespread regulatory strategy than previously recognized" (p. 16).

²⁴ Although simple memory states may not be stored over longer periods of time (this too is a matter of degree) the presence of simple forms of memory implies—pace Millikan (1989)—that even simple organisms have some capacity to store representations.

that simple organisms such as yeast have a form of memory. However, it has been well established that even some bacteria such as *E. coli* exhibit a simple form memory (Macnab and Koshland 1972; Lambert and Kussell 2014).²⁵ Furthermore, the idea of immune memory in invertebrates has gained recent support (see Milutinovic and Kurtz 2016). Unlike the memory states that arise as one approaches more complex representational states at the strong pole of the continuum, the sole function of memory at the weak end of the intentionality continuum is to guide adaptive behaviour in the here and now. As such, weak intentional content exhibiting low degrees of future-oriented-ness cannot be decoupled from behaviour.²⁶

Accuracy conditions at the weak intentionality pole of the continuum coincide with the contribution that content makes to an organism successfully adapting to its environment. Inaccuracy for *S. cerevisiae* may be understood as a case in which the induction of the galactose pathway occurs due to a small decrease in the sensed glucose concentration despite the overall tendency of the environment to be glucose abundant. *S. cerevisiae*'s weak intentional content in such a case diverges from the actual environmental state changes that it will encounter and, more importantly, such content results in a maladaptive behavioural investment. Accuracy and inaccuracy are respectively manifested in an organism's continued homeostasis or dys-homeostasis and as such can be the basis for acquiring more complex rules of correlation or "unlearning" those rules that fail to effectively serve homeostatic norms.

With this picture of weak intentionality to hand, I would like to argue that weak intentionality is implicated in BA. As we've seen above, Lyon takes anticipation to be fundamental to cognition. Recall also that BA conceptualizes cognition functionally, in terms of the adaptive behavioural capacities underwriting a system's homeostatic maintenance. As such, cognition according to BA is both fundamentally anticipatory and geared towards controlling adaptive behaviour. It should be clear by this point how weak intentionality is implied in BA: weak intentionality describes the future-oriented internal dynamics which underwrite the kind of cognitively driven behaviour that is central to BA.

From weak intentionality to strong intentionality

Located on the opposite end of the continuum is the kind of intentionality that Adams identifies as a distinctive feature of cognition. Strong intentionality may be characterized as the *aboutness of mental states* rather than the aboutness of behaviour-driving internal dynamics of the organism. It is at this level of intentionality that representations with conceptual-level content (e.g., beliefs, desires, thought, etc.) arise. Rather than coordinating behaviour with environmental states that would be encountered over the course of, say, minutes, or hours, mental states provide

²⁵ For other forms of interesting behaviour-guiding cellular memory in non-neuronal organisms see Burrell and Silver (2010).

²⁶ This behaviour-guaranteeing aspect of weak intentionality along with the fact that it is directed at environmental states suggests something very similar to the kind of undifferentiated indicative and imperative content of Millikan's (2004) "pushmi-pullyu representations". See also Clark and Toribio's (1994) "action-oriented representations".

organisms with a means to anticipate yet-to-be-encountered environmental states over the course of months and years (e.g., the environmental states that accompany seeding the soil in this spring will result in the states that accompany a bountiful harvest in autumn and hence avoiding hunger in winter).

At this pole of the continuum, anticipatory behaviour is no longer dependent upon an organism's learning (or having evolutionarily acquired an internal model of) correlations between exact environment state *relata* types. Moving from weak intentionality to strong intentionality, correlation rules become generalizable. Because environmental state type A' is similar to type A, content will be directed towards B subsequent to encountering state types A or A'. Importantly, this generalization is not merely due to the fact that sensory systems may respond similarly to stimuli that are alike; rather, it occurs in virtue of the fact that agents have access to their mental states and may compare like properties of perceived objects/events. Such access allows complex organisms to both metacognitively evaluate new beliefs against their other held beliefs and opens up the possibility of choosing the most appropriate behaviour that is consistent with one's set of held beliefs or desires. Such access means that content at the end of strong intentionality becomes more reliant upon learning at faster, ontogenetic timescales, allowing for living systems to elaborate upon the rules that have been acquired at slower phylogenetic timescales.

Moreover, content at the level of strong intentionality may be generated offline; it (at least in principle) may be fully insensitive to the influences of currently impinging environmental stimuli. As such strong intentional content may serve to both control and select a range of future actions in virtue of representing action outcomes over longer timescales. As organism and their environments complexify, rather than it being the case that encounters with A type states lead to content's being directed towards B type states, organisms can self-generate hypothetical B states that are desired (i.e., self-imposed goals) and work backwards towards the kinds of A-involving behaviours which would bring B about. This naturally places those living systems that have evolutionarily developed the capacity to engender future-oriented mental states in the advantageous position of using aboutness to engage in counterfactual reasoning (e.g., were I to rent a flat nearer to the university campus, I would spend less time in transit next term).

Another distinctive feature of aboutness at the level of strong intentionality is the fact that its effects upon maintaining homeostasis occurs at longer timescales; as intentional content increases in strength, not only does content's time horizon increase, but both content and content-driven behaviour *can* become further removed from needing to directly answer to an organism's homeostatic norms. Having the belief that "black widows are poisonous" whilst living in Edinburgh is valuable to one's adaptive behaviour even when such spiders are not native to the UK. For having such a belief could make the difference between careful and careless behaviour were one to encounter (and recognize) a black widow when visiting Los Angeles next year. At the far end of the continuum, complex animals such as ourselves may plan to bring out preferred states of affairs that very well may temporally fall beyond our own lifespans. For example, we are able to systematically formulate thoughts as to how interaction between humans and artificial intelligent agents *should* look in the best future case—a vision that corresponds to certain ethical, economic, and

scientific values. We can engage in such systematic thought despite the fact that the development of such artificially intelligent agents may be well beyond our lifetimes. It is because the aboutness of mental states may operate over increasingly longer timescales in the future (more future-oriented-ness) that complex organisms such as ourselves to have the capacity to plan and bring about the shape of distant events despite the fact that such events may have no direct affect upon our own homeostatic maintenance.

It is important to stress that although all content across the continuum of intentionality is future-oriented, it is only at the level of strong intentionality that content may come apart from causally bringing about behaviour; it is not behaviour-guaranteeing. Similarly, strong intentional content may be relevant to a range of different behaviours across varying timescales despite an agent never having to initiate any one of those behaviours (see Sterelny 2003 for a similar emphasis upon the relationship between what I am calling strong intentionality and content relevant to a range of behaviour).

I would like to suggest that since weak and strong intentionality are located at opposite ends of a continuum, there is no sharp cut-off between them.²⁷ Strong intentionality may be seen as grounded in the same anticipation-driven, self-preserving processes that underwrite weak intentionality. To use a phrase coined by Godfrey-Smith (1996), that weak intentionality “shades off into” strong intentionality does not imply that the former does not play a role in underwriting the latter. On the contrary, without the kind of weak intentionality that drives adaptive behaviour at shorter timescales, it would seem that strong intentionality could not arise. Consistent with BA’s evolutionary assumption that ‘higher’ cognition has evolved from simpler cognitive capacities, strong intentionality may be seen as an evolutionary achievement. I suggest that BA should construe strong intentionality as a variable trait that was *built upon a foundation of weak intentionality* and selected for because it allowed complex organisms to respond to selective pressures across progressively longer timescales and environments with more behavioural complexity. The daily behaviour of complex organisms is a continuous shifting between higher and lesser degrees of strong intentionality, but weak intentional states always subtend even the most complex behaviour.

One may reasonably speculate that the arrival of strong intentionality and the biological hardware underpinning it on the biological scene was closely linked to the fact that meeting selective pressures placed new requirements on organisms as they began to inhabit environments with more complexity (or their niches grew in complexity) (Godfrey-Smith 1996). One such requirement may have been being able to understand causal relationships without requiring that the agent itself be the cause. Intentional content that could be generated and monitored without entailing any

²⁷ That this is the case does not imply that there is not a distinction to be had between non-conceptual content and conceptual content. Conceptual content may be of a format that is a more recent evolutionary development, but as an expression of strong intentionality it is based upon the same weak intentionality that non-conceptual content is. I am grateful to an anonymous reviewer for pushing me to clarify this point.

behaviour at all, by allowing an agent to influence long-timescale dynamics and efficiently influence and predict other agents' behaviour made complex environments less hostile. The conditions underpinning the transition from lesser to increased environmental-organismal complexity might very well be a matter of what Sterelny (2003) describes as move from inhabiting "informationally transparent environments" (i.e., environments in which there are one-to-one mappings between sensory cues and environmental resources) to inhabiting "informationally translucent environments" (i.e., environments in which there are many-to-one or one-to-many mappings of sensory cues to resources).²⁸

Strong intentionality and the conceptual representations that often accompany it, as I am envisioning them, allow organisms to exert influence upon the protracted and often informationally ambiguous environmental dynamics that are specific to complex environments. They allow such organisms to utilize control states (i.e., conceptual-level representations) that become further and further removed from influencing (and the influence of) current environmental states. On the other hand, when simpler organisms are able to meet the selection pressures in their lesser complex environments, using conceptual representations that are suited to long-timescale environmental dynamics to guide behaviour in the here and now falls short of being an adaptive strategy. For too much organismal complexity (relative to the level of environmental complexity that 'satisfices' for self-preservation) renders an organism's responses inefficient, slow, and often detrimental to its survival. Although this may be the case, a continuum of intentionality helps to illustrate how responding to shorter timescale dynamics and less complexity nonetheless involves behaviour that is driven by weak intentionality.

We can therefore agree with Adams that conceptual representations are distinct from sensory representation while rejecting that only conceptual representations demarcate the domain of the cognitive. As one moves from weak to strong intentionality, the content and the kinds of representations change from having content that is directed at target yet-to-be-encountered objects/gradients causally underpinning sensorimotor behaviour to having content that can be directed at one's own mental states (e.g., a belief about one's own belief—that it is false). As one moves along the continuum from weak to strong intentionality, there is also a move from representations that are behaviour entailing and influenced (but not fully) by impinging sensory stimuli to representations that are behaviour contingent and possibly stimulus insensitive (i.e., content being immune to the influenced the impinging sensory stimuli). As the continuum approaches stimulus insensitive representation, agents acquire the capacity to shift from behaviourally exploiting their environment to exploring their environment with no goal other than knowledge gain (i.e., epistemic foraging) (see Fig. 2).²⁹

²⁸ Importantly, Sterelny suggests that the increased complexity of informationally translucent environments is partly due to presence of hostile agents (i.e., prey, predators, and competitors) which pollute environmental information for a given agent via the subversion of that agent's efforts or concealment. This induces the possibility of false negatives, making the obtaining of resources expensive.

²⁹ A further manner of understanding this transition across the continuum (at least along part of the continuum) is as a passage from content that is about changes in meaningful gradient intensity to content

Overly permissive or not?

The initial worry which set the stage for Adams' classification methodology and its upshot, the argument from intentionality, was that BA is overly permissive. Now given what has been said regarding the intentionality continuum, is there any reason for thinking that this worry has been defanged? I would like to suggest that the actual worry behind the claim that BA is overly permissive is this: accepting BA means that everything which an organism does now turns out to be cognition. I think this worry *would* be a real one if all biological processes were underwritten by weak intentionality. Recall however that weak intentionality regulates homeostatic mechanisms but is not itself homeostasis; there is intentionality *simpliciter* only from allostasis up. This being the case means that there is a range of ubiquitous biological processes (homeostatic control processes) which, when not subsumed under weak intentional dynamics, fail to rise to the status of cognitive processes. The possibility that all organisms engage in *some behaviour* that is driven by weak intentional dynamics does not imply that *every behaviour* of all organisms is so driven. As such, BA when buttressed by the intentionality continuum does not result in making every biological process cognitive.

Moreover, even if all organisms engage in some kind of cognitively driven behaviour (which very well might be the case if the continuum thesis is correct), this needn't be seen as evidence for a theory's being explanatorily barren. This might indeed be the case if, as a result of cognition being a ubiquitous feature of living systems, all accounts of cognition ended up being somehow indistinguishable (e.g., explaining cognition in bacteria turned out to be no different than explaining cognition in bonobos). If the intentionality continuum thesis is on track, however, it provides BA theorists with not only a conceptual manner of distinguishing different forms of cognition but a way of recognizing that phenomena located across the continuum are in fact in need of different explanations. Importantly, the continuum thesis suggests a way in which these explanations, although distinct, are nonetheless related and honours the biological facts as BA demands.

Conclusion

From the perspective of BA, a complete investigation of cognition does not involve relinquishing AA full stop, but rather rejecting the notion that all forms of cognition must satisfy human-based demarcation criteria. If what I have argued in this paper is correct, weak intentionality may serve as a common feature of cognition for both BA and AA; a feature that in ensuring that both approaches are not talking past one another, allows for the possibility of a complete investigation of cognition in its various expressions and degrees of complexity. Far from being

Footnote 29 (continued)

that is about the sources of the gradients (i.e., objects). This is a passage from what has been called "minimal perception" to maximal perception (see Sims 2019).

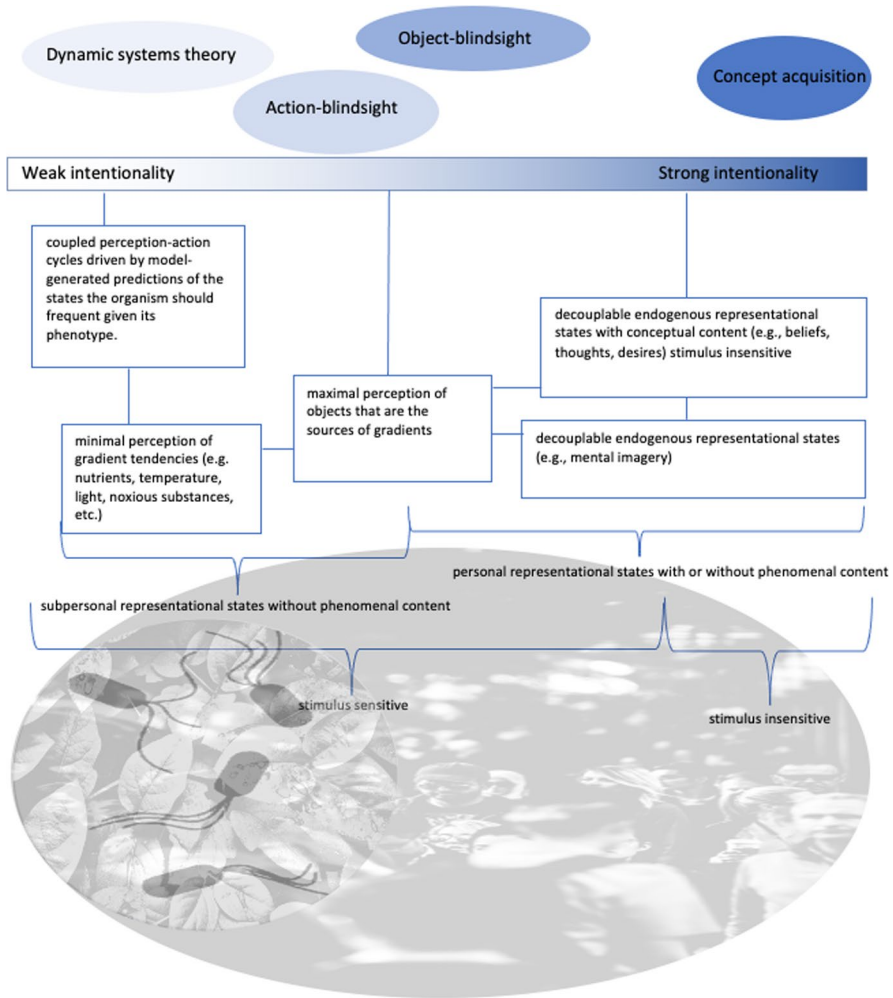


Fig. 2 A continuum of intentionality: *above* are some of the various research programmes which might be seen to investigate forms of cognitive phenomena associated with a particular degree of intentionality; *below* are some of the primary attributes of representational content as they are located at various places upon the continuum and capacities associated with them. When going from the weak end to the strong end of the continuum there is an increase in future-oriented-ness or how far into the future representational content can be directed

harmful to cognitive scientific enquiry, as Adams suggests, BA offers a valuable peek inside the possible evolution of biological cognition, shedding light on its various and surprising forms as located across a continuum of intentionality.

Acknowledgements I would like to thank the following people for both their incredibly useful feedback and the many discussions that lead to the sharpening of the ideas that populate this paper: Dave Ward, Julian Kiverstein, Andy Clark, Pamela Lyon, Michael Levin, Carrie Figdor, Nick Brancazio, Alistair Isaac, Eva Jablonka, Tobias Schlicht, Giovanni Pezzulo, Tobias Starzak, Marco Faccin, Nina Poth,

Kathryn Nave, Beren Millidge, Lee Wilson, Jonny Lee, Krzysztof Dolega, Francios Kammerer, Bartosz Radomski, and Benjamin Little. I would also like to express my deepest gratitude to two anonymous reviewers for their generous feedback and detailed comments, all of which have tremendously improved the quality of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL. This work was funded by the Alexander Von Humboldt-Stiftung.

Declarations

Conflict of interest The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adams F (2018) Cognition wars. *Stud Hist Philos Sci* 68:20–30. <https://doi.org/10.1016/j.shpsa.2017.11.007>
- Arbib MA (1992) Schema theory. In: Shapiro S (ed) *Encyclopedia of artificial intelligence*, vol 2, 2nd edn. Wiley, Chichester, pp 1427–1443
- Alfieri N (2008) Cellular stress response: from homeostatic to allostatic perspective. *Biosci Hypotheses* 1:330–331
- Baluška F, Levin M (2016) On having no head: cognition throughout biological systems. *Front Psychol* 7:902. <https://doi.org/10.3389/fpsyg.2016.00902>
- Baluška F, Mancuso S (2009) Deep evolutionary origins of neurobiology: turning the essence of “neural” upside-down. *Commun Integrat Biol* 2(1):60–65. <https://doi.org/10.4161/cib.2.1.7620>
- Bar M (2007) The proactive brain: using analogies and associations to generate predictions. *Trends Cogn Sci* 11(7):280–289
- Bartlett FC (1932) *Remembering: a study in experimental and social psychology*. Cambridge University Press, Cambridge
- Barrett N (2019) On the nature and origins of cognition as a form of motivated activity. *Adapt Behav*. <https://doi.org/10.1177/1059712318824325>
- Bear RD (2000) Dynamical approaches to cognitive science. *Trends Cogn Sci* 4(3):91–99
- Ben Jacob E, Shapira Y, Tauber A (2006) Seeking the foundation of cognition in bacteria: from Schödinger's negative entropy to latent information. *Physica A* 369(1):495–524
- Bechtel W (2014) Cognitive biology: surprising model organisms for cognitive science. In: *Proceedings of the 36th annual conference of the cognitive science society: cognitive science society, Austin, pp 158–163*
- Bechtel W, Bich L (2021) Grounding cognition: heterarchical control mechanisms in biology. *Phil Trans R Soc B* 376:20190751. <https://doi.org/10.1098/rstb.2019.0751>
- Bickhard M (2009) The interactivist model. *Synthese* 166:547–591. <https://doi.org/10.1007/s11229-008-9375>
- Bongard J, Levin M (2021) Living things are not (20th century) machines: updating mechanism metaphors in light of the modern science of machine behavior. *Front Ecol Evol* 9:650726. <https://doi.org/10.3389/fevo.2021.650726>

- Brentano F (1874/1995) *Psychology from an empirical standpoint*. Transl by A.C. Rancurello, D.B. Terrell, and L. McAlister. Routledge, London, 1973.
- Burge T (2010) *Origins of objectivity*. Oxford University Press, Oxford
- Burrill DR, Silver PA (2010) Making cellular memories. *Cell* 140(1):13–18. <https://doi.org/10.1016/j.cell.2009.12.034>
- Calvo P, Gagliano M, Souza GM, Trewavas A (2020) Plants are intelligent, here's how. *Ann Bot* 125:11–28. <https://doi.org/10.1093/aob/mcz155>
- Cannon WB (1929) Organization for physiological homeostasis. *Physiol Rev* 9(3):399–431
- Carruthers P (2019) *Human and animal minds: consciousness questions laid to rest*. Oxford University Press, Oxford
- Castelfranchi C (2005) Mind as an anticipatory device: for a theory of expectations. In De Gregorio M, Di Maio V, Frucci M, Musio C (eds) *BVAI 2005. LNCS, vol 3704*. Springer, Heidelberg, pp 258–276
- Chemero A (2009) *Radical embodied cognitive science*. MIT Press, Cambridge
- Christensen WD (2004) Self-directedness: a process approach to cognition. *Axiomathes* 14:171–189
- Clark A (1997) *Being there*. MIT Press, Cambridge
- Clark A (2008) *Supersizing the mind: embodiment, action, and cognitive extension*. Oxford University Press, Oxford
- Clark A, Toribio J (1994) Doing without representing? *Synthese* 101:401–431. <https://doi.org/10.1007/BF01063896>
- Cisek P, Kalaska JF (2010) Neural mechanisms for interacting with a world full of action choices. *Annu Rev Neurosci* 33:269–298
- Craik KJW (1943) *The nature of explanation*. Cambridge University Press, Cambridge
- Conant R, Ashby W (1970) Every good regulator of a system must be a model of that system. *Int J Syst Sci* 1:89–97
- Corcoran A, Pezzulo G, Hohwy J (2020) From allostatic agents to counterfactual cognisers: active inference, biological regulation, and the origins of cognition. *Biol Philos* 35:32. <https://doi.org/10.1007/s10539-020-09746-2>
- Crain T (2016) *The mechanical mind: a philosophical introduction into minds, machines and mental representation*. Routledge, London and New York
- Darwin C (1871) *The descent of man, and selection in relation to sex, vol 2, 1st edn*. John Murray, London
- Deacon T (2012) *Incomplete nature: how mind emerged from matter, 1st edn*. Norton, New York
- Dennett DC (1983) Intentional systems in cognitive ethology: the panglossian paradigm deferred. *Behav Brain Sci* 6:343–390
- Dennett DC (1984) *Elbow room: the varieties of free will worth wanting*. MIT Press/Bradford Books, Cambridge
- Di Paolo EA (2005) Autopoiesis, adaptivity, teleology, agency. *Phenom Cogn Sci* 4:429. <https://doi.org/10.1007/s11097-005-9002-y>
- Drescher GL (1991) *Made-up minds: a constructivist approach to artificial intelligence*. MIT Press, Cambridge
- Dretske F (1981) *Knowledge and the flow of information*. MIT/Bradford, Cambridge
- Engel AK, Maye A, Kurthen M, König P (2013) Where's the action? The pragmatic turn in cognitive science. *Trends Cogn Sci* 17(5):202–209
- Figdor C (2017) On the proper domain of psychological predicates. *Synthese* 194:4289–4310
- Figdor C (2018) *Pieces of mind: the proper domain of psychological predicates*. OUP, Oxford
- Fodor J (1975) *The language of thought*. Harvard University Press, Cambridge
- Fodor J (1987) *Psychosemantics*. MIT Press, Cambridge
- Fodor J (1990) *A theory of content and other essays*. MIT Press, Cambridge
- Friston K, Rosch R, Parr T, Price C, Bowman H (2017) Deep temporal models and active inference. *Neurosci Biobehav Rev* 77:388–402
- Ginsburg S, Jablonka E (2019) *The evolution of the sensitive soul: Learning and the origins of consciousness*. The MIT Press, Cambridge
- Godfrey-Smith P (1996) *Complexity and the function of mind in nature*. Cambridge University Press, Cambridge
- Godfrey-Smith P (2002) Environmental complexity and the evolution of cognition. In: Sternberg R, Kaufman J (eds) *The evolution of intelligence*. Lawrence Erlbaum, Mahwah, pp 233–249

- Goodson FE (2003) *The evolution and function of cognition*. Lawrence Erlbaum Associates Inc., Mahwah
- Grice HP (1957) Meaning. *Philos Rev* 66:377–388
- Grush R (2004) The emulation theory of representation: motor control, imagery, and perception. *Behav Brain Sci* 27(3):377–442. <https://doi.org/10.1017/s0140525x04000093>
- Hurley S (1998) *Consciousness in action*. Harvard University Press, Cambridge
- Jeannerod M (2006) *Motor cognition: What actions tell the self*. Oxford University Press, Oxford
- Kelso S (2016) On the self-organizing origins of agency. *Trends Cogn Sci* 20(7):490–499
- Kentridge RW, Heywood CA, Weiskrantz L (1999) Attention without awareness in blindsight. *Proc R Soc B Biol Sci* 266:1805–1811. <https://doi.org/10.1098/rspb.1999.0850>
- Kirsh D (2010) Thinking with external representations. *AI Soc* 25:441–454
- Kiverstein J, Sims M (2021) Is free-energy minimisation the mark of the cognitive? *Biol Philos* 36:25. <https://doi.org/10.1007/s10539-021-09788-0>
- Lakoff G, Johnson M (1980) *Metaphors we live by*. University of Chicago Press, Chicago
- Lambert G, Kussell E (2014) Memory and fitness optimization of bacteria under fluctuating environments. *PLoS Genet* 10(9):e1004556. <https://doi.org/10.1371/journal.pgen.1004556>
- Levin M (2019) The computational boundary of a “self”: developmental bioelectricity drives multicellularity and scale-free cognition. *Front Psychol* 10:2688. <https://doi.org/10.3389/fpsyg.2019.02688>
- Lyon P (2006) The biogenic approach to cognition. *Cogn Process* 7:11–29
- Macnab RM, Koshland DE (1972) The gradient-sensing mechanism in bacterial chemotaxis. *PNAS* 69(9):2509–2512
- Matrana HR (1970/1980) *Biology of cognition*. In: Maturana HR, Varela FJ (eds) *Autopoiesis and cognition: the realization of the living*, vol 42. D. Reidel Publishing Company, Dordrecht, pp 1–18
- Maturana HR, Varela FJ (1980) *Autopoiesis and cognition: the realization of the living*. D. Reidel Publishing Company, Dordrecht
- McBeath MK, Shaffer DM, Kaiser MK (1995) How baseball outfielders determine where to run to catch fly balls. *Science (new York, NY)* 268(5210):569–573
- McLeod P, Reed N, Dienes Z (2006) The generalized optic acceleration cancellation theory of catching. *J Exp Psychol Hum Percept Perform* 32(1):139–148. <https://doi.org/10.1037/0096-1523.32.1.139>
- Millikan RG (1989) *Biosemantics*. J Philos 86:281–297
- Millikan RG (2004) *Varieties of meaning*. MIT Press, Cambridge
- Milner MA, Goodale AD (1992) Separate visual pathways for perception and action. *Trends Neurosci* 15(1):20–25
- Milutinovic B, Kurtz J (2016) Immune memory in invertebrates. *Semin Immunol* 28:328–342
- Neisser U (1976) *Cognition and reality: principles and implications of cognitive psychology*. Freeman, New York
- Newell KM, Liu Y-T (2012) Landscape dynamics of motor learning and development. *Crit Rev Biomed Eng* 40:519–534
- Novoplansky A, Cohen D, Sachs T (1990) How portulaca seedlings avoid their neighbors. *Oecologia* 82:490–493
- O’Regan JK (1992) Solving the “real” mysteries of visual perception: the world as an outside memory. *Can J Psychol* 46:461–488
- Piaget J (1970) *Genetic epistemology*. Columbia University Press, New York
- Pezzulo G, Rigoli F, Friston KJ (2015) Active inference, homeostatic regulation and behavioural control. *Prog Neurobiol*. <https://doi.org/10.1016/j.pneurobio.2015.09.001>
- Pezzulo G (2008) Coordinating with the future: the anticipatory nature of representation. *Mind Mach* 18:179–220
- Pezzulo G (2011) Grounding procedural and declarative knowledge in sensorimotor anticipation. *Mind Lang* 26:78–114
- Popper KR (1965/1972) Of clouds and clocks: an approach to the problem of rationality and the freedom of man. In: Popper KR (ed) *Objective knowledge. An evolutionary approach*. Clarendon, Oxford, pp 206–655
- Riegler A (2001) The role of anticipation in cognition. In: Dubois DM (ed) *Computing anticipatory systems*. AIP Proceedings, Melville, pp 534–541
- Rodaki A, Bohovych IM, Enjalbert B, Young T, Odds FC, Gow NA, et al (2009) Glucose promotes stress resistance in the fungal pathogen *Candida albicans*. *Mol Biol Cell* 204845–4855
- Rosen R (1985/2012) *Anticipatory systems: philosophical, mathematical, and methodological foundations*. Pergamon, Oxford

- Saigusa T, Tero A, Nakagaki T, Kuramoto Y (2008) Amoebae anticipate periodic events. *Phys Rev Lett* 100:18101
- Schlicht T (2018) Does separating intentionality from mental representation imply radical enactivism? *Front Psychol* 9:1497. <https://doi.org/10.3389/fpsyg.2018.01497>
- Schulz AW (2018) *Efficient cognition: the evolution of representational decision making*. MIT Press, Cambridge
- Sims M (2019) Minimal-perception: responding to the challenges of perceptual constancy and veridicality with plants. *Philos Psychol*. <https://doi.org/10.1080/09515089.2019.1646898>
- Sims M, Pezzulo G (2021) Modelling ourselves: what the free energy principle reveals about our implicit notions of representation. *Synthese*. <https://doi.org/10.1007/s11229-021-03140-5>
- Stephan KE, Manjaly ZM, Mathys CD, Weber LAE, Paliwal S, Gard T, Tittgemeyer M, Fleming SM, Haker H, Seth AK, Petzschner FH (2016) Allostatic self-efficacy: a metacognitive theory of dyshomeostasis-induced fatigue and depression. *Front Hum Neurosci* 10(550):1–27
- Sterelny K (2001) *The evolution of agency and other essays*. Cambridge University Press, Cambridge
- Sterelny K (2003) *Thought in a hostile world: the evolution of human cognition*. Blackwell Publishing Ltd, Malden
- Sterling P, Eyer J (1988) Allostatics: A new paradigm to explain arousal pathology. In: Fisher S, Reason J (eds) *Handbook of life stress, cognition, and health*. Wiley, Chichester, UK, pp. 629–649
- Tagkopoulou I, Liu Y-C, Tavazoie S (2008) Predictive behavior within microbial genetic networks. *Science* 320(5881):1313–1317
- Thelen E, Smith LB (1994) *A dynamic systems approach to the development of cognition and action*. MIT Press, Cambridge
- Thompson E (2007) *Mind in life*. Harvard University Press, Cambridge
- Tversky B (2019) *Mind in motion: how action shape thought*. Hachette, New York
- van Duijn M, Keijzer F, Franken D (2006) Principles of minimal cognition casting cognition as sensorimotor coordination. *Adapt Behav* 14(2):157. <https://doi.org/10.1177/105971230601400207>
- van Gelder T (1995) What might cognition be, if not computation? *J Philos* 92:345–381
- Vertosick FTJ (2002) *The genius within: discovering the intelligence of every living thing*. Harcourt Inc, New York
- Villalobos M, Ward D (2015) Living systems: autonomy, autopoiesis and enaction. *Philos Technol* 28(2):225–239
- Walsh DM (2015) *Organisms, agency, and evolution*. Cambridge University Press, Cambridge
- Wang J, Atolia E, Hua B, Savir Y, Escalante-Chong R, Springer M (2015) Natural variation in preparation for nutrient depletion reveals a cost-benefit tradeoff. *PLoS Biol* 13(1):e1002041. <https://doi.org/10.1371/journal.pbio.1002041>
- Wheeler M (2010) In defense of extended functionalism. In: Menary R (ed) *The extended mind*. Life and mind series: philosophical issues in biology and psychology. MIT Press, Cambridge, MA, USA, pp. 245–270. <http://mitpress.mit.edu/books/extended-mind>
- Wiener N (1961) *Cybernetics; or control and communication in the animal and the machine*. MIT Press, New York

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Matthew Sims¹ 

✉ Matthew Sims
matthew.sims-m4e@ruhr-uni-bochum.de

¹ Institute for Philosophy II, Ruhr-Universität Bochum, Universitätsstr. 150, GA3/29, Bochum 44780, Germany