MARCO VIOLA
IUSS Pavia and Vita-Salute San Raffaele University
marco.viola@iusspavia.it

CARVING MIND AT BRAIN'S JOINTS. THE DEBATE ON COGNITIVE ONTOLOGY[1]

abstract

Since neuroimaging methods allow researchers to study the human brain at work, the vexed mind-brain problem ceased to be just a metaphysical issue, and became a practical concern for Cognitive Neuroscientists: how could they carve mind and brain into distinct entities, and what is the relation between these two sets? In this paper, I discuss the classical model of one-to-one mappings between mental and neural entities, inherited from phrenology, and make its assumptions explicit. I then examine the shortcomings of this "new phrenology", and explore two solutions to them: the first accepts many-to-many mappings, whereas the second proposes a radically rethinking of the relata of this correspondence.

keywords

philosophy of neuroscience, philosophy of psychology, cognitive ontology, one-to-one mapping, Mind-Body problem

Phenomenology and Mind, n. 12 - 2017, pp. 162-172
DOI: 10.13128/Phe_Mi-21115
Web: www.fupress.net/index.php/pam

ISSN 2280-7853 (print) - ISSN 2239-4028 (on line)

In some cases, when thorny epistemological or foundational issues concerning some specific[1] science are addressed, the boundaries between philosophy and (that specific) science are blurred, and philosophical contributions cannot be distinguished from scientific discourse anymore – thus realizing a sort of Quinean ideal of continuity.
This is currently happening in the ongoing debate on cognitive ontology, where neuroscientists (e.g. Pessoa, 2014; Poldrack, 2010) and philosophers (e.g. the contributions in the special issue edited by Janssen *et al.*, 2017) are discussing the proper role of neuroscientific evidence in determining what mental entities there are. In this paper I sketch a brief reconstruction of this debate, focusing on the controversial assumption of one-to-one mapping between mental and neural entities. I will address the shortcomings of this assumption, and briefly describe an alternative framework that does without it. I will also discuss some attempts to save the one-to-one mapping hypothesis by revising the *relata*, i.e. by proposing that mappings occur between different entities from those that are usually conceived.

1. Basic Ontological Assumptions of Cognitive Neuroscience

In some senses, the cognitive ontology debate is but the heir of the vexed Mind-Body problem. What kind of relationship holds between "the mental" and "the physical" is one of the most debated topics throughout the history of Western philosophy. However, as modern tools and techniques for studying the human brain and mind were developed, the Mind-Body problem has been (re)framed as a methodological question concerning the proper relation between psychology and neuroscience. Inheriting phrenologists Gall and Spurzheim's bold thesis that some mental faculties are localized in some areas of the brain, filtered and refined by figures such as Broca, twentieth century neuropsychologists appealed to the co-occurrence between cognitive impairments and brain lesions (observed post-mortem) to establish some link between mental processes and neural areas. They also looked for counter-intuitive cases in which two (allegedly) unified systems could be selectively impaired in order to establish that they are independent on functional and (usually, though not always) structural grounds. The most famous case was that of memory: once thought to be a single capacity, the observation of

1 I am grateful to Joe Dewhurst, Anneli Janssen and Colin Klein for their valuable comments, as well as for our many stimulating discussions. I am also indebted to three anonymous referees who helped me to ameliorate the intelligibility of the paper.

brain lesioned patients led to its being *split* into different kinds, e.g. "declarative" *vs* "implicit" memory (see for instance Squire, 1992).
However, by the end of the twentieth century, several techniques were developed, that allowed scientists to study living human brains at work, either by measuring electrical signals on the scalp (EEG, MEG) or by assessing the metabolic activity of cortical areas (PET, fMRI). More recent techniques also allow researchers to produce temporary interferences on the activity of some brain areas (TMS, tDCS). Each of these techniques represents a new window[2] through which the psychologist can peek at the brain at work. By accepting some bridge-principles, she can then use this new evidential source to add further constraints to her cognitive theories – each of them positing slightly different carvings-up of the mind (that is, different taxonomies of mental entities; see below). And so many did, thus giving rise to Cognitive Neuroscience.
A common feature of cognitive neuroscience is that neuroscience is deemed a legitimate arbiter for refining cognitive ontology, i.e. for choosing the right set of mental entities. This features rely on the following assumptions:

(1) the mind can be decomposed into distinct mental entities;
(2) the brain can be decomposed into distinct neural entities;
(3) some systematic correspondence holds between mental and neural entities.[3]

Assumptions (1) and (2) have both been challenged on the ground that mental and neural processes are too tightly interconnected to be studied in isolation (see e.g. van Orden and Paap, 1997). While not completely settled, a certain degree of decomposability of mind seems reasonably warranted by the successes of Cognitive Psychology, whereas the decomposability of the brain seems reasonable in virtue of the fact that neuropsychology has proved that the brain cortex is not equipotential. Moreover, albeit mechanistic explanations (regarded by most researchers as the proper kind of explanations in cognitive neuroscience) often begin with decomposition into parts (Bechtel & Richardson, 1993), they need not be insensitive to mutual interactions between such parts. Quite on the contrary, sophisticated mechanistic frameworks involve both a *de*composition and a *re*composition of phenomena, as well as a detailed characterization of the overall context (Bechtel, 2009).
Mind and brain can also be carved up independently of one another. This may result in several conflicting taxonomies. For example, the faculties posited by Thomas Reid differ from the systems described in modern psychology textbooks, and both are different from the taxonomies of folk psychology, which nonetheless underlie much of our everyday language about mind and behavior. Brains too can be carved along many joints: for instance, brain cartographers can distinguish among lobes, on the basis of gross spatial properties; or they can distinguish brain regions, according to either a single or a mix of properties such as their histology or their connectivity profiles (Mundale, 2009). Assumption (3) offers a possible way out from the underdetermination of both kinds of ontology, providing a (further) criterion for choosing one out of these many possible carvings: namely, it prescribes that we should choose

2 Though they are very opaque windows. While we cannot address these topics here, it is worth remembering that each technique has its limitations, and that the assumptions necessary in order to interpret the data are far from being theoretically neutral. Rather than conceiving of them as tools for *seeing* brain activity, it is thus more prudent to conceive of neuroimaging techniques as tools for *inferring* brain activity. See e.g. Roskies (2008) and Klein (2010).
3 Notice that, while obviously easier to reconcile with a materialist metaphysics, assumption (3) does not entail it: all that is required is some kind of psycho-neural bridge principle. See Nathan & Del Pinal (2016).

those that warrant systematic mappings between entities of each domain.[4]
Whilst assumptions (1-3) are arguably shared by every Cognitive Neuroscientist, they are often construed in different ways. In the following sections I distinguish and compare between two such ways, giving rise to two different ontological frameworks.

2. The Neo-Phrenological Framework of Cognitive Neuroscience

During the early years, and until recent criticisms have cast shadows upon them, the ontological desiderata underlying most studies in cognitive neuroscience were efficaciously stated by neuroscientists Price & Friston (2005), who claim that we should aim for "a systematic definition of structure–function relations whereby structures predict functions and functions predict structures" (p. 263). They stress that both functions and structures can be described at multiple levels of abstraction, and propose that the best level of abstraction is that which allows us to assign a single function to each structure (see also Rathkopf, 2013). Within this framework, assumptions (1-3) are thus specified as such:

(1a) the mind can be decomposed into distinct mental entities $(m_1, m_2, \dots m_n)$, i.e. *mental functions*;
(2a) the brain can be decomposed into distinct neural entities $(n_1, n_2, \dots n_n)$, i.e. *neural structures*;
(3a) a *one-to-one mapping* holds between each mental entity and a given neural entity $(m_1 \leftrightarrow n_1, m_2 \leftrightarrow n_2, \dots m_n \leftrightarrow n_n)$.

This framework has been derogatorily dubbed *The New Phrenology* by some critics (notably, this is the title of Uttal's 2001 book). Indeed, this label is not totally unreasonable: the practice of mapping mental entities to some neural entity was introduced by the phrenologists, and continued through the twentieth century thanks to the work of physicians such as Broca and Ferrier, who in turn passed it to modern scientists. Therefore, I shall call this ontological framework "neo-phrenology".
However, modern neo-phrenologists differ from their ancestors in that they have thoroughly refined their ontology; they no longer try to associate the thickness of bumps on the skull to some of the disputable mental faculties posited by Gall, but rather seek to map inner neural regions to some mental function, i.e. usually psychological constructs that best fit the models of Cognitive Psychologists (Zawidzki & Bechtel, 2004).
Still, in our current ontology, virtually any attempt to map mental functions onto neural structures revealed a many-to-many mapping. Price and Friston are obviously aware that current theories are far from being as well ordered as they desire. But their claim is not meant to describe the current ontology. Rather, it is meant to play a heuristic function (Bechtel & McCauley, 1999), i.e. to prescribe how future ontologies should look (*fig. 1*).

4 Usually, reductionist philosophers assume that the relation between psychology and neuroscience is merely or mostly bottom-up, that is, it is only/mainly neuroscience that influences the psychological categories (e.g. Bickle, 2006). However, historically brain cartographies have been widely driven by considerations concerning the function, i.e. ultimately the psychological role of brain regions (see Hatfield, 2000). In drawing his notorious brain maps, Brodmann relied on histological criteria because "tissue elements of uniform specific structure, whether they are limited to a large or small cortical field or diffusely distributed over the whole cortex, must also have a uniform physiological function, and thus [...] such elements are to be regarded as not only morphologically but also functionally equivalent" (2006, p. 5). Functional significance, along with anatomical criteria, is still at play in modern mappings (notably, Glasser *et al.*, 2016).

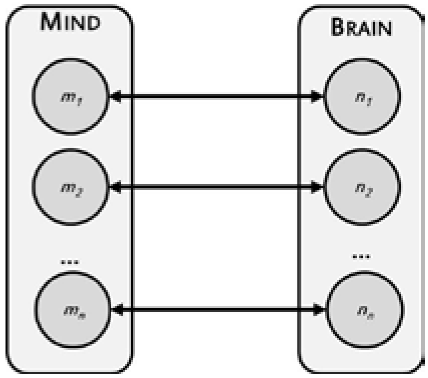

Figure 1. A stylized representation of the ideal ontology envisaged by the neo-phrenology: both Mind and Brain are decomposed into a set of discrete entities, each one standing in a one-to-one mapping with a member of the other set.

The moral of their story then is that, whenever there is no function-structure one-to-one relation, either the function or the structure (or both) are to be reformed or discarded. To put it into Lakatos's (1970) terms, rather than giving up the core assumptions (1a-3a), we can "put the blame" on one of the following auxiliary assumptions:

(4) the correct set of *mental functions is* $M = \{m_1, m_2, ... m_n\}$;
(5) the correct set of *brain structures* is $N = \{n_1, n_2, ... n_n\}$.

Either (4) or (5) or both are then replaced by either one or both the following auxiliary assumptions:

(4*) the correct set of *mental functions is* $M^* = \{m^*_1, m^*_2, ... m^*_n\}$;
(5*) the correct set of *brain structures* is $N^* = \{n^*_1, n^*_2, ... n^*_n\}$,

where 'the correct set' means 'the set that is better at preserving (3)' (*fig. 2*).

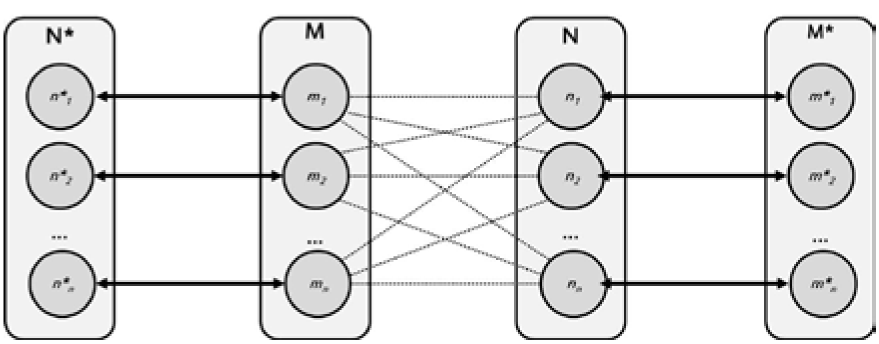

Figure 2. Whenever one-to-one mappings cannot be established between the members of a set of mental entities *M* and those of a set of neural entities *N*, neo-phrenology tells us to discard either (or both), in favor of different sets *M** and/or *N**, in order to reestablish one-to-one mappings between mental and neural entities.

For instance, the left posterior lateral fusiform gyrus has been associated with processing of visual words, and was thus called Visual Word Form Area. However, as Price and Friston report, this area is also engaged in processing the visual attributes of animals, and in tactile-visual integration. Thus, they propose that a better functional label for that area is one that characterizes its working at a more abstract level: *sensorimotor integration*. Similarly, since the left inferofrontal gyrus (Broca's area) is found to be activated in many domains outside the linguistic one, Tettamanti & Weniger (2006) propose to reconceive its function as a "supramodal hierarchial processor".

Yet, as Klein (2012) observes, a functional label as vague as 'sensorimotor integration' is not very informative, since "at some level of abstraction, that's what nearly all of the cortex does" (p. 5). Thus, rather than some too-abstract-to-be-purposeful definition, Klein argues that functional labels for brain regions are only meaningful when they are considered in the light of some given context. He then construes this context as neural context, i.e. the set of areas that are co-activated. In the end, he comes up with a prudent endorsement of an ongoing shift of emphasis about the level at which functions and structures ought to be mapped: from single regions to many regions gathered into functional networks (Bressler & Menon, 2010).

Such a shift was also fostered by the development of Multi-Variate Pattern Analysis (MVPA), statistical techniques that allow scientists to compare activation patterns spread across the whole brain, instead of single regions (Haxby *et al.*, 2014). The take-home message of such techniques is that the functional activity of the brain is better understood when complex arrays of activity are taken into account. Most of these MVPA employ machine learning techniques to perform what has come to be known as brain reading: in order to train a classifier to predict which cognitive function is performed from the corresponding neural activation, researchers "feed" it with functionally labelled patterns. After that, classifiers are shown unlabeled patterns of neural activation, and are asked to "guess" which one (among a given set) is the correct functional label, based on the similarity to the training set; a task in which they often succeed far above the chance level, sometimes even across different subjects. By applying such a method to a large database, Lenartowicz *et al.* (2010) investigated whether some cognitive constructs from the domain of control functions were mapped consistently to some specific and distinguishable activation pattern. As the classifier struggled to discriminate the pattern associated with "task switching" from those associated with some other constructs, the researchers suggest replacing "task switching" with some other constructs that map more neatly with some discriminable neural basis.

3. A Post-Phrenological Framework

Notwithstanding the reforming attempts of some researchers, many-to-many function-structure mappings are still ubiquitous. While Bechtel & Mundale (1999) claimed that Multiple Realizability vanished in the light of the empirical success of comparative Neuroscience, many mental functions turned out to be *degenerate*, i.e. they can be implemented by multiple brain structures – which, according to Figdor (2010), might sometimes count as multiple realizability. Meanwhile, while not *equipotent*, all brain regions are found to exhibit *pluripotency*, i.e. to be involved in the implementation of several (apparently) different mental functions. Degeneracy and pluripotency contradict the one-to-one mapping assumption (3a), and therefore neo-phrenology regards them as two anomalies to deal with by revising the ontology. Yet, despite many years of reformistic efforts, including the shift from regions toward networks, degeneracy and pluripotency seem to be here to stay. The neuroscientist Pessoa (2014) pessimistically upholds that

> the attempt to map structure to function on a one-to-one manner in terms of networks will be fraught with similar difficulties as the one based on brain regions […] – the

> problem is simply passed along to a higher level. Thus, two distinct networks may generate similar behavioral profiles ([...] many-to-one); a given network will also participate in several behaviors (one-to-many) (p. 408).

Thus, some thinkers bite the bullet and try to sketch an ontological framework that does without one-to-one mappings. The most mature formulation up to date is presented in Anderson's 2014 book *After Phrenology* (see also Barrett & Saptute, 2013). Stressing the often-neglected phenomenon of neural plasticity, both at long and at short timescales, Anderson argues that brain structures have no such things as "intrinsic functions". Rather, they get their functional significance depending on their structural characteristic, and on the partnership they institute with other neural structures, gathering into transient neural assemblies. Structures that usually play some role are thus commonly redeployed for other functions – a principle that Anderson dubs *neural reuse*.

Since neural reuse can hardly if ever be made consistent with one-to-one mapping, Anderson sketches a new protocol for linking mind and brain. The relation between brain and mind, he claims, could be established by measuring the functional dispositions of each neural structure (be it a region or a network), i.e. the likelihood that such structure gets (significantly) activated when some task is undertaken that bears the label of a given function or cognitive domain. In other words, (3a) is rejected and (3) is rather construed as (3b):

> (3b) for each structure *s*, and for each function (or domain) $f_1, f_2, ... f_n$, there is a probability *P* that *s* is engaged ($P(s,f_1), P(s,f_2), ... P(s,f_n)$).

By exploiting available databases of neuroimaging data, it is possible to provide robust measurements of each structure's disposition to be engaged in various functions/domains, obtaining what Anderson calls a *functional fingerprint*. Statistical dimension techniques could then be used to "dig out" similarities, and cluster related functions into what he calls *Neuroscientifically Relevant Psychological (NRP) factors* – which Anderson presents as an analogue of "personality traits" for neural structures.

An interesting feature of this approach is that, because the only direction of ontological revision is from brain data to mental categories, it might spur radical revisions in psychological taxonomies, leading to notice similarities and distinctions that could have been otherwise foreshadowed by our folk psychological prejudices.

Groundbreaking as it is, this post-phrenological framework has its weak spots too. For instance, it is very conservative with respect to neural ontology (McCaffrey & Machery, 2016), i.e. it does not offer incentives to find new and more interesting neural entities (see below). Moreover, whenever different functions are inadvertently given the same functional label (which is frequent across different labs; see Sullivan, 2016), Anderson's big data approach seems ill suited to recognize the distinction (Kaplan & Craver, 2016). Nonetheless, since it is still in its infancy, no doubt post-phrenology will arguably address these and other shortcomings. Would neo-phrenology be completely supplanted by post-phrenology, or can the old framework still be saved somehow?

4. What Reforms (If Any) Can Save Neo-Phrenology?

While many are pessimistic about the fate of one-to-one mapping (3a), I think that it is worth attempting to rescue it, in order not to throw the baby out with the bath water. As stressed by Bechtel & McCauley (1999), even gross function-structure mappings, inasmuch as they approximate the truth, may provide the basis for posing further research questions. For instance, while a claim such as "vision happens in the striate cortex" is largely imprecise, it was a starting point for further finer-grained models. Thus, once properly refined, one-to-one

mapping might still represent a viable heuristic for prompting ontological revisions.
Whilst much emphasis has been placed upon revising mental entities, revising the neural entities might also be possible. Indeed, while both lesion studies and neuroimaging biased researchers toward assuming contiguous brain regions as the basic entities, due to their spatial resolutions, it is possible that two (or more) functionally coherent and intertwined neural populations can co-exist within what is currently classified as "the same" brain region. McCaffrey (2015) stresses that various areas labelled as multifunctional (i.e. pluripotent) might differ in their mechanistic organization, thus requiring various explanatory strategies: some areas contain different structures, and thus require the abovementioned "divide-and-conquer" approach; others conserve the same role in different tasks, thus qualifying for a more abstract redefinition of their function such as that advocated by Price & Friston (2005, see §3); but for yet other areas, context-sensitive mappings are unavoidable.
Therefore, perhaps replacing some specific set of either or both of mental function and/or of neural structure (i.e., replacing either or both assumptions 4a and 5a with 4b and 5b) might be insufficient. A more productive approach may be that to stop seeking one-to-one mappings *between structures and functions*, and rather pick some wholly different kind of *neural* and/or *mental entities*. After all, this is how phrenology turned into *neo*-phrenology. For instance, after having stressed that a same set of brain regions could support several functions, Pessoa (2014) speculates that different functions can be due to differences in the strengths of the connections between them, or by the different time course of their activation.
Similarly, in Viola & Zanin (forthcoming) I proposed another kind of reform in neural ontology, i.e. dropping

> (2a) the brain can be decomposed into distinct neural entities $(n_1, n_2, \ldots n_n)$, i.e. *neural structures*,

and replacing it with

> (2b) the brain can be decomposed into distinct neural entities, i.e. *neural processes* $(p_1, p_2, \ldots p_n)$,

and further specifying that

> (6) each neural process *p* is defined as the activity of some (set of) neural structure(s) *n* when acting according some way of working *w* $(n{=}s,w)$.

The notion of a "way of working" is introduced in order to account for the intuition that the same part (or set of parts) of a mechanism can implement multiple functions (Bechtel & Abrahamsen, 2005). This proposal thus implies that the functions of brain mechanisms, rather than their parts, should be the *relata* of one-to-one mappings.
Admittedly, "way of working" is a vague notion. Given that, paraphrasing Price and Friston, we want to be able to predict functions from structures *plus something* (and vice versa), we cannot discriminate distinct ways of working on the basis of their contribution to mental functions, since that would entail a circularity. Rather, we need some purely neural marker for discriminating ways of working. In Viola & Zanin (forthcoming) I tentatively adopted the suggestion advanced by Siegel *et al.* (2012, p. 121) that "frequency-specific correlated oscillations in distributed cortical networks may provide indices, or 'fingerprints', of the network interactions that underlie cognitive processes". Therefore, I proposed that different oscillatory patterns represent distinct ways of working.

The expected outcome of such framework might be represented as in *fig. 3*:

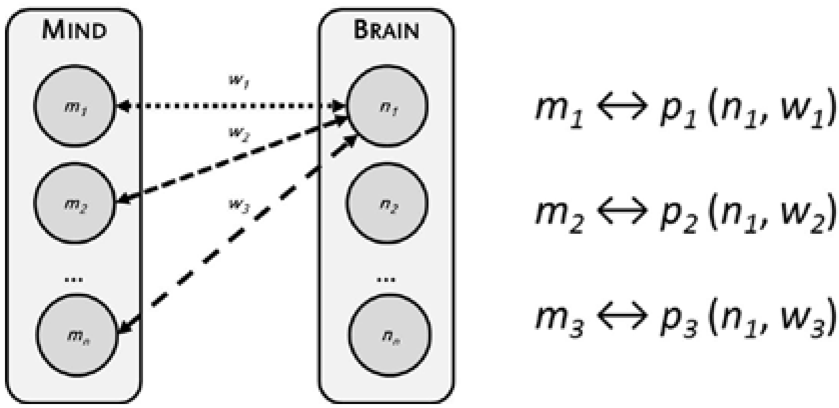

Figure 3. The revision of the neo-phrenology proposed in Viola & Zanin (forthcoming). The *m*s, *n*s,*w*s, and *p*s represent, respectively, mental functions, neural structures (or sets of neural structures) and ways of workings, which I propose to construe in term of neural oscillations. Mental functions cannot be put into one-to-one correspondences with (sets of) neural structures anymore; rather, (sets of) neural structures become one of the two properties that define neural processes, along with a specific way of working.

Such a strategy allows us to reconcile pluripotency and one-to-one mapping. However, notwithstanding this specific proposal, I am more concerned with demonstrating that, notwithstanding its anomalies, neo-phrenology (of some evolved version, say neo-phrenology*) can still play a role in Cognitive Neuroscience.

Ultimately, as suggested by Anderson and Pessoa (among others), it is possible that the Brain-Mind relation turns out to be so complex that no one-to-one mapping whatsoever will be tenable for organizing existent knowledge. In other words, it is not implausible that one-to-one mappings will eventually vanish from the context of justification. Even so, hypothesizing and testing some sophisticated one-to-one mapping might still play some role in the context of discovery, e.g. by challenging the researchers to reshape the taxonomical landscapes of mind and brain.

REFERENCES

Anderson, M.L. (2014). *After phrenology*. Cambridge, MA: MIT Press.

Barrett, L.F., & Satpute, A.B. (2013). Large-scale brain networks in affective and social neuroscience: Towards an integrative functional architecture of the brain. *Current opinion in neurobiology*, 23(3), 361-372.

Bechtel, W. (2009). Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology*, 22(5), 543-564.

Bechtel, W. & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), 421-441.

Bechtel, W. & McCauley, R.N. (1999). Heuristic identity theory (or back to the future): The mind-body problem against the background of research strategies in cognitive neuroscience. In Hahn, M., & Stoness, S.C. (Eds.), *Proceedings of the 21st annual meeting of the cognitive science society*. Mahwah, NJ: Erlbaum, 67-72.

Bechtel, W. & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of science*, 66(2), 175-207.
Bechtel, W. & Richardson, R.C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Boston, MA: MIT Press.
Bickle, J. (2006). Reducing mind to molecular pathways: Explicating the reductionism implicit in current cellular and molecular neuroscience. *Synthese*, 151(3), 411-434.
Bressler, S.L., & Menon, V. (2010). Large-scale brain networks in cognition: emerging methods and principles. *Trends in cognitive sciences*, 14(6), 277-290.
Brodmann, K. (2006). *Localization in the Cerebral Cortex* (Eng. Transl. by L. J. Garvey.). New York, NJ: Springer. (Original work published in 1909).
Figdor, C. (2010). Neuroscience and the multiple realization of cognitive functions. *Philosophy of Science*, 77(3), 419-456.
Glasser, M.F., Coalson, T., Robinson, E., Hacker, C., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jankinson, M., Smith, S.M., & Van Essen, D.C. (2016). A Multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), 171-178.
Hatfield, G. (2000). The brain's "new" science: Psychology, neurophysiology, and constraint. *Philosophy of Science*, 67, S388-S403.
Janssen, A., Klein, C., & Slors, M. (2017). What is a cognitive ontology, anyway?. *Philosophical Explorations*, 20(2), 123-128.
Kaplan, D.M., & Craver, C.F. (2016). A registration problem for functional fingerprinting. *Behavioral and Brain Sciences*, 39, 15-16.
Klein, C. (2010). Images are not the evidence in neuroimaging. *The British Journal for the Philosophy of Science*, 61(2), 265-278.
Klein, C. (2012). Cognitive ontology and region-versus network-oriented analyses. *Philosophy of Science*, 79(5), 952-960.
Lakatos, I. (1970). Falsification and the Methodology of Scientific Research Programmes. In Lakatos, I., Musgrave, A. (Eds.), *Criticism and the Growth of Knowledge*, Cambridge: Cambridge University Press, 91-195.
Lenartowicz, A., Kalar, D.J., Congdon, E., & Poldrack, R.A. (2010). Towards an ontology of cognitive control. *Topics in Cognitive Science*, 2(4), 678-692.
McCaffrey, J.B. (2015). The brain's heterogeneous functional landscape. *Philosophy of Science*, 82(5), 1010-1022.
McCaffrey, J.B. & Machery, E. (2016). The reification objection to bottom-up cognitive ontology revision. *Behavioral and Brain Sciences*, 39, 16-18.
Mundale, J. (2009). Epistemic Preliminaries: Normative Priorities and Neuropsychological Kinds. *Humana Mente*, 11, 1-9.
Nathan, M.J., & Del Pinal, G. (2016). Mapping the mind: bridge laws and the psycho-neural interface. *Synthese*, 193(2), 637-657.
Pessoa, L. (2014). Understanding brain networks and brain organization. *Physics of life reviews*, 11(3), 400-435.
Poldrack, R.A. (2010). Mapping mental function to brain structure: how can cognitive neuroimaging succeed?. *Perspectives on Psychological Science*, 5(6), 753-761.
Rathkopf, C.A. (2013). Localization and intrinsic function. *Philosophy of Science*, 80(1), 1-21.
Roskies, A.L. (2008). Neuroimaging and inferential distance. *Neuroethics*, 1(1), 19-30.
Siegel, M., Donner, T.H., & Engel, A.K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience*, 13(2), 121-134.
Squire, L.R. (1992). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. *Journal of cognitive neuroscience*, 4(3), 232-243.
Sullivan, J. (2016). Neuroscientific kinds through the lens of scientific practice. In C. Kendig

(Ed.), *Natural Kinds and Classification in Scientific Practice*. New York: Routledge, 47-56.
Tettamanti, M. & Weniger, D. (2006). Broca's area: a supramodal hierarchical processor?. *Cortex*, 42(4), 491-494.
Uttal, W.R. (2001). *The new phrenology: The limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT press.
van Orden, G.C., & Paap, K.R. (1997). Functional neuroimages fail to discover pieces of mind in the parts of the brain. *Philosophy of Science*, 64, S85-S94.
Viola, M. & Zanin, E. (forthcoming). The standard ontological framework of cognitive neuroscience: Some lessons from Broca's area. *Philosophical Psychology*, DOI: 10.1080/09515089.2017.1322193.
Zawidzki, T. & Bechtel, W. (2005). Gall's legacy revisited. Decomposition and localization in cognitive neuroscience. In Emering, C.E. & Johnson, D.M. (Eds.), *The mind as a scientific object: Between brain and culture*. New York: Oxford University Press, 293-318.