



# Beyond the Platonic Brain: facing the challenge of individual differences in function-structure mapping

Marco Viola<sup>1</sup>

Received: 20 November 2019 / Accepted: 14 September 2020  
© The Author(s) 2020

## Abstract

In their attempt to connect the workings of the human mind with their neural realizers, cognitive neuroscientists often bracket out individual differences to build a single, abstract model that purportedly represents (almost) every human being's brain. In this paper I first examine the rationale behind this model, which I call 'Platonic Brain Model'. Then I argue that it is to be surpassed in favor of multiple models allowing for patterned inter-individual differences. I introduce the debate on legitimate (and illegitimate) ways of mapping neural structures and cognitive functions, endorsing a view according to which function-structure mapping is context-sensitive. Building on the discussion of the ongoing debate on the function(s) of the so-called Fusiform "Face" Area, I show the necessity of indexing function-structure mappings to some populations of subjects, clustered on the basis of factors such as their expertise in a given domain.

**Keywords** Neuroscience · Cognitive ontology · Individual differences · Functional specialization · Fusiform face area · Contextualism

## 1 Cognitive neuroscience and the Platonic Brain

Cognitive neuroscience handbooks often contain sentences such as: "Cognitive neuroscientists attempt to understand the relationship between the brain and mind" (Banich and Compton 2018: p. 3). Other handbooks recount "how mental processes such as thoughts, memories and perceptions are organized and implemented by the brain" (J. Ward 2015: p. 1) and introduce the methods to investigate such processes. An early step in many neuroscientific enterprises is functional localization: the ascription of a psychological label to some neural structure based on its purported involvement

---

✉ Marco Viola  
marco.viola@unito.it

<sup>1</sup> Department of Philosophy and Education, University of Turin, via Sant'Ottavio 20, Turin, Italy

in some behavioral task. In most cases, functional ascriptions are a matter of sharp contention. However, a few functional ascriptions are taken as received wisdom.

Take the striate cortex, also known as V1. Indeed, almost everybody agrees that V1 is in the business of processing visual information—hence its alias *primary visual cortex*. This is unsurprising, given that V1 receives afferents fibers from the geniculate nucleus, which in turn receives fibers from the retina. It would thus be safe to say that V1 deals with visual information *in most cases*.

But what functional role does this *primary visual cortex* play in blind subjects? Since they do not process any visual information, were brain areas rigidly ‘tied’ to their functional destiny, we would expect that V1 simply does *nothing*. However, this seems not to be the case. In fact, fMRI scans reveal activity in blind subjects’ V1 (and in other cortices) when they read Braille; and impairment in their performance may be obtained by perturbing V1 via transcranial stimulation (Kupers et al. 2007). A growing number of such cases suggest that several ‘visual’ areas are way too busy in non-sighted subjects to be actually and uniquely *visual*.

To account for such findings, some scholars have been ready to retract the ascription of visual functionality to V1. Notably, according to the meta-modal hypothesis (Hamilton and Pascual-Leone 2001),<sup>1</sup> V1 and other purportedly ‘visual’ cortices still perform the same computation, e.g. shape recognition, across various modalities. Having described striking similarities in the neural activity of the ‘visual’ cortices of sighted and congenitally blind subjects on several circumstances, Ricciardi and Pietrini (2011) argue for a more abstract characterization of the function of these structures.

However, this ‘digging deeper’ strategy (which I shall examine more closely in §3.2 under the term “functional abstraction”) may sometimes be unsatisfactory. In fact, in blind (as opposed to non-blind) individuals, several “visual cortices” activate during tasks such as listening to a spoken language or doing math equations, which are hardly accountable in terms of “shape recognition” (Bedny 2017).

The story might get even more complex—and intriguing. As bizarre as it may seem, some congenitally blind individuals devise an ingenious strategy to compensate for the lack of visual input when navigating space: echolocation. They repeatedly make a click sound with their tongue and use the echo to identify obstacles and other objects in their surroundings (Downey 2016). Notably, fMRI reveals that echolocation activates areas in the vicinity of V1 (Thaler et al. 2011).

Now, undoubtedly several metamodal labels can be devised to account for the activity of some visual areas in tasks such as reading Braille or echolocating. Yet, by dismissing the differences between such tasks, these labels are likely to miss something quite important: namely, that this very area is involved in doing something different in the two cases. However, it is worth stressing that the area’s engagement in visual processing simply *cannot* occur in blind subjects; nor can its engagement in echolocation occur in sighted subjects.

This simplified story was meant to emphasize that *brains and cognitive abilities exhibit important differences across human beings*. This is a well-known fact, of course. Yet, as it often happens to facts when they are too well-known, I contend that its implications have received less attention than they deserve. Recall the quotes from

<sup>1</sup> Sometimes also called *supramodal* hypothesis.

the handbook that open of this section: the task they set for cognitive neuroscience is to “understand the relationship between *the* brain and mind” (Banich and Compton 2018: p. 3. Italics mine); to unravel “how mental processes [...] are organized and implemented by *the* brain” (J. Ward: 1. Italics mine)—as if *one* “Platonic” model of the brain might explain everything we need to know about the functional organization of *every* actual (and possible) human *brains*.

In this paper, I have three aims. First, I aim to explain the rationale behind the “Platonic Brain Model” that treats individual differences in brain functional organization as noise to be zeroed out (§2). Second, I highlight some of its limitations. To do so, I provide an upshot of the extant philosophical literature on (the difficulty of) mapping cognitive functions onto neural structures (§3), and present a case study in which a Platonic Brain Model falls short of accounting for all the relevant function-structure mappings, namely the case of the Fusiform ‘Face’ Area (§4). Third and last, I discuss the virtue of, and begin to sketch a recipe for, a post-platonic neuroscience, in which the function-structure mappings are indexed to certain populations of subjects, e.g. subjects who share a certain degree of expertise in some domain (§5).

## 2 Dealing with individual differences: the Platonic Brain Model

Individual differences in human brains are particularly evident when we consider extreme pathological cases like those hydrocephalic subjects<sup>2</sup> who exhibit no sign of mental impairment (Nahm et al. 2017), or when we compare children and adults’ brains (e.g. Wilke, Schmithorst and Holland 2003). However, aided by the development of tools such as Voxel-Based Morphometry (Whitwell 2009), neuroscientists have become growingly interested in investigating differences even outside the clinical or the developmental domain. In a now classic study, Taxi drivers in London were shown to exhibit a peculiar increase and decrease in volume in the posterior and anterior hippocampus, respectively, which correlates with the amount of time spent driving taxi (Maguire et al. 2000). Hippocampus’ thickness also seems positively correlated with the socioeconomic status, like many other cortical and subcortical structures (Farah 2017).

But individual differences are not limited to anatomy. Instead, they extend to physiology, altering the functional landscape of brain organization, sometimes profoundly. As people get older, the same tasks that used to elicit a right-lateralized activity of the prefrontal cortex tend to bring about a more bilateral activation. In other tasks, aging correlates with increased activity in the anterior regions, and diminished activity in posterior regions, even when performance rests unaltered (see De Brigard 2017 and the references therein). In people who played Pokémon videogames in their youth, as opposed to those who did not, looking at Pokémon consistently activates a portion of the occipitotemporal sulcus whose localization is consistent among subjects (Gomez et al. 2019). Similarly, merely looking at *this* very sentence (or at any string of letter) should suffice to engage a part of the left ventrotemporal cortex in literate subject-

---

<sup>2</sup> Hydrocephalus indicates a clinical condition in which the cranium is enlarged due to excess cerebrospinal fluid pressure.

s—hence its alias visual word form area. But the same activity would not show up in illiterate subjects (Dehaene et al. 2010).

On the one hand, psychological science never ignored individual differences. Indeed, according to Danziger (Danziger 1990/1994), investigating individual differences was precisely the core business (sometimes literally!) of one of the most prominent tradition during the early decades of scientific psychology, namely the Galtonian tradition (after Francis Galton). Setting the stage for modern psychometrics, they studied the relationships between indicators of different nature—sociological, anatomical, and psychological (e.g. performance in tests)—in individual subjects and groups (sometimes yielding to questionable conclusions). However, throughout the history of scientific psychology, and later also in the history of neuroscience, individual differences were mainly conceived *quantitatively*. That is, the difference between subjects is represented by scoring differently on several scales, e.g., IQ scores, the Big Five personality traits, the volume of some brain areas such as the posterior hippocampus.

Far less attention has been paid to *qualitative* differences in cognitive skillset or brain organization. Most studies in cognitive neuroscience are taken to be generalizable for *almost* every human brain. When they are too dissimilar from the average, “outliers” are treated as exceptions or noise, whose explanation is left to clinicians.

A rationale for this framework can be found in a well-known methodological paper by Caramazza (1986). Before the rise of neuroimaging techniques that allow studying the human brain *in vivo*, lesion studies were the bread and butter of mind and brain sciences. Cognitive neuropsychologists studied the cognitive impairments of patients who suffered a focal brain lesion not merely for the sake of understanding a clinical case, but also to unravel the workings of healthy brains.

As a neuropsychologist, Caramazza reflects upon the inferences that can be drawn from impaired to unimpaired cognitive systems, making the underlying assumptions explicit. He states that

A crucial feature of [any model of human mind/brain] is the assumption of universality; that is, the assumption that [the model] is true of “normal” human mind/brains in general and, therefore, of any individual normal mind/brain (Caramazza 1986: p. 49).

Caramazza clearly does not ignore individual differences. Instead, he seeks them as something to *rule out* for the sake of generality, claiming that “we are going to have to place some restrictions on what will count as ‘normal human mind/brain’” (*ibid.*). He goes as far as to claim that “if we were not to accept the assumption of universality, we would negate the possibility of scientifically investigating the mind/brain” (*ibid.*).

Crucially, the assumption of universality is the cornerstone of the single-patient method, for it justifies a generalization of the model (however falsifiable) based on evidence gathered from a single-patient and then extended to every “normal” subject. For Caramazza, the single-patient method is superior to the group-method, given that the latter may engender a spurious clustering of patients based on syndromes rather than on etiology.<sup>3</sup> Indeed, what might *prima facie* appear as the *same* impairment

<sup>3</sup> Although an extensive comparison of single-case versus group-case methodology is beyond the scope of this paper, for the present discussion it is worth noting that a major matter of contention concerns the selection of a methodology that would avoid the confounds of idiosyncratic performances. For a focus on the controversy see Shallice (1988: ch. 9–10; 2015); Bartolomeo and colleagues (2017).

due to common symptoms might be caused by lesions in different sites, and these important differences are easily overlooked when averaging across patients.

However, Caramazza clearly stresses that the universality assumption also holds for group studies:

The justification for using the performance of groups of subjects in our experimental investigations is based on the assumption that the averaged performance of the group essentially reflects the performance of any individual in the reference population from which the group was drawn. Thus, any conclusions arrived at for the group of subjects tested will be assumed to be true of all individuals in the reference population. This argument is only valid if the assumption of universality is true (Caramazza 1986: p. 50).

Let us now make a twenty-years leap forward in time. The rise of hemodynamic neuroimaging techniques has reshaped the scientific landscape. Lesion studies are no longer the favorite tool of cognitive neuroscientists (that is how they get called now), as researchers can now glimpse at brain activity in vivo in healthy subjects. As these subjects are healthy, the syndrome-etiology worry that concerned Caramazza ceases to apply. However, the universality assumption is still there, albeit in a new guise.

At the end of a lengthy discussion aimed at formalizing and justifying the inferences drawn from neuroimaging data to function-structure mappings, Henson (2005) addresses the possible objection according to which “different developmental trajectories (based, for example, on the order of exposure to different stimuli, such as language) may result in different ‘final’ function–structure mappings” (Henson 2005: p. 224). His reply reads as follows:

the only solution is the nomothetic solution: to define “normal” psychological functions—in the healthy adult, for example—and to assume that there is [a] single “normal” mapping of these functions to the brain. This then becomes an empirical question: if the function–structure mapping changed [...] so dramatically as a function of different developmental trajectories, one would not observe reliable differences in imaging data when averaging over random samples of normal individuals. The fact that neuroimaging can produce reliable and reproducible differences suggests that there is a normal (default) function–structure mapping. (Henson 2005: pp. 224–225).

Researchers working under such an assumption thus look at individual differences as a possible threat to generalizability. Two strategies are employed for ensuring generalization: (a) adopting some *screen-off criteria* to exclude “abnormal subjects”, and (b) establishing the perimeter of normal results by *averaging* among data gathered from normal subjects.

*Screen-off the abnormal*—Usually, the “participants” section of neuroscientific articles reassures the reader that screen-off criteria have been respected. Typically, the description of participants boils down to a few details, e.g. whether they are healthy, adults, right-handed, and normal-sighted (to be able to see the stimuli in the scanner). As long as they are respected, the researchers go on describing the findings as if they pertain to a Platonic Brain, suggesting the reader that they generalize to virtually any human being—unless she fails to meet the screen-off criteria.

*Averaging*—Even after the ‘abnormal’ brains are screened off, normal brains do come with different sizes, shapes, and several other idiosyncratic features. Researchers have come up with several different techniques for aggregating data from multiple sub-

jects. A common strategy, that Zina Ward (2019) dubbed the *cartographic* approach, consists in plotting activations from multiple subjects' brain onto a common reference space or template. But other approaches exist, that employ functional criteria together with anatomical ones to align brain activations of multiple subjects (cf. fn. 10).

Working under the assumption of the Platonic Brain Model was not unreasonable. After all, brain development is often based on somewhat similar genetic blueprints, and many experiences are similar enough for the whole humankind to result in the same brain organization.<sup>4</sup> This theoretical bet has proven fruitful not only for phylogenetically old mental activities, but also for relatively recent ones: in fact, the location of the visual word form area is stable across subjects (Dehaene et al. 2010; see Rathkopf (forthcoming) for a discussion of some implications).

Nevertheless, working under the assumption of a Platonic Brain Model comes with some theoretical costs and risks. For instance, Zina Ward (2019) shows how, when different brains are aligned with respect with some features (e.g. microanatomical properties) in neuroimaging studies, this may result in misalignment in other features (e.g. functional properties), and concludes that the choice of aligning technique depends upon the question one has in mind—i.e. there is no one-size-fits-all techniques to align brains.

Since the topic has many more implications that one can account for with a single paper, I will restrict my discussion to a specific topic. Like many neuroscientific endeavors start by hypothesizing a simple function-structure mapping (Bechtel and Richardson 2010), the following discussion will scrutinize Henson's claim that "there is a normal (default) function-structure mapping".


### 3 Mapping cognitive functions onto neural structures

#### 3.1 The quest for functional specialization

The question of whether the cortex is equipotential (i.e., its parts are functionally interchangeable) or functionally specialized (i.e., each part plays a specific function) dates back at least to the nineteenth century (Mundale 2002). Since then, scientists have been oscillating between the two positions. The longevity of the debate suggests that empirical data do not speak for themselves: they can be interpreted differently to defend one viewpoint or the other. The advocates of the functional specialization hypothesis have strenuously defended the following claim:

Functional Specialization [FS]: for each neural structure  $s$ , there is one functional description  $f$ , so that the process  $f$  occurs if and only if  $s$  activates.

<sup>4</sup> To what extent the whole mankind shares similar experience is matter of controversy. Indeed, in a now classical paper, Henrich, Heine and Norenzayan (2010) lament that most psychological studies, while making universal claims, are uniquely based on WEIRD subjects (Western, Educated, from Industrialized, Rich and Democratic countries). In response, a new wave of cross-cultural research in cognitive science flourished during the following decade, aimed at studying the diversities in human minds (Clark Barrett 2020).



**Fig. 1** Functional Specialization [FS] entails that the relationship between mental functions  $f$  and (the activity of some) neural structures  $s$  will result in a set of biconditionals  $s_x \leftrightarrow f_y$ . Counterevidence can affect both directions: when a structure underpins several functions (multifunctionality), it contradicts the statement that  $s_x \rightarrow f_y$ . When a function is conjunctively (distributed processing) or disjunctively (degeneracy) realized by several structures or set of structures, it contradicts  $f_y \rightarrow s_x$

*Prima facie*, FS is wiped out by current evidence: the literature reports several cases in which one or the other side of the biconditional fails to obtain. Such cases may be sorted into two major classes: cases of many-to-one function-structure relationships (a)<sup>5</sup>; and cases of one-to-many function-structure relationships—which come in two versions (b, c) (for a schematic recap, see Fig. 1). Let us briefly examine these cases in turn:

(a) Multifunctionality: one neural structure implements multiple functions.

Multifunctionality seems to be a ubiquitous property of brain structures, albeit some have more functions than others (Anderson and Pessoa 2011).<sup>6</sup> A few examples: the insula seems to be involved in a variety of functions, ranging from sensorimotor control to the experience of disgust and its perception in others, to the processing of salience for novel stimuli (Uddin et al. 2017). Broca’s area, once thought to be eminently involved in language production and/or in syntactic processing, has been reported to activate also during motor planning and music processing (Tettamanti and Weniger 2006). Left posterior lateral fusiform, corresponding to the abovementioned visual word form area, also seems to be involved in the interpretation of visual attributes of animals, as well as in tactile-visual interface (Price and Friston 2005).

(b) Distributed processing: one cognitive function relies upon the coordinated work of multiple brain structures.

No brain region operates in isolation from others: instead, each cognitive task seems to elicit the orchestrated activity of a coalition of neural regions. This principle has

<sup>5</sup> The expectation that each neural structure would have one *cognitive* function might also fail because some structures have functions that are non-cognitive in nature. Haueis (2018) has vindicated the existence and relevance of *non-cognitive functions* whose job is to maintain and protect the nervous system rather than processing behaviorally relevant information.

<sup>6</sup> Recognizing the pervasive multifunctionality of most brain structures famously led Poldrack (2006) to warn against the misuse of reverse inference, i.e. inferring that a cognitive process is ongoing because of the activation of a neural area traditionally associated to it. However, it is useful to point out that at least three distinct kinds of mechanistic organizations lurk beneath the (alleged or real) multifunctionality of biological structures (McCaffrey 2015; cf. later in the main text).

been discussed at length in the literature on neural correlates of specific emotions: each one correlates with a set of regions, some of which are shared by multiple emotions (Lindquist et al. 2012). More generally, distributed processing is gaining traction due to the predictive success of the multivoxel pattern analysis techniques, whose underlying assumption is that neural representations are codified by the activity of vast, sparse networks of neurons.

(c) Degeneracy: one cognitive function might be implemented by different neural underpinnings.

Another reason why functions may fail to map neatly onto a single neural structure is that they are *degenerate*, i.e. that they can be realized by different neural underpinnings (Noppeney et al. 2004). For instance, it has been argued that reading familiar words might be processed either lexically (reading whole words) or phonologically (reading letter-by-letter).<sup>7</sup>

While both distributed processing and degeneracy represent cases of many-to-one function-structure relationships, in the former a single function maps onto a *conjunction* of structures, whereas in the latter it maps onto their *disjunction*. It is also worth noting that these three issues are not mutually exclusive. On the contrary, they intersect with one another, thereby casting multiple shadows upon FS.

However, (a—c) do not necessarily *falsify* FS. Indeed, far from being a purely empirical claim about the brain, FS is best understood as a working hypothesis, a bet whose heuristic power guides the development of an integrated cognitive ontology (McCauley and Bechtel 2001). The history of psychology shows that the operations of the mind can be variously described, and multiple hypotheses exist as to how these operations can be decomposed. Moreover, although less obvious, there are several ways to partition brain structures. All things considered, in a case of putative falsification of FS *given a certain set of mental and neural categories*, a researcher has the option of preserving FS, and rather dispense with some of the mental or neural categories that fail to map onto each other, in favor of more brain-friendly mental categories and/or more mind-friendly neural categories.

Such an ontological refurbishing promises a huge payoff: namely, the possibility of predicting a function straight out of a structure's activation, and vice versa. The degree of success of this search for predictive power “can be embodied operationally in a cost function that reflects the prediction error in going from function to structure and back again” (Price and Friston 2005: p. 272).

For instance, some of the threats posed by (a—c) might be addressed renegotiating the neural ontology. To begin with, a putative multifunctional structure may sometimes be subdivided in smaller structures, each with different functions. So, for instance, the seemingly multifunctional insula turns out to be decomposable into four distinct sub-regions, each one with more specific functions (Uddin et al. 2017: p. 301; McCaffrey 2015). Also, Broca's Area has been reconceptualized as consisting in two structures: “one frontotemporal language-selective network and a second that

<sup>7</sup> One might object that what Noppeney and colleagues describe as a single function (reading) with many neural underpinnings is more naturally characterized as two distinct functions (phonological reading and lexical reading), each implemented by their own neural machinery. I address this objection later in the main text (§3.2).



belongs to the domain-general frontoparietal [multiple-demand] network” (Fedorenko and Blank 2020: p. 270), whose distinction was largely overlooked in group studies due to their anatomical proximity (Z. Ward 2019).

Reforms of the neural ontology can also go the other way, i.e. clustering structures. Indeed, in recent times, to address the fact that virtually every function correlates with the activity of multiple and scattered brain areas, the focus of neural ontology progressively shifted from regions to networks, where ‘network’ is often simply regarded as a set of brain areas (Klein 2012). More recently, De Brigard (2017) suggested that topological properties—as opposed to localization—are the relevant properties to individuate structures (and to link them to functions). However, for the purpose of this paper, we are mainly interested in a strategy to reform of cognitive ontology, that I discuss in the next sub-section.

### 3.2 Functional Abstraction and Neural Reuse

How to characterize the function of a neural structure has been a hot topic in many philosophical and scientific discussions. In many contexts (e.g. Sterelny 1990; Price and Friston 2005; Bergeron 2007; Klein 2012; Rathkopf (2013), forthcoming; Burnston 2016a, b; Genon et al. 2018; Barack 2019; see Fig. 2), a distinction has been proposed between (a) functional labels that are closer to folk psychological terms, that represent a specific kind of psychological information, or that are task/domain-bounded, and transparent with regard to the behavioral outcome, e.g. “reading”; and (b) domain-neutral functional labels indicating the operation/computation a neural structure performs, irrespectively of the outcome it subserves, e.g. “normalization” (Carandini and Heeger 2012). Despite their similarities, most notions introduced to account for the multi-layered nature of functional ascriptions carry some theoretical baggage with themselves, which makes them not perfectly overlapping. Thus, to avoid committing to any theory-laden lexicon, in the remainder of the paper I will employ the rather uncommitted notions ‘surface-functions’ and ‘deep-functions’ to refer to (a) and (b), respectively. The deliberate vagueness of these terms is meant to capture the heterogeneity of functional ascriptions among cognitive neuroscientists and philosophers. They are relational, contrastive terms, i.e. it only makes sense to speak of a surface-functions in opposition to a deep(er)-function.

However interpreted, the distinction between surface-functions and deep-functions comes in handy for defenders of FS: confronted with evidence that a given structure gets activated across different functions, they often admit that multifunctionality obtains at the level of surface-functions, only to withdraw into defending FS at the level of deep-functions. In Burnston’s (2016a) terms, they retreat from *traditional absolutism* to *computational absolutism*. The rationale of this strategy, which I dub *functional abstraction*, is expressed by Price and Friston: “structure–function relationships can be described at multiple levels [...] Each level may be appropriate in a different context [...] but it is more useful to label a region with a function that explains all patterns of activation” (2005: p. 268). Thus, they suggest that the surface-multifunctionality of the left posterior lateral fusiform are but various manifestations of deep-monofunctionality, i.e. sensorimotor integration. Similarly, in the face of the

**Function**  
WHAT a brain structure does

WHY it does that	Reference	HOW it does that
Ecological level	Sterelny (1990)	Computational level
Role	Bergeron (2007)	Workings
Task-bound function	Rathkopf (2013)	Intrinsic function
Face (of a functional polyhedron)	Genon et al. (2018)	Core (of a functional polyhedron)
Processing Model	Barack (2019)	Function
Surface-Function	<i>This paper</i>	Deep-function

**Fig. 2** A visual array of the lexicon employed by various scholars to characterize the distinction between ‘surface-functions’ and ‘deep-functions’. Notice that, as this comparison abstracts away from many theoretical details, terms of each column are not perfectly synonymous

surface-multifunctionality of Broca’s Area, Tettamanti and Weniger (2006) propose a single underlying deep-function, i.e. hierarchical processing. The reader might recall that a similar strategy was invoked by the proponents of meta-modal hypothesis at the beginning of this paper (§1).

The distinction between surface-functions and deep-function also allows preserving FS in the light of degeneracy. In the dual mode of reading example from the previous subsection, a broadly defined surface-function such as reading is exploited by means of two distinct (sets of) deep-functions, each one associated with its neural machinery (Figdor 2010; Polger and Shapiro 2016).

The pervasiveness of surface-multifunctionality (nicely documented by Anderson et al. 2013) is predicted by, and in fact inspired, several theories that fall under the umbrella-term neural reuse (Anderson 2010). Neural reuse boils down to the claim that complex behaviors do not hinge upon modules sculpted by evolution, but rather emerge by repurposing the same old neural structures into playing new roles by gathering into different functional allegiances (Anderson 2014).

Noticeably, while Anderson used to present neural reuse in terms of surface-multifunctional structures that nonetheless may preserve the same deep-function,<sup>8</sup> he later becomes skeptic that FS can be preserved even at the level of deep-function (Anderson 2014; for a discussion see Zerilli 2019). In his most recent views, not even the deep-functions of a neural entity are strictly predetermined: the fact that some neural structures tend to play the same (either deep and surface-)functional roles across

<sup>8</sup> For instance, in (Anderson 2010) he borrows the distinction between cognitive workings and cognitive roles from Bergeron (2007). Notice further that, while reuse of a same deep-function may nicely explain some instances of surface-multifunctionality, in other cases that are best kept separate surface-multifunctionality may be due to a same structures performing different deep-functions (Barack 2019; McCaffrey 2015). Still other putative cases of multifunctionality can be artifact due to failures in recognizing distinct structures, as it sometimes happens in neuroimaging with neural populations intertwined within the same voxels (Klein 2010; McCaffrey 2015).

individuals depend upon their *functional biases*—a weaker, dispositional counterpart of FS,—which may nonetheless end up being differently exploited depending on ontogeny. In a similar vein, the current neuroscience literature increasingly recognizes that the etiology of cortical functions depends upon an interplay of intrinsic facts pertaining cytoarchitectonic and connectivity and the experiences that sculpt its plasticity (see for instance Spunt and Adolphs 2017; de Beeck, Pillet and Ritchie 2019).

### 3.3 Contextualism

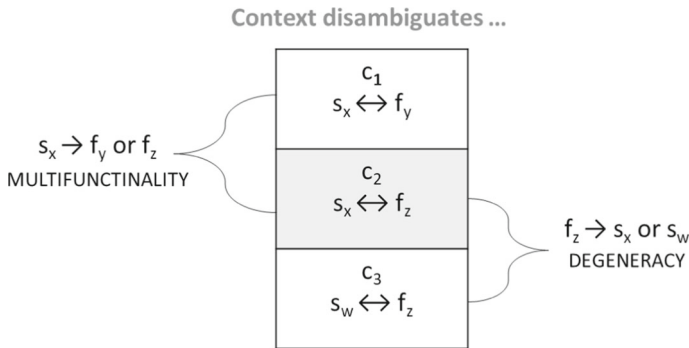
By shifting the focus from surface-functions to deep-functions, i.e. analyzing the inner workings of some neural structures irrespectively of the behavior they contribute to, the functional import of their structural properties is likely to become more transparent—eventually to the point that we may bridge the mind-brain gap and predict functions out of structures, as was hoped by Price and Friston (2005; cfr. Rathkopf 2013). However, a perfect knowledge of deep-functions will not suffice for cognitive neuroscience.

To see why, recall Price and Friston’s functional abstraction upon the left posterior lateral fusiform: the three surface-functions whose involvement they report (perception of words, of animal parts, and tactile-visual interface) are subsumed under the deep-functional label “sensorimotor integration”. However, such an abstract functional description is hardly specific to a given structure: apart from those neural regions most proximal to sensory receptors or motor effectors, such a description applies to the whole cortex. And even if from the activity of left posterior lateral fusiform we were to infer that some sensorimotor integration is occurring, this will leave us clueless about what this mean at the level of behavior: is the subject reading words? Or is she at the zoo, focusing on some animal parts? Or perhaps she is engaged in some new tasks we did not expect? We cannot tell (for an extensive discussion, see Klein 2012; Burnston 2016a,b). While brain anatomy and physiology may also be interesting *regardless of their contribution to the behavior* (Haueis 2014), cognitive neuroscience seeks to establish how they contribute to behavior. Deep-functions alone are not enough to predict behavior: surface-functions are also in order. We need them for planning and interpreting experiments, which are often made of *tasks*; and, *a fortiori*, to understand the import of cognitive neuroscience in everyday life.<sup>9</sup>

So, rather than aiming for deep-function labels that accounts for the activation of some structure across different contexts, several scholars argued for context-sensitive function-structure mappings.

Contextualism has been first introduced in philosophical debates by Klein (2012), though its rationale has been defined and defended at length by Burnston (2016b). Roughly speaking, contextualism boils down to a simple recipe: whenever one variable, i.e. a given function or structure, is not enough to predict the other, further variable(s) can be added, i.e. context(s) (Fig. 3).

<sup>9</sup> This does not imply that the ontology of surface-functions needs be constrained by folk ontology: they have different scopes (Dewhurst 2020). More likely, translations should be established between scientific and folk ontologies (Francken and Slors 2017).



**Fig. 3** Contextualism allows us to disambiguate many-to-many (surface-)function-structure mappings by reducing them to one-to-one mappings indexed to a specific context

According to Burnston (2016a), computational absolutism (i.e., the quest for establishing a single deep-function for each brain area) owes its charm to ambitious promises of generalizability and projectability. Supported by a thorough discussion of the context-dependent sensitivity of the MT area (which responds either to motion, depth, or both, according to the availability of the stimuli), Burnston (2016a,b) argues that such promises are empty. While absolutists regard the ascription of an open-ended conjunction of possible (surface-)functions to an area as a bug, Burnston suggests embracing it as a *feature* of the mapping endeavor. In the framework that he sketches, known functional ascriptions are still generalizable in the sense that they guide the generation of hypotheses about unknown functions. But “the level of generality of each of the conjuncts in the theory is precisely a matter for empirical investigation—discovering the limits of generalizability for a certain conjunct is just as important an advance as discovering that it holds in many instances” (Burnston 2016b: pp. 546–547).

A few notions of context have been proposed in the philosophical literature. To assess their usefulness, at least in principle, one may adapt a revised version of the criterion put forward by Price and Friston (2005: p. 272):

that function should be predicted from anatomical activation and conversely that anatomical activation should predict [surface-]function [given a certain context]. This can be embodied operationally in a cost function that reflects the prediction error in going from [surface-]function [plus context] to structure and back again.

Building on the works of McIntosh (2004), Klein introduced the notion of *neural context*, defined as “the overall network in which a region is participating” (2012: p. 957). As a matter of fact, an area can be activated while performing various sorts of tasks, to the point that the ongoing task is not predictable upon the observation of the area’s activity in isolation. Yet, the relevant area is accompanied by different sets of areas in different cases. Actually, even the very same set of areas may implement different surface-functions, due to changes in connection strength between regions, or due to their fine-grained temporal sequence of activation, which is unlikely to be revealed by fMRI scans alone (Pessoa 2014). In a similar spirit, Viola and Zanin (2017) proposed that some physiological features of neural structures’ activity, such as their oscillatory frequency, may help discriminate their current (surface-, but pos-

sibly even deep-)function. This idea is further developed by Burnston (2019: p. 12), who painstakingly describes how “physiological activity within a given brain unit—cell, population of cells, or brain area—is not dependent solely on its inputs, but is modulated by background variables” such as the Local Field Potential.

Khalidi (2017) argued that function-structure relationships should also consider non-neural variables concerning the *environmental-etiological context*. His strategy would allow us to preserve some folk psychological distinctions even when they do not map neatly onto dedicated neural substrates. For instance, given the striking similarity in terms of neural activity between remembering the past and imagining it (or the future), a conflation between memory and imagination has been proposed. However, doing so would prevent us from distinguishing correct memories and incorrect ones (as you cannot *misimagine* something). To avoid this issue, Khalidi claims that knowledge of distal etiological context must come into play when individuating (surface-)functions. Namely, when the recalled trace and some actual mental representation held in the subject’s past are congruent, we speak of memory. Otherwise we speak of imagination.

### 3.4 Population as a context for function-structure mapping

In the light of the discussion of the previous sub-section, I can now state my theoretical commitments with the regard to the debate on function-structure mapping. In general, I endorse contextualism: I do not think that cognitive neuroscience should (nor could) do without surface-functions. Quite on the contrary, I think that cognitive neuroscientists should mainly care about the surface-functional properties of neural structures. However, I also think that the notion of deep-function has some merits, with the caveat that it should not be pursued as a “ultimate functional description” of a given neural structure, but more deflationary as a heuristic device to unravel new surface-functions. Conversely, discovering new surface-functions may lead to revision in deep-functional ascriptions (see §5.1).

That being said, while all the notions of context presented in the previous sub-section may prove useful to bridge structures and surface-function despite multifunctionality and degeneracy, none of them address the very challenge I am addressing in this paper: namely, that different neural structures can play different surface-functions (e.g. processing visual information or subserving echo-location) in some individuals, but not in others.

Previously, individual differences in function-structure mapping were invoked to justify skepticism about the very idea of cognitive neuroscience. For instance, inspired by Karl Lashley’s seminal work on cortical plasticity, Jerry Fodor remarkably claimed that “is entirely possible that the nervous system of higher organisms characteristically achieves a given psychological end by a wide variety of neurological means. If so, then the attempt to pair neurological structures with psychological functions is foredoomed” (1974: p. 105).

However, *pace* Fodor, rejecting the (quasi-)universal ambitions of the Platonic Brain Model does not foredoom neuroscience to turn into a mere collection of idiosyncratic token-correspondences about the brain and mind of single individuals, blocking any

possible generalizations. Patterned function-structure mappings can provide a middle-ground between the Platonic Brain Model and Fodor's *aporia*. Similarly to Kim (1992), who proposed to split "jades" into two categories "jadeite" and "nephrite", I suggest sorting individuals into distinct neuroscientifically relevant populations based on some factors that drive them to share a same functional organization. To be fair with Carrazza, he himself envisaged the possibility "to scientifically investigate some domain of natural phenomena where we could not make the assumption of universality (e.g., if there were several kinds of human minds)" (1986: 49, fn. 3). However, unlike jadeite and nephrite, human brains cannot be sorted out into natural kinds based on a single, underlying essence. Quite on the contrary, membership to some population is due to the joint work of several factors steering functional organization. As such, the same individual may belong to several populations at once, and its membership can be altered in time (see below in §5.2).

Before proposing some criteria for sorting subjects into populations, in order to stress *why* this proposal is in order, in the next section (§4) I discuss the case of the so-called "Fusiform Face Area".

## 4 The fusiform face (or flexible?) area

### 4.1 A dedicated module for face perception

Going somehow against the tide of contemporary emphasis on distributed processing (e.g. Pessoa 2014), Nancy Kanwisher (e.g. 2010, 2017) strenuously defends the existence of cortices functionally specialized for some category of stimuli. Among them, the most notable is the so-called Fusiform Face Area: a portion of the right fusiform gyrus (in temporal lobe) specialized in processing *faces* so to determine their identities.

The hypothesis that a dedicated cognitive system exists for recognizing faces (but not other objects) circulated in psychology way before neuroscience kicked in. Early findings showed that presenting a stimulus upside-down impairs the recognition of faces far more than other objects (Yin 1969). At the same time, the identity of two juxtaposed half-faces is harder to recognize when the two halves are matched than when they are misaligned. Interestingly, the recognition of the two halves is better when these composite faces are presented upside-down (Young, Hellawell and Hey 1987). Bodamer (1947) reported a specific syndrome in which face perception is selectively impaired, dubbing it *prosopagnosia*. In the following decades, double dissociations have been reported between (rare) patients who suffer only from impairment in face recognition (e.g. Rossion et al. 2003; Riddoch et al. 2008) and patients with preserved face recognition, who are nonetheless impaired with other objects (Rumiati and Humphreys 1997). However, prosopagnosic subjects present the same recognition accuracy when faces are presented upright or upside-down (Busigny and Rossion 2010). These and other findings support the idea that, in most cases, face recognition is based on holistic processing, i.e., on the relative positions of the elements of a face rather than on the detailed analysis of specific facial details.

Enter neuroscience. Inspired by evidence of a specific module for face processing, some laboratories undertook the quest for *localizing* it. The most notable result was achieved by Kanwisher and colleagues (1997). With a fMRI study, they found that a portion of the occipitotemporal cortex, in the right fusiform gyrus, was consistently more responsive to stimuli depicting faces than other objects—including scrambled faces.<sup>10</sup> Kanwisher’s team thus claimed that that area, which they dubbed FFA, is selective for processing faces. More specifically, for holistic processing of faces, as suggested by lesser activation with scrambled face.

Their hypothesis was supported by other evidence. For instance, FFA’s activation in subjects suffering from prosopagnosia differ from that of controls (e.g. Hadjikhani and de Gelder 2002). And acquired prosopagnosia often depend upon lesions in the FFA. Yet, in some patients FFA is spared, and the lesion is instead located in the left frontal cortex (Cohen et al. 2019). Does this speak against the one-to-one mapping between face processing and FFA?

#### 4.2 Is face processing distributed?

Contrary to Kanwisher, some scholars held that face processing should not be narrowly localized, as it rather depends on a distributed processing of several neural structures spread across the ventral temporal cortex. In what became the proof of concept for MultiVariate Pattern Analysis, Haxby and colleagues (2001) exposed subjects to visual stimuli from six categories (including faces) during an fMRI scan. Based on a wide, non-localized activation pattern corresponding to each category from half of their dataset, they predicted the category of the stimulus of the unlabeled activations for the other half of the dataset with astonishing accuracy (>90%). Crucially, when the prediction was repeated leaving out the voxels that were maximally responsive to each category, accuracy dropped only very slightly. The following moral was drawn: the representation of object categories (including faces) depends on a widespread coding and not on the activity of a small area.

While Kanwisher grants that face recognition, like any cognitive task, may depend upon the joint activation of several areas, she also vindicates FFA’s centrality by employing a “correlation is not causation” argument. In fact, since BOLD fMRI only provides indirect measures of brain activity, the signal employed by researchers (or by classifiers) to infer mental states from neural activations is not necessarily the same signal playing a causal role in the neural implementation of that mental state (Ritchie et al. 2019). Her argument is further supported by the finding that, through direct cortical stimulation of FFA, a neurosurgical patient reported seeing illusory faces on top of ordinary objects (Schalk et al. 2017).

In the next subsection, I examine some evidence showing that FFA has multiple functions in different subjects. Beside highlighting the need for a population-based

<sup>10</sup> A further aspect that makes FFA an interesting case study for individual differences pertains its localization. In fact, despite being consistently involved in face processing across subjects, its precision location within the fusiform gyrus cannot be determined by purely anatomical means. To circumvent this misalignment, Kanwisher advocates the use of *functional localizers*, i.e. determining the localization of a region of interest by checking which portion of cortex is selectively responsive to some stimuli (as opposed to others).

contextualism (§5.1), this discussion will also suggest some criteria for sorting subjects into populations (§5.2).

### 4.3 A Fusiform “Flexible” Area?

While not denying that FFA is crucial for face recognition, the advocates of the so-called *expertise hypothesis* argue that this functional description is but the tip—however big—of an iceberg. In my lexicon, the hypothesis claims that FFA has a deep-function that is *holistic processing of familiar stimuli*. This deep-function is typically put at the service of face recognition because faces are a kind of stimulus with which *most* human beings are familiar. Still, the same deep-function could be recruited for other surface-functions. Hence, they propose to preserve the acronym “FFA”, but to reinterpret it as “Fusiform *Flexible* Area” (Tarr and Gauthier 2000). The higher levels of accuracy registered in recognizing faces from one’s own ethnicity and age may lend further support to this hypothesis (Gross 2009), and so would the recent claim that, after all, most people are good at recognizing only some faces – namely, those they are familiar with (Young and Burton 2017).

During the last two decades, different research groups seeking confirmation for the expertise hypothesis designed an experiment to demonstrate the involvement of FFA for other stimuli domain. Gauthier and colleagues (1999) trained some subjects to recognize some artificial objects called *greebles*. fMRI scans of both trained and untrained participants during a greebles recognizing task revealed greater FFA activity in the trained participants. Interestingly, this engagement vanished when greebles were presented upside down, similarly to what happens with faces. While *greebles* proved to be scarcely significant due to their similarity to faces, other categories of stimuli have also been reported to undergo a distinctive (holistic) processing style in experts of some domain, yielding to more accurate and faster recognition, and to activate FFA. These domains include cars and birds (Gauthier et al. 2000; Xu 2005); regular chess games, but not isolated chess pieces (Bilalić et al. 2011); and radiological images (Bilalić et al. 2014).

However suggestive, findings of the involvement of FFA in new surface-functions does not establish that its putative FS for faces was just a matter of expertise. Further support for this claim comes from Arcaro et al. (2017). In their study, some rhesus monkeys were raised in a face-deprived environment since birth. After 2 months, these monkeys did not exhibit the same face-centric gaze fixation patterns, nor the same cortical activation patterns, of controls.

Yet, still other studies suggest that the special status of face, as well as its privileged relation with FFA, may be at least partially independent from visual experience, and instead depend upon the anatomy of FFA and top-down connections with cortices implied in social cognition (Powell, Kosakowski and Saxe 2018; Kamps et al. 2020).<sup>11</sup>

<sup>11</sup> For instance, Buiatti and colleagues (2019) employed specific EEG adapted for testing 1- to 4-days old human newborns. They found increased activity in newborns’ temporooccipital cortex (their “proto-FFA”) for upright face-like stimuli when compared to inverted or scrambled face-like stimuli. Moreover, van den Hurk, Van Baelen and de Beeck (2017) report that either visual and auditory stimuli of “facial actions” (e.g. kissing or whistling) activate similar areas in the ventral-temporal cortex (where the FFA stands) of sighted individuals, and that auditory face stimuli result in bilateral fusiform activations in congenitally blind



So, is the selectivity for face a surface-function that simply exploits FFA's deep-functional predisposition for holistic processing? Or perhaps faces are so tightly tied to that spot of cortex that they themselves influence its deep-functional workings, making it suitable for being reused in recognizing other stimuli such as chess games? As the debate is ongoing, it is prudent to stay agnostic on this question. In either case, it is undisputed that FFA is sometimes reused for processing objects other than faces. But this reuse is conditional on learning. In other words: the FFA *becomes* a surface-multifunctional area in some subjects with an expertise in certain domains, while in others it seemingly remains surface-monofunctionally attuned to face. In the next section, I will leverage on this observation to argue for a shift from a Platonic function-structure mapping framework toward a population-bounded contextualism.

## 5 Beyond the Platonic Brain Model

### 5.1 The limits of the Platonic Brain Model

Recall: one aim of a Platonic neuroscience is that of finding “a normal (default) function–structure mapping” (Henson 2005: p. 225) of the utmost generalizability. Within that framework, it seems safe to claim that the FFA is functionally specialized in processing faces. The fact that the FFA-face processing mapping could fail to obtain in congenital prosopagnosic subject is not a nuisance for the Platonic Brain Model. After all, the model is meant to generalize to *healthy* adults, and prosopagnosic subjects hardly count as such.

However, the Platonic Brain Model is *committed to overlook* the latent surface-multifunctionality of FFA, as it only shows up in certain individuals, i.e. expert perceivers of some category of objects like cars or radiological images.

A defender of the Platonic Brain Model may well respond that these are just idle details about surface-functions, and rather settles for specifying the *deep*-function of each neural structure. Thus, since neuroscientists indeed came up with a relatively solid deep-functional characterization of FFA as an area involved in *holistic shape processing*, why bother specifying to which domains this processing applies?

Now, while legitimate, such a reply comes with a non-negligible implication: namely, that the Platonic Brain Model becomes rather uninformative of behavior. As discussed in §3.3, deep-functions underdetermine behavior. Recall that according to Price and Friston (2005: p. 272), the hallmark of a good ontology is that “function should be predicted from anatomical activation and conversely that anatomical activation should predict function”. If we take “function” to include surface-function, we can observe that, by knowing that someone's FFA is getting relevantly activated, you cannot infer whether she is just looking at some face or she is e.g. a veteran radiologist trying to diagnose some lung infection. However, we have seen (in §3.3) that contextual information may ameliorate the lack of robustness of reverse inference from

---

Footnote 11 continued

individuals. In a recent preprint (Ratan Murty and colleagues *manuscript*), Kanwisher's team describes an experiment where they found that in congenitally blind subjects, a same area in the fusiform gyrus responds to both auditory facial stimuli (inspired by the van den Hurk and colleagues' study), and haptic stimuli (3D-printed faces).

structures to surface-functions. In this case, the relevant contextual information is the expertise of the subject in some specific visual domains: by knowing that she is not an expert radiologist, for instance, we may exclude that the activation of FFA in her case underlies the observation of radiological images.

While knowing about a structure's deep-functions is insufficient to predict its contribution to behavior, deep-functional ascriptions can play an important heuristic role. Indeed, when investigating FFA's involvement in expert perception beyond faces, they did expect a same deep-functional style of processing, that is, configural processing. To verify this, they checked for inversion effects in the perception of greebles (Gauthier et al. 1999) or radiological images (Bilalić et al. 2014). Based on such expectations about the deep-function of FFA, future investigations about its possible redeployment in recognizing objects from other domains will likely privilege classes of objects which can be discriminated by their overall shape rather than classes of objects that can be discriminated by paying attention on some details.

Conversely, knowledge of possible surface-functions may help to understand the deep-functions an area may play. In localizing a portion of cortex selective for Pokémons across various former players, Gomez and colleagues (2019) were not merely checking an oddity. Rather, they were wondering what deep-functional semantic property best accounted for the surface-functional stimulus specificity that portion of cortex could acquire: animacy, rectilinearity, real-world size, or eccentricity of the stimuli (i.e. whether it is typically perceived in foveal or in peripheral vision). By comparing subjects' response to Pokémon characters with other classes of objects like faces, animals, body parts, and corridors, they argued in favor of eccentricity (but see de Beeck et al. 2019: p. 794).

Other than across-subject (surface-)multifunctionality, the Platonic Brain Model is poorly equipped to account for across-subject degeneracy. Some cognitive tasks, especially those involving complex higher-cognitive faculties, admit multiple solutions – that is, subjects can solve them by leveraging on alternative sets of deep-functions, resulting in the involvement of different neural circuits. Miller and colleagues (2002) report wide across-subject differences in patterns of fMRI activation of six subjects performing a retrieval task. However, they also noted a good across-time stability within the same individual—suggesting these idiosyncrasies should not be dismissed as noise. Since these idiosyncratic activations might be overlooked by pooling the activations of different subjects, they warn that “group analysis alone, particularly for higher order cognitions like episodic retrieval, may be incomplete and, in some cases, misleading” (Miller et al. 2002: 1211).

Accepting that individuals accomplish tasks in different ways does not imply that no generalizability is possible. On the contrary, in a later study Miller and colleagues (2012) try to interpret this variety of activations as differences in cognitive style (i.e. whether people are mainly *visualizers* or *verbalizers*) and encoding strategies, measured with some questionnaires. They found that both factors accounted for a significant part of the variance in the similarity of brain activity between individuals. In a similar vein, Noppeney and colleagues (2006) had 17 subjects performing a verbal decision task about some auditory stimulus (spoken word or sound), primed by either congruent or incongruent visual stimuli (written word or picture). Besides some regions commonly activated across all subjects, the experimenters were able

to distinguish two subgroups, based on the set of further regions activated beside the common ones. They speculate on the deep-functional meaning of different activations, although unfortunately, unlike Miller and colleagues (2012), they did not verify them with independent measures such as questionnaires.

## 5.2 Population of subjects as a context for function-structure mappings

According to Burnston, within a contextualist approach toward function-structure mapping, “discovering the limits of generalizability for a certain conjunct is just as important an advance as discovering that it holds in many instances” (Burnston 2016b: p. 547). Thus far, the allure of the Platonic Brain Model may have misled neuroscientists. They could either have overlooked some aspects of some structure’s deep-function or failed to appreciate some of their unusual surface-function. For instance, the comparisons between sighted and blind subjects allowed to reveal supramodal deep-functions in areas traditionally (mis)conceived as merely visual, as well as to reveal further surface-functions that they may acquire due to plastic reorganization (Pascual-Leone and Hamilton 2001; Ricciardi and Pietrini 2011; Cecchetti et al. 2016; Bedny 2017). On the other hand, Simons, Shoda and Lindsay (2017) suggest that to mitigate the dreaded replication crisis in psychological science, the section “Constraints on generality” may be added in research papers. Here researchers should answer questions concerning the generalizability they expect for their findings, including: “to which populations do you expect those findings to apply?”

To answer this question, and more in general to develop a proper contextualist function-structure mapping that still admits some generalizability, criteria shall be found for sorting subjects in different *populations*. The carvings of such populations are not self-evident. Rather, “the definition of a reference population is theory-laden” (Caramazza 1986: p. 51). So, for instance, blind subjects may be a population if we are studying vision, but other criteria will possibly drive population-clustering when we turn to study for instance smell. Moreover, Hochstein (2016) reminds us that, like the rest of science, the sciences of the mind have many purposes, and that the best categorization of the mental depends on what purpose one is pursuing. The same point holds for population-based contextualism about structure function mappings. For instance, neurosurgeons are happy to sacrifice generalizability for the sake of accuracy, which is why before the surgical intervention they often directly assess causal role of some specific spots of cortex is often assessed by means of direct electrical stimulation (Duffau 2017). To them, the ideal population is composed by only one individual (at a given time).

For researchers who want to retain some generality, developing an ontology of neuroscientifically relevant populations will require a significant amount of conceptual and empirical effort. Factors predictive of a certain functional organization of the brain must be predicted, and their boundaries and mutual interactions scrutinized, in the light of certain aims and questions. Subjects will simultaneously belong to multiple populations, and each of these memberships is likely of interest only *a few* function-structure mappings. To complicate things, we have reason to expect that these memberships will interact non-linearly, as surface-functions often compete for

neural real estate: developmental studies show that the FFA's preference for face is lateralized in the right hemisphere *after that*, and possibly *because of*, the ipsilateral region establish its preference for words (Dundas, Plaut and Behrmann 2013; Dehaene-Lambertz, Monzalvo and Dehaene 2018).

That being said, I can sketch with broad strokes the contour of at least one kind of population, for which I propose the following working definition:

[Expertise-based Population] Structure S playing the deep-function F is recruited for some processing related to the surface-functional domain D in F-way if and only if a subject is an expert in D.

This definition captures the empirical insight that, in some cases, prolonged experience with some kinds of stimuli or cognitive activities does not only yield quantitative changes in brain anatomy (as in Maguire et al. 2000). Rather, when the relevant amount of expertise is achieved, expert performance may differ not only quantitatively, e.g. because it engages neural structure more or less intensively (which can be easily accounted for by the Platonic Brain Model), but also qualitatively. When this is the case, a same task may recruit different deep-functions and hence different neural structures (Roepstorff, Niewöhner, and Beck 2010; Guida et al. 2012; Bilalić 2017). The transition between non-experts and expert may occur smoothly, and sharp boundaries may thus be hard to establish (Buckner 2016). But just like vague boundaries between brain areas do not discourage scientists from speaking about brain areas (Haueis 2012), neither should vague boundaries between experts and non-experts in a certain domain discourage them to speak about expert populations.

Being an expert in the respective domains makes the difference for the FFA's involvement in holistic perception of objects like greebles, chess games, and radiological images. Moreover, if the expertise hypothesis is correct, expertise is also at the root of FFA's surface-functional specialization in face recognition. The same logic applies equally well for other cases such as echolocators and literate subjects: it is only thanks to the acquaintance with written letters that the left fusiform gyrus of literate subjects acquires the surface-function of reading. Moreover, a sufficiently liberal notion of 'expertise' can also account for the difference between the functional vocation of the (Platonically speaking) "visual" areas in sighted as opposed to non-sighted subjects the former being conceived as 'experts' in the (broad) domain of vision.

Other kinds of populations may be individuated based on demographic (e.g. socio-economic status, ethnicity or gender) or clinical factors (e.g. suffering from a psychopathological disease).

Admittedly, "demographical factors" include disparate things, from Socioeconomic Status (SES) to sex. A review of the neural correlates of SES (Farah 2017) shows that, on the one hand, patterned differences can be found in the neural activity underlying certain tasks based on the subjects' SES, which are not only quantitative but also qualitative. However, on the other hand, Farah highlights that, despite being a proxy measure for *predicting* these differences, SES per se is not necessarily *producing* them. On the contrary, SES is a composite index that collapses many factors like stress and literacy, whose causal relationships with brain organization can follow multiple paths. Indirect measures such as SES can be at best *proxies* toward the causally relevant factors driving brain organization. In other words, it is likely that, at least in some cases,

function-structure mappings with a nice correlation with some demographic factor turn out to correlate even better with some underlying factors that can be accounted in term of expertise.<sup>12</sup>

One may think that this line of cautionary reasoning applies to certain factors, like SES, but not to others, like sex. After all, unlike SES, an eminently social phenomenon, sex is rooted in biology. Thus, we can perhaps hope to distinguish “the female brain” and “the male brain”. This line of reasoning, however, fails to appreciate that belonging to a particular sex (and gender) has several consequences that span throughout the life history of an individual, and that are highly dependent on cultural factors. In simple terms, in virtually every society of human history, being born female or male makes it much more likely to routinely perform some activities rather than others. Thus, even if differences that are driven on genetic or hormonal bases exist, they are so tightly intertwined with culturally based learning-driven factors that distinguishing them is not straightforward (Joran-Young and Rumiati 2012; Joel et al. 2015; but cf. Del Giudice 2019).

Finally, seen in the light of population contextualism, many clinical uses of functional neuroimaging are seen as inferences from structure *plus* function *to* a certain clinical population. Similarly to the Platonic Brain Model strategy to zero-out the contextual variable of population by checking only “normal subjects”, many clinical studies aim at zeroing-out the possible confounds of an externally-driven function by using a resting state design (e.g. van den Heuvel et al. 2013). Defining and diagnosing psychopathological and neurological disorders presents several challenges (e.g., recall Caramazza’s concern about etiology- versus symptom-based grouping), and relevant ethical implications, that cannot be addressed in this paper. However, abandoning the Platonic Brain Model may set the stage for a more fruitful discussion. To mention but one example, think of the recent claims that, at least in some cases, Asperger Syndrome is not a syndrome at all, it is just *neurodiversity*—not differently from other cases of biological dimorphism (for a discussion, see Jaarsma and Welin 2012). Whatever one think about this claim, approaching it from a post-platonic framework will probably help to take this claim with the seriousness that it deserves, because such a framework does not prescribe a monolithic normativity about the brain.

## 6 Conclusion: sketching the agenda for a post-Platonic neuroscience

In this paper, I have urged to take into account inter-individual variability in cognitive skills and brain organization as a possible source of *qualitative* diversity in function-structure mappings. I have described a popular strategy for addressing this variability: excluding abnormal subjects, and take the (averaged) results of normal subjects as informative about “the Platonic Brain Model”, from which inferences can be drawn about any normal brain. Building on the discussion surrounding the functional role(s) of the FFA, I showed how failure to relativize functional ascriptions to specific populations may result in overlooking surface-multifunctionality (and degeneracy). I thus

<sup>12</sup> Despite being useful in flagging the problem, the label “WEIRD” (Henrich et al. 2010) is too large a category to drive population grouping within the contextualist framework I am advocating (see also Roepstorff et al. 2010).

invoked the necessity to develop an ontology of neuroscientifically relevant notions of populations. Finally, I suggested one promising (albeit not exhaustive) criterion for sorting population, namely expertise-based population.

My plea goes on a par with those coming from an increasing number of researchers who invite to treat inter-subject diversity as signal, not as noise (e.g. Falk et al. 2013; Thiebault de Schotten and Shallice 2017; Seghier and Price 2018; Clark Barrett 2020). The Platonic Brain Model has served us quite well in many cases, providing a baseline for many claims about brain organization. But as cognitive neuroscience is evolving into a mature science, replacing it with the more context-sensitive framework of a post-platonic neuroscience would repay in terms of accuracy in predicting behavior. Moreover, on the grounds of philosophical anthropology, it would allow us to break free from the tethers of a concept of human nature largely predetermined by evolution (see also Rathkopf forthcoming), and to widen the scope of neuroscience from *traditional* behaviors toward *peculiar* ones. The time is ripe to look beyond what the *ideal* brain typically does, toward what *real* brains can do in their beautiful diversity.

**Acknowledgements** Many core ideas for this paper were developed during my time as a Visiting Scholar at Macquarie University, Sydney, in Fall 2016, and especially thanks to conversations with Colin Klein, Greg Downey, Alex J. Gillett, and the *Culture and Cognition* reading group. I then developed these ideas extensively and presented them to several audiences in Genoa, Rijeka, Florence, Moscow, Novara, Pavia, Lucca, and Turin, to whom I am grateful for the feedback. I am also grateful to Dan Burnston, Fabrizio Calzavarini, Sabrina Coninx, Joe Dewhurst, Sarah Genon, Alex J. Gillett, Michele Giuli, Davide Serpico, Valentina Petrolini and Zina Ward for their precious comments on earlier versions of this manuscript. Finally, I owe much to the clever comments of two anonymous reviewers. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No 819649 - FACETS).

**Funding** Open access funding provided by Università degli Studi di Torino within the CRUI-CARE Agreement.

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

## References

- Anderson, M. (2010). Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33(4), 245–266.
- Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. Cambridge: MIT Press.
- Anderson, M. L., Kinnison, J., & Pessoa, L. (2013). Describing functional diversity of brain regions and brain networks. *Neuroimage*, 73, 50–58.
- Anderson, M. L., & Pessoa, L. (2011). Quantifying the diversity of neural activations in individual brain regions. In L. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 2421 – 2426). Austin, TX: Cognitive Science Society.
- Arcaro, M. J., Schade, P. F., Vincent, J. L., Ponce, C. R., & Livingstone, M. S. (2017). Seeing faces is necessary for face-domain formation. *Nature Neuroscience*, 20(10), 1404.

- Banich, M. T., & Compton, R. J. (2018). *Cognitive neuroscience*. Cambridge: Cambridge University Press.
- Barack, D. L. (2019). Cognitive recycling. *The British Journal for the Philosophy of Science*, *70*(1), 239–268.
- Bartolomeo, P., Malkinson, T. S., & De Vito, S. (2017). Botallo's error, or the quandaries of the universality assumption. *Cortex*, *86*, 176–185.
- Bechtel, W., & Richardson, R. C. (2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge: MIT press.
- Bedny, M. (2017). Evidence from blindness for a cognitively pluripotent cortex. *Trends in Cognitive Sciences*, *21*(9), 637–648.
- Bergeron, V. (2007). Anatomical and functional modularity in cognitive science: Shifting the focus. *Philosophical Psychology*, *20*(2), 175–195.
- Bilalić, M. (2017). *The neuroscience of expertise*. Cambridge: Cambridge University Press.
- Bilalić, M., Grottenhaler, T., Nägele, T., & Lindig, T. (2014). The faces in radiological images: fusiform face area supports radiological expertise. *Cerebral Cortex*, *26*(3), 1004–1014.
- Bilalić, M., Langner, R., Ulrich, R., & Grodd, W. (2011). Many faces of expertise: fusiform face area in chess experts and novices. *Journal of Neuroscience*, *31*(28), 10206–10214.
- Bodamer, J. (1947). Die Prosop-Agnosie. *Archiv für Psychiatrie und Nervenkrankheiten*, *179*, 6–53.
- Buckner, C. (2016). Transitional gradation in the mind: rethinking psychological kindhood. *The British Journal for the Philosophy of Science*, *67*(4), 1091–1115.
- Buiatti, M., Di Giorgio, E., Piazza, M., Polloni, C., Menna, G., Taddei, F., et al. (2019). Cortical route for face like pattern processing in human newborns. *Proceedings of the National Academy of Sciences*, *116*(10), 4625–4630.
- Burnston, D. C. (2016a). Computational neuroscience and localized neural function. *Synthese*, *193*(12), 3741–3762.
- Burnston, D. C. (2016b). A contextualist approach to functional localization in the brain. *Biology and Philosophy*, *31*(4), 527–550.
- Burnston, D. C. (2019). Getting over atomism: Functional decomposition in complex neural systems. *British Journal for Philosophy of Science*. <https://doi.org/10.1093/bjps/axz039>.
- Busigny, T., & Rossion, B. (2010). Acquired prosopagnosia abolishes the face inversion effect. *Cortex*, *46*(8), 965–981.
- Caramazza, A. (1986). On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. *Brain and Cognition*, *5*(1), 41–66.
- Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*(1), 51–62.
- Cecchetti, L., Kupers, R., Püto, M., Pietrini, P., & Ricciardi, E. (2016). Are supramodality and cross-modal plasticity the yin and yang of brain development? From blindness to rehabilitation. *Frontiers in Systems Neuroscience*, *10*, 89.
- Clark Barrett, H. (2020). Towards a cognitive science of the human: cross-cultural approaches and their urgency. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2020.05.007>.
- Cohen, A. L., Soussand, L., Corrow, S. L., Martinaud, O., Barton, J. J., & Fox, M. D. (2019). Looking beyond the face area: Lesion network mapping of prosopagnosia. *Brain*, *142*(12), 3975–3990.
- Danziger, K. (1990/1994). *Constructing the subject: Historical origins of psychological research*. Cambridge University Press, Cambridge.
- de Beeck, H. P. O., Pillot, I., & Ritchie, J. B. (2019). Factors determining where category-selective areas emerge in visual cortex. *Trends in Cognitive Sciences*, *23*(9), 784–797.
- De Brigard, F. (2017). Cognitive systems and the changing brain. *Philosophical Explorations*, *20*(2), 224–241.
- de Schotten, M. T., & Shallice, T. (2017). Identical, similar or different? Is a single brain model sufficient? *Cortex*, *100*(86), 172–175.
- Dehaene, S., Pegado, F., Braga, L. W., Ventura, P., Nunes Filho, G., Jobert, A., et al. (2010). How learning to read changes the cortical networks for vision and language. *Science*, *330*(6009), 1359–1364.
- Dehaene-Lambertz, G., Monzalvo, K., & Dehaene, S. (2018). The emergence of the visual word form: Longitudinal evolution of category-specific ventral visual areas during reading acquisition. *PLoS Biology*, *16*(3), e2004103.
- Del Giudice, M. (2019). *Measuring sex differences and similarities. Gender and sexuality development: Contemporary theory and research*. New York, NY: Springer.

- Dewhurst, J. (2020). Folk psychological and neurocognitive ontology. In F. Calzavarini, M. Viola (Eds.), *Neural Mechanisms: New Challenges in the Philosophy of Neuroscience*. Available at <http://philsci-archive.pitt.edu/17269>.
- Downey, G. (2016). Sensory enculturation and neuroanthropology: the case of human echolocation. In J. Y. Chiao, S. C. Li, R. Seligman, & R. Turner (Eds.), 2016 (pp. 41–55). The Oxford handbook of cultural neuroscience: Oxford University Press.
- Duffau, H. (2017). A two-level model of interindividual anatomo-functional variability of the brain and its implications for neurosurgery. *Cortex*, 86, 303–313.
- Dundas, E. M., Plaut, D. C., & Behrmann, M. (2013). The joint development of hemispheric lateralization for words and faces. *Journal of Experimental Psychology: General*, 142(2), 348.
- Falk, E. B., Hyde, L. W., Mitchell, C., Faul, J., Gonzalez, R., Heitzeg, M. M., et al. (2013). What is a representative brain? Neuroscience meets population science. *Proceedings of the National Academy of Sciences*, 110(44), 17615–17622.
- Farah, M. J. (2017). The neuroscience of socioeconomic status: Correlates, causes, and consequences. *Neuron*, 96(1), 56–71.
- Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision Research*, 35(14), 2089–2093.
- Fedorenko, E., & Blank, I. A. (2020). Broca's area is not a natural kind. *Trends in cognitive sciences*, 24(4), 270–284.
- Figdor, C. (2010). Neuroscience and the multiple realization of cognitive functions. *Philosophy of Science*, 77(3), 419–456.
- Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, 28(2), 97–115.
- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience*, 2(6), 568–573.
- Genon, S., Reid, A., Langner, R., Amunts, K., & Eickhoff, S. B. (2018). How to characterize the function of a brain region. *Trends in Cognitive Sciences*, 22(4), 350–364.
- Gomez, J., Barnett, M., & Grill-Spector, K. (2019). Extensive childhood experience with Pokémon suggests eccentricity drives organization of visual cortex. *Nature Human Behaviour*, 3(6), 611–624.
- Gross, T. F. (2009). Own-ethnicity bias in the recognition of Black, East Asian, Hispanic, and White faces. *Basic and Applied Social Psychology*, 31(2), 128–135.
- Guida, A., Gobet, F., Tardieu, H., & Nicolas, S. (2012). How chunks, long-term working memory and templates offer a cognitive explanation for neuroimaging data on expertise acquisition: a two-stage framework. *Brain and Cognition*, 79(3), 221–244.
- Hadjikhani, N., & De Gelder, B. (2002). Neural basis of prosopagnosia: an fMRI study. *Human Brain Mapping*, 16(3), 176–182.
- Haueis, P. (2012). The fuzzy brain. Vagueness and mapping connectivity of the human cerebral cortex. *Frontiers in Neuroanatomy*, 6, 37.
- Haueis, P. (2014). Meeting the brain on its own terms. *Frontiers in human neuroscience*, 8, 815.
- Haueis, P. (2018). Beyond cognitive myopia: a patchwork approach to the concept of neural function. *Synthese*, 195(12), 5373–5402.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293(5539), 2425–2430.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83.
- Henson, R. (2005). What can functional neuroimaging tell the experimental psychologist? *The Quarterly Journal of Experimental Psychology Section A*, 58(2), 193–233.
- Hochstein, E. (2016). Categorizing the mental. *The Philosophical Quarterly*, 66(265), 745–759.
- Jaarsma, P., & Welin, S. (2012). Autism as a natural human variation: Reflections on the claims of the neurodiversity movement. *Health Care Analysis*, 20(1), 20–30.
- Joel, D., Berman, Z., Tavor, I., Wexler, N., Gaber, O., Stein, Y., et al. (2015). Sex beyond the genitalia: The human brain mosaic. *Proceedings of the National Academy of Sciences*, 112(50), 15468–15473.
- Jordan-Young, R., & Rumiat, R. I. (2012). Hardwired for sexism? *Approaches to sex/gender in neuroscience. Neuroethics*, 5(3), 305–315.



- Kamps, F. S., Hendrix, C. L., Brennan, P. A., & Dilks, D. D. (2020). Connectivity at the origins of domain specificity in the cortical face and place networks. *Proceedings of the National Academy of Sciences*, *117*(11), 6163–6169.
- Kanwisher, N. (2010). Functional specificity in the human brain: a window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences*, *107*(25), 11163–11170.
- Kanwisher, N. (2017). The quest for the FFA and where it led. *Journal of Neuroscience*, *37*(5), 1056–1061.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*(11), 4302–4311.
- Khalidi, M. A. (2017). Crosscutting psycho-neural taxonomies: the case of episodic memory. *Philosophical Explorations*, *20*(2), 191–208.
- Kim, J. (1992). Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research*, *52*(1), 1–26.
- Klein, C. (2010). Redeployed functions versus spreading activation: A potential confound. *Behavioral and Brain Sciences*, *33*(4), 280–281.
- Klein, C. (2012). Cognitive ontology and region-versus network-oriented analyses. *Philosophy of Science*, *79*(5), 952–960.
- Kupers, R., Pappens, M., de Noordhout, A. M., Schoenen, J., Ptito, M., & Fumal, A. (2007). rTMS of the occipital cortex abolishes Braille reading and repetition priming in blind subjects. *Neurology*, *68*(9), 691–693.
- Lindquist, K. A., Wager, T. D., Kober, H., Bliss-Moreau, E., & Barrett, L. F. (2012). The brain basis of emotion: a meta-analytic review. *The Behavioral and Brain Sciences*, *35*(3), 121–143.
- Maguire, E. A., Gadian, D. G., Johnsrude, I. S., Good, C. D., Ashburner, J., Frackowiak, R. S., et al. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *Proceedings of the National Academy of Sciences*, *97*(8), 4398–4403.
- McCaffrey, J. B. (2015). The brain's heterogeneous functional landscape. *Philosophy of Science*, *82*(5), 1010–1022.
- McCauley, R. N., & Bechtel, W. (2001). Explanatory pluralism and heuristic identity theory. *Theory & Psychology*, *11*(6), 736–760.
- McIntosh, A. R. (2004). Contexts and catalysts: a resolution of the localization and integration of function in the brain. *Neuroinformatics*, *2*(2), 175–182.
- Miller, M. B., Van Horn, J. D., Wolford, G. L., Handy, T. C., Valsangkar-Smyth, M., Inati, S., et al. (2002). Extensive individual differences in brain activations associated with episodic retrieval are reliable over time. *Journal of Cognitive Neuroscience*, *14*(8), 1200–1214.
- Miller, M. B., Donovan, C. L., Bennett, C. M., Aminoff, E. M., & Mayer, R. E. (2012). Individual differences in cognitive style and strategy predict similarities in the patterns of brain activity between individuals. *Neuroimage*, *59*(1), 83–93.
- Mundale, J. (2002). Concepts of localization: Balkanization in the brain. *Brain and Mind*, *3*(3), 313–330.
- Nahm, M., Rousseau, D., & Greyson, B. (2017). discrepancy between cerebral structure and cognitive functioning: a review. *The Journal of Nervous and Mental Disease*, *205*(12), 967–972.
- Noppeney, U., Friston, K. J., & Price, C. J. (2004). Degenerate neuronal systems sustaining cognitive functions. *Journal of Anatomy*, *205*(6), 433–442.
- Noppeney, U., Price, C. J., Penny, W. D., & Friston, K. J. (2006). Two distinct neural mechanisms for category-selective responses. *Cerebral Cortex*, *16*(3), 437–445.
- Pascual-Leone, A., & Hamilton, R. (2001). The metamodal organization of the brain. In C. Casanova, M. Ptito (Eds.), *Vision: from Neurons to Cognition* (Vol. 134, pp. 427–445). Elsevier.
- Pessoa, L. (2014). Understanding brain networks and brain organization. *Physics of Life Reviews*, *11*(3), 400–435.
- Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in cognitive sciences*, *10*(2), 59–63.
- Polger, T. W., & Shapiro, L. A. (2016). *The multiple realization book*. Oxford: Oxford University Press.
- Powell, L. J., Kosakowski, H. L., & Saxe, R. (2018). Social origins of cortical face areas. *Trends in Cognitive Sciences*, *22*(9), 752–763.
- Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology*, *22*(3–4), 262–275.
- Ratan Murty, N.A., Teng, S., Beeler, D., Mynick, A., Oliva, A., & Kanwisher, N. (manuscript). Visual Experience is not Necessary for the Development of Face Selectivity in the Lateral Fusiform Gyrus. <https://doi.org/10.1101/2020.02.25.964890>.

- Rathkopf, C. A. (2013). Localization and intrinsic function. *Philosophy of Science*, 80(1), 1–21.
- Rathkopf, C. A. (forthcoming). Neural Reuse and the Nature of Evolutionary Constraints. In F. Calzavarini, M. Viola M (Eds.), *Neural Mechanisms: New Challenges in the Philosophy of Neuroscience*. Available at <http://philsci-archive.pitt.edu/17195/>.
- Ricciardi, E., & Pietrini, P. (2011). New light from the dark: What blindness can teach us about brain function. *Current Opinion in Neurology*, 24(4), 357–363.
- Riddoch, M. J., Johnston, R. A., Bracewell, R. M., Boutsen, L., & Humphreys, G. W. (2008). Are faces special? A case of pure prosopagnosia. *Cognitive Neuropsychology*, 25(1), 3–26.
- Ritchie, J. B., Kaplan, D. M., & Klein, C. (2019). Decoding the brain: Neural representation and the limits of multivariate pattern analysis in cognitive neuroscience. *The British Journal for the Philosophy of Science*, 70(2), 581–607.
- Roepstorff, A., Niewöhner, J., & Beck, S. (2010). Enculturing brains through patterned practices. *Neural Networks*, 23(8–9), 1051–1059.
- Rossion, B., Caldara, R., Seghier, M., Schuller, A. M., Lazeyras, F., & Mayer, E. (2003). A network of occipito-temporal face-sensitive areas besides the right middle fusiform gyrus is necessary for normal face processing. *Brain*, 126, 2381–2395.
- Rumiati, R. I., & Humphreys, G. W. (1997). Visual object agnosia without alexia or prosopagnosia. *Visual Cognition*, 4, 207–217.
- Schalk, G., Kapeller, C., Guger, C., Ogawa, H., Hiroshima, S., Lafer-Sousa, R., et al. (2017). Facephens and rainbows: Causal evidence for functional and anatomical specificity of face and color processing in the human brain. *Proceedings of the National Academy of Sciences*, 114(46), 12285–12290.
- Seghier, M. L., & Price, C. J. (2018). Interpreting and utilising intersubject variability in brain function. *Trends in cognitive sciences*, 22(6), 517–530.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.
- Shallice, T. (2015). Cognitive neuropsychology and its vicissitudes: The fate of Caramazza's axioms. *Cognitive Neuropsychology*, 32(7–8), 385–411.
- Simons, D. J., Shoda, Y., & Lindsay, D. S. (2017). Constraints on generality (COG): A proposed addition to all empirical papers. *Perspectives on Psychological Science*, 12(6), 1123–1128.
- Spunt, R. P., & Adolphs, R. (2017). A new look at domain specificity: insights from social neuroscience. *Nature Reviews Neuroscience*, 18(9), 559–567.
- Sterelny, K. (1990). *The representational theory of mind: An introduction*. Oxford: Basil Blackwell.
- Tarr, M. J., & Gauthier, I. (2000). FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience*, 3(8), 764–769.
- Tettamanti, M., & Weniger, D. (2006). Broca's area: a supramodal hierarchical processor? *Cortex*, 42(4), 491–494.
- Thaler, L., Arnott, S. R., & Goodale, M. A. (2011). Neural correlates of natural human echolocation in early and late blind echolocation experts. *PLoS ONE*, 6(5), e20162.
- Uddin, L. Q., Nomi, J. S., Hébert-Seropian, B., Ghaziri, J., & Boucher, O. (2017). Structure and function of the human insula. *Journal of clinical neurophysiology: official publication of the American Electroencephalographic Society*, 34(4), 300–306.
- van den Heuvel, M. P., Sporns, O., Collin, G., Scheewe, T., Mandl, R. C., Cahn, W., et al. (2013). Abnormal rich club organization and functional brain dynamics in schizophrenia. *JAMA Psychiatry*, 70(8), 783–792.
- van den Hurk, J., Van Baelen, M., & de Beeck, H. P. O. (2017). Development of visual category selectivity in ventral visual cortex does not require visual experience. *Proceedings of the National Academy of Sciences*, 114(22), E4501–E4510.
- Viola, M., & Zanin, E. (2017). The standard ontological framework of cognitive neuroscience: Some lessons from Broca's area. *Philosophical Psychology*, 30(7), 945–969.
- Ward, J. (2015). *The student's guide to cognitive neuroscience*. Hove: Psychology Press.
- Ward, Z. B. (2019). Registration pluralism and the cartographic approach to data aggregation across brains. *The British Journal for the Philosophy of Science*. <https://doi.org/10.1093/bjps/axz027>.
- Whitwell, J. L. (2009). Voxel-based morphometry: an automated technique for assessing structural changes in the brain. *Journal of Neuroscience*, 29(31), 9661–9664.
- Wilke, M., Schmithorst, V. J., & Holland, S. K. (2003). Normative pediatric brain data for spatial normalization and segmentation differs from standard adult data. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 50(4), 749–757.

- Xu, Y. (2005). Revisiting the role of the fusiform face area in visual expertise. *Cerebral Cortex*, *15*(8), 1234–1242.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*(1), 141–145.
- Young, A. W., & Burton, A. M. (2017). Recognizing faces. *Current Directions in Psychological Science*, *26*(3), 212–217.
- Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception*, *16*(6), 747–759.
- Zerilli, J. (2019). Neural reuse and the modularity of mind: where to next for modularity? *Biological Theory*, *14*(1), 1–20.

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