Registration Pluralism and the Cartographic Approach to Data Aggregation Across Brains

Zina B. Ward

Abstract

Neuroscience has become increasingly reliant on multi-subject research in addition to studies of unusual single patients. This research has brought with it a challenge: how are data from different human brains to be combined? The dominant strategy for aggregating data across brains is what I call 'the cartographic approach', which involves mapping data from individuals to a spatial template. Here I characterize the cartographic approach and argue that one of its key steps, registration, should be carried out in a way that is sensitive to the target of investigation. Because registration aims to align homologous brain locations, but not all homologous locations can be simultaneously aligned, a multiplicity of registration methods is required to meet the needs of researchers investigating different phenomena. I call this position 'registration pluralism'. Registration pluralism has potential implications for neuroscientific practice, three of which I discuss here. This work shows the importance of reflecting more carefully on data aggregation methods, especially in light of the substantial individual differences that exist between brains.

1 Introduction: Data Aggregation in Neuroscience
2 The Contemporary Cartographic Approach to Aggregation Across Brains
3 Registration Pluralism
4 Homology and the Goal of Registration
5 Organizational Variation and Failures of Simultaneous Alignment
6 The Scope of Registration Pluralism
7 Potential Methodological Implications of Registration Pluralism
7.1 Purpose-sensitive selection of registration methods
7.2 Functional registration
7.3 Standardized preprocessing pipelines
8 Registration Pluralism and the Study of the Brain
9 Conclusion

1 Introduction: Data Aggregation in Neuroscience

Neuropsychology has a long history of drawing conclusions from lone, anomalous cases (think of Phineas Gage or H.M.), but most neuroscientists agree that collecting data from multiple subjects is preferable when it is feasible. The brain is a notoriously noisy organ, and using multiple subjects helps to distinguish the signal from the noise. It also ensures that one's findings are not hostage to the idiosyncrasies of a single brain. But the use of multiple subjects brings with it the challenge of data aggregation: how are the data from different people to be combined and analysed?

One dominant strategy for dealing with data aggregation in neuroscience is what I will call 'the cartographic approach'. On the cartographic approach, cross-brain comparison and aggregation are accomplished by placing whole-brain data from multiple subjects into a common reference frame or onto a template (Toga *et al*. [2006]). This mapping (or 'registering') of individual data into a shared space allows brains to be compared and group-level statistics to be computed. Although alternative aggregation strategies exist,[1] the cartographic approach has been widely adopted since at least the mid-twentieth century. In that time, it has evolved substantially: from the visual inspection of paper atlases, constructed from post mortem examination of stained sections; through the invention of stereotactic spaces and early landmark-based alignment methods, which were designed for neurosurgery but coopted for neuroimaging research in the 1980s; to the construction of digital brain atlases and the proliferation of automated and semi-automated registration methods, which continues to the present (Toga and Mazziotta [2002]; Toga *et al*. [2006]; Evans *et al*. [2012]).

In this paper, I will characterize the present-day cartographic approach and argue against a tempting view about registration, one of its key components. The view I reject, which I call 'registration monism', maintains that all brain data should be registered to spatial templates in the same way. The registration monist takes it to be a problem that different researchers currently use different registration methods and believes that eventually neuroscientists should or will converge on the single best one. I'll argue that this view of the cartographic approach is mistaken. Instead, we ought to embrace 'registration pluralism', which claims that the best way to register data to a brain template depends on the phenomenon under investigation. Registration pluralism asserts the in-principle impossibility of ever finding a single spatial mapping suitable for all neuroscientific purposes. This impossibility is a consequence of the substantial individual differences that exist in the organization of the human brain.

I begin in Section 2 by describing the fundamental components of the cartographic approach. Section 3 introduces registration pluralism and Sections 4 and 5 defend it. I then flesh out its scope in Section 6. Finally, I explore its possible methodological consequences in Section 7 and its broader philosophical significance in Section 8.

2 The Contemporary Cartographic Approach to Aggregation Across Brains

True to its name, the cartographic approach to aggregation involves mapping information about the brains of individuals to a shared spatial reference frame: whole-brain data from different subjects are projected onto a two- or three-dimensional template or into a stereotactic space (Fig. 1).[2] Statistical analysis is then conducted on the aggregate data. Often this involves the use of an

[1] Alternative strategies include the functional localizer approach (Poline *et al*. [2010]), temporal alignment (Zhang *et al*. [2017]), and hyperalignment (Haxby *et al*. [2011]).

[2] The cartographic approach, and the problem of aggregation which it solves, can be found in many medical imaging contexts, not just in neuroscience (Crum *et al*. [2003]). The issues to be discussed here may therefore have analogues in other parts of physiology and biomedical science. How far the analogies extend, including whether registration pluralism applies to organs besides the brain, is an empirical matter (see Sections 5 and 6).

atlas to divide the template brain into distinct regions. To understand the cartographic approach in its contemporary form, it is important to understand these central components.

A stereotactic space is a coordinate system used for specifying locations in the brain relative to internal or external landmarks (Roland and Zilles [1994]; Rahman *et al*. [2009]). Stereotactic spaces may be two- or three-dimensional (Tucholka *et al*. [2012]), and they may be applicable to the entire brain or just a part of it. Each stereotactic space comes with rules about how a brain is to be positioned in the space: it specifies where the origin lies and how the axes are oriented. Once a brain is mapped to a stereotactic space, specific points can be labeled with stereotactic coordinates. A groundbreaking stereotactic coordinate system in neuroscience was introduced by Talairach *et al*. ([1967]). The three-dimensional 'Talairach space' uses the inter-hemispheric fissure and the anterior and posterior commissures, subcortical structures that are relatively invariant across individuals, for orientation.

A brain template is a representation of a brain onto which other brains are mapped. Templates may also be called targets, references, or baseline images (Crum *et al*. [2004]). Some studies use a brain scan from one subject chosen at random to be the image to which all others are normalized. Others use templates that have been constructed by averaging images from multiple subjects. Many stereotactic spaces are associated with templates. Talairach and Tournoux ([1988]), for example, published a template to go along with the Talairach space some twenty years after it was introduced. Like stereotactic spaces, a template can be two- or three-dimensional (Saad and Reynolds [2012]).

An atlas is distinct from a template in that it has labeled parts. Atlases are representations of the brain that partition the volume (if it is three-dimensional) or surface (if it is two-dimensional) into discrete, labeled regions (Gholipour *et al*. [2007]). The regions that atlases pick out may be cytoarchitectural, macroanatomical, functional, histological, or chemoarchitectural. There are a variety of brain atlases in use today: MNI, ICBM, Harvard-Oxford, and Freesurfer, to name a few (Evans *et al*. [2012]). Even though some authors use 'template' and 'atlas' interchangeably (Toga [1998]; Dickie *et al*. [2017]), it is important to distinguish them because the construction of atlases raises a host of issues about how to divide the brain into parts that templates alone do not. Selecting a representative template is a different problem from partitioning a template in a scientifically useful way.

Registration is the process of transforming a target image in order to relate positions in the target to positions in a template or stereotactic space. This paper will deal only with cross-subject registration, not the co-registration of multiple images from the same subject, so I will use the term 'registration' interchangeably with 'normalization'. To register or normalize an image is to determine a 'mapping', a 'warping', or a 'spatial transformation' from the image to the template.[3] A registration method typically consists of three components: a similarity measure, an optimization measure, and a mapping (Crum *et al*. [2004]). The similarity measure provides a way of assessing how well the image matches the template. (For example, a very crude similarity measure for two

[3] In what follows I will primarily discuss registration to templates, but my conclusions apply equally well to registration to reference spaces.

images of the same size would be the sum of the difference in intensity values of every corresponding pair of pixels.) An optimization procedure is used to choose a transformation that maximizes the similarity measure. The transformation that is selected is then applied to the image to register the image to the template.

It is common to distinguish between two broad kinds of registration methods: intensity-based and feature-based approaches.[4] Intensity-based approaches employ a similarity measure that assesses the difference in image intensity between the target and the template. Feature-based approaches, by contrast, represent distinct 'elements in each of the scans to be matched […] includ[ing] functionally important surfaces, curves, and point landmarks. [These] elements are each parameterized and matched with their counterparts in the target scan, and their correspondences guide the volumetric transformation' (Toga [1998], p. 4). In other words, intensity-based approaches select transformations that make the image look visually similar to the template, while feature-based approaches aim to bring specific landmarks into alignment. Intensity-based approaches have gradually been losing ground to feature-based approaches over the last several decades, partly because the range of usable features has expanded (Ashburner [2012]). Early feature-based registration was based on gross macroanatomical landmarks; now, there are feature-based methods that incorporate information about curves, major sulci/gyri, microstructure, and even function (see Section 7.2).

I will only be discussing registration methods that map data onto a two- or three-dimensional template with a standard spatial interpretation. (I'll sometimes call this 'spatial registration'.) There are other kinds of registration, used by non-cartographic approaches, to which my arguments do not apply. For instance, one can align some kinds of neuroscientific and psychological data temporally (Zhang [2017]). Another non-cartographic data aggregation method is hyperalignment, which involves projecting individual data into an abstract, high-dimensional space (Haxby *et al*. [2011]). Both alternative kinds of registration fall outside the scope of my discussion.

Philosophers of neuroscience have done a good deal of work on topics related to brain atlases, such as the problem of identifying the brain's parts or functions (van Orden [1997]; Klein [2012]; Anderson [2014]). They have paid relatively little attention to apparently prior questions about how neuroscientific data is aggregated in the first place.[5] Despite its neglect by philosophers, registration has been the object of intense scientific activity. Over the last several decades, scientists have developed an increasingly diverse and sophisticated set of registration methods. The simplest transformations that one can apply to a brain image include translation (moving the image up/down or left/right), scaling (changing the image's overall size), and rotation (rotating the image around the origin). Early registration techniques used these kinds of simple transformations.

[4] This distinction goes by many names, including 'geometric vs. intensity approaches' (Crum *et al*. [2004]), 'model-based vs. intensity-based approaches' (Toga [1998]), 'label-based vs. non-label-based approaches' (Friston *et al*. [1995]), and 'photometric vs. geometric approaches' (Hellier *et al*. [2003]).

[5] As we will see below, the task of alignment only appears to be prior to the identification of parts in the brain. Registration pluralism implies that selecting an appropriate registration method requires one to pick out certain brain locations as salient, which would seem to require some type of neuroscientific ontology.

The original Talairach method, for example, divides an individual brain image into 12 rectangular regions, each of which is individually scaled and positioned using piecewise affine transformations (Toga [1998]; Chau and McIntosh [2005]). By contemporary standards, this is a very crude procedure. Simple rigid transformations have now been largely replaced by affine and non-linear transformations, which have significantly more degrees of freedom (Crum *et al*. [2004]; Toga and Thompson [2007]). Unlike the Talairach method, which can only reposition and rescale the brain in limited ways, newer registration techniques allow neuroscientists to move, dilate, stretch, scale, and rotate brain images in a highly complex fashion.

3 Registration Pluralism

There is now a wide array of spatial registration methods available to the neuroscientist who adopts the cartographic approach. Some researchers have expressed concern about this variety of methods (van Essen and Dierker [2007]) or implicitly assumed that it is temporary, to be winnowed down over time as we move toward the single best registration method. Such an attitude is arguably implicit, for example, in the literature on the validation and comparison of different registration methods. Many papers that compare multiple methods conclude with an overall recommendation about which method provides 'the optimal alignment' or is 'the best registration method', without obviously indexing these phrases to a specific neuroscientific context (Klein *et al*. [2009]; Robinson *et al*. [2014]). Improvements to the optimization methods used to select transformations, the increasing mathematical complexity of mappings, and the inclusion of more features in cost functions may be taken to suggest that we are getting ever closer to the ideal, universally applicable registration procedure. Let's call the view that there is a single best registration method toward which all current methods strive 'registration monism'.[6]

In what follows, I aim to show that registration monism is mistaken. There can be no universally applicable registration procedure because, given the nature of the brain, different target phenomena require different methods. I intend to defend registration pluralism:

> Registration pluralism: There is more than one appropriate way to register a brain to a spatial template. The best registration method depends on the phenomenon under investigation.

Registration pluralism is an analog of pluralism about atlases, a far more visible and popular view. Neuroscientists frequently acknowledge that there is no single best way of partitioning the cortex, and that different atlases may be appropriate for different purposes (Arslan *et al*. [2018]; Dickie *et al*. [2017]). Bohland *et al*. ([2009]), for example, argue that, 'it is highly unlikely that the neuroscience community will, or even should, adopt a single scheme for partitioning the brain or for labeling its pieces' because 'the motivations underlying the construction of one atlas can be different from another' (p. 11). Many express this ecumenical attitude about atlases (Toga and

[6] I will try to show in Section 7 that a monistic attitude toward registration is at least implicit in several neuroscientific practices. I hope that this will satisfy the skeptical reader that, even if there are not many scientists who explicitly endorse registration monism, it is still a worthwhile target for criticism.

Thompson [2001]; Brett *et al*. [2002]; Shattuck *et al*. [2008]; Amunts *et al*. [2014]). Registration pluralism is in the same spirit. Different registration methods, as well as different partitioning schemes, ought to be applied in different scientific contexts.

The plausibility of registration pluralism can be illustrated with a simple example. Consider two hypothetical research projects. The first concerns the relationship between sleep/wake cycles and functional activation in areas of the brain that specialize in language processing. The research question is: how does the BOLD signal in language processing regions change over the course of the day? The second research project examines the effects of mercury on the brain (Azevedo *et al*. [2012]). This researcher asks: do people who have had more mercury exposure have more GABA-A receptor activity in cortex surrounding the calcarine sulcus? Let's imagine that both researchers are committed to using the cartographic approach to aggregate their multi-subject data. My contention is that the two researchers should register their data in different ways. For the first project, the researcher should try to align brain regions thought to be involved in the same language-related functions. For the second, the researcher should try to align the calcarine sulcus and its surrounding brain tissue. Given that brains are slightly different from one another, a transformation that aligns language processing regions will not perfectly align points surrounding the calcarine sulcus. Hence, different registration methods are required in the two cases. Diversity in our neuroscientific projects necessitates registration pluralism.

The following two sections will generalize the reasoning in this simple example to provide an argument for registration pluralism. Section 4 will characterize the goal of registration and Section 5 will show that the goal cannot be achieved with a single method.

4 Homology and the Goal of Registration

The first step of the argument for registration pluralism requires understanding what registration aims to do. Fortunately, neuroscientists who employ the cartographic approach are quite explicit on this point. Nearly every author who writes about registration claims that the goal of registration is 'to maximize the genuine homology of points that are brought into correspondence by the transformations' (Mazziotta *et al*. [2001], p. 1301). Or, to put it a different way, 'the objective is to warp the images such that homologous regions of different brains are moved as close together as possible' (Ashburner *et al*. [1997], pp. 350–1). Registration to a template succeeds to the extent that 'homologous cortical regions in different subjects have been brought into register by the registration transform' (Toga and Thompson [2001], p. 4). I see no reason to doubt this consensus, so I take it that the aim of registration is to align homologous locations across brains.

The real challenge is understanding what this means. It may be tempting to think that the homology concept at work here is simply a homology concept from biology applied to cognitive science. I believe this is mistaken: neuroscientists discussing the cartographic approach do not use

'homology' in the same way as biologists or philosophers of biology (Brigandt [2002]; Wagner [1994], [2014]).[7]

The concept of homology originated with Richard Owen ([1843]), who famously defined a homologue as 'the same organ in different animals under every variety of form and function'. Owen was a comparative anatomist, interested in how animals from different species seem to instantiate the same 'archetypes' (Panchen [1994]). Although the concept of homology has changed since Owen, biologists continue to understand homology primarily as a cross-species notion. This is the first reason that a biological interpretation of 'homology' in the context of brain registration is inappropriate. Registration is used to align the brains of individuals from the same species. When neuroscientists talk about points in the brain being homologous across people, they are not using 'homology' to talk about inter-specific relationships, as a biologist would. Even the concept of serial homology in biology bears no resemblance to the neuroscientist's concept, since it involves the repetition of a part within one and the same animal. Hence, there is no inter-individual but intra-specific sense of 'homologous' that neuroscientists could be importing directly from biology.

One might think it would be perfectly natural to apply biologists' homology concept(s) to within-species relationships even if biologists do not. After all, it seems possible to talk about homologies across different dog breeds, even though all dogs belong to the same species. Even granting this, there are other indications that neuroscientists are using 'homology' differently from biologists. In contemporary biology and philosophy of biology, a distinction is frequently made between two homology concepts: genealogical homology and developmental homology (Wagner [1989]; Brigandt [2002]; Ramsey and Peterson [2012]). Homologous parts in the genealogical sense share a phylogenetic origin, while homologous parts in the developmental sense are subject to the same developmental constraints and underwritten by common ontogenic mechanisms (Wagner [1994]). Neither a genealogical nor a developmental homology concept is reflected in the neuroscientific usage. Very few of the neuroscientists who adopt the cartographic approach are investigating the evolution of the nervous system, so it would be strange to interpret their use of the phrase 'homologous parts' in terms of phylogeny. When neuroscientists use the cartographic approach, they do not discuss the evolutionary history of the places in the brain they are attempting to align, nor do they try to establish that the places share a common phylogenetic origin. The genealogical homology concept is therefore ill suited to the context of registration. A developmental homology concept can be ruled out for similar reasons. Neuroscientists do not need to know anything about the development of the brain locations they intend to co-register. Their aim during registration is to align present structure or function across individuals. They make no effort to show that the co-registered locations share a developmental trajectory or are subject to the same constraints, as would be expected if they were deploying a developmental homology concept.

[7] This is not to say that neuroscientists never use 'homology' this way (Liebeskind *et al*. [2016]). It also may be the case that neuroscientists' usage of 'homologous' is historically related to biologists' usage.

For these reasons, the phrase 'homologous locations in the brain' should not be understood by appeal to biologists' homology concept(s). Instead, neuroscientists should be read simply as using homology as a synonym for sameness. Neuroscientists call places in different brains 'homologous' when they are the same. Importantly, however, sameness of brain locations across people cannot be specified tout court because there are many different kinds of sameness that one might be interested in (Goodman [1972]). In the hypothetical case above, one researcher cares about points that have the same language-processing functions while the other cares about points in the same position relative to the calcarine sulcus. Sameness of functional capacity is distinct from sameness of sulcus-centered positioning. Because sameness is always sameness-in-some-respect, which type of sameness is relevant in any particular scientific context is determined by the target of investigation. Sameness of location in the brain, and therefore homology of brain locations, is purpose-relative.[8]

This should not surprise us since it is just an instance of a general phenomenon. Sameness of location within wholes is purpose-relative whenever the wholes are qualitatively different from one another. Consider a simple inanimate example: imagine there are two similar houses, and a location in one house has been singled out as a point of interest. We might ask the question: where is the same location in the other house? The answer clearly depends on what we are interested in. There are many places in the second house that one could identify as the 'same' as the place in the first: the location that is the same absolute distance from the front door; the location that is the same distance, proportionally, from the two ends of the house; the location that contains the same piece of furniture; and so on. Sameness of location in houses is purpose-relative. Brains are no different from houses in this respect.

In line with this, several neuroscientists acknowledge that different types of homologies in the brain exist simultaneously. Mazziotta *et al*. ([2001]) explain that 'various criteria can be used to define homology' (p. 1301). They distinguish between 'anatomical [and] functional homologues' and then proceed to make finer grained distinctions within these categories, arguing that 'homologies based on function and cytoarchitectonics are more fundamental to neuroscience […] than homologies based on sulcal and gyral anatomy' (Mazziotta *et al*. [2001], p. 1316, p. 1301). Likewise, Uylings *et al*. ([2005]) claim that a 'critical issue' for any study is 'defining the criterion of correspondence, i.e. homology' (p. 424). This implies that different criteria are available. At least some neuroscientists agree, then, that homology in the brain is purpose-relative. In what follows, I use the term 'homologous' to mean sameness-in-some-respect, in keeping with this standard scientific usage.

[8] Some biologists have claimed that homology is actually 'context dependent' in just this way (Abouheif [1997]). Wagner ([1994]) argues that 'homology is a scientific conceptualization of th[e] perception of 'sameness'' (p. 274). He thinks the reason homology has been such a tricky concept is that there are different 'aspects' of sameness that are prioritized by different biologists: 'the same structural organization, the same developmental origin, the same developmental constraints, the same (genetic) information, and common phylogenetic origin' (Wagner [1994], p. 274). If Wagner is right, the analogy between biological and neuroscientific uses of 'homology' need not be rejected after all: both are purpose-relative.

5 Organizational Variation and Failures of Simultaneous Alignment

We have now seen that the goal of registration is to align locations across brains that are the same. Furthermore, there are different ways for brain locations to count as 'the same'. By itself, this does not establish registration pluralism. To understand why, let's return to the house analogy. Imagine that that an architect is trying to register two houses' blueprints to a single template. If the houses were built by the same developer in a cookie-cutter suburban neighborhood, they may be essentially identical. In this situation, it may be possible to register the blueprints to the template in a way that preserves all the homologous relationships the architect cares about. A single transformation will be sufficient to align points that are the same distance from the front door, locations with the same furniture, walls with the same load-bearing capacity, rooms with the same practical functions, and so on. Registration, for the architect, need not be sensitive to the feature of interest.

Human brains are too different from one another for this to work in neuroscience. Spatial alignment of one kind of homologous brain location will not align homologous locations of all other types. A registration method that brings the sulci and gyri of two different brains into alignment, for instance, will not perfectly align cytoarchitectural regions, and vice versa. This is because the brain's organization is variable across people: different types of regions do not stand in constant spatial relationships to one another. For example, the position of one person's occipitotemporal sulcus relative to his Wernicke's area may not be the same as the relative position of another person's occipitotemporal sulcus and Wernicke's area. Consequently, by aligning two subjects' occipitotemporal sulci, you may not succeed in aligning their Wernicke's areas, and vice versa. Organizational variation prevents there being a single way of spatially aligning all homologous brain regions at once.[9]

Evidence for this claim comes from studies showing individual differences in the relative positioning of different types of brain regions. First, it is well known that macroanatomical features like sulci and gyri are variably positioned relative to cytoarchitectural boundaries (Uylings *et al*. [2005]; Amunts *et al*. [2007]). Amunts *et al*. ([1999]) demonstrate this variation in their classic examination of Brodmann's areas 44 and 45. They use a computerized technique to identify cytoarchitectural borders on stained brain sections and then compare the cytoarchitectural regions identified with macroanatomical landmarks. The authors find significant inter-individual differences in the location of cytoarchitectural regions relative to macroanatomical features: 'one and the same cytoarchitectural border was located in a sulcal fundus in some [individuals'] hemispheres but on one or the other wall of the sulcus or at the top of the gyrus in others' (Amunts *et al*. [1999], p. 335). Scheperjans *et al*. ([2008]) reach the same conclusion about superior parietal

[9] As some authors have noted (Klein *et al*. [2009]), a neuroscientist using the cartographic approach has to assume that the locations he is attempting to align are present in all brains. This assumption might not always be justified. For example, Nieto-Castañón and Fedorenko ([2012]) argue that it is unlikely that exact correspondences exist between ocular dominance columns in V1 of different individuals. They claim that it would be a mistake to try to co-register such fine-grained functional regions. When is it safe to assume that the same location exists across brains, and that it is therefore appropriate to apply the cartographic approach? This is an important question that deserves more attention.

cortex. Using a similar technique, they find that 'the locations of [cytoarchitectonic] borders are not reliably associated with macroanatomical landmarks' (Scheperjans *et al*. [2008], p. 2152).

There is also considerable evidence that macroanatomy and function are not predictably related to one another. For instance, Watson *et al*. ([1993]) examine intersubject variability in the location of V5, a visual motor area, in relation to sulcal and gyral patterns. They define V5 functionally by comparing PET data collected while subjects saw a moving or stationary checkerboard. They find that the position of V5 can 'vary by as much as 27 mm in the left hemisphere and 18 mm in the right' relative to macroanatomical features (Watson *et al*. [1993], p. 79). Others have stressed that although some functional areas, like the frontal eye fields, are strongly related to macroanatomy, others, like the fusiform face area, are not (Frost and Goebel [2012]). The divergences are emphasized by critics of the cartographic approach, who argue that registration methods based on anatomical features often fail to align functional regions of interest (Fedorenko and Kanwisher [2009], Nieto-Castañón and Fedorenko [2012]).

What of the relationship between cytoarchitectural and functional regions? It is usually thought that functional differences between brain regions are underwritten by cytoarchitectural differences, and hence that there is a close correspondence between cytoarchitectural and functional areas. If this is correct, cytoarchitectural regions do generally stand in constant spatial relationships with functional regions; indeed, they are coextensive. However, not all functional divisions are marked by changes in cytoarchitecture. Weiner *et al*. ([2017]) examine the relationship between cytoarchitectural regions and functional regions in human ventral temporal cortex (VTC). As predicted, they observe that face- and place-selective regions in VTC have different cytoarchitectural properties. They also find, however, that there is a 'many-to-one mapping' between functional regions of interest (fROIs) and cytoarchitectural regions (cROIs), with several fROIs contained within a single cROI (Weiner *et al*. [2017], p. 155).

Hence, there is considerable empirical evidence that macroanatomical, cytoarchitectural, and functional brain regions do not stand in constant spatial relationships with one another. This indicates that organizational variation in the brain prevents the simultaneous spatial alignment of every type of homologous region. Given that registration aims to align homologous regions, different registration methods should be used in different scientific contexts. Registration pluralism follows. Note that this argument is not based on the shortcomings of our current methods. That there is no way to spatially align all homologous brain regions at once is a consequence of the nature of the brain, not the limits of our current tools for studying it. As such, registration pluralism captures a permanent feature of the cartographic approach rather than a temporary obstacle.

The importance of purpose-relative registration is supported by work aimed at comparing different registration methods. Crivello *et al*. ([2002]), for example, register the same fMRI and PET data to a Human Brain Atlas template using four different registration methods and then assess the methods' success using several different metrics. They calculate the degree of spatial overlap between the template and the normalized individual MRI volumes for grey matter, white matter, and cerebrospinal fluid (CSF). Two of the normalization procedures Crivello *et al*. use are

the procedure implemented in the 1996 Statistical Parametric Mapping software (SPM), a popular software package for neuroimagers, and a multi-grid technique based on Navier-Lamé continuum mechanics theory (FMG). The details of the SPM and FMG methods do not concern us. What matters is that they excelled in different respects: the FMG method was the best of the four methods at aligning anatomical landmarks, while the SPM method was best at CSF alignment (Crivello *et al*. [2002], p. 237). None of the methods dominated the others.

Although the authors do not discuss the implications of this result (indeed, it is not even highlighted as an important finding), it lends support to registration pluralism. It suggests that if a researcher is interested in using the cartographic approach to answer a question about CSF, she ought to use SPM; if the alignment of macroanatomy is more important given her project, she ought to use FMG. Surprisingly, despite the obvious way in which their findings support purpose-relative registration, Crivello *et al*. do not endorse registration pluralism. At the end of their paper, they argue that FMG is the 'normalization procedure providing the highest degree of accuracy' and recommend that researchers adopt it (Crivello *et al*. [2002], p. 248). They seem to think that different registration techniques are simply better or worse at providing 'the most accurate brain', ignoring the fact that accuracy is purpose-relative (Crivello *et al*. [2002], p. 248).

Papers like this one also show that the errors introduced by failing to register one's data in a context-appropriate way can be large enough to scuttle a statistical analysis. One might have thought that neuroimaging is so noisy that minor failures of alignment produced by use of a general-purpose registration method would not make a difference to one's ability to obtain statistically significant results. But the literature on the validation of registration methods suggests otherwise. It is not uncommon for researchers to compare different methods using statistical tests on data that have been registered in different ways. For instance, Nenning *et al*. ([2017]) perform a group-level activation analysis of task-based fMRI data that have been registered with two different methods. They find that one method results in central regions with higher *t*-values, meaning that it allows more sensitive region detection than the other. Such results show that the use of context-appropriate registration methods can increase statistical power and thus appreciably improve experimental outcomes, despite the noisiness of neuroimaging.

There are some neuroscientists who, unlike Crivello *et al*. ([2002]), embrace registration pluralism by acknowledging that different projects demand different mappings (Friston *et al*. [1995]; Evans *et al*. [2012]). Hellier *et al*. ([2003]) claim that 'the 'ideal' transformation surely depends on the application' (p. 1120). Crum *et al*. ([2003]) agree that 'the kind of correspondence, the manner of achieving it, and the acceptable accuracy are application dependent […] the scientific question defines the kinds of correspondence that *should* be sought' (p. 1434). And Dubois and Adolphs ([2016]) recommend 'that investigators try more than one approach to alignment, and report all of them, so we can see which might work best for which kinds of questions' (p. 427). As we will see in Section 7, however, even researchers who endorse registration pluralism in the abstract may not have fully considered its potential methodological implications.

6 The Scope of Registration Pluralism

I have argued that the goal of registration in the cartographic approach is to align homologous locations across brains, but that individual differences make it impossible to achieve this goal with a single registration method. When I offered empirical evidence of individual differences, I focused on macroanatomical, cytoarchitectural, and functional regions because they provide paradigmatic examples of the kinds of homologous locations to be aligned in different research contexts. Registration pluralism would be of limited interest, however, if it held only across these broad categories. A critic might grant that researchers interested in macroanatomy need to use different registration methods than those interested in function, but insist that all macroanatomists should use the same method. The same could be said of all functional researchers and all cytoarchitectural researchers. None of the evidence discussed above entails that registration has to be purpose-sensitive even within broad research areas.

This objection hints at a broader question about the scope of my view. I have so far spoken of registration pluralism as an all-or-nothing thesis asserting that, given the nature of the human brain, there cannot be single registration method appropriate for all neuroscientific projects. So conceived, registration pluralism is true. However, such a thesis is of limited utility by itself. What we really want to know is which projects require different registration methods and which do not. After all, the claim that no method will suffice for all purposes doesn't imply that there are no two purposes for which a single method will suffice. So, just how pluralist should we be about registration? A pithy answer is: as pluralist as the evidence requires. The more variation there is in the positioning of different types of brain locations relative to one another, the greater the variety of registration methods that are needed.

There is in fact empirical evidence of organizational variability even within the three major brain modalities, implying that the scope of registration pluralism is wider than the critic above suggests. Much of the research cited in Section 5 in support of the idea that cytoarchitectural and macroanatomical regions do not stand in constant spatial relationships also shows that different cytoarchitectural regions are variably positioned relative to one another. In their microstructural study, Scheperjans *et al*. ([2008]) find that which cytoarchitectural regions border one another is different across brains. For example, area hIP3 only borders hIP1 in half of the hemispheres they examine. They conclude that 'a considerable number of cytoarchitectonic borders are not present in every brain' (Scheperjans *et al*. [2008], p. 2152). It is also well-known that the three-dimensional size and shape of cytoarchitectural regions differ between people. This suggests that a scientist needing to align, say, Brodmann's area (BA) 17 across subjects may require a different registration method than a scientist needing to align BA44. Moreover, there are sub-regions within cytoarchitectural areas whose relative positioning varies. Amunts *et al*. ([1999]) describe substantial variability within BA44 and BA45, including lamina whose distribution patterns differ substantially between people. This evidence suggests that different registration methods are needed for research concerning different cytoarchitectural features.

A similar conclusion is arguably true of functional research as well: different kinds of locations count as functionally homologous, and the relative positioning of different functional

homologues is such that they cannot all be simultaneously aligned. The reasons for this are somewhat more theoretical. First, given the assumedly tight relationship between cytoarchitecture and function, if different types of cytoarchitectural regions cannot be simultaneously aligned, as suggested above, the same is likely true of different functional regions. Second, recent theorizing about neural reuse supports the idea that the functional organization of the brain precludes the use of a single registration method for all functional purposes. The neural reuse hypothesis states that 'individual neural elements (at multiple spatial scales) are used and reused for multiple cognitive and behavioral ends' (Anderson [2016], p. 1). Anderson ([2014], [2016]) argues that each brain region has a functional 'fingerprint' or 'profile' but can be recruited for a diverse array of tasks. The brain is constantly re-organizing and functional partnerships between brain regions are not fixed. Which regions are recruited for a particular task depends on current activation patterns and other functional demands. Anderson's dynamic, fragmented view of the brain's functional organization supports the idea that locations that are functionally 'the same' may be highly variable, both across time and across people. Such neural flexibility and complexity make it impossible to align all functionally homologous points at once.

All of this suggests a wide scope for registration pluralism. The critic imagined above conceded that different neuroscientific contexts require different registration methods, but claimed that contexts should be individuated quite coarsely: we need only differentiate between a 'cytoarchitectural context', a 'macroanatomical context', and a 'functional context', because within each of these, a single registration method suffices. The evidential and theoretical considerations raised here, however, suggest that contexts need to be considerably more fine-grained. I have not specified the level of grain precisely, partly because to do so would be premature. We understand relatively little about individual differences in the spatial positioning of the brain's different features. Our ideas about which projects require different registration methods and which don't must therefore be continually refined in light of new empirical findings.[10]

7 Potential Methodological Implications of Registration Pluralism

Defending registration pluralism in principle is easier than figuring out what it means for neuroscientists in practice. The challenge is not just to delimit its scope, but also to determine the broader impacts of the methodology it seems to recommend. The use of different registration procedures in different scientific contexts might well increase alignment accuracy, but it could also have adverse effects on other scientific desiderata. Deciding whether such trade-offs are worth making requires data- and simulation-driven assessments of various methodological approaches along multiple dimensions. In this section, I'll gesture at three potential methodological consequences of registration pluralism that are deserving of further study.

[10] This section has focused on questions about the scope of registration pluralism with respect to the granularity of neuroscientific contexts. But scope questions are also spatial: it is possible that registration pluralism holds with respect to some brain regions but not others (namely, those that are more variable across individuals).

7.1 Purpose-sensitive selection of registration methods

First, registration pluralism seems to imply that researchers who adopt the cartographic approach should select a registration method in a purpose-sensitive manner. In practice, this is rare in studies with neurotypical individuals.[11] Neuroscientists usually do not justify their choice of registration method in print, and when they do, it is in general terms that do not engage with the specific features of their research question. They may, for example, explain that the registration method was chosen because it came as a default in a software package or because it was not computationally demanding. Registration pluralism suggests that neuroscientists who adopt the cartographic approach should explicitly identify the locations they are hoping to align and then select a registration method likely to align those homologous locations (Crum *et al*. [2003]). There are several worries one might have about this recommendation: one could object that choosing a registration method in a purpose-sensitive manner is not practically feasible; that it presents an obstacle to comparability across studies; or that it could be a source of bias. While all three are serious concerns, I believe they do not decisively undermine purpose-sensitive registration.

The first objection is that it is too much to ask of researchers that they tailor registration to their research question. Neuroscientists, on this view, have too little information about the performance of different methods to make a purpose-sensitive choice. Luckily, this is not the situation that neuroscientists find themselves in. Researchers can choose a registration method by consulting the rapidly expanding literature on the validation of registration methods, which I discussed briefly in Sections 3 and 5 (Woods *et al*. [1998]; Crivello *et al*. [2002]; Hellier *et al*. [2003]; Crum *et al*. [2004]; Ng *et al*. [2009]; Klein *et al*. [2009]; Conroy *et al*. [2013]; Robinson *et al*. [2014]). As many authors have noted, it is usually impossible to directly assess how well a registration method aligns homologous regions because there is no 'gold standard' against which to compare (Woods *et al*. [1998]; Brett *et al*. [2002]; Gholipour *et al*. [2007]). It is, however, possible to use indirect measures of evaluation. Many of these metrics measure some dimension of accurate alignment (Gholipour *et al*. [2007]). Klein *et al*. ([2009]), in one of the most comprehensive validation efforts, compare fourteen registration methods along eight different dimensions. Since they are interested in anatomical alignment, they use manually labeled structural images as a 'silver standard' for comparison. For each of the fourteen methods, they compare the registered images of individual subjects with the manually labeled images. They measure volume overlap agreement (three different measures), volume overlap error (two measures), surface overlap agreement (one measure), volume similarity (one measure), and distance error (one measure) between the source and target images. The data show that the fourteen registration methods perform differently on these eight metrics.

We ought to think of such papers as providing information for choosing registration methods in a purpose-sensitive way rather than identifying the all-around best method. The variety of evaluation metrics on offer permit an individual neuroscientist to select the method that does

[11] There is more discussion of registration in research on aging and non-neurotypical populations (Ganzetti *et al*. [2018]).

the best on the metrics that are most relevant to his project. Which locations he needs to align will determine which metrics are important. At least some of the authors working on validation do seem to think about their findings in this way. Conroy *et al*. ([2013]), for example, who propose a new registration algorithm and compare it to two alternatives, claim that their method is especially good at aligning prefrontal regions (which are particularly difficult to co-register), and should therefore be used in studies of social cognition.

Second, one might worry that the use of different registration methods by different researchers presents an obstacle to comparing results across studies. This concern is present in the literature. Van Essen and Dierker ([2007]) claim that 'unintended biases may be introduced when comparing datasets registered by different algorithms to different templates' (p. 1052). They argue that, to provide 'apples-to-apples comparisons', researchers should use simple linear registration methods (Van Essen and Dierker [2007], p.1052). Brett *et al*. ([2002]) similarly claim that, 'if we have used a different template or a different normalization method, then […] meta-analysis might have low spatial resolution and power' (p. 248). These are reasonable worries, especially concerning templates. There is a trade-off between selecting a template that fits the population at hand and choosing one that permits easy generalization and cross-study comparison (Evans *et al*. [2012]). However, the problem is less acute in the case of registration. When one conducts a meta-analysis, the data one is compiling typically concern the same phenomenon. Since the phenomenon of interest determines which registration method is appropriate to apply, the data being aggregated for meta-analysis usually should not have been registered in wildly different ways.

Third, one might worry that, given how little we know about the neural basis of many cognitive processes, adopting purpose-sensitive registration will introduce bias into research. On this line of thinking, making registration decisions based on a partial or incorrect understanding of the phenomenon being studied will cause systematic errors of alignment, biasing the results. Using a one-size-fits-all registration method also leads to alignment errors, so the objection goes, but at least they are theory-free or random. Better to introduce random noise than theory-driven bias. I believe this concern, too, is overblown. We are not as ignorant of the brain's functioning as is suggested. To select a purpose-sensitive registration method, we only need some idea of the areas that are homologous given the phenomenon under investigation; we do not need to know exactly how cognitive capacities are realized in the brain. Moreover, it is not the case that the misalignments that occur under a system of purpose-sensitive registration are directional while the misalignments resulting from the use of a single, general-purpose method are random. When everyone uses the same registration method, researchers make similar alignment errors, leading to systematic biases. The status quo of general-purpose registration fares no better than purpose-sensitive registration from the perspective of bias.[12]

[12] An anonymous reviewer raised a related concern: when we apply the cartographic approach in neuroimaging, we want to be able to identify activity we were not expecting to find. Purpose-sensitive registration, with its emphasis on aligning locations we already know to be involved in the function under study, might make this less likely. This is a reasonable concern. However, it is an open empirical question whether the use of a general-purpose registration method will make it more likely that unexpected activation will be uncovered than a method known to excel at aligning at least some of the implicated areas.

Hence, there is reason to think that registration methods should be chosen in a way that is sensitive to the phenomenon under investigation, despite legitimate concerns about feasibility, comparability, and bias. Entertaining this implication of registration pluralism opens the door to several other possible methodological consequences.

7.2 Functional registration

Functional registration is a type of feature-based registration that uses functional rather than structural features to select a mapping between an individual brain image and a template (Sabuncu *et al*. [2010]; Conroy *et al*. [2013]; Nenning *et al*. [2017]). Most functional registration methods use fMRI data for alignment. In those that rely on task-based fMRI, subjects are presented with a specific stimulus or task while functional data are collected. Researchers then find a transformation that aligns the functional signals from different subjects. Functional registration is a relatively new tool that has generated substantial interest among neuroscientists. Researchers investigating brain function no longer need to register brains to a template using structural information and hope that structure and function correlate. Instead, functional data can directly drive registration. (Though the data driving registration must not be the data to be analysed, on pain of circularity [Sabuncu *et al*. [2010], p. 139]).[13]

Registration pluralism warns us to watch out for a potential complication: it may be a mistake to think that a single task-based functional registration method can serve all functional purposes. I suggested above that there might be different kinds of functional regions that cannot be simultaneously aligned. If so, different functional registration methods could be needed in different scientific contexts. However, most current methods are intended to be generally applicable. In pioneering work on functional registration, researchers asked subjects to watch a full-length action movie, *Raiders of the Lost Ark*, while they collected fMRI data to be used for alignment (Sabuncu *et al*. [2010]; Haxby *et al*. [2011]; Conroy *et al*. [2013]). The researchers explained that they chose a movie, a 'complex and dynamic natural stimulus', because it sampled a 'diverse variety of representational states' (Haxby *et al*. [2011], p. 411) and because 'neural activity during a movie viewing is synchronized across subjects in a large percentage of the cerebral cortex' (Sabuncu *et al*. [2010], p. 131).

The discussion in Section 6 suggests that trying to find a generic stimulus to be used in a universal task-based functional registration method may be misguided. Given the variable and dynamic functional organization of the brain, a registration algorithm that aligns brain regions involved in passively watching a movie may not align the neural substrates of, say, social cognition. If a researcher is ultimately interested in analyzing brain activity during social interaction, then, she may be best off not using functional data collected during a movie-watching

[13] Some functional registration methods begin by aligning images with a simple anatomical registration procedure before fine-tuning the alignment with functional information (Sabuncu *et al*. [2010]; Conroy *et al*. [2013]). Others are explicitly 'multi-modal': they use functional and structural information simultaneously (Robinson *et al*. [2014]). It is interesting to note that the registration methods called 'functional' typically rely on other types of information as well.

task as the basis for registration. Instead, registration pluralism seems to suggest that functional registration should be purpose-sensitive. Researchers ought to consider aligning brains using fMRI data collected while subjects are performing a task that is related to the domain or phenomenon of interest. This would generate functional data that could bring the areas that matter into alignment. The researcher interested in social cognition, for example, could base her registration method on functional data collected while subjects engage in a social task to ensure that areas essential to social interaction are aligned.[14] There is at least one example in the literature of (something resembling[15]) purpose-sensitive functional registration. Langs *et al*. ([2010]) aim to identify and ultimately align brain regions involved in language-processing across tumor patients. They propose a registration method based on fMRI data collected while subjects are engaged in antonym generation, a language task. Hence, they ask subjects to perform a task that produces activation in the areas they need to align and then use the data collected for registration. It may be fruitful for other researchers to follow their example.

Alternatively, one might argue we should turn away from task-based functional registration altogether. Some authors have claimed that task-based methods rely on the implausible assumption that different people's brains are performing the same functions at precisely the same times (Jiang *et al*. [2013]). They argue that we ought instead to use resting-state fMRI to build a functional connectivity profile for each subject and then find transformations that bring the individual connectivity patterns into alignment (Jiang *et al*. [2013]; Zhou *et al*. [2017]; Nenning *et al*. [2017]; Chen *et al*. [2017]). I am somewhat skeptical of these methods, partly because I share others' doubts about what resting-state fMRI really tells us (Buckner *et al*. [2013]; McCaffrey and Danks [forthcoming]), and partly because functional connectivity-based methods are meant to be general-purpose. However, such qualitative considerations are far from decisive. A resolution awaits systematic quantitative comparison of connectivity-based methods with purpose-sensitive task-based methods of the kind I have proposed.

7.3 Standardized preprocessing pipelines

Finally, registration pluralism may undermine projects aimed at standardizing the preprocessing of neuroscientific data. There is currently concern among neuroimagers about the lack of uniformity in early data processing. They worry that preprocessing is usually ad hoc, making neuroimaging less replicable (Esteban *et al*. [2019]) or generalizable (Gabard-Durnam *et al*. [2018]), and forcing individual labs to 'repeatedly reinvent the wheel' (Freeman [2015], p. 156). Such concerns have led to efforts to construct standardized preprocessing protocols for different

[14] Some of the researchers who use a movie stimulus for functional registration recognize that it may not be appropriate for all scientific contexts. Sabuncu *et al*. ([2010]) explain that 'it is possible that function-based normalization based on neural activity evoked by more controlled experiments could be more effective for a specific functional region […] The key question is whether a single, optimal warp exists for the cerebral cortex or for sectors of the cerebral cortex – or will overlapping topographic maps for different functions be aligned optimally by different warps' (p. 138). They are betting on the former and I am betting on the latter.

[15] The example is, strictly speaking, not an instance of the cartographic approach because Langs *et al*.'s ([2010]) technique does not use a standard spatial reference frame.

imaging modalities (Bigdely-Shalmo *et al*. [2015]; Gabard-Durnam *et al*. [2018]). One example is fMRIPrep, a preprocessing workflow recently introduced by Esteban *et al*. ([2019]) to provide 'robust and reproducible preprocessing' for data from task-based or resting state fMRI (p. 111). fMRIPrep involves a spatial normalization step in which individual data are registered to an ICBM template using the ANTs registration procedure. Esteban *et al*. claim that fMRIPrep is 'analysis-agnostic' in the sense that it supports a wide range of analysis and works across many kinds of input datasets (Esteban *et al*. [2019], p. 112).

Registration pluralism suggests that any preprocessing pipeline that includes spatial registration is not analysis-agnostic. Which registration method is appropriate depends on the phenomenon of interest, and hence on the analysis to be performed. Uniformity of preprocessing therefore comes at a cost: sometimes different projects really do call for particularized preprocessing steps. Some researchers recognize this. Freeman [2015], for instance, stresses the importance of balancing 'standardization and scalability with […] flexibility and interactivity' (p. 157). Registration pluralism seems to imply that registration should be highlighted as a locus of deliberate choice in standardized workflows. Researchers using tools like fMRIPrep ought to be encouraged to think about which registration method suits their purposes best.

8 Registration Pluralism and the Study of the Brain

The previous section suggested that registration pluralism may have implications for how scientists select registration procedures, what tasks are used in functional registration, and which parts of preprocessing can be standardized. These issues, I claimed, are deserving of more systematic analysis by neuroscientific methodologists. Zooming out from such concrete methodological questions, we can also inquire into the broader significance of registration pluralism and the empirical evidence of variability on which it is based. What does registration pluralism mean for neuroscience in general?

First, registration pluralism represents an obstacle to the transfer of evidence across contexts. Data that have been collected and registered to a template in one way, with one purpose in mind, may need to be re-registered in order to apply to a new research question. This is because the initial registration may have aligned a different set of features than those requiring alignment for the new project. This is resonant with current philosophical work discussing the difficulty of sharing scientific data across epistemic contexts (Leonelli [2016]; Boyd [2018]). To use Boyd's ([2018]) terminology, empirical results may be maladapted to a scientific theory, meaning that they cannot serve as a constraint on that theory because of how data were collected or processed. However, sometimes data can be repurposed and become well adapted to theories to which they were initially maladapted. According to Boyd, this repurposing is accomplished (in part) using 'workflow metadata', that is, auxiliary information about the data processing that had been performed. Information about the registration procedure initially used to aggregate neuroscientific data is a kind of workflow metadata that is important to preserve and deploy when we want to repurpose neuroscientific results. Just like in the sciences that Leonelli ([2016]) and Boyd ([2018])

discuss, there is danger in transferring neuroscientific data from one context to another without taking account of how they were produced.

The argument that supports registration pluralism also suggests that there is a type of theory-ladenness in data aggregation that is often overlooked by philosophers of science. One might have thought that gathering all of one's data in the same place once it has been collected is relatively straightforward. But in neuroscience, this apparently simple task can rely on prior information about localization, information that counts as 'theory' on a minimal understanding of that term. When we deploy the cartographic approach, best practice requires us to know something about which locations matter given our explanatory target, or so I have argued. Since our goal is accurate alignment, the conditions of success for aggregation depend on the target of investigation. Aggregation succeeds when the locations we care about are aligned, but to accomplish this, it is helpful to have background neuroscientific theory about which locations those are. As I pointed out in Section 7.1, this doesn't require an extensive understanding of the brain's workings, but it does make data aggregation using the cartographic approach somewhat theory-laden.

The theory-ladenness of the cartographic approach strengthens an existing philosophical argument about the epistemic status of neuroimages. Pushing back against the hype surrounding neuroimaging, several authors have claimed that neuroimages are quite unlike standard photographs. Klein ([2010]) highlights two primary dissimilarities: unlike ordinary pictures, neuroimages are laden with theoretical assumptions and present the results of statistical tests rather than raw data. Klein explains that neuroimages cannot be interpreted without understanding the experimental design that produced the data and the specific tasks that the subjects performed. Roskies ([2007]) calls this property 'belief-opaqueness': neuroimages are belief-opaque because 'the information needed for the[ir] interpretation is not present in the resultant image' (p. 868). Registration pluralism reinforces these claims. Neuroimages depicting aggregate results are theory-laden in an additional way not discussed by Klein ([2010]) or Roskies ([2007]), since aggregation itself is another substantive and theory-laden data processing step involved in visualizing neuroscientific data. Moreover, one cannot tell from looking at an aggregate neuroimage how the data have been registered to the template. This is yet another way in which neuroimages are belief-opaque.

An interesting philosophical question is whether similar conclusions hold in other domains. Aggregation of data would seem to require additional input whenever it isn't obvious which data points count as 'the same'. One (but by no means the only) source of such ambiguity, as I have shown, is individual differences. Variation complicates the task of figuring out which observations or measurements are directly comparable across subjects. Data aggregation may therefore be purpose-driven and theory-laden in interesting ways throughout the psychological and social sciences.

9 Conclusion

This paper has aimed to characterize, defend, and explore the implications of a pluralist view about the cartographic approach to aggregation across brains. Because registration aims to align

homologous regions, but not all homologous regions can be simultaneously aligned, there can be no single spatial registration method that suffices for all neuroscientific purposes. A few neuroscientists endorse the thesis of registration pluralism outright (Hellier *et al*. [2003]), some acknowledge it implicitly (Dubois and Adolphs [2016]), but at least a few seem to reject it (Robinson *et al*. [2014]). Even those who recognize the truth of registration pluralism may not have fully grappled with its methodological and philosophical implications, which I have begun to explore.

Neuroscience has undoubtedly benefited from its increasing reliance on multi-subject studies in addition to single, pathological cases: it has been able to study normal functioning in non-clinical populations, enjoyed an increase in statistical power, and produced findings that are more generalizable. But the field is still wrestling with the additional challenges posed by the aggregation of data across brains. As I hope to have shown here, reflection on the strategies used for aggregation is needed to appropriately deploy neuroscientific tools in multi-subject research and fully understand the results they produce.

Acknowledgements

My thanks to two anonymous reviewers, whose thoughtful feedback helped improve this paper. I also am grateful for comments from Mazviita Chirimuuta, David Danks, Shivam Patel, Marco Viola, Naftali Weinberger, Jim Woodward, and (especially) Edouard Machery. Thanks to Aaron Novick, Lena Kästner, and Jessey Wright for helpful conversations about homology, big data neuroscience, and fMRIPrep, respectively. Audiences at Pitt HPS's WIP series, the Society for Philosophy of Science in Practice 2018 meeting, the British Society for the Philosophy of Science 2018 meeting, and the Philosophy of Science Association 2018 meeting also provided generous and useful feedback.

Zina B. Ward
Department of History & Philosophy of Science
University of Pittsburgh
Pittsburgh, PA, USA
zina.b.ward@pitt.edu

References

Abouheif, E., Akam, M., Dickinson, W. J., Holland, P. W., Meyer, A., Patel, N. H., Raff, R. A., Roth, V. L. and Wray, G. A. [1997]: 'Homology and Developmental Genes', *Trends in Genetics: TIG*, **13**, pp. 432–3.

Amunts, K., Hawrylycz, M. J., Van Essen, D. C., Van Horn, J. D., Harel, N., Poline, J.-B., De Martino, F., Bjaalie, J. G., Dehaene-Lambertz, G., Dehaene, S., Valdes-Sosa, P., Thirion,

B., Zilles, K., Hill, S. L., Abrams, M. B., Tass, P. A., Vanduffel, W., Evans, A. C. and Eickhoff, S. B. [2014]: 'Interoperable Atlases of the Human Brain', *NeuroImage,* **99**, pp. 525–32.

Amunts, K., Schleicher, A., Bürgel, U., Mohlberg, H., Uylings, H. B. and Zilles, K. [1999]: 'Broca's Region Revisited: Cytoarchitecture and Intersubject Variability', *The Journal of Comparative Neurology*, **412**, pp. 319–41.

Amunts, K., Schleicher, A. and Zilles, K. [2007]: 'Cytoarchitecture of the Cerebral Cortex--More than Localization', *NeuroImage,* **37**, pp. 1061–1065.

Anderson, M. L. [2014]: *After Phrenology: Neural Reuse and the Interactive Brain*, Cambridge MA: MIT Press.

Anderson, M. L. [2016]: 'Précis of After Phrenology: Neural Reuse and the Interactive Brain', *The Behavioral and Brain Sciences*, **39**, pp. 1–10.

Arslan, S., Ktena, S. I., Makropoulos, A., Robinson, E. C., Rueckert, D., and Parisot, S. [2018]: 'Human Brain Mapping: A Systematic Comparison of Parcellation Methods for the Human Cerebral Cortex', *NeuroImage*, **170**, pp. 5–30.

Ashburner, J., Neelin, P., Collins, D. L., Evans, A. and Friston, K. [1997]: 'Incorporating Prior Knowledge into Image Registration', *NeuroImage*, **6**, pp. 344–52.

Ashburner, J. [2012]: 'SPM: A History', *NeuroImage*, **62**, pp. 791–800.

Azevedo, B.F., Furieri, L. B., Peçanha, F. M., Wiggers, G. A., Vassallo, P. F., Simões, M. R., Fiorim, J., de Batista, P. R., Fioresi, M., Rossoni, L., Stefanon, I., Alonso, M. J., Salaices, M. and Vassallo, D. V. [2012]: 'Toxic Effects of Mercury on the Cardiovascular and Central Nervous Systems', *Journal of Biomedicine and Biotechnology*, **2012**, pp. 1–11.

Bigdely-Shamlo, N., Mullen, T., Kothe, C., Su, K.-M. and Robbins, K. A. [2015]: 'The PREP Pipeline: Standardized Preprocessing for Large-Scale EEG Analysis', *Frontiers in Neuroinformatics*, **9**, pp. 1–20.

Bohland, J. W., Bokil, H., Allen, C. B. and Mitra, P. P. [2009]: 'The Brain Atlas Concordance Problem: Quantitative Comparison of Anatomical Parcellations', *PLOS ONE*, **4**, pp. 1–18.

Boyd, N. M. [2018]: 'Evidence Enriched', *Philosophy of Science,* **85**, pp. 403–421.

Brett, M., Johnsrude, I. S. and Owen, A. M. [2002]: 'The Problem of Functional Localization in the Human Brain', *Nature Reviews Neuroscience*, **3**, pp. 243–9.

Brigandt, I. [2002]: 'Homology and the Origin of Correspondence', *Biology and Philosophy*, **17**, pp. 389–407.

Buckner, R. L., Krienen, F. M. and Yeo, B. T. T. [2013]: 'Opportunities and Limitations of Intrinsic Functional Connectivity MRI', *Nature Neuroscience*, **16**, pp. 832–7.

Chau, W. and McIntosh, A. R. [2005]: 'The Talairach Coordinate of a Point in the MNI Space: How to Interpret It', *NeuroImage*, **25**, pp. 408–16.

Chen, H., Zhao, Y., Li, Y., Lv, J. and Liu, T. [2017]: 'Inter-Subject fMRI Registration Based on Functional Networks', *2017 IEEE 14th International Symposium on Biomedical Imaging,* pp. 863–67.

Conroy, B. R., Singer, B. D., Guntupalli, J. S., Ramadge, P. J. and Haxby, J. V. [2013]: 'Inter-Subject Alignment of Human Cortical Anatomy Using Functional Connectivity', *NeuroImage*, **81**, pp. 400–11.

Crivello, F., Schormann, T., Tzourio-Mazoyer, N., Roland, P. E., Zilles, K. and Mazoyer. B. M. [2002]: 'Comparison of Spatial Normalization Procedures and Their Impact on Functional Maps', *Human Brain Mapping*, **16**, pp. 228–50.

Crum, W. R., Griffin, L. D., Hill, D. L. G. and Hawkes, D. J. [2003]: 'Zen and the Art of Medical Image Registration: Correspondence, Homology, and Quality', *NeuroImage*, **20**, pp. 1425–37.

Crum, W. R., Hartkens, T. and Hill, D. L. G. [2004]: 'Non-Rigid Image Registration: Theory and Practice', *The British Journal of Radiology*, **77**, pp. S140–53.

Dickie, D. A., Shenkin, S. D., Anblagan, D., Lee, J., Cabez, M. B., Rodriguez, D., Boardman, J. P., Waldman, A., Job, D. E. and Wardlaw, J. M. [2017]: 'Whole Brain Magnetic Resonance Image Atlases: A Systematic Review of Existing Atlases and Caveats for Use in Population Imaging', *Frontiers in Neuroinformatics*, **11**, pp. 1–15.

Dubois, J. and Adolphs, R. [2016]: 'Building a Science of Individual Differences from fMRI', *Trends in Cognitive Sciences*, **20**, pp. 425–43.

Esteban, O., Markiewicz, C., Blair, R. W., Moodie, C., Isik, A. I., Aliaga, A. E., Kent, J., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S., Wright, J., Durnez, J., Poldrack, R. and Gorgolewki, K. J. [2019]: 'FMRIPrep: A Robust Preprocessing Pipeline for Functional MRI', *Nature Methods*, **16**, pp. 111–6.

Evans, A. C., Janke, A. L., Collins, D. L. and Baillet., S. [2012]: 'Brain Templates and Atlases', *NeuroImage,* **62**, pp. 911–22.

Fedorenko, E. and Kanwisher, N. [2009]: 'Neuroimaging of Language: Why Hasn't a Clearer Picture Emerged?', *Language and Linguistics Compass*, **3**, pp. 839–65.

Freeman, J. [2015]: 'Open Source Tools for Large-Scale Neuroscience', *Current Opinion in Neurobiology*, **32**, pp. 156–63.

Friston, K. J., Ashburner, J., Frith, C. D., Poline, J.-B., Heather, J. D. and Frackowiak, R. S. J. [1995]: 'Spatial Registration and Normalization of Images', *Human Brain Mapping*, **3**, pp. 165–89.

Frost, M. A. and Goebel, R. [2012]: 'Measuring Structural-Functional Correspondence: Spatial Variability of Specialised Brain Regions after Macro-Anatomical Alignment', *NeuroImage*, **59**, pp. 1369–81.

Gabard-Durnam, L. J., Leal, A. S. M., Wilkinson, C. L. and Levin, A. R. [2018]: 'The Harvard Automated Processing Pipeline for Electroencephalography (HAPPE): Standardized Processing Software for Developmental and High-Artifact Data', *Frontiers in Neuroscience*, **12**, pp. 1–24.

Ganzetti, M., Liu, Q., Mantini, D. and Alzheimer's Disease Neuroimaging Initiative [2018]: 'A Spatial Registration Toolbox for Structural MR Imaging of the Aging Brain', *Neuroinformatics*, **16**, pp. 167–79.

Gholipour, A., Kehtarnavaz, N., Briggs, R., Devous, M. and Gopinath, K. [2007]: 'Brain Functional Localization: A Survey of Image Registration Techniques', *IEEE Transactions on Medical Imaging*, **26**, pp. 427–51.

Goodman, N. [1972]: 'Seven Strictures on Similarity', in *Problems and Projects*, Indianapolis: Bobbs-Merrill.

Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., Hanke, M. and Ramadge, P. J. [2011]: 'A Common, High-Dimensional Model of the Representational Space in Human Ventral Temporal Cortex', *Neuron*, **72**, pp. 404–16.

Hellier, P., Barillot, C., Corouge, I., Gibaud, B., Le Goualher, G., Collins, D. L., Evans, A., Malandain, G., Ayache, N., Christensen, G. E. and Johnson, H. J. [2003]: 'Retrospective Evaluation of Intersubject Brain Registration', *IEEE Transactions on Medical Imaging*, **22**, pp. 1120–30.

Jiang, D., Du, Y., Cheng, H., Jiang, T. and Fan, F. [2013]: 'Groupwise Spatial Normalization of fMRI Data Based on Multi-Range Functional Connectivity Patterns', *NeuroImage*, **82**, pp. 355–72.

Klein, A., Andersson, J., Ardekani, B. A., Ashburner, J., Avants, B., Chiang, M-C., Christensen, G. E., Collins, D. L., Gee, J., Hellier, P., Song, J. H., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R. P., Mann, J. J. and Parsey, R. V. [2009]: 'Evaluation of 14 Nonlinear Deformation Algorithms Applied to Human Brain MRI Registration', *NeuroImage*, **46**, pp. 786–802.

Klein, A., Ghosh, S. S., Avants, B., Yeo, B. T. T., Fischl, B., Ardekani, B., Gee, J. C., Mann, J. J. and Parsey, R. V. [2010]: 'Evaluation of Volume-Based and Surface-Based Brain Image Registration Methods', *NeuroImage*, **51**, pp. 214–20.

Klein, C. [2010]: 'Philosophical Issues in Neuroimaging', *Philosophy Compass*, **5**, pp. 186–98.

Klein, C. [2012]: 'Cognitive Ontology and Region- Versus Network-Oriented Analyses', *Philosophy of Science*, **79**, pp. 952–960.

Langs, G., Golland, P., Tie, Y., Rigolo, L. and Golby, A. J. [2010]: 'Functional Geometry Alignment and Localization of Brain Areas', *Advances in Neural Information Processing Systems*, **1**, pp. 1225–33.

Leonelli, S. [2016]: *Data-Centric Biology: A Philosophical Study*, Chicago: University of Chicago Press.

Liebeskind, B. J., Hillis, D. M., Zakon, H. H. and Hofmann, H. A. [2016]: 'Complex Homology and the Evolution of Nervous Systems', *Trends in Ecology & Evolution*, **31**, pp. 127–35.

Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Le Goualher, G., Boomsma, D., Cannon, T., Kawashima, R. and Mazoyer, B. [2001]: 'A Probabilistic Atlas and Reference System for the Human Brain: International Consortium for Brain Mapping (ICBM)', *Philosophical Transactions of the Royal Society of London Series B*, **356**, pp. 1293–1322.

McCaffrey, J., and Danks, D. [forthcoming]: 'Mixtures and Psychological Inference with Resting State fMRI', *The British Journal for the Philosophy of Science*.

Nenning, K.-H., Liu, H., Ghosh, S. S., Sabuncu, M. R., Schwartz, E. and Langs, G. [2017]: 'Diffeomorphic Functional Brain Surface Alignment: Functional Demons', *NeuroImage*, **156**, pp. 456–65.

Ng, B., Abugharbieh, R. and McKeown, M. J. [2009]: 'Adverse Effects of Template-Based Warping on Spatial fMRI Analysis', *Medical Imaging 2009: Biomedical Applications in Molecular, Structural, and Functional Imaging*, **7262**, pp. 1–12.

Nieto-Castañón, A. and Fedorenko, E. [2012]: 'Subject-Specific Functional Localizers Increase Sensitivity and Functional Resolution of Multi-Subject Analyses', *NeuroImage*, **63**, pp. 1646–69.

Orden, G. C. van. [1997]: 'Functional Neuroimages Fail to Discover Pieces of Mind in the Parts of the Brain', *Philosophy of Science Supplement*, **64**, pp. 85–94.

Owen, R. [1843]: *Lectures on the Comparative Anatomy and Physiology of the Invertebrate Animals*, London: Longman, Brown, Green, and Longmans. https://www.biodiversity library.org/bibliography/6788.

Panchen, A. L. [1994]: 'Richard Owen and the Concept of Homology', in B. Hall (*ed.*), *Homology: The Hierarchical Basis of Comparative Biology*, San Diego: Academic Press.

Poline, J.-B., Thirion, B., Roche, A. and Merlaux, S. [2010]: 'Intersubject Variability in fMRI Data: Causes, Consequences, and Related Analysis Strategies', in S. J. Hanson and M. Bunzl (*eds*), *Foundational Issues in Human Brain Mapping*, Cambridge MA: MIT Press.

Rahman, M., Murad, G. J. A. and Mocco, J. [2009]: 'Early History of the Stereotactic Apparatus in Neurosurgery', *Neurosurgical Focus*, **27**, pp. 1–5.

Ramsey, G. and Peterson, A. [2012]: 'Sameness in Biology', *Philosophy of Science*, **79**, pp. 255–75.

Robinson, E. C., Jbabdi, S., Glasser, M. F., Andersson, J., Burgess, G. C., Harms, M. P., Smith, S. M., Van Essen, D. C. and Jenkinson, M. [2014]: 'MSM: A New Flexible Framework for Multimodal Surface Matching', *NeuroImage*, **100**, pp. 414–26.

Roland, P. E. and Zilles, K. [1994]: 'Brain Atlases - a New Research Tool', *Trends in Neurosciences*, **17**, pp. 458–67.

Roskies, A. L. [2007]: 'Are Neuroimages Like Photographs of the Brain?', *Philosophy of Science*, **74**, pp. 860–72.

Saad, Z. S. and Reynolds, R. C. [2012]: 'SUMA', *NeuroImage*, **62**, pp. 768–73.

Sabuncu, M. R., Singer, B. D., Conroy, B., Bryan, R. E., Ramadge, P. J. and Haxby, J. V. [2010]: 'Function-Based Intersubject Alignment of Human Cortical Anatomy', *Cerebral Cortex*, **20**, pp. 130–40.

Scheperjans, F., Eickhoff, S. B., Hömke, L., Mohlberg, H., Hermann, K., Amunts, K. and Zilles, K. [2008]: 'Probabilistic Maps, Morphometry, and Variability of Cytoarchitectonic Areas in the Human Superior Parietal Cortex', *Cerebral Cortex*, **18**, pp. 2141–57.

Shattuck, D. W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K. L., Poldrack, R. A., Bilder, R. M. and Toga, A. W. [2008]: 'Construction of a 3D Probabilistic Atlas of Human Cortical Structures', *NeuroImage*, **39**, pp. 1064–80.

Talairach, J., Szikla, G., Tournoux, P., Prosalentis, A., Bordas-Ferrier, M., Covello, L., Iacob, M. and Mempel, E. [1967]: *Atlas D'anatomie Stereotaxique Du Telencephale*, Paris: Masson.

Talairach, J. and Tournoux, P. [1988]: *Co-Planar Stereotaxic Atlas of the Human Brain*, New York: G. Thieme Medical Publishers.

Toga, A. W. [1998]: *Brain Warping*, San Diego: Academic Press.

Toga, A. W. and Mazziotta, J. C. [2002]: *Brain Mapping: The Methods*, 2nd ed., Boston: Academic Press.

Toga, A. W. and Thompson, P. M. [2001]: 'Maps of the Brain', *The Anatomical Record*, **265**, pp. 37–53.

Toga, A. W. and Thompson, P. M. [2007]: 'What Is Where and Why It Is Important', *NeuroImage*, **37**, pp. 1045–68.

Toga, A. W., Thompson, P. M., Mori, S., Amunts, K. and Zilles, K. [2006]: 'Towards Multimodal Atlases of the Human Brain', *Nature Reviews Neuroscience*, **7**, pp. 952–66.

Tucholka, A., Fritsch, V., Poline, J.-B. and Thirion, B. [2012]: 'An Empirical Comparison of Surface-Based and Volume-Based Group Studies in Neuroimaging', *NeuroImage*, **63**, pp. 1443–53.

Uylings, H. B. M., Rajkowska, G., Sanz-Arigita, E., Amunts, K. and Zilles, K. [2005]: 'Consequences of Large Interindividual Variability for Human Brain Atlases: Converging Macroscopical Imaging and Microscopical Neuroanatomy', *Anatomy and Embryology*, **210**, pp. 423–31.

Van Essen, D. C. and Dierker, D. [2007]: 'On Navigating the Human Cerebral Cortex Response to "In Praise of Tedious Anatomy"', *NeuroImage*, **37**, pp. 1050–68.

Wagner, G. P. [1994]: 'Homology and the Mechanisms of Development', in B. Hall (*ed.*), *Homology: The Hierarchical Basis of Comparative Biology*, New York: Academic Press.

Wagner, G. P. [2014]: *Homology, Genes, and Evolutionary Innovation*, Princeton: Princeton University Press.

Watson, J. D. G., Myers, R., Frackowiak, R. S. J., Hajnal, J. V., Woods, R. P., Mazziotta, J. C., Shipp, S. and Zeki, S. [1993]: 'Area V5 of the Human Brain: Evidence from a Combined Study Using Positron Emission Tomography and Magnetic Resonance Imaging', *Cerebral Cortex*, **3**, pp. 79–94.

Weiner, K. S., Barnett, M. A., Lorenz, S., Caspers, J., Stigliani, A., Amunts, K., Zilles, K., Fischl, B. and Grill-Spector, K. [2017]: 'The Cytoarchitecture of Domain-Specific Regions in Human High-Level Visual Cortex', *Cerebral Cortex*, **27**, pp. 146–61.

Woods, R. P., Grafton, S. T., Watson, J. D., Sicotte, N. L. and Mazziotta, J. C. [1998]: 'Automated Image Registration: II. Intersubject Validation of Linear and Nonlinear Models', *Journal of Computer Assisted Tomography*, **22**, pp. 153–65.

Zhang, Q., Borst, J. P., Kass, R. E. and Anderson, J. R. [2017]: 'Inter-Subject Alignment of MEG Datasets in a Common Representational Space', *Human Brain Mapping*, **38**, pp. 4287–301.

Zhou, Y., Yap, P.-T., Zhang, H., Zhang, L., Feng, Q. and Shen, D. [2017]: 'Improving Functional MRI Registration Using Whole-Brain Functional Correlation Tensors', *Medical Image Computing and Computer-Assisted Intervention*, **10433**, pp. 416–23.

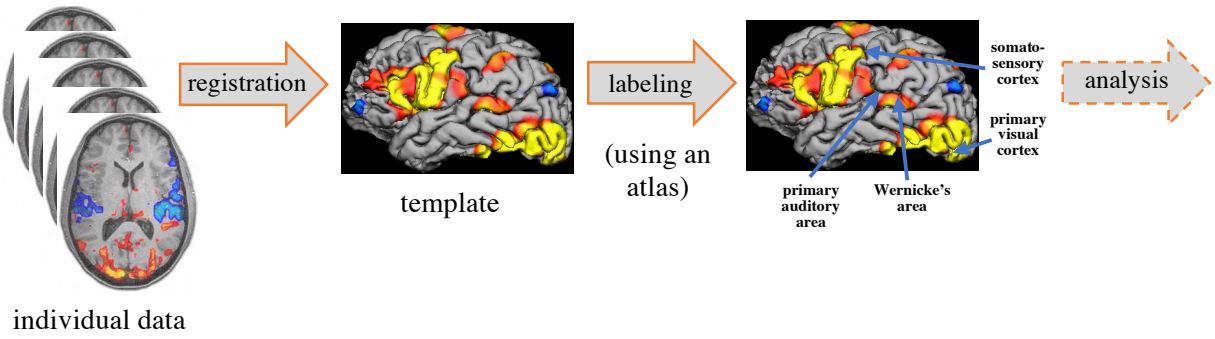

Figure 1. Schematic representation of the cartographic approach.