

REVIEW ARTICLE

Linking brain imaging signals to visual perception

ANDREW E. WELCHMAN^{1,2} AND ZOE KOURTZI^{1,2}

¹School of Psychology, University of Birmingham, Birmingham, UK

²Laboratory for Neuro- and Psychophysiology, K.U. Leuven, Leuven, Belgium

(RECEIVED April 4, 2013; ACCEPTED September 17, 2013; FIRST PUBLISHED ONLINE October 29, 2013)

Abstract

The rapid advances in brain imaging technology over the past 20 years are affording new insights into cortical processing hierarchies in the human brain. These new data provide a complementary front in seeking to understand the links between perceptual and physiological states. Here we review some of the challenges associated with incorporating brain imaging data into such “linking hypotheses,” highlighting some of the considerations needed in brain imaging data acquisition and analysis. We discuss work that has sought to link human brain imaging signals to existing electrophysiological data and opened up new opportunities in studying the neural basis of complex perceptual judgments. We consider a range of approaches when using human functional magnetic resonance imaging to identify brain circuits whose activity changes in a similar manner to perceptual judgments and illustrate these approaches by discussing work that has studied the neural basis of 3D perception and perceptual learning. Finally, we describe approaches that have sought to understand the information content of brain imaging data using machine learning and work that has integrated multimodal data to overcome the limitations associated with individual brain imaging approaches. Together these approaches provide an important route in seeking to understand the links between physiological and psychological states.

Keywords: fMRI, Adaptation, pattern classification, 3D shape perception, Perceptual learning

Using human brain imaging to understand the neural basis of behavior

The development of magnetic resonance imaging (MRI) technology since the early 1990s has led to rapid advances in mapping the human brain at high spatial resolution using noninvasive methods. Although significant progress has been made in mapping the functions of different brain regions, linking this understanding of basic brain organization to the fine-grained local computations that underlie complex human behaviors remains a core challenge in neuroscience. Here we discuss the potential for these methods to contribute to understanding the links between perceptual and physiological states within Teller’s linking hypothesis framework (Teller & Pugh, 1983; Teller, 1984). Exploiting brain imaging methods necessitates bringing to mind an additional set of *linking hypotheses* between the measurement signals and the underlying patterns of cortical activity than those originally discussed by Teller and Pugh (1983). In particular, noninvasive brain imaging methods (e.g., fMRI: functional magnetic resonance imaging) involve remote sensing that aggregates the activity of large populations of neurons. Such transforms of neural activity complicate the mapping between perceptual states

and their physiological basis. Although these issues are present—at least in some form—for direct measures of neural activity, and further complicated when making comparisons between species, the indirect nature of the measurements used in human brain imaging makes it important to understand the constraints imposed on linking neural activity to behavior. Moreover, within Teller’s framework, we can frame the notion of equivalence classes in neural responses and their measurement *via* imaging techniques. In particular, it is possible that different sensory inputs might evoke the same neuronal responses, but it is also quite possible that a measured equivalence between inputs is “apparent” in that the measurements are insufficiently differentiable to reveal different underlying neuronal states. This is a familiar problem in psychophysical research that has long grappled with the possibility that equivalent performance on a particular task (for instance, judgments of distance) does not preclude different underlying perceptual states (e.g., depth evoked by different combinations of cues).

fMRI relies on the link between neuronal activity and changes in blood flow in the brain. Such hemodynamic changes (and in particular changes in blood oxygenation) that underlie the blood oxygenation level-dependent (BOLD) response are linked to changes in neuronal activity in an indirect manner. In particular, the chain of events that changes BOLD in response to the metabolic demands of neuronal activity is a complex one involving a cascade of processes (increased synaptic activity, changes in glucose consumption, increase in venous blood volume, and changes in magnetic susceptibility

Address correspondence to: Andrew E. Welchman, University of Birmingham, School of Psychology, Edgbaston, Birmingham B15 2TT, UK. E-mail: a.e.welchman@bham.ac.uk and Zoe Kourtzi. E-mail: z.kourtzi@bham.ac.uk

that can be contrasted for BOLD) that are not fully understood (Heeger & Ress, 2002). Moreover, BOLD changes caused by the metabolic requirements of neurons are slow in relation to neuronal activity, typically taking 4 to 6 s after changes in neuronal activity.

Further uncertainty surrounds the neuronal events that give rise to BOLD changes. In particular, the fMRI signal almost certainly not only reflects the firing rates of very large numbers of neurons but is likely to include subthreshold (nonspiking) activity, pooled excitatory and inhibitory neural signals and feedback from higher portions of the processing stream (Logothetis et al., 2001; Mukamel et al., 2005; Viswanathan & Freeman, 2007). Using Teller's terminology, there is a large *equivalence class* of fMRI responses in that the same BOLD response could be evoked by large changes in a small number of active neurons or small changes in a large number of neurons (Scannell & Young, 1999). Although BOLD responses to sensory stimuli are thought to follow reliably from the linear summation of neural signals across space and time (Boynton et al., 1996), the relationship between changes in neuronal activity and changes in BOLD may reflect a complex nonlinear function: small changes in activity may evoke large changes in BOLD, but responses may become saturated at higher response levels, such that changes in neuronal activity based on changes in blood oxygenation become difficult to measure (Miller et al., 2001). Further, task-related components (e.g., anticipation, awareness, and attention) introduce nonlinear influences that disrupt the correspondence between BOLD and neural activity (Boynton, 2011; Cardoso et al., 2012).

Despite these concerns, there appears to be a reasonable range over which stimulus-evoked BOLD responses can be characterized using a linear transducer function that assumes fMRI signals reflect the local average of neuronal activity. For instance, Boynton et al. (1996) tested the dependence of visually evoked BOLD responses on the (a) presentation duration and (b) stimulus contrast, finding evidence that these dimensions are separable and importantly that hemodynamics are similar at different contrast levels. Thereafter, Boynton et al. (1999) were able to test for similarities (i.e., Tellerian *analogies*) between fMRI-measured contrast response functions and psychophysical contrast increment thresholds. By systematically varying stimulus contrast, they were able to show similarities between the contrast response functions in early visual areas and psychophysical judgments of contrast increments. This link between activity in early portions of the visual processing stream and psychophysical task performance was further strengthened by Ress and Heeger (2003), who contrasted fMRI responses evoked while participants performed a demanding contrast-detection task. They found that higher fMRI responses were evoked when participants reported detecting a target, whether the target was present or not. Specifically, BOLD responses were higher when the target was presented and detected (i.e., a "hit"), and when the signal was absent, but the participant reported seeing it (i.e., a "false alarm"); by contrast, "correct rejections" (stimulus absent, reported to be absent) and "misses" (stimulus present but not reported) gave rise to lower responses. This result is particularly elegant in within Teller's framework (Teller, 1984) in that it provides an illustration of a case in which the same perceptual state (e.g., believing a target is present) is associated with similar fMRI responses (i.e., a *converse proposition*), as well as demonstrating that different perceptual states give rise to different responses, despite the physical similarity of the viewed stimulus (i.e., a *contrapositive proposition*).

Thus, despite important caveats on the uncertainty of the mapping between neuronal activity and fMRI measurements in particular, and brain imaging signals more generally, manipulations of simple

sensory features (e.g., contrast), whose physiological bases have been studied extensively, provide good reasons to believe that these methods provide a basis from which relationships between physiological and psychological states can be tested. From this foundation, we review work that has sought to understand the link between human neuronal activity and perception for higher-level visual features. We focus on studies using methodological approaches (i.e., fMRI adaptation, multivoxel pattern classification) that are thought to be more sensitive than conventional fMRI analysis methods to weak, but informative, brain signals related to neural preference. We compare fMRI signals revealed by these methods to behavioral performance to test for links (analogies) between human brain signals and behavior. To illustrate these links, we review studies examining the link between (1) brain activity and 3D perception, (2) brain plasticity and visual learning.

Linking fMRI responses to the perception of 3D shapes

Understanding the neural basis of mid-level visual processes requires the careful selection of stimulus conditions as part of the fMRI design. In particular, identifying particular aspects of functional processing requires that we can identify a contrast between different conditions. In practice, it can be very difficult to find such a contrasting stimulus, so forms of spatial scrambling of the image features are often employed (e.g., Kanwisher et al., 1996; Georgieva et al., 2008). This approach has been successful when localizing cortical activity related to high-level visual features; however, low-level image differences can constitute a concern. For example, do responses in a particular area relate to the aspect we would like to study (e.g., 3D shape from shading)? Or is contrasting activity evoked because the different stimuli differ in their low-level statistics? If low-level image features differ between "depth" and "nondepth" stimuli, it can be difficult to determine whether an fMRI response relates to low-level feature differences or the depth percept (see Georgieva et al., 2008 for a systematic investigation of lower-level image changes related to shading cues).

An alternative approach is to make parametric stimulus manipulations and evaluate changes in the fMRI response to look for similarities between changes in the fMRI response and psychophysical judgments (Chandrasekaran et al., 2007; Georgieva et al., 2009). Chandrasekaran et al. (2007) used such an approach to investigate the perception of 3D shape defined by disparity. In particular, they presented observers with random dot patterns (Fig. 1a) to ensure that the 3D structure of the display was only apparent after global binocular correspondence had been established. To test for cortical areas whose activity related to processing of disparity signals for 3D shape perception, shapes were systematically degraded to measure the effects on both fMRI response and perceptual discrimination. Specifically, the spatial integrity of the disparity-defined shapes was parametrically manipulated by randomly relocating the spatial positions of dots with a given disparity in the display (Fig. 1a). This affected observers' performance in judging the global 3D shape (Fig. 1b) while keeping the overall histogram of dot disparities unchanged.

The experimental logic of the study of Chandrasekaran et al. was that reliable perceptual judgments of disparity-defined shape would be associated with strong fMRI responses in areas that encode disparity signals for 3D shape. The study demonstrated different patterns of results in different cortical areas. In particular, responses in early visual areas were largely unaffected by the manipulation of shape coherence (Fig. 1c). This suggests a local processing of

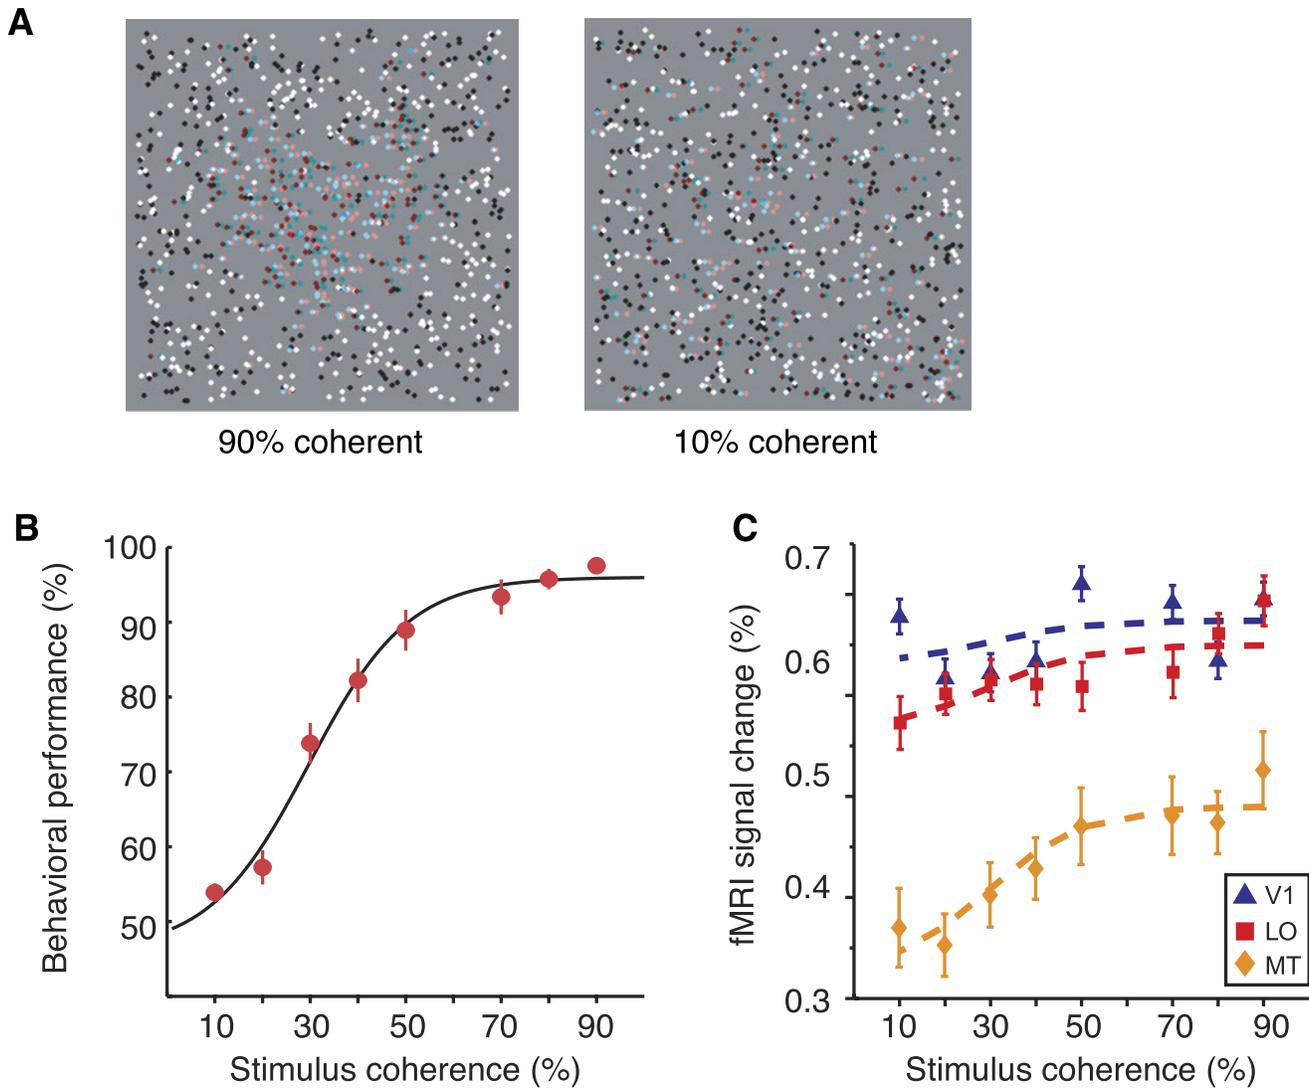


Fig. 1. Linking fMRI and behavioral judgments of 3D shape. (A) Stimuli: RDS of a 3D shape at different levels of spatial coherence (90%, 10%) for use with red–green anaglyph glasses. RDS provide a means of isolating disparity signals from monocular cues to shape. Manipulating the spatial coherence of the disparities in the shape makes it harder to see the global 3D shape depicted. (B) Psychometric function showing performance in judging the properties of the 3D shape (here it is axis of symmetry: horizontal vs. vertical) as a function of stimulus coherence. (C) fMRI responses in three representative areas of interest (V1, MT, and LO) as a function of stimulus coherence. The psychometric function shown in part B is scaled to fit the fMRI response.

disparity signals in early areas that is blind to the global coherence of the shape. In contrast, responses in higher ventral (lateral occipital: LO) and dorsal (hMT+/V5) areas increased significantly as the shapes became more coherent, suggesting the encoding of more extensive surfaces important for shape judgments. Further support for a link between activity in these areas and perceptual judgments of 3D shape was provided by significant correlations between individual observers' psychophysical judgments and individual mean fMRI responses. While correlating average perceptual judgments with mean fMRI responses suggests a link between neural activity and behavior, this coupling is more compelling if provided on a trial-by-trial basis. The approach of exploiting between-trial variation is well established in macaque electrophysiology (Britten et al., 1992), and was adopted by Chandrasekaran et al. for the analysis of a subset of fMRI responses. In particular, they asked whether it was possible to determine the 3D shape viewed by a participant based only on fMRI responses within a particular area of cortex. Using

the distributions of fMRI responses obtained across trials for different types of 3D shape (horizontal vs. vertical axis of symmetry), they computed receiver operating characteristic (ROC) curves, revealing fMRI responses diagnostic of the viewed shape in cortical areas V3, V3B/KO, hMT+/V5, and LO. This analysis did not allow them to dissociate the stimulus from its perceptual interpretation (Bradley et al., 1998), as wholly ambiguous stimuli were not employed. However, the idea of predicting the viewed stimulus, given only the fMRI response, has been explored further in subsequent work using machine learning fMRI data analysis, which we return to below.

fMRI measurement and analysis—unpicking apparent “equivalence classes” in responses

fMRI measurements are crude in relation to the underlying neural responses that they seek to quantify. Basic units of measurements

(fMRI voxels) are typically $3 \times 3 \times 3$ mm in size and therefore pool the activity of (at least) hundreds of thousands of individual neurons. This limited measurement resolution means that the neuronal responses evoked by two different classes of stimuli may be equivalent at the fMRI measurement level but clearly separable at the scale of individual neural responses (Fig. 2a). For instance, consider the problem of identifying whether a visual area responds to two different cues (disparity and motion) on the basis of processing each cue separately or conjointly (Preston et al., 2009; Ban et al., 2012). The two cues might be processed by independent neuronal populations that lie within the same fMRI voxel. By averaging the responses of these subpopulations, the voxel response might appear identical when the different depth cues are presented. This could be taken to suggest a convergence of depth signals, when in actuality, information is processed in parallel and follows different routes through the cortex (Ponce et al., 2008). Thus, there might be an apparent *equivalence class* in that fMRI measurements give rise to similar patterns of activity, but actually the physiological basis of that activity differs (see Ban et al., 2012 for a means of separating these different hypotheses for this example).

One means of addressing this potential for apparent equivalence between responses is to alter the fMRI design to maximize sensitivity to different responses by subpopulations. One approach to doing this is based on the logic of neural adaptation and repetition suppression effects whereby neural activity is lower for stimuli that have been viewed recently than for stimuli that have not (Grill-Spector & Malach, 2001; Kourtzi & Kanwisher, 2001). In particular, the premise of this approach is that showing the same stimulus twice (“test–test”) results in an adapted fMRI response, while pairing the test stimulus with a stimulus that engages a different set of neurons (“test–comparison”) will not produce an adapted response (Fig. 2b). The difference in the fMRI response between these two conditions (“the rebound effect”) thus acts as an index of the sensitivity of the underlying neural population to differences between the test and comparison stimuli. Although the direct link between fMRI adaptation and neural selectivity is complicated by the complexity of the BOLD signal (Krekelberg et al., 2006; Sawamura et al., 2006), adaptation has been a useful tool for discerning large neural populations as revealed by fMRI (Grill-Spector & Malach, 2001; Grill-Spector et al., 2006).

For example, imagine presenting a 3D shape defined by disparity and motion. If an area responded to the information specified individually by disparity and motion, we would have expected a rebound effect whereby changes in the component cues gave rise to higher fMRI responses. In contrast, an area that encoded 3D shape would show no rebound effect if changes in the individual cues were such that the impression of 3D shape remained constant. Welchman et al. (2005) used such an approach to study 3D shape from disparity and perspective information, introducing small differences between the individual cues to effect changes in individual cues and changes in the estimate of shape based on combined cues. Thereby, they were able to demonstrate responses in higher ventral area Lateral Occipital Complex (LOC) that related to the 3D shape based on integrated cues. Based on the small differences between the presented stimuli, an approach based on simply contrasting fMRI responses from different conditions would have been unlikely to be sufficiently sensitive.

A second approach to resolving the potential for mistaken equivalence between measured responses is based on adopting more sensitive fMRI analysis methods. In particular, data treatment often involves fairly extensive spatial smoothing, exacerbating the problem of differentiating the convergence of signals from separate encoding by independent neuronal subpopulations described above.

Moreover, fMRI activity within a particular region of cortex is often treated as equivalent, and thus quantified on the basis of the grand average (i.e., univariate) response. Adopting multivariate analysis methods for fMRI data (Cox & Savoy, 2003; Haynes & Rees, 2006; Norman et al., 2006) is therefore useful in allowing a finer-scale analysis of fMRI data that can differentiate responses that would appear equivalent based on average responses across a cortical area (Fig. 2c). There is an ongoing debate about the origin of the preferences that are exploited by multivoxel pattern classification analysis (MVPA) approaches, with contributions likely from random spatial irregularities in fine columnar maps (Swisher et al., 2010) as well as coarser scale signals that reflect the systematic (i.e., retinotopic) organization of stimulus responses to properties such as orientation and motion direction (Raemaekers et al., 2009; Freeman et al., 2011; Beckett et al., 2012). This debate reflects the complex correspondence between BOLD and neural selectivity (Kriegeskorte et al., 2010) and emphasizes the importance of work comparing directly physiology to imaging signals. Nevertheless, these techniques have increased the analysis sensitivity to help unpick apparent equivalence classes of responses.

Decoding approaches to test the link between fMRI responses and the perception of depth from disparity

Understanding the link between physiological responses and perceptual processes represents a considerable challenge for visual features that are encoded in a widespread manner in the visual cortex. For instance, neurons sensitive to binocular disparity information have been found throughout visual, temporal, and parietal regions (Cumming & DeAngelis, 2001; Orban et al., 2006; Parker, 2007), making it a challenge to determine the functional significance of such widespread representations.

Using the multivoxel analysis approach, Preston et al. (2008) sought to test the relationship between fMRI responses and the perception of disparity-defined depth, using a simple low-level stimulus manipulation. In particular, they rendered depth planes with different disparities in random dot stereograms (RDS) and altered the visibility of the disparity-defined form by manipulating the contrast of the dots presented to the two eyes (Fig. 3a). In cases where the contrast of the dots presented to each eye was the same (correlated RDS: a white dot in the left eye matches a white dot in the right eye), observers had no difficulty in determining the depth position of a target plane. However, reversing the contrast of dots in the two eyes (anticorrelated RDS: a black dot in the left eye matches a white dot in the right eye) abolished the perception of disparity-defined depth (Cogan et al., 1995; Cumming et al., 1998). Nevertheless, the disparity information contained in anticorrelated RDS stimuli is encoded by neurons in primary visual cortex of the macaque (Cumming & Parker, 1997). This simple manipulation therefore provided a tool with which to identify responses to binocular disparity that were *analogous* to perceptual judgments (see also Bridge & Parker, 2007 and Cottreau et al., 2011, 2012).

Preston et al. (2008) manipulated stimulus correlation to test for areas that contained information about depth structure only when observers could perceive the disparity-defined structure of the stimuli. They investigated a basic property of disparity encoding: whether a target is positioned near (crossed disparity) or far (uncrossed disparity) with respect to the fixation point (Fig. 3b). Measuring the fMRI responses evoked by crossed *versus* uncrossed disparities using both correlated and anticorrelated stimuli allowed them to determine which cortical areas represent disparity signals

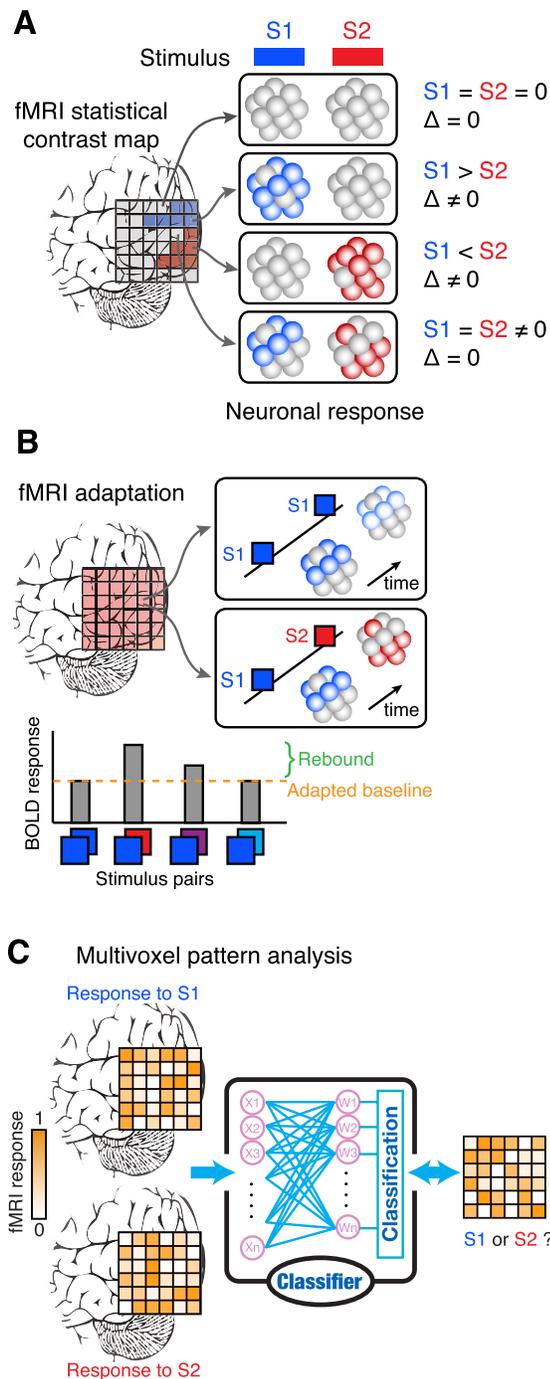


Fig. 2. Illustrations of the potential for apparent equivalence in responses at the fMRI voxel level. **(A)** Illustration of the standard fMRI contrast analysis approach. Four different scenarios are presented for the response of neuronal populations within an individual voxel (represented as a cluster of spheres) to two different stimuli: S1 and S2. Scenario 1: the neuronal population does not respond to either stimulus, so a contrast (correctly) suggests no differential activity. Scenario 2, 3: the neuronal population responds preferentially to S1 or S2 respectively, with the fMRI contrast revealing these differential selectivities. Scenario 4: the voxel contains a mixed population—some neurons respond to S1 and some to S2; the response of the voxel to both stimulus classes means the contrast yields no difference, a result that is identical to scenario 1. **(B)** The fMRI adaptation approach can help identify mixed populations within individual voxels. The logic is that presenting the same stimulus twice will evoke a lower neuronal response, and therefore lower BOLD response, relative to presenting two different stimuli that

that are constrained by matching features with the same luminance contrast between the eyes. Based on results from neurophysiology, neurons sensitive to crossed and uncrossed disparities are closely collocated within cortex (DeAngelis & Newsome, 1999). Therefore, a standard fMRI analysis that compared responses to targets in front of *versus* behind the fixation point would be unlikely to yield reliably different responses as neurons selective to near and far stimuli would be collocated within the same voxels. Thus, Preston et al. designed their study for multivariate data analysis and trained a support vector machine (SVM) to associate the pattern of fMRI responses evoked by RDS that depicted depth positions behind or in front of the fixation plane. Then, using an independent data set, they tested how well the SVM classifier could determine what the participant had been looking based on their fMRI data (Fig. 2c). The classifier’s performance in discriminating whether the subject had viewed a near or far stimulus thus provided an index of the information contained in a particular region of cortex.

Decoding accuracies suggested a differential pattern of responses for correlated and anticorrelated stimuli across the visual hierarchy. The classifier’s ability to predict the viewed stimulus was significantly above chance across the visual cortex when subjects viewed correlated stimuli (Fig. 3c). This is consistent with previous imaging and neurophysiological data that shows disparity-selective responses widespread throughout cortex. However, the same was not true for anticorrelated stimuli. In early visual areas, the classifier could also decode the disparity signals contained in anticorrelated stimuli. This is consistent with electrophysiological recordings in suggesting dissociation between early cortical responses and the perception of disparity-defined depth (Cumming & Parker, 1997). Results in higher areas were somewhat surprising in that activity in intermediate areas in the ventral stream (V3v, V4), such as early visual cortex, supported similar decoding performance for correlated and anticorrelated stimuli. In contrast, responses in higher ventral LO area supported significant decoding of correlated, but not anticorrelated, stimuli. This suggests that, like perceptual judgments, responses in LO are selective for binocular matches that involve similar features in the two eyes. This finding is in agreement with neurophysiological recordings in inferior temporal cortex (analogous to LO), where responses to anticorrelated stimuli are not observed (Janssen et al., 2003).

Preston et al. also found that responses in the dorsal visual cortex supported decoding that was preferential for correlated over anticorrelated responses, indicating multiple loci in which the constraint of binocular matching by contrast similarity is implemented. However, a further experiment in which they manipulated disparity parametrically suggested different types of responses in higher dorsal *versus* ventral cortex. In particular, their data suggested that dorsal areas had responses that were quite specific to the disparity being

stimulate different neuronal populations. This rebound effect can be used to index the way in which the underlying neuronal population encodes stimulus information: different stimuli that are encoded by the same populations will give rise to a minimal rebound effect. **(C)** The multivoxel analysis approach. Information about the stimulus is distributed across the pattern of responses of individual voxels: some voxels will show similar responses to the different stimuli, others will give rise to subtly different responses. The machine learning algorithm learns which features are informative in separating the two stimulus classes. Once trained, the classifier is tested by predicting the stimulus for a voxel pattern that it was not trained on. The prediction performance of the classifier is typically used to index the amount of information about the stimulus that is contained in the pattern of voxel responses.

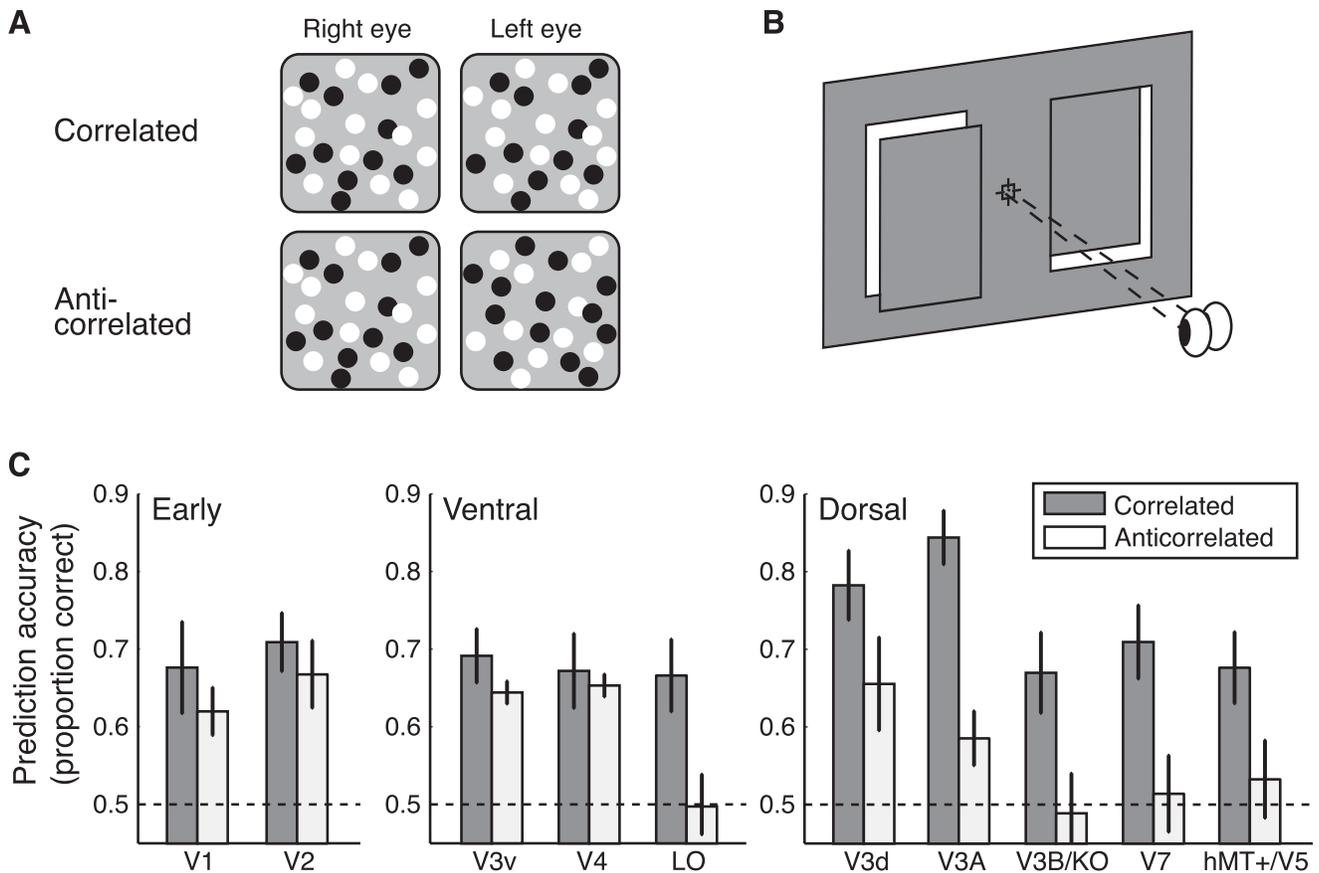


Fig. 3. Decoding disparity signals contained in correlated and anticorrelated RDS. (A) An illustration of correlated random dot patterns (white dots match with white dots between the two eyes, and black dots with black dots) and anticorrelated patterns (white dots match with black dots and *vice versa*). Free-fusing the stereopairs (by crossing the eyes) yields the impression of a central square region in front of a background for the correlated pattern only; depth is incoherent for the anticorrelated pattern. (B) An illustration of the crossed *versus* uncrossed disparity configuration used in the study by Preston et al. (2008). (C) The pattern of decoding accuracies across visual cortex for correlated and anticorrelated disparity stimuli. The bar graphs represent the prediction accuracy of the classifier, where chance is at 0.5 (binary classification task). Error bars show SEM.

viewed, whereas ventral responses were less specific to the particular disparity of the scene, but rather might represent the local stimulus configuration. Thus, responses analogous to perceptual interpretations were widespread, but further testing indicated different forms of representation in dorsal *versus* ventral visual areas. Reflecting Teller, this highlights the need for multiple explorations of the possible parameter space of responses when seeking to understand the links between measures of human cortical activity and perceptual states.

Linking brain plasticity to human learning

Next, we consider evidence for links between perceptual states and brain responses through studies that show how these can change through learning. Learning is known to result in improved behavioral performance in a variety of perceptual and motor tasks, but what are the concomitant changes in brain states that support it? Brain imaging provides a useful tool with which to assess the neural signatures that support learning-dependent behavioral improvements, with recent work demonstrating both functional and structural brain changes as a result of training (Zatorre et al., 2012). Frequently, links between behavioral improvements and brain changes have been quantified using correlations of the magnitude of functional

or structural brain changes and improvement in task performance after training. Although such approaches can provide insight into the relationship between brain signals and behavior, they are often limited by small sample sizes, outlier observations, and bias in data sampling (Vul & Pashler, 2012). Here, we discuss recent work using multivoxel pattern classification methods to investigate the link between BOLD and learning-dependent changes in neural tuning and behavior. These methods rely on cross-validating patterns of results based on independent data set and therefore provide a sensitive tool for uncovering neural preferences and studying the link between brain patterns and behavior in an unbiased manner.

Learning-dependent changes in neural signals as revealed by fMRI

Understanding the link between behavioral improvement with training and learning-dependent changes in fMRI signals is fraught with complexity. First, learning-dependent changes measured by BOLD may reflect different underlying changes in neural selectivity (i.e., sharpening of neuronal tuning to a visual stimulus): there may be enhanced responses to the preferred stimulus, decreased responses to nonpreferred stimuli, or a combination of the two. This clearly complicates the inferences that can be drawn from fMRI data; that is *apparent*

equivalence in responses might have different underlying causes. A number of studies have shown that learning changes the overall responsiveness (i.e., increased or decreased fMRI responses) of visual areas to trained stimuli (Kourtzi et al., 2005; Sigman et al., 2005; Op de Beeck et al., 2006; Mukai et al., 2007; Yotsumoto et al., 2008). However, considering only the overall magnitude (i.e., univariate fMRI signal) does not allow us to discern whether learning-dependent changes in fMRI relate to changes in the overall magnitude of neural responses or changes in the selectivity of neural populations.

A second difficulty relates to making comparisons between overall activity between sessions—that is for brain imaging measurements made before *versus* after training. Such differences may have generalized causes such as differences in the state of the observers (e.g., alertness, attention, and interest) or differences in the measurement equipment (e.g., different levels of scanner noise) rather than training *per se*. Therefore, it is preferable to compare brain activity between different stimulus conditions within each session, using an experimental design that looks for changes in the relative response to different conditions before *versus* after training to index underlying changes in neural representations due to training.

Third, to quantify changes in (1) perceptual performance and (2) brain responses due to learning, it is useful to vary stimuli parametrically and test changing levels of performance. However, as discussed above, simple univariate fMRI measures may be insufficient to understand subtle changes in the neuronal responses and the links between perceptual and neuronal states. Recent work has sought to increase analysis sensitivity using MVPA to test training-dependent changes in the neural representations of visual stimuli (Op de Beeck et al., 2006; Zhang et al., 2010; Jehee et al., 2012). Next, we describe work that has looked for Tellerian analogies between psychometric and brain imaging measurements by comparing the choices of human observers and pattern classifiers to derive similar functions from behavioral and fMRI data.

For example, in recent work Zhang et al. (2010) tested the mechanisms that mediate visual perceptual learning by combining behavioral judgments with high resolution imaging (i.e., smaller voxels for finer-grained measurements) and MVPA tools (Haynes & Rees, 2005; Kamitani & Tong, 2005). To gain insight into the way in which responses to visual forms change with training, they employed a parametric stimulus space in which they could morph systematically between two different stimulus classes (concentric *vs.* radial glass patterns; Fig. 4a). By adding external noise to the displays, and using fine-scale variations in the morphing space, they were able to characterize observers' behavioral discrimination performance before and after training on the task (Fig. 4b). Using fMRI measurements concurrent with stimulus presentation, Zhang et al. (2010) were able to evaluate the ability of the machine-learning algorithm to decode the presented stimuli. They were particularly interested in the choices of the pattern classifier in predicting each of the presented stimuli. Using the distribution of choices for each stimulus, they were able to generate fMRI-based voxel tuning functions that described the distribution of choices made by the classifier when given data were measured in different visual areas (Fig. 4c). Thereby, they sought to link learning-based changes in behavioral responses and fMRI responses by comparing psychometric functions and fMRI pattern-based tuning functions before and after training.

Comparing the performance of human observers and classifiers demonstrated that learning altered the observers' sensitivity to visual forms and the tuning of fMRI activation patterns in visual areas selective for task-relevant features. For high signal stimuli (i.e., 80% of the dots in the display were aligned to the stimulus

shape, whereas 20% were randomly positioned), a tuned response across visual areas was observed, that is, the classifiers mispredicted similar stimuli more frequently than dissimilar ones. Consistent with the behavioral results, these tuning functions had higher amplitude for high-signal stimuli than low-signal stimuli (i.e., only 45% of the dots in the display were aligned to the stimulus shape) before training in higher visual areas (Fig. 4c). However, training on low-signal stimuli enhanced the amplitude while it reduced the width of pattern-based tuning functions in higher dorsal and ventral visual areas (Fig. 4c). Increased amplitude after training indicates higher stimulus discriminability that may relate to enhanced neural responses for the preferred stimulus category at the level of large neural populations. Reduced tuning width after training indicates fewer classification mispredictions, suggesting that learning decreases neural responses to nonpreferred stimuli. Thus, these findings suggest that learning of visual patterns is implemented in the human visual cortex by enhancing the response to the preferred stimulus category while reducing the response to nonpreferred stimuli.

Linking BOLD changes to behavioral improvement after training

Learning is known to improve performance in perceptual tasks when there is considerable uncertainty, such as when detecting stimuli presented in cluttered backgrounds or discriminating between highly similar stimuli (Goldstone, 1998; Gilbert et al., 2001; Goldstone et al., 2001; Fine & Jacobs, 2002). These training effects can be manifested in two different ways: training can improve behavioral sensitivity for target detection (i.e., learning makes psychometric functions steeper) or alter the decision criterion that represents the perceptual boundary between stimulus categories (i.e., learning can shift the midpoint, or bias, of the psychometric function). To investigate the brain mechanisms that support such changes, Li et al. (2009) tested how training paradigms altered the pattern-based analysis of single-trial fMRI signals. In particular, they employed a categorization task in which observers were presented with stimuli from a morphing space generated by interpolating between radial and concentric glass patterns (Fig. 4a). Observers were asked to decide whether the viewed stimulus was radial or concentric. Using feedback, they were able to alter the location of the observers' categorical boundary in the morphing space (Fig. 5a), thereby dissociating the physical stimuli from their categorical interpretation. Using fMRI, they asked which cortical regions carry the neural signature of this learning-dependent flexibility in categorical decisions. Li et al. (2009) reasoned that such regions would show trial-by-trial variations that correspond to the behavioral choice of the observers and change to reflect the shift in the decision criterion (i.e., category boundary).

To test this prediction, Li et al. (2009) computed psychometric functions based on the observers' behavioral choice on each stimulus trial and "fMR-metric" functions based on the classifier's choice on single-trial fMRI signals (Fig. 5b). In particular, they trained a linear SVM to classify fMRI signals based on the observer's behavioral choice (radial *vs.* concentric) on each trial and evaluated the SVM's prediction performance in predicting the observers' choice for an independent dataset. Comparing the classifier's choices with the observer's choices showed that fMR-metric functions resemble psychometric functions, suggesting a link between behavioral and neural responses (Fig. 5b). Importantly, Li et al. (2009) showed that learning-dependent changes in the behavioral decision criterion (i.e., categorical boundary) are reflected by changes in the threshold

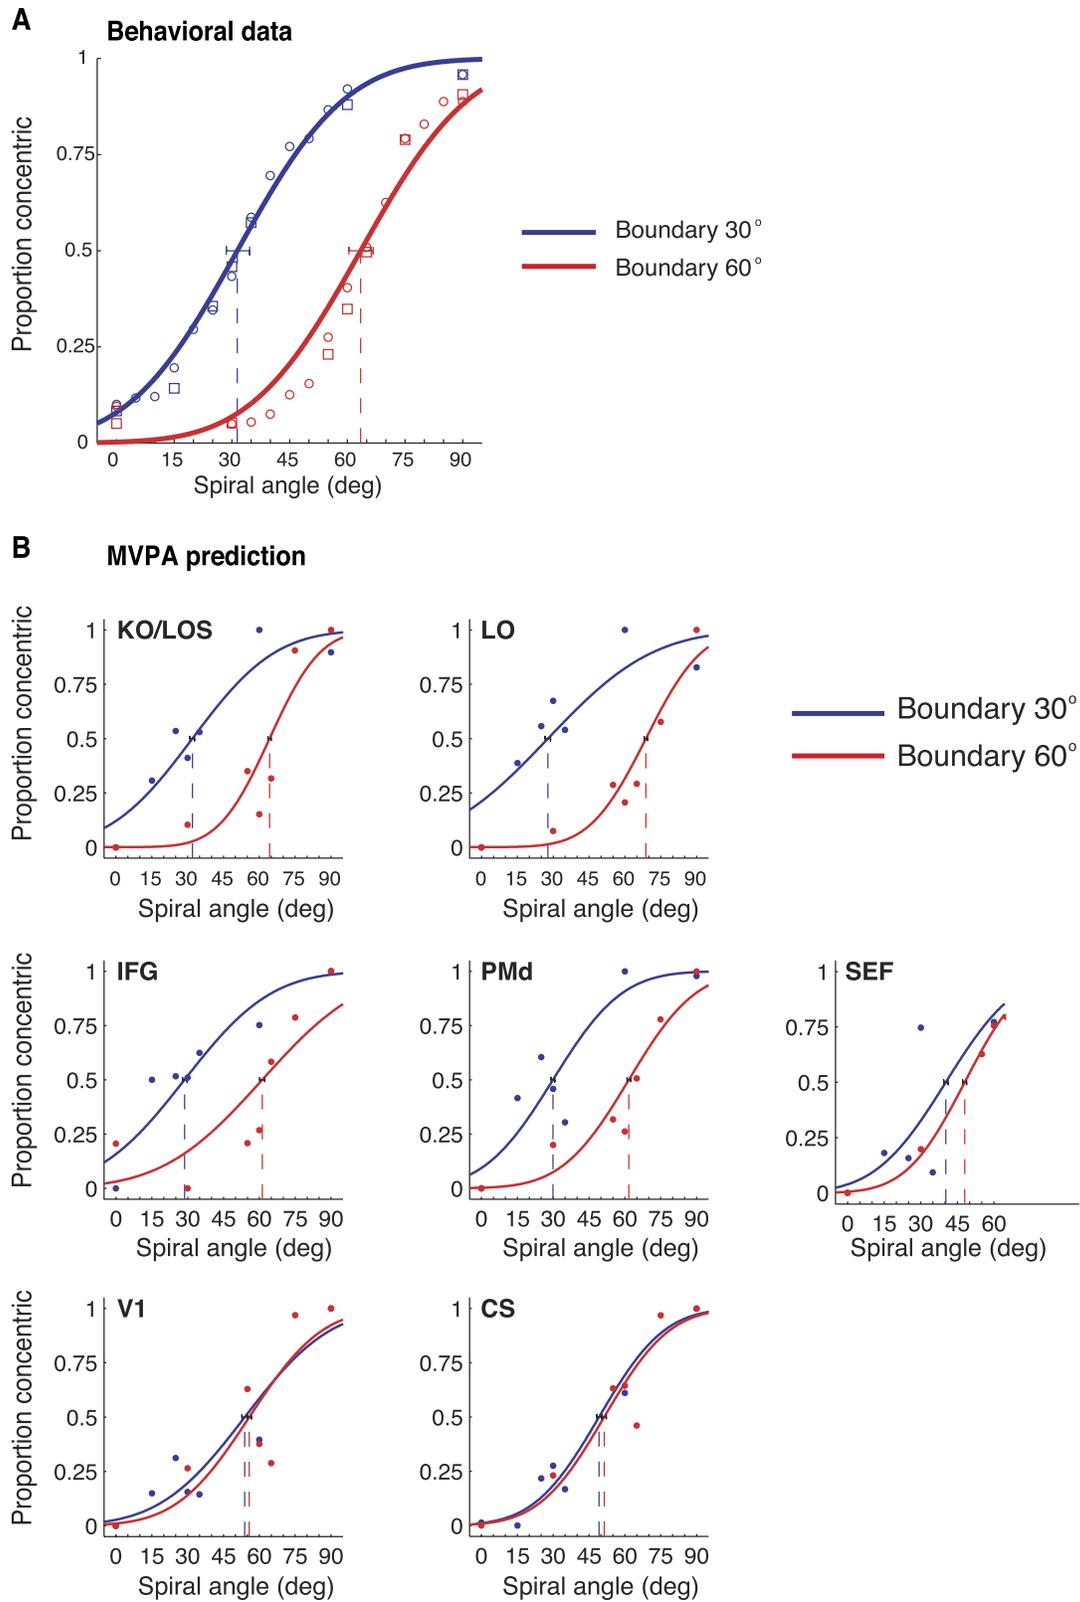


Fig. 5. Learning for flexible categorical decisions. **(A)** Behavioral data for each of the two trained category boundaries (30 and 60° boundary). Lines indicate the cumulative Gaussian fits of the behavioral data from the lab. Error bars indicate the 95% confidence interval at 50% concentric threshold. **(B)** fMR-metric curves based on the observers' behavioral choice for each task (30 and 60° boundary) and ROI. Average classifier prediction data across observers are scaled from 0 to 1 and fitted with cumulative Gaussian functions. Error bars indicate the 95% confidence interval at 50% concentric threshold.

(50% point) of the fMR-metric functions in ventrolateral prefrontal, premotor, and higher occipitotemporal regions. Analysis of the slopes of the psychometric and fMR-metric functions suggested that learning changes processes related to decision variables (criterion for stimulus comparison, stimulus uncertainty relative to criterion). In contrast, no significant shift was observed in the fMR-metric functions in V1, CS, or SEF suggesting that learning a new criterion for categorization does not affect the physical stimulus- or response-related representations in these regions that may support signal detection and motor execution, respectively.

Thus, comparing psychometric and fMR-metric functions allows to reliably decode (a) the observers' choice on single trials and (b) learning-dependent changes on the decision criterion used by the observers for classifying sensory input. These findings provide evidence that category learning shapes neural representations in a circuit involving frontal and higher occipitotemporal regions to reflect the observers' behavioral choice during categorical decisions. This link between behavioral and brain changes due to learning was supported by two further analyses methods.

First, fitting the fMRI data using a scaled version of the psychometric function obtained during scanning showed that fMR-metric functions in frontal (IFG, PMd) and extrastriate visual (KO/LOS, LO) areas (but not SEF, V1 or CS) were shifted in correspondence with the learned categorization boundary. Second, Li et al. (2009) used a choice probability analysis [e.g., (Britten et al., 1996; Uka & DeAngelis, 2003)] to quantify the relationship between the observers' choice and fMRI responses. They used voxel pattern responses (i.e., classification accuracy) evoked by stimuli near threshold performance (i.e., stimuli at the boundary), as observers' performance included a sufficient number of errors for these stimuli. In particular, for each observer, trials were labeled based on the observer's behavioral choice and the classifier's performance in predicting the observer's choice per stimulus trial was tested on an independent data set. The area under this ROC curve signified the choice probability for each observer and cortical region. Significant choice probabilities were observed in frontal and higher occipitotemporal areas, suggesting that pattern-based fMRI responses in these regions were strongly analogous to observers' behavioral choices.

Comparing spatiotemporal brain patterns to human behavior

The work so far reviewed has focused on identifying spatial brain patterns (i.e., brain regions) that change with learning, implicating occipitotemporal and frontoparietal areas. However, little is known about the interactions within this network that mediate learning-dependent improvement in complex perceptual tasks. Some studies argue that learning alters early sensory processing (Adini et al., 2002; Teich & Qian, 2003), whereas others propose that learning alters later decision-related processes (Doshier & Lu, 1999; Li et al., 2004; Law & Gold, 2008; Jacobs, 2009).

Using fMRI alone would make it difficult to identify cortical circuits related to different temporal processes involved in visual learning due to the low temporal resolution of the technique. In some recent work, Li et al. (2012) and Mayhew et al. (2012) took advantage of the complementary high spatial and temporal resolution of simultaneous electroencephalogram (EEG)-fMRI to identify the spatiotemporal dynamics between cortical networks involved in perceptual learning under uncertainty. Using EEG-informed fMRI and pattern classification analysis methods, they tested for learning-dependent changes in EEG-fMRI activation patterns that related to the observers' enhanced performance in

discriminating global forms (i.e., radial vs. concentric patterns) after training. These multimodal measurements allowed testing the hypothesis that changes in neural representations due to training may occur at different timescales related to different computational processes (e.g., stimulus processing vs. perceptual judgment).

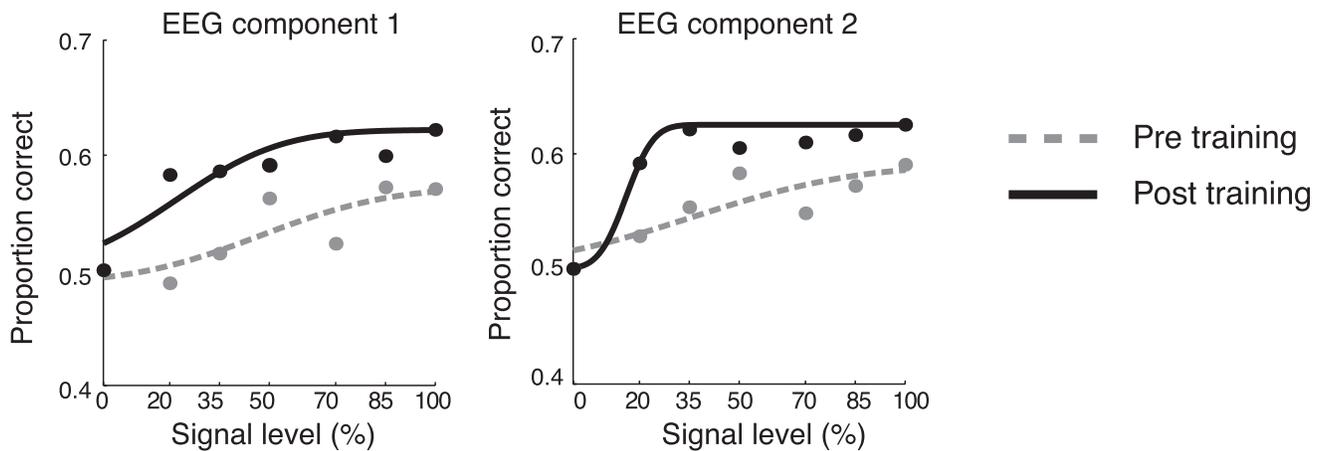
To test this hypothesis, Li et al. (2012) and Mayhew et al. (2012) used EEG to identify components with distinct latencies that are thought to relate to different processes (Johnson & Olshausen, 2003; Ohla et al., 2005; Pei et al., 2005; Tanskanen et al., 2008; Das et al., 2010): a first component around 120 ms known to relate to visual form integration and a second component around 240 ms known to relate to fine perceptual judgments. To identify brain regions that are involved in the different temporal processes related to these EEG components, they conducted an EEG-informed fMRI analysis, using the EEG time course to generate regressors corresponding to the two different components for the General Linear Model (GLM) analysis on the fMRI signals (Debener et al., 2005; Eichele et al., 2005; Philiastides & Sajda, 2007). To test the link between spatiotemporal brain activations (i.e., brain areas activated for each of the two EEG components) and behavior, they measured the choices of a pattern classifier when predicting each stimulus from EEG-fMRI signals. That is, similar to the fMR-metric functions, Li et al. (2012) and Mayhew et al. (2012) generated EEG-metric functions (Philiastides & Sajda, 2006; Das et al., 2010) for each of the two EEG components. They tested whether decoding radial *versus* concentric glass patterns from single-trial EEG data improved after training.

Comparing the choices of the classifier algorithm (EEG/fMR-metric functions) with the observer's choices (psychometric functions) demonstrated that learning improves our ability to discriminate visual forms and EEG-fMRI activation patterns related to distinct visual recognition processes. In particular, when observers discriminated radial from concentric patterns in noise, training improved observers' sensitivity. This behavioral improvement was associated with neural changes in (a) early processes involved in the integration of global forms engaging occipitotemporal and posterior parietal areas and (b) later processes related to categorical judgments engaging frontal circuits (Fig. 6a). In contrast, when observers performed categorical judgments on morphed stimuli (i.e., radial vs. concentric patterns), training altered the decision criteria (i.e., category boundary). These learning-dependent changes in the decision criteria were associated only with changes in later EEG processes engaging higher occipitotemporal and frontoparietal circuits: category learning did not modulate early processes in the medial frontotemporal network that are thought to support the coarse interpretation of visual scenes (Fig. 6b).

Summary

Teller's formalized framework (Teller & Pugh, 1983; Teller, 1984) for understanding the links between perceptual and physiological states was originally discussed in relation to the activity of small numbers of neurons. Here we reviewed the potential contributions from human brain imaging in understanding the links between neuronal responses and visual perceptual states. These imaging approaches aggregate information from large numbers of neurons and involve complex transducer functions that necessitate careful considerations in their interpretation. Nevertheless, there are good reasons to believe that human imaging data can provide a mid-level analysis of physiological processes that can link to single-unit recording work as well as facilitating the use of perceptual tasks whose complexity approximates many of the challenges of everyday human vision. In particular, we have

A Detection in noise



B Categorical judgments

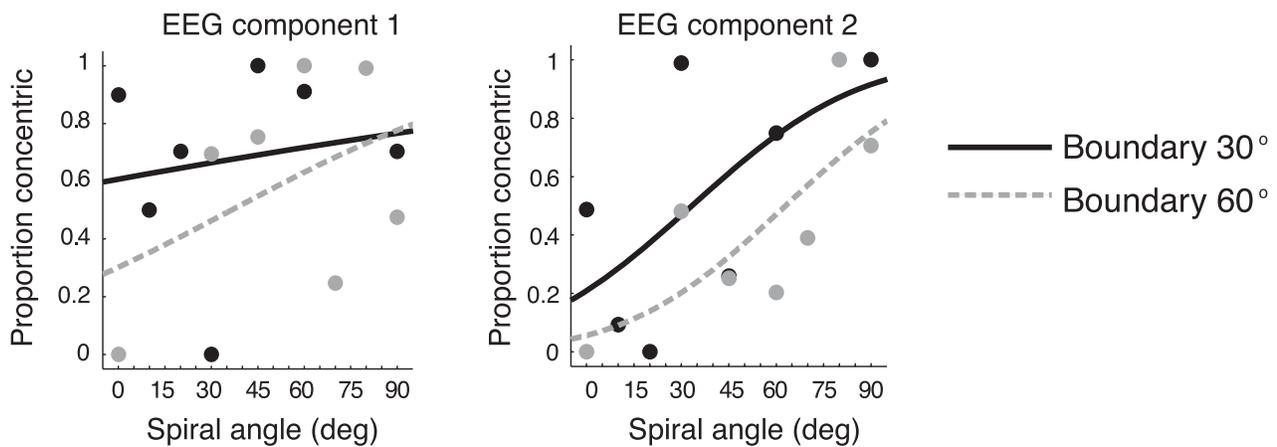


Fig. 6. Learning-dependent changes in spatiotemporal brain patterns. EEG-metric curves before and after training when (A) discriminating shapes in noise and (B) discriminating and categorizing similar shapes. For each of the two EEG components, the classifier performance at each condition was averaged across observers and fitted with cumulative Gaussian functions for each session: pre- (gray dotted lines) and post- (black solid lines) training.

highlighted the potential to use parametric stimulus manipulations coupled with fMRI design approaches that seek to increase the sensitivity of data acquisition (e.g., high-resolution acquisition; adaptation designs) and analysis (e.g., machine-learning approaches and information theory analysis to quantify the information contained in rich data sets). These data are best understood as providing analogies between perceptual states and physiological signals within Teller’s terminology. The strength of these analogies become more compelling as different data converge to the same conclusions to promote links between activity within particular circuits and what we see.

References

ADINI, Y., SAGI, D. & TSODYKYS, M. (2002). Context-enabled learning in the human visual system. *Nature* **415**, 790–793.
 BAN, H., PRESTON, T.J., MEESON, A. & WELCHMAN, A.E. (2012). The integration of motion and disparity cues to depth in dorsal visual cortex. *Nature Neuroscience* **15**, 636–643.

BECKETT, A., PEIRCE, J.W., SANCHEZ-PANCHUELO, R.M., FRANCIS, S. & SCHLUPPECK, D. (2012). Contribution of large scale biases in decoding of direction-of-motion from high-resolution fMRI data in human early visual cortex. *Neuroimage* **63**, 1623–1632.
 BOYNTON, G.M. (2011). Spikes, BOLD, attention, and awareness: A comparison of electrophysiological and fMRI signals in V1. *Journal of Vision* **11**, 12.
 BOYNTON, G.M., DEMB, J.B., GLOVER, G.H. & HEEGER, D.J. (1999). Neuronal basis of contrast discrimination. *Vision Research* **39**, 257–269.
 BOYNTON, G.M., ENGEL, S.A., GLOVER, G.H. & HEEGER, D.J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *The Journal of Neuroscience* **16**, 4207–4221.
 BRADLEY, D.C., CHANG, G.C. & ANDERSEN, R.A. (1998). Encoding of three-dimensional structure-from-motion by primate area MT neurons. *Nature* **392**, 714–717.
 BRIDGE, H. & PARKER, A.J. (2007). Topographical representation of binocular depth in the human visual cortex using fMRI. *Journal of Vision* **7**, 1–14.
 BRITTEN, K.H., NEWSOME, W.T., SHADLEN, M.N., CELEBRINI, S. & MOVSHON, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience* **13**, 87–100.
 BRITTEN, K.H., SHADLEN, M.N., NEWSOME, W.T. & MOVSHON, J.A. (1992). The analysis of visual motion: A comparison of neuronal

- and psychophysical performance. *The Journal of Neuroscience* **12**, 4745–4765.
- CARDOSO, M.M., SIROTTIN, Y.B., LIMA, B., GLUSHENKOVA, E. & DAS, A. (2012). The neuroimaging signal is a linear sum of neurally distinct stimulus- and task-related components. *Nature Neuroscience* **15**, 1298–1306.
- CHANDRASEKARAN, C., CANON, V., DAHMEN, J.C., KOURTZI, Z. & WELCHMAN, A.E. (2007). Neural correlates of disparity-defined shape discrimination in the human brain. *Journal of Neurophysiology* **97**, 1553–1565.
- COGAN, A.I., KONTSEVICH, L.L., LOMAKIN, A.J., HALPERN, D.L. & BLAKE, R. (1995). Binocular disparity processing with opposite-contrast stimuli. *Perception* **24**, 33–47.
- COTTEREAU, B.R., MCKEE, S.P., ALES, J.M. & NORCIA, A.M. (2011). Disparity-tuned population responses from human visual cortex. *The Journal of Neuroscience* **31**, 954–965.
- COTTEREAU, B.R., MCKEE, S.P., ALES, J.M. & NORCIA, A.M. (2012). Disparity-specific spatial interactions: Evidence from EEG source imaging. *The Journal of Neuroscience* **32**, 826–840.
- COX, D.D. & SAVOY, R.L. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* **19**, 261–270.
- CUMMING, B.G. & DEANGELIS, G.C. (2001). The physiology of stereopsis. *Annual Review of Neuroscience* **24**, 203–238.
- CUMMING, B.G. & PARKER, A.J. (1997). Responses of primary visual cortical neurons to binocular disparity without depth perception. *Nature* **389**, 280–283.
- CUMMING, B.G., SHAPIRO, S.E. & PARKER, A.J. (1998). Disparity detection in anticorrelated stereograms. *Perception* **27**, 1367–1377.
- DAS, K., GIESBRECHT, B. & ECKSTEIN, M.P. (2010). Predicting variations of perceptual performance across individuals from neural activity using pattern classifiers. *Neuroimage* **51**, 1425–1437.
- DEANGELIS, G.C. & NEWSOME, W.T. (1999). Organization of disparity-selective neurons in macaque area MT. *The Journal of Neuroscience* **19**, 1398–1415.
- DEBENER, S., ULLSPERGER, M., SIEGEL, M., FIEHLER, K., VON CRAMON, D.Y. & ENGEL, A.K. (2005). Trial-by-trial coupling of concurrent electroencephalogram and functional magnetic resonance imaging identifies the dynamics of performance monitoring. *The Journal of Neuroscience* **25**, 11730–11737.
- DOSHER, B.A. & LU, Z.L. (1999). Mechanisms of perceptual learning. *Vision Research* **39**, 3197–3221.
- EICHELE, T., SPECHT, K., MOOSMANN, M., JONGSMA, M.L., QUIROGA, R.Q., NORDBY, H. & HUGDAHL, K. (2005). Assessing the spatiotemporal evolution of neuronal activation with single-trial event-related potentials and functional MRI. *Proceedings of the National Academy of Sciences of the United States of America* **102**, 17798–17803.
- FINE, I. & JACOBS, R.A. (2002). Comparing perceptual learning tasks: A review. *Journal of Vision* **2**, 190–203.
- FREEMAN, J., BROUWER, G.J., HEEGER, D.J. & MERRIAM, E.P. (2011). Orientation decoding depends on maps, not columns. *The Journal of Neuroscience* **31**, 4792–4804.
- GEORGIEVA, S.S., PEETERS, R., KOLSTER, H., TODD, J.T. & ORBAN, G.A. (2009). The processing of 3D shape from disparity in the human brain. *The Journal of Neuroscience* **29**, 727–742.
- GEORGIEVA, S.S., TODD, J.T., PEETERS, R. & ORBAN, G.A. (2008). The extraction of 3D shape from texture and shading in the human brain. *Cerebral Cortex* **18**, 2416–2438.
- GILBERT, C.D., SIGMAN, M. & CRIST, R.E. (2001). The neural basis of perceptual learning. *Neuron* **31**, 681–697.
- GOLDSTONE, R.L. (1998). Perceptual learning. *Annual Review of Psychology* **49**, 585–612.
- GOLDSTONE, R.L., LIPPA, Y. & SHIFFRIN, R.M. (2001). Altering object representations through category learning. *Cognition* **78**, 27–43.
- GRILL-SPECTOR, K., HENSON, R. & MARTIN, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences* **10**, 14–23.
- GRILL-SPECTOR, K. & MALACH, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica* **107**, 293–321.
- HAYNES, J.D. & REES, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience* **8**, 686–691.
- HAYNES, J.D. & REES, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews. Neuroscience* **7**, 523–534.
- HEEGER, D.J. & RESS, D. (2002). What does fMRI tell us about neuronal activity? *Nature Reviews. Neuroscience* **3**, 142–151.
- JACOBS, R.A. (2009). Adaptive precision pooling of model neuron activities predicts the efficiency of human visual learning. *Journal of Vision* **9**, 22, 1–15.
- JANSSEN, P., VOGELS, R., LIU, Y. & ORBAN, G.A. (2003). At least at the level of inferior temporal cortex, the stereo correspondence problem is solved. *Neuron* **37**, 693–701.
- JEHEE, J.F.M., LING, S., SWISHER, J.D., VAN BERGEN, R.S. & TONG, F. (2012). Perceptual learning selectively refines orientation representations in early visual cortex. *The Journal of Neuroscience* **32**, 16747–16753.
- JOHNSON, J.S. & OLSHAUSEN, B.A. (2003). Timecourse of neural signatures of object recognition. *Journal of Vision* **3**, 499–512.
- KAMITANI, Y. & TONG, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience* **8**, 679–685.
- KANWISHER, N., CHUN, M.M., MCDERMOTT, J. & LEDDEN, P.J. (1996). Functional imagining of human visual recognition. *Brain Research. Cognitive Brain Research* **5**, 55–67.
- KOURTZI, Z., BETTS, L.R., SARKHEIL, P. & WELCHMAN, A.E. (2005). Distributed neural plasticity for shape learning in the human visual cortex. *PLoS Biology* **3**, e204.
- KOURTZI, Z. & KANWISHER, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science* **293**, 1506–1509.
- KREKELBERG, B., BOYNTON, G.M. & VAN WEZEL, R.J.A. (2006). Adaptation: From single cells to BOLD signals. *Trends in Neurosciences* **29**, 250–256.
- KRIEGESKORTE, N., CUSACK, R. & BANDETTINI, P. (2010). How does an fMRI voxel sample the neuronal activity pattern: Compact-kernel or complex spatiotemporal filter? *Neuroimage* **49**, 1965–1976.
- LAW, C.T. & GOLD, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience* **11**, 505–513.
- LI, S., MAYHEW, S.D. & KOURTZI, Z. (2009). Learning shapes the representation of behavioral choice in the human brain. *Neuron* **62**, 441–452.
- LI, S., MAYHEW, S.D. & KOURTZI, Z. (2012). Learning shapes spatiotemporal brain patterns for flexible categorical decisions. *Cerebral Cortex* **22**, 2322–2335.
- LI, W., PIECH, V. & GILBERT, C.D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience* **7**, 651–657.
- LOGOTHETIS, N.K., PAULS, J., AUGATH, M., TRINATH, T. & OELTERMANN, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature* **412**, 150–170.
- MAYHEW, S.D., LI, S. & KOURTZI, Z. (2012). Learning acts on distinct processes for visual form perception in the human brain. *The Journal of Neuroscience* **32**, 775–786.
- MILLER, K.L., LUH, W.M., LIU, T.T., MARTINEZ, A., OBATA, T., WONG, E.C., FRANK, L.R. & BUXTON, R.B. (2001). Nonlinear temporal dynamics of the cerebral blood flow response. *Human Brain Mapping* **13**, 1–12.
- MUKAI, I., KIM, D., FUKUNAGA, M., JAPEE, S., MARRETT, S. & UNGERLEIDER, L.G. (2007). Activations in visual and attention-related areas predict and correlate with the degree of perceptual learning. *The Journal of Neuroscience* **27**, 11401–11411.
- MUKAMEL, R., GELBARD, H., ARIELI, A., HASSON, U., FRIED, I. & MALACH, R. (2005). Coupling between neuronal firing, field potentials, and fMRI in human auditory cortex. *Science* **309**, 951–954.
- NORMAN, K.A., POLYN, S.M., DETRE, G.J. & HAXBY, J.V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences* **10**, 424–430.
- OHLA, K., BUSCH, N.A., DAHLEM, M.A. & HERRMANN, C.S. (2005). Circles are different: The perception of glass patterns modulates early event-related potentials. *Vision Research* **45**, 2668–2676.
- OP DE BEECK, H.P., BAKER, C.I., DICARLO, J.J. & KANWISHER, N.G. (2006). Discrimination training alters object representations in human extrastriate cortex. *The Journal of Neuroscience* **26**, 13025–13036.
- ORBAN, G.A., JANSSEN, P. & VOGELS, R. (2006). Extracting 3D structure from disparity. *Trends in Neurosciences* **29**, 466–473.
- PARKER, A.J. (2007). Binocular depth perception and the cerebral cortex. *Nature Reviews. Neuroscience* **8**, 379–391.
- PEI, F., PETTET, M.W., VILDAVSKI, V.Y. & NORCIA, A.M. (2005). Event-related potentials show configural specificity of global form processing. *Neuroreport* **16**, 1427–1430.

- PHILIASTIDES, M.G. & SAJDA, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex* **16**, 509–518.
- PHILIASTIDES, M.G. & SAJDA, P. (2007). EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. *The Journal of Neuroscience* **27**, 13082–13091.
- PONCE, C.R., LOMBER, S.G. & BORN, R.T. (2008). Integrating motion and depth via parallel pathways. *Nature Neuroscience* **11**, 216–223.
- PRESTON, T.J., KOURTZI, Z. & WELCHMAN, A.E. (2009). Adaptive estimation of three-dimensional structure in the human brain. *The Journal of Neuroscience* **29**, 1688–1698.
- PRESTON, T.J., LI, S., KOURTZI, Z. & WELCHMAN, A.E. (2008). Multivoxel pattern selectivity for perceptually relevant binocular disparities in the human brain. *The Journal of Neuroscience* **28**, 11315–11327.
- RAEMAEKERS, M., LANKHEET, M.J.M., MOORMAN, S., KOURTZI, Z. & VAN WEZEL, R.J.A. (2009). Directional anisotropy of motion responses in retinotopic cortex. *Human Brain Mapping* **30**, 3970–3980.
- RESS, D. & HEEGER, D.J. (2003). Neuronal correlates of perception in early visual cortex. *Nature Neuroscience* **6**, 414–420.
- SAWAMURA, H., ORBAN, G.A. & VOGELS, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: A single-cell study of the fMRI adaptation paradigm. *Neuron* **49**, 307–318.
- SCANNELL, J.W. & YOUNG, M.P. (1999). Neuronal population activity and functional imaging. *Proceedings of the Royal Society of London. Series B, Biological Sciences* **266**, 875–881.
- SIGMAN, M., PAN, H., YANG, Y., STERN, E., SILBERSWEIG, D. & GILBERT, C.D. (2005). Top-down reorganization of activity in the visual pathway after learning a shape identification task. *Neuron* **46**, 823–835.
- SWISHER, J.D., GATENBY, J.C., GORE, J.C., WOLFE, B.A., MOON, C.H., KIM, S.G. & TONG, F. (2010). Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *The Journal of Neuroscience* **30**, 325–330.
- TANSKANEN, T., SAARINEN, J., PARKKONEN, L. & HARI, R. (2008). From local to global: Cortical dynamics of contour integration. *Journal of Vision* **8**, 11–12.
- TEICH, A.F. & QIAN, N. (2003). Learning and adaptation in a recurrent model of V1 orientation selectivity. *Journal of Neurophysiology* **89**, 2086–2100.
- TELLER, D.Y. (1984). Linking propositions. *Vision Research* **24**, 1233–1246.
- TELLER, D.Y. & PUGH, E.N. Jr. (1983). Linking propositions in color vision. In *Colour Vision: Physiology and Psychophysics*, ed. MOLLON, J.D. & SHARPE, L.T., London, UK: Academic Press.
- UKA, T. & DEANGELIS, G.C. (2003). Contribution of middle temporal area to coarse depth discrimination: Comparison of neuronal and psychophysical sensitivity. *The Journal of Neuroscience* **23**, 3515–3530.
- VISWANATHAN, A. & FREEMAN, R.D. (2007). Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience* **10**, 1308–1312.
- VUL, E. & PASHLER, H. (2012). Voodoo and circularity errors. *Neuroimage* **62**, 945–948.
- WELCHMAN, A.E., DEUBELIUS, A., CONRAD, V., BÜLTHOFF, H.H. & KOURTZI, Z. (2005). 3D shape perception from combined depth cues in human visual cortex. *Nature Neuroscience* **8**, 820–827.
- YOTSUMOTO, Y., WATANABE, T. & SASAKI, Y. (2008). Different dynamics of performance and brain activation in the time course of perceptual learning. *Neuron* **57**, 827–833.
- ZATORRE, R., FIELDS, R. & JOHANSEN-BERG, H. (2012). Plasticity in gray and white: Neuroimaging changes in brain structure during learning. *Nature Neuroscience* **15**, 528–536.
- ZHANG, J., MEESON, A., WELCHMAN, A.E. & KOURTZI, Z. (2010). Learning alters the tuning of fMRI multi-voxel patterns for visual forms. *The Journal of Neuroscience* **30**, 14127–14133.