NeuroImage 197 (2019) 699–706

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Laminar fMRI and computational theories of brain function

K.E. Stephan [a,b,*], F.H. Petzschner [a], L. Kasper [a,c], J. Bayer [a,c], K.V. Wellstein [a], G. Stefanics [a,d], K.P. Pruessmann [c], J. Heinzle [a]

[a] Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich & ETH Zurich, 8032 Zurich, Switzerland
[b] Wellcome Trust Centre for Neuroimaging, University College London, London, WC1N 3BG, UK
[c] Institute for Biomedical Engineering, ETH Zurich and University of Zurich, 8092 Zurich, Switzerland
[d] Laboratory for Social and Neural Systems Research (SNS), Dept. of Economics, University of Zurich, 8006 Zurich, Switzerland

ARTICLE INFO

Keywords:
Cortical layers
Predictive coding
Effective connectivity
Neuromodeling
Computational psychiatry
Computational psychosomatics

ABSTRACT

Recently developed methods for functional MRI at the resolution of cortical layers (laminar fMRI) offer a novel window into neurophysiological mechanisms of cortical activity. Beyond physiology, laminar fMRI also offers an unprecedented opportunity to test influential theories of brain function. Specifically, hierarchical Bayesian theories of brain function, such as predictive coding, assign specific computational roles to different cortical layers. Combined with computational models, laminar fMRI offers a unique opportunity to test these proposals non-invasively in humans.

This review provides a brief overview of predictive coding and related hierarchical Bayesian theories, summarises their predictions with regard to layered cortical computations, examines how these predictions could be tested by laminar fMRI, and considers methodological challenges. We conclude by discussing the potential of laminar fMRI for clinically useful computational assays of layer-specific information processing.

Introduction

Laminar fMRI represents one of the most innovative areas of functional neuroimaging that has been rapidly developing over the past few years (e.g De Martino et al., 2015; Goense et al., 2012; Kok et al., 2016; Koopmans et al., 2010; Muckli et al., 2015; Petridou and Siero, 2017; Polimeni et al., 2010; Siero et al., 2011; Yu et al., 2014). One of the reasons why laminar fMRI is perceived as important and exciting is that influential hierarchical Bayesian theories of brain function, such as predictive coding (Clark, 2013; Feldman-Barrett and Simmons, 2015; Friston, 2005; Rao and Ballard, 1999), make claims about the specific computational roles of different cortical laminae – and laminar fMRI as the currently only non-invasive method for probing layer-specific activity in humans provides a unique opportunity to test these predictions.

In this short review, we summarise these theories, the predictions they provide with regard to layered computations, and how these predictions could be tested by laminar fMRI. In what follows, we first provide a very brief sketch of the historical roots of contemporary concepts of layered computations, followed by a brief outline of predictive coding and related theories. We then examine how predictions from these theories could be tested concretely and consider methodological challenges that need to be addressed. We conclude by discussing the potential clinical utility of laminar fMRI, particularly in the context of computational assays of layer-specific information processing.

Notably, this short review does not cover all existing theories of cortical computations that consider the roles of different layers. Instead, in order to maintain accessibility and focus, it is restricted to what is arguably the most influential concept of perception in cognitive neuroimaging at present: predictive coding, in the particular variant that originates from the concepts of Rao and Ballard (1999) and Friston (2005). By contrast, we will only briefly refer to other formulations of predictive coding (for overview of different predictive coding variants, see Spratling, 2017) and related concepts, such as hierarchical filtering (Mathys et al., 2011, 2014). For other theories of layered computations in cortex, such as "sequence seeking" in cortical counter-streams (Ullman, 1995) or belief propagation (Deneve and Jardri, 2016; Friston et al., 2017a; Lee and Mumford, 2003), the reader is referred to the original papers.

Throughout the paper, we deliberately adopt a simple and conceptual perspective, leaving aside all mathematical aspects of the theories (and glossing over many of their details) in order to keep the exposition brief and accessible for fMRI researchers. Those interested in the mathematical

* Corresponding author. Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich & ETH Zurich, 8032 Zurich, Switzerland.
E-mail address: stephan@biomed.ee.ethz.ch (K.E. Stephan).

https://doi.org/10.1016/j.neuroimage.2017.11.001
Received 4 August 2017; Received in revised form 29 October 2017; Accepted 1 November 2017
Available online 2 November 2017

details of predictive coding are referred to Friston (2005, 2008), Rao and Ballard (1999), Spratling (2017), and an excellent tutorial by Bogacz (2017).

A very brief history of theories of layered cortical computations

Contemporary concepts of the computational roles of different cortical layers have their origin in neuroanatomy. The increasing availability of highly sensitive and specific neuroanatomical tract tracing studies since the 1970s (for reviews, see Kobbert et al., 2000; Lanciego and Wouterlood, 2011) and their collation in databases since the 1990s (Felleman and Van Essen, 1991; Scannell et al., 1995; Stephan et al., 2001) made it possible to identify systematic rules of layer-specific connectivity patterns across the cortex. The seminal paper by Felleman and Van Essen (1991) conducted the first systematic analysis of these connectional patterns, with a particular emphasis on the visual system. They demonstrated that three types of connections, distinguished by different patterns of origin and target layers, allow for defining a neuroanatomical hierarchy among cortical areas that matches many longstanding physiological findings, such as receptive field sizes; for critical discussions of this hierarchical concept, see Hilgetag et al. (1996, 2000) and Markov et al. (2013, 2014).

Two of these connection types are of particular relevance for the theories discussed below (see Fig. 1). Ascending (forward or bottom-up) connections primarily originate from supragranular layers (II/III) and project to the granular layer IV of a hierarchically higher target region. Conversely, descending (backward or top-down) connections from higher to lower areas originate and terminate in extragranular layers (i.e., outside layer IV). In particular, they arise from infragranular layer V and project to layers I and VI of a hierarchically lower target region.

Soon after the publication of the Felleman & Van Essen hierarchy, David Mumford (1992) published a theoretical proposal about the computational role of this hierarchy that assigned specific roles to these connections and the associated cortical layers. In brief, he suggested that cortical hierarchies encode models, where higher levels provides abstract "templates" via descending connections that regularise and predict inputs from lower levels, while the latter only broadcast those parts of the signals up the hierarchy that do not fit the template. In his words: "The higher area attempts to fit its abstractions to the data it receives from lower areas by sending back to them from its deep pyramidal cells a template reconstruction best fitting the lower level view. The lower area attempts to reconcile the reconstruction of its view that it receives from higher areas with what it knows, sending back from its superficial pyramidal cells the features in its data which are not predicted by the higher area." (Mumford, 1992, p. 421).

A few years later, Rao and Ballard (1999) cast this idea more formally in a probabilistic (Bayesian) framework. Specifically, they suggested a hierarchical formulation of the longstanding idea of "predictive coding". Predictive coding is based on the simple notion, originally from engineering and information theory (the term was coined by Elias, 1955), that an efficient way to transmit a signal is to remove all components that can be predicted, leaving only a predictor (e.g., a model and parameter values) and data residuals to be encoded. Predictive coding was first imported to neuroscience in order to explain centre-surround antagonism in the retina (Srinivasan et al., 1982).

Capitalising on the proposal by Mumford (1992), Rao and Ballard (1999) suggested a hierarchical formulation of predictive coding in which higher areas encode predictions about the state (activity) of lower areas and signal this prediction via descending connections. At the lower level, a prediction error (PE) is computed by comparing the actual against the predicted activity; this PE is sent back up to the higher level where it serves to update the prediction. Rao and Ballard (1999) showed that this principle can be cast as a Bayesian belief updating process, where a prior belief (held by the higher area) is updated by PE input (in statistical terms, the likelihood) from the lower area. (A related hierarchical Bayesian model based on belief propagation was proposed by Lee and Mumford, 2003).

Rao and Ballard implemented their predictive coding approach as an artificial neural network that was designed to learn the statistics of static natural images. In a series of papers, Karl Friston generalised the idea of predictive coding and suggested a neurobiologically plausible implementation, describing a minimal neuronal model where each level of the

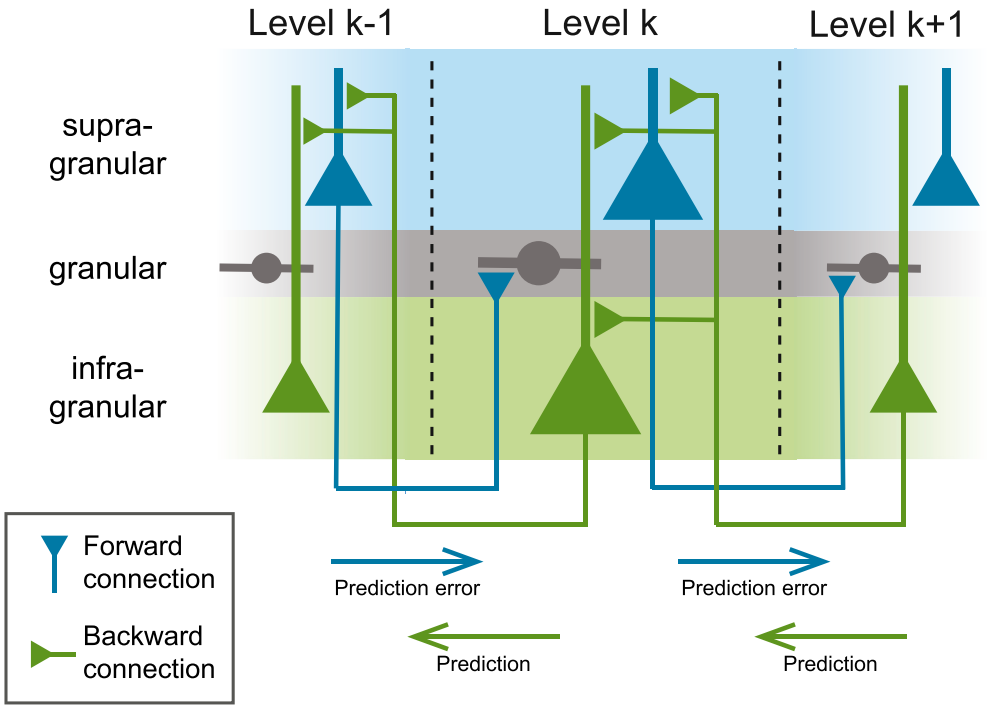

Fig. 1. Schematic illustration of forward and backward connections in cortex and their relation to predictive coding. Backward connections (green) travel down the hierarchy and convey prediction signals to both supra- and infragranular layers. Forward connections (blue) travel up the hierarchy and signal prediction errors, targeting the granular layer which, in circuit models of predictive coding (Bastos et al., 2012), is assumed to relay PE signals to supragranular cells (connection not shown).

hierarchy consists of separate neuronal units for predictions and PEs (Friston, 2003, 2005, 2008). These units are assumed to be composed of pyramidal cells located in supragranular layers II/III (for PEs) and infragranular layers V/VI (for predictions), and exchange messages with each other (via intrinsic connections, within a cortical area) and with neuronal populations at the next lower and higher levels (extrinsic connections, between areas). Their dynamics is described in terms of differential equations that prescribe a gradient descent on free energy, an approximation to information-theoretic (Shannon) surprise. This means that a sensory perturbation of the hierarchical network, for example by an unpredicted stimulus, triggers a change of beliefs (posterior distributions) at all levels such that the overall surprise is minimised (a tutorial introduction to this model can be found in Bogacz (2017)). Depending on the formulation, this predictive coding model not only enables inference on a static cause of sensory input (Friston, 2005), but can deal with a dynamically changing environment (Friston, 2008).

Friston's model of predictive coding incorporated the anatomical constraints from the earlier work by Mumford (1992) and Rao and Ballard (1999), that is, ascending connections (from supragranular to granular layers) are assumed to convey PEs and descending connections (from infragranular to extragranular layers) are thought to signal predictions. However, he finessed this concept in two important ways. First, he emphasised the importance of precision (inverse variance or uncertainty). Put simply, the magnitude of a belief update decreases with the precision of the prior belief and increases with the precision-weight of the PE (the likelihood); the latter can be understood as the signal-to-noise ratio (SNR) that is ascribed to a sensory channel. Neurobiologically, the precision of a neuronal population response distribution depends on gain (Marreiros et al., 2008). In turn, gain is known to be shaped by neuromodulatory transmitters such as dopamine and acetylcholine that regulate slow afterhyperpolarization mediated by calcium-dependent potassium channels (McCormick et al., 1993; McCormick and Williamson, 1989; Thurley et al., 2008) and represents a candidate mechanism for attentional processes (Feldman and Friston, 2010). These biological links are important as they point to suitable experimental interventions for perturbing predictive coding processes in laminar fMRI studies (see below).

Second, Friston's model made explicit that predictions and PEs are not only exchanged between cortical areas, but also within each area. For example, to compute prediction errors, supragranular neurons at a given level of the hierarchy must be informed about the current state (encoded by infragranular neurons) via intrinsic connections. Building on the seminal “canonical microcircuit” by Douglas and Martin (Douglas and Martin, 1991; Douglas et al., 1989) and subsequent work by Thomson et al. (2002) and Haeusler and Maass (2006), a canonical microcircuit model was proposed that essentially describes a prototypical cortical column and specifies the dynamics within nodes at each level of a predictive coding hierarchy (Bastos et al., 2012). This model not only assigns specific types of neurons in different cortical layers to specific computational subprocesses of predictive coding but specifies the intrinsic (within-node) connections by which predictions and PEs are signalled within a cortical column (Bastos et al., 2012).

Notably, there are alternative formulations of predictive coding that propose different algorithmic and neuronal implementations (see Spratling, 2017 for an overview). For example, in contrast to the predictive coding schemes described above, which assume that ascending connections signal PEs, the PC/BC (predictive coding/biased competition) model by Spratling (2010, 2012) proposes that ascending connections convey predictions. These opposing hypotheses can, in principle, be addressed by laminar fMRI (see next section).

While predictive coding is a theory restricted to inference (perception) and learning, it has a twin concept that specifically focuses on action or control. This theory is known as “active inference” (Adams et al., 2013a; Friston et al., 2010) and assumes the same type of objective function as predictive coding, i.e., the minimisation of prediction errors (and thus surprise) about sensory inputs. Put simply, in active inference, beliefs or expectations about (current or future) environmental states are not updated, as in predictive coding, but serve to specify actions that fulfil these beliefs. The assignment of computational quantities to different cortical layers is the same as in predictive coding; here, a particular emphasis is on pyramidal cells in layer V and VI of cortex that not only send predictions to other cortical areas, but also to subcortical structures that host reflex arcs for motor, autonomic and endocrine actions, for example, in brainstem, hypothalamus, or spinal cord (Adams et al., 2013a; Shipp et al., 2013; Stephan et al., 2016).

Together, hierarchical Bayesian models like predictive coding and active inference provide a comprehensive description of the afferent and efferent branches of inference-control loops (Petzschner et al., 2017). Of particular interest for laminar fMRI, both theories refer to the same small set of computational quantities (predictions, PEs, and precisions) that are postulated to be signalled by distinct connection types (ascending versus descending) with specific laminar patterns of origin and termination. This allows for formulating concrete predictions about layer-specific expressions of activity and connectivity in response to experimental manipulations of predictions, PEs, or precisions.

Predictions and experimental tests

Previous experimental tests of predictive coding at the circuit and area level

Predictions derived from predictive coding about brain responses to experimental manipulations can be specified at different levels of spatial and temporal resolution and with different degrees of conceptual sophistication. Spatially, one could test predictions at the levels of circuits, areas, columns and laminae, or even single neurons; temporally, predictions can be tested using responses averaged over trials (Wacongne et al., 2012), trial-by-trial responses (Lieder et al., 2013a), or even with respect to within-trial activity (Lieder et al., 2013b). Conceptually, one can perform categorical comparisons of trials with high vs. low predictability of stimuli (Alink et al., 2010; Kok et al., 2012) or with high vs. low precision of predictions (Kok et al., 2011); alternatively, one can employ computational models that provide quantitative trial-by-trial (Diaconescu et al., 2017) or within-trial (Lieder et al., 2013b) estimates of prediction errors and predictions.

Generally, predictions derived from predictive coding concern the consequences of manipulating the probability of a particular event or stimulus, or the precision by which a prediction can be made. So far, these predictions have typically been formulated at the level of cortical circuits. For example, one of the most obvious predictions by predictive coding is that a particular cortical processing hierarchy should exhibit PE responses when an unexpected sensory stimulus is administered. This has been demonstrated numerous times and in different sensory systems (e.g., den Ouden et al., 2010; den Ouden et al., 2009). One of the most famous paradigms in this regard is the so-called mismatch negativity (MMN), where unexpected stimuli are embedded into a stimulus sequence with strong statistical regularities (Garrido et al., 2009; Naatanen et al., 2007). While classically an EEG paradigm, the MMN is increasingly used in fMRI as well (Gaebler et al., 2015). A particularly striking prediction by predictive coding is that not only the unexpected presence but also the unexpected *absence* of a stimulus should elicit activation; indeed, this has been demonstrated experimentally (den Ouden et al., 2009; Naselaris et al., 2015; Wacongne et al., 2011). Conversely, predictive coding also postulates that activity should be diminished when a stimulus can be predicted (“explaining away”). This has been confirmed in a number of experiments, particularly in the visual and auditory domain (Harrison et al. 2007; Alink et al., 2010).

A third prediction is that precision should critically affect the magnitude of evoked PE responses. That is, higher precision of predictions should decrease PE responses, while higher precision assigned to sensory channels should increase them. While perhaps the least studied of the three core computational quantities of predictive coding, these effects have also been confirmed by studies using manipulations of

attention and neuromodulatory systems, respectively (e.g., Kok et al., 2011; Moran et al., 2013; Vossel et al., 2015).

As illustrated by this short overview, established fMRI methods have already been extremely useful for testing many of the implications of predictive coding. So far, however, all available fMRI studies that tested implications of predictive coding have focused on one core quantity (PE, prediction or precision) alone, or on the precision-weighting of either predictions or PEs. This is in contrast to invasive human electrophysiological studies that simultaneously examined PEs (surprise), predictions, and precision, albeit using a non-hierarchical model (Sedley et al., 2016). To our knowledge, there is no study yet that tested the general nature of the hierarchical Bayesian message passing *in toto* and demonstrated that all variables act as proposed by predictive coding, with regard to temporal expression and their specificity for certain types of connections.

Furthermore, the existing experimental literature has almost exclusively focused on predictions related to activations by computational quantities and their modulation of connectivity at the level of interacting cortical areas. By contrast, the spatial resolution of conventional fMRI has prevented testing existing proposals about within-area mechanisms of predictive coding, for example, which layers compute PEs and predictions, respectively, and how these quantities are communicated and integrated across layers. Given appropriate experimental designs, this is the domain where high-resolution laminar fMRI can make a unique contribution.

Potential future experiments using laminar fMRI

The plausibility of predictive coding as a general theory of inference and learning cannot be corroborated conclusively by merely examining activations and connectivity at the level of cortical areas. This is seen easily by noting that the suggested inter-area connectivity patterns do not constitute an uninterrupted chain of information processing unless one adds intra-areal relays via intrinsic connections that link different layers. For example, PE-computing neuronal units with the proposed location in supragranular layers need access to input from the hierarchical level below; this input, however, is suggested by predictive coding schemes (and its precursors) to enter granular layer IV (Friston, 2005; Mumford, 1992; Rao and Ballard, 1999).

Specific suggestions about the neuroanatomical layout of intrinsic connections between cortical layers within a given hierarchical level exist, such as the canonical microcircuit model proposed by (Bastos et al., 2012). At present, laminar fMRI arguably represents the most powerful method to directly test these proposals in humans. While computational methods for inferring the laminar contributions to non-invasive electrophysiological measurements in humans are also under development (Bonaiuto et al., 2017), these are less advanced and require more *a priori* knowledge. Either approach will greatly benefit from the information provided by studies in non-human primates that can record activity from individual layers with exquisite spatial and temporal precision (e.g., Bastos et al., 2015; Michalareas et al., 2016; Self et al., 2012, 2017; van Kerkoerle et al., 2014).

Theoretical proposals of predictive coding microcircuits can be tested with high-resolution laminar fMRI in several ways. For example, the canonical microcircuit model by Bastos et al. (2012) makes specific predictions about how PE signals entering the granular layer spread within the target region, reaching supragranular layers first and being relayed subsequently to infragranular layers. Provided laminar fMRI data with sufficient spatial and temporal resolution and SNR are available, this proposal might be testable using generative models of fMRI data (for review, Stephan et al., 2015). For example, dynamic causal modelling (DCM; Friston et al., 2003) that partitions blood oxygen level dependent (BOLD) data into neuronal and vascular contributions (an important topic we revisit below) provides powerful opportunities for statistical model comparison, both with regard to neuronal (Penny et al., 2004) and hemodynamic (Stephan et al., 2007) model components. DCMs with laminar resolution (Heinzle et al., 2016) would allow one to test which layers are likely targeted by forward and backward connections and perhaps also evaluate the relative plausibility of different possibilities about the directionality of signals within cortical microcircuits. Similarly, one could test competing hypotheses whether a change in laminar responses results from changes in extrinsic inputs or is due to a change in synaptic gain.

However, the importance of laminar fMRI for experimental tests of predictive coding already applies at an even more fundamental level than examining intra-areal signal flow and microcircuit dynamics (compare the benchmarks for laminar fMRI proposed by Self et al., 2017). For example, in humans, some of the anatomical assumptions of the predictive coding concepts by Rao and Ballard (1999) and Friston (2005), namely that PE units are located in supragranular layers and that prediction units are situated in infragranular layers, still require conclusive experimental verification. While first laminar fMRI studies have begun addressing this issue (e.g., Kok et al., 2016; Muckli et al., 2015), the available results are based on categorical comparisons of experimental conditions with more versus less predictable stimuli. By contrast, to our knowledge, a parametric relationship between lamina-specific BOLD signals and quantitative, trial-by-trial estimates of PEs and predictions has not been reported so far. Furthermore, the basic assumption of the predictive coding concepts by Rao and Ballard (1999) and Friston (2005) that PE signals are conveyed via forward connections (originating from supragranular layers and terminating in granular layers) and that predictions are signalled via backward connections (from infragranular to extragranular layers) still requires testing. This is important, not least because another variant of predictive coding, the PC/BC model by Spratling (2010, 2012), assumes a different circuit layout and postulates that forward connections signal predictions.

Testing these opposing hypotheses requires embedding the respective computational models into laminar DCMs, using trial-wise estimates of PEs and predictions as modulators of ascending or descending connections (compare den Ouden et al., 2010; den Ouden et al., 2009). In a subsequent step, Bayesian model selection (Stephan et al., 2009) could be used to evaluate whether forward connections more likely convey PEs (the Rao & Ballard/Friston model) or predictions (the Spratling model). This model comparison could also include simpler models that do not represent predictive coding or other models based on hierarchical Bayesian inference. This allows for testing whether there is any evidence at all that the dynamics of cortical connection strengths reflect hierarchical Bayesian inference, or whether a simpler explanation suffices (for a similar model comparison approach using EEG, see Lieder et al., 2013a).

One particularly elegant demonstration for the plausibility of predictive coding would be provided by a generative model that is simultaneously fitted to fMRI and behavioural data and explains trial-by-trial fluctuations of both BOLD signal and behaviour as the consequence of PEs and prediction signals being relayed along the suggested connections within and between areas, while taking into account potential vascular confounds (see below). For this type of hybrid generative model only simple formulations exist that are presently outside the realm of laminar fMRI (Rigoux and Daunizeau, 2015); compare the discussion in Stephan et al. (2017). Another hybrid modelling approach could be supported by simultaneous EEG-fMRI measurements (Scheeringa et al., 2016). This is particularly relevant given results from non-human primate studies indicating that the putative signaling of PEs and predictions along ascending and descending connections, respectively, may operate predominantly in different frequency bands (Bastos et al., 2015; Michalareas et al., 2016; van Kerkoerle et al., 2014).

Methodological challenges

Laminar fMRI poses a number of methodological challenges, many of which are discussed in other articles of this special issue. Here, we would like to discuss three of them in more depth that are particularly pertinent to testing predictive coding models.

First and foremost, the strongest contribution to the BOLD signal comes from venous compartments; these, however, do not show a uniform distribution across the cortical depth. Instead, the blood supply of cortex exhibits a highly regular and similarly layered structure as cortex itself: descending arterioles penetrate the cortex in a perpendicular fashion, with capillaries branching off to supply individual cortical layers, and venous blood returns to the pial surface in ascending venules (Duvernoy et al., 1981; Weber et al., 2008). Due to this vascular architecture, BOLD signals in upper cortical layers are affected by blood draining effects, that is, venous blood from lower cortical layers contributes to the BOLD signal measured in upper cortical layers. This complicates the analysis of layer-specific activations by computational quantities. Perhaps most importantly, it cautions against using connectivity analysis approaches that are agnostic about laminar differences in hemodynamics, such as simple functional connectivity indices, for analysing the signaling of PEs or predictions across areas and layers.

Instead, in order to infer mechanisms of interest for predictive coding, we require modelling approaches that can partition laminar signals into neuronal and vascular components (cf. David et al., 2008) and distinguish different factors that may cause activation in a specific lamina: remote neuronal inputs (extrinsic connections from other areas), local neuronal connections across layers (intrinsic connections), local hemodynamics (within a given layer), and local blood draining effects (across layers). A first model of this sort, essentially representing a single-region DCM equipped with an extended hemodynamic model that accounts for venous draining effects, was proposed recently (Heinzle et al., 2016); see Fig. 2. Simulations showed that this model can, under certain conditions (e.g., sufficient SNR), disambiguate neuronal and vascular sources of fMRI signals in upper cortical layers and distinguish between the effects of intrinsic synaptic connectivity and across-layer draining effects. While this model represents an important first step, its application to questions of predictive coding requires extension to a multi-region setting. An alternative model of potential utility is a recent DCM variant that combines a canonical microcircuit model with a hemodynamic forward model (Friston et al., 2017b). While this model does not presently contain a layer-specific formulation of BOLD signal generation, it could be extended to do so. Further avenues of development for models of laminar high-field fMRI data can be found in the review article by Uludağ and Blinder (2017).

A second methodological challenge is that thorough layer-specific tests of predictive coding processes depend on tight control over the key computational quantities from predictive coding, i.e., PEs, predictions, and precision. This requires (i) experimental manipulations by which these quantities are altered and (ii) predictive coding models that can be applied to empirical data and infer subject-specific trajectories of the quantities of interest (i.e., precision-weighted PEs and predictions). Neither of these requirements is trivially fulfilled. Concerning experimental manipulations of PEs and predictions, these are straightforward in the exteroceptive domain (e.g., using cues that predict trial outcomes probabilistically), but represent a non-trivial problem for interoception; compare the discussions in Khalsa and Lapidus (2016) and Petzschner et al. (2017). In either interoceptive or exteroceptive domains, manipulating precision can be achieved by manipulating the statistical structure of stimulus sequences (e.g., Lecaignard et al., 2015; Vossel et al., 2014), by changing attention, or by pharmacological interventions that affect neuromodulatory systems. Concerning models for practical data analysis, the existing implementations of predictive coding are useful for simulations, but cannot presently be used to infer trajectories of trial-wise precision-weighted PEs and predictions from measured data. In practice, one can resort to closely related hierarchical Bayesian models, such as the hierarchical Gaussian filter (Mathys et al., 2011, 2014). These hierarchical filtering models rest on precision-weighted PEs and predictions in the same way as predictive coding, but are not fully identical: in these models, higher levels do not predict the state of the level below, as in predictive coding, but the rate of change.

Finally, it is worth noting that recent neuroanatomical studies point

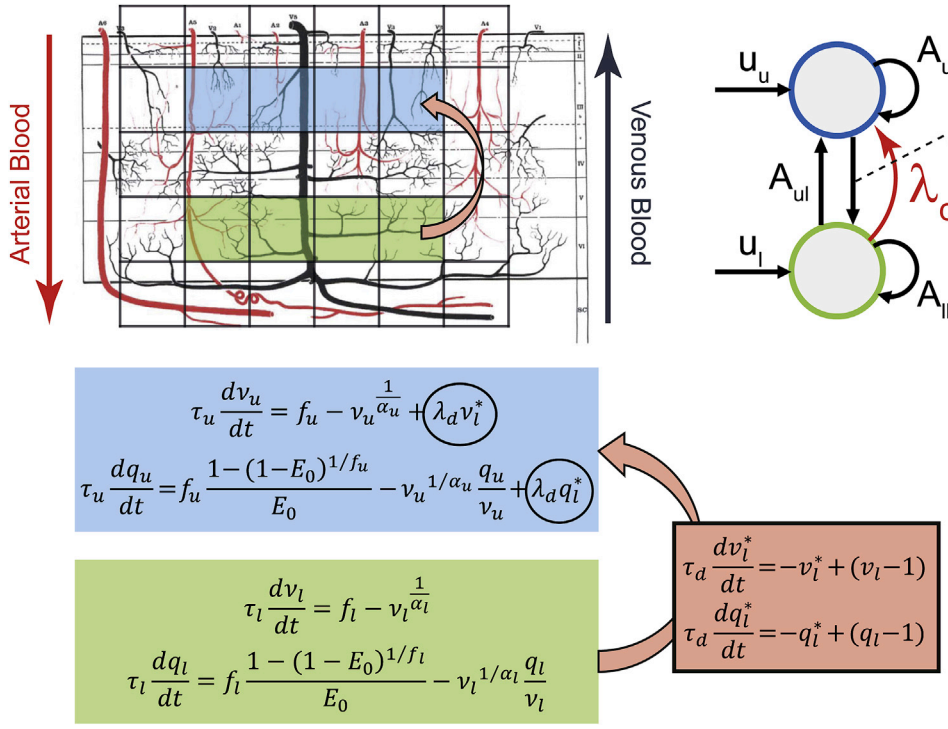

Fig. 2. Summary of a generative model of laminar fMRI data that corresponds to a single-region DCM (Heinzle et al., 2016) and accounts for hemodynamic confounds (venous blood draining effects in upper layers). Top left: Illustration of venous blood draining effects in layered fMRI. Neuronal activation in infragranular layers (green) leads to changes in blood volume and deoxyhemoglobin concentration that, due to upward draining, affects signals from supragranular layers (blue), thus potentially confounding inference on supragranular neuronal activation. Top right: Sketch of a two-layer DCM with full neuronal connectivity. The effect of blood draining is indicated by the red arrow. Bottom: Schematic illustration of an extended Balloon-Windkessel model of hemodynamics that takes into account blood draining effects; for details, see Heinzle et al. (2016). Figure adapted, with permission, from Heinzle et al. (2016) and Duvernoy et al. (1981).

to a more complicated layout of hierarchy-defining connections in cortex than assumed previously (Markov et al., 2013, 2014). While not calling into question the general definition of forward and backward connections, these studies demonstrated the existence of so-called "cortical counter streams". These are pathways consisting of forward and backward connections that run in parallel to each other within supragranular and infragranular layers, respectively. While the functional significance of these pathways is hitherto unknown, future revisions of predictive coding concepts are challenged to take these new findings into account; cf. Shipp (2016). In any case, the discovery of these more complex layer-specific pathways of communication further underscore the importance of high-resolution laminar fMRI for elucidating principles of cortical computations.

Potential clinical utility of laminar fMRI

Predictive coding and related hierarchical Bayesian theories of brain function are not only amongst the most influential concept in basic neuroscience, but represent important elements of contemporary theories of mental disorders. For example, predictive coding has taken centre stage in theories of schizophrenia (Friston et al., 2016; Stephan et al., 2006), particularly with regard to positive symptoms like hallucinations and delusions (Corlett et al., 2009, 2011; Fletcher and Frith, 2009; Powers et al., 2017). Similarly, predictive coding and hierarchical Bayesian concepts have begun exerting notable influence on disease theories of autism (Haker et al., 2016; Lawson et al., 2014) and depression (Seth and Friston, 2016; Stephan et al., 2016), as well as psychosomatic disorders including pain and placebo effects (Buchel et al., 2014; Farb et al., 2015; Wiech, 2016). In brief, predictive coding currently represents the leading concept for understanding disturbances of perception in computational psychiatry (Friston et al., 2014) and computational psychosomatics (Petzschner et al., 2017).

In combination with generative models and suitable experimental perturbations, laminar fMRI might represent a promising technique to assay, in individual patients, failures of hierarchical Bayesian message passing in disease-relevant circuits. Such computational assays could potentially aid differential diagnosis, particularly in psychosomatics, where the closed-loop nature of inference-control circuits poses considerable diagnostic problems (Petzschner et al., 2017). A model-based and physiologically interpretable readout of hierarchical Bayesian message passing could, for example, help to distinguish between disturbances of afferent (interoception), efferent (control), and metacognitive (self-monitoring) branches of circuits for homeostatic/allostatic regulation of bodily processes (Stephan et al., 2016). It could also contribute to predicting or monitoring the patient-specific efficacy of psychotherapeutic approaches that target core computational quantities of predictive coding, such as precision-reweighting by contemplative approaches (Farb et al., 2015).

Finally, laminar models of predictive coding that disambiguate between failures of PEs, predictions or precisions could also potentially be useful for pharmacological considerations because each of these quantities may be related to different classes of neurotransmitters. For example, while PE signaling along forward connections mainly seems to be utilising glutamatergic AMPA receptors, predictions via backward connections appear to be relying on NMDA receptors (Self et al., 2012). Precision-weighting is likely implemented by changes in gain which, in turn, depends on neuromodulatory (e.g., dopaminergic and cholinergic) regulation of excitability (McCormick et al., 1993) as well as GABAergic mechanisms. For reviews of these potential relations and their relevance for psychiatry, see Adams et al. (2013b); Corlett et al. (2011, 2010).

These few selected examples may already be sufficient to illustrate why high-resolution laminar fMRI is an essential component within a computational neuroimaging strategy for establishing clinically useful single-patient predictions (Stephan et al., 2017). Provided we can address the methodological challenges discussed above and further enhance the resolution and SNR of laminar fMRI data, the combination of computational models and laminar fMRI have considerable promise to result in diagnostic tools for psychiatry and psychosomatics.

Acknowledgements

We are very grateful for generous support by the René and Susanne Braginsky Foundation and the University of Zurich.

References

Adams, R.A., Shipp, S., Friston, K.J., 2013a. Predictions not commands: active inference in the motor system. Brain Struct. Funct. 218, 611–643.
Adams, R.A., Stephan, K.E., Brown, H.R., Frith, C.D., Friston, K.J., 2013b. The computational anatomy of psychosis. Front. Psychiatry 4, 47.
Alink, A., Schwiedrzik, C.M., Kohler, A., Singer, W., Muckli, L., 2010. Stimulus predictability reduces responses in primary visual cortex. J. Neurosci. 30, 2960–2966.
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, 695–711.
Bastos, A.M., Vezoli, J., Bosman, C.A., Schoffelen, J.M., Oostenveld, R., Dowdall, J.R., De Weerd, P., Kennedy, H., Fries, P., 2015. Visual areas exert feedforward and feedback influences through distinct frequency channels. Neuron 85, 390–401.
Bogacz, R., 2017. A tutorial on the free-energy framework for modelling perception and learning. J. Math. Psychol. 76, 198–211.
Bonaiuto, J.J., Rossiter, H.E., Meyer, S.S., Adams, N., Little, S., Callaghan, M.F., Dick, F., Bestmann, S., Barnes, G.R., 2017. Non-invasive Laminar Inference with MEG: Comparison of Methods and Source Inversion Algorithms. bioRxiv, 147215.
Buchel, C., Geuter, S., Sprenger, C., Eippert, F., 2014. Placebo analgesia: a predictive coding perspective. Neuron 81, 1223–1239.
Clark, A., 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behav. Brain Sci. 36, 181–204.
Corlett, P.R., Frith, C.D., Fletcher, P.C., 2009. From drugs to deprivation: a Bayesian framework for understanding models of psychosis. Psychopharmacol. Berl. 206, 515–530.
Corlett, P.R., Honey, G.D., Krystal, J.H., Fletcher, P.C., 2011. Glutamatergic model psychoses: prediction error, learning, and inference. Neuropsychopharmacology 36, 294–315.
Corlett, P.R., Taylor, J.R., Wang, X.J., Fletcher, P.C., Krystal, J.H., 2010. Toward a neurobiology of delusions. Prog. Neurobiol. 92, 345–369.
David, O., Guillemain, I., Saillet, S., Reyt, S., Deransart, C., Segebarth, C., Depaulis, A., 2008. Identifying neural drivers with functional MRI: an electrophysiological validation. PLoS Biol. 6, 2683–2697.
De Martino, F., Moerel, M., Ugurbil, K., Goebel, R., Yacoub, E., Formisano, E., 2015. Frequency preference and attention effects across cortical depths in the human primary auditory cortex. Proc. Natl. Acad. Sci. 112, 16036–16041.
den Ouden, H.E., Daunizeau, J., Roiser, J., Friston, K.J., Stephan, K.E., 2010. Striatal prediction error modulates cortical coupling. J. Neurosci. 30, 3210–3219.
den Ouden, H.E., Friston, K.J., Daw, N.D., McIntosh, A.R., Stephan, K.E., 2009. A dual role for prediction error in associative learning. Cereb. Cortex 19, 1175–1185.
Deneve, S., Jardri, R., 2016. Circular inference: mistaken belief, misplaced trust. Curr. Opin. Behav. Sci. 11, 40–48.
Diaconescu, A.O., Mathys, C., Weber, L.A., Kasper, L., Mauer, J., Stephan, K.E., 2017. Hierarchical prediction errors in midbrain and septum during social learning. Soc. Cogn. Affect Neurosci. 12, 618–634.
Douglas, R.J., Martin, K., 1991. A functional microcircuit for cat visual cortex. J. Physiology 440, 735–769.
Douglas, R.J., Martin, K.A., Whitteridge, D., 1989. A canonical microcircuit for neocortex. Neural Comput. 1, 480–488.
Duvernoy, H.M., Delon, S., Vannson, J., 1981. Cortical blood vessels of the human brain. Brain Res. Bull. 7, 519–579.
Elias, P., 1955. Predictive coding–I. IRE Trans. Inf. Theory 1, 16–24.
Farb, N., Daubenmier, J., Price, C.J., Gard, T., Kerr, C., Dunn, B.D., Klein, A.C., Paulus, M.P., Mehling, W.E., 2015. Interoception, contemplative practice, and health. Front. Psychol. 6, 763.
Feldman-Barrett, L.F., Simmons, W.K., 2015. Interoceptive predictions in the brain. Nat. Rev. Neurosci. 16, 419–429.
Feldman, H., Friston, K.J., 2010. Attention, uncertainty, and free-energy. Front. Hum. Neurosci. 4, 215.
Felleman, D.J., Van Essen, D.C., 1991. Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex 1, 1–47.
Fletcher, P.C., Frith, C.D., 2009. Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. Nat. Rev. Neurosci. 10, 48–58.
Friston, K., 2003. Learning and inference in the brain. Neural Netw. 16, 1325–1352.
Friston, K., 2005. A theory of cortical responses. Philos. Trans. R. Soc. Lond B Biol. Sci. 360, 815–836.
Friston, K., 2008. Hierarchical models in the brain. PLoS Comput. Biol. 4, e1000211.
Friston, K., Brown, H.R., Siemerkus, J., Stephan, K.E., 2016. The dysconnection hypothesis (2016). Schizophr. Res. 176, 83–94.
Friston, K.J., Daunizeau, J., Kilner, J., Kiebel, S.J., 2010. Action and behavior: a free-energy formulation. Biol. Cybern. 102, 227–260.
Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19, 1273–1302.

Friston, K.J., Parr, T., de Vries, B., 2017a. The graphical brain: belief propagation and active inference. Netw. Neurosci. 0, 1–78.
Friston, K.J., Preller, K.H., Mathys, C., Cagnan, H., Heinzle, J., Razi, A., Zeidman, P., 2017b. Dynamic causal modelling revisited. Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.02.045 (in press).
Friston, K.J., Stephan, K.E., Montague, R., Dolan, R.J., 2014. Computational psychiatry: the brain as a phantastic organ. Lancet Psychiatry 1, 148–158.
Gaebler, A.J., Mathiak, K., Koten Jr., J.W., König, A.A., Koush, Y., Weyer, D., Depner, C., Matentzoglu, S., Edgar, J.C., Willmes, K., 2015. Auditory mismatch impairments are characterized by core neural dysfunctions in schizophrenia. Brain 138, 1410–1423.
Garrido, M.I., Kilner, J.M., Stephan, K.E., Friston, K.J., 2009. The mismatch negativity: a review of underlying mechanisms. Clin. Neurophysiol. 120, 453–463.
Goense, J., Merkle, H., Logothetis, N.K., 2012. High-resolution fMRI reveals laminar differences in neurovascular coupling between positive and negative BOLD responses. Neuron 76, 629–639.
Haeusler, S., Maass, W., 2006. A statistical analysis of information-processing properties of lamina-specific cortical microcircuit models. Cereb. Cortex 17, 149–162.
Haker, H., Schneebeli, M., Stephan, K.E., 2016. Can bayesian theories of autism spectrum disorder help improve clinical practice? Front. Psychiatry 7, 107.
Harrison, L.M., Stephan, K.E., Rees, G., Friston, K.J., 2007. Extra-classical receptive field effects measured in striate cortex with fMRI. NeuroImage 34, 1199–1208.
Heinzle, J., Koopmans, P.J., den Ouden, H.E., Raman, S., Stephan, K.E., 2016. A hemodynamic model for layered BOLD signals. Neuroimage 125, 556–570.
Hilgetag, C.C., O'Neill, M.A., Young, M.P., 1996. Indeterminate organization of the visual system. Science 271, 776–777.
Hilgetag, C.C., O'Neill, M.A., Young, M.P., 2000. Hierarchical organization of macaque and cat cortical sensory systems explored with a novel network processor. Philos. Trans. R. Soc. Lond B Biol. Sci. 355, 71–89.
Khalsa, S.S., Lapidus, R.C., 2016. Can interoception improve the pragmatic search for biomarkers in psychiatry? Front. Psychiatry 7, 121.
Kobbert, C., Apps, R., Bechmann, I., Lanciego, J.L., Mey, J., Thanos, S., 2000. Current concepts in neuroanatomical tracing. Prog. Neurobiol. 62, 327–351.
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, 371–376.
Kok, P., Jehee, J.F., De Lange, F.P., 2012. Less is more: expectation sharpens representations in the primary visual cortex. Neuron 75, 265–270.
Kok, P., Rahnev, D., Jehee, J.F., Lau, H.C., de Lange, F.P., 2011. Attention reverses the effect of prediction in silencing sensory signals. Cereb. Cortex 22, 2197–2206.
Koopmans, P.J., Barth, M., Norris, D.G., 2010. Layer-specific BOLD activation in human V1. Hum. Brain Mapp. 31, 1297–1304.
Lanciego, J.L., Wouterlood, F.G., 2011. A half century of experimental neuroanatomical tracing. J. Chem. Neuroanat. 42, 157–183.
Lawson, R.P., Rees, G., Friston, K.J., 2014. An aberrant precision account of autism. Front. Hum. Neurosci. 8, 302.
Lecaignard, F., Bertrand, O., Gimenez, G., Mattout, J., Caclin, A., 2015. Implicit learning of predictable sound sequences modulates human brain responses at different levels of the auditory hierarchy. Front. Hum. Neurosci. 9.
Lee, T.S., Mumford, D., 2003. Hierarchical Bayesian inference in the visual cortex. J. Opt. Soc. Am. A Opt. Image Sci. Vis. 20, 1434–1448.
Lieder, F., Daunizeau, J., Garrido, M.I., Friston, K.J., Stephan, K.E., 2013a. Modelling trial-by-trial changes in the mismatch negativity. PLoS Comput. Biol. 9, e1002911.
Lieder, F., Stephan, K.E., Daunizeau, J., Garrido, M.I., Friston, K.J., 2013b. A neurocomputational model of the mismatch negativity. PLoS Comput. Biol. 9, e1003288.
Markov, N.T., Ercsey-Ravasz, M., Van Essen, D.C., Knoblauch, K., Toroczkai, Z., Kennedy, H., 2013. Cortical high-density counterstream architectures. Science 342, 1238406.
Markov, N.T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., Barone, P., Dehay, C., Knoblauch, K., Kennedy, H., 2014. Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. J. Comp. Neurol. 522, 225–259.
Marreiros, A.C., Daunizeau, J., Kiebel, S.J., Friston, K.J., 2008. Population dynamics: variance and the sigmoid activation function. Neuroimage 42, 147–157.
Mathys, C., Daunizeau, J., Friston, K.J., Stephan, K.E., 2011. A Bayesian foundation for individual learning under uncertainty. Front. Hum. Neurosci. 5, 39.
Mathys, C.D., Lomakina, E.I., Daunizeau, J., Iglesias, S., Brodersen, K.H., Friston, K.J., Stephan, K.E., 2014. Uncertainty in perception and the hierarchical Gaussian filter. Front. Hum. Neurosci. 8, 825.
McCormick, D.A., Wang, Z., Huguenard, J., 1993. Neurotransmitter control of neocortical neuronal activity and excitability. Cereb. Cortex 3, 387–398.
McCormick, D.A., Williamson, A., 1989. Convergence and divergence of neurotransmitter action in human cerebral cortex. Proc. Natl. Acad. Sci. U. S. A. 86, 8098–8102.
Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J.M., Kennedy, H., Fries, P., 2016. Alpha-beta and gamma rhythms subserve feedback and feedforward influences among human visual cortical areas. Neuron 89, 384–397.
Moran, R.J., Campo, P., Symmonds, M., Stephan, K.E., Dolan, R.J., Friston, K.J., 2013. Free energy, precision and learning: the role of cholinergic neuromodulation. J. Neurosci. 33, 8227–8236.
Muckli, L., De Martino, F., Vizioli, L., Petro, L.S., Smith, F.W., Ugurbil, K., Goebel, R., Yacoub, E., 2015. Contextual feedback to superficial layers of V1. Curr. Biol. 25, 2690–2695.
Mumford, D., 1992. On the computational architecture of the neocortex. II. The role of cortico-cortical loops. Biol. Cybern. 66, 241–251.
Naatanen, R., Paavilainen, P., Rinne, T., Alho, K., 2007. The mismatch negativity (MMN) in basic research of central auditory processing: a review. Clin. Neurophysiol. 118, 2544–2590.
Naselaris, T., Olman, C.A., Stansbury, D.E., Ugurbil, K., Gallant, J.L., 2015. A voxel-wise encoding model for early visual areas decodes mental images of remembered scenes. Neuroimage 105, 215–228.
Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004. Comparing dynamic causal models. Neuroimage 22, 1157–1172.
Petridou, N., Siero, J.C.W., 2017. Laminar fMRI: what can the time domain tell us? Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.07.040 (in press).
Petzschner, F.H., Weber, L.A.E., Gard, T., Stephan, K.E., 2017. Computational Psychosomatics and Computational Psychiatry: towards a joint framework for differential diagnosis. Biol. Psychiatry 82, 421–430.
Polimeni, J.R., Fischl, B., Greve, D.N., Wald, L.L., 2010. Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1. Neuroimage 52, 1334–1346.
Powers, A.R., Mathys, C., Corlett, P.R., 2017. Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors. Science 357, 596–600.
Rao, R.P., 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.
Rigoux, L., Daunizeau, J., 2015. Dynamic causal modelling of brain-behaviour relationships. Neuroimage 117, 202–221.
Scannell, J.W., Blakemore, C., Young, M.P., 1995. Analysis of connectivity in the cat cerebral cortex. J. Neurosci. 15, 1463–1483.
Scheeringa, R., Koopmans, P.J., van Mourik, T., Jensen, O., Norris, D.G., 2016. The relationship between oscillatory EEG activity and the laminar-specific BOLD signal. Proc. Natl. Acad. Sci. 113, 6761–6766.
Sedley, W., Gander, P.E., Kumar, S., Kovach, C.K., Oya, H., Kawasaki, H., Howard, M.A., Griffiths, T.D., 2016. Neural signatures of perceptual inference. Elife 5, e11476.
Self, M.W., Kooijmans, R.N., Super, H., Lamme, V.A., Roelfsema, P.R., 2012. Different glutamate receptors convey feedforward and recurrent processing in macaque V1. Proc. Natl. Acad. Sci. U. S. A. 109, 11031–11036.
Self, M.W., van Kerkoerle, T., Goebel, R., Roelfsema, P.R., 2017. Benchmarking laminar fMRI: neuronal spiking and synaptic activity during top-down and bottom-up processing in the different layers of cortex. Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.06.045 (in press).
Seth, A.K., Friston, K.J., 2016. Active interoceptive inference and the emotional brain. Philos. Trans. R. Soc. Lond B Biol. Sci. 371.
Shipp, S., 2016. Neural elements for predictive coding. Front. Psychol. 7, 1792.
Shipp, S., Adams, R.A., Friston, K.J., 2013. Reflections on agranular architecture: predictive coding in the motor cortex. Trends Neurosci. 36, 706–716.
Siero, J.C., Petridou, N., Hoogduin, H., Luijten, P.R., Ramsey, N.F., 2011. Cortical depth-dependent temporal dynamics of the BOLD response in the human brain. J. Cereb. Blood Flow. Metab. 31, 1999–2008.
Spratling, M.W., 2010. Predictive coding as a model of response properties in cortical area V1. J. Neurosci. 30, 3531–3543.
Spratling, M.W., 2012. Unsupervised learning of generative and discriminative weights encoding elementary image components in a predictive coding model of cortical function. Neural comput. 24, 60–103.
Spratling, M.W., 2017. A review of predictive coding algorithms. Brain Cogn. 112, 92–97.
Srinivasan, M.V., Laughlin, S.B., Dubs, A., 1982. Predictive coding: a fresh view of inhibition in the retina. Proc. R. Soc. Lond. B Biol. Sci. 216, 427–459.
Stephan, K.E., Baldeweg, T., Friston, K.J., 2006. Synaptic plasticity and dysconnection in schizophrenia. Biol. Psychiatry 59, 929–939.
Stephan, K.E., Iglesias, S., Heinzle, J., Diaconescu, A.O., 2015. Translational perspectives for computational neuroimaging. Neuron 87, 716–732.
Stephan, K.E., Kamper, L., Bozkurt, A., Burns, G.A., Young, M.P., Kötter, R., 2001. Advanced database methodology for the collation of connectivity data on the macaque brain (CoCoMac). Philos. Trans. R. Soc. Lond B Biol. Sci. 356, 1159–1186.
Stephan, K.E., Manjaly, Z.M., Mathys, C., Weber, L.A.E., Paliwal, S., Gard, T., Tittgemeyer, M., Fleming, S., Haker, H., Seth, A., Petzschner, F.H., 2016. Allostatic self-efficacy: a metacognitive theory of dyshomeostasis-induced fatigue and depression. Front. Hum. Neurosci. 10, 550.
Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., Friston, K.J., 2009. Bayesian model selection for group studies. Neuroimage 46, 1004–1017.
Stephan, K.E., Schlagenhauf, F., Huys, Q.J., Raman, S., Aponte, E.A., Brodersen, K.H., Rigoux, L., Moran, R.J., Daunizeau, J., Dolan, R.J., Friston, K.J., Heinz, A., 2017. Computational neuroimaging strategies for single patient predictions. Neuroimage 145, 180–199.
Stephan, K.E., Weiskopf, N., Drysdale, P.M., Robinson, P.A., Friston, K.J., 2007. Comparing hemodynamic models with DCM. Neuroimage 38, 387–401.
Thomson, A.M., West, D.C., Wang, Y., Bannister, A.P., 2002. Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2–5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. Cereb. Cortex 12, 936–953.
Thurley, K., Senn, W., Lüscher, H.-R., 2008. Dopamine increases the gain of the input-output response of rat prefrontal pyramidal neurons. J. neurophysiology 99, 2985–2997.
Ullman, S., 1995. Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. Cereb. Cortex 5, 1–11.
Uludağ, K., Blinder, P., 2017. Linking brain vascular physiology to hemodynamic response in ultra-high field MRI. Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.02.063 (in press).
van Kerkoerle, T., Self, M.W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., Roelfsema, P.R., 2014. Alpha and gamma oscillations characterize feedback

and feedforward processing in monkey visual cortex. Proc. Natl. Acad. Sci. 111, 14332–14341.
Vossel, S., Mathys, C., Daunizeau, J., Bauer, M., Driver, J., Friston, K.J., Stephan, K.E., 2014. Spatial attention, precision, and Bayesian inference: a study of saccadic response speed. Cereb. Cortex 24, 1436–1450.
Vossel, S., Mathys, C., Stephan, K.E., Friston, K.J., 2015. Cortical coupling reflects bayesian belief updating in the deployment of spatial attention. J. Neurosci. 35, 11532–11542.
Wacongne, C., Changeux, J.P., Dehaene, S., 2012. A neuronal model of predictive coding accounting for the mismatch negativity. J. Neurosci. 32, 3665–3678.
Wacongne, C., Labyt, E., van Wassenhove, V., Bekinschtein, T., Naccache, L., Dehaene, S., 2011. Evidence for a hierarchy of predictions and prediction errors in human cortex. Proc. Natl. Acad. Sci. U. S. A. 108, 20754–20759.
Weber, B., Keller, A.L., Reichold, J., Logothetis, N.K., 2008. The microvascular system of the striate and extrastriate visual cortex of the macaque. Cereb. Cortex 18, 2318–2330.
Wiech, K., 2016. Deconstructing the sensation of pain: the influence of cognitive processes on pain perception. Science 354, 584–587.
Yu, X., Qian, C., Chen, D.-y., Dodd, S.J., Koretsky, A.P., 2014. Deciphering laminar-specific neural inputs with line-scanning fMRI. Nat. methods 11, 55–58.