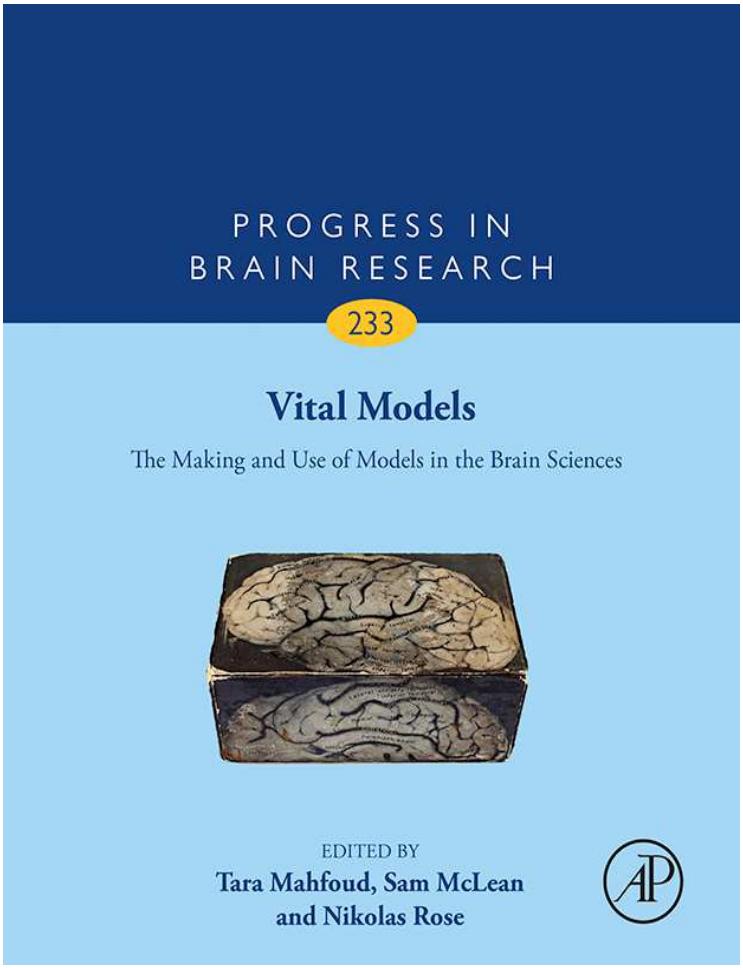

From Robert Turner and Daniel De Haan, Bridging the gap between system and cell: The role of ultra-high field MRI in human neuroscience. In: Tara Mahfoud, Sam McLean and Nikolas Rose, editors, *Progress in Brain Research, Vol. 233*, Amsterdam: Academic Press, 2017, pp. 179-220.
ISBN: 978-0-12-804215-1

CHAPTER

8

Bridging the gap between system and cell: The role of ultra-high field MRI in human neuroscience

Robert Turner[*,1], **Daniel De Haan**[†]

[*]*Max-Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany*

[†]*Translational Cognitive Neuroscience Lab, University of Cambridge, Cambridge, United Kingdom*

[1]*Corresponding author: e-mail address: turner@cbs.mpg.de*

Abstract

The volume of published research at the levels of systems and cellular neuroscience continues to increase at an accelerating rate. At the same time, progress in psychiatric medicine has stagnated and scientific confidence in cognitive psychology research is under threat due to careless analysis methods and underpowered experiments. With the advent of ultra-high field MRI, with submillimeter image voxels, imaging neuroscience holds the potential to bridge the cellular and systems levels. Use of these accurate and precisely localized quantitative measures of brain activity may go far in providing more secure foundations for psychology, and hence for more appropriate treatment and management of psychiatric illness. However, fundamental issues regarding the construction of testable mechanistic models using imaging data require careful consideration. This chapter summarizes the characteristics of acceptable models of brain function and provides concise descriptions of the relevant types of neuroimaging data that have recently become available. Approaches to data-driven experiments and analyses are described that may lead to more realistic conceptions of the competences of neural assemblages, as they vary across the brain's complex neuroanatomy.

Keywords

MRI, Imaging neuroscience, Mechanism, Teleological models, Cognitive science, Neuroanatomy, Cortical parcellation, Myeloarchitecture, Causal directionality

Progress in Brain Research, Volume 233, ISSN 0079-6123, http://dx.doi.org/10.1016/bs.pbr.2017.05.005

1 INTRODUCTION

The central aim of this chapter is to sketch a methodological proposal for bridging the gap between cellular and systems levels in human neuroscience, by making use of the technological advances provided by anatomical and functional magnetic resonance imaging (MRI) at the ultra-high field strength of 7 T, together with rigorous data analysis, integrated into a reflective conceptual framework combined with detailed multilevel mechanistic explanations.

Many brain researchers and psychiatrists agree that there is something of a crisis in brain science, extending from the exploration of neural connectivity to the definition and treatment of mental illnesses (Canino and Alegria, 2008; Kendler, 2016; Kendler et al., 2011; Murphy, 2006). While relatively large and increasing research funding is becoming available across this entire spectrum of research, one notes with concern the rows over how the European Human Brain Project should be steered and directed, the distress regarding "voodoo" brain imaging results (Vul et al., 2009) caused by bad practice in data analysis (resurfacing in a further panic about estimates of spatial extent; Eklund et al., 2016; Shifferman, 2015), and reactions to the author's own critique of neuroimaging analyses which rely on drastic spatial smoothing of reasonably high-resolution raw data (Turner and Geyer, 2014b). Furthermore, Uttal (2011), Bennett and Hacker (2003), and Anderson (2014) criticize sharply the tendency to ascribe very specific roles to specific brain areas.

Some words are needed concerning the focus of this chapter on modeling human brains, rather than those of other mammalian species. In the first place, the potential clinical value of vastly improved understanding of brain processes clearly applies mostly to humans, and experimental results on the brain are obviously easiest to interpret if they derive from the same species. Besides this, the issue of fundamental importance is the link between function and anatomy. The range of well-characterized brain tasks that can be reliably performed by conscious humans, often needing only a few words of instruction, is far larger than that feasible for animal models, which may have to endure many days of rigorous training that actually change the brain structures of interest. What the advent of MRI has created in recent years, empowering a huge expansion of brain research on humans, is the possibility of comparing specific functions with specific individual brain anatomy.

A number of preliminary issues pertain to the uncritical adoption and employment of particular conceptual frameworks or theoretical assumptions in the descriptions, interpretations, and explanations of experimental data. Many findings of experimental psychology have been shown to depend strikingly on the cultural and indeed population genetic background of the human populations studies (Chiao and Immordino-Yang, 2013; Henrich et al., 2010), calling into question the conceptual framework of this discipline (Turner, 2012). At the level of psychiatry, the current uncertainty is attested by the far-reaching revision by the American Psychiatric Association of the categories of mental illness (DSM—Diagnostic and Statistical Manual of Mental Disorders), and the decision in 2013 of the National

Institute for Mental Health to refuse to consider grant applications which rely on DSM categories, instead recognizing value only in applications which deal with Research Domain Criteria rather than patient labels. Concern has grown enormously regarding the overuse and misuse of psychopharmaceutical drugs, many of which have much lower effectiveness than initially advertised (Bentall, 2009). Neuroimaging research in psychiatric medicine is currently quite widespread, but it is strikingly difficult to discern characteristic and reproducible differences in structure or in functional organization between the brain of a normal person and that of someone whose social behavior is extremely abnormal.

At the root of this unease lies the fundamental problem of relating the mind to the brain. For more than 2 centuries, many researchers (among them Sigmund Freud) have been motivated by the hope that an area of discourse could be developed in which mind and brain are both comfortably included. This would hopefully enable mechanistic explanatory and predictive modeling of human behavior, and facilitate the analysis of the complex multilevel causal and constitutional mechanisms—including molecular, genetic, neural cells, circuits, and systems, endocrine systems, psychological, ecological, social–cultural—that produce, underlie, and sustain psychiatric syndromes (Kendler et al., 2011). The research field of cognitive neuroscience is focused around this central question, attempting to find brain-based explanations of human and animal cognitive abilities. At the cellular, subcellular, and biochemical levels, using animal models, good theories of phenomena such as experience-driven changes in synaptic efficiency have been established, which have testable predictions. This would hopefully enable mechanistic predictive modeling of human behavior, and facilitate the treatment of mental illness as a brain disease.

It is not yet clear whether relabeling "mind" as "cognition" or treating them as synonyms, in moving the field of enquiry from psychology to cognitive science, brings our understanding further forward (Bennett and Hacker, 2003, 2012; Chemero, 2009). Both terms and their ubiquitous application in psychology and cognitive neuroscience bring with them a variety of conceptual problems. For instance, if these concepts are to be representative of common sense psychology (or folk psychology) then they need to be exhibited across cultures, and not be restricted to the atypical populations most often studied, that have been characterized as WEIRD (Western, educated, industrialized, rich, and democratic, Henrich et al., 2010). Unfortunately "mind" and its cognates are only awkwardly translated into many other languages and cultures, and thus their ontological validity remains in serious doubt (Turner, 2012).

Additional difficulties arise when the terminology of psychology comes to be applied to the processes which take place in the brain. Many researchers take the English word "cognition" to mean the processes internal to the brain that culminate in the encoding of memories, planning of action, or directly as immediate actions. This common practice must be regarded as metaphorical, and thus not really scientific (Bennett and Hacker, 2003, 2012). For these reasons, the introductory and

discussion sections of many publications in cognitive neuroscience can be regarded as speculative and tendentious.

A somewhat less ambitious area of enquiry deals with the relationship between cellular and systems neuroscience. Systems neuroscience is concerned with patterns of activity taking place across entire brains and the peripheral nervous system. While most of the research in this field considers only macroscale spatial distributions of brain activity, sporadic efforts are made to model system-level activity in terms of known properties of the constituent neurons and other relevant types of brain cell, such as astrocytes. Attaching psychological labels to patterns of activity at this system level are seen to be less far-fetched than ascribing memory, volition, intention, perception, and so forth to component neurons. And systems neuroscience in the human brain prospers, due to dramatic technological developments made since 1980 largely by physicists.

But the magnitude of the task of modeling the brain should never be underestimated. Lohmann et al. (2013) writes:

> *Most (if not all) researchers today will agree that the human brain is a complex adaptive system … (that has) a large number of components … that interact and adapt or learn.*
>
> *One of the key features of complexity is thus the sheer size of the system. Small systems with few constituents can never display a considerable complexity by this definition … Geometric studies suggest that this property is more dominant in systems with a large number of constituents … Systems with many constituents have a higher probability of developing an adaptation process to increase complexity.*
>
> *Complexity is quite distinct from chaos. Chaos is characterized by high sensitivity to initial conditions resulting in unpredictable behaviour (not to be confused with random noise), while the hallmark of complexity is emerging behaviour and a tendency toward self-organization. Complex systems exhibit an interplay of chaos and complexity, whereas fully chaotic systems are unpredictable in their entire phase space. Very small systems can be fully chaotic (e.g., a double pendulum), while complexity requires positive feedback mechanisms as well as a large number of nonlinearly interacting agents. If our brains were merely chaotic, but not complex, we would certainly not be able to survive.*
>
> *Another aspect to be considered is that the complex system inside our brain is not static, but changes over time and is thus termed an adaptive system. We need to distinguish adaptability due to the fast dynamics of activation patterns versus the slower changes in synaptic coupling strength. Both aspects come into play and leave traces in the data that we measure.*

In the early 1980s, the leading noninvasive method for studying the localization of brain function was positron emission tomography (reviewed in Turner and Jones, 2003). Most studies involved averaging results obtained using a particular task paradigm across several human subjects. In order to match corresponding cortical areas across spatially normalized brains, a smoothing kernel of 10–15 mm was normally applied to the image data before further statistical analysis. In any case, the intrinsic resolution of PET was only about 5 mm. In order to identify where activation occurred in the averaged brain, statistical maps were generated which could then be thresholded to reveal compact patches of increased blood flow (Friston et al., 1991).

Compared with earlier approaches to functional brain mapping, the results from this strategy were deemed highly satisfactory, and few researchers worried whether the resultant large brightly colored blobs supposedly showing areas of significant activation, superimposed on nonquantitative T1-weighted structural MR images, gave an adequate picture of the brain areas involved in particular tasks. Instead, most brain scientists were happy that this analysis appeared to show that brain activity was strongly segregated into specialized regions, as opposed to the hitherto popular connectionist view of Lashley (1929) and others (see below). When blood oxygenation level-dependent (BOLD) contrast fMRI was discovered by Ogawa et al. (1990), Turner et al. (1991), and Kwong et al. (1992), almost all users of this higher resolution technique processed their much more abundant data to generate similar highly smoothed maps.

It was inevitable that such maps would foster the illusion of cortical modules. Even though Fodor, the most influential promulgator of the idea of cognitive modularity (Fodor, 1983) himself repudiated the idea that cognitive modules would be necessarily manifested as compact spatial groupings of neurons, it became highly popular, with thousands of research studies claiming that this or that psychological task was actually performed by a specific cortical area. The raw functional MRI data were rarely shown in the research publications of cognitive psychology. Instead, unnecessary spatial smoothing of the functional images was commonly performed before further processing, which simplified the statistical analysis but essentially ruled out the possibility of associating brain functional activity with specific neural substrates. The misleading pictures of activated brains thus derived also fostered the erroneous notion that modeling functional activity as graphs with relatively few nodes and edges, for instance dynamic causal modeling (Friston et al., 2003) could be useful in explaining brain function (Lohmann et al., 2012).

We should distinguish between distinct kinds of cognitive modularity from distinct kinds of neural modularity and recognize the real challenges of bringing these two together.

> *A frequently made assumption is that the mind can be subdivided into modules or parts whose activity can then be studied with fMRI. If this assumption is false, then even if the brain's architecture is modular, we would never be able to map mind modules onto brain structures, because a unified mind has no components to*

> *speak of. Even if true, the challenge remains in coming up with the correct recursive decompositions—in each of which any given cognitive capacity, however abstract, is divided into increasingly smaller functional units that are localized to specific brain parts, which in turn can be detected and studied with fMRI. This is not a neuroimaging problem but a cognitive one. Hierarchical decompositions are clearly possible within different sensory modalities and motor systems. Their mapping, which reflects the brain's functional organization, is evidently possible and certainly meaningful beyond any reasonable doubt.*
>
> **Logothetis (2008)**

Despite the potentially excellent spatial resolution of fMRI, very little research has been done to discover the data-driven natural scales of granularity in human brain. Now, using ultra-high field MRI at 7 T, the spatial resolution can be as high as 0.5 mm for functional scanning, and 350 μm for structural images. Such structural images can show layers of myelinated axons within the cerebral cortex, enabling a direct evaluation of the relationship between structure and function. Using the observable myeloarchitecture as a guide, in principle one can deduce the cytoarchitecture (Nieuwenhuys, 2013; Van Essen and Glasser, 2014), and thus address the relationship between local neuronal circuits and specific cortical competences. Furthermore, novel methods for assessing cerebral blood volume, rather than changes in blood oxygenation, promise to provide a reliable quantitative marker of layer-dependent activation.

With such tools at one's disposal, the way is open toward greatly improved models of human brain function, which may show promise in bridging the gap between systems neuroscience—the province of cognitive neuropsychology—and cellular neuroscience. Novel concepts and techniques are becoming available that can provide insight at this level of explanation: the study of neural reuse at a microscopic scale (Anderson, 2014) using repetition priming experimental designs; detailed voxel encoding (Stansbury et al., 2013) using very large stimuli sets; and data-driven modeling of phenomenological characterizations of brain function (Gallagher et al., 2013) using ecologically valid experimental strategies.

In order to justify the approach just outlined, this chapter describes a well-established perspective on explanation and modeling in systems-level neuroscience and includes the addressing of another important issue—the level of detail needed to provide adequate naturalistic modeling.

2 MODELING THE BRAIN

2.1 MECHANISTIC MODELS

The brains of mammals contain some 10^8–10^{11} neurons, organized into well-defined groupings, and each neuron possesses up to thousands of synaptic connections with other neurons. Statistically, despite the huge number of neurons, a signal only needs to cross about five synapses to travel from any given neuron to any other. Besides this

structural complexity, which is continually modified as development and experience drive axonal myelination, changes in the strength of synapses, the formation of new synapses, and the loss of existing synapses, the mode of transmission of signals between neurons can vary greatly. In addition to the principal neurotransmitters, excitatory glutamate, and inhibitory gamma-aminobutyric acid (GABA), there are 4 major modulatory neurotransmitters, and up to 20 less-common neurotransmitters relating to appetite and other regulatory functions. Nitric oxide, hormones, and several neuropeptides (such as oxytocin) also play important modulatory roles in brain activity, adjusting the "set points" of activity locally and even globally. The high density and interconnectedness of neurons found in the brain are also found in the spinal cord. In the form of peripheral nerves, neurons eventually terminate at muscle fibers and sensory receptor cells. Surrounding and supporting the neurons there is an interconnected network of a similar number of astrocytes, whose electrochemical connections are typically of shorter range but are known to have a crucial role in neuronal signal transmission.

In the face of this amazing complexity, how can human scientists possibly make sense of brain function? Scientific understanding only increases when simple models can be devised to describe quantitatively the operations of a complex system, so that given appropriate boundary and initial conditions, the future behavior of the system can be predicted with a nontrivial degree of accuracy. A contemporary "new mechanistic" approach to biology, neuroscience, and psychology is led by Carl Craver, among other exponents well versed in the philosophy of science (Bechtel, 2006, 2008; Bechtel and Richardson, 2010; Cartwright, 1999; Craver and Darden, 2013; Piccinini and Craver, 2011). Craver (2007) provides a robust account of well-accepted strategies for developing satisfactory mechanistic models in neuroscience, based on a broad reading of influential research publications and his own expertize in the philosophy of science—amounting to something resembling an ethnology of brain modeling cultures.

> *Mechanisms are entities and activities organized such that they exhibit the explanandum phenomenon.*
>
> **Craver (2007, p. 6)**

The phenomenon to be explained is some behavior or capacity such as a mental illness, episodic memory, spatial navigation, generation of spatial maps in the hippocampus, circadian rhythms, neurotransmitter release, action potentials, activation of NMDA receptors, and so on. Mechanisms produce, underlie, or maintain phenomenon.

> *Mechanistic explanations are constitutive or componential explanations: they explain the behaviour of the mechanism as a whole in terms of the organized activities of its components.*
>
> **Craver (2007, p. 128)**

Mechanisms are constituted by the spatial, temporal, active organization of their component entities, and activities. The component parts of a mechanisms comprise

various "entities (ion channels, active zones, a host of intracellular molecules, vesicles containing neurotransmitters, fusion pores, and neural membranes) and their various activities (opening, clamping, diffusing, docking, fusing, incorporating, phosphorylating, and priming)" (Craver, 2007, p. 5).

Most credible biological mechanistic explanations, especially in neuroscience, are not centrally concerned with the explanatory reduction of the activities of higher levels to those of the lower, but aim instead to "bridge multiple levels and that require the disciplinary expertise of multiple fields" (Craver, 2007, p. 265).

> *Appeal to mechanisms is not necessarily reductionistic. Mechanisms are often described as multi-level, with activities at different levels being equally essential to how a mechanism works. Mechanistic explanations might look, up, down, or around depending on the choice of an explanandum and the presuppositions of the explanatory context.*
>
> **Bechtel (2009)**

> *Mechanists can be reductionists or anti-reductionists. That said, many mechanists opt for some form of explanatory anti-reductionism, emphasizing the importance of multilevel and upward-looking explanations, without rejecting the central ideas that motivate a broadly physicalist world-picture.*
>
> **Craver and Darden (2013)**

Craver uses the integrated interlevel experiments on the mechanisms of spatial memory, among many others, to illustrate the way that "the phenomena at each of these levels—NMDA receptor function, LTP [long term potentiation], spatial map formation, and spatial memory—is constitutively relevant to the next" (Craver, 2007, p. 265). By experimentally knocking out causal interactions of component entities at a lower level, activities at a higher level can be inhibited or undermined that are enabled constitutionally by the organized activities of the mechanism's components. Crucially, in most cases:

> *the operations within a mechanism are different from the phenomenon produced by the mechanism. Within a neuron, for example, neurotransmitters perform such operations as diffusing across a synapse and binding to a receptor; but the neuron itself generates action potentials. The point of organizing component parts and operations into a mechanism is to accomplish something that cannot be performed by the individual components. Hence, assuming a homunculus with the same capacities as the agent in which it is posited to reside clearly produces no explanatory gain. The recognition that it is problematic to assume that operations within a mechanism perform the same type of operations as the mechanism itself may be a major reason many find problematic Fodor's (1975) proposal of a language of thought to explain language and thought.*
>
> **Bechtel (2009)**

New mechanists provide many insights into the questions, constraints, and strategies that guide scientists searching for mechanisms. They distinguish the search for biological mechanisms into four stages: characterizing the phenomenon, constructing a mechanistic schema, evaluation, and revision.

> *Biologists use mechanism schemas to describe, explain, explore, organize, predict, and control phenomena.*
>
> **Craver and Darden (2013, p. 91)**

> *The aim of scientific investigation is to transform a black box sketch of some mechanism and its organized components into a grey box sketch that can eventually become a transparent glass box schema that completely or at least adequately represents the ways in which all the levels of a mechanism are constituted from the interactions of their components. These boxes are gradually filled in by investigating bit-by-bit how all the components in the mechanisms work and function together often through a decomposition of a mechanism into its lower level components. Understanding the interaction of the components of a mechanism often involves moving from how-possible models to how-actual models of the organized components of a mechanism. We begin constructing our how-possible model of a mechanism by considering all of its possible constraints; the components of mechanisms have a variety of constraints such as their locations, structures, abilities, activities, functions, what they produce, underlie, maintain, and their overall organization. What does this neural assembly do? Where is it located? What can the cytoarchitecture tell us about the abilities of these neurons? What features of these components make a difference to the phenomenon of interest? How-plausible models are distinguished from how-possible models by evaluating the degree to which some models can accommodate more of the known constraints of the mechanism that come to light through further investigation.*
>
> **Craver and Darden (2013)**

> *The central criterion of adequacy for a mechanistic explanation is that it should account for the multiple features of the phenomenon, including its precipitating conditions, manifestations, inhibiting conditions, modulating conditions, and nonstandard conditions…. The model of a mechanism does not describe capacities of the mechanism as a whole; it describes the activities of the mechanism's components. How-possibly models can be composed of fictional components, but how-actually models describe real components that have multiple properties, that are detectable with multiple techniques, that are utilizable for the purposes of intervention, and that are physiologically relevant. The model of the mechanism also describes the causal relations (activities) that compose the mechanism.*
>
> **Craver (2007, p. 139)**

It is thus vitally important for the success of mechanistic models of neurons, neural assembles, and neural systems that they should be consistent with what we know these neurons can actually do Craver's detailed account of multilevel mechanistic explanations, interfield integration, the process of discovering, evaluating, and revising how-actually models from how-possible models, and a surfeit of illustrations from neuroscience, merits a close consideration. While Craver's examples are generally drawn from experimental research at the scale of neurons or small assemblages of neurons, this discussion will apply the principles that Craver enunciates to systems-level attempts to formulate tractable models, and thus provide criteria for the likely success of such modeling. The appropriateness of Craver's characterization of successful mechanistic explanation is borne out in numerous examples from the literature of cellular-level neuroscience, which do not need to be rehearsed here. What are scarcer, however, are successful examples of mechanistic explanation at the level of cognitive neuroscience. Two severe difficulties are at once encountered.

The first difficulty arises in the task of defining components (in Craver's sense as organized entities and activities: see "Field Guide to Levels" in Craver, 2007, chap. 5). Models constructed at the scale of the entire brain that specify individual neurons as components are obviously unusable. While much is known regarding the typical characteristics of typical neurons, no method has yet been established for mapping every neuron in an actual mammalian brain, together with its activity over time. Even if such data could be obtained, the computational power available in the most powerful computers yet available would be wholly inadequate to the task of making testable predictions of the future states of such a model system.

In the physics of matter, models are successful when they employ principles of statistical mechanics to characterize the collective behavior of very large numbers of microcomponents that can be considered as identical, very often describing such behavior as emergent properties such as pressure, temperature, and so forth. But neurons cannot be regarded as identical in their properties. Even the 305 neurons of the nematode *Caenorhabditis elegans* are separately identifiable, and each performs a separate function, in combination with the activities of many of the other neurons. Ablation of even one of these neurons can have striking specific effects on the animal's behavioral repertoire.

Some authors have attempted to explain mammalian brain function using the premise of "neural mass modeling." Here the activity of a group of neurons is supposed to be approximated by their average properties and interactions, located at a point in the brain, together with postulated interactions with a small number of other groups of neurons located at other points in the brain. However, this model entails the unlikely presupposition that neurons in a mammalian brain are far less specific and differentiated than those comprising the nervous system of a nematode. A further difficulty arises in the way that this approach has been applied, which due to drastic spatial smoothing typically takes little account of neuroanatomical details, indeed conflating gray matter and white matter, let alone grouping together separate brain areas on either side of sulci. Classifying brain tissue on the basis of the cognitive or perceptual tasks supported (e.g., Haxby et al., 2014) can provide some insight into

the basic competences of each neural assembly, but far more creative work is yet needed to find realistic and useful labels for each assembly (see below).

The question then arises, is there any other way of grouping neurons, so that they form physiologically relevant components of a model that is computationally tractable? Craver (2007) describes levels of mechanisms as follows:

> *Organization is the inter-level relation between a mechanism as a whole and its components. Lower-level components are made up into higher-level components by organizing them spatially, temporally and actively into something greater than a mere sum of these parts.*

Thus, one might hope to progress downward through levels by analyzing the structure of the top-level components as assemblies of already well-characterized parts, such that the ascribed properties of the top-level components are consistent with the arrangement, properties, and activities of the elements of these assemblies. In the context of systems neuroscience, the first task is to define the top-level components. Several possibilities offer themselves. One might take brain regions, defined as the collection of identifiable cortical areas (Brodmann, 1909; Glasser et al., 2016; Turner, 2016; Vogt and Vogt, 1919) together with deep brain nuclei (Keuken et al., 2014). This is discussed further in a later section of this chapter.

2.2 DESCRIPTION AND ATTRIBUTION OF FUNCTION

The second major difficulty lies in defining the phenomena to be explained by the model—the physiological activities of the brain. At the cognitive level, one is generally dealing with the constructs formulated over the last 150 years within Western psychology. The ontological validity of many such constructs is seldom seriously investigated (see Turner, 2012) and their cross-cultural robustness is far from certain. Much work has been invested into "how-possibly" models, using boxes representing brain modules that are labeled with terms drawn from Western psychology, which interact with each other and the outside world along posited pathways and causal directions. However, few efforts have been made to justify the plausibility of each box by showing how the combined activity of the neurons within each box could collectively demonstrate the assumed behavior. It can also be strongly argued (for instance by Anderson, 2014, p. 104) that it is not rare for local neural assemblies to be "multifunctional," taking part in a range of brain systems that each accomplish different tasks. Thus the spatial compactness of an apparent cognitive representation cannot be taken as a criterion of the ontological stability of that cognitive faculty.

Poldrack and Yarkoni (2016) have called for a far more empirical approach to the classification of cognition, via a "Cognitive Atlas" that is intended to capture two primary forms of knowledge:

> *First, it aims to define psychological constructs in order to provide consensus definitions that can serve as the basis for accurate scientific communication and discussion. A fundamental distinction made within the Cognitive Atlas is between*

> *mental concepts, which refer to putative but unobservable psychological processes or structures, and mental tasks, which are the objective operations used to measure those putative constructs. … Second, the project aims to establish a knowledge base of the relations within and between mental tasks and mental concepts. … Within the Cognitive Atlas, mental tasks are described in terms of three primary features: (a) conditions (which specify different conditions of measurement), (b) contrasts (which specify either comparisons between conditions or relationships with continuous variables), and (c) indicators (which specify variables that are measured within the task; these could reflect behavioral, neural, or other physiological measurements). In order to capture the relations between tasks and concepts, a novel ontological relationship (measured-by) was defined that denotes the fact that a specific concept is measured by a specific task. Importantly, concepts are not related to the overall representation of a task but rather to specific contrasts.*

While this approach is well intentioned, it may perpetuate a crypto-Cartesian dualism between mind and brain. Furthermore, it makes no effort to deal with the ethnocentrism of Western psychology, mentioned above and brilliantly exposed by Henrich et al. (2010) "The weirdest people in the world?," because its main purpose is to find more empirical definitions of psychological concepts that are already part of a received lexicon, which may have very little relevance in widely different human cultural settings.

A potentially more fruitful approach to the categorization of cognition and its role in mechanistic brain modeling can be found in work by Maley and Piccinini (2017) (henceforth MP), where they emphasize the concept of "teleological function" in regard to the goals of organisms:

> *The paradigmatic goals of organisms are survival and inclusive fitness, although organisms may have additional goals.*

> *What a trait or part of an organism is for, as opposed to the others things it does, is its teleological function. When a trait fails to perform its teleological function at the appropriate rate in an appropriate situation, it malfunctions.*

> *A teleological function in an organism is a stable contribution by a trait (or component, activity, property) of organisms belonging to a biological population to an objective goal of those organisms.*

> *Construed generally, a trait's function (and sometimes the successful performance of that function) depends on some combination of the organism and its environment … the truthmakers for attributions of functions to an organism's trait are facts about the organism and its environment.*

These authors go on to address the implications of this formalism in regard to cognition and neuroscience:

> *When psychologists posit the performance of cognitive functions within organisms, they offer a sketch of a mechanism. Such a sketch can be completed at that level by specifying which structures perform those functions. That level of explanation can then be combined with other levels by showing how each structure perform its functions in virtue of its lower level organization as well as how each structure fits within a larger containing mechanism. Going down one level involves adding details about a given structure and how it performs its functions; going up one level involves abstracting away from lower level details and fitting a structure into its mechanistic context. When all relevant levels and their mutual relations are understood, the upshot is a unified, integrated explanation of cognition.*

In attempting to formulate a vocabulary of teleological functions, it is important to recognize that the term "cognition" is inherently vague. As was noted, many researchers take this word to mean the processes internal to a brain that culminate in the encoding of memories, planning of action, or directly as immediate actions. However, as far as the electrochemical activity of individual neurons is concerned, the terms "cognition," "action," "perception," "volition," and "emotion" have no distinctive meanings. It is likely that there are no specific markers, in regard to spike trains or patterns of membrane polarization, which discriminate the type of cognitive activity taking place in any particular neuronal assembly. In what follows, therefore, all of these mental activities will be considered collectively.

In the MP framework just outlined, the challenge for brain modeling is to robustly identify teleological functions at the level of the organism (the human person) which can then be mechanistically described at the next lower level in terms of measurable interactions between well-defined components (in Craver's sense) of the organism.

Traditional psychology offers a motley collection of candidates for such functions, some more firmly rooted in reproducible and cross-culturally valid experimental methods than others. A preferable alternative would be to formulate a cognitive ontology from the set of brain tasks that can be inferred from consideration of task-specific brain networks. To ensure that a wide range of human behavior could be mechanistically explained using such a set of functions, a very large number of contrasting tasks would need to be explored, while recognizing the computational difficulties arising from the frequent overlaps (Anderson, 2014) of such networks. It must be considered an empirical question whether the resultant enormous set of components can possibly form a tractable basis for mechanistic explanation. If there is substantial clustering of networks, such that they can be grouped into a relatively small number of types, the types themselves might be usable. Evidence that this might be the case has been provided by the study of so-called "resting-state networks" (Beckmann et al., 2005; Biswal et al., 1995), observable using functional MRI with subjects instructed not to perform any particular task. Independent component analysis (ICA) has found that no more than about 100 temporally independent

spatial networks capture most of the variance in the fMRI time course data. These networks can be roughly associated with brain areas shown using PET or fMRI to support action competences, such as visual perception, limb movement, and auditory perception (Smith et al., 2009). Whether any of these definable and relatively reproducible networks can be mapped onto more psychological concepts such as intention, volition, agency, attention, and emotional valence remains unclear. This approach leads straightforwardly toward the concept of "representational models" (Kriegeskorte and Diedrichsen, 2016) which will be discussed more fully below.

Recently, the concept of "affordance," introduced by Gibson (1966, 1979) has been considered in regard to itemizing the brain's competences. Anderson (2014) summarizes this approach: "the brain should be understood as an action-oriented system" and "perception should be thought of as the assessment of the values of salient organism-environment relationships and the detection of opportunities for changing those values through action."

This conceptualization converges with that provided by the veteran visual neuroscientist Purves et al. (2015), who argues that "the solution (to understand visual perception) depends on: (1) rejecting the assumption that the goal of vision is to recover, however imperfectly, properties of the world; and (2) replacing it with a paradigm in which perceptions reflect biological utility based on past experience rather than objective features of the environment." If the brain can be regarded as the organ that prepares the organism for what happens next, par excellence, this formulation appears to make a great deal of sense. The question then becomes one of assigning identifiable competences to brain areas or networks that work together in performing this fundamental task: in short, discovering neural embodiments of learned affordances. We shall return to this topic in a later section, dealing with the themes of population receptive fields, voxel encoding, and deep learning.

2.3 PREDICTION AND PREDICTIVE CODING

Whatever choices are made regarding a useful listing of teleological functions, one guiding principle is relevant. In addition to receiving input from neural sensors and controlling motor outputs, any assemblage of neurons complex enough to be described as a brain must have the capability of predicting future events, and planning responses accordingly. It is obvious that such a competence adds greatly to the survival potential of the organism, and it can easily be argued that the capacity of organisms to remember previous experience (usually by virtue of having a brain) is contingent on the biological need to predict what may happen next. For instance, identification of objects of sensory experience can be performed with high efficiency when neural assemblies can rapidly compare the afferent signals arriving from sensory organs (such as the eye) with a restricted set of possibilities stored in memory—Bayesian priors—and report accordingly to other brain areas even when only a small number of discriminating signals has been received. This approach, first hinted at by Gregory (1966) has been formalized under the description of "predictive coding" (Rao and Ballard, 1999), and it has been embodied in neural network models, such

as recurrent neural networks (Bastos et al., 2012; Bitzer and Kiebel, 2012). Brain activity that can experimentally be demonstrated to provide predictive coding can clearly be considered as teleological, in the MP sense of the term. From this perspective, elaborating a list of domains of experience in which predictive coding can be observed to operate—perception and action-related—might greatly assist the development of plausible brain models.

In this context, the concept of prediction error, as a brain signal of importance to the organism, has been explored in some depth. Friston (2010) has formulated a global model of the relationship of a brain to its environment, which gives a central place to this concept, which is equivalent to surprise, using the statistical mechanical metaphor of "free energy." He claims that many measurable characteristics of brain function can be viewed as a process of minimization of free energy. Whether this formulation is essentially tautologous, vacuous, or can provide genuinely mechanistic modeling is not yet clear, largely because Friston's approach to systems neuroscience neglects detailed neuroanatomy, and thus fails to relate measurable brain activity to its observable neural substrate.

2.4 CONNECTIONISM AND SPATIAL MODULARITY

The intense neural connectivity mentioned above encouraged many researchers (Hinton et al., 1986; Lashley, 1929) to reject the idea of specialized brain regions for specific tasks, arguing instead that "each entity is represented by a pattern of activity distributed over many computing elements, and each computing element is involved in representing many different entities" (Hinton et al., 1986, p. 77). This activity, parallel distributed processing, leads to a theoretical perspective known as connectionism, in which it makes little sense to divide the brain into specialized functional regions. However, while a brain based on such principles can be imagined and modeled, there is enormous empirical evidence that many brain areas do not have the equipotentiality called for by strict connectionism, instead showing quite distinct preferences for specific types of task or stimulus. Much of this evidence comes from invasive experiments on a wide range of animals, where highly localized lesions of gray matter cause highly specific functional deficits. Much has also been inferred from comparable lesion studies in humans, where brain damage has been measured using structural imaging techniques such as MRI, and compared directly with specific task performance (Bates et al., 2003).

Moreover, there are now thousands of studies using the methods of BOLD contrast functional MRI that show striking spatial variations in brain activity for a great range of tasks and stimulus presentations. Using analysis techniques that have become standardized and appropriate thresholding of the resulting t-maps, cognitive subtraction studies often show relatively small numbers of activated cortical or subcortical areas, with well-defined boundaries. Such findings appeared to be entirely incompatible with connectionism, until the results of recent careful studies were reported. These studies attacked the problem of the relatively low signal to noise available in functional MRI, by the straightforward means of repeating the

experiment many times with the same volunteer subjects. The researchers avoided the use of the more popular analysis methods, concatenating about 10 h-worth of single-subject functional MRI data obtained across many imaging sessions (Gonzalez-Castillo et al., 2015). Strikingly, they found that a significant time-locked response to the quite simple task paradigm performed by the subject could be detected in most of the voxels in the brain. Thus, the tidy localization of function found in studies using poorer data may sometimes be explained by the choice of analysis method. Absence of evidence is not evidence of absence.

Nonetheless, fMRI researchers very often observe dramatic spatial variations in the amplitude of functional activity, even amounting to reversals of sign, for instance in center-surround suppression in the visual cortex (Shmuel et al., 2002). Boundaries of functionally discriminable brain areas can be experimentally defined and are normally conserved across a range of relevant stimuli. Such boundaries can even be precisely delineated in task-absent BOLD functional MRI studies (see Glasser et al., 2016). Clearly an extreme connectionist view in which the function of a given area of cortex is entirely dependent on the behavioral task has no experimental justification. Neuroanatomical constraints play a crucial role in the range of competences that can be deployed by local neural assemblies. How important more subtle time-locked variations in brain activity might be in explaining task performance still remains to be explored. To a first approximation, statistically thresholded data using relatively short periods of data acquisition may be enough to provide robust brain modeling—but this is unlikely to be achieved unless the level of spatial granularity of the data is appropriate.

2.5 ISSUES OF GRANULARITY: LEVEL OF EXPLANATION AND CONSISTENCY

The term "granularity" is used in physics to refer to the level of detail at which a phenomenon is described and analyzed (Gell-Mann, 1995). The natural granularity, or graininess, of the known universe is defined by the size of Planck's constant h, the minimum quantity of action that may be exerted by anything on any other thing. Explanations and models at this level of granularity, when they can be formulated, are apparently the best that one can hope to achieve. However, due to the physical and chemical properties of the world we live in at the surface temperatures of our planet, we have some natural kinds, such as the physical elements, with their chemical interactions, and consequently objects exist. Objects have size and mass, and often well-defined boundary surfaces.

This has meant that many satisfactory explanations have been established, dealing with the behavior of objects and physical fields, at a much coarser level of granularity than Planck's constant. Sciences such as geology and neuroanatomy rely on the possibility of unambiguous and reliable identification of durable components of systems whose behavior can be predicted using mechanistic modeling and measured using reproducible equipment.

As regards brain science at the mesoscopic scale, granularity becomes a crucial question. Here the concept of level needs to be discussed. Craver (2007, p. 189) argues strongly from his ethnographic evidence that the components of a successfully conceived mechanism can be analyzed into lower-level components, which are "organized together" to form the higher-level components. He asserts that this analysis need not be on the basis of object size, but on the basis of behaving components that are unified by their organization in an activity, which may not indeed correspond to individually definable objects. The levels at which these components are local, defined only within a given compositional hierarchy.

Such a broad characterization of levels of mechanism offers considerable scope for theorists of brain function. However, certain restrictions must still apply. Craver (2007, p. 131) states that the parts of a successful mechanism "have a stable cluster of properties, they are robust, they can be used for intervention, and they are physiologically plausible in a given pragmatic context." This last characteristic can be restated as a condition of consistency or appropriateness: the component must be comprised of known elements at a lower level that plausibly can be organized together to form the higher level component, and at this level causes can be proposed that plausibly can have effects that are measurable at the higher level. Considering the brain as a system, it would seem appropriate to define its elements as those parts that have an unambiguous neuroanatomical identity.

Before providing a list of such components that fulfill the requirements for mechanistic modeling of brain function at the systems level, it is worth briefly reviewing the types of object that make up brains.

3 NEUROANATOMY

3.1 THE NEURON DOCTRINE OF CAJAL

The neuron doctrine is the universally accepted concept that the nervous system is made up of discrete individual cells, the neurons, supported by astrocytes and by other glial cells. This discovery was due to the brilliant neuroanatomical work of Santiago Ramón y Cajal (Finger, 2000). By comparison with all other types of cell, what is remarkable about neurons is their ability to communicate with each other over long distances (compared with the size of their cell bodies) via the highly extended neurites, or neural processes, which carry pulses of electrical voltage (action potentials). These electrical pulses generally cause changes in the neurons that receive them, and can thus be formally regarded as signals. There are two types of neurites, dendrites and axons, where dendrites mainly convey action potentials toward the parent cell body (or "soma") and axons conveying action potentials away from the soma toward other neurons, at which they terminate, forming synaptic junctions with dendrites belonging to the target neurons.

Neuroanatomists have defined about 100 types of neuron that can be found in the human brain, two important primary classes being excitatory pyramidal cells and

inhibitory interneurons. It has long been known (Brodmann, 1909) that the cerebral cortex shows macroscopic areas with fairly uniform organization of neurons, typically into six layers within the 2–4-mm thick cortex. The local pattern of neuronal organization is described as "cytoarchitecture." More than 45 such distinguishable areas have been identified in cadaver brain tissue.

3.2 MYELIN

It is a feature of many axons that over the course of development they become surrounded with a multilayer sheath of phospholipid membrane known as myelin. This sheath dramatically increases the velocity of action potentials, confers greater physical strength to the axon, reduces its energy requirements, and provides electrical insulation against cross-talk with other neurons. In the gray matter of the brain axonal myelination also inhibits the formation of new synapses. The formation of myelin is often driven by the organism's experience, whereby the passage of action potentials along unmyelinated axons stimulates neighboring specialized glial cells to begin the process of myelination.

The axonal pathways within the white matter of the brain are normally heavily myelinated by the time a human is about 3 years old, which gives rise to the whitish color of this tissue. Within the cortex, different cortical areas show layer-specific variations in the degree to which axons are myelinated, some prefrontal areas being very lightly myelinated even in adulthood, while primary sensory and motor areas are generally very heavily myelinated. The pattern of myelination is termed the myeloarchitecture (Flechsig, 1920; Vogt and Vogt, 1919). Typically, myelinated fibers within the cortex are either radial to the cortical surface or run tangentially in well-defined layers—the so-called "bands of Baillarger." Recently, Micheva et al. (2016) have discovered that a majority of myelinated cortical axons connect to inhibitory synapses.

It is widely believed (Nieuwenhuys, 2013) that the boundaries of distinguishable cytoarchitectural areas correspond closely with those of myeloarchitectural areas. Nieuwenhuys et al. (2015) have recently transferred the areas discovered by the Vogt laboratory (1919) onto a standard 3D reference brain in digital format, in concert with the ongoing massive research effort to map human cortical areas in vivo using their MRI-visible myeloarchitecture (Bazin et al., 2014; Dinse et al., 2015; Geyer et al., 2011; Sereno and Huang, 2014; Tardif et al., 2015; Turner and Geyer, 2014a,b, Van Essen and Glasser, 2014; Waehnert et al., 2016). Interestingly, about 190 myeloarchitecturally distinct areas were already reported by Vogt and Vogt, a number which closely corresponds to that recently published by Glasser et al. (2016) using a range of incommensurable imaging techniques. Glasser et al. have claimed that their discovery of additional areas, beyond 47 areas mapped by Brodmann, is novel, but they fail even to refer to the work of Oskar and Cecile Vogt.

While the cerebral cortex everywhere has the same general appearance and requires detailed examination to reveal distinct areas, this is not always the case for the subcortex—those areas of gray matter deep in the human brain which usually

have homologous counterparts in other mammalian species. Here the hippocampus, the amygdala, the basal ganglia, and even the tiny habenula can be easily distinguished from neighboring tissue. Brain nuclei in the pons and medulla can also be easily identified in cadaver brain. Other regions such as the thalamus and hypothalamus are somewhat less easily discriminated, but histology has shown these to be composed of many subnuclei. Typically, the mesoscale parts of the subcortex have spatial dimensions of millimeters to centimeters. A taxonomy of these structures (Alkemade et al., 2013) gives a total of 455 named nuclei, of which perhaps 10% may be identifiable using current MRI techniques.

3.3 COMPACTNESS OF BRAIN AREAS

Largely due to how they develop, according to the fundamental biophysical rules of brain morphogenesis (Nieuwenhuys and Puelles, 2016), many anatomical components of the brain have identifiable boundaries. Although these can vary widely in shape—the long corticospinal tract of efferent and afferent neurons stretches tens of centimeters between the somatomotor cortex and neurons in the spinal cord, while the brainstem's red nucleus is an almost spherical ball about 5 mm across—neurons with similar functions appear to be often grouped together. Within areas of relatively uniform neural anatomy, maps have been discovered (Sereno et al., 2013)—in visual cortex, spatial maps of the visual scene; in auditory cortex, maps of sound frequency; in motor and somatosensory cortices and in cerebellum, spatial maps of the body; and in the caudate nucleus, maps of the level of abstraction (Mestres-Missé et al., 2012). The existence of such maps within identifiable areas may be regarded as a useful clue in the task of listing teleological functions, and thus brain model building, and the fact that these maps themselves have boundaries provides reassurance that the brain areas where they are located can be used as model components. For instance, excellent congruence has been found between the borders of primary visual cortex V1 as defined by structural MRI visualization of its unique anatomy (see below) and the spatial map provided by the functional MRI signal in the occipital cortex excited by sweeping a visual stimulus across the visual field (Bridge et al., 2005; Sanchez-Panchuelo et al., 2012).

3.4 NETWORKS AND CONNECTIVITY

3.4.1 Intracortical networks

Characteristically, brain areas contain populations of excitatory pyramidal cells and various types of interneurons, which may have inhibitory or excitatory synapses onto the pyramidal cells. This cortical circuitry is now the subject of intense study at nanometer resolution, using a range of techniques (Denk et al., 2012; Markram et al., 2015). Networks are modified as the result of experience (Chandrasekaran et al., 2015) as dendrites sprout and shrink, and new synaptic connections are made and lost. But the basic architecture is laid down during development, before and after birth. Extensive research with laboratory animals has revealed area-specific patterns

of connection, in which neurons in layer IV very often receive projections from thalamic nuclei that transmit sensory signals from the peripheral nervous system, and neurons in other layers receive and send signals to other cortical areas or subcortical nuclei. As will be discussed further on, if brain activity could be measured at the relevant spatial scale, such circuit diagram information may be highly valuable for building mesoscale mechanistic models.

3.4.2 Systems level

As already remarked, long-range brain connections are immensely complex. The Human Connectome programme (Van Essen et al., 2013) has absorbed substantial funding and manpower in order to make progress in understanding the macroscopic wiring diagram. The two types of information that are currently sought are the brain's "structural connectivity," discovered using diffusion-weighted MRI (Turner et al., 1990; Wedeen et al., 2012); and its "functional connectivity," which reveals itself as spatial correlations within time series of functional images obtained while no overt task is performed (so-called "resting-state fMRI," Biswal et al., 1995). The chief virtue of these techniques is that they can be applied in a living human subject, so that brain functional activity can be correlated directly, subject by subject, with its neuroanatomical substrate and putative neural network. For genetic, developmental, nutritional, environmental, and experiential reasons, all human brains differ somewhat in their anatomy and organization. In healthy brains, the main anatomical divisions are well conserved, but the shapes and sizes of the components can differ strikingly. To develop realistic models of brain function some precision regarding pathways is surely required, and so connectivity data drawn from a different brain, perhaps that of a cadaver or an average brain summarizing data from many subjects, must be regarded as insufficient. Mechanisms only operate in living individual subjects.

However, neither of these methods (tractographic) for discovering brain connections is close to a gold standard. The first method, using diffusion-weighted imaging (DWI) to infer axonal pathways in white matter, has very serious limitations due to the relatively poor spatial resolution (at best 1.5 mm using conventional 3 T MRI scanners) and the difficulty of discriminating the directions of different axonal fiber bundles in the frequent regions where they cross each other (Jones et al., 2013). These issues are to a large extent problems of inadequate data, and improvement is found when more specialized MRI equipment is used—higher field strength, such as 7 T (Heidemann et al., 2012a), and more powerful magnetic field gradient coils (Setsompop et al., 2013). Nevertheless, there may well be configurations of axonal pathways that MRI techniques can never successfully trace (Thomas et al., 2014). By comparing their high-resolution DWI animal brain data with the ground truth derived from axonal chemical tracer techniques, these authors show that the best result one can achieve is a compromise between substantial numbers of false-positive tracts and false-negative tracts, depending on the statistical thresholds used in computing the pathways. The best possible information about a typical cadaver brain, obtained using gold standard methods such as the use of tracers, can help to resolve ambiguities in the DWI data.

Wedeen et al. (2012) have proposed basic principles of white matter organization, in which fibers are arranged in sheets, which cross close to orthogonally and can be grouped roughly into posterior/anterior, left/right, and superior/inferior categories. Evidence for these concepts can be found from developmental neuroanatomy (Nieuwenhuys and Puelles, 2016). If they can be fully confirmed by experiment, the task of establishing the connectome for each individual human brain will become much easier.

The second method, functional connectivity, is more problematic at the conceptual level, although the raw data are relatively easy to obtain in a few minutes of scanner time. As yet there are no firmly grounded explanations of the observed phenomenon that regions of the brain which are spatially separate exhibit strong temporal correlations at time scales of a few seconds. Thus inference of connections from resting-state fMRI (rs-fMRI) data remains speculative. For fairly obvious reasons, the DWI-based connectome is not likely to match the rs-fMRI connectome very precisely, so the one cannot be used to cross-validate the other.

In summary, at this point in time, many of the major white matter pathways can be identified in individual living human brains, and some of these can be tracked with confidence into their terminations in gray matter. Methods are under development to quantify the uncertainty of the location of these terminations. (Note that probabilistic tractography (Behrens et al., 2003) does not provide this information—see discussion in Jones et al., 2013.)

4 MRI, BRAIN FUNCTION, AND NEUROANATOMY

Since the discovery of functional MRI in the early 1990s (Kwong et al., 1992; Ogawa et al., 1990; Turner et al., 1991), MRI has progressed from its earliest role as a radiological diagnostic modality to become an indispensable tool for cognitive neuroscience. Perhaps not surprisingly, the excitement that functional MRI could noninvasively provide usefully high spatial and temporal resolution maps of human brain activity was balanced by a comparatively lesser degree of interest in MRI's ability to provide excellent details of individual brain anatomy. Normal neuroanatomical differences tended to be regarded as nuisance variables, and techniques were developed to warp individual brains into the same template fiducial brain, assuming that once the gross anatomical differences between brains were taken care of, the positions of their functionally specialized regions would more-or-less coincide. But with the introduction of higher magnetic field MRI scanners, especially at 7 T, submillimeter spatial resolution became available, and fine anatomical details have been demonstrated with astounding clarity and precision, to the point that in vivo MRI images can sometimes bear a close resemblance to stained histological sections. Turner has summarized progress in delineating these neuroanatomical details in an encyclopedia article (2015), from which the following paragraphs are excerpted.

4.1 SOURCES OF MRI CONTRAST

MRI sequences sensitive to the longitudinal relaxation time T1 of water protons provide brain images with excellent gray/white matter contrast. The pioneering work by Koenig (1991) and Kucharczyk et al. (1994) showed that the large difference in T1 between gray matter and white matter could be explained almost entirely by the greater concentration of myelin in white matter. Furthermore, this difference was shown to arise very largely from specific components in myelin—not from proteins and phospholipids but instead from the abundant membrane lipids cholesterol and cerebroside.

Not all longitudinal (T1) relaxation can easily be explained in this way. Paramagnetic contrast agents, such as manganese chloride and ferritin, relax spin magnetization through the local rapidly varying magnetic fields with frequency components at the Larmor frequency that they provide due to thermal vibrations. Where such materials occur naturally in brain tissue, for instance, ferritin in the basal ganglia and in cortical myelin (Fukunaga et al., 2010), they will also contribute to T1 relaxation. Recent work by Stüber and coworkers (Stüber, 2014) comparing quantitative MRI of T1 and T2 with proton beam microscopy of the elemental composition of slices from the same block of cadaver brain tissue has revealed that T1 depends only weakly on iron content but strongly on myelin content…. (Thus) measurements of the relaxation rate R1 (1/T1) provide a quite reliable guide to the amount of myelin present in a voxel of healthy cortex or white matter, with rather less confidence in iron-rich deep brain regions such as the putamen, red nuclei, and substantia nigra.*

In recent years … attention has been paid by researchers to the assessment of myelin density using the techniques of MTR and imaging of myelin water fraction (MWF). MTR uses the phenomenon, discovered by Wolff and Balaban (1989), that excitation of protons attached to large molecules within the tissue causes more rapid relaxation of neighbouring free water proton spins. Because this process also depends on good coupling between the relevant macromolecules and free water molecules, one can expect a very similar spatial dependence to that of T1. The highly relaxing molecules of cholesterol and cerebroside are again likely to dominate (Ceckler, 1992; Koenig, 1991). Maps of MTR (e.g., Dortch, 2013) appear to be almost identical to maps of R1 (Weiskopf, 2013). Efficient MTR mapping has been developed extensively by Helms et al. (2010), but the signal to noise available per unit time is inevitably lower than that achievable with quantitative T1 mapping methods, such as MP2RAGE (Marques et al., 2010).

But both of these methods may turn out to be less accurate for estimation of myelin content than the third possibility, MWF (MacKay, 1994; Zhang, 2015). Myelin water fraction mapping uses T2 relaxometry, in which spin-echo images are acquired over a wide range of echo times. The data are analyzed to reveal a spectrum of T2 values, which cluster into relatively narrow peaks that can readily be identified as arising, respectively, from relatively free extracellular and

cytoplasmic water and water trapped between the layers in the myelin sheaths. The partial volumes of these compartments can be estimated quite accurately, and one can be confident that the myelin water fraction corresponds closely to the total amount of myelin. When the maps that are expected to correspond to myelin density acquired by each of these methods are compared, a problem appears. Maps of MWF show much greater inhomogeneity throughout the white matter than do T1 maps or MTR maps. Histograms of T1 and MTR are very sharply peaked at values corresponding to white matter, with a variance of perhaps 20%, but MWF values can vary systematically by a factor of up to two from region to region (e.g., Zhang et al., 2015). Clearly, these methods are measuring different aspects of myelin. However, once a suitable model has been developed, combining results from MWF and T1/MTR mapping may offer a simple opportunity to infer maps of other aspects of white matter myelin, such as the mean number of wraps and perhaps the mean axonal diameter.

In addition to mapping myelin noninvasively, MRI is also a powerful means of mapping brain iron (Drayer et al., 1986). The putamen, the substantia nigra, and the habenula (Strotmann et al., 2014) are rich in iron and are distinguished using MRI sequences which use a gradient echo to prepare the signal for acquisition. Such sequences deliver maps of the MRI parameter T2*, which describes the speed of the free induction decay, and also quantitative maps of magnetic susceptibility (QSM) (Marques and Bowtell, 2005) each of which provides interpretable information regarding the concentration and the spatial distribution of iron atoms within brain tissue. Ferritin is also often colocated with myelinated axons within the cortex, hence iron-sensitive images can have a synergetic effect with myelin in discriminating cortical areas (Duyn et al., 2007).

4.2 SPATIAL RESOLUTION

Returning to the question of granularity, it is important to recognize the spatial scale at which useful information can be available using the conveniently noninvasive techniques of MRI. The spatial resolution achievable in human and other animal brains depends on the desired signal-to-noise ratio (SNR), the magnetic field strength, the quality of radiofrequency receiver coils used, the type of MRI pulse sequence employed, and the scan duration.

To provide interpretable data, the SNR normally must exceed 10, and for some purposes a value of 100 or greater is needed. The SNR increases with magnetic field, with an exponent measured at 1.65 (Pohmann et al., 2016). Currently, the highest magnetic field strength that is routinely available for scanning human brain in vivo is 7 T, although good results are starting to be achieved using 9.4 T. However, engineering and safety challenges discourage the use of field strengths above 9.4 T for systematic imaging neuroscience purposes.

The good performance just mentioned depends critically on high quality radiofrequency hardware. In particular, receiver coils comprising many small elements

(up to 128), each separately connected to the required preamplifiers and filtering circuits, can provide enhanced performance. Such multielement coils also allow much more rapid data acquisition because the inhomogeneous receptive fields of each coil can provide additional spatial data for image formation (Setsompop et al., 2016).

To take fullest advantage of the intrinsic SNR of the MRI scanner, the pulse sequence used to acquire image data must be wisely chosen. The current trend is to select sequences that provide quantitative maps of the tissue parameters accessible to MRI: the water proton density, the longitudinal and transverse relaxation times, the effective diffusion constant of water molecules, the magnetic susceptibility, and the voxel-by-voxel volume of cerebral blood. Sequences with a relatively high spin flip angle, and which use most of the scanning time in performing data acquisition, give the highest SNR. This indicates the use of high spatial resolution, high flip-angle techniques such as echo-planar imaging (EPI) and gradient echo and spin echo (GRASE) imaging, both for functional brain imaging (Goense et al., 2016; Heidemann et al., 2012b; Huber et al., 2015) and also for structural imaging (Renvall et al., 2016; Trampel et al., 2014).

Scan duration is also a factor in voxel size. In principle, given enough time, most MRI scanners can produce images with better than 100 μm resolution, but to harvest enough SNR at such a resolution can take several days of scanning, feasible only with cadaver brain tissue. With advanced methods of avoiding effects of head motion during scanning (Zaitsev et al., 2006), the option has become available of splitting the acquisition into several sessions, to improve the SNR, but so far this has been little used. More typically, a total of about 90 min scanning time is seen as acceptable for human volunteer subjects. To achieve the best results, the blurring effects of involuntary head motion, even as little as a fraction of a millimeter, must be corrected. Several methods have now been successfully demonstrated. Markers can be attached to the head, and tracked with infrared cameras (Schulz et al., 2012; Zaitsev et al., 2006), and the resulting movement parameters can be fed back to the scanner prior to each excitatory radiofrequency pulse, enabling an update of the slice positioning to follow the head and thereby providing prospective motion correction. The received NMR signal from the scalp fat itself carries information about head movement, which can be used for the same purpose (Federau and Gallichan, 2016).

The current resolution capability in human brain at the field strength of 7 T stands at about 300 μm isotropic resolution for structural brain images (Federau and Gallichan, 2016), and about 500 μm resolution for functional brain images. Measurement of water diffusion in the brain, useful for assessing brain connectivity, has intrinsically lower sensitivity, and the best resolution yet achieved in living human brain is about 800 μm (Heidemann et al., 2012a).

Crucially, for these three variables, structural, functional, and connective, the resolution is substantially finer than the cortical thickness. This level of technological progress constitutes a tipping point (in our view), a breakthrough in the feasibility of realistic and effective modeling of human brain function. Sections 4.3 and 4.4 describe why imaging at this spatial scale is so important.

4.3 IN VIVO HISTOLOGY

Recalling the earlier remarks about cortical cytoarchitecture and myeloarchitecture, it is important to stress that the features characterizing the identity of specific patches of cortex are generally discernible on the scale of a few tens of microns. While higher resolution (better than 1 μm) is normally used for cortical parcellation with microscopy of cadaver brains, the cytoarchitectonic techniques used by Zilles (Schleicher et al., 1999) and the Jülich neuroanatomy laboratory effectively take local averages of neural density, comparing the moments across the cortical thickness of this "gray value" density along the cortical surface in order to detect discontinuities that betoken a cortical area boundary. While the myeloarchitectonic distinctions elaborated by Vogt and his team (1919) relied on qualitative changes of the proportion of radial and tangential fibers, and the extent to which the radial fibers penetrate the cortex as they course outward from the white matter, most of these features can also be visualized at a spatial scale of tens of microns. The most salient feature visible in the entire cerebral cortex is the heavily myelinated Stria of Gennari, discovered by the Italian neuroanatomist Gennari in 1782 (Gennari, 1782), which forms a sheet located halfway through the cortical thickness in the primary visual cortex, located in the posterior occipital lobe. This feature can be seen with the naked eye. Thus it need not be regarded as hopelessly ambitious to consider many cortical areas to be distinguishable in vivo with MRI, provided that this technique can be made sensitive to the relevant markers.

Section 4.1 described the striking capability of MRI for delineating patterns of myelination and tissue iron. Together, without further technical developments, mapping of these substances can discriminate dozens of brain areas in vivo, provided that spatial resolution is better than about 0.5 mm (Alkemade et al., 2013; Deistung et al., 2013; Tardif et al., 2015). Myelinated cortical layer structure has been observed using MRI in primary visual cortex (e.g., Trampel et al., 2011), extrastriate visual areas V5/MT and V3a, anterior cingulate cortex, and primary somatosensory cortex. Differences in their layer dependence of myelination is very likely to greatly extend the number of areas identifiable in vivo using myelin mapping, beyond those enumerated by Glasser and Van Essen (2011), Sereno and Huang (2014), and Tardif et al. (2015).

To extend the classification of brain tissue yet further, further strategies have been suggested. Use of diffusion-weighted MRI techniques can provide a rough measure of dendritic density (Jespersen et al., 2007), which has been implemented in the form of the NODDI method (Zhang et al., 2012). Nagy et al. (2013) and Morris et al. (2016) have shown that diffusion imaging can also parcellate the cortex. Cortical layers can be detected in cadaver brain using very high-resolution tractography (Leuze et al., 2014), and recent efforts by Aggarwal et al. (2015) have revealed region-specific diffusion signatures for Brodmann areas 9, 4, 3b, 17, and 18. Cerebral blood volume varies between cortical areas (Michaloudi et al., 2005; Zheng et al., 1991), and it can be straightforwardly mapped using MRI techniques (Uh et al., 2009). Measurements by LaManna et al. (1992) and others of capillary angiogenesis

show that cortical capillary density is driven by demand, through the secretion of vascular endothelial growth factor when brain tissue becomes hypoxic. However, no systematic effort has yet been made to employ CBV mapping for cortical parcellation.

To summarize, at the present time, one can be cautiously optimistic that many more of the cortical areas identified in cadaver brain tissue by Vogt, Brodmann, and more recently the Zilles laboratory in Jülich, will be observed using structural MRI methods in the brains of living subjects.

4.4 FUNCTIONAL MRI, LOCAL CONNECTIVITY, AND CAUSAL DIRECTIONALITY

The main purpose of the study of neuroanatomy is to provide a basis for the understanding of brain function. It is not enough to tease out the detailed structure of some part of the brain: it is vitally necessary to discover what task it performs. Thus functional mapping of the brain should aspire to the same level of detail as the structural images obtained. At the field strength of 7 T, as previously mentioned, a voxel size of 0.5 mm is practicable for functional BOLD images, which is not much greater than the 0.35 mm voxel size for anatomical images.

Contrast in BOLD images arises from variations in the blood concentration of deoxyhemoglobin, which is more paramagnetic than surrounding tissue and therefore decreases the MRI signal. This occurs when the signal is formed by a gradient echo, or by a spin echo with a sufficiently long echo time. There is debate on the comparative merits of GRASE BOLD (Sanchez-Panchuelo et al., 2015). Spin echo appears to provide better localization, but at the cost of signal to noise (SNR). The gradient echo BOLD signal is sensitive to oxygenation changes over a wide range of vein diameters, even in pial veins lying on the surface of the brain some distance downstream from the neuronal tissue supporting the activation (Turner, 2002). Thus its dependence on cortical depth is not easy to interpret as resulting from layer-dependent neural activity (Heinzle et al., 2016; Markuerkiaga et al., 2016). Greater spatial confidence can perhaps be placed in functional MRI data using the spin-echo technique (Boyacioğlu et al., 2014; Goense and Logothetis, 2006), which is preferentially sensitive to oxygenation changes in smaller vessels.

Of especial note is an alternative technique known as VASO (vascular space occupancy), which provides a usefully quantitative measure of changes in cerebral blood volume (Lu et al., 2003), and is insensitive to changes in blood oxygenation. Besides the great advantage of being quantitative, allowing comparison within and across subjects and sites, the VASO functional signal appears to reflect the highly localized control of blood volume provided by pericytes (Hall et al., 2014), cells that surround arterioles and capillaries and regulate their blood flow (Hamilton et al., 2010). Thus in functional studies, the VASO signal is often maximal in the middle layers of the cortex, where the neural activity is maximal (Huber et al., 2015), as was found earlier using fMRI with iron-based contrast agents indicating blood volume (Goense et al., 2016; Zhao et al., 2006).

The revolutionary implications of the ability to evaluate cortical layer-dependent brain activity in vivo can be summarized as follows. Prior neuroanatomical knowledge of neuronal circuitry is often available, for each cortical area (e.g., Mao et al., 2011). Histology and animal brain research defines the cortical layer specificity of input and output pathways. For instance, input axons arriving from thalamic nuclei generally form their main synaptic connections with dendrites belonging to neurons in cortical layer IV. Activity in input layers can be driven by experimental conditions, and the effects of activity in output layers reveal themselves in experimentally observable behavior. The evidence of brain activity measured using functional MRI has been firmly associated with electrical activity taking place postsynaptically in the dendritic arborization (Logothetis, 2002). Whenever fMRI spatial resolution can discriminate input and output cortical layers, causal relationships between brain areas might thus be empirically validated.

Examples open to empirical experimentation include: motor imagery vs actual motion (primary motor area M1), visual imagery vs actual vision (primary visual cortex V1), auditory imagery vs actual hearing (primary auditory cortex A1), and touch experienced vs touch observed (a subset of primary somatosensory cortex, BA1).

Beyond the primary sensory brain areas for which thalamocortical inputs dominate, one can also consider cortical regions where the input layers of top-down and bottom-up afferents are spatially separated within the cortical thickness. Here also a causal direction might be established for the neural activity corresponding to a given task. There are many experimentally accessible instances: cross-modal vs unimodal stimulation, top-down vs bottom-up attentional modulation, and self-motion vs other-motion in mirror-system brain areas.

In their pioneering study, Trampel and coworkers (Trampel et al., 2012, Turner, 2016) measured activation in the hand area of human primary motor cortex. They used BOLD fMRI to study activation for three motor tasks: finger tapping, finger movement without touch, and motor imagery. The experimental design was based on the fact that the output to the corticospinal output tract of motor nerves from agranular primary motor cortex M1 arises almost entirely from large pyramidal neurons in layer V. The primary motor cortex was unambiguously identified by its anatomical location and high myelin content, indicated by its characteristically short T1. At 7 T structural data were obtained with 0.5 mm isotropic resolution, and fMRI data with 0.75 mm isotropic resolution. Cortical activation profiles specific to each motor condition were computed from the data, and averaged across the activated area at four different cortical depths, and across nine human volunteer subjects. The crucial finding from this study is that during the motor imagery condition, the BOLD signal at a depth corresponding to cortical layer V was comparatively smaller than the signal from other cortical layers in this condition. In this condition, with no motor output, this output layer would be expected to show a relatively lower activation signal—as was indeed observed. Fortunately, the cortical thickness in M1 is unusually large, about 4 mm, which facilitated the discrimination of specific cortical layers using fMRI.

In the more challenging context of primary visual cortex V1, Kok et al. (2016), also using gradient echo BOLD, have been able to discriminate layer-dependent activation associated with veridical visual perception from that related to illusory contors produced by the Kanitza illusion.

As mentioned earlier, however, the BOLD signal represents the history of blood oxygenation changes as blood travels from the pial arteries into the diving arterioles and hence into capillaries and veins. Its cortical profile represents a spatial convolution of task-driven changes in oxygen extraction with local blood flow, modulated by changes in blood volume, which blurs out the layer dependence of underlying neural activity. As such, this signal cannot provide a precise layer-specific indication of oxygen extraction. In the study just described, the statistically highly significant difference found in cortical profiles between the tapping and motor imagery conditions is noteworthy, but should not be overinterpreted.

More confidence can be placed in assessment of functional activity by means of observing changes in CBV, for instance using VASO. Cortical layer-dependent high-resolution VASO appears to capable of distinguishing afferent and efferent functional connectivity. In a very recent report, Huber et al. (2016) hypothesize that individual layers should show different resting-state signal fluctuations (Polimeni et al., 2010) and hence different brain connectivity patterns. They used VASO fMRI acquisition and a clever analysis method to measure layer-dependent resting-state fluctuations, to show directional functional connectivity from primary sensory cortex (S1) to primary motor cortex (M1), which was validated with task-based fMRI measurements. Such studies offer considerable optimism that the causal directionality of links between many brain areas can be objectively evaluated.

5 DATA-DRIVEN FUNCTIONAL CATEGORIZATION

Having described strategies for empirically identifying the material components of the brain treated as a system, and experimental techniques that offer the possibility of defining the causal links between these components, we return now to the task of defining their teleological functions, as discussed in Section 2.2. Given the existing confusion arising from ascribing functions drawn from traditional psychology, folk psychology, and psychoanalysis, it is attractive to replace these functional ontologies with a set of functions that is more data driven.

One possibility is to assemble a catalog of the transformational capabilities of cortical networks, once the anticipated results of the European Human Brain Programme have become available, perhaps an enrichment of the types of logical operations performed by silicon-based electronic components (AND gates, etc.). However, these data are not yet available, and in any case such an ontological catalog may still be far from useful in building the desired bridge between system and cell. It is becoming clear that a more practical opportunity for data-driven typologies of function can be generated using the methods of imaging neuroscience.

5.1 REPRESENTATIONAL MODELS

Representational models (Diedrichsen and Kriegeskorte, 2017) are intended to explain how activity patterns in populations of neurons (or, more generally, in multivariate brain activity measurements) relate to sensory stimuli, motor actions, or cognitive processes. Experimentally, in imaging neuroscience, representational models can be defined as probabilistic hypotheses about what profiles of activations across conditions are likely to be observed. At this point in time there are basically three methods to test such models—encoding models (otherwise known as voxelwise modeling, VM), pattern component modeling (PCM), and representational similarity analysis (RSA).

> *The hallmark of VM is an explicit model of representation, known as an encoding model. Formally, an encoding model proposes a set of sensory or cognitive features and specifies how these features are transformed into a prediction of brain activity for the experiment under consideration. A given set of features represents an explicit hypothesis about the representation encoded in the brain. This hypothesis is tested by evaluating how much variance in measured activity the encoding model explains. Competing hypotheses can be adjudicated by comparing the amount of variance explained by different encoding models. Alternatively, hypotheses can be assessed by comparing how well a representational similarity matrix (e.g., a matrix with correlations between pairs of experimental conditions) constructed from a set of features matches the representational similarity matrix constructed from the measured activity. This approach, called 'representational similarity analysis', imposes fewer constraints on the mapping between features and brain activity [12]. In both cases, hypotheses are tested by evaluating explicit models of representation.*
>
> **Naselaris and Kay (2015)**

This approach for modeling fMRI data typically begins with providing the experimental subject with a very large number of related stimuli or tasks, often naturalistic. These are analyzed to generate a large set of features, which comprise a training dataset then used to compile a separate model for each recorded voxel in the functional brain images. Essentially, the goal is to determine the functional repertoire of each gray matter voxel, as encompassed by a model that characterizes the "feature space" of the stimuli. The correctness and completeness of the model in predicting brain activity to new stimuli can be tested on a separate validation dataset (Naselaris et al., 2011). Remarkable cortical maps, for instance, depicting the space of semantic categories (Huth et al., 2016) have been generated using these methods. For such purposes, spatial smoothing would be quite unacceptable. Hence this technique lends itself to research in which myeloarchitecture, cytoarchitecture, and functional repertoire can be directly compared.

In the present formulation, the proposed (or inferred) sensory or cognitive features that best predict neuroimaging data can be considered to embody the teleological functions of the brain components that support them. These features themselves

can be suggested by the use of deep learning analysis (Marblestone et al., 2016), using multilayer neural network methods to provide parsimonious accounts of what usefully distinguishes one experience or action from another. The explanatory power of these features can be explicitly tested using neuroimaging data. In this framework, the catalog of such features can be regarded as minimizing some definition of cost—neatly conforming to the requirement of Piccinini that a teleological function must contribute to the survival of the organism. Cognitive neuroscience may benefit from deeper understanding of these data-driven insights into the categorization of experience and action, which may avoid the Procrustean tendency to force our experience into predefined inherited conceptual frameworks that may have little affinity with how brains actually operate (Turner, 2012).

5.2 POPULATION RECEPTIVE FIELDS (DUMOULIN AND WANDELL)

The population receptive field mapping approach (Dumoulin and Wandell, 2008) estimates a model of the population receptive field for voxels in visual cortex that best explains measured fMRI responses resulting from a series of various visual stimuli. This can be regarded as a special case of voxel encoding, applying specifically to visual stimuli and visual cortex.

Results using these approaches reveal that specific features of experience often have widely distributed spatial representations in the brain. However, clustering can also be noted, often in accordance with linguistic or common-sense categorization of experience and action. The features of lived experience—sensory, decision making, emotional, and rational—that emerge from this research strategy may well be best characterized as the affordances of such experiences—because it can be argued that affordances are what we are evolved to remember. And features which are affordances are good candidates for being teleological, in Piccinini's sense.

6 CONCLUSIONS: THE PROSPECTS OF LINKING CELL AND SYSTEM

In this section, we review the above-discussed criteria that characterize a satisfactory biological mechanistic model. We show that current modeling practice in systems neuroscience very often fails to meet these criteria, due to an absence of insight regarding the importance of neuroanatomy and an ensuing misguided strategy for image data preprocessing and statistical analysis. We summarize the proposed new strategy, now feasible using 7 T MRI scanners, for developing mechanistic models at the systems level: identifying the experimentally accessible components of the brain, which can incorporate prior knowledge of their cellular makeup; the links between these components; and the ascribing of teleological functions to these components. We conclude with a discussion of the extent to which a model formed in this way conforms to the recommendations of Piccinini and Craver.

The criteria for effective mechanistic modeling, described in detail in Sections 2.1 and 2.2, can be summarized as follows:

- **(i)** Mechanisms are multilevel.
- **(ii)** Mechanisms explain the behavior of the mechanism as a whole in terms of the organized activities of its component entities.
- **(iii)** The operations within a mechanism are different from the phenomena produced by the mechanism.
- **(iv)** The spatially, temporally, and actively organized components of a lower-level mechanism can constitute the components of a higher-level mechanism, which is itself something greater than a mere sum of these component parts.
- **(v)** A mechanistic explanation should account for the multiple features of the phenomenon, including its precipitating conditions, manifestations, inhibiting conditions, modulating conditions, and nonstandard conditions.
- **(vi)** Mechanistic brain modeling must robustly identify teleological functions at the level of the organism which can be described at the next lower level in terms of measurable interactions between well-defined components.
- **(vii)** Mechanistic models of the brain must be consistent with what neurons can actually do.

Current practice in human systems neuroscience was developed by leading imaging neuroscience laboratories in the late 1990s, and attempted to link brain location, neuroanatomy, and function at a spatial scale of about 8 mm, about as close as anyone then dared to expect that corresponding cortical areas could be located across brains. A strategy of spatial smoothing of the raw fMRI images, using a Gaussian smoothing kernel of typically about 8 mm, was fundamental to this approach. This had the following very important benefits: (a) it often considerably improved the signal to noise (SNR) of functional data; (b) after structural brain images had been spatially normalized into a standard template brain registered within MNI space, it allowed for the residual mismatch of actual cortical areas, so that positive results could be anticipated from group averaging across normalized brains; and (c) it enabled very simple analytic equations (Worsley et al., 1992) to be used for assessing the statistical significance of measured brain activity, and thus for thresholding the resulting group images to provide spatial activation maps.

However, this strategy has two severe failings (Stelzer et al., 2014). The spatial smoothing of functional imaging data prior to further analysis leads to a prevalence of false assignments of "activated" voxels to anatomical locations where there is no credible evidence of activity, and to a dramatic loss of knowledge regarding the relationship of brain activity to its neural substrate.

> *Numerous attempts have been made to develop mechanistic models (for instance Dynamic Causal Modelling, Friston et al., 2003) at the level of the statistical parametric maps generated by this analysis strategy. Lip-service is paid to the details of neural circuitry, which by default are considered to be homogeneous throughout the cortex, following the assumption that no further information*

can be obtained in vivo. The concept of 'neural mass modelling' is invoked to take care of the role of individual neurons—these are lumped together in their millions, and substituted by a single equivalent neuron at the centroid of the large fuzzy blob vaguely demarcating the site of the brain activity.

Recall that a successful mechanistic model requires definition of components, their separate functions, and the interactions between them. Because current models are multilevel only in intention, but ignore the details of individual brain neuroanatomy, a cursory glance at the above list of criteria for mechanistic models reveals that they are obviously incapable of satisfying most of them. To spell this out:

(a) The model components are typically ill-defined, relying on probabilistic atlases of brain anatomy produced using incompletely validated methods of cortical and subcortical parcellation of cell-stained histological slices from no more than 10 cadaver human brains. It is only occasionally possible to identify the activity detected using fMRI methods with a specific neural substrate. Much of the time, assignment of activity to named brain regions is no more than informed guesswork.

(b) The network of interaction between the components, the way in which they are organized, is also poorly defined. Researchers use putative connections derived from much older cadaver brain studies, connections derived from tractography diffusion-weighted MRI data, with a single-tensor approximation (Jones et al., 2013), and functional connectivity presumed from rs-fMRI data, often smoothed and averaged across many subjects. It is unlikely that the important connections in any particular subject's brain actually follow the course assumed by any of these strategies.

Furthermore, the directionality of any of these connections is very hard to deduce. Attempts have been made using the methods of Granger causality to infer directionality from the time course of the BOLD signal in different brain areas (Roebroeck et al., 2005), but the lag of several seconds between neural activity and the ensuing BOLD signal, which varies spatially, makes such inferences very hard to sustain (Smith et al., 2012). Dynamic causal modeling also attempts to deduce causal directionality, but this method relies for computability on a far too simplistic graph-theory model of brain activity, in which only a very small number of nodes participate for any given task (Lohmann et al., 2012). The precisely delineated features of human neuroanatomy are summarily neglected.

(c) In current common practice, the functional role attributed to the model components (to the extent that they are unambiguously defined) often exhibits the mereological fallacy, which consists of ascribing to portions of the brain psychological concepts that only make sense when ascribed to whole human persons (Bennett and Hacker, 2003). While studies of the effects of brain lesions on task performance (such as the classic injury to Broca's patient, leading to the idea that there is a delimited "language area") encourage the idea that mental tasks can be localized in specialized brain tissue, it is obviously wiser to try to

find less question-begging descriptors of brain areas until a mechanistic model is validated that demonstrates the entire system whereby a human being produces the behavior of interest. See criteria (iii) and (iv) in the list above.

So can the beautiful submillimeter data provided by much higher field MRI data provide surer ground for novel attempts to model the relationship between human behavior, linking the systems level, and the cellular level? Once again, we consider the structure of a successful mechanistic model:

Brain components: In vivo, at the mesoscopic scale, submillimeter resolution quantitative MRI (in particular mapping myelin and iron) can be used to parcellate the cerebral cortex into territories that can be associated with known cytoarchitecture. This parcellation can be guided by high quality histological studies and studies of chemoarchitecture. Subcortical structures can also be discriminated with high precision using MRI techniques. All subsequent fitting of models should be performed on an individual brain basis. Because the voxels are small enough to contain only thousands of neurons, rather than millions, the opportunity exists for relating the cortical structure within a voxel to its computational competence. Thus at the level of components, criteria (ii) and (iv) are fulfilled.

Links between components: In vivo tractography using crossing fiber analysis, guided by myelin stain histology and polarized light imaging in cadaver brain, can be used to establish the major long distance neural pathways. Histological and tracer studies in animals should be used to distinguish wherever possible the cortical layers where efferent and afferent pathways terminate, for each identifiable cortical area. Mechanistic models with causal relationships posited between the components can thus be proposed. Such links satisfy criteria (ii), (iv), and (vii).

Teleological functions of each component: As far as possible, large batteries of naturalistic stimuli for volunteer or clinical subjects should be used, together with cost-minimization deep learning algorithms to propose relevant feature spaces. The teleological functions for each brain component should be inferred from careful analysis of the imaging data, in which clustering and tightly cross-correlated activity play an important role. It should not be assumed that the inferred operations of each component can be described by any of the competences of the system as a whole. The conceptual framework of traditional experimental psychology should be considered as at best a source of inspiration. The granularity of different brain areas can be established empirically using submillimeter resolution fMRI. Causal directionality between regions with well-established histology can be investigated using layer-dependent VASO. All of the criteria listed are met when teleological functions are ascribed at this spatial scale and level of organization.

If it is now granted that use of the available in vivo, submillimeter resolution, neuroanatomical, functional, and connectivity MRI data allows finer grained and far more plausible mechanistic models of the operations of the human brain at the systems level, several questions remain.

(a) In regard to the level of detail, how much resolution is really required? The MRI best resolution achievable is smaller than the typical width of a cortical column.

But imaging neuroscience has been thriving with effectively 8 mm resolution, although the variance explained in fMRI data by current modeling attempts is often highly disappointing.

(b) Once techniques such as representational modeling and voxel encoding are applied more systematically, will it always be possible to create realistic feature spaces that reflect the self-organization of the brain resulting from probabilistic learning?

(c) How many anatomically distinct brain areas, observable in vivo, will be needed to provide relatively complete mechanistic models?

(d) How much improvement, in regard to model fits with greatly reduced variance, will be found when brain functional activity and neural substrate are much better aligned?

Addressing such questions, using models of brain function based on the proposed components, interactions, and teleological functions and thereby finding robust MRI-based biomarkers for observable abnormalities in patterns of cortical and subcortical competences, should provide much firmer scientific ground for the understanding of psychiatric disorders.

REFERENCES

Aggarwal, M., Nauen, D.W., Troncoso, J.C., Mori, S., 2015. Probing region-specific microstructure of human cortical areas using high angular and spatial resolution diffusion MRI. Neuroimage 105, 198–207.

Alkemade, A., Keuken, M.C., Forstmann, B.U., 2013. A perspective on terra incognita: uncovering the neuroanatomy of the human subcortex. Front. Neuroanat. 7, 40.

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

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. Neuron 76 (4), 695–711.

Bates, E., Wilson, S.M., Saygin, A.P., Dick, F., Sereno, M.I., Knight, R.T., Dronkers, N.F., 2003. Voxel-based lesion-symptom mapping. Nat. Neurosci. 6 (5), 448–450.

Bazin, P.L., Weiss, M., Dinse, J., Schäfer, A., Trampel, R., Turner, R., 2014. A computational framework for ultra-high resolution cortical segmentation at 7Tesla. Neuroimage 93 (2), 201–209.

Bechtel, W., 2006. Discovering Cell Mechanisms: The Creation of Modern Cell Biology. Cambridge University Press, Cambridge.

Bechtel, W., 2008. Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience. Routledge, London.

Bechtel, W., 2009. Constructing a philosophy of science of cognitive science. Top. Cogn. Sci. 1 (3), 548–569.

Bechtel, W., Richardson, R.C., 2010. Neuroimaging as a tool for functionally decomposing cognitive processes. In: Hanson, S.J., Bunzl, M. (Eds.), Foundational Issues in Human Brain Mapping. MIT Press, Cambridge, MA, pp. 241–262.

Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. Philos. Trans. R. Soc. Lond. B Biol. Sci. 360 (1457), 1001–1013.

Behrens, T.E., Johansen-Berg, H., Woolrich, M.W., Smith, S.M., Wheeler-Kingshott, C.A., Boulby, P.A., Barker, G.J., Sillery, E.L., Sheehan, K., Ciccarelli, O., Thompson, A.J., Brady, J.M., Matthews, P.M., 2003. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. Nat. Neurosci. 6 (7), 750–757.

Bennett, M.R., Hacker, P.M.S., 2003. Philosophical Foundations of Neuroscience. Blackwell Publishing, Malden, MA.

Bennett, M.R., Hacker, P.M.S., 2012. History of Cognitive Neuroscience. Wiley Blackwell, London.

Bentall, R.P., 2009. Doctoring the Mind: Why Psychiatric Treatments Fail. Penguin, London.

Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34 (4), 537–541.

Bitzer, S., Kiebel, S.J., 2012. Recognizing recurrent neural networks (rRNN): Bayesian inference for recurrent neural networks. Biol. Cybern. 106 (4–5), 201–217.

Boyacioğlu, R., Schulz, J., Müller, N.C., Koopmans, P.J., Barth, M., Norris, D.G., 2014. Whole brain, high resolution multiband spin-echo EPI fMRI at 7 T: a comparison with gradient-echo EPI using a color-word Stroop task. Neuroimage 97, 142–150.

Bridge, H., Clare, S., Jenkinson, M., Jezzard, P., Parker, A.J., Matthews, P.M., 2005. Independent anatomical and functional measures of the V1/V2 boundary in human visual cortex. J. Vis. 5 (2), 93–102.

Brodmann, K., 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. J.A. Barth, Leipzig.

Canino, G., Alegria, M., 2008. Psychiatric diagnosis: is it universal or relative to culture? J. Child Psychol. Psychiatry 49, 237–250.

Cartwright, N., 1999. The Dappled World. A Study of the Boundaries of Science. Cambridge University Press, Cambridge.

Chandrasekaran, S., Navlakha, S., Audette, N.J., McCreary, D.D., Suhan, J., Bar-Joseph, Z., Barth, A.L., 2015. Unbiased, high-throughput electron microscopy analysis of experience-dependent synaptic changes in the neocortex. J. Neurosci. 35 (50), 16450–16462.

Chemero, A., 2009. Radical Embodied Cognitive Science. MIT Press, Cambridge, MA.

Chiao, J.Y., Immordino-Yang, M.H., 2013. Modularity and the cultural mind: contributions of cultural neuroscience to cognitive theory. Perspect Psychol. Sci. 8 (1), 56–61.

Craver, C.F., 2007. Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience. Clarendon Press, Oxford.

Craver, C.F., Darden, L., 2013. In Search of Mechanisms: Discoveries Across the Life Sciences. University of Chicago Press, Chicago.

Deistung, A., Schäfer, A., Schweser, F., Biedermann, U., Turner, R., Reichenbach, J.R., 2013. Toward in vivo histology: a comparison of quantitative susceptibility mapping (QSM) with magnitude-, phase-, and R2*-imaging at ultra-high magnetic field strength. Neuroimage 65, 299–314.

Denk, W., Briggman, K.L., Helmstaedter, M., 2012. Structural neurobiology: missing link to a mechanistic understanding of neural computation. Nat. Rev. Neurosci. 13 (5), 351–358.

Diedrichsen, J., Kriegeskorte, N., 2017. Representational models: a common framework for understanding encoding, pattern-component, and representational-similarity analysis. PLoS Comput. Biol. 13 (4), e1005508. http://dx.doi.org/10.1371/journal.pcbi.1005508.

Dinse, J., Härtwich, N., Waehnert, M.D., Tardif, C.L., Schäfer, A., Geyer, S., Preim, B., Turner, R., Bazin, P.L., 2015. A cytoarchitecture-driven myelin model reveals area-specific signatures in human primary and secondary areas using ultra-high resolution in-vivo brain MRI. Neuroimage 114, 71–87.

Drayer, B., Burger, P., Darwin, R., Riederer, S., Herfkens, R., Johnson, G.A., 1986. MRI of brain iron. Am. J. Roentgenol. 147 (1), 103–110.

Dumoulin, S.O., Wandell, B.A., 2008. Population receptive field estimates in human visual cortex. Neuroimage 39 (2), 647–660.

Duyn, J.H., van Gelderen, P., Li, T.Q., de Zwart, J.A., Koretsky, A.P., Fukunaga, M., 2007. High-field MRI of brain cortical substructure based on signal phase. Proc. Natl. Acad. Sci. U.S.A. 104 (28), 11796–11801.

Eklund, A., Nichols, T.E., Knutsson, H., 2016. Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. Proc. Natl. Acad. Sci. U.S.A. 113 (28), 7900–7905.

Federau, C., Gallichan, D., 2016. Motion-correction enabled ultra-high resolution in-vivo 7T-MRI of the brain. PLoS One 11 (5)e0154974.

Finger, S., 2000. Chapter 13: Santiago Ramon y Cajal. From nerve nets to neuron doctrine. In: Minds Behind the Brain: A History of The Pioneers and Their Discoveries. Oxford University Press, New York, pp. 197–216.

Flechsig, P., 1920. Anatomie des menschlichen Gehirns und Rückenmarks auf myelogenetischer Grundlage. G. Thieme, Leipzig.

Fodor, J.A., 1983. The Modularity of Mind. MIT Press, Cambridge, MA.

Friston, K.J., 2010. The free-energy principle: a unified brain theory? Nat. Rev. Neurosci. 11 (2), 127–138.

Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S., 1991. Comparing functional (PET) images: the assessment of significant change. J. Cereb. Blood Flow Metab. 11 (4), 690–699.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19 (4), 1273–1302.

Fukunaga, M., Li, T.Q., van Gelderen, P., de Zwart, J.A., Shmueli, K., Yao, B., et al., 2010. Layer-specific variation of iron content in cerebral cortex as a source of MRI contrast. Proc. Natl. Acad. Sci. U.S.A. 107 (8), 3834–3839.

Gallagher, S., Hutto, D.D., Slaby, J., Cole, J., 2013. The brain as part of an enactive system. Behav. Brain Sci. 36 (4), 421–422.

Gell-Mann, M., 1995. The Quark and the Jaguar: Adventures in the Simple and the Complex. St. Martin's Press, New York.

Gennari, F., 1782. De peculiari structura cerebri nonnulisque ejus morbis. Ex Regio, Parma, Italy.

Geyer, S., Weiss, M., Reimann, K., Lohmann, G., Turner, R., 2011. Microstructural parcellation of the human cerebral cortex—from Brodmann's post-mortem map to in vivo mapping with high-field magnetic resonance imaging. Front. Hum. Neurosci. 5, 19.

Gibson, J.J., 1966. The Senses Considered as Perceptual Systems. Houghton Mifflin, Boston.

Gibson, J.J., 1979. The Ecological Approach to Visual Perception. Lawrence Erlbaum, New Jersey.

Glasser, M.F., Van Essen, D.C., 2011. Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. J. Neurosci. 31 (32), 11597–11616.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen,

D.C., 2016. A multi-modal parcellation of human cerebral cortex. Nature 536 (7615), 171–178.
Goense, J.B., Logothetis, N.K., 2006. Laminar specificity in monkey V1 using high-resolution SE-fMRI. Magn. Reson. Imaging 24 (4), 381–392.
Goense, J., Bohraus, Y., Logothetis, N.K., 2016. fMRI at high spatial resolution: implications for BOLD-models. Front. Comput. Neurosci. 10, 66.
Gonzalez-Castillo, J., Hoy, C.W., Handwerker, D.A., Roopchansingh, V., Inati, S.J., Saad, Z.S., Cox, R.W., Bandettini, P.A., 2015. Task dependence, tissue specificity, and spatial distribution of widespread activations in large single-subject functional MRI datasets at 7T. Cereb. Cortex 25 (12), 4667–4677.
Gregory, R., 1966. Eye and Brain: The Psychology of Seeing. Weidenfeld and Nicolson, London.
Hall, C.N., Reynell, C., Gesslein, B., Hamilton, N.B., Mishra, A., Sutherland, B.A., O'Farrell, F.M., Buchan, A.M., Lauritzen, M., Attwell, D., 2014. Capillary pericytes regulate cerebral blood flow in health and disease. Nature 508 (7494), 55–60.
Hamilton, N.B., Attwell, D., Hall, C.N., 2010. Pericyte-mediated regulation of capillary diameter: a component of neurovascular coupling in health and disease. Front. Neuroenergetics 21, 2.
Haxby, J.V., Connolly, A.C., Guntupalli, J.S., 2014. Decoding neural representational spaces using multivariate pattern analysis. Annu. Rev. Neurosci. 37, 435–456.
Heidemann, R.M., Anwander, A., Feiweier, T., Knösche, T.R., Turner, R., 2012a. k-Space and q-space: combining ultra-high spatial and angular resolution in diffusion imaging using ZOOPPA at 7 T. Neuroimage 60 (2), 967–978.
Heidemann, R.M., Ivanov, D., Trampel, R., Fasano, F., Meyer, H., Pfeuffer, J., Turner, R., 2012b. Isotropic submillimeter fMRI in the human brain at 7 T: combining reduced field-of-view imaging and partially parallel acquisitions. Magn. Reson. Med. 68 (5), 1506–1516.
Heinzle, J., Koopmans, P.J., Den Ouden, H.E.M., Raman, S., Stephan, K.E., 2016. A hemodynamic model for layered BOLD signals. Neuroimage 125, 556–570.
Helms, G., Dathe, H., Dechent, P., 2010. Modeling the influence of TR and excitation flip angle on the magnetization transfer ratio (MTR) in human brain obtained from 3D spoiled gradient echo MRI. Magn. Reson. Med. 64 (1), 177–185.
Henrich, J., Heine, S.J., Norenzayan, A., 2010. The weirdest people in the world? Behav. Brain Sci. 33 (2–3), 61–83.
Hinton, G.E., McClelland, J.L., Rumelhart, D.E., 1986. Distributed representations. In: Rumelhart, D.E., McClelland, J.L. (Eds.), Parallel Distributed Processing: Explorations in the Microstructure of Cognition. vol. 1. Foundations, MIT Press, Cambridge, MA, pp. 77–109.
Huber, L., Goense, J., Kennerley, A.J., Trampel, R., Guidi, M., Reimer, E., Ivanov, D., Neef, N., Gauthier, C.J., Turner, R., Möller, H.E., 2015. Cortical lamina-dependent blood volume changes in human brain at 7 T. Neuroimage 107, 23–33.
Huber, L., Handwerker, D.A., Gonzalez-Castillo, J., Jangraw, D.C., Guidi, M., Ivanov, D., Poser, B.A., Jozien Goense, J., Bandettini, P.A., 2016. Effective connectivity measured with layer-dependent resting-state blood volume fMRI in humans. Proc. Intl. Soc. Mag. Reson. Med. 24 (p), 948.
Huth, A.G., de Heer, W.A., Griffiths, T.L., Theunissen, F.E., Gallant, J.L., 2016. Natural speech reveals the semantic maps that tile human cerebral cortex. Nature 532 (7600), 453–458.

Jespersen, S.N., Kroenke, C.D., Østergaard, L., Ackerman, J.J., Yablonskiy, D.A., 2007. Modeling dendrite density from magnetic resonance diffusion measurements. Neuroimage 34 (4), 1473–1486.

Jones, D.K., Knösche, T.R., Turner, R., 2013. White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI. Neuroimage 73, 239–254.

Kendler, K.S., 2016. The nature of psychiatric disorders. World Psychiatry 15 (1), 5–12.

Kendler, K.S., Zachar, P., Craver, C., 2011. What kinds of things are psychiatric disorders? Psychol. Med. 41, 1143–1150.

Keuken, M.C., Bazin, P.L., Crown, L., Hootsmans, J., Laufer, A., Müller-Axt, C., Sier, R., van der Putten, E.J., Schäfer, A., Turner, R., Forstmann, B.U., 2014. Quantifying inter-individual anatomical variability in the subcortex using 7 T structural MRI. Neuroimage 94, 40–46.

Koenig, S.H., 1991. Cholesterol of myelin is the determinant of gray-white contrast in MRI of brain. Magn. Reson. Med. 20 (2), 285–291.

Kok, P., Bains, L.J., van Mourik, T., Norris, D.G., de Lange, F.P., 2016. Selective activation of the deep layers of the human primary visual cortex by top-down feedback. Curr. Biol. 26 (3), 371–376.

Kriegeskorte, N., Diedrichsen, J., 2016. Inferring brain-computational mechanisms with models of activity measurements. Philos. Trans. R. Soc. Lond. B Biol. Sci. 371 (1705).

Kucharczyk, W., Macdonald, P.M., Stanisz, G.J., Henkelman, R.M., 1994. Relaxivity and magnetization transfer of white matter lipids at MR imaging: importance of cerebrosides and pH. Radiology 192 (2), 521–529.

Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., Cheng, H.-M., Brady, T.J., Rosen, B.R., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proc. Natl. Acad. Sci. U.S.A. 89, 5675–5679.

LaManna, J.C., Vendel, L.M., Farrell, R.M., 1992. Brain adaptation to chronic hypobaric hypoxia in rats. J. Appl. Physiol. 72 (6), 2238–2243.

Lashley, K.S., 1929. Brain Mechanism and Intelligence. The University Press, Chicago.

Leuze, C.W., Anwander, A., Bazin, P.L., Dhital, B., Stüber, C., Reimann, K., Geyer, S., Turner, R., 2014. Layer-specific intracortical connectivity revealed with diffusion MRI. Cereb. Cortex 24 (2), 328–339.

Logothetis, N.K., 2002. The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philos. Trans. R. Soc. Lond. B Biol. Sci. 357 (1424), 1003–1037.

Logothetis, N.K., 2008. What we can do and what we cannot do with fMRI. Nature 453 (7197), 869–878.

Lohmann, G., Erfurth, K., Müller, K., Turner, R., 2012. Critical comments on dynamic causal modelling. Neuroimage 59 (3), 2322–2329.

Lohmann, G., Stelzer, J., Neumann, J., Ay, N., Turner, R., 2013. "More is different" in functional magnetic resonance imaging: a review of recent data analysis techniques. Brain Connect. 3 (3), 223–239.

Lu, H., Golay, X., Pekar, J.J., Van Zijl, P.C., 2003. Functional magnetic resonance imaging based on changes in vascular space occupancy. Magn. Reson. Med. 50 (2), 263–274.

Maley, C.J., Piccinini, G., 2017. A unified mechanistic account of teleological functions for psychology and neuroscience. In: Kaplan, D.M. (Ed.), Explanation and Integration in Mind and Brain Science. Oxford University Press, Oxford.

Mao, T., Kusefoglu, D., Hooks, B.M., Huber, D., Petreanu, L., Svoboda, K., 2011. Long-range neuronal circuits underlying the interaction between sensory and motor cortex. Neuron 72 (1), 111–123.

Marblestone, A., Wayne, G., Kording, K., 2016. Towards an integration of deep learning and neuroscience. Front. Comput. Neurosci. http://arxiv.org/abs/1606.03813.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M.W., Abdellah, M., et al., 2015. Reconstruction and simulation of neocortical microcircuitry. Cell 163 (2), 456–492.

Markuerkiaga, I., Barth, M., Norris, D.G., 2016. A cortical vascular model for examining the specificity of the laminar BOLD signal. Neuroimage 132, 491–498.

Marques, J.P., Bowtell, R., 2005. Application of a Fourier-based method for rapid calculation of field inhomogeneity due to spatial variation of magnetic susceptibility. Concepts Magn. Reson. B 25B, 65–78.

Marques, J.P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P.F., Gruetter, R., 2010. MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. Neuroimage 49 (2), 1271–1281.

Mestres-Missé, A., Turner, R., Friederici, A.D., 2012. An anterior–posterior gradient of cognitive control within the dorsomedial striatum. Neuroimage 62 (1), 41–47.

Michaloudi, H., Grivas, I., Batzios, C., Chiotelli, M., Papadopoulos, G.C., 2005. Areal and laminar variations in the vascularity of the visual, auditory, and entorhinal cortices of the developing rat brain. Brain Res. Dev. Brain Res. 155 (1), 60–70.

Micheva, K.D., Wolman, D., Mensh, B.D., Pax, E., Buchanan, J., Smith, S.J., Bock, D.D., 2016. A large fraction of neocortical myelin ensheathes axons of local inhibitory neurons. eLife 5. pii: e15784.

Morris, L.S., Kundu, P., Dowell, N., Mechelmans, D.J., Favre, P., Irvine, M.A., Robbins, T.W., Daw, N., Bullmore, E.T., Harrison, N.A., Voon, V., 2016. Fronto-striatal organization: defining functional and microstructural substrates of behavioural flexibility. Cortex 74, 118–133.

Murphy, D., 2006. Psychiatry in the Scientific Image. MIT Press, Cambridge, MA.

Nagy, Z., Alexander, D.C., Thomas, D.L., Weiskopf, N., Sereno, M.I., 2013. Using high angular resolution diffusion imaging data to discriminate cortical regions. PLoS One 8 (5), e63842.

Naselaris, T., Kay, K.N., 2015. Resolving ambiguities of MVPA using explicit models of representation. Trends Cogn. Sci. 19, 551–554.

Naselaris, T., Kay, K.N., Nishimoto, S., Gallant, J.L., 2011. Encoding and decoding in fMRI. Neuroimage 56, 400–410.

Nieuwenhuys, R., 2013. The myeloarchitectonic studies on the human cerebral cortex of the Vogt–Vogt school, and their significance for the interpretation of functional neuroimaging data. Brain Struct. Funct. 218 (2), 303–352.

Nieuwenhuys, R., Puelles, L., 2016. Towards a New Neuromorphology. Springer, Heidelberg.

Nieuwenhuys, R., Broere, C.A., Cerliani, L., 2015. A new myeloarchitectonic map of the human neocortex based on data from the Vogt–Vogt school. Brain Struct. Funct. 220 (5), 2551–2573.

Ogawa, S., Lee, T.M., Nayak, A.S., Glynn, P., 1990. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. Magn. Reson. Med. 14, 68–78.

Piccinini, G., Craver, C.F., 2011. Integrating psychology and neuroscience: functional analyses as mechanism sketches. Synthese 183, 283–311.

Pohmann, R., Speck, O., Scheffler, K., 2016. Signal-to-noise ratio and MR tissue parameters in human brain imaging at 3, 7, and 9.4 Tesla using current receive coil arrays. Magn. Reson. Med. 75 (2), 801–809.

Poldrack, R.A., Yarkoni, T., 2016. From brain maps to cognitive ontologies: informatics and the search for mental structure. Annu. Rev. Psychol. 67, 587–612.

Polimeni, J.R., Witzel, T., Fischl, B., Greve, D.N., Wald, L.L., 2010. Identifying common-source driven correlations in resting-state FMRI via laminar-specific analysis in the human visual cortex. Proc. Intl. Soc. Mag. Reson. Med. 18, 353.

Purves, D., Morgenstern, Y., Wojtach, W.T., 2015. Perception and reality: why a wholly empirical paradigm is needed to understand vision. Front. Syst. Neurosci. 9, 156.

Rao, R.P.N., Ballard, D.H., 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat. Neurosci. 2, 79–87.

Renvall, V., Witzel, T., Wald, L.L., Polimeni, J.R., 2016. Automatic cortical surface reconstruction of higher-solution T1 echo planar imaging data. Neuroimage 134, 338–354.

Roebroeck, A., Formisano, E., Goebel, R., 2005. Mapping directed influence over the brain using Granger causality and fMRI. Neuroimage 25 (1), 230–242.

Sanchez-Panchuelo, R.M., Francis, S.T., Schluppeck, D., Bowtell, R.W., 2012. Correspondence of human visual areas identified using functional and anatomical MRI in vivo at 7 T. J. Magn. Reson. Imaging 35 (2), 287–299.

Sanchez-Panchuelo, R.M., Schluppeck, D., Harmer, J., Bowtell, R., Francis, S., 2015. Assessing the spatial precision of SE and GE-BOLD contrast at 7 Tesla. Brain Topogr. 28 (1), 62–65.

Schleicher, A., Amunts, K., Geyer, S., Morosan, P., Zilles, K., 1999. Observer-independent method for microstructural parcellation of cerebral cortex: a quantitative approach to cytoarchitectonics. Neuroimage 9 (1), 165–177.

Schulz, J., Siegert, T., Reimer, E., Labadie, C., Maclaren, J., Herbst, M., Zaitsev, M., Turner, R., 2012. An embedded optical tracking system for motion-corrected magnetic resonance imaging at 7T. MAGMA 25 (6), 443–453.

Sereno, M.I., Huang, R.S., 2014. Multisensory maps in parietal cortex. Curr. Opin. Neurobiol. 24 (1), 39–46.

Sereno, M.I., Lutti, A., Weiskopf, N., Dick, F., 2013. Mapping the human cortical surface by combining quantitative T(1) with retinotopy. Cereb. Cortex 23 (9), 2261–2268.

Setsompop, K., Kimmlingen, R., Eberlein, E., Witzel, T., Cohen-Adad, J., McNab, J.A., Keil, B., Tisdall, M.D., Hoecht, P., Dietz, P., Cauley, S.F., Tountcheva, V., Matschl, V., Lenz, V.H., Heberlein, K., Potthast, A., Thein, H., Van Horn, J., Toga, A., Schmitt, F., Lehne, D., Rosen, B.R., Wedeen, V., Wald, L.L., 2013. Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. Neuroimage 80, 220–233.

Setsompop, K., Feinberg, D.A., Polimeni, J.R., 2016. Rapid brain MRI acquisition techniques at ultra-high fields. NMR Biomed. 29 (9), 1198–1221.

Shifferman, E., 2015. More than meets the fMRI: the unethical apotheosis of neuroimages. J. Cogn. Neuroethics 3, 57–116.

Shmuel, A., Yacoub, E., Pfeuffer, J., Van de Moortele, P.F., Adriany, G., Hu, X., Ugurbil, K., 2002. Sustained negative BOLD, blood flow and oxygen consumption response and its coupling to the positive response in the human brain. Neuron 36 (6), 1195–1210.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F., 2009. Correspondence of the brain's functional architecture during activation and rest. Proc. Natl. Acad. Sci. U.S.A 106 (31), 13040–13045.

Smith, S.M., Bandettini, P.A., Miller, K.L., Behrens, T.E., Friston, K.J., David, O., Liu, T., Woolrich, M.W., Nichols, T.E., 2012. The danger of systematic bias in group-level FMRI-lag-based causality estimation. Neuroimage 59 (2), 1228–1229.

Stansbury, D.E., Naselaris, T., Gallant, J.L., 2013. Natural scene statistics account for the representation of scene categories in human visual cortex. Neuron 79 (5), 1025–1034.

Stelzer, J., Lohmann, G., Mueller, K., Buschmann, T., Turner, R., 2014. Deficient approaches to human neuroimaging. Front. Hum. Neurosci. 8, 462.

Strotmann, B., Heidemann, R.M., Anwander, A., Weiss, M., Trampel, R., Villringer, A., Turner, R., 2014. High-resolution MRI and diffusion-weighted imaging of the human habenula at 7 Tesla. J. Magn. Reson. Imaging 39 (4), 1018–1026.

Tardif, C.L., Schäfer, A., Waehnert, M., Dinse, J., Turner, R., Bazin, P.L., 2015. Multi-contrast multi-scale surface registration for improved alignment of cortical areas. Neuroimage 111, 107–122.

Thomas, C., Ye, F.Q., Irfanoglu, M.O., Modi, P., Saleem, K.S., Leopold, D.A., Pierpaoli, C., 2014. Anatomical accuracy of brain connections derived from diffusion MRI tractography is inherently limited. Proc. Natl. Acad. Sci. U.S.A. 111 (46), 16574–16579.

Trampel, R., Ott, D.V., Turner, R., 2011. Do the congenitally blind have a stria of Gennari? First intracortical insights in vivo. Cereb. Cortex 21 (9), 2075–2081.

Trampel, R., Bazin, P.-L., Schäfer, A., Heidemann, R.M., Ivanov, D., Lohmann, G., Geyer, S., Turner, R., 2012. Laminar-specific fingerprints of different sensorimotor areas obtained during imagined and actual finger tapping. Proc. Intl. Soc. Mag. Reson. Med. 20, 663.

Trampel, R., Reimer, E., Huber, L., Ivanov, D., Heidemann, R.M., Schäfer, A., Turner, R., 2014. Anatomical brain imaging at 7T using two-dimensional GRASE. Magn. Reson. Med. 72 (5), 1291–1301.

Turner, R., 2002. How much cortex can a vein drain? Downstream dilution of activation-related cerebral blood oxygenation changes. Neuroimage 16 (4), 1062–1067.

Turner, R., 2012. The need for systematic ethnopsychology: the ontological status of mentalistic terminology. Anthropol. Theory 12 (1), 29–42.

Turner, R., 2015. Myelin Imaging. In: Toga, A.W. (Ed.), Brain Mapping: An Encyclopedic Reference. Elsevier, Amsterdam.

Turner, R., 2016. Uses, misuses, new uses and fundamental limitations of magnetic resonance imaging in cognitive science. Phil. Trans. R. Soc. Lond. B Biol. Sci. 371 (1705).

Turner, R., Geyer, S., 2014a. Introduction to the NeuroImage special issue: "In vivo Brodmann mapping of the human brain". Neuroimage 93 (Pt 2), 155–156.

Turner, R., Geyer, S., 2014b. Comparing like with like: the power of knowing where you are. Brain Connect. 4 (7), 547–557.

Turner, R., Jones, T., 2003. Techniques for imaging neuroscience. Br. Med. Bull. 65, 3–20.

Turner, R., Le Bihan, D., Maier, J., Vavrek, R., Hedges, L.K., Pekar, J., 1990. Echo-planar imaging of intravoxel incoherent motion. Radiology 177 (2), 407–414.

Turner, R., Le Bihan, D., Moonen, C.T., Despres, D., Frank, J., 1991. Echo-planar time course MRI of cat brain oxygenation changes. Magn. Reson. Med. 22, 159–166.

Uh, J., Lewis-Amezcua, K., Varghese, R., Lu, H., 2009. On the measurement of absolute cerebral blood volume (CBV) using vascular-space-occupancy (VASO) MRI. Magn. Reson. Med. 61 (3), 659–667.

Uttal, W.R., 2011. Mind and Brain: A Critical Appraisal of Cognitive Neuroscience. MIT Press, Cambridge MA.

Van Essen, D.C., Glasser, M.F., 2014. In vivo architectonics: a cortico-centric perspective. Neuroimage 93 (Pt 2), 157–164.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, 2013. The WU-Minn Human Connectome project: an overview. Neuroimage 80, 62–79.

Vogt, C., Vogt, O., 1919. Allgemeinere Ergebnisse unserer Hirnforschung. J. Psychol. Neurol. 25, 279–468.

Vul, E., Harris, C., Winkielman, P., Pashler, H., 2009. Puzzlingly high correlations in fMRI studies of emotion, personality and social cognition. Perspect. Psychol. Sci. 4, 274–324.

Waehnert, M.D., Dinse, J., Schäfer, A., Geyer, S., Bazin, P.L., Turner, R., Tardif, C.L., 2016. A subject-specific framework for in vivo myeloarchitectonic analysis using high resolution quantitative MRI. Neuroimage 125, 94–107.

Wedeen, V.J., Rosene, D.L., Wang, R., Dai, G., Mortazavi, F., Hagmann, P., Kaas, J.H., Tseng, W.Y., 2012. The geometric structure of the brain fiber pathways. Science 335 (6076), 1628–1634.

Worsley, K.J., Evans, A.C., Marrett, S., Neelin, P., 1992. A three-dimensional statistical analysis for CBF activation studies in human brain. J. Cereb. Blood Flow Metab. 12, 900–918.

Zaitsev, M., Dold, C., Sakas, G., Hennig, J., Speck, O., 2006. Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. Neuroimage 31 (3), 1038–1050.

Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. Neuroimage 61 (4), 1000–1016.

Zhang, J., Kolind, S.H., Laule, C., MacKay, A.L., 2015. Comparison of myelin water fraction from multiecho T2 decay curve and steady-state methods. Magn. Reson. Med. 73 (1), 223–232.

Zhao, F., Wang, P., Hendrich, K., Ugurbil, K., Kim, S.G., 2006. Cortical layer-dependent BOLD and CBV responses measured by spin-echo and gradient-echo fMRI: insights into hemodynamic regulation. Neuroimage 30 (4), 1149–1160.

Zheng, D., LaMantia, A.S., Purves, D., 1991. Specialized vascularization of the primate visual cortex. J. Neurosci. 11 (8), 2622–2629.

FURTHER READING

Ceckler, T.L., Wolff, S.D., Yip, V., Simon, S.A., Balaban, R.S., 1992. Dynamic and chemical factors affecting water proton relaxation by macromolecules. J. Magn. Reson. 98, 637–645.

Dortch, R.D., Moore, J., Li, K., Jankiewicz, M., Gochberg, D.F., Hirtle, J.A., et al., 2013. Quantitative magnetization transfer imaging of human brain at 7 T. Neuroimage 64, 640–649.

MacKay, A., Whittall, K., Adler, J., Li, D., Paty, D., Graeb, D., 1994. In vivo visualization of myelin water in brain by magnetic resonance. Magn. Reson. Med. 31 (6), 673–677.

Schmidhuber, J., 2015. Deep learning in neural networks: an overview. Neural Netw. 61, 85–117.

Stüber, C., Morawski, M., Schäfer, A., Labadie, C., Wähnert, M., Leuze, C., Streicher, M., Barapatre, N., Reimann, K., Geyer, S., Spemann, D., Turner, R., 2014. Myelin and iron concentration in the human brain: a quantitative study of MRI contrast. Neuroimage 93 (Pt 1), 95–106.

Weiskopf, N., Suckling, J., Williams, G., Correia, M.M., Inkster, B., Tait, R., et al., 2013. Quantitative multi-parameter mapping of R1, PD(*), MT, and R2(*) at 3 T: a multi-center validation. Front. Neurosci. 7, 95.

Yamins, D.L., DiCarlo, J.J., 2016. Using goal-driven deep learning models to understand sensory cortex. Nat. Neurosci. 19 (3), 356–365.