

THE EXPLORATORY DIMENSION OF FMRI EXPERIMENTS

A. NICOLÁS VENTURELLI

<https://orcid.org/0000-0003-0745-2643>
Instituto de Humanidades (UNC/ CONICET)
Department of Philosophy
Córdoba
Argentina
nventurelli@unc.edu.ar

Article info

CDD: 153.4

Received: 23.06.2020; Revised: 23.09.2020; Accepted: 08.10.2020

<https://doi.org/10.1590/0100-6045.2021.V44N1.AV>

Keywords

Exploratory Experimentation
Functional Magnetic Resonance Imaging
Cognitive Neuroscience
Experimental Practices
Functional Neuroimaging

Abstract: Driven by an appreciation of the field's early stage of development, I apply the concept of exploratory experimentation, originally put forward in the late 90s philosophy of biology, to current research in cognitive neuroscience. I concentrate on functional magnetic resonance imaging and how this wide-spread technique is used, from experimental design to data analysis. I claim that, although subject to certain significant modifications with respect to the concept's original rendering, the exploratory character of neuroimaging experiments can be appreciated considering their goals, centered on the stabilization of experimental systems for phenomenological description, and the relevance of their methodological facet. Although I do not claim that there is a specific kind of experiment that one can single out as definitely exploratory, exploration can be seen as a general trait imbuing fMRI-based experimentation.

1. Introduction

Although still mostly uncharted territory, the philosophical analysis of experimental practices in cognitive neuroscience has a deep theoretical relevance, mainly due to their known complexity and their rapidly changing dynamics. In particular, functional magnetic resonance imaging (fMRI) is presently the most widely used technique to study the healthy human brain in action. The dominant tendency, as Wright (2017) has recently noted, has been one of skeptical, critical evaluation on fMRI's potential to study how cognitive function relates to the brain. While this critical stance comprises firmly skeptical positions on the interpretation of experimental results (such as Uttal [2001], or Hardcastle and Stewart [2002], and Stewart and Hardcastle [2005]) as well as more moderate warnings on the current use of the technique and its limitations (such as Klein [2010], or Aktunç [2014]), it is only recently that philosophers have begun to adopt a more constructive philosophical approach to their analysis of fMRI.

One of the most commonly noted faults of fMRI experiments is the fact that they are based on a markedly indirect measure. That is, fMRI measures changes in blood-oxygenation levels (the BOLD signal), and scientists use those measures as the basis for claims about neural functioning (Logothetis et al. [2001], Bogen [2001, 2002], and Roskies [2008] discuss this in detail). Cognitive neuroscientists use fMRI in conjunction with tasks and behavioral manipulations that control the cognitive processes subjects have while in the scanner. Treating this fMRI's indirect nature as an epistemic fault has motivated and grounded much of the skeptical literature (Aktunç [2014]; Uttal [2001, 2011]; even van Orden and Paap [1997] refers to the indirect nature of the measures). Here, I join more recent efforts to articulate the positive aspects of

neuroimaging experiments on the hope that, through a better understanding of fMRI's current use, a richer assessment of its role in advancing cognitive neuroscience will be possible.

The viewpoint I will articulate revolves around the concept of exploratory experimentation, a view of experimental practices conceived not simply as a tool to test theories but as mainly centered on phenomenological description and the assessment of the very experimental approach's strengths and limitations. I will argue that certain local and idiosyncratic features of BOLD fMRI-based experimentation in contemporary cognitive neuroscience give it an exploratory character. To support this, I will focus on two related questions: In the first place, following Steinle (1997), in what sense experiments aim at stabilizing phenomena as their main objective, and, secondly, following Burian (1997)'s ideas on less established or still poorly known techniques, in what sense part of the researcher's interest rests on methodological aspects of the ongoing research. Although I do not claim that there is a specific kind of experiment that one can single out as definitely exploratory, exploration can be seen as a general trait imbuing most of contemporary fMRI-based experimentation. It is also important to state at the outset that I do not think the original concept(s) of exploratory experimentation, as proposed by Steinle (1997, 2002) and Burian (1997), can be applied to neuroscience as they have originally been proposed. As I hope to show, these proposals nevertheless constitute a rich starting point in order to ponder the presence and nature of an exploratory dimension in contemporary fMRI-based experimentation.

Two previous suggestions along the lines I will be pursuing here can be found in the literature. Towards the end of her article, Franklin (2005) mentions functional imaging instruments adopted in cognitive science as one among other fields which are undergoing what she sees as an

exploratory shift in methodology. Klein (2010b) also touches on the issue, contrasting interpretations of neuroimaging as either confirmatory or exploratory, but doesn't lay out precisely in what sense and to what extent we are understanding it as exploratory. While these are relevant, early suggestions, to my knowledge there haven't been any detailed philosophical analyses of the concept of exploratory experimentation as applied to cognitive neuroscience and functional neuroimaging.

I first present the original concept developed by Steinle and Burian as well as some relevant later refinements (section 2). I then adapt and revise the concept, attending to the specific conditions under which BOLD fMRI experiments are planned and conducted, following my two main guiding themes. Hinging on Steinle's contribution, in section 3 I describe cognitive neuroscientists' goals in terms of phenomena stabilization through extensive probing of parameter space, characterizing a particular sort of balance between the different stages at which idiosyncratic choices can be made by experimenters (in experimental designs, protocols, and data analysis techniques). Then, elaborating on Burian's take on exploratory experimentation, in section 4 I positively interpret fMRI opaqueness or inferential distance in terms of a search to define and consolidate guidelines to direct experimental design, leveraging on the technique's flexibility and compensating contributions from other existing techniques. The early motivation behind this effort was an understanding of cognitive neuroscience as a young field and, particularly, the limited scope that theories of neural function still have. Nevertheless, I will elaborate on the idea of a characteristic exploratory profile of fMRI-based experimentation, that draws it away from experimentation classically conceived as a tool to test highly elaborated hypotheses via some associated theory.

2. Exploratory experimentation: The original concept

The original concept of exploratory experimentation, as proposed by Steinle (1997, 2002) and Burian (1997), was intended to capture a kind of experimental practice that up to that moment had not been thoroughly discussed in the history and philosophy of science tradition. As Schickore (2016, p. 22) points out, they coined the expression independently, but following different objectives and emphasizing different aspects of experimental practice. Moreover, Steinle elaborates the concept drawing on examples from the early years of electromagnetism, particularly the work of Ampère and Faraday on electromagnetic induction, while Burian focuses on the emerging research on protein synthesis carried out by Brachet.

Steinle distinguishes two main characteristics that go into the exploratory aspect of his case studies: On the one hand, as already mentioned, the fact that in these cases the experimental work was not a tool for theory testing and, on the other hand, centrally, the systematic variation of experimental parameters as a defining element (Steinle 2002, p. 419) directed towards the identification of empirical regularities. For his part, Burian, besides acknowledging the absence or reduced influence of scientific theories, focuses on the presence of a diverse cluster of techniques, individually still not well-established but directed towards the same phenomenon; in this way, Burian's proposal is akin to the ideas of triangulation and experimental robustness discussed in the literature (for instance, Bechtel [2002], Weber [2005], Stegenga [2009], Roskies [2010]). Although phenomenological description as a basic goal is a common general theme, the defining features pointed out are different and independent. However, as I will try to show, both present themselves in a particular, idiosyncratic way in the

functional neuroimaging based experimentation typical of cognitive neuroscience.

(Another central characteristic that has been originally associated to exploratory experimentation, and later elaborated, among others, by Feest [2012] with an eye on psychology, is its role on concept formation, associated to the primary aim of identifying empirical regularities. Although I think this is fertile ground for future philosophical reflection on neuroscience, in what follows I am leaving this particular aspect aside, being one of the later or definitive contributions of this kind of experiments, to focus instead on the experimental process itself, its nature, and motivations.)

Following the initial proposals by these authors, there has been a series of later developments surrounding the concept of exploratory experimentation. Specially during the first decade of the new century, we find a number of contributions aimed at refining the concept or adapting it to a particular field of research. A theme of special interest concerns the kind of instruments that tends to be involved in exploratory experimentation. The contributions of Franklin (2005) and O'Malley (2007) point in this direction: According to these authors, these are instruments that (a) can be characterized as wide or high-throughput, that is, that produce or generate large amounts of data, and (b) that admit a broad range of variations and possible results. As we will see, these features are both present in the case of fMRI, and in a very prominent way. Additionally, exploratory experiments have been defined by the flexibility and breadth of the associated conceptual framework as well as by the fact that their results are interpreted on the basis of chains, series, or nets of experiments (Feest and Steinle 2016); again, we will be able to verify that these points apply in the case of the neuroscientific experimentation considered here.

As long as this is the case, particular facets of exploration could fruitfully describe at least part of the experimental practices in present day cognitive neuroscience. Assessing in what sense and to what extent the positive aspects of what makes experimentation exploratory can be found in cognitive neuroscience, and particularly in research based on functional neuroimaging, is the aim of the rest of this paper. I distinguish two facets of the exploratory character of fMRI-based experiments: one related to the experiment's goals and a second facet, concerning methodological issues surrounding the technique's use. These facets correspond to the two main features I single out, respectively, from Steinle's and Burian's proposals, that is, a focus on parameter variation aimed at phenomenological description and the presence of a variety of techniques directed towards the same phenomenon.

It is worth stressing that my intended contribution to a more thorough understanding of fMRI's use in contemporary neuroscience hinges on an elaboration of these two features, according to the way they unfold in the field. Nonetheless, this is not to say that I fully endorse neither Steinle's nor Burian's accounts as directly applicable to current neuroscientific experimental practices. As I remarked, these accounts stem from the analysis of historical case studies that are very far apart from today's cognitive neuroscience; however, they are also the result of a critical assessment of the restrictive understanding of experimentation as exclusively aimed at theory testing, an assessment that, as I intend to show, is very fruitful when discussing functional neuroimaging.

3. On fMRI-based experiments' goals

Let's recall the main feature singled out by Steinle in order to characterize explanatory experimentation: The systematic variation of experimental parameters mainly directed to the identification of empirical regularities. Steinle developed his concept by considering the simple, planned variation of a single parameter –the selected case of Ampère's 'astatic magnetic needle' had the relative position between a needle and a wire as its parameter of systematic variation. In the case of fMRI, on the one side, the situation is very different and much more complex, but, on the other side, its underlying rationale is analogous to the simpler cases considered by Steinle. I am going to maintain that many first-rate functional neuroimaging (and, particularly, fMRI) based experiments have phenomena stabilization as their primary objective: More precisely, they aim at establishing stable experimental arrangements for phenomenological description through the careful modification of experimental parameters as well as other relevant aspects of the experimental setting. Let's lay out these ideas carefully.

Firstly, there are a number of reasons why obtaining and identifying empirical regularities may be the main aim of this kind of experiments. We can group them within two central problems the experimenter has to face: (a) The problem of the high variability of target phenomena and (b) the problem of the great flexibility we find at the level of experimental setting as well as in the subsequent treatment of the obtained data. Both problems underpin the relevance that the control of variables and the associated exploration and stabilization of phenomena¹ under study assume in a field such as

¹ I am adopting Bogen and Woodward's (1988) account of the notion of 'phenomenon', that is, understood as a set of general and

cognitive neuroscience and, especially, functional neuroimaging-fueled research. Let's see how this is so.

A central issue surrounding phenomena variability in the case at hand has to do with the fact that phenomena of interest in cognitive neuroscience are both cognitive and neural. This first point calls for further clarification. In fact, one might also think that target phenomena in cognitive neuroscience are either cognitive or neural, depending on the particular line of research under consideration: Crucially, one might think that in direct (or forward) inference studies – such as the common localization studies, aimed at determining the anatomical structures responsible for a specific cognitive operation, or implementation studies, aimed at determining the parameters of neural activity that mediate in a particular cognitive process– we are looking at a properly neural phenomena of interest, while in the case of reverse inference studies, where on the basis of what is known about the neural correlates of certain cognitive processes a given behavior and the set of operations needed to realize it are the focus of research, we face a cognitive or psychological target phenomenon. However, from a more global outlook, if one takes direct and reverse inference studies as providing partial perspectives on the phenomena neuroscientists are interested in, the resulting view is that of a complex target phenomenon, at the same time neural and cognitive. This view also takes directly into account, on the one hand, the fact that standard contemporary reverse inference approaches, such as multi-voxel pattern analysis – at least since landmark studies such as Mitchell et al. (2008) and Kay et al. (2008)–, use data collected from localization or other direct inference studies in order to constrain the decoding and cognitive interpretation of brain activation

relatively stable aspects of the world that the relevant scientific community aims at explaining or predicting.

patterns, and, on the other hand, the more general fact that mixed or hybrid approaches, directed at leveraging on the relative strengths of different experimental techniques, are commonplace (an aspect I'll go back to in section 4).

As is well known, fMRI's main appeal (as is also the case for other functional neuroimaging techniques) is that it allows for the non-invasive study of neuronal activity based on the cognitive performance of human subjects. Now, for the one part, obtaining clarity and accuracy concerning cognitive phenomena has proven to be more difficult than could have originally appeared to be. Feest (2011) focuses on the concept of 'phenomenon' in the cognitive sciences and stresses the importance of the production of stable phenomena in the laboratory. She understands this involves stabilizing the fit between what she calls surface phenomena (i.e., behavioral regularities) and hidden phenomena (i.e., underlying cognitive operations and, ultimately, neural processes).

I think part of Feest's contribution here has to do with minimizing scientific explanation as a dominant topic in recent philosophical work (a central claim I advanced in *Venturelli* [2016]): One of the main upshots this imbalance in recent philosophy of neuroscience has is the tendency to take phenomena that are objects of study in neuroscience as mere *explananda*, already there to be tackled, leaving aside the processes through which those are delimited and produced, in favor of a restricted focus on the corresponding appropriate kinds of *explanantia*. The central point here is that, in general, neither cognitive nor neural phenomena are clearly delimited, and a good deal of experimental work aims to identify the boundaries of the phenomena of interest.

Feest (2011, 2017) takes as an example the 'chunking' effect in short-term memory studies: That is, the fact the in experimental studies on this phenomenon, subjects can recall more items when they chunk the information they are

instructed to recall. In order to determine short-term memory capacity, it is crucial to consider how chunking works and, in this sense, the experimental question becomes one of identifying the chunks and how they are exploited by subjects: 'None of these are explanatory questions. They are, rather, theoretical and methodological questions about how to investigate and describe a given object of research' (Feest 2017, p. 1168). We can mention similar cases such as working memory (Ma et al. 2014), episodic memory (Cheng and Werning 2016), or attention (Taylor 2017), each one counting different and sometimes divergent models being developed and used. Behind central concepts such as these, we find stories filled with difficulties to circumscribe and experimentally isolate the associated phenomena, previously to any attempt at explaining them. Figdor (2013) and Sullivan (2015, p. 37) have framed worries along these same lines hinging on the typically loose, underconstrained use of cognitive constructs by working neuroscientists:

In the absence of discipline-wide criteria for guiding cognitive inferences made from an everexpanding basis of kinds of brain-behavior associations, it is left to individual researchers to determine the nature of the cognitive constructs inferred in a given study—that is, to choose which cognitive or mental concepts and terms are appropriate for interpreting and reporting their results and drawing implications. (Figdor 2013, p. 107)

Moreover, a related difficulty concerns almost any experiment with functional neuroimaging in the field, regardless of the direction of inference: 'To some degree, it is necessary to rely on the previous decomposition of a given task in order to identify the neural bases responsible for the

one operation we assume has to be necessarily realized for its performance. This is most notable if we look at a still widespread (although no longer dominant) localizationist paradigm, the subtraction paradigm, based as it is on the differences in neural activation found between two conditions that are assumedly identical, with the exception of only one of the operations that collectively compose a given task. In this way, apart from a number of additional issues and boundary conditions to be considered, the success of the experiment is going to depend, crucially, on the previous decomposition of the task into individual operations according to the information processing deemed necessary for its successful performance.

Now, in the case of neural phenomena the problem of their identification and stabilization is, as could be expected, also very pressing if we just consider the basic fact that fMRI experimental data are obtained through the detection of the level of metabolic activity: A complex variable related to blood flow, volume, and oxygen demand, which is compared to a baseline condition, generally with the subject in a resting state. A very eloquent example is that of the considerable differences that exist in the vasculature (blood supply) in different regions of any subject's skull as well as between different experimental subjects. Bellgowan et al. (2003), for example, is an early effort to compare the temporal properties in hemodynamic response between different areas in the brain (in this particular case, on a lexical decision making task) precisely in order to discard the differences in detected activity which could bear no relation to the underlying neural function but only to the differential hemodynamic responses between brain areas.

If one considers the problem of identifying a functionally relevant neural event, one of the main and most discussed difficulties lies on the temporal hiatus between hemodynamic activity, which is actually detected by the

scanner (via the reaction of oxyhemoglobin to a strong magnetic field), and neuronal electric activity, which is thought to be functionally relevant. This is one of the most prominent and known grounds behind the idea of fMRI affording a markedly indirect kind of detection. Constable spells this out clearly:

From a single event, say, for example, a flash of bright light lasting only 100 milliseconds, the blood-flow response in the primary visual cortex may start to increase some two seconds after the event, it may peak between six to eight seconds and return to baseline by 16 to 20 seconds after this single 100-millisecond flash of light. (Constable 2010, pp. 81–82)

As a prominent example, a vexing consequence of this, common in fMRI research, has to do with delimiting regions or networks, which includes their localization through differing strategies such as cytoarchitecture via Brodmann's areas, functional localization, or anatomical landmarking, and even defining what a 'region' or 'network' actually consists of. As Yarkoni and collaborators put it: "There is currently little consensus about how to classify or group different brain regions, networks, experimental tasks or cognitive functions, let alone how to develop mappings between different levels of description" (Yarkoni et al. 2010, p. 491).

Another issue behind the relevance of variable control and neural phenomena stabilization has been pointed out by Sullivan (2009), who claims that even small differences between experimental protocols can direct towards entirely different phenomena. Building on this, Sullivan also lays out a warning call on the fact that more attention should be paid to the multiplicity of experimental protocols, understood as

a pressing problem in need to be more overtly tackled. Now, for my present interests, the take-home message here is that a host of diverse assumptions –concerning, on the one hand, the three-way relation between the signals detected by the scanner, the underlying hemodynamic patterns, and the electrophysiological activity, and, on the other hand, the structural and functional architecture responsible for neural activity– will inevitably structure experimental design, data pre-processing, and the analytic approach to the obtained data.

This takes us to the second issue I pointed out, that is, the problem of the great flexibility we find in experimental systems in the field –and it is here where Steinle’s emphasis on parameter variation comes into play, although, as we will see, in a different, idiosyncratic way. This flexibility has mainly to do with the control of variables and the stabilization of the phenomena under experimental study, as we saw, a prime objective of the kind of experimentation I am considering. The central point here concerns the fact that experimental systems in functional neuroimaging based research admit a very large number of variations, which can obtain within each one of the many stages that compose the planning and execution of an experiment or of a series of experiments as well as the results’ analysis. This fact, despite being deceptively plain in itself, is philosophically critical to my contribution.

What I call flexibility in experimental systems can be found in data production as well as in data interpretation processes; it is a product of the complexity of techniques such as fMRI in the way relevant parameters are fixed and in the treatment of the obtained data. Regarding data production, let us consider some issues concerning experimental design. Early on, design selection implies addressing the experiment’s temporal structure: Designs can be defined depending on whether the experimental subject

is at rest or active during detection, and in this last case they can be, among other possible solutions, blocked designs, where the activation condition is generated at regular intervals, or event-related designs, where researchers can determine the timing of the activation condition based on the controlled presentation of stimuli. Favoring the presentation of stimuli at regular over irregular intervals, or vice versa, can respond to a number of reasons: For instance, although blocked designs in general terms grant a greater statistical power, if the focus is set on the temporal dimension of a given effect or if the researcher's interest rests mainly on the cognitive aspect, then an event-related design may confer a greater degree of control to the experimenter and be more advantageous. This in turn can always be further assessed *vis-à-vis* the kind of cognitive phenomenon at hand –consider, for instance, cases where the intervention of the experimental subject's expectation as an additional factor could be disdained or, conversely, where it would make the obtained results inadmissible. The fact that, as Aguirre (2010, p. 61) points out, blocked and event-related designs are better conceived as extremes on a continuum, with several mixed and intermediate solutions available, further contributes to the flexibility of experimental arrangements concerning this particular aspect.

This leads us to the related issue of the manipulation and control of cognitive processes, so much for their direct study as for the study of their neuronal bases or implementation conditions. Now, the already mentioned cognitive subtraction paradigm and the related cognitive conjunction paradigm are only two popular variants for this, with the same underlying logic: Through the subtraction or addition of the obtained data in two conditions that crucially differ in the one cognitive process of interest, the idea is to attempt to estimate the relevant differences at the level of oxygen consumption, which are then associated to this particular

process, always on the basis of a previous decomposition of the task to be realized by the experimental subjects.

Letting aside the various critiques that have been raised at this kind of paradigms –concerning, for instance, the underlying cognitive models, the assumed existence of discrete and successive stages, the question of what exactly is being subtracted, and so on–, there can be no doubt that they are grounded on previously defined cognitive models as well as on a battery of different assumptions, among which we can mention the assumption of space-time separability (Alexander et al. 2015, p. 367) and, probably the most debated on, the assumption of pure insertion, that is, the idea that there exist no relevant interactions between the different processes that compose a given cognitive task and that, in this sense, these can be subtracted or added without risk of producing undesired effects or artifacts (for an early example, see Friston et al. [1996]). In this last case, we are dealing with a general assumption that determines, or is at least very influential on, the use of the technique. Moreover, there are other kinds of designs: For instance, parametric designs, which focus on the variation of a range of different levels of the same parameter, with the aim of identifying correlations between the imaging signal and the parameter values, or factorial and adaptation designs, which are more centered on the task and the array of stimuli to be presented. These different kinds of experimental designs also offer disparate benefits and reflect different assumptions concerning the phenomena at hand as well as the scanner's detection process.

The global conclusion to be extracted from the above is that pondering a design over another carries with it a number of related decisions that, generally speaking, have to be taken situatedly, considering the phenomena of interest, and based on the internal practices of a given laboratory and the degree of progress made so far in the experimental work that is

being carried on. Savoy expresses something along these lines, as follows:

Experimental design is not as self-consistent nor as theory-based as most of the other aspects of fMRI-based research. While there are some general principles that can be taught [...], many aspects of experimental design can best be thought of as a collection of tricks and techniques that are useful in specific situations. (Savoy 2005, p. 263)

This, I think, is a very powerful appreciation adding to the ensuing considerations revolving the concept of exploratory experimentation as applied to cognitive neuroscience. In the same vein, Churchill et al. (2011) argue for individual subject-based preprocessing and analysis pipelines for fMRI. Let's then turn now to data analysis and interpretation.

We find here an analogue situation of flexibility and high variability. A first point of variation in the treatment of fMRI data has to do with setting activation thresholds. The criteria to define whether a given threshold is adequate so as to take into account only significant activations are not extensively widespread inasmuch as they depend on several factors, such as the (generally high) number of voxels to be considered, the (generally low) number of experimental subjects for each study, the multifarious sources of artifacts and noise to be controlled, and so on (see, for example, Bogen [2001, 2002], and Yarkoni [2009]). This has been and continues to be a much discussed, relatively arbitrary aspect that, as it may well be expected, strongly influences the results' variability (Constable 2010, p. 81). Once again in tune with Savoy's opinion, cited earlier, the fact that a particular level of activation is considered significant is going to depend in

some measure on the research program a researcher may be working in, under the general purpose to avoid false positives or, conversely, false negatives (that is, leaving aside relevant activations for the task or effect under study).

The selection of analytic approach is another crucial factor of variation. The main cut-off point is the one between univariate type of analysis, where the focus is set on signal's activity or intensity, and multivariate analyses, where the focus is set on interaction. In the first, more common case, thresholds are applied to data-sets before analysis and, once this is done, statistical tests are applied to each voxel taken independently; in the second case, all produced patterns (of interest) are analyzed (if we exclude strategies to avoid statistical issues arising from multiple comparisons). In general terms, while the first approach to data analysis is oriented towards the average activation in a region or group of neurons, the second one is oriented towards the pattern of activation in a region (Poldrack 2012, p. 1218).

Favoring one kind of analytic approach over the other is decisive for entire lines of experimental research and their ultimate set of goals, as can be appreciated if we consider research on short-term visual memory. The bulk of research in this subfield has been led by the assumption that the retention of visual information corresponds to high levels of activity during the delay period of the task; an assumption that crystallizes in the almost exclusive use of univariate analyses. Although this situation is starting to reverse, this particular bias has defined the recent history of neuroscientific studies on this particular phenomenon (Postle 2015). To end with a compelling point on flexibility, Carp (2012) encountered a substantial methods-related variability in activations resulting from a single experiment in an attempt to estimate the flexibility of a diverse range of analysis procedures commonly used in fMRI experiments.

With this tour, partial as it may be, across the different stages that make up an fMRI experiment, I have tried to illustrate the complexity and, especially, the number of variations that each one of these stages admits. This kind of technique, together with the way it is used in the field, represents a particular version of exploratory experimentation, exhibiting many of the features that have been variously (and more recently) highlighted regarding the concept: The experimental systems associated with this technique produce large amounts of data, allow for a wide range of variations and results, involve series or nets of experiments, and, given experimental complexity and the absence of a robust theory to support it, present a marked strand of parameter variation.

Every one of these aspects presents itself in an idiosyncratic way. Particularly, this last aspect originally stressed by Steinle is admittedly far from the simple, systematic variation of a single experimental parameter for the identification of empirical regularities. It is important to make clear that, on account of the highlighted complexity in contemporary fMRI research, it would in this sense be very difficult to find the kind of systematicity in parameter variation that very simple experimental setups –such as the ones Steinle focused on– enable. This central difference notwithstanding, this certainly need not imply that we may not find cases of this kind of simple parameter variation in the history or present of neuroscience. One example worthy of thorough attention is H. K. Hartline's early study (published in 1938 and 1940) on the receptive field properties of the frog's ganglion cells: The discovery of on-, on-off, and off-fibers was the result of a series of experiments where he used a tiny light spot of 0,05 mm and consecutive movements in steps of 0,01 mm to map the area where an electrical response could be elicited.

Now, this kind of simple variation and concomitant systematicity—in Hartline’s particular case, systematicity just being an expected, unplanned result of the low number of experimental parameters—is patently not to be found in the complex, multistage experiments I am reviewing. The kind of parameter variation typical of contemporary fMRI-based research is more akin to a probing, a wide breadth exploration, of parameter space. A telling example of an exploratory strategy strictly set within the domain of data analysis is Hindy et al. (2016)’s study on the role of pattern completion in the hippocampus to ground memory-based expectations in visual cortex. In their use of multivariate pattern analysis so as to identify representations in the relevant brain areas, the authors have performed almost every sensible statistical analysis in order to make the best of the obtained data, thus revealing a strong reliance on parameter variation to help overcome the inferential weaknesses of fMRI.

I suggest the kind of exploratory strategies deployed in functional neuroimaging experiments can be usefully seen as a sort of orchestration between the different stages of the execution of an experiment, from its design to data analysis and interpretation, and not only on one or more experimental parameters. Building on an understanding of fMRI-based experimentation as an inherently complex process, composed of many related, individually multifaceted, stages, the idea of orchestration points to the fact that the decisions made at any one stage are bound to influence, and are made based on, other decisions made at other stages. The various decisions involved in this process are made so much on the basis of the theoretical, albeit more or less unspecific, goals behind a given series of experiments than based on the multiple assumptions of different scope and generality maintained at the different stages. In the already mentioned case, pointed out by Postle (2015), a

single, general assumption concerning peaks of activity during the delay period interpreted in cognitive terms (i.e., as retention of visual information) has altogether set aside a whole spectrum of multivariate methods for statistical data analysis.

As I have tried to show, each one of these stages is a world of its own: It presents a complex situation, admitting a high internal variability, which exhibits a compromise between different factors, generally associated with requirements concerning complexity, anatomical and physiological adequacy, and interpretability, in addition to mathematical and tractability requirements. Of course, and this is crucial, it is not complexity *per se* which makes exploratory experimentation plausible in the case at hand, but rather how it unfolds on the backdrop of the limited knowledge of brain function. It is important to bear in mind that, just as exploratory experimentation has traditionally been associated with early, “theory light” stages of a particular field of research –as is the case of the inaugural accounts by Steinle and Burian as well as later elaborations by Feest and Steinle (2016) or Radder (2003)–, exploration in cognitive neuroscience can only be pondered *vis-à-vis* an accompanying appraisal of this particular field as a young and accordingly still immature one².

One of the main idiosyncratic aspects of this for functional neuroimaging research is the fact that, given the general stage of development of the field, there are only partial restrictions available on the assumptions that operate at each stage of the production and analysis of data. Bogen

² Here “immature” is to be understood in a loosely Kuhnian sense inasmuch as it can be said that the field of cognitive neuroscience is not structured around dominant, overarching theories being further developed and articulated. It should in no way be understood in a pejorative sense.

(2002, pp. S66–S67) has emphasized this particular issue, pointing out that the use of fMRI and other functional neuroimaging techniques such as positron emission tomography includes several aspects that are beyond any possible control by the experimenter and, furthermore, cannot be verified against independently obtained results. This takes us to the second facet of the exploratory character of fMRI-based experiments I want to stress.

4. On fMRI-based experiments' methodological focus

Considering the extensive use of functional neuroimaging in the field, its exploratory character, such as I portray it, may put the actual goals pursued in current research under a different light: In particular, how relevant the identification of empirical regularities is, understood here as a primary experimental goal, over and above explanatory goals. This is in line with the idea that explanation in cognitive neuroscience is one among many of its goals, whose urgency is reduced *vis-à-vis* its degree of development and the complexity of the experimental (and modeling) practices that define it (Venturelli 2016). One may think of this reframing of experimental goals as setting the groundwork in order to formulate more precise hypotheses to be tested and plan more systematic future investigations. As Waters (2007, p. 279) puts it: '[T]he aim of exploratory experimentation is to generate significant findings about phenomena without appealing to a theory about these phenomena for the purpose of focusing experimental attention on a limited range of possible findings.'

Again, although we are dealing with a very different kind of practice compared to the systematic variation of experimental parameters emphasized by Steinle, the central idea of an exploratory search to identify and describe

phenomena lingers on. In this case, the wide-ranging goal is stabilizing the experimental system as such, *vis-à-vis* the high variability of results and flexibility the use of a technique such as fMRI admits. Accordingly, we can appreciate how part of the interest in the experiments rests on the methodological aspects of the research themselves, along the lines of what, in turn, Burian pointed out as one defining feature of exploratory experimentation.

It can be hardly doubted that, in the case of fMRI, we are dealing with a not fully established technique. Technical and methodological issues have been, and still are, widely debated, at varying degrees of generality, as Roskies (2008, pp. 25-26) pointed out regarding data analysis: ‘Technical debates about proper statistical and data analysis abound in the literature, ranging from questions about how to correct for multiple comparisons to whether analysis should be hypothesis-driven or whether brute-force statistics suffice.’ Common disputes surrounding the adoption of subtraction paradigms, the use of functional localizers (for instance, Saxe et al. [2006], Friston et al. [2006]), or the use of multivariate over univariate techniques (for instance, Haxby [2010], Serences and Saproo [2011]) are just a few far-reaching examples of this. Adding to this point, innovations, especially in how data is analyzed, are happening constantly at both big and small scales, an aspect which Wright (2017) has forcefully stressed, pointing out the need for a much more thorough attention on it by philosophers of neuroscience.

Generally speaking, and building on my previous argument on flexibility, the guidelines for experimental practice are relatively unspecific inasmuch as they allow for a variety of experimental arrangements depending on the relative reliability associated with one or more of the involved assumptions as well as the specific interests that guide research. This in turn leads to favoring some possible

aspects to be focused on over others –such as temporal over spatial resolution, the estimation of a quantity over its detection, or the extent to which the obtained results can be generalized.

Although this may possibly be a transitory situation, due to the state of the art in the field and the available knowledge on this kind of techniques, the sort of exploratory experimentation revolving functional neuroimaging is mostly oriented towards settling on stable experimental designs and protocols, in order to assess the relevance of the different known limitations (Constable 2010) that ultimately affect the obtained results. Along these lines, one can think about the search-space of exploratory experiments in terms of rules to direct experimental design and data analysis or, as Franklin (2005, p. 895) points out, in terms of heuristics for experimental methods.

This last point is especially relevant for the case at hand. The kind of experimentation considered here is exploratory not only because it doesn't mainly aim to confirm results but because the methodological dimension is *per se* an important part of the theoretical interest driving it. Two recent tendencies strengthen this idea: On the one hand, a group of recent wake-up calls concerning the reliability and validity of results extracted from functional neuroimaging experiments (Button et al. 2013; Eklund et al. 2016) and, on the other hand, increasingly frequent studies aimed at checking results obtained with other techniques against those obtained with fMRI (Nunez and Silberstein 2000; Strangman et al. 2002; Scheeringa et al. 2016). In this way, achieving stable results and managing the experimental setting to reproduce reliably certain results are in themselves prevailing objectives –see for example Gilmore et al. (2017) and Poldrack et al. (2017) for recent assessments and proposed solutions surrounding reproducibility of results and transparency of scientific practices.

On a related note, it bears mentioning that the importance of triangulation from multiple techniques in neuroscience cannot be overstated. In this sense, Burian's take on the concept of exploratory experimentation is very telling for the case at hand since, besides the general goal of phenomenological description also stressed by Steinle, the focus here is on a varied cluster of still poorly or not fully established as well as not fully understood techniques pointed towards the same phenomenon. This kind of strategy, although reaches an additional aspect with respect to those I focused on concerning specifically fMRI use, is very common in recent experimental practices in the field, as Bandettini straightforwardly puts it:

The scope and success of the best fMRI research and applications are defined by how well the paradigms, processing, and interpretation of results are integrated with other brain function assessment techniques. This integration has yielded far greater insights into human brain function than any that fMRI could provide on its own. (Bandettini 2006, p. 195)

Elliott has also highlighted this particular aspect of exploratory experimentation, which strongly resonates in contemporary cognitive neuroscience broadly conceived, defining it as 'an attempt to study a phenomenon using as many tools and techniques as possible so as to understand it more fully and to gain more solid epistemic access to it' (Elliott 2007, p. 328). Although this is one among other central features of how fMRI is deployed in the field, it certainly presents an accurate outlook of a strong trend in recent neuroimaging research –one that, although certainly

not new, is surely strengthening and will most probably continue to do so in the near future.

On a final note, I think necessary to stress once again the wide, fertile ground still to be covered, from a philosophical point of view, if we abandon the tendency to conflate experimental aims with some notion of underlying theory. I think this is particularly so in the case of cognitive neuroscience. My views here go beyond this restrictive focus on experimental practices in the field, towards an approach much closer to Ian Hacking's ideas on the self-justification and stability of laboratory sciences: The still emerging path towards the stabilization of experimental systems, understood as an essential goal in functional neuroimaging research, is embraced by my views on the exploratory character of this kind of experimental work.

5. Conclusion

I laid out my views on how the concept of exploratory experimentation can help characterize the experimental side of contemporary functional neuroimaging based cognitive neuroscience. I have taken the concept into consideration on the basis of two main strands: The experiment's goals and the relevance of its methodological facets. I focused on fMRI as a prominent example of an experimental technique widespread across the field, in order to assess in what sense there can be said to be exploratory experimentation in this particular case. Although it is important to stress that we won't find anything like a *sui generis* kind of experiment—and, in this sense, I can endorse Schickore (2016)'s claim that the notion fails to identify a specific kind of experimental practice—the concept is telling and delves deep in the experimental practices characterizing this particular area of neuroscience.

I maintained that the primary goal of fMRI-based experimentation is, in line with the understanding of the concept by Steinle (1997, 2002), the variation of parameters as well as relevant aspects of the experimental system for phenomenological description. In terms of Steinle himself, it can be said that, in the case of functional neuroimaging experiments, ‘the relation between instrumentation, theory, and experimentation has distinct features...’ (Steinle 1997, p. S65), which, as I have argued, can be associated to an exploratory character that here presents itself in a particular and idiosyncratic way. The complexity of this kind of technique and, specially, the flexibility in the way parameters are fixed as well as in the treatment of data, entails a considerable effort dedicated to what I referred to as an orchestration between the different stages of the execution of an experiment. On the other hand, I highlighted Burian’s (1997) idea that exploratory experimentation finds fertile ground in the case of still poorly known techniques, where an important part of the experiment’s interest rests on the methodological aspects involved in establishing experimental paradigms.

On a final note, a clarification on the scope of my contribution is in order. Although we have been dealing with an extensive approach in the cognitive neurosciences, exploratory experimentation, in the particular fashion described here, does not characterize the whole experimental side of the field. There are many research lines, particularly in advanced stages of continued studies, which are guided by very specific hypotheses, rooted in local theories. Under illustrative purposes, we could take the well-known example of research on spatial orientation in the hippocampus – particularly, the lines of research led by 2014 Nobel winners John O’Keefe, May-Britt Moser, and Edvard Moser – led, for instance, by questions concerning the kind of representation (allocentric or egocentric) or the specific contribution of the

firing rate of a particular cell to a given representation (see Bechtel [2016] for a philosophical survey). This sort of experimental questions stems from the postulation of a system, localized in the hippocampus, for spatial navigation in the rat, following a number of previous discoveries (for instance, concerning the existence of types of place cells, grid cells, head-direction cells, and so on). Again, my considerations here apply specially to the case of fMRI due to the highly indirect character of the detection it allows as well as the status of the knowledge regarding the technique and the phenomena this is directed to.

References

- Aguirre, Geoffrey K. 2010. "Experimental Design and Data Analysis for fMRI." In *Bold/fMRI. A Guide to Functional Imaging for Neuroscientists*, ed. Scott H. Faro, and Feroze B. Mohamed, 55–70. New York: Springer Science.
- Aktunç, M. Emrah. 2014. "Severe Tests in Neuroimaging: What We Can Learn and How We Can Learn It." *Philosophy of Science* 81: 961–973.
- Alexander, David M., Chris Trengove, and Cees van Leeuwen. 2015. "Donders is Dead: Cortical Traveling Waves and the Limits of Mental Chronometry in Cognitive Neuroscience." *Cognitive Process* 16: 365–375.
- Bandettini, Peter. 2006. "Functional Magnetic Resonance Imaging." In *Methods in Mind*, ed. Carl Senior, Tamara Russell, and Michael Gazzaniga, 193–235. Cambridge: The MIT Press.

- Bechtel, William. 2002. "Aligning Multiple Research Techniques in Cognitive Neuroscience." *Philosophy of Science* 69: S48–S58.
- Bechtel, William. 2016. "Investigating Neural Representations: The Tale of Place Cells." *Synthese* 193/5: 1287–1321.
- Bellgowan, Patrick S.F., Ziad S. Saad, and Peter Bandettini. 2003. "Understanding Neural System Dynamics through Task Modulation and Measurement of Functional MRI Amplitude, Latency, and "Width." *Proceedings of the National Academy of Sciences* 100/3: 1415–1419.
- Bogen, James. 2001. "Functional Image Evidence: Some Epistemic Hot Spots." In *Theory and method in neuroscience*, ed. Peter Machamer, Peter McLaughlin, and Rick Grush, 173–199. Pittsburgh: University of Pittsburgh Press.
- . 2002. "Epistemological Custard Pies from Functional Brain Imaging." *Philosophy of Science* 69/3:59–71.
- Bogen, James, and James Woodward. 1988. "Saving the Phenomena." *Philosophical Review* 97/3: 303–352.
- Burian, Richard. 1997. "Exploratory Experimentation and the Role of Histochemical Techniques in the Work of Jean Brachet, 1938–1952." *History and Philosophy of the Life Sciences* 19/1: 27–45.
- Button, Katherine S., John P.A. Ioannidis, Claire Mokrysz, Brian A. Nosek, Johnathan Flint, Emma S.J. Robinson, and Marcus Munafò. 2013. "Power Failure: Why Small Sample Size Undermines the Reliability of Neuroscience." *Nature Reviews Neuroscience* 14: 365–376.

- Carp, Joshua. 2012. “On the Plurality of (Methodological) Worlds: Estimating the Analytic Flexibility of fMRI Experiments.” *Frontiers in Neuroscience* 6: 149. DOI 10.3389/fnins.2012.00149
- Cheng, Sen, and Markus Werning. 2016. “What is Episodic Memory if it is a Natural Kind?” *Synthese* 193: 1345–1385.
- Churchill, Nathan W., Anita Oder, Hervé Abdi, Fred Tam, Wayne Lee, Christopher Thomas, Jon E. Ween, Simon J. Graham, and Stephen Strother. 2011. “Optimizing Preprocessing and Analysis Pipelines for Single-Subject fMRI. I: Standard Temporal Motion and Physiological Noise Correction Methods.” *Human Brain Mapping* 33/3: 609–627.
- Constable, R. Todd. 2010. “Challenges in fMRI and its Limitations.” In *Bold fMRI. A Guide to Functional Imaging for Neuroscientists*, ed. Scott H. Faro, and Feroze B. Mohamed, 71–92. New York: Springer Science.
- Eklund, Anders, Thomas E. Nichols, and Hans Knutsson. 2016. “Cluster Failure: Why fMRI Inferences for Spatial Extent Have Inflated False-Positive Rates.” *Proceedings of the National Academy of Sciences* 113: 7900–7905.
- Elliott, Kevin. 2007. “Varieties of Exploratory Experimentation in Nanotoxicology.” *History and Philosophy of the Life Sciences* 29/3: 313–336.
- Feest, Uljana. 2011. “What Exactly is Stabilized When Phenomena are Stabilized?” *Synthese* 182/1: 57–71.
- . 2012. “Exploratory Experiments, Concept Formation, and Theory Construction in Psychology.” In *Scientific Concepts and Investigative Practice*, ed. Uljana Feest, and Friedrich Steinle, 167–189. Berlin:

DeGruyter.

- _____. 2017. “Phenomena and Objects of Research in the Cognitive and Behavioral Sciences.” *Philosophy of Science* 84/5: 1165–1176.
- Feest, Uljana, and Friedrich Steinle. 2016. “Experiment.” In *The Oxford Handbook of Philosophy of Science*, ed. Paul Humphreys, 275–295. Oxford: Oxford University Press.
- Figdor, Carrie. 2013. “What is the ‘Cognitive’ in Cognitive Neuroscience?” *Neuroethics* 6/1: 105–114.
- Franklin, Laura. 2005. “Exploratory Experiments.” *Philosophy of Science* 72/5: 888–899.
- Friston, Karl J., Cathy Price, Paul Fletcher, Rachel Moore, Richard Frackowiak, and Ray Dolan. 1996. “The Trouble with Cognitive Subtraction.” *NeuroImage* 4/2: 97–104.
- Friston, Karl J., Pia Rotshtein, Joy J. Geng, Philipp Sterzer, and Rik Henson. 2006. “A Critique of Functional Localizers.” *NeuroImage* 30/4: 1077–1087.
- Gilmore, Rich O., Michele T. Diaz, Brad A. Wyble, and Tal Yarkoni. 2017. “Progress toward Openness, Transparency, and Reproducibility in Cognitive Neuroscience.” *Annals of the New York Academy of Sciences* 1396: 5–18.
- Hardcastle, Valerie, and C. Matthew Stewart. 2002. “What Do Brain Data Really Show?” *Philosophy of Science* 69: 72–82.
- Haxby, James. 2010. “Multivariate Pattern Analysis of fMRI Data: High-Dimensional Spaces for Neural and Cognitive Representations.” In *Foundational issues in human brain mapping*, ed. Stephen J. Hanson, and

- Martin Bunzl, 55–68. Cambridge: The MIT Press.
- Hindy, Nicholas C., Felicia Y. Ng, and Nicholas B. Turk-Browne. 2016. “Linking Pattern Completion in the Hippocampus to Predictive Coding in Visual Cortex.” *Nature Neuroscience* 19: 665–667. doi:10.1038/nn.4284.
- Kay, Kendrick N., Thomas Naselaris, Ryan J. Prenger, and Jack L. Gallant. 2008. “Identifying Natural Images from Human Brain Activity.” *Nature* 452/718: 352–355.
- Klein, Colin. 2010a. “Images Are not the Evidence in Neuroimaging.” *The British Journal for the Philosophy of Science* 61/2: 265–278.
- Klein, Colin. 2010b. “Philosophical Issues in Neuroimaging.” *Philosophy Compass* 5/2: 186–198.
- Logothetis, Nikos, Jon Pauls, Mark Augath, Torsten Trinath, and Axel Oeltermann. 2001. “Neurophysiological Investigation of the Basis of the fMRI Signal.” *Nature* 412: 150–157.
- Ma, Wei J., Masud Husain, and Paul Bays. 2014. “Changing Concepts of Working Memory.” *Nature Neuroscience* 17/3: 347–356.
- Mitchell, Tom M., Svetlana V. Shinkareva, Andrew Carlson, Kai-Min Chang, Vicente L. Malave, Robert A. Mason, and Marcel A. Just. 2008. “Predicting Human Brain Activity Associated with the Meanings of Nouns.” *Science* 320/5880: 1191–1195.
- Nunez, Paul L., and Richard Silberstein. 2000. “On the Relationship of Synaptic Activity to Macroscopic Measurements: Does Co-Registration of EEG with fMRI Make Sense?” *Brain Topography* 13/2: 79–96.
- O’Malley, Maureen. 2007. “Exploratory Experimentation

- and Scientific Practice: Metagenomics and the Proteorhodopsin Case.” *History and Philosophy of the Life Sciences* 29: 337–360.
- Poldrack, Russell. 2012. “The Future of fMRI in Cognitive Neuroscience.” *NeuroImage* 62: 1216–1220.
- Poldrack, Russell, Chris I. Baker, Joke Durnez, Krzysztof J. Gorgolewski, Paul M. Matthews, Marcus R. Munafò, Thomas E. Nichols, Jean-Baptiste Poline, Edward Vul, and Tal Yarkoni. 2017. “Scanning the Horizon: Towards Transparent and Reproducible Neuroimaging Research.” *Nature Reviews Neuroscience* 18/2: 115–126.
- Postle, Bradley. 2015. “The Cognitive Neuroscience of Visual Short-Term Memory.” *Current Opinion in Behavioral Sciences* 1: 40–46.
- Radder, Hans. 2003. “Technology and Theory in Experimental Science.” In *The Philosophy of Scientific Experimentation*, ed. Hans Radder, 152–173. Pittsburgh: University of Pittsburgh Press.
- Rosa, Maria J., Jean Daunizeau, and Karl Friston. 2010. “EEG/fMRI Integration: A Critical Review of Biophysical Modeling and Data Analysis Approaches.” *Journal of Integrative Neuroscience* 9/4: 453–476.
- Roskies, Adina. 2008. “Neuroimaging and Inferential Distance.” *Neuroethics* 1/1: 19–30.
- Roskies, Adina. 2010. “Saving Subtraction: A Reply to van Orden and Paap.” *British Journal for the Philosophy of Science* 61/3: 635–665.
- Savoy, Robert. 2005. “Experimental Design in Brain Activation MRI: Cautionary Tales.” *Brain Research*

Bulletin 67: 361–367.

- Saxe, Rebecca, Matthew Brett, and Nancy Kanwisher. 2006. “Divide and Conquer: A Defense of Functional Localizers.” *NeuroImage* 30/4: 1088–1096.
- Scheeringa, René, Peter J. Koopmans, Tim van Mourik, Ole Jensen, and David Norris. 2016. “The Relationship between Oscillatory EEG Activity and the Laminar-Specific BOLD Signal.” *Proceedings of the National Academy of Sciences* 113/24: 6761–6766.
- Schickore, Jutta. 2016. “‘Exploratory Experimentation’ as a Probe into the Relation between Historiography and Philosophy of Science.” *Studies in History and Philosophy of Science* 55/1: 20–26.
- Serences, John T., and Sameer Saproo. 2012. “Computational Advances towards Linking BOLD and Behavior.” *Neuropsychologia* 50/4: 435–446.
- Stegenga, Jacob. 2009. “Robustness, Discordance, and Relevance.” *Philosophy of Science* 76: 650–661.
- Steinle, Friedrich. 1997. “Entering New Fields: Exploratory Uses of Experimentation.” *Philosophy of Science* 64: S65–S74.
- _____. 2002. “Experiments in History and Philosophy of Science.” *Perspectives on Science* 10/4: 408–432.
- Stewart, C. Matthew, and Valerie Hardcastle. 2005. “Localization in the Brain and Other Illusions.” In *Cognition and the Brain*, ed. Andrew Brook, and Kathleen Akins, 27–39. Cambridge: Cambridge University Press.
- Strangman, Gary, Joseph P. Culver, John H. Thompson, and David Boas. 2002. “Quantitative Comparison of Simultaneous BOLD fMRI and NIRS Recordings

- during Functional Brain Activation.” *NeuroImage* 17: 719–731.
- Sullivan, Jacqueline. 2009. “The Multiplicity of Experimental Protocols: A Challenge to Reductionist and Non-Reductionist Models of the Unity of Neuroscience.” *Synthese* 167: 511–539.
- Sullivan, Jacqueline. 2015. “Experimentation in Cognitive Neuroscience and Cognitive Neurobiology.” In *Handbook of Neuroethics*, ed. Jens Clausen, and Neil Levy, 31–47. Dordrecht: Springer.
- Taylor, Henry. 2017. “Attention, Psychology and Pluralism.” *The British Journal for the Philosophy of Science* axx030. DOI 10.17863/CAM.8295
- Uttal, William. 2001. *The New Phrenology*. Cambridge: The MIT Press.
- van Orden, Guy, and Kenneth Paap. 1997. “Functional Neuroimages Fail to Discover Pieces of Mind in the Parts of the Brain.” *Philosophy of Science* 64: S85–S94.
- Venturelli, A. Nicolás. 2016. “A Cautionary Contribution to the Philosophy of Explanation in Cognitive Neuroscience.” *Minds and Machines* 26/3: 259–285
- Waters, C. Kenneth. 2007. “The Nature and Context of Exploratory Experimentation: An Introduction to Three Case Studies of Exploratory Research.” *History and Philosophy of the Life Sciences* 29/3: 275–284.
- Weber, Marcel. 2005. *Philosophy of Experimental Biology*. Cambridge: Cambridge University Press.
- Wright, Jessey. 2017. “The Analysis of Data and the Evidential Scope of Neuroimaging Results.” *The British Journal for the Philosophy of Science*. DOI 10.1093/bjps/axx012

- Yarkoni, Tal. 2009. “Big Correlations in Little Studies: Inflated fMRI Correlations Reflect Low Statistical Power. Commentary on Vul et al. (2009).” *Perspectives on Psychological Science* 4/3: 294–298.
- Yarkoni, Tal, Russell Poldrack, David C. van Essen, and Tor Wager. 2010. “Cognitive Neuroscience 2.0: Building a Cumulative Science of Human Brain Function.” *Trends in Cognitive Sciences* 14/11: 489–496.

