



# Blinded by Biology: Bio-inspired Tech-Ontologies in Cognitive Brain Sciences

Paola Hernández-Chávez<sup>(✉)</sup>

Center for Philosophy of Science, University of Pittsburgh, Pittsburgh, USA  
hcpaola@gmail.com

**Abstract.** In his pioneering paper on neuromorphic systems, Carver Mead conveyed that: “Biological information-processing systems operate on completely different principles from those with which most engineers are familiar” (Mead 1990: 1629). This paper challenges his assertion. While honoring Mead’s exceptional contributions, specific purposes, and correct conclusions, I will use a different line of argumentation. I will make use of a debate on the classification and ordering of natural phenomena to illustrate how background notions of causality permeate particular theories in science, as in the case of cognitive brain sciences. This debate shows that failures in accounting for concrete scientific phenomena more often than not arise from (1) characterizations of the architecture of nature, (2) singular conceptions of causality, or (3) particular scientific theories – and not rather from (4) technology limitations *per se*. I aim to track the basic bio-inspiration and show how it spreads bottom-up throughout (1) to (4), in order to identify where bioinspiration started going wrong, as well as to point out where to intervene for improving technological implementations based on those bio-inspired assumptions.

**Keywords:** Natural kinds · Cognitive brain sciences · Ontology · BioInspiration · Technology

*Listen to the technology and find out what it's telling you.*

Carver Mead

## 1 Introduction

Ontological conceptions are background ideas that pervade the practice of science, technology, and their contrivances. A typical example is that, for a system to be explained scientifically, it must be the kind of thing that admits a mechanical account. Ontological conceptions are ways of framing a problem that we often take for granted. One can go to a psychotherapist to reflect upon the hidden motivations for their behaviors. But we do not go to the ontologist to analyze the assumptions we initially embrace about the structure of the world. Since we do not often reflect upon these

assumptions, the consequences of adopting them are not addressed. We can step back to draw attention to them when facing technological implementation difficulties such as how stable systems are in general, or the extent to which humans can track microscopic states.

Classifying natural phenomena into general regularities reflects a human proclivity to optimize our understanding of the world. Classification and categorization are recurring practices throughout the history of science. Metascientific studies of how scientists allocate different systems of classifications go back and forth in physics, biology, and recently in cognitive brain sciences. This is not surprising, given the impact these ones have in terms of technological resources, numerous results, and the public attention it draws. Yet, few analyses have been run regarding the origins of neuroscience classifications and the bio-inspired causal assumptions they involve.

As we will see, the bio-inspired causality assumptions at work in current cognitive brain sciences can be traced back to a causal agent-based model coming mainly from nineteenth-century classical biology – which is very much alive.

The following section gives an account of natural kinds and the multiple realization debates in philosophy. The purpose of discussing these concepts is to state the basic terminology and general background for addressing the problem at hand, as well as to draw awareness to the way ontological assumptions and conceptualizations permeate technological tools for explaining particular natural phenomena. Talking about the organization of nature will allow us to transition from biology to medicine, neurology, and finally to cognitive brain sciences. The third section is a quick overview of cognitive brain sciences' range of classification schemes. The fourth section surveys the roots of common inspiration between biology and cognitive brain sciences. The fifth and final section offers recommendations for improving bio-inspiration.

The main contribution of this paper consists in tracking the inspirational explanatory patterns common to biology and brain science technologies, showing why it is necessary to stop importing causal assumptions from classical biology, and where to intervene to improve technological implementations based on those bio-inspired assumptions. To get to these conclusions, we must start with some basic notions from philosophy.

## **2 Some Philosophical Background: Natural Kinds and Multiple Realizability**

The purpose of this section is to define a basic terminology and the general background ideas needed to address the various ways in which natural phenomena could be described. There is an old philosophical debate regarding our most basic notions of the organization of nature and of causality, known as the natural kinds (NK) debate. If you are among those endorsing the idea that our basic scientific taxonomies depict the exact organization of nature as it really is, a philosopher might claim you are a “universalist,” a “realist,” or an “essentialist.” This side of the debate encompasses the intuition that the labor of science is to group phenomena based on their properties, causal

relationships, and governing laws. On the other side is instrumentalism, or pragmatism. This is a slightly looser perspective. Kinds and scientific categories function as useful tools to grouping phenomena only for the sake of providing explanations, predictions, or elaborating reliable inferences.

Both sides have persuasive points. On the one hand, an instrumentalist might say that belief in natural kinds exhibits a (wrong-headed) faith in the order and regularity of nature. On the other, a merely instrumental use of scientific classifications does not guarantee accuracy, since the history of science teaches us that scientific practice can prevail for centuries, even millennia, grounded on spurious categories and utterly wrong beliefs. The shift from geocentrism to heliocentrism is the best exemplar of this case.

There is another way that ontology can inspire our causal assumptions: through multiple realizability (MR). MR is the possibility of achieving the same goal by causally different routes. For example, having two distinct entities (ex. parts of the body, components, neural substrates) performing the same function (ex. mental state, cognitive task, instruction) by different operation modes could count a case for MR.

For twentieth century philosophers of mind, it was almost a truism that psychological states are multiply realizable (Putnam 1967) - that a mental state like feeling pain could be realized in different (physical) ways across species, and even between individuals of the same species. In the same vein, a cognitive state could be accomplished by different brain substrates. This is similar to having two subjects performing the same cognitive task while a brain scanning shows that they elicit different brain activity patterns.

MR has some resemblance to redundancy in biology, a common event where a gene is duplicated within a genome of a complex organism. When a system is interrupted by a backup condition or a compensatory response, for example, redundancy helps to facilitate the central functioning and maintenance of the system. In such cases, two or more genes in the genome of an organism can be performing the same function but the activation in one of them happens to have no effect on the phenotype. There are three types of genetic redundancy: true, generic, and ‘almost’. In true redundancy, a subject with a redundant genotype AB happens to be no fitter than a second one who has redundant genes being mutated or knocked out. In generic redundancy, an AB subject could be occasionally fitter. However, in an ‘almost’ redundancy case, the redundant genotype is slightly fitter than any other genotype where one of the redundant genes has been mutated or knocked out (see Nowak et al. 1997).

Several authors have written against MR. One of its most remarked inconveniences is that, as Kim (1992) noticed decades ago, multiply realizable functional kinds are not projectible –owning a predicate to project properties upon it–, which makes it complicated to nominate them as candidates of scientific kinds. This is, when kind members lack an underlying causal basis for membership in that kind, we cannot make inductive generalizations about the nature of that kind. From here, Kim (1992) conjectured that kinds with different physical realization are distinct kinds - structure independent kinds that do not count as causal kinds - so they are not proper scientific kinds.

It is easy to anticipate why MR is ontologically compromising for scientists. Systematizing experiments with possibly ‘multiple’ outcomes makes science less defensible since experimental configurations would be difficult to model. In spite of this, it is essential to consider that such an uncomfortable possibility might be the case.

Neglecting NKs to avoid ontological commitments became common ground. Philosophers moved on from arguing over natural kinds and so did scientists at stopping debating the reality of natural kinds to avoid ontological commitments. The discussion settled on a softened, naturalized or more scientific friendly conception of “scientific kinds,” the hallmark of which was a profound reliance on prosperous scientific practice. This is, further formulations derived from successful (or productive) scientific practices derived from analyzing how scientists use classifications in practice. A sophisticated approach was Richard Boyd’s discussion of homeostatic property clusters (Boyd 1990, 1991, 1999a, 1999b, 2000; Boyd 2003). Boyd contended that nature doesn’t neatly divide into well-delimited sorts of things. Instead, groupings might be modified in light of new observations or when inferences fail, since definitions of kinds are *a posteriori* (Boyd 2000, 54). For him, kinds used in sciences are not features of the world but products of our engagement with it. As he puts it: “the theory of natural kinds is about how schemes of classification contribute to the formulation and identification of projectible hypotheses” (Boyd 1999a: 147). This means that NKs are groups of entities that share a cluster of projectable properties sustained by homeostatic causal mechanisms, which are understood as anything that causes a repetitive clustering of properties. Consequently, if we cannot make projections and inductive generalizations of a case – say, if the case is multiply realizable – that case cannot be considered as a NK.

Ereshefsky and Reydon (2015) proposed an influential reformulation of the debate that shifted the focus to scientific practice. There they aimed to track a variety of classificatory practices of successful science. They pursued a NKs account that recognizes diverse scientists’ aims when constructing classifications. Yet, they denied that any classification offered by scientists correspond to NKs.

Ereshefsky and Reydon (2015) introduce the idea of a “classificatory program”, understood as the part of a discipline that produces a classification. According to them, a classificatory program contains three parts: sorting principles, motivating principles, and classifications. Classifications describing NKs are marked by: internal coherence, empirical testability, and progressiveness. The primary virtue of Ereshefsky and Reydon’s (2015) proposal is their lack of commitment to the existence of ontological NKs. They introduced several criteria for what makes a classification of natural kinds valid (its internal coherence, empirical testability, etc.), which shifts the frame for analyzing classifications away from ‘are these true natural kinds?’ to ‘do these natural kinds work in the context of scientific practice?’.

But some have argued that, in their effort to encompass all successful classificatory practices, Ereshefsky and Reydon distort some heuristics and common practices in science. However, scrutinizing such critiques is beyond the purposes of this work. Suffice it to say that the possibility of MR in the context of natural kinds raises the issue of whether natural kinds are ontologically real or just an instrumental/methodological guide. Ideally, kind members should be caused by the real properties of the world, and we should be able to identify and specify the causal mechanisms responsible for the grouping of kind members. Nevertheless, arbitrary classifications may occur.

It is crucial to be aware of the fact that an Essentialist, who by definition is committed to the ontological existence of NKs, in his willingness to avoid the possibility of MR, would need to postulate an *overflowing* amount of kinds – as many as one

for almost every new phenomenon (each corresponding to a particular law of Nature). But closed cause-effect events are scarce. The paradox is that for the Essentialist the ontological search for NKs as consisting of a system of neat, simple, and well-organized frames of Nature loses its original appeal: that of unraveling the laws of an engineered Nature. By contrast, since an Instrumentalist only needs to describe Nature in such a way that it works for science, he would readily allow the incursion of MR. He would then indirectly promote a more *parsimonious* amount of causes, admitting only those coming from the most basic physics, for example. Ironically, for an Instrumentalist Nature would be less chaotic than for an Essentialist.

The take-home lesson from this survey on NK and MR is the importance of being aware of the ontological commitments brought by our intuitions about how the world and natural phenomena are arrayed. If your intuition tells you, for example, that the human mind is modular, and that the most basic operations of it are computational, this will permeate the technological implementation of your experiments, perhaps by means of scaling up the power of computational outcomes to make them the most fundamental part of brain cognitive operations – regardless of whatever related processes remain unaccounted for. As another example: if you have sympathy for NKs, it is more likely that you start a paradigm that correlates a brain region to remembering as a broad category, then looking for additional brain regions for particular tasks, associating particular regions to each thing we humans do. On the contrary, if you accept MR and instrumentalism in principle, it would be plausible to claim that there is nothing essential to remembering, so you will find no troubles transiting from that nothingness, to game theory, to modeling accordingly. Overall, starting from scratch might look like a challenging and complex outset, but it will leave you the option of using any explanatory theory or methods that allows you to understand and predict phenomena. On the contrary, starting from within a NKs framework might lead you to a very chaotic multiplication of categories. This option seems to be less defensible, since there has been no long-lasting scientific system of categories – or maybe there are just a very few ones. This is most likely because, in the effort to control, predict and discover new natural phenomena, science constantly invents emerging classification systems.

With this philosophic background in mind, we can turn to a comparison of explanatory patterns in brain sciences and in biology. I will pursue this comparison by first elaborating on the landscape of explanatory frameworks in the brain sciences.

### 3 Cognitive Brain Sciences Landscapes and Classifications

Cognitive accounts of the brain are abundant. Diverse descriptions are found depending on the different collection of components, hierarchies, and type of interactions between them. Agreement about which are the correct set of components of cognitive processes in the brain remains far away.

Many different lists of cognitive brain sciences components might be found. For the sake of space and with the purpose of illustrating the point, let us grant that there exist only two basic models. The most basic level of brain cognition, what I will call the “atomistic model,” would include the following components:

a molecular scale  
 ion channels + neurotransmitters  
 synapses  
 neurons  
 neural networks  
 neural systems

On the other hand, we could elaborate another list for the “cognitive perspective model.” This would include the following components:

brains  
 a physical body (embodied experience)  
 subjects  
 interacting subjects  
 culture / environmental elements (extended cognition)

From this perspective, the cognitive brain and its components would be literally extended in multifactorial ways.

Why is integration among this dualistic simplified version of brain and cognition still so challenging for theorists, engineers, and scientists in general? Following the discussion in section two, I suspect this is because different models imply a different ontological array of natural components.

The ontological components comprising a scientist’s favorite system will determine the level to which each one pays more attention. Among the possible elements of composition that a scientist could emphasize are: the scale of interest (a fine-grained one, or even a nanotechnological perspective, as opposed to big data collaborative projects)<sup>1</sup>, the multiscale (a computer simulation of interactions vs. engineering a whole neural systems), the method (ex. reverse engineering as a tool for emulating micro-structures interactions), the elements to simulate (ex. mimicking synaptic transmission arrays), the spatial and temporal resolution, the logic of the behavior (decision making vs. game theoretic), or the underlying rules (representational vs. connectionist responses; or connectionist vs. modular responses), among others.

---

<sup>1</sup> Notable approaches to the study of the brain are the Human Connectome Project (USA) and the Human Brain Project Initiatives (Europe). Despite their refractory differences, both concur that the fundamental puzzles in Neuroscience are:

- Deciphering the primary language of the brain
- Understanding the rules governing how neurons organize into circuits;
- Understanding how the brain communicates information from one region to another, and which circuits to use in a given situation;
- Understanding the relation between brain circuits, genes, and behavior;
- Developing new techniques for analyzing and observing brain function.
- Disentangling the essential elements of neural computation.

Having sketched the connection between ontology and cognitive brain sciences, and the relevance of this connection, the following section addresses the common explanatory patterns that cognitive brain sciences share with biology and where these patterns come from.

## 4 Blinded by Biology

The search for ideological principles guiding cognitive brain sciences refers us to the roots of scientific medical thinking, itself so much inspired by a notion of causality coming from nineteenth-century biology. There, an atomistic logic pattern prevails: a single agent ‘S’ causes infection ‘I’. Until recently, biology had focused on genes as the definers of structure and function, where the agent is understood as only a physical realization of the genetic program, without regard for the dynamics, interactions and breakout patterns of genes.

Causal agent-based explanatory models in biology and medicine enjoyed a glorious era in the 19<sup>th</sup> centuries. Illnesses such as cholera are great examples of successful causal pathogen identification: bacterium *vibrio cholera* causes cholera, *treponema pallidum* causes syphilis, *H1N1* causes swine flu. Under this line of thought, once you identify the pathogen, it is possible to track its dynamics and predict its course of action, so that subsequent prognosis and models of intervention arise.

The same atomistic causal explanatory practices have lasted right up to the present day. Approaches to the study of cognition and its brain substrates use tools such as: (1) the review of empirical evidence in subjects suffering from brain damage or selective cognitive impairments as isolated components; (2) the study of selective deficits in atypical subjects, i.e., with a neurodevelopmental disorder or a brain damage, also taken as units; and (3) the use of neuroimaging techniques (PET, fMRI, etc.) aimed to register brain activity related to the performance of very particular cognitive tasks – again, taken as units. These tools or techniques embody a linear causal pattern coming from classical biology. These approaches could be seen as trying to test counterfactual situations – like with those who suffer from neurological dysfunctions, being the logic of counterfactuals the same kind of linear causal logic.

The operating logic of classical biology proved to be quite successful for over three centuries. But, for several reasons, it became indisputable that the gene-reductionist paradigm in biology was not a good way to start accounting for the function of a vast number of illnesses. As those authors claim, classical questions in biology as to what life is, what nature is, what an organism is, how they are organized, what their regulating functions are, etc., were considered metaphysical issues (Laubichler 2000; Cornish-Bowden 2006; Nicholson 2014). That is why favoring gene-centered accounts received so much attention, provided the extensive deployment of empirical tools they invoked –leaving the metaphysical voids aside–.

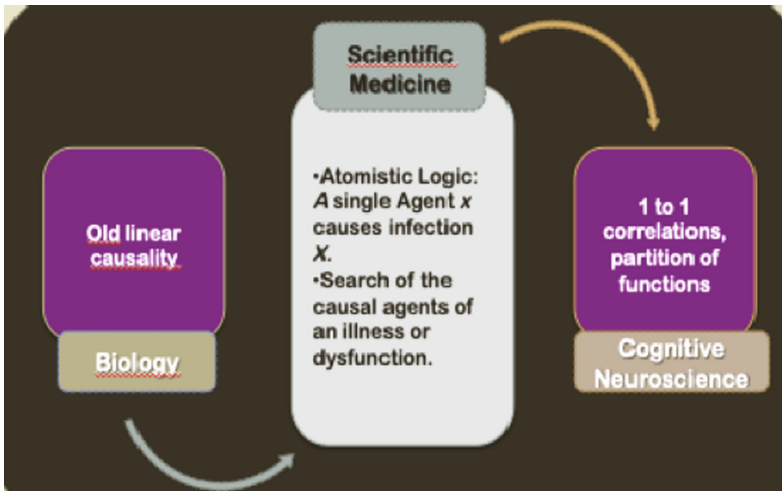
Similarly, as far as cognitive brain science goes, that logic might not be the best route to approach the understanding of the brain and cognitive phenomena provided the conceptual gaps between a brain activation elicited by a cognitive task, the problems with the basic framework assumptions in regards to brain and cognition, the instrumental difficulties elicited by current brain scanning techniques, among others (see Hernández Chávez 2019, forthcoming). But how to overcome this biology-inspired atomistic causality?

## 5 Improving Bio-inspiration

What kind of biology could cognitive brain scientists look to for inspiration? I argue that systems biology offers a better way forward, far beyond nineteenth-century biology, thus representing a better source of inspiration.

A shift took place in biology when some biologists stopped focusing only on genes. As some authors noticed, the term ‘organism’ almost disappeared from mainstream literature as a fundamental explanatory concept in biology (Webster and Goodwin 1982; Laubichler 2000; Huneman and Wolfe 2010; Cornish-Bowden 2006; Nicholson 2014). The shift accomplished by systems biology focused on the following facts: genes live within an organism, and they form complex systems; organisms, not genes, are the agents of evolution; phenotypic plasticity is the rule (genotypes generate different phenotypes depending on the environmental circumstances); phenotypic innovations can be genetically inherited; and, organisms are heterogeneous, and their dynamicity and variability more precisely characterize them. This additionally promoted niche construction approaches, which underlined how organisms modify their environment and also inherit ecological changes (for a quick overview, see Laland et al. 2016).

Given the relationships depicted in Fig. 1, one would expect innovations in systems biology to penetrate reciprocally into cognitive brain sciences. Yet, cognitive brain sciences accounts have not been updated to take into account systems biology conceptualizations. This is my plea for cognitive brain scientists to do so. A systems approach for cognitive brain sciences would promote putting the brain back into the center of the discussion, thus modeling brain operation and functioning as a complex phenomenon.



**Fig. 1.** Represents the possible causal route going from assumptions in classical biology to the atomistic logic of scientific medicine, and then to cognitive brain sciences. As the figure tries to show, the mindset still operating in current cognitive brain sciences was most likely imported from nineteenth-century biology.



Noble's (2012) modeling of heart cells is an interesting case of going beyond a simplistic, classical bio-inspired causality framework. Noble found that as it occurs in all functions that require cellular structural inheritance as well as genome inheritance, there is not a program for cardiac rhythm in the genome. Those findings led him to accept that "we cannot yet characterize all the relevant concentrations of transcription factors and epigenetic influences. [so that] It is ignorance of all those forms of downward causation that is impeding progress" (Noble 2012: 60). As this case shows, a move away from classical bio-inspiration concerning the architecture of natural organization brought him to acknowledge that there is no privileged scale at which biological functions are determined. Why this is a good thing? What was gained from the recognition that there is no privileged scale at which biological functions similarly applies as the upshot for practitioners of cognitive scientists as well.

Another example of overcoming simplistic causality is Marr and Hildreth's (1980) model of edge detection based on intensity changes occurring within an image over a range of scales, where each of them are detected separately at different frames. They discovered that intensity changes tend to emerge from surface discontinuities, reflectance, or illumination boundaries. The most notable characteristic is that they all share the property of being spatially localized, so that there are crossing segments coming from different but non-independent channels – so much so that the operating rules can be deduced from a sketch-like combination of images. The relevant payout of this theory is that it succeeded in advancing several psychophysical discoveries as to how oriented segments are formed.

Some additional cases from studies of insects might be of help for building improved causal frameworks. Recent experiments demonstrate insects' capacity to learn and memorize complex visual arrays that eventually carrying with a modular brain processing. These experiments are salient because they make use of descriptions of the content learned through visual stimuli in combination with a generalized understanding of learned and unlearned routines such that the memory of 'what' and 'where' are stored differently than what is supposed in traditional accounts (see Sztarker and Tomsic 2011).

Similarly, it has been documented in *Drosophila* discrimination and remembering of visual landmarks that select patterns as size and color that are stored according to particular parameter values. They recognize patterns independently of the retinal position and acquisition of the pattern. It has been shown that they also contain something similar to a network-mediated visual pattern recognition. Short-term memory traces of elevation and contour orientation have also been documented, among other findings (see Liu et al. 2006).

Experimental models based on machine learning also offer additional routes out of traditional and simplistic causal assumptions. In those cases, models are carried on trained neural networks to study the propagation, modes of learning, speed, gradients, dimensional patterns, algorithms, processing, and many other features. Neural networks themselves are approximating different types of learning through the manipulation of variables like propagation and speed. More recently, deep learning is

becoming quite influential in bringing together engineering and reinforcement learning. Within those frameworks, complex control problems can be successfully explained not by invoking basic, general operation rules of causality, but rather stochastic, nonlinear, and autonomous dynamics.<sup>2</sup>

Many different research programs now provide improved perspectives for systems dynamics models by going beyond traditional approaches so to include: reinforcement of learning, positive reinforcement, distributed agency, cooperation in game theories, decision making, action control, systems resolutions, alignment of utilities, to name a few. Modeling brain sciences after these efforts would be as fruitful. Those are thus better sources of bio-inspiration when elaborating models for cognitive brain sciences that promote new emerging assemblies and properties.

## 6 Concluding Remarks

In this work, I highlighted how theoretical models are often more powerful than we think, insofar as they influence explanations and technological models of phenomena. To a significant extent, cognitive brain science in particular is bio-inspired by causal assumptions in classical nineteenth-century biology. I argued that an approach inspired by systems biology would be more suitable for understanding human brain dynamics and cognition. This would output back integration, regulation, and organization of the living phenomenon, innovation, among others. In other words, approaches inspired by systems biology would put the focus back on investigating the integration, regulation, and organization of living phenomena *as a whole* –instead of treating organisms as hierarchical assemblies of generic basic components.

I made use of the natural kinds and multiple realizability debate on the ordering of natural phenomena (i.e. discussions of whether nature is a messy place, or if it is well organized; whether phenomena are subject to regularities, etc.) to illustrate how background notions of causality permeate theories in science (ex. atomistic/agent based ones, complex phenomena mindset, indeterminism). In the case of cognitive brain sciences, I examined those background notions of causality evident in different initiatives that focus on learning and connectionist patterns, modular explanations, and input-output processing outlets. I argued that, more often than not, the failures of these initiatives to account for concrete scientific phenomena most likely arise from (1) characterizations of the architecture of natural phenomena, (2) singular conceptions of causality, and (3) particular scientific theories, and not rather from (4) technology limitations *per se*. So I tracked the basic inspiration for these background notions of causality from classical biology, and how this inspiration spread bottom-up throughout (1) to (4). Thus, it suggested it is possible to identify where to intervene to improve technological implementations.

In general, humans and other animals master perceiving the world as a set of uniform and organized structures. As we improve, some aspects become less meaningful. A substantial number of reductions take place – such as ignoring tridimensional

---

<sup>2</sup> See for example, Ng et al. (2006).

arrays, color palette nuances, minimizing complexities, making complex systems to appear as binary, or any form of dualism – which promotes the blatant reduction of representations. To that extent, brain cognitive operations are “crafted functional abstractions” (Cauwenberghs 2013: 15513). These strategies are not necessarily mistaken, since they promote efficient learning of routines and fast processing, among other virtues. Yet, it is paramount to remain aware of those simplifications by re-assessing the levels of simplification/complexity involved and the overall size of the full scale, so as to know: where and how a model is scaling up, what exactly is being computed or simulated, which level of abstraction is being taking for granted and (more importantly) at what cost. Such is the importance of analyzing seriously the quality of bio-inspiration operative models in the cognitive brain sciences.

Being aware of our most basic assumptions going from (1) to (4) wards off contamination in technology-based explanatory models in science, as in the case of brain and cognition functioning and dynamics. We need accuracy in technologies, but we also need real awareness of our ontological commitments in order to have technologies that are well-matched to the explanatory task at hand.

## References

- Boyd, R.N.: Realism, approximate truth, and philosophical method. In: Savage, C.W. (ed.) *Scientific Theories*, pp. 355–391. University of Minnesota Press, Minneapolis (1990)
- Boyd, R.N.: Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philos. Stud.* **61** (1–2), 127–148 (1991)
- Boyd, R.N.: Kinds, complexity, and multiple realization. *Philos. Stud.* **95**(1–2), 67–98 (1999a)
- Boyd, R.N.: Homeostasis, species, and higher taxa. In: Wilson, R.A. (ed.) *Species: New Interdisciplinary Essays*, pp. 141–185. MIT Press, Cambridge (1999b)
- Boyd, R.N.: Kinds as the “Workmanship of Men”: realism, constructivism, and natural kinds. In: Nida-Rümelin, J. (ed.) *Rationalität, Realismus, Revision: Vorträge des 3. Internationalen Kongresses der Gesellschaft für Analytische Philosophie*, pp. 52–89. De Gruyter, Berlin (2000)
- Boyd, R.N.: Finite beings, finite goods: the semantics, metaphysics and ethics of naturalist consequentialism, Parts I & II. *Philos. Phenomenol. Res.* **LXVI**, 505–553 & **LXVII**, 24–47 (2003)
- Cauwenberghs, G.: Reverse engineering the cognitive brain. *PNAS*, **110**(39), 15512–15513 (2013)
- Cornish-Bowden, A.: Putting the systems back into systems biology. *Perspect. Biol. Med.* Autumn **49**(4), 475–489 (2006)
- Putnam, H.: Psychological predicates. In: Capitan, W.H., Merrill, D.D. (eds.) *Art, Mind, and Religion*, pp. 37–48. University of Pittsburgh Press, Pittsburgh (1967)
- Ereshefsky, M., Reydon, T.A.C.: Scientific kinds. *Philos. Stud.* **172**(4), 969–986 (2015)
- Hernández Chávez, P.: *Disentangling Cognitive Dysfunctions: When Typing Goes Wrong*, (2019, forthcoming)
- Huneman, P., Wolfe, C.: The concept of organism: historical philosophical, scientific perspectives. *Hist. Philos. Life Sci.* **32**(2–3), 147–154 (2010)
- Kim, J.: Multiple realization and the metaphysics of reduction. *Philos. Phenomenol. Res.* **52**, 1–26 (1992)

- Laland, K., Matthews, B., Feldman M.: An introduction to niche construction theory. *Evol. Ecol.* **30**, 191–202 (2016)
- Laubichler, M.D.: The organism is dead. long live the organism! *Perspect. Sci.* **8**(3), 286–315 (2000)
- Liu, G., et al.: Distinct memory traces for two visual features in the *Drosophila* brain. *Nature* **439**, 551–556 (2006)
- Marr, D., Hildreth, E.: Theory of edge detection. *Proc. R. Soc. Lond. B Biol. Sci.* **207**(1167), 187–217 (1980)
- Mead, C.: Neuromorphic electronic systems. *Proc. IEEE* **78**(10), 1629–1636 (1990)
- Ng, A.Y., et al.: Autonomous inverted helicopter flight via reinforcement learning. In: Ang, Marcelo H., Khatib, O. (eds.) *Experimental Robotics IX*. STAR, vol. 21, pp. 363–372. Springer, Heidelberg (2006). [https://doi.org/10.1007/11552246\\_35](https://doi.org/10.1007/11552246_35)
- Nicholson, D.J.: The return of the organism as a fundamental explanatory concept in biology. *Philos. Compass* **9**, 347–359 (2014). <https://doi.org/10.1111/phc3.12128>
- Noble, D.: A theory of biological relativity: no privileged level of causation. *Interface Focus* **2**, 55–64 (2012)
- Nowak, M.A., Boerlijst, M.C., Cooke, J., Smith, J.M.: Evolution of genetic redundancy. *Nature* **388**, 167–171 (1997)
- Sztarker, J., Tomsic, D.: Brain modularity in arthropods: individual neurons that support “what” but not “where” memories. *J. Neurosci.* **31**(22), 8175–8180 (2011). <https://doi.org/10.1523/JNEUROSCI.6029-10.2011>
- Webster, G., Goodwin, B.C.: The origin of species: a structuralist approach. *J. Soc. Biol. Struct.* **5**(1), 15–47 (1982). [https://doi.org/10.1016/S0140-1750\(82\)91390-2](https://doi.org/10.1016/S0140-1750(82)91390-2)