

On (not) defining cognition

Colin Allen^{1,2,3} 

Received: 7 October 2014 / Accepted: 20 May 2017 / Published online: 3 June 2017
© Springer Science+Business Media Dordrecht 2017

Abstract Should cognitive scientists be any more embarrassed about their lack of a discipline-fixing definition of cognition than biologists are about their inability to define “life”? My answer is “no”. Philosophers seeking a unique “mark of the cognitive” or less onerous but nevertheless categorical characterizations of cognition are working at a level of analysis upon which hangs nothing that either cognitive scientists or philosophers of cognitive science should care about. In contrast, I advocate a pluralistic stance towards uses of the term ‘cognition’ that eschews the urge to treat cognition as a metaphysically well-defined “natural” kind.

Keywords Cognition · Definition · Animals · Extended mind · Group cognition · Pluralism

1 Introduction

What is cognition? An answer to this question has seemed to many philosophers to be the *sine qua non* of cognitive science—or at least of its philosophical foundations. Others treat such understanding as the *sine fine qua non* (a bad Latin pun, but roughly that without which there is no end goal) for cognitive science. I will argue that the term “cognition” is at a level of description that is too general to be worth defining precisely or characterizing categorically. This view is sometimes assented to in vivo,

✉ Colin Allen
prof.colin.allen@gmail.com

¹ Department of Philosophy, Xi’an Jiatong University, Xi’an, China

² Department of History and Philosophy of Science, Indiana University, Bloomington, IN, USA

³ Department of History and Philosophy of Science, University of Pittsburgh,
1017 Cathedral of Learning, 4200 Fifth Avenue, Pittsburgh, PA 15260, USA

when philosophers of cognitive science meet face-to-face, but not fully represented *in folio*, where philosophers commit their arguments for all to read (but see [Chemero 2009](#), p. 201 n. 8 for a precedent).

To define cognition has seemed important to philosophers for a number of reasons. These reasons have varying degrees of relevance to the topics taken up by cognitive scientists. For instance, one distinction that has exercised philosophers, particularly those bearing the burdens of epistemology, is that between perception and cognition. This empirical case for a separation of perception from cognition also remains a live issue (see, e.g., [Firestone and Scholl 2015](#)). But given that areas such as perceptual psychology and machine perception fall fully within the scope of cognitive science, it is not a distinction that affects the range of phenomena properly studied by cognitive scientists.

While I would be equally critical of those seeking to expand the range of cognitive science by definitional fiat, my primary targets in this paper are those who insist on defining ‘cognition’ as part of a rearguard action against more expansive claims about the scope of cognitive science. One source of expansive claims springs from increasing scientific appreciation of the capacities of organisms (bacteria, plants, etc.) and collectives (insect swarms, human groups) previously assumed to be too different from prototypical human cognition to warrant serious attention by cognitive scientists. I describe some examples in the next section of this paper. Another source of expansion springs from emerging but still controversial trends in the cognitive sciences such as embodied and embedded cognition (promoted among philosophers by, e.g., [Clark and Chalmers 1998](#); [Clark 2010](#)), or dynamical systems theory (harnessed for “radical” embodied cognitive science by [Chemero 2009](#)). Bearing metaphysical burdens, which some attempt to cast off ([Chemero and Silberstein 2008](#)), these philosophers find no metaphysically defensible reason to limit the term ‘cognition’ to processes occurring only within the nervous system or bounded by the skin of an individual human being.

I have observed philosophers’ *in vivo* reactions to various expansionary claims that range from frowns and shudders, to strong expressions of emetic urges. It harder to find traces of these expressions committed to print, although Searle is on record as saying, “the theory of extended mind makes me want to throw up” ([Boag 2014](#)). A more commonly recorded response by philosophers who have strong intuitions about the nature of cognition is the response to state a set of conditions on cognition that are strong enough to rule out all and only the offensive cases.

Thus, [Adams and Aizawa \(2001\)](#) introduced the notion of “the mark of the cognitive” to undergird a distinction between processes that are constitutive of cognition and those that merely contribute causally. As their mark they proposed representations with “non-derived” intentional content. Yet the notions from which this suggestion is constructed are no better defined than cognition itself. Although the term ‘representation’ is widely used in cognitive science, it arguably does not mean to scientists what philosophers assume ([Ramsey 2007](#); also, [Ramsey 2015](#), this issue). Furthermore, the distinction between derived and non-derived intentional content has no status within the science, seeming instead to stem from introspective claims about the nature of thought and meaning that are of questionable probative force. Other attempts to characterize the mark of the cognitive fare little better. For example, with somewhat different objectives than Adams and Aizawa, Rowlands proposes a set of criteria for

the mark of the cognitive that includes information processing of representations by a process that “belongs to a subject” (Rowlands 2010, p. 111). Here, in addition to relying on a notion of representation, the notion of information may be problematic (discussed further below), and as Rupert (2011) points out in his review of Rowlands’ book, the notion of “belonging to a subject”, which has roots in phenomenology, is difficult to spell out clearly, and not well-grounded in the modeling practices of actual cognitive science.

A more scientifically-grounded motive for defining cognition comes from a close look at practices within comparative cognition, wherein cognitive explanations of animal behavior are contrasted with the “merely associative” (Buckner 2015). Buckner argues that there are inclusionary and exclusionary versions of the distinction between cognition and association, and that the exclusionary version can be made good by conceiving cognition as a natural kind underlying behavioral flexibility. By extension, one might argue (although Buckner does not do so explicitly) that the scope of cognitive science is limited to organisms with the right kind of behavioral flexibility. The notion of ‘natural kind’ here incurs a metaphysical burden that I shall return to among my replies to objections below. Meanwhile, Buckner’s proposal also highlights another reason that philosophers have thought it necessary to define cognitive processes in an exclusionary sense, in order to distinguish cognitive science from the psychological behaviorism that preceded it (see also Aizawa 2015).

In this paper I advocate a pluralistic stance towards uses of ‘cognition’ that eschews the urge to treat cognition as a category that would benefit from the kind of regimentation proposed by some philosophers. In the next section I describe some empirical findings concerning “simple” organisms that challenge the idea of a sharp line between “genuine” cognition merely metaphorical extensions of cognitive terminology. I then discuss the organizing role that the term ‘cognition’ plays for cognitive scientists, particularly focusing on the more specific capacities that are actually studied experimentally and observationally. This leads me to articulate the stance I call “relaxed pluralism” that is nevertheless an industrious pluralism. Finally, I reply to various objections to this stance.

2 Prokaryote preamble

More than four decades ago, with an elegant experiment, Macnab and Koshland (1972) argued for some sort of “memory” mechanism in a species of bacterium, *Salmonella typhimurium*. Bacterial chemotaxis is a well-known characteristic of many different species of bacteria, and it was also known that bacterial motion had two distinct modes: an apparently random tumbling mode, and directional movement up a chemical gradient, which Macnab and Koshland called “superordinated swimming”.

Naturally the question arises how bacteria manage to swim in the right direction when a gradient is present. An obvious suggestion is that different numbers of molecules binding to chemoreceptors distributed across the bacterium’s outer membrane can be used to drive directional swimming towards the side that has bound the most molecules of the relevant type. Macnab and Koshland argued, however that the size of a bacterial cell is too small and the concentrations in which superordi-

nated swimming appears are too low for them to be able to detect such differences in concentration spatially. They proposed, instead, that the bacterium makes a temporal comparison: detection of a temporal gradient. They created a device in which bacteria tumbling in a uniform medium would experience a simultaneously global increase in the concentration of a chemotactic attractant. When they ran the experiment, tumbling bacteria indeed switched to superordinated swimming when the concentration was raised, but bacteria headed off in different directions depending which way they had been moving just prior to the concentration being raised. As Macnab and Koshland put it, “The apparent detection of a spatial gradient by the bacteria therefore involves an actual detection of a temporal gradient experienced as the result of movement through space.” Macnab and Koshland’s use of scare quotes in calling this a “memory” mechanism most likely reflects a combination of influences. The general behavioristic *Zeitgeist* of the time was against using psychological terminology for any non-human organism. Additionally, specific concerns about attributing such properties to such “simple” organisms as bacteria are not hard to generate. It is noteworthy that they did not feel the need to mark “experienced”, however.

In fact, bacteria are far more complex than most philosophers (or psychologists) have imagined, and with each revelation about this complexity, biologists have become increasingly willing to drop the scare quotes around “memory”, as a Google search for the combination of terms “memory” and “bacteria” will reveal. For example, in a paper titled, “Memory in Microbes: Quantifying History-Dependent Behavior in a Bacterium”, Wolf et al. (2008) write, in the paper’s abstract, “Memory is usually associated with higher organisms rather than bacteria. However, evidence is mounting that many regulatory networks within bacteria are capable of complex dynamics and multi-stable behaviors that have been linked to memory in other systems.” In the discussion section of their paper, they write, “in this work we propose that the familiar phenomenon of history-dependent behavior in microbes reflects a form of memory worth studying systematically and quantifying.” None of the occurrences of the term “memory” here is marked by scare quotes, and it is important to note that Wolf et al. justify their talk of bacterial memory by drawing an explicit connection between it and other systems with memory—a connection that they ground in the dynamics of complex systems and test in *Bacillus subtilis*.

It is tempting to argue that extending the notion of memory in any cognitive sense to bacteria is merely metaphorical, for *genuine* cognition requires *genuine* intentionality (the word I have italicized appearing ubiquitously in the philosophical literature). Among those who sport such intuitions, this genuineness often requires that intentional content be accessible to a conscious subject (for example, Searle 1992; Ludwig 1996). Such authors are inclined to insist on direct, conscious knowledge of their own original intentionality while asserting that no one, themselves included, has any clue how to explain consciousness. The connection between genuine intentionality and consciousness cannot be articulated beyond the assertion that it seems to them to exist.

The well of certainty that bacterial memory has nothing to do with “genuine” cognition is reinforced by the apparent ease by which philosophers think they can imagine a mechanism for “simple” capacities such as chemotaxis; an ease that is sometimes expressed by people, even (or, especially) those with limited software

programming experience, who say, “I could program a computer to do that”. I take this to be an instance of what Rozenblit and Keil (2002, p. 521) have termed “the illusion of explanatory depth”, wherein people “feel they understand complex phenomena with far greater precision, coherence, and depth than they really do.” Combine this with the pervasive sense of mystery about consciousness, and the conditions are right for a perfect storm of intuition-driven line drawing.

Many scientists (and some philosophers) are, however, beginning to understand that the molecular binding mechanisms regulating communicative and adaptive capacities of bacteria are deeply homologous, structurally and functionally, to the neuronal NMDA-receptors central to synaptic long term potentiation in learning and memory (Kuryatov et al. 1994; Baker and Stock 2007; Paoletti and Neyton 2007; see also Stotz and Allen 2011). Brains are, in the end, large assemblies of individual cells acting through the even larger assemblies of cells we call “bodies” so as to produce coordinated, adaptive behavior in interaction with the world. It should come as no surprise that the cellular basis for such adaptation reuses molecular mechanisms that govern the individual and collective behaviors of single-cell organisms. Indeed, Iyer et al. (2004) hypothesize that cell-signalling mechanisms in eukaryotes were acquired via lateral gene transfer from prokaryotes.

3 From generic cognition to specific cognitive capacities

Claims on behalf of bacterial cognition (previous section), slime mold intelligence (e.g., Nakagaki et al. 2004; Reid et al. 2012), and plant cognition (Garzón and Keijzer 2011), raise boundary questions about the scope of cognitive science. The urge to settle such boundary questions by defining the high-level notion of ‘cognition’ is understandable, but any attempt to do so using philosophical precepts, such as ‘intrinsic intentionality’, or (quasi) theoretical terms such as ‘representation’ (critiqued by Ramsey 2007, 2015/*this issue*) is likely, in my view, to be rendered otiose by scientific developments. In this section I rely on a comparison to biological science to make the case that a definition of cognition is less important than attention to theories and models of specific capacities. Scientific fields aggregate around common questions and techniques more than well-defined terms delineating exact domains of application. Biologists don’t have and don’t need an exact definition of ‘life’ and ichthyologists don’t fret about the fact that ‘fish’ is a phylogenetically impure (“paraphyletic”) term, or that the category of fish cannot be defined in terms of intrinsic features—most are ectothermic, some are endothermic; most are oviparous, some are ovoviviparous, and a few are even placental; most have scales, some do not; etc. Biologists don’t need a definition of ‘fish’ that abstracts from all this variety. Likewise, I maintain, cognitive scientists need neither an abstract definition of ‘cognition’ nor a theoretically pure conception of ‘cognitive system’.

By pinning my critique on what is useful to cognitive science, it might be thought that I am giving up on doing philosophy for its own sake. All participants to this debate might be willing to concede the point that biologists and cognitive scientists do not need definitions of ‘life’ and ‘cognition’ in advance of their research. But Adams and others might suggest that when discussing extended cognition we are no longer

doing science but philosophy. For the sake of argument, I may accept that there are questions that philosophical reflection, and philosophical reflection alone, can answer. However, I maintain that the boundary between doing science and doing philosophy is not precise, and questions about extended cognition straddle the putative boundary. A further refinement of the suggestion might be that when discussing whether cognitive processes extend into the environment, we are no longer doing science but trying to identify the philosophical implications of certain scientific views—or, as argued by Aizawa (2015), we are concerned not with normal science but with revolutionary science. I agree that philosophical reflection is indispensable for investigating the philosophical implications of scientific findings or theories, and extraordinary science may deserve extraordinary attention to conceptual issues. Such work is also potentially valuable to practicing scientists insofar as it helps to clarify the commitments of the theories they are developing and to identify potential conflicts among theories. But the specific claim I wish to advance in this paper is narrower in scope: it is not that we can or should dispense with philosophical reflection in general, but specifically that ‘cognition’ is too much an umbrella term for it to be usefully defined either for philosophical or scientific purposes (see also, Ramsey 2015). Furthermore, even if the present moment is a revolutionary moment for cognitive science, I would argue that terms and concepts demanding extraordinary attention are those which give the new alternative its substance, such as embodiment, embeddedness, and the key concepts of dynamical systems theory.

Detailed scientific investigation of the adaptive information processing capacities of a variety of systems—organized at different scales, from bacteria to humans to honeybee swarms to groups of humans—is typically conducted at the level of concepts such as learning, memory, problem solving, and decision making rather than at the more abstract level of “cognition” (Theiner et al. 2010). Phenomena specific to learning, memory, etc., are targeted by experiments and models. Although the term “cognition” is used in scientific contexts, in my view it serves as a kind of umbrella term under which more specific capacities are grouped, whereas (with one subfield of cognitive science providing an exception) all serious modeling efforts and investigation of mechanisms pertain to the specific capacities. The exception concerns those subareas of cognitive science that concern themselves with models that attempt to combine the various capacities in an overall “cognitive architecture”, such as parts of cognitive neuroscience, AGI (Artificial General Intelligence), and robotics. Such general modeling efforts do not, however, proceed by first using a general concept of cognition to determine what should or should not be included in the model. Rather, they attempt to combine elements of other theories and models into an architecture that can solve more complex and realistic tasks. For example, John Anderson’s seminal work in this area extends theories of problem solving (esp. Newell and Simon 1972) with ideas from learning theory, by treating production rules as the targets of a basic learning mechanism (Anderson 1983, 1993).

In the context of such general models, the notion of cognition is not precisely defined, although general characterization of a cognitive system is sometimes attempted. For example, Simon and Newell (1970) describe human thinking as an information-processing system, but they note that details of the system are “elusive” (1970, p. 149) because of its adaptive nature. If “adaptive information processing” is

a reasonable working definition of cognition—a definition suitable, that is, for orienting newcomers to phenomena of potential interest—it is inadequate for the kind of ontological purposes that have driven philosophers to seek the mark of the cognitive. This is because the key notions of adaptiveness and information are no less a source of ambiguity or controversy than ‘cognition’. One might, for example, raise questions about the timescale over which a cognitive system should manifest adaptive processing: momentary, lifetime, multigenerational, or evolutionary? This issue is entirely germane to the question of whether plants engage in adaptive information processing, since their responses to changes in environmental conditions tend to happen much more slowly than those of animals. The notion of information proves just as problematic. Some have embraced Shannon’s (1948) mathematical theory of information (e.g., Dretske 1981, who simultaneously disavows a connection between Shannon’s conception and the ordinary meaning of the term). Others consider the application of Shannon’s theory of information to cognitive science to be limited due to its lack of semantics (Piccinini and Scarantino 2011) or because it lacks the resources to capture important structure in psychological stimuli (Luce 2003). Even Shannon himself urged caution about the application of a mathematical formalism to an area for which it was not originally conceived, urging that the appropriateness of doing so depends on “the slow tedious process of hypothesis and empirical verification” (Shannon 1956). Thus, to characterize cognition in terms of adaptive information processing brings us no closer to resolving boundary questions.

Given these issues (and others that could be added), how then does ‘adaptive information processing’ manage to function as a working definition? In one sense, a simple answer is just to say that it does so function insofar as it gives new students some relatively familiar terms that help orient them to the phenomena of interest. Of course, one might quibble with the individual terms; plenty of cognitive science has focused on the apparently *maladaptive* aspects of human psychology, for example, but even here it seems clear that the challenge is to explain such anomalies against the background of a set of cognitive capacities that are generally assumed to work to the organism’s advantage. The term ‘information’ serves to indicate a level of abstraction between the circuit-level details of brains or computers (and perhaps groups of organisms, or individual cells) and the typically adaptive functions served by a variety of mechanisms. And the notion of processing serves to indicate the structured and active aspects of such mechanisms. To insist on precisely defining the terms that make up a working definition is to put the cart before the horse inasmuch orienting towards phenomena worthy of further investigation does not depend upon the kind of precisification that follows from detailed study of those phenomena—just as fruitful investigation of samples of soft, yellow metal did not depend on providing, in advance, precise definitions of the terms making up that working definition of ‘gold’.

At this stage of empirical enquiry, the search for a defining mark of cognition that can serve as a touchstone for genuine cognition is, I maintain, akin to insisting (incorrectly) that biology needs a general definition of “life”. For the meaning of ‘life’, Cleland (2012) argues that all definitional strategies are flawed, and Machery (2012) argues that the exercise is either impossible or pointless. Biology has achieved enormous success without such a definition, and it is far from clear what would be gained by a definitional ruling whether (say) viruses are alive or not. Whether or not

viruses are stipulated to be alive, biologists will continue to study viruses (but not rocks) using the same kinds of techniques and models that they use to study cellular forms of life. Villareal (2004) argues that there would be heuristic advantages to giving an affirmative answer to the question, “Is pondering the status of viruses as living or non living more than a philosophical exercise?” Specifically, he thinks that viruses deserve more attention from biologists working in non-medical fields. Analogously, cognitive science gains from taking a broad perspective on cognition, but given that the usual motivation of those who argue for a mark of the cognitive is to restrict the scope of cognitive science, rather than broaden it, their philosophical exercise does not afford the same opportunities.

Unhampered by philosophical scruples, cognitive scientists are increasingly applying their models and analyses to systems above and below the level of a prototypical, multicellular animal with a brain, spanning unicellular organisms and systems with just a few hundred neurons below, and to the collective action of colonies or groups of organisms above. Some philosophers (but not all) remain conservatively committed to historically favored notions of cognition involving rational judgment, intentionality, or consciousness, leading them to resist such applications. But conservatism about the meanings of terms in a science has a poor track record (see also Figdor 2014, for discussion on non-prototypical uses of cognitive predicates). I will return to the question further below of whether the kind of liberalism I am allowing for here amounts to changing the question.

4 The range of cognitive science

Cognitive scientists study everything from the control of behavior in the 302-neuron nervous system of the roundworm *Caenorhabditis elegans* (Izquierdo and Beer 2013) to the neural dynamics of the human connectome (Sporns 2012), and from the learning and categorization abilities of invertebrates (Srinivasan 2010) to cognition in groups of organisms, whether swarms of honeybees or other animals (Couzin 2009; Seeley et al. 2012) or groups of humans (Theiner et al. 2010). And while notions of cognition in plants and prokaryotes remain niche interests among cognitive scientists, ideas about quorum sensing and signaling between single-celled organisms (Fuqua et al. 1994) are feeding back into attempts to understand how billions of neurons can coordinate the behavior of large, cognitively sophisticated animals. Likewise, models of cross-inhibition which can be applied both at the level of interactions among neurons and to the group dynamics among dancing bees can explain how a bee colony manages to coordinate the decision to swarm to one new nesting site rather than another (Seeley et al. 2012).

At the heart of the broad-tent approach to cognitive science that I espoused above is a commitment to letting the productivity of research programs in cognitive science guide the extension of language to new contexts. The slogan that “cognition is whatever cognitive scientists study” is partial and inadequate, however. Quite apart from its failure to recognize differences of opinion among cognitive scientists about what to study, it also suggests something too unidirectional in the relationship between science and philosophy. Historically and epistemically, science has been the dominant partner

to philosophy of mind: driving perceived problems and theoretical innovations in the latter. But philosophical insight and creativity also can and does contribute to advances in science, and even philosophical skepticism about science has its place as a healthy counterweight to the naive acceptance of scientific pronouncements.

Those who wish to extend the language of cognition to cover nonstandard cases may not owe a “story” about why it is appropriate and useful to do so—such practices will live or die according to their empirical successes or failures. Nevertheless, the provision of such story may usefully go beyond the interpretation of the terms of a theory or model in specific experimental contexts, to link those theories and models to the phenomena motivating scientific enquiry in the first place (Hartmann 1999). This means explaining how models and mechanisms of putatively cognitive capacities falling outside the traditional core of rational human judgment nevertheless connect to the phenomena motivating the investigation of cognition in the first place. Here, notions such as memory, problem solving and the like, play a dual role. They connect nonstandard cases, such as bee swarms, to the decision making that arises from interactions among neurons in the large-brained creatures more typically studied by psychologists. They also connect better to vernacular English terms such as “memory” and “problem solving”, than to “cognition” which is much more a piece of philosophical and scientific jargon.

5 Relaxed pluralism, but industrious all the same

If there’s a label for the stance I am advocating, it might be called “relaxed pluralism” about cognition. It is a stance in the sense of van Fraassen (2002); i.e., a set of attitudes towards ways of characterizing the relationship between metaphysics and the science. It is pluralistic in the sense that it tolerates different ways of selecting which natural phenomena are appropriate targets for investigation within the science, even when they make incompatible judgments about cases. And while relaxed it is not “lazy”—that is, not just anything goes. Rather, the point is that enquiry should not be stifled by a conservatism about terms and their meanings that insists on stipulating what we are studying before we study it, especially when this conservatism is coupled with an introspectively-based claim to knowledge of the subject of enquiry that is highly resistant to empirical adjustments. Nonetheless, the relaxed but industrious pluralist must offer an explanation for why these apparently incompatible ways of carving up the phenomena do cohere, rather than driving the discipline toward disintegration. The task of providing such an explanation is beyond the scope of this paper (but see also reply to objection 3 below). However, I favor a model-centric approach (descriptions of mechanisms are models too) that builds on the examples provided below.

Relaxed but industrious pluralism thus allows multiple targets and lines of enquiry while insisting on rigor where it can be provided, such as in formal and mechanistic models of particular capacities related to memory, decision making, etc. My claim is that many extensions of the cognitive umbrella to non-standard targets are best understood as the application of models and methods from cognitive science to the investigation of processes that may, and often do, turn out to be shared with more standard examples of cognition. In this respect, they are neither metaphorical, or the

result of a lazy, anything-goes approach to the use of terminology. It takes hard work to discover and understand these connections.

Can the relaxed pluralist do more than note that models are diverse and often have broader than intended utility? The challenge is to explain why different systems at different scales composed of different parts converge on similar but not identical solutions. One promising approach to meeting this challenge recognizes that there is an interaction between the informational structure of environments and the structure of adaptive information processing systems. The nature of the task places more constraints on the structure of adaptive information processing systems than is generally understood, particularly by philosophers who have acquired strong intuitions about multiple realizability. (See [Polger and Shapiro 2016](#) for challenges to those intuitions.) But it is a two-way street because organisms also shape their environments. The quest to understand the diversity of cognitive systems without losing sight of the commonalities is aided by an understanding of principles by which such systems display self-organizing emergent properties. By “emergent” here, I mean non-aggregativity in Wimsatt’s (1986) sense (see also [Theiner et al. 2010](#)). By “self-organizing” I mean that the relevant organization emerges reliably from local interactions among the units making up the whole, and such organization is relatively robust once it emerges. For example, Sur and colleagues ([Angelucci et al. 1997](#); [Sharma et al. 2000](#); [Newton and Sur 2005](#)) have shown that after the optic nerves of neonatal ferrets were rewired to primary auditory cortex (A1), the neural organization of A1 converged on a network topology called “pinwheel structures” similar to those found in primary visual cortex area V1 in surgically unmodified animals. The structure is emergent in that it cannot be characterized as the sum of individual, non-relational, neural properties. And it is self-organized in that both A1 and V1 reliably and robustly produce the pinwheel structures given typical visual input. It is (part of) a cognitive mechanism insofar as these structures help explain how organisms process visual input to guide behavior.

The time scale for this emergent organization is a developmental one of several weeks. But cognitive systems also self-organize over both shorter and longer time scales. On a much shorter time scale, the previously mentioned study by [Seeley et al. \(2012\)](#) of collective decision making in honeybees also illustrates the robustly self-organizing capacity of populations of bees that use cross inhibition between sub-populations to “integrate noisy evidence for alternatives” (2012, p. 108). At a much longer time scale one might point to the convergent evolution of neural structures similar to mammalian hippocampus in a variety of recognized “intelligent” invertebrates such as the common octopus and the honeybee. For instance, [Hochner et al. \(2006, p. 315\)](#), write, “The findings emerging from recent electrophysiological studies in the octopus suggest that a convergent evolutionary process has led to the selection of similar networks and synaptic plasticity in evolutionarily very remote species that evolved to similar behaviors and modes of life. ... Features are found in both the octopus MSF-VL system and the hippocampus, ... it would appear that they are needed to create a large capacity for memory associations.”

An important feature of these examples concerns similarity of structure at different scales. In the mushroom body of honey bees, MFS-VL of octopi, and the mammalian hippocampus, very different numbers of neurons are organized into similar structures with similar functions. In the cross-inhibition example, the units comprising the com-

peting populations are at rather different scales: populations of neurons in one case and populations of bees, each containing roughly a million neurons, in the other. Both are organized to similar effect, although these similarities should not be caricatured as strictly identical in either form or function (Polger and Shapiro 2016). Self-organizing similarity at different scales is also apparent in comparing the learning capacities of the mammalian spinal cord to the learning capacities of the mammalian neocortex. Even though the spinal cord is much more restricted in important respects, it is organized in such a way that it exhibits classical (Pavlovian) and instrumental learning, and exhibits phenomena such as latent inhibition and overshadowing—phenomena which have been used by learning theorists such as Rescorla (1988) to argue that a cognitive rather than strict behaviorist approach to Pavlovian learning was necessary (see, also, Allen et al. 2009). Another way in which cognitive models may apply at different scales is provided by Thomas Hills and colleagues (Hills 2003; Hills et al. 2008). Hills (2003) argues that the physiological mechanisms controlling the search for food are deeply conserved in evolutionary history, from *C. elegans* to modern mammals, adaptively adjusting the organisms' exploration and exploitation strategies according to how patchily food is distributed in the environment. Hills et al. (2008) conducted experiments with human subjects and found that persistence in a memory search task was affected by the degree of patchiness in a prior task involving spatial search for unrelated resources.

These examples indicate that the relaxed pluralist has materials with which to do more than simply accept the utility of models in various domains. The recognition of principles underlying the emergence of similar structures for similar functions in different systems motivates the search for deeper connections among the different domains, by behavioral experiments and through the attempt to identify common models and mechanisms. I have tended to emphasize similarity in these examples, but the proper mantra is that of *similarities and differences*. The structural and functional diversity of adaptive information-processing systems does not strictly eliminate the possibility of a “mark” belonging to all and only cognitive systems. It does, however, call into question the utility of seeking such a mark when so much remains to be learned about the variety and complexity of intelligent systems. Rather than trying to settle metaphysical issues or terminological disputes conducted a high-level of abstraction from the experiments and observations, cognitive science is better served by remaining open to the idea that swarms of bees and networks of neurons may reach consensus about a collective decision in similar ways, even though their decision-making capacities differ in many respects. The notion of similarity remains to be fully explicated here, but advances in network science and in information theory are providing new measures (Newman 2010; Beer and Williams 2015). Nevertheless, these measures reduce complex multidimensional phenomena to lower-dimensional representations and their applicability to the broad range of phenomena under the cognitive umbrella does not entail that any one feature (or feature dimension) can supply a “mark of the cognitive”.

6 Objections and replies

1. *Objection: Uses of “cognition” outside the paradigmatic cases (maybe only humans, but perhaps other large-brained animals too) is a Pickwickian misuse of language. (See Ludwig 2015 for an accusation of this kind with respect to “group cognition”).* Reply: This is not so much an objection as a restatement of the conservative view that words have meanings that are fixed by prior usage. I am inclined to resist the assumption that language has such determinate meanings that competent users know and must respect. Are dwarf planets planets? Are pronghorns (*Antilocapra americana*) antelope? Is the crowd-sourced Wikipedia authoritative about the common or scientific meanings of taxonomic terms? What is cognition? Competence with the word “cognition” does not entail knowing a fixed meaning, but rather being competent with a range of uses, which may diverge from paradigmatic human cognition in different ways, and which may reasonably diverge from someone else’s preferences about how to use the word. Stipulation (about the meaning of ‘planet’ or ‘antelope’, e.g.) may be useful for certain purposes, but pointless in other contexts. (And note that “antelope” is not actually considered a *bona fide* taxonomic group among evolutionary biologists, who prefer a naming scheme that tracks clades, i.e., all the descendants of a common ancestor.) However, if someone insists on stipulating that cognition is essentially related to conscious understanding of meaning (see discussion related to Searle 1992; Ludwig 1996, above), let them have the term but point out that by that standard, much of cognitive science is not about cognition, and neither would it be useful to constrain the science by so stipulating.
2. *Objection: Cognitive science has changed the subject then.* Reply: Science always changes the subject. Or, to put it less tendentiously, the kinds of changes in conception that are on the table under the broad umbrella do not amount to changing the subject in any pernicious sense. Human understanding of the world, including philosophy, is changed by the complex interplay between empirical findings and scientific theorizing. If anyone thought that Darwin changed the subject by turning the notion of species away from one of fixed, immutable kinds to something that, over a century and a half later, is still resistant to precise definition, so much the worse for whatever intuitions they started with about the nature of species, and the worldview underlying those intuitions. If anyone complained that Descartes changed the subject by adopting a mechanistic approach to physiology that rejected *a prioristic*, vitalist assumptions about Aristotelian teleology, so much the worse for the teleology.
3. *Objection: if cognition is an umbrella concept, then there is no theoretical unity to cognitive science and cognition is not a natural kind term.* Reply: Even physics hasn’t achieved theoretical unity! The question, “What is a physical phenomenon?” looks like a question only a philosopher could love. There are many physical phenomena, and although physicists strive to unify them all within a single explanatory framework, it’s not a *sine qua non* for doing good physics that one start with a definition that unifies all the various phenomena. There is, as yet, no grand unified theory that unites gravity with the other three forces, and all attempts to treat gravity as the others have been treated, by positing a kind of particle (the “graviton”) as the

bearer of the force, founder mathematically or empirically. And while the forces are “unified” through conservation principles (e.g., allowing electromagnetic and gravitational energy to be interconvertible), these principles are phenomenological rather than fundamental. Is “force” a natural kind? Who knows, and why should physicists answer this question in advance? Perhaps, after a lot more research, the question will be answered, but the *post hoc* determination that force was (or was not) a natural kind is of no practical import to the science. It is no wonder that scientists become puzzled when philosophers tell them that they (the scientists) are in the business of discovering natural kinds. The metaphor of carving nature at its joints is just that, a metaphor, and one can find many examples of terms that scientists use where the joints are fuzzy at best. Despite the fact that fish are not a monophyletic group (equivalently, not a clade, thus outside the preferred naming structure for evolutionary taxonomists), there are still fish biologists. Worse still, marine mammalogy spans porpoises and polar bears. That “marine mammal” is a junk category from a taxonomic perspective does not mitigate the scientific utility of housing those who study cetaceans and polar bears under the same roof. Porpoises and polar bears both have adaptations for life at sea, and the experts who study them are usefully housed together for many reasons, both intellectual and practical. But this doesn’t motivate a hunt for a mark of the marine mammal, and even if a very generic definition of marine mammal can be given, it is of no theoretical interest.

4. *Objection: You are resisting the attempt to give a substantive account of “cognition” because you are working with an outmoded notion of natural kind, rather than the more sophisticated, less metaphysically-loaded Boydian notion of a homeostatic property cluster (HPC) (Boyd 1991).* Reply: Maybe it is correct that the variety of mechanisms that underlie paradigmatic cases of cognition are causally-related in a non-accidental “homeostatic” way (assuming that the notion of homeostasis is well-defined in this usage). Maybe there is already enough evidence from comparative cognition (reviewed by Buckner 2015) that cognition is an HPC cluster of more specific kinds—i.e., memory, learning, problem solving, etc.—and that shared mechanisms reliably cause the elements of this package to cluster, thus qualifying it as a “superordinate” HPC kind (Buckner 2015). Perhaps there are general reasons why these things non-accidentally cluster, and even share mechanisms. But perhaps, alternatively, those reasons are no more profound than the reason why radios, electronic fuel injection, and adjustable wing mirrors non-accidentally cluster in automobiles, and share a common electrical system. At any rate, while I regard the notion of a homeostatic property cluster as an improvement over essentialism in metaphysics, it does not provide any particular guidance about how to pursue cognitive science. It appears rather to be another *post hoc* labeling of something that would arise from the research anyway given the pluralistic pursuit of cognitive science I have adumbrated. Namely, an approach focused on models of specific capacities and architectures capable of integrating these capacities into systems in which adaptive information processing occurs at multiple timescales, from milliseconds to millennia.
5. *Objection: Doesn’t replacing general debates about whether or not something is cognitive with debates about whether or not something is a specific kind of*

process—such as memory, concept-formation, decision-making, and so on—lead to the same kinds of issues and worries at this level also? For example, could not Clark and Chalmers’ case of Otto be run as an argument for extended memory without even mentioning cognition? Reply: Clark & Chalmers are not my target here, except insofar as I think they run their argument at a level of abstraction that leads other philosophers to search for the ‘mark of the cognitive’. Insofar as the debate would shift to being about whether Otto’s interaction with his notebook are explained by the same scientifically-validated models of memory as neurologically-intact Inga, that would be progress.

7 Concluding remarks

Cognitive scientists study adaptive information processing at multiple scales, both spatial and temporal. Intuition-based arguments that attempt to establish the impropriety of cognitive scientists broadening their attention to non-paradigmatic cases are unconvincing. Demands that a definition of cognition can and should be stipulated in advance are not true to the actual operation of science. Although *post hoc* identification of high level categories can help organize research, such efforts typically have little theoretical importance. Cognitive scientists should no more be embarrassed about their lack of an all-purpose definition or categorical description of cognition than biologists are about their inability to define “life”, “species” or any other number of terms, or for physicists to be able to give a unified account of such fundamental notions as “force” and “matter”. Philosophers seeking a unique “mark of the cognitive” or less onerous but nevertheless categorical characterizations of cognition are working at a level of analysis upon which nothing of substance hangs, or using concepts which are no less obscure than those they are intended to analyze. The real action lies, as I have indicated, in the pursuit of explanations of more specific capacities. Having an answer to what makes the list of things such as learning, memory, concept formation, etc. “cognitive” is, according to Adams and Garrison (2013, p. 339), “one very important reason to be interested in the mark of the cognitive.” But this rather begs the question of whether it is important to the science. In my view, cognitive science will proceed perfectly well without such enquiry, and in the course of so doing it may profoundly change philosophical conceptions of mind and ourselves. Amen to that!

Acknowledgements I am grateful for invitation to present these ideas and discussion that ensued from members of the audience at the conference “What is Cognition?” organized by Cameron Buckner, Ellen Fridland, Albert Newen, and Michael Pauen at the Ruhr University, Bochum, in 2013. I am also very grateful for the patience of the editors of this special issue, which I have tested to the extreme. I have also benefited greatly from the very patient and constructive criticisms of reviewers for this journal, including the most patient of all, reviewer number 3. Earlier drafts also benefitted from comments by students and colleagues at Indiana University. Any further deficits in this paper serve as evidence of my own failure to possess any mark of cognition.

References

- Adams, F., & Aizawa, K. (2001). *The bounds of cognition*. Oxford: Wiley-Blackwell.
 Adams, F., & Garrison, R. (2013). The mark of the cognitive. *Minds and Machines*, 23(3), 339–352.

- Allen, C., Grau, J. W., & Meagher, M. W. (2009). The lower bounds of cognition: What do spinal cords reveal? In J. Bickle (Ed.), *The Oxford handbook of philosophy of neuroscience* (pp. 129–142). Oxford: Oxford University Press.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1993). Problem solving and learning. *American Psychologist*, 48, 35–44.
- Angelucci, A., Clasca, F., Bricolo, E., Cramer, K. S., & Sur, M. (1997). Experimentally induced retinal projections to the ferret auditory thalamus: Development of clustered eye-specific patterns in a novel target. *The Journal of Neuroscience*, 17, 2040–2055.
- Aizawa, K. (2015). Cognition and behavior. *Synthese*, 1–20. doi:10.1007/s11229-014-0645-5.
- Baker, M. D., & Stock, J. B. (2007). Signal transduction: Networks and integrated circuits in bacterial cognition. *Current Biology*, 17, R1021–R1024.
- Beer, R. D., & Williams, P. L. (2015). Information processing and dynamics in minimally cognitive systems. *Cognitive Science*, 1, 1–38.
- Boag, Z. (2014). Interview of John Searle. *New Philosopher Magazine* (2). January 25, 2014. Retrieved from August 4, 2015. <http://www.newphilosopher.com/articles/john-searle-it-upsets-me-when-i-read-the-nonsense-written-by-my-contemporaries/>.
- Boyd, R. (1991). Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies*, 61, 127–148.
- Buckner, C. (2015). A property cluster theory of cognition. *Philosophical Psychology*, 28, 307–336. doi:10.1080/09515089.2013.843274.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: The MIT Press.
- Chemero, A., & Silberstein, M. (2008). After the philosophy of mind: Replacing scholasticism with science. *Philosophy of Science*, 75, 1–27.
- Clark, A. (2010). *Supersizing the mind: Embodiment, action, and cognitive extension*. New York: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 10–23.
- Cleland, C. E. (2012). Life without definitions. *Synthese*, 185(1), 125–144.
- Couzín, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1), 36–43.
- Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, MA: The MIT Press.
- Figdor, C. (2014). On the proper domain of psychological predicates. *Synthese*, 1–22. doi:10.1007/s11229-014-0603-2.
- Firestone, C., & Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for ‘top-down’ effects. *Behavioral and Brain Sciences*, 39, 1–77.
- Fuqua, W. C., Winans, S. C., & Greenberg, E. P. (1994). Quorum sensing in bacteria: The LuxR–LuxI family of cell density-responsive transcriptional regulators. *Journal of Bacteriology*, 176, 269–275.
- Garzón, P. C., & Keijzer, F. (2011). Plants: Adaptive behavior, root-brains, and minimal cognition. *Adaptive Behavior*, 19, 155–171.
- Hartmann, S. (1999). Models and stories in hadron physics. In M. S. Morgan & M. Morrison (Eds.), *Models as mediators. Ideas in context* (pp. 326–346). Cambridge: Cambridge University Press.
- Hills, T. T. (2003). Toward a unified theory of animal event timing. In W. H. Meck (Ed.), *Functional and neural mechanisms of interval timing* (pp. 77–111). Boca Raton, FL: CRC Press.
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2008). Search in external and internal spaces evidence for generalized cognitive search processes. *Psychological Science*, 19(8), 802–808.
- Hochner, B., Shomrat, T., & Fiorito, G. (2006). The octopus: A model for a comparative analysis of the evolution of learning and memory mechanisms. *The Biological Bulletin*, 210, 308–317.
- Iyer, L. M., Aravind, L., Coon, S. L., Klein, D. C., & Koonin, E. V. (2004). Evolution of cell-cell signaling in animals: Did late horizontal gene transfer from bacteria have a role? *Trends in Genetics*, 20, 292–299.
- Izquierdo, E. J., & Beer, R. D. (2013). Connecting a connectome to behavior: An ensemble of neuroanatomical models of *C. elegans* klinotaxis. *PLoS Computational Biology*, 9(2), e1002890. doi:10.1371/journal.pcbi.1002890.
- Kuryatov, A., Laube, B., Betz, H., & Kuhse, J. (1994). Mutational analysis of the glycine-binding site of the NMDA receptor: Structural similarity with bacterial amino acid-binding proteins. *Neuron*, 12, 1291–1300.
- Luce, R. D. (2003). Whatever happened to information theory in psychology? *Review of General Psychology*, 7(2), 183–188.
- Ludwig, K. (1996). Singular thought and the Cartesian theory of mind. *Nous*, 30(4), 434–460.

- Ludwig, K. (2015). Is distributed cognition group level cognition? *Journal of Social Ontology*, 1(2), 189–224.
- Machery, E. (2012). Why I stopped worrying about the definition of life.. and why you should as well. *Synthese*, 185(1), 145–164. doi:10.1007/s11229-011-9880-1.
- Macnab, R., & Koshland, D. (1972). The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences of the United States of America*, 69, 2509–2512.
- Nakagaki, T., Kobayashi, R., Nishiura, Y., & Ueda, T. (2004). Obtaining multiple separate food sources: Behavioural intelligence in the Physarum plasmodium. *Proceedings of Biological Science*, 271, 2305–2310.
- Newman, M. (2010). *Networks: An introduction*. New York: Oxford University Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Newton, J. R., & Sur, M. (2005). Rewiring cortex functional visual plasticity in the auditory cortex during development. In J. Syka & M. M. Merzenich (Eds.), *Plasticity of the central auditory system and processing of complex acoustic signals* (pp. 127–138). New York: Springer.
- Paoletti, P., & Neyton, J. (2007). NMDA receptor subunits: Function and pharmacology. *Current Opinion in Pharmacology*, 7(1), 39–47.
- Piccinini, G., & Scarantino, A. (2011). Information processing, computation, and cognition. *Journal of Biological Physics*, 37(1), 1–38. doi:10.1007/s10867-010-9195-3.
- Polger, T. W., & Shapiro, L. A. (2016). *The multiple realization book*. New York: Oxford University Press.
- Ramsey, W. D. (2007). *Representation reconsidered*. Cambridge University Press.
- Ramsey, W. D. (2015). Must cognition be representational? *Synthese*, 1–18. doi:10.1007/s11229-014-0644-6.
- Reid, C. R., Latty, T., Dussutour, A., & Beekman, M. (2012). Slime mold uses an externalized spatial “memory” to navigate in complex environments. *Proceedings of the National Academy of Sciences*, 109(43), 17490–17494.
- Rescorla, R. A. (1988). Pavlovian conditioning: It’s not what you think it is. *American Psychologist*, 43(3), 151.
- Rowlands, M. (2010). *The new science of the mind: From extended mind to embodied phenomenology*. Cambridge, MA: The MIT Press.
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 26, 521–562.
- Rupert, R. (2011). Review of Mark Rowlands, *The new science of the mind: From extended mind to embodied phenomenology*. *Notre Dame Philosophical Reviews*, 2011-03-35. Retrieved from August 3, 2015. <https://ndpr.nd.edu/news/24671-the-new-science-of-the-mind-from-extended-mind-to-embodied-phenomenology/>.
- Searle, J. (1992). *The rediscovery of mind*. Cambridge, MA: MIT Press.
- Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. R. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335, 108–111.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27, 379–423–623–656.
- Shannon, C. E. (1956). The bandwagon. *IRE Transactions-Information Theory*, 2(1), 3.
- Sharma, J., Angelucci, A., & Sur, M. (2000). Induction of visual orientation module in auditory cortex. *Nature*, 404, 841–847.
- Simon, H. A., & Newell, A. (1970). Human problem solving: The state of the theory in 1970. *American Psychologist*, 26(2), 145–159.
- Sporns, O. (2012). *Networks of the brain: Discovering the human connectome*. Cambridge, MA: The MIT Press.
- Srinivasan, M. V. (2010). Honey bees as a model for vision, perception, and cognition. *Annual Review of Entomology*, 55, 267–284.
- Stotz, K., & Allen, C. (2011). From cell-surface receptors to higher learning: A whole world of experience. In K. S. Plaisance & T. A. C. Reydon (Eds.), *Philosophy of behavioral biology, Boston studies in the philosophy of science* (pp. 85–123). Berlin: Springer.
- Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, 11(4), 378–395.
- van Fraassen, B. (2002). *The empirical stance*. New Haven, CT: Yale University Press.
- Villareal, L. (2004). Are viruses alive? *Scientific American*, 291(6), 96–102.

- Wimsatt, W. C. (1986). Forms of aggregativity. In M. G. Grene, A. Donagan, A. N. Perovich, & M. V. Wedin (Eds.), *Human nature and natural knowledge* (pp. 259–291). Dordrecht: Reidel.
- Wolf, D. M., Fontaine-Bodin, L., Bischofs, I., Price, G., Keasling, J., & Arkin, A. P. (2008). Memory in microbes: Quantifying history-dependent behavior in a bacterium. *PLoS ONE*, 3, e1700. doi:[10.1371/journal.pone.0001700](https://doi.org/10.1371/journal.pone.0001700).