Why can't we say what cognition is (at least for the time being)

Marco Facchin [**corresponding author**]
Istituto Universitario di Studi Superiori IUSS Pavia
Department of Human and Life Sciences; Pavia, Italy
Palazzo del Broletto, Piazza della Vittoria n. 15, 27100, Pavia
marco.facchin@iusspavia.it
https://orcid.org/0000-0001-5753-9873

Abstract:

Some philosophers search for the *mark of the cognitive*: a set of individually necessary and jointly sufficient conditions identifying all and only the instances of cognition. They claim the mark is necessary to answer difficult questions concerning the nature and distribution of cognition. Here, I will argue that, as things stand, given the current landscape of cognitive science, we are not able to identify a mark of the cognitive. I proceed as follows. First, I clarify some factors motivating the search for the mark of the cognitive, thereby highlighting the desiderata the mark is supposed to satisfy. Then, I highlight a tension in the literature over the mark. Given the literature, it is not clear whether the search aims for a mark capturing the intuitive notion of cognition or a genuine scientific kind. I then consider each option in turn, claiming that,either way, no mark satisfying the desiderata can be provided. I then deflect a foreseeable objection and highlight some implications of my view.

***Keywords*:** Mark of the cognitive, Cognition wars, Definition, Necessary conditions, Minimal cognition, Embodied cognition, Representation

***Wordcount*:** 9159, reference excluded

Submission Status *Currently Under Review*

Declarations:

Founding: Not applicable

Conflict of Interests/competing interests: The author declares no conflict of interests.

Availability of data and Material: Not applicable.

Code availability: Not applicable.

Authors' contribution: Marco Facchin is the sole author of the paper.

Acknowledgments: This paper owes much to some observations made in correspondence by Beate Krickel, and to some other observations made by Joel Dewhurst after my talk at the *Society for Metaphysics of Science* in 2021. Indeed, their observations motivated me to write this paper and I wish to thank them. I Wish also to thank Giacomo Zanotti, Marco Viola, Bruno Cortesi and Arianna Beghetto for their useful feedback on a number of previous drafts of this paper.

1 - Introduction

Some philosophers search for the *mark of the cognitive* (MOC). They aim at providing a set of individually necessary and jointly sufficient conditions whose satisfaction identifies all and only the instances of cognition (e.g. Adams & Aizawa 2001; Rowlands 2009; Adams 2019). They claim the MOC is necessary to allow cognitive science to develop correctly. Should cognitive science investigate distributed brain-body-world systems, as argued by the extended mind thesis (cf. Walter 2010; Wheeler 2011; 2019)? Clearly the answer to this question depends on whether brain-body-world systems qualify as cognitive systems. Should botany and microbiology be considered as parts of cognitive science? This depends on whether plants and bacteria qualify as cognitive systems too (cf. Adams 2010; 2018). And to know whether these things qualify, we need to know the MOC. Other philosophers disagree: cognitive science can, and will, successfully develop without the explicit guidance of a MOC (Clark 2008; Allen 2017).

Here, I argue that we should hope that philosophers of the second group are right - for, *at least at present*, no MOC can be provided. I will argue as follows. In §2, I examine the literature searching for the MOC. In §2.1, I discuss some reasons motivating the search, thereby identifying some *desiderata* the MOC should satisfy. In §2.2, I highlight a tension present in that literature: whereas the reasons motivating the search suggest the MOC should capture a scientific (most likely, natural) kind, the role intuitions play in the search suggest that philosophers are actually trying to capture our *intuitive notion* of cognition. I propose to tease the two projects (defining a scientific kind *vs* defining our intuitive notion) apart, to then argue that, *as things stand*, both projects are bound to fail. In §3, I claim that our intuitive notion of cognition (if it exists), cannot be captured by a MOC. I also claim that even if such a notion were to be captured by a MOC, it would not satisfy the *desiderata* motivating the search. The project of finding a MOC capturing our intuitive notion of cognition should thus be abandoned. In §4, I claim that, *as things stand*, we cannot provide a MOC capturing cognition as a scientific kind. This is because cognitive science is splintered in a number of competing, but equally legitimate, research traditions, each suggesting a MOC (even if implicitly). Since they compete, they jointly fail to specify a coherent MOC. But since they're also equally legitimate, we have no grounds to privilege a specific proposed MOC over any other. §5 concludes the essay considering an objection and providing further clarifications to my position.

2 - Searching for the Mark of Cognition

2.1 - Some *desiderata*...

Everyone agrees we can safely distinguish paradigmatic instances of cognition (e.g. recalling something) from non paradigmatic ones (e.g. sneezing; cf. Adams 2019). So, *why* should we bother searching the MOC? Answering is important: knowing *why* we search for a MOC tells us what we'd like it to accomplish, suggesting the *desiderata* the MOC should satisfy.

Inspired by Akagi (2016; 2018) I identify *three* motivations fueling the search - and so three *desiderata*. Notice: I do not, and need not, claim these three reasons are *all* the reasons motivating the search. The same goes for the *desiderata*. Here I provide only what is necessary for my argument. The list may be incomplete, and the door is open for further elaboration.

Motivation #1: as a whole, cognitive science has expanded, and partially shifted, its focus away "higher thought", towards skilled sensorimotor interactions (cf. Dennett 1978; Clark 2001). Early in its development, cognitive science was mainly interested in "high level", perhaps exclusively human, phenomena. Early AI researchers, for example, were interested in making computers able to play checkers (e.g. Samuel 1959). And they were sure that their procedures held the keys to *thought*: indeed, they thought that computers in the '60s were able to think, but unable to perceive and act swiftly and skillfully (cf. Selfridge and Neisser 1960). The *zeitgeist* seems now inverted: AI researchers focus on sensorimotor interactions (e.g. Tani 2016), and the consensus seems to be that whilst computers might perceive like us, they *definitely* don't think like us (Mitchell 2019). And whilst "higher thought" is still an *explanandum* of cognitive science, the emphasis often is now placed on its sensorimotor roots - for instance highlighting the number of ways in which the cortical structures for "higher thought" depend on the ones in charge of our sensorimotor couplings (cf. Barsalou 1999; Dehaene and Cohen 2007; Anderson 2014; Cisek & Hayden 2022). Sensorimotor interactions precede "higher cognition" not just in phylogeny and ontogeny, but now also in the order of explanation.

Such a shift in focus can easily generate worries in the distribution of cognition. Only humans (and some computers) play checkers. Only humans (and *perhaps* some computers) understand natural languages. If *these* are the central cases of cognition, then cognizers are relatively few: a restricted set of mammals, maybe some computers. In contrast, if the central cases of cognition consist in some sensorimotor interaction, the number of cognizers is way bigger: all multicellular animals, and arguably some simple robots (cf. Braitenberg 1984), plants (Calvo Garzon 2007) and single celled organisms (Lyon 2015). Maybe even some planetary scale processes could be construed as cognitive processes (Frank *et al*. 2022). So, who's in? *Which* (kind of) systems should cognitive science study? The MOC should enable us to answer. It should give us an *extensionally adequate* definition of cognition. Hence the first desideratum.

> ***Desideratum #1***: The MOC should be an extensionally adequate definition of cognition: i.e. a set of individually necessary and jointly sufficient conditions the satisfaction of which identifies all and only cognitive systems/states/processes

This seems an important desideratum, whose centrality is greatly emphasized in the literature concerning extended cognition (Adams and Aizawa 2001; 2008; Adams 2010; 2019; Rowlands 2009; 2010).

Motivation #2: cognitive science is *extremely* fragmented. Not only the paradigmatic *explananda* of cognitive science have changed, the *explanantia* have changed too, and dramatically so. Yet, "change" might not be the right word - it might

suggest a *gradual maturation*. But, looking at the history of cognitive science, one has a hard time seeing a *gradual* maturation. What one sees is the *splintering* of a (relatively well defined) research tradition into a myriad of *different and competing* research traditions, each rhetorically presenting itself as a "kuhnian revolution" able to supplant all other research traditions and let cognitive science run free from the shackles of ignorance.[1] Notice that here I'm using "research tradition" technically, to name a (fairly well-defined) set of theoretical assumptions, modeling tools, experimental procedures and other research practices a group of scientists use to investigate a set of phenomena of interest (cf. Laudan 1977: 81).

Here's a - simplified, incomplete, and doubtlessly whiggish[2] - rendition of how cognitive science historically developed. It all began in the '50s with the *cognitive revolution*: a multi-, or even inter-, disciplinary enterprise guided by an operative definition of cognition as *symbolic* (digital) *computation* (cf. Newell and Simon 1976). Then the connectionist revolution: connectionists proposed new and different computational models loosely inspired by cortical architecture (Rumelhart *et al*. 1986), inadvertently redefining cognition as *subsymbolic* computation (cf. Churchland 1992). As these models grew in complexity, cognitive scientists discovered they were often better off using a different branch of math to deal with them; namely *dynamical system theory*. Hence the dynamicist revolution: computation faded in the background while cognition became the swirl of activity of a self-organizing system (cf. Van Gelder 1995; Thelen and Smith 1999). A system which might, but *need not* be identical with the brain, as the "4E" revolution quickly claimed (e.g. Clark 1997). The discovery that such a swirl of activity is a form of Bayesian inference (Friston 2012), caused another revolution, accompanied by an appropriate redefinition of cognition as inferential prediction (Corcoran *et al*. 2020; Kiverstein & Sims 2021). In parallel, "old" computational ideas have been revamped by the cognitive neuroscience revolution (Boone and Piccinini 2016).

The above are all different, competing and equally legitimate research traditions. They are *different*, for they all endorse different sets of theoretical assumptions, use different models and modeling techniques, and resort to different explanatory strategies abiding to different explanatory standards (cf. Lamb and Chemero 2018; Piccinini 2020). They *compete*, for they aim (or, at least, publicly declare to aim) at explaining the same thing - namely cognition.[3] And they are *all equally legitimate*, at least to the extent that none of them is obviously false and they are all able to generate result counting as genuine progress within the boundaries of the tradition. But which is right? The MOC should help us answer. By telling us

[1] As Steiner (2021) convincingly argues, that the "khunian rhetoric" does not capture the relevant conceptual changes in cognitive science. In his view, cognitive science (or, at least, the passage from "classic" to "embodied" cognitive science) is best described as the shift in the balance of power between two longstanding and competing research traditions. I agree, and I think the analysis should be expanded to *all* the supposed "revolutions" in cognitive science.

[2] For a *real* history of cognitive science, see (Boden 2008).

[3] Admittedly, there have been some calls for integration (e.g. Eliasmith 2013, ch. 1) and/or statement that different approaches may be complementary (Kaplan and Bechtel 2011). But these are not just few and far between, they also exhibit some bias towards some privileged research tradition and/or model of explanation. Thus, for example, Eliasmith still suggests a sub-symbolic *cognitive* level below the symbolic one, and Kaplan and Bechtel still think that the explanatory power of dynamical models depends on them being "mappable" on mechanistic explanations.

what cognition is, it should identify a scientific (perhaps natural) kind supporting relevant scientific generalizations and principles. (Adams and Aizawa 2001; 2008; Buckner 2015; Newen 2017). And by so doing it will indicate us which research tradition(s) to pursue and which to abandon:

> ***Desideratum #2***: The MOC should have enough explanatory power to allow us to identify which research tradition(s) are worth pursuing in the study of cognition

Motivation #3: disciplinary boundary disputes. All the research traditions mentioned above agree in construing cognitive science as a multi-, or even inter-, disciplinary enterprise. But the agreement stops here. *Which* disciplines should be allowed to take part in the enterprise in a hotly debated manner. Sure, "classic" cognitive science had some clear ideas - icastically represented in the "cognitive exagon". Cognitive science was construed as an interdisciplinary enterprise animated by philosophy, psychology, neuroscience, linguistics, anthropology and computer science/A.I. (Gardner 1987). But, to put it mildly, the hexagon was a *very idealized* description of the endeavor, or perhaps had an aspirational role. For, anthropology has contributed comparatively little to classical cognitive science. And classical cognitive science was not exactly *keen on neuroscience*: indeed, "classicism" is often mocked as the view that the best way to study the *mind* is to systematically ignore *the brain* (cf Fodor 1999).

Moreover, as "classic" cognitive science splintered in the variety of competing research traditions briefly indicated above, new disciplines were put in contact with, and included in, the forming research traditions. These include: *engineering* (Pfeiffer and Bongard 2007), material science (McGivern 2019; Tripaldi, *forthcoming*), physics and complex system science (Hacken 1978; Kelso 1995), plant biology (Calvo Garzon 2007), microbiology (Yakura 2019) archeology (Malafouris 2013) and more. Are *all* of them rightful contributors to cognitive science? This question is important to answer for at least two reasons. First, the regularities and generalization about cognition cognitive science will discover depend largely on which individual disciplines constitute it. The more inclusively one construes cognitive science, the more these principles and generalizations will tend to be minimal and behavior based (e.g. Lyon 2006; Sims 2021). The more one construes cognitive science demandingly, the more the principles and generalization will be demanding and concept-focused (e.g. Adams and Aizawa 2008; Adams 2016). Secondly, it is important to determine the disciplinary boundaries of cognitive science to use our material and intellectual resources correctly. To make the point bluntly: if microbiology was "in", we should create reliable informational channels connecting microbiologists to psychologists and linguists/psycholinguists, allowing them to share ideas, models, methods of inquiry and results. This isn't easily done, nor is it something that can be done *for free*. It will require intellectual elaboration and monetary funds. These are limited resources, which we shouldn't waste. Hence the third desideratum:

> ***Desideratum #3***: The MOC should determine the disciplinary boundaries of cognitive

> science, allowing us to allot our intellectual and non-intellectual resources in an appropriate manner[4]

Few words about these *desiderata*. First, as said above, I don't think the list just given is complete. There may be more relevant *desiderata*. But they won't matter here, so I'll stay silent on those. Secondly, I take these *desiderata* to function as a metric determining which candidate MOC(s) we should *at least pro tempore* accept (or keep considering) and which we must flat out reject. Thirdly, and relatedly, I don't take the satisfaction of these *desiderata* as an all-or-nothing affair. It might be possible, for example, that when comparing two candidate MOCs *a* and *b* the former will turn out to have more explanatory power than the latter.[5] Lastly, notice that all these desiderata indicate that the MOC should capture a genuine *scientific kind*; that is, a kind supporting the genuine generalization and principles of a science of cognition. Thus, "cognition" should be a theoretical term with some explanatory purchase in the sciences of cognition. It should be a term like "strontium" in chemistry, or "ribosome" in biology. In fact, extensional adequacy, explanatory power and the capacity of defining the boundaries of a scientific endeavor seem all properties of a theoretical term naming a genuine scientific kind (see also Adams and Aizawa 2001; Wheeler 2011; 2019). We *wouldn't* expect "folksy" terms (such as, say, "animal") to have such features.

And yet, "the shadow of the folk" looms large on the MOC, generating a tension.

2.2 ... and a tension

To feel the tension, consider the following three, importantly related and interconnected, features of the search for the MOC.

Feature #1: the appeal to (more or less commonsensical) intuitions[6] is rampant (cf. Elpidorou 2014). Examples abound. Bermudez (2010: 415) and Shapiro (2013: 363) simply assert, basically without argument, that cognition *must* involve representations, stating they cannot *see* how it could be otherwise. Adams and Garrison (2013) do the exact same thing when they take personal-level reasons to be *necessary* for cognition, and offer no argument in support of their claim. Similarly, Aizawa (2017: 16) claims without argument that cognition *must* in a sense be centrally unified; that is, that a cognitive agent *cannot* be built out of the interaction of special purpose mechanisms.

Notice what it is happening: philosophers are indicating part of the MOC - namely, an *individually necessary* condition - based on their intuitions. This typically isn't how we go about searching for scientific kinds. Indeed, our intuitions often *stood in the way* of us discovering genuine scientific kinds. Our intuitions clumped together jadeite and nephrite as Jade. Dante's

---

[4] This might be partially redundant in respect to the second *desideratum*: knowing which research tradition is right will most likely tell us which disciplines constitute cognitive science and how to allocate resources. Yet it was worth making the point explicitly.

[5] This suggests that it could be possible to build a "MOC adequacy scale" by identifying all the relevant desiderata and using them to rank candidate MOCs. This may be an interesting project to pursue somewhere else.

[6] I will adopt a very unsophisticated view of intuitions: they are judgments we're prone to make and report (if asked). Thus, for example, most westerners today share the intuition that the Earth revolves around the Sun, but very few westerners shared that intuition before the "scientific revolution".

intuition told him that the sun is a planet (*Inferno*, canto I) and that to move upwards from the center of the Earth one must turn 180° (*Inferno,* canto XXXIV). Alchemists found it compelling to think that nitric acid and hydrochloride were species of water (called *Aqua Regalia* and *Aqua Fortis*). Yet, when it comes to cognition, the care these examples invite seems to get thrown out of the window.

Now, one could perhaps adopt a broadly hermeneutical standpoint, arguing naive intuitions (and all sorts of biases) are *always* informing our scientific practice. They lay in the background, silently skewing our research in certain directions. True. Yet notice the intuitions above do *not* lurk in the background. They are stated in the main text of the papers. Their influence is upfront and direct. For they play an *essential* role in the philosophical literature on the MOC, which brings us to the second feature.

Feature #2: intuitions often counteract the influence of legitimate and well established research traditions. To continue with the examples above: Barmudez and Shapiro deem representation necessary *despite* the presence and successes of anti-representationalist research traditions in cognitive science (e.g. Beer 1995; 2000). Aizawa takes a "central processor" to be necessary, *despite* the successes of the massive modularity research tradition (cf. Charrutes 2006). Adams and Garrison's case is even more puzzling: it seems to me that *no* research tradition in cognitive science even *mentions* personal-level reasons! Examples proliferate easily: as Chemero (2009: ch. 1) notices, arguments of that sort are fairly common in cognitive science, and indeed pre-date the whole debate on the MOC. Thus Fodor and Pylyshyn (1988), finding it intuitive that all cognition must be systematic, claimed that artificial neural networks are at best how-possibly models depicting the *implementation* of (independently studied and characterized) cognitive capacities. Earlier still, Searle (1980) purported to show the untenability of an entire research tradition with a *thought experiment*; that is, appealing to our intuitive reactions to an imaginary scenario. Closer to us, the "dark room" argument against predictive processing views of cognition is based on the intuitive idea that, if all our brain tries to do is to predict the incoming inputs as accurately as possible, our brains (and thus *us*) should crave very predictable and boring environments. But we clearly don't crave them, so predictive processing is wrong (Sims 2017; Ryan *et al* 2021 for discussion).

Notice that intuitions typically *do not* play this role when it comes to examine the merits, and eventually attack, research traditions.[7] Nor intuitions typically play this role when it comes to capturing scientific kinds. When it comes to these matters, we (typically) do not take intuitions and the outputs of scientific research traditions to be epistemic peers.

Feature #3: The *pertinence* of certain scientific findings is openly contested Notice: *the findings themselves* are uncontested. No

[7] To get a sense of the atypicality, compare the positive reception of the Chinese Room argument against classical cognitivism to the *extremely negative reception* of Fodor and Piattelli-Palmarini's (2011) arguments against adaptationism (cf. Godfrey-Smith 2010; Block and Kitcher 2010).

one claims that a certain experiment never happened, or that such-and-such an observation was not really made, or that certain data have been "rigged" to favor a specific research tradition. To the contrary, the purely factual aspect of discoveries and findings is left uncontested. What is contested is that such discoveries and findings give us insight about what cognition is and/or how it operates. Bluntly: it is contested *whether such findings matter to understand cognition*.

Some examples to clarify. If Fodor and Pylyshyn (1988) are right, then we can happily ignore artificial neural networks and other such neurocomputational models, for these models won't tell us anything relevant about *what cognitive processes are*. They will, at best, tell us how (independently studied and characterized) cognitive processes are *implemented*. Similarly, if inspired by Searle (1980; 1983) one concludes *non-derived content* is a necessary ingredient of cognition, one can outright dismiss lines on research concerning human-artifact interactions, including research on sensory substitution devices (Bach-Y-Rita and Kercell 2003; Loomis *et al*. 2012) or the usage of external tools in problem solving (Risiko and Gilbert 2016; Bocanegra *et al*. 2019). Presumably, neither sensory substitution devices nor external tools host non-derived content. So they *can't* tell us anything relevant about cognition. Their status *as evidence* concerning the nature of cognition is nullified. And typically this happens in virtue of some necessary condition imposed by intuition.

Now, unless one thinks that our intuition is tailored to capture scientific kinds (and there are good reasons *not* to believe this, see Churchland 1987; Aikins 1996), we should regard these features as generating an important tension in the project of searching for the MOC. On the one hand, the motivation for the search, and thus the *desiderata* the MOC is called to satisfy, suggest that the MOC should define a technical term capturing a scientific kind. On the other hand, the widespread to intuition and the fact that intuition are taken to have the same epistemic standing of scientific result, to the point that they can challenge their *evidential* status, suggest the MOC is aimed to capture something different; namely what we'd normally call "cognition" in our everyday lives. The MOC would thus elaborate upon, and make explicit, an important piece of our "manifest image".

Whilst both legitimate, the two projects are clearly regulated by different epistemic norms and standards.[8] I thus propose to disentangle them, and consider them separately. So, how are the prospects of these projects?

3 - Everyday cognition

Consider first the project of providing a MOC capturing the "everyday" notion of cognition. Such a MOC aims to define a *folk notion* - i.e., it aims to define what the layperson thinks cognition is. The prospects of this project appear extremely grim.

First: *who* is the folk? "the folk" is an abstraction. People are different, and intuitions vary. It's at best *naive* to think there is *one* folk conception of cognition everyone on Earth shares. Indeed, some empirical data suggest that the folk conception of

[8] Scarantino (2012) seems to notice that two similar projects are also often entangled in affective science.

cognition changes across cultures. Trovato and Eyseel (2017) for example, provide data indicating that Italian and Japanese highschool students attribute mental property to androids differently. Now, not all the mental properties in the study were cognitive properties, *but some were* (for example, the capacity to plan and act accordingly). And the attribution of these properties revealed significant differences too. Further, when probed via the *IDAQ* (individual differences animism questionnaire), Italian and Japanese students show different *patterns* of attribution of mental properties (among which cognitive properties): Italians are more prone to ascribe them to artifacts, whereas Japanese are more prone to ascribe them to animals and natural phenomena. On the fairly uncontroversial assumption that high school students' ascriptions are "folk" ascriptions, the data suggest that the folk concept of cognition differs cross-culturally.

So maybe it is prudent to consider "folks" having a common cultural background. But even *within* a culture it is doubtful there is *a single shared* conception of cognition. Well educated westerners define cognition in a variety of different ways (cf Bayne *et al*. 2018 for a sample). Swiss people are *very* divided on what counts as cognitive: 44% of Swiss think robots are *genuinely intelligent* and 56% think they are not (Arras and Cerqui 2005). Almost a 50/50 split.

Of course, it is always possible to claim that, were a more fine grained subdivision considered (e.g. or a single city, or a single quarter) a single and unified MOC would be found. Surely, at the level of the smallest subdivision possible (i.e. the single individual) we *will* find a single and unified MOC - one's concept of cognition. So, the project of providing a MOC capturing our "folksy and intuitive" notion of cognition need not be abandoned. It's just that we have to spell out more MOCs than expected; maybe even 7.9 billions more.

Yet, there are I think good reasons not to hold an individual's folksy and intuitive conception of cognition might not constitute a MOC in any reasonable sense of the term. As hinted at above, MOCs are sets of individually necessary and jointly sufficient conditions. And the appeal to intuition is often used to impose necessary conditions. Yet, there are good reasons to think that our "folksy and naive" concept of cognition does not prescribe *individually necessary* conditions. For one thing, attempts at capturing ordinary concepts by sets of individually necessary and jointly sufficient conditions have traditionally been crippled by counterexamples (cf. Fodor 1981). And whilst *some* concepts can be spelled out in that way (e.g. *x* is a triangle *iff x* has exactly three sides and exactly three edges) it is typically easy to do so, and the definitions provided are uncontested (cf Machery 2012). Surely this isn't the case of "cognition".

Further, our psychological theories of concepts suggest that "folksy and intuitive" concepts can hardly be *adequately* captured by sets of individually necessary and jointly sufficient conditions, as they do not seem to include any individually necessary condition (cf. Machery 2009 Ch. 4). According to *prototype theory* (Rosch and Mervis 1975) concepts are statistical centroids: representations encoding the *statistically typical* features of a class of items. But, of course, since typical features

are in no way *necessary* (birds typically fly and sing, but ostrich do neither), no individual feature is necessary, and so the possession of any individual feature does not constitute a necessary condition for the inclusion in that class of item. What matters is possessing a sufficiently large number of features, rather than any individually necessary one. According to *exemplar theory* (Medin and Schaffer 1978) "folk concepts" consist in representations of individual members of a class, and, again, in order to be included in a class what matters is not the possession of necessary features, but rather the overall similarity with the exemplars. Others still conceive concepts as somewhat analogous to scientific theories, which encode some knowledge about the causal and/or functional regularities that tie together different properties. On this view, categorization is basically a form of inference to the best explanation: an object is categorized as a member of a class just in case belonging to that class best explains its properties. Again, no individual property or feature is deemed necessary (cf Murphy and Medin 1978). Now, my point is *not* that one of these theories is definitely right and so we cannot capture our "folksy and intuitive" conception of cognition in a MOC. My point is that, currently, our best theories of concepts *en masse* strongly suggest that that is the case, and it is advisable to trust them.

But perhaps such theories are radically misleading, and possessing the *true theory* of concept will support the idea a MOC capturing our "folksy and intuitive" conception of cognition can be provided. Perhaps. Yet, unless one is able to point at such *true theory*, the objection is vacuous, and points at a *mere* logical possibility. Worse still, even if such a "true theory" was provided there would *still* be reasons to give up the search for the MOC capturing our "folksy and intuitive" conception of cognition. For it would fail to satisfy *desiderata* #1 to #3.

As noticed above, our folksy intuitions are not tailored to the discovery of scientific kinds and the definition of technical terms. Not even when it comes to *mental*, *psychological* or *cognitive* terms and kinds. Examples of this come by easily. We "folksy conceive" vision as a single process, while it most likely consists of at least two different sets of processes (Milner and Goodale 1992). Our folksy kind "memory" has been subdivided in a myriad of ways (working memory, semantic memory, procedural memory, long term memory, *etc*). And, more generally, it is not clear whether our folksy psychological categories provide a solid foundation for the sciences of the mind (Buzaki 2019: ch. 1; Pessoa *et al*. 2022). Indeed, if our "folksy intuitive" conceptions about the mind weren't wrong, it would be very hard to make sense of the history of psychology - why did it take so long to become a *real* science?

So, providing a MOC capturing satisfying *desiderata #1* to *#3* cannot be the reason why we should provide a MOC capturing the folksy concept of cognition. But then what reasons are left for such a search? The search for the MOC was (partially) motivated by concerns regarding the extension of "cognition" and how to study it. Indeed, even before the search for the MOC "officially" started (in Adams and Aizawa 2001), the attempt to identify *essential* features of cognition (i.e.

individually necessary conditions partially constituting the MOC) was driven by the development of alternative research traditions in cognitive science (cf. Akagi 2018). So it seems that the reasons *driving* the search are importantly connected with the troubled development of cognitive science, and the current absence of a central scientific kind *cognition* driving the enterprise (cf. Rupert 2013). Providing a definition capturing our commonsense intuitions about cognition would do little to nothing to meet the goals of the research. So, why bother?

Perhaps I've been unfair. *Of course* providing a MOC capturing our commonsensical notion of cognition would do little to aid cognitive science. Haven't I conceded that much at the end of §2.2, when I suggested there are two distinct projects searching for the MOC? Yes, I did. I proposed to *disentangle* two different projects; the first aimed at a MOC spelling out a scientific kind, the second aimed at a MOC spelling out our folksy intuitions. But, as things stand, this second project seems to me entirely unmotivated. Indeed, as I noticed in that very same section, intuitions are currently brought to bear when it comes to finding a MOC of the first kind. This is indeed *why* the two projects are entangled. And even the philosophers who bring these intuitions to bear take themselves to be searching for a MOC of the *first* kind. Thus, for example, Adams and Garrison wrote about the MOC to safe cognitive *science* from the embarrassment of not knowing what cognition is (Adams and Garrison 2013: 340), Searle (1980) proposed his thought experiment to identify the *right* research tradition for artificial intelligence (i.e. weak AI), Fodor and Pylyshyn (1988) were interested in determining the role of neurocomputational models and Aizawa (2017) aimed at evaluating "4E" cognitive science. None of them seems particularly interested in spelling out our folksy intuition of cognition. Maybe I'm simply blind to some important reasons that may motivate such a search. And while at present I know of no such reasons, that does not mean that such reasons are altogether missing. Yet, until they are provided, it seems we have just no reason to search for a MOC capturing our "folksy" notion of cognition.

Summarizing: it is not clear whether there is a *single* folk notion of cognition - that notion may exhibit significant cultural variability. And even if there were such a notion, our best theories of concepts suggest it wouldn't provide us with a *mark* of the cognitive; that is, it wouldn't spell out a set of individually necessary and jointly sufficient conditions. And even if such a mark *were* to be spelled out, it is unclear (to put it mildly) that it would meet the relevant *desiderata* motivating the search. And, at present, there seems to be no *other desiderata* that would motivate it. So, at present, the prospects of finding a MOC capturing our folk notion of cognition aren't rosy. But what about a MOC capturing a scientific kind?

4 - Scientific cognition

Suppose the MOC is best conceived as a technical term whose definition captures a scientific kind. The definition cannot be *stipulative*. We can stipulate the meaning of terms, deciding, say, that *x* is a feline *iff x* is a piece of furniture, built before

1947, and weights less than 3 Kg. This definition, whilst extremely idiosyncratic, is perfectly fine, and satisfied by at least some pieces of furniture (which are thereby revealed to be felines in a sense of the term).[9] Yet the MOC *cannot* be like that. For a stipulative definition fails to meet the relevant *desiderata*. Surely wouldn't have any explanatory power (*desideratum #2*), nor would it help us determine which disciplines are part of the multidisciplinary enterprise called "cognitive science" (*desideratum #3*). And it would not help us determine the extension of "cognition" (*desideratum #1*). Under this construal of the MOC, "cognition" names a scientific kind, and presumably a *natural* kind.[10] These kinds are typically not up for stipulation. They're discovered, rather than created.

Now, if we should discover what cognition is, a good way to discover it is *via* a dedicated scientific endeavor; namely cognitive science. And here lies the rub: we'd like the relevant MOC to come out of cognitive science, but cognitive science is in a state of disarray, splintered in many competing research traditions at least implicitly suggesting different MOCs. And to bring cognitive science out of this state of disarray we need a MOC. Indeed, it is *because* cognitive science is in disarray that we're looking for the MOC (§2.1).

Notice: it is *crucial* the various research traditions cognitive science is currently splintered in are not just competing, but also implicitly defining *mutually exclusive* MOCs. For this prevents us from adopting a form of "happy pluralism" according to which cognition itself is so multifaceted and complex that each of these candidate MOCs is partially correct. Cognition, as complex as it may be, cannot have *contradictory* properties. Yet the MOCs suggested by these research traditions ascribe it contradictory properties. The methodological solipsist (Fodor 1980) proposes a MOC that cannot peacefully coexist with the one suggested by the ecological psychologist (Chemero 2009), the enactivist (Hurley 2001), or even an externalistically minded connectionist (Clark 1993). Otherwise put: some candidate MOCs are *mutually exclusive*. They cannot be accepted together.

One might object we *could* accept all these candidate MOCs together without giving rise to a contradiction. It is sufficient to knit them together *via inclusive disjunctions*: cognition is as the solipsist describes, or as the enactivist describes, or as the connectionist describes, etc. But this falls short of the relevant *desiderata*. Maybe this procedure could yield us the true extension of cognition as *desideratum #1* wishes. But it surely *won't* reveal us which research tradition(s) is(are) worth pursuing (*desidertaum #2*). And, arguably, it tells us little (if anything) about the disciplinary boundaries of cognitive science (*desideratum #3*). It's hard to see how it could be used to determine, say, whether microbiology or hematology are

[9] Whilst, to my knowledge, *no one* has ever stipulated a MOC in this quite explicit fashion, some such stipulation might have been proposed. Keijizer's (2021) suggestion to substitute the study of cognition with the study of a newly defined kind called *cobolism* seems to qualify. And, if Lyons (2020) is right, "minimal cognition" is a term introduced by stipulation too.

[10] Whilst even dualists concede that cognitive processes are fully natural (cf. Chalmers 1996), it's not uncontested that cognitive science is a natural science. For example, the European Council considers cognitive science a social science. So maybe *its* kinds are not natural after all.

parts of cognitive science.[11]

Once could contend that a MOC created via inclusive disjunctions as hinted above actually *satisfies* the desiderata. It satisfies *desideratum #1* because it gives us the true extension of cognition: everything cognitive scientists study. It satisfies *desideratum #2*, because it tells us which research traditions to pursue: namely, all of them. And it satisfies *desideratum #3* because it tells us the disciplinary boundary of cognitive science: these boundaries include all disciplines one might use to study cognition. Yet, it seems to me that arguing in this way leads the MOC searcher to a pyrrhic victory (at best). After all, the MOC thus provided *makes no difference* to the current state of cognitive science. It delivers the boons the MOC should bring about in name only.

Thus, we cannot hold all candidate MOCs together; *we must choose*. How? Choosing arbitrarily would amount to *stipulating* a MOC. So, we need some *reason in principle* to choose a MOC (or at least few mutually consistent ones)

It's tempting to think such reasons might be extrapolated or derived from the relevant empirical evidence. That's how Einstein prevailed on Newton. Why can't Gibson prevail on Gregory the same way? Yet, it is hard to see *how* empirical evidence could decide for one of the many research traditions (and associated MOC) over any other. For, as highlighted in §2.2, *what counts as evidence* is in itself contested. And whilst in §2.2 I've pointed out that what counts as evidence is contested based on intuitions, these intuitions are often closely aligned to fairly specific and recognizable research traditions.

Moreover, often the relevant evidence is very hard to interpret. Consider one of the coarsest divisions in cognitive science; namely the one between representationalist and anti-representationalist research traditions. The former claim cognition *requires* representations; the latter claim it doesn't. Some think the debate can be easily solved by looking at the evidence: can we explain all *prima facie* intelligent behaviors *just* in stimulus-response terms? If yes, anti-representationalism wins the day, otherwise, it loses (Churchland 2002). Alas, things are not even remotely so simple. For anti-representationalists *can* (and *do*) hold that the machinery internal to the organism matters to cognition. It's just that such machinery is not in the business of representing (Chemero 2009). And it seems *very hard*, if not impossible, to "take a peek inside" and simply see whether there are representations or not. For, paradigmatically non-representational systems *can* be easily interpreted in representational terms (Bechtel 1998; Shapiro 2013), and, equally easily, paradigmatic cases of representation can be "delfated" and accounted for in non representational terms (Ramsey 2007; Facchin 2021). The same holds if instead of toy models we look directly at the brain (cf. Kriegskorte and Kevit 2013; Ritchie *et al*. 2019; Gessel *et al*. 2021). Both camps

[11] Notice that the MOC proposed by Akagi (2021) is something of this sort: it consists of a rigid structure of interconnected variables, each of which is able to assume a range of values. Each variable represents a *locus* of contention in regards to the definition of cognition, and each value represents a position actively engaged in the dispute. But, as Akagi notes, such a MOC does not tell us what cognition *is*, and so it does little to satisfy the desiderata identified above. What such a MOC does is to capture in an orderly manner the extent of the disagreement concerning cognition. And that is Akagi's primary purpose.

can, without deceit, look at the same piece of evidence and conclude it supports them. It is thus hard to see how scientific evidence could provide us with a principled ground for our choice.

Maybe, then, clever reasoning will succeed where the appeal to evidence fails. We could design sophisticated arguments showing that one, or more, research tradition(s) ought to be abandoned. Chomsky managed to provide one such argument against behaviorism, and there seems to be no reason as to why, say, enactivism should be immune to such arguments. So we could search the MOC by *elimination*: narrowing down the set of research traditions (and thus candidate MOCs) down to one, of few mutually consistent ones.

Though this way of proceeding is viable *in principle*, it likely won't be viable *in practice*. Even Chomsky's famous arguments against behaviorism failed to force a wholesale abandonment of behaviorism (cf. Staddon 2014). Minsky and Papert's (1969) analysis, whilst rigorous and on the point, (thankfully) failed to force a wholesale abandonment of connectionism. And the arguments offered by Chomsky, Minky and Papert are not just strong and well constructed: they are (and have been) *persuasive*. They impacted the day-to-day research practice of numerous cognitive scientists, and left a *sizable* mark on cognitive science. Most other arguments aimed at motivating the abandonment of a specific research tradition are neither as strong nor as persuasive as these ones (cf. Chemero 2009, Ch. 1). This suggests such a process of elimination is very hard, if not impossible, to translate into practice.

And even if it were translated into practice, *it might not be translated successfully*. For, even *within* a single research tradition, various disciplines (and their philosophical spokesperson) might propose and defend different candidate MOCs. The point is nicely exemplified by the exchange between (Corcoran *et al*. 2020) and (Kiverstein and Sims 2021). Both candidate MOC "came out" the same research tradition; namely *Active Inference*.[12] According to this tradition, cognition is best studied deploying a complex set of modeling tools (cf Andrews 2021) allowing us to *construe* cognitive activity as a self-organizing process whereby a system brings about sensory states consistent with (and confirming) its own prolonged existence through time. Whilst Corcoran and colleagues and Kiverstein and Sims agree on that much, they still not only propose different MOCs, but also MOCs whose difference *matters* given the *desiderata* highlighted above. In short, Corcoran and colleagues suggest that cognition *is* a rather sophisticated form of counterfactual inference which is *not* universally possessed by living systems. Conversely, Kiverstein and Sims suggest that cognition *is* a form of anticipation all living systems exhibit. So, they disagree concerning the extension of "cognition" (*Desideratum #1*). They also disagree on the disciplinary boundaries of cognitive science (*Desideratum #3*): whereas Kiverstein and Sims suggest that biological sciences are *en masse* part of cognitive science, Corcoran and colleagues resist the suggestion.

[12] Or, at a coarser level of analysis, "Predictive Processing" (cf. Clark 2016). The point is not really relevant here.

So, there are reasons to believe neither "the evidence" nor reasoning will provide us the MOC. Why don't we ask history? Let the research traditions compete. Some will blossom, some won't, and at the end they will deliver a coherent picture of cognition.

Whilst sympathetic to this wait-and-see strategy, the MOC searcher should be wary about it. For it suggests either postponing or abandoning (depending on how cognitive science will develop) the search for the MOC. It also suggests that cognitive science *can develop correctly* without a MOC[13] - thus casting some doubt about the usefulness of the whole search for the MOC project. Lastly, it assumes cognitive science will converge to a single research tradition (or few mutually consistent ones) with a single MOC (or few mutually consistent ones). But *will it*? What Justifies such faith? Indeed, as things stand the *individual disciplines* within cognitive science seem to require, and thus to push for, different MOCs. Microbiology and plant science often focus on the way in which (comparatively simple) biological systems cope with its immediate environment, focusing on relatively small-scale sensorimotor interactions (Lyon 2015; Baluska and Levin 2016). These disciplines - or, better, their philosophical spokespeople - push for fairly minimal and liberal MOCs, that can be easily applied to the system they are interested in studying (e.g. Van Duijn *et al*. 2006; Lyons 2006; Kiverstein and Sims 2021). But robotics and AI push for more restrictive and demanding MOCs (cf. Webb's piece in Bayne *et al*. 2018; Webb 2006; Nolfi 2002; Tani 2007; 2016). It's not hard to understand why: they know that *comically simple* systems can skillfully interact with the environment (cf. Braitenberg 1984). So they favor demanding MOCs justifying their claim that (certain) robots and computers *really* cognize. Note that here I've considered two disciplines that, in the current landscape of cognitive science, are often quite close and willing to cooperate with each other (cfr. Beer *et al*. 1997; Keijzer 2001). And yet, they do not seem to agree on the MOC.

Does any of this *prove* that no MOC defining the scientific kind cognition can be provided? Of course not. Maybe tomorrow there will be a real revolution in cognitive science, establishing a single MOC. But this isn't the most likely of scenarios. And, if cognitive science continues to stay splintered the way it currently is, then no MOC capturing the scientific kind cognition can be provided. Or so, at least, I argued.

5 - Where from here?

Tying things up: §2 highlighted some of the motivations behind the search for the MOC, as well as the *desiderata* a MOC should satisfy (§2.1). It also highlighted a tension in tension due to the current search for the MOC, which is due to the massive role played by intuitions (§2.2). Hence the proposal of distinguishing two different projects: (i) that of providing a

[13] It is also likely to underplay a host of non-truth-aimed factors that might impact the development of sciences for the better or the worst (e.g. distribution of funds).

MOC capturing the "folksy" notion of cognition and (ii) that of providing a MOC capturing a scientific kind. I then moved a variety of considerations to claim that, *at least at things stand now*, none of these two projects can be carried out successfully (§§3-4). Am I right? and if yes, what would follow?

Concerning the "am I right?" point, I can foresee an obvious objection: the focus on necessary and sufficient conditions sets the bar too high.[14] The idea that genuine scientific kinds (amongst which, natural kinds) are identified by a set of individually necessary and jointly sufficient conditions is wrongheaded. *In lieu* of these conditions, we should search for homeostatic property clusters; that is the mechanism in virtue of which a number or relevant properties (in this case, cognitive ones) cluster together (cf. Boyd 1991). Buckner (2015) did it and *his* MOC seems safe from many of the problems I raised. True. But *it does not* meet the relevant *desiderata* I identified.[15] As a matter of contingent fact, that MOC does not determine *which* research tradition(s) is (are) right, nor which disciplines are a legitimate part of the multidisciplinary endeavor of cognitive science. Nor it identifies *all and only* instances of cognition. Indeed, it cannot: homeostatic property clusters leave, and must leave, a "gray area" around them. By *not* indicating any individually necessary condition, they *always* leave the door open to borderline cases: cases that are neither definitely "inside" nor definitely "outside" the kind defined *via* the cluster.[16] Thus *all* cluster-based MOCs fail (and must fail) to satisfy *desideratum #1* (cf. Walter and Kastner 2012: 21). As a result, they all fail to satisfy *at least some* relevant desiderata guiding and motivating the search for the MOC.

On the "what follows" point, opinions vary. Some think the absence of a MOC does not matter (Allen 2017). Others, instead, paint apocalyptic scenarios. Adams and Aizawa (2008: 79-83 and Aizawa 2017), in particular, contend that, absent a MOC, we are all drawn towards a nasty form of operationalism. And that is undesirable for a wide number of reasons. First, it allows us to identify cognition only in reference to some paradigmatic cognitive processing, without knowing what it *really* is. Secondly, it leads us to *overdistribute* cognition. Many outcomes of cognitive processes can be brought about by non-cognitive means. Thirdly, operationalism is empirically inaccurate: rarely (if ever) cognitive scientists use an operationalistic lexicon. Lastly, operationalism leaves the door open to the return of behaviorism, and surely no one wants behaviorism to return, right?[17]

This pessimism is unjustified. Behaviorism is not returning, partially because it never left, and partially because it does not seem to be gaining popularity. *Anti-representationalism* might be gaining popularity, but anti-representationalism is not

[14] Noteice, however, that is *not me* who focused the discussion on necessary and sufficient conditions, see (Adams and Aizawa 2001; Rowlands 2009; Walter and Kastner 2012).
[15] And, to be fair to Buckner, his MOC is not *supposed* to meet them. He proposed it *only* to distinguish cognition from association in a limited series of cases.
[16] The same applies to all MOCs (or ways to identify the cognitive) that do not identify individually necessary conditions (e.g. Rowlands 2009; 2010; Newen 2017)
[17] These last two worries are more evident in (Adams and Aizawa 2001).

behaviorism. Reading (Kelso 1995; Beer 2000; Chemero 2009; Anderson 2014) and others, one *does not* find any reference to classical or operant conditioning, stimulus response chains, or skinner boxes. Further, not having a MOC is not a point in favor of behaviorism. One could stay agnostic about what cognition is but model it as a computational process for purely pragmatic reasons (cf von Neuman 1958). Whilst the truth (or appropriateness) of a behavioristic MOC *would* entail the truth (or appropriateness) of behaviorism, the *absence* of a MOC doesn't. Indeed, it's entirely unclear *how* the absence of a MOC would support a research tradition *over* any other research tradition.

Moreover, the charge of operationalism is surely overblown (see also Rupert 2013). Operationalism is the view that the meaning of theoretical terms consists in observations/measurement outputs. According to operationalism, a statement such as "the temperature of the substance *a* is *x*" means roughly "you will read *x* if you probe *a* with a thermometer". Operationalism is a view on the semantics of theoretical terms. Such a semantic hypothesis is *not entailed* by our inability of defining cognition. We may be unable to define cognition, but still hold that "cognition" refers to a real kind. Or to a platonic essence. Or to a social construct. No contradiction arises. Nor our inability to define cognition supports operationalism. Indeed, if our inability to define terms were evidence in favor of operationalism, *operationalism should be our default semantic view*. It seems that we're unable to define most terms (cf. §3). If that were supported for operationalism, operationalism should be extremely well supported. But it isn't, so our inability to define terms (amongst which "cognition") does not support operationalism.

One might wonder whether our inability to define cognition should incline us towards a form of *mysterianism*. In the philosophy of mind, mysterianism is a claim concerning phenomenal consciousness. According to mysterianism, phenomenal consciousness is a physical phenomenon. But due to the limitations of our cognitive architecture, we *cannot* figure out how it metaphysically depends on the brain. (cf. McGinn 1989). The thesis can be easily tinkered with to be applied to cognition. According to the mysterianist about cognition, cognition is a physical phenomenon, but, due to some feature of our cognitive architecture, we *cannot* figure out how it metaphysically depends on the brain.

My view is different from mysterianism. According to mysterianism, the mystery is due to our cognitive architecture. If the mysterian philosopher is right, in order to solve the mystery we don't need arguments. We need evolution to push us in a fairly specific direction. But my argument makes no mention of cognitive architectures.[18] For this reason, in my view our inability to provide a MOC is *both* more contingent and less contingent than mysterianism would entail. It is *less* contingent because, even if we were able to produce a MOC capturing our intuitive notion of cognition, it won't satisfy the

[18] Not even the point about our folksy conception of cognition not specifying individually necessary conditions. For, that point was intended to hold *regardless of the* relevant conceptual format deployed, and thus *regardless of the* specific architecture relying on them.

relevant desiderata - it won't tell us what cognition is *in the relevant sense* (§3). So there just is *no chance* of providing a MOC that way. Yet, according to the mysterianist *there is always* a chance to solve the mystery by evolving in the "right" direction. But when it comes to a MOC capturing a theoretical term, chances seem *way higher* than that. For, in that case, our inability to provide a MOC depends on the current status of the relevant science(s). That they will converge towards a MOC seems way more likely than the "right" mutation the mysterianist encourages us to hope for.

Whilst my view does not entail mysterianism, it suggests that "cognition" is *vague* in at least some way. On the face of it, my position is minimally compatible with an epistemicist position (e.g. Williamson 1994). Take a vague term *x*. The epistemicist claims: (a) every thing is either determinately *x* or *not x*; (b) there's a sharp cut-off point between the two and (c) when do not *know* where the cut-off (b) mentions is, and so for some things we are unable to tell whether they are *x* or *not x*. This is entirely consistent with what I argued above. I've argued that, *as things stand*, given the current state of cognitive science (and what we know about our folksy concepts) we cannot define cognition. So, we're not in a position to *know*, for every system, state or process whether it is cognitive or not. Moreover, it seems correct to say that MOC searcher *must be* epistemicists. For they search necessary and sufficient conditions to identify *all* and *only* cognitive processes. This search makes sense only if one thinks that everything is either determinately cognitive or not cognitive and that we do not (yet) know where the cutoff point between the two is. And, to me, that seems the core of epistemicism. Yet notice: the MOC searcher is (and must be) a *contingent* epistemicist. If the search is fruitful, we *will* discover the relevant cut-off point in (b). The same element of contingency is present in the view I've offered here. My claim is that *as things stand* no MOC can be provided. I'm thus *not* committed to the view that no MOC can *ever* be provided - things may change.

What about the idea that *cognition itself* is vague; that is, that there are systems, states or processes that are definitely neither cognitive nor non-cognitive? This claim is compatible with mine: clearly, if there are systems, states or processes that are neither cognitive nor non-cognitive, then we cannot hope to find a MOC drawing a crisp line between the two. Yet, whereas my claim is consistent with this form of vagueness, my claim *does not* support it above the alternatives. Indeed, it is positively compatible with alternative views such as epistemicism.

This suggests an intriguing direction for future research: is *cognition itself* vague? I'm tempted to say yes, and I think that potent arguments could be leveraged to defend it (arguments *highly inspired* by Dennett 1991; Schwitzgebel *unpublished*). But of course, defending it is well beyond the scope of this paper. I just want to point out that answering that question seems crucial, for it might have important implications for "the mental broadly understood". Would discovering that cognition itself is vague suggest that the property of having a mind is a *fuzzy* property, and that there are systems that are neither determinately minded nor determinately mindless? What would be the implications for extended cognition and

other “vehicle externalist” views - would this entail that “extended” systems are neither definitely cognitive nor definitely non-cognitive? What would the ethical implications be (e.g. for robot and animal rights)? Time (and a different paper) will tell.

References

Adams, F. (2010). Why we still need a mark of the cognitive. *Cognitive Systems Research*, 11(4), 324-331.

Adams, F. (2016). Information and cognition. In L. Floridi (Ed.), *The Routledge Handbook of Philosophy of Information* (pp.321 - 332). NY: Routledge.

Adams, F. (2018). Cognition wars. *Studies in History and Philosophy of Science A*, 68, 20-30.

Adams, F. (2019). The elusive extended mind. In M. Colombo, E. Irvine, M. Stapleton (Eds.), *Andy Clark and His Critics* (pp. 21-31). New York: Oxford University Press.

Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*(1), 43-64.

Adams, F., & Aizawa, K. (2008). *The Bounds of Cognition*. Oxford: Blackwell.

Adams, F., & Aizawa, K. (2010). The value of cognitivism in thinking about extended cognition. *Phenomenology and the cognitive sciences*, *9*(4), 579-603.

Adams, F., & Garrison, R. (2013). The mark of the cognitive. *Minds and Machines*, *23*(3), 339-352.

Aikins, K. (1996). Of sensory systems and the aboutness of mental states. Journal of Philosophy, 93(7), 337-372

Aizawa, K. (2017). Cognition and behavior. *Synthese*, *194*(11), 4269-4288.

Akagi, M. (2016). Cognition in practice: conceptual development and disagreement in cognitive science. *Ph.D. Dissertation*. University of Pittsburgh. Retrieved at http://d-scholarship.pitt.edu/27590/1/akagimsm_etd.pdf last accessed 02/03/2022

Akagi, M. (2018). Rethinking the problem of cognition. *Synthese*, *195*(8), 3547-3570.

Akagi, M. (2021). Cognition as the sensitive management of an agent’s behavior. *Philosophical Psychology*, https://doi.org/10.1080/09515089.2021.2014802

Anderson, M. (2014). *After Phrenology*. Cambridge, MA.: The MIT Press.

Andrews, M. (2021). The math is not the territory. *Biology & Philosophy*, *36*(3), 1-19.

Allen, C. (2017). On not defining cognition. *Synthese*, *194*(11), 4233-4249.

Arras, K.O., & Cerqui, D. (2005). Do we want to share our lives and bodies with robots? A 2000 people survey. *Swiss Federal Institute of Technology*, report 605.

Bach-y-Rita, P. & Kercel, S. W: (2003). Sensory substitution and the human-machine interface. *Trends in Cognitive Sciences*, *7*(12), 541-546.

Baluska, F., & Levin, M. (2016). On having no head: cognition throughout biological systems. *Frontiers in Psychology*, *7*: 902.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*(4), 577-560.

Bayne, T. *et al.* (2018). What is cognition?. *Current Biology*, *29*(13), R608-R615.

Bechtel, W. (1998). Representations and cognitive explanations. Assessing the dynamicists challenge in cognitive science. *Cognitive Science*, *22*(3), 295-317.

Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, *72*(1-2), 173-215.

Beer, R. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, *4*(3), 91-99.

Beer, R., *et al*. (1997). Biologically inspired approaches to robotics: what can we learn from insects? *Communications of the ACM*, *40*(3), 30-38.

Bermudez, J. (2010). *Cognitive Science: An introduction to the sciences of the mind*, Cambrdige, Cambridge University Press.

Blok, N., Kitcher, P. (2010). Misunderstanding Darwin. *Boston Review*. Accessed at: https://bostonreview.net/articles/ned-block-philip-kitcher-misunderstanding-darwin-natural-selection/ last accessed 03/03/22

Bocanegra, B. R. *et al*. (2019). Intelligent problem solvers externalize cognitive operations. *Nature Human Behavior*, *3*(2), 136-142.

Boden, M. (2008). *Mind as Machine: A History of Cognitive Science*. New York: Oxford University Press.

Boone, W., & Piccinini, G. (2016). The cognitive neuroscience revolution. *Synthese*, *193*(5), 1509-1534.

Boyd, R. (1991). Realism, anti-foundationalism, and the enthusiasm for natural kinds. *Philosophical studies*, *61*(1), 127-148.

Buckner, C. (2015). A property cluster theory of cognition. *Philosophical Psychology*, *28*(3), 307-336.

Buzaki, G. (2019. *The Brain from Inside Out*. New York: Oxford University Press.

Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA.: The MIT Press.

Calvo Garzon, F. (2007). The quest for cognition in plant neurobiology. *Plant Signaling & Behavior*, *2*(4), 208-211.

Chalmers, D. (1996). *The Conscious Mind*. New York: Oxford University Press.

Charrettes, P. (2006). *The Architecture of the Mind*: *Massive modularity and the flexibility of thought*. Gloucestershire: Clarendon Press

Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, MA.: The MIT Press.

Churchland, P. S. (1987). Epistemology in the age of neuroscience. The Journal of Philosophy, 84(10), 544-553

Churchland, P. S. (2002). *Brainwise*. Cambridge, MA.: The MIT Press.

Churchland, P. M. (1992). *A Neurocomputational Perspective*. Cambridge, MA.: The MIT Press.

Cisek, P., & Hayden B. Y. (2022). Neuroscience needs evolution. *Philosophical Transactions of the Royal Society B*, *377*(1844), 20200518.

Clark, A: (1993). *Associative Engines*. Cambridge, MA.: The MIT Press.

Clark, A. (1997). *Being There.* Cambridge, MA.: The MIT Press.

Clark, A. (2001). Reasons, robots, and the extended mind. *Mind & Language*, *16*(2), 121-145.

Corcoran, A., *et al*. (2020). From allostatic agents to counterfactual cognizers: active inference, biological regulation, and the origins of cognition. *Biology and Philosophy*, *35*(3), 1-45.

Dehaene, S., & Cohen, L. (2007). Cultural recycling of cortical maps. *Neuron*, *56*(2), 384-392.

Dennett, D. C. (1978). Why not the whole iguana? *Behavioral and Brain Sciences*, *1*(1), 103-104.

Dennett, D. C. (1991). *Consciousness Explained*. Boston: Little Brown.

Eliasmith, C. (2013). *How to Build a Brain*. New York: Oxford University Press.

Elpidorou, A. (2014). Reasoning about the mark of the cognitive. *Minds and Machines*, *24*(2), 201-211.

Facchin, M. (2021). Structural representations do not meet the job description challenge. *Synthese*, *199*(3-4), 5479-5505.

Fodor, J. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences*, *3*(1), 63-73.

Fodor, J. (1981). The present state of the innateness controversy. In J. Fodor (Ed.), *Representations: philosophical essays on the foundations of cognitive science* (pp. 127-137). Cambridge, MA.: The MIT Press.

Fodor, J. (1999). Why the brain? *London Review of Books*, *19*(30), accessed at https://www.lrb.co.uk/the-paper/v21/n19/jerry-fodor/diary last accessed 03/03/2022

Fodor, J., & Palmarini, M. P. (2011). *What Darwin Got Wrong*. Profile Books.

Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition*, *28*(1-2), 3-71.

Frank, A. *et al*. (2022). Intelligence as a planetary scale process. *International Journal of Astrobiology*. https://doi.org/10.1017/S147355042100029X.

Friston, K. (2012). A free-energy principle for biological systems. *Entropy*, *14*(11), 2100-2121.

Gardner, H. (1987). *The Mind's New Science*. Basic books.

Gessel, B., *et al*. (2021). Multivariate pattern analysis and the search for neural representations. *Synthese*, *199*(3-4), 12869-12889.

Godfrey-Smith, P. (2010). It got eaten. *London Review of Books*, *32*(13).

Hacken, H. (1978). *Synergistics: an introduction.* Berlin/New Work: Springer.

Hurley, S. (2001). Perception and action: alternative views. *Synthese*, *129*(1), 3-40.

Kaplan, D., & Bechtel, W. (2011). Dynamical models: an alternative or complement to mechanistic explanations? *Topics in Cognitive Science*, *3*(2), 438-444.

Keijzer, F. (2001). *Representation and Behavior*. Cambridge, MA.: The MIT Press.

Keijzer, F. (2021). Demarcating cognition: the cognitive life sciences. *Synthese*, *198*(1), 137-157.

Kelso, S. (1995). *Dynamic Patterns*. Cambridge, MA.: The MIT Press.

Kiverstein, J., & Sims, M. (2021). Is free-energy minimization the mark of the cognitive?, *Biology and Philosophy*, *36*(2), 1-27.

Kriegskorte, N. & Kevit, R. (2013). Representational Geometry: Integrating Cognition, Computation, and the Brain. *Trends in Cognitive Sciences*, *17*, 401–12.

Lamb, M., & Chemero, A. (2018). Interaction in the Open. In A. Newen, L. De Bruin, S. Gallagher (Eds.), *The Oxford Handbook of 4E Cognition* (pp. 147-162). New York: Oxford University Press.

Laudan, (1977). *Progress and its problem: Towards a theory of scientific growth*. Berkeley: University of California Press.

Loomis, J. *et al*. (2012). Sensory substitution of vision: importance of perceptual cognitive processing. In R. Manduchi, S. Kurniawan (Eds.), *Assistive technology for blindness and low vision* (179-210) CRC Press.

Lyon, P. (2006). The biogenic approach to cognition. *Cognitive Processing*, *7*(1), 11-29.

Lyon, P. (2015). The cognitive cell: bacterial behavior revisited. *Frontiers in Microbiology*, *6*: 264.

Lyon, P. (2020). Of what is "minimal cognition" the half-backed version? *Adaptive Behavior*, *28*(6), 407-424.

Machery, E. (2009). *Doing Without Concepts*. New York: Oxford University Press.

Machery, E. (2012). Why I stopped worrying about the definition of life... and why you should as well. *Synthese*, *185*(1), 145-164.

Malafouris, L. (2013). *How things shape the mind*. Cambridge, MA.: the MIT Press.

McGinn, C. (1989). Can we solve the mind body problem? *Mind, 98*(391), 349-366.

McGivern, P. (2019). Active materials: minimal models of cognition? *Adaptive Behavior*, *28*(6), 441-451.

Medin, D. L., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.

Minky, M., & Papert, S. (1969). *Perceptrons*. The MIT Press(??)

Milner, D., & Goodale, M. A. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20-25.

Mitchell, M. (2019). *Artificial Intelligence: a guide for thinking humans*. UK: Penguin.

Murphy, G. L., & Medin, D. L. (1978). The role of theories in conceptual coherence. *Psychological Review*, *92*(3), 289.

Newell, A. & Simon, H. (1976). Computer science as an empirical inquiry: symbols and search. Reprinted in J. Haugeland (Ed.), *Mind Design II* (81-110). Cambridge, MA.: The MIT Press.

Newen, A. (2017). What are cognitive processes? An example-based approach. *Synthese*, *194*(11), 4251-4268.

Nolfi, S., (1999). The power and limits of purely reactive agents. *Neurocomputing*, *42*(1-4), 119-145.

Pessoa, L. *et al* (2022). Refocusing neuroscience: moving away from mental categories and towards complex behaviors. *Philosophical Transactions of the Royal Society B*, *377*(1844), 20200534.

Pfeifer, R., & Bongard, J. (2006). *How the Body Shapes the Way We Think*. Cambridge, MA.: The MIT Press.

Piccinini, G. (2020). *Neurocognitive Mechanisms: explaining biological cognition.* New York: Oxford University Press.

Ramsey, W. (2007). *Representation Reconsidered*. Cambridge: Cambridge University Press.

Ritchie, J.B., iet al. (2019). Decoding the brain: neural representations and the limits of multivariate pattern analysis in cognitive neuroscience. *The British Journal of Philosophy of Science*, *70*(2), 581-607.

Risko, E. F. & Gilbert, S. J. Cognitive offloading. Trends Cogn. Sci. 20, 676–688 (2016).

Rosch, E. & Mervis, C. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573-603.

Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, *22*(1), 1-19.

Rowlands, M. (2010). *The New Sciences of the Mind*. Cambridge, MA.: The MIT Press.

Rumelhart, D., McClelland, J., & the PDP group (1986). *Parallel Distributed Processing*. *Vol I and II*. Cambridge, MA.: The MIT Press.

Rupert, R. (2013). Memory, natural kinds, and cognitive extension. *Review of Philosophy and Psychology*, *4*(1), 25-47.

Ryan, S., *et al* (2021). Active inference models do not contradict folk psychology. *PsyArXiv Preprint 10.31234/osf.io/kr5xf*

Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, *3*(3), 210-229.

Scarantino, A. (2012) . How to define emotions scientifically. *Emotion Review*, *4*(4), 358-368.

Schwitzgebel, E. (*unpublished manuscript*). Borderline Consciousness, When It's Neither Determinately True nor Determinately False That Experience Is Present. Accessed at: http://www.faculty.ucr.edu/~eschwitz/SchwitzPapers/BorderlineConsciousness-211220.htm Last accessed 03/03/2023

Searle, J. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, *3*(3), 417-424.

Searle, J. (1984). Intentionality and its place in nature. *Dialectica*, *38*(2-3), 87-99.

Selfridge, O. G., & Neisser, U. (1960). Pattern recognition by machines. *Scientific American*, *203*(2), 60-69.

Shapiro, L. (2013). Dynamics and cognition. *Minds and Machines*, *23*(3), 355-375.

Sims, A. (2017). The problems with prediction: the dark room problem and the scope dispute. In T. Metzinger, W. Wiese (Eds.), Philosophy and Predictive Processing: 23. Frankfurt am Main: The MIND Group. https://doi.org/10.15502/9783958573246

Sims, M. (2021). A continuum of intentionality: linking biogenic and anthropogenic approaches to cognition. *Biology and Philosophy 36*(6), 1-31.

Sprevak, M. (2009). Extended cognition and functionalism. *The Journal of Philosophy*, *106*(9), 503-527.

Steiner, P. (2021). Radical views of cognition and the dynamics of scientific change. *Synthese*, *198*(1), 547-569.

Staddon, J. (2014). *The New Behaviorism*, Psychology Press.

Tani, J. (2007). On the interactions between top-down anticipation and bottom-up regression. *Frontiers in Neurorobotics*, *1*, 2.

Tani, J. (2016). *Exploring Robotic Minds*. NY: Oxford University Press.

Thelen, E:, & Smith, L. (1996). *A Dynamic System Approach to the Development of Cognition and Action*. Cambridge, MA.: The MIT Press.

Tripladi, L. (*Forthcoming*). *Parallel Minds*. Cambridge, MA.: The MIT Press.

Trovato, G., & Eyssel, F. (2017). Mind attribution to androids: a comparative study with Italian and Japanese adolescents. *26th IEEE international symposium on robot and human interactive communication* (Ro-MAN), (pp. 561-566), IEEE.

Van Dujin, M., *et al*. (2006). Principles of minimal cognition: casting cognition as sensorimotor coordination. *Adaptive Behavior*, *14*(2), 157-170.

Van Gelder, T. (1995). What might cognition be, if not computation?. *The Journal of Philosophy*, *192*(5), 345-381.

von Neumann, J. (1958). *The Computer and the Brain.* New Haven, Connecticut: Yale University.

Walter, S. (2010). Cognitive extension, the parity argument, functionalism, and the mark of the cognitive. *Synthese*, *177*(2), 285-300.

Walter, S., & Kasner, L. (2012). The where and what of cognition: the untenability of cognitive agnosticism and the limits of the Motley Crew Argument. *Cognitive Systems Research*, *13*(1), 12-23.

Webb, B. (2006). Transformation, encoding, representation. *Current Biology*, *16*(6), R184-R185.

Wheeler, M. (2011). In search of clarity about parity. *Philosophical Studies*, *152*(3), 417-425.

Wheeler, M. (2019). Breaking the waves. In M. Colombo, E. Irvine, M. Stapleton (Eds.), *Andy CLark and His Critics* (pp. 81-95). NY:Oxford University Press.

Williamson, T. (1994). *Vagueness*. New York: Routledge.

Yakura, H. (2019). A hypothesis: CRISPR-Cas as a minimal cognitive system. *Adaptive Behavior*, *27*(3), 167-173.