

Review

Seeing patterns through the hemodynamic veil – The future of pattern-information fMRI

Elia Formisano ^{a,b,*}, Nikolaus Kriegeskorte ^{c,*}

^a Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University, The Netherlands

^b Maastricht Brain Imaging Center, Maastricht University, The Netherlands

^c Medical Research Council, Cognition and Brain Sciences Unit, Cambridge, UK

ARTICLE INFO

Article history:

Accepted 27 February 2012

Available online 6 March 2012

ABSTRACT

Pattern-information fMRI (pi-fMRI) has become a popular method in neuroscience. The technique is motivated by the idea that spatial patterns of fMRI activity reflect the neuronal population codes of perception, cognition, and action. In this commentary, we discuss three fundamental outstanding questions: (1) What is the relationship between neuronal patterns and fMRI patterns? (2) Does pattern-information fMRI benefit from hyperacuity, enabling the investigation of columnar-level neuronal information, even at low resolution? (3) Do high-resolution and high-field fMRI increase sensitivity to pattern information? The empirical answers will enable us to optimize pi-fMRI data acquisition and to understand the ultimate potential and appropriate interpretation of pi-fMRI results. Furthermore, considering the relationship between neuronal activity and fMRI at the level of spatiotemporal patterns provides a novel and important perspective on the basis of the fMRI signal.

© 2012 Elsevier Inc. All rights reserved.

Contents

Introduction	1249
From classification to regression, and on to computational modeling: how can we best test theories about brain information processing?	1250
Pattern classification	1250
Pattern regression	1250
Computational models in pattern analysis	1251
Seeing through the hemodynamic veil: does pi-fMRI have hyperacuity?	1251
If voxels sample neuronal patterns like compact kernels, fMRI hyperacuity is unlikely	1251
A large-scale organization of selectivity might underlie orientation decodability	1252
Actually, voxels sample neuronal patterns with complex spatiotemporal filters – which might afford fMRI hyperacuity.	1252
Optimal pi-fMRI acquisition: do high-resolution and high-field fMRI give more pattern information?	1253
From spatial SNR to pattern information	1253
Spatial SNR	1253
Temporal SNR	1253
FCNR	1253
Pattern information	1253
The promise of high-field fMRI	1254
Conclusion	1255
Acknowledgments	1255
References	1255

Introduction

In an increasing number of fMRI studies, the analysis of informative spatial response patterns is complementing – or even replacing –

* Corresponding authors.

E-mail addresses: e.formisano@maastrichtuniversity.nl (E. Formisano), nikolaus.kriegeskorte@mrc-cbu.cam.ac.uk (N. Kriegeskorte).

conventional analysis based on voxel-by-voxel statistical tests of activation levels. Pattern-information fMRI (pi-fMRI) is motivated by the idea that perceptual, cognitive, and action information is represented in neuronal population codes (Haxby et al., 2001; Kriegeskorte and Kreiman, 2011). While regional-average activation can tell us what regions are “involved” in a task, pi-fMRI promises to reveal the representational content (Mur et al., 2009; Norman et al., 2006). Compared to the univariate approach, pi-fMRI can more sensitively detect information about the stimuli (or tasks) in fMRI patterns because: (1) it considers not merely the regional mean activation, but also fine-grained patterns and (2) it statistically combines weak, spatially distributed single-voxel effects across all voxels within a region of interest or even across the whole brain (De Martino et al., 2008; Mourao-Miranda et al., 2005). In conventional activation analysis, fine-grained pattern information can go undetected because it is lost in regional averaging and/or spatial smoothing of neighboring voxels with opposite effects (Kriegeskorte et al., 2006).

Pi-fMRI has been very successful, bringing a range of insights into brain function that would not have been possible with the classical methods. To illustrate, let us recall a few examples: It has shown that visual object-category information is widely distributed throughout the human visual ventral stream and not confined to category selective regions (Haxby et al., 2001; see also Carlson et al., 2003; Cox and Savoy, 2003) and, similarly, that auditory object-category information is distributed throughout the auditory cortex (Staeren et al., 2009). It has shown that a grating orientation that is attended or held in short-term memory is represented in V1 (Harrison and Tong, 2009; Kamitani and Tong, 2005; see also Haynes and Rees, 2005). It has shown that individual faces are represented in anterior temporal cortex (Kriegeskorte et al., 2007; see also Nestor et al., 2011). It has shown that the inferior temporal representational space of objects is remarkably similar between man and monkey in both its categorical divisions and its within-category similarity structure (Kriegeskorte et al., 2008b). It has shown the presence of activity patterns in auditory cortex that are informative of speaker identity with invariance to which of several vowels was voiced and – vice-versa – activity patterns informative of vowels with invariance to speaker identity (Formisano et al., 2008a). It has shown that activation patterns in early auditory areas reflect the subjective perception of a sound, rather than its acoustic content (Kilian-Hutten et al., 2011). This is just a small (and unavoidably biased) selection of examples from vision and audition. But the contributions of pi-fMRI are far more numerous and span multiple fields, e.g. memory (e.g. Polyn et al., 2005; Xue et al., 2010), learning (e.g. Li et al., 2009), emotion (e.g. Ethofer et al., 2009) and many other brain functions.

This commentary is divided into three sections. The first section gives a brief overview of available pi-fMRI analysis methods. These methods provide powerful means for detecting stimulus information in fMRI patterns. The presence of fMRI pattern information provides strong evidence for neuronal effects at some spatial scale. However, it is not well understood how exactly neuronal pattern differences are reflected in fMRI pattern differences. At a small spatial scale, fMRI patterns might provide only scrambled images of neuronal patterns, but could still reflect *changes* of the neuronal pattern from one experimental condition to another (Kriegeskorte and Bandettini, 2007). Moreover, it has been suggested that pi-fMRI might have the power to reveal sub-voxel-scale columnar pattern information through small biases in the sample each voxel takes of the underlying neuronal pattern (Kamitani and Tong, 2005). The second section examines the current evidence for and against such “fMRI hyperacuity”, concluding that this intriguing idea has not been either established or disproven (Kriegeskorte et al., 2010). In the third section, we consider optimal acquisition schemes for pi-fMRI. fMRI at high-resolution and high field (7 T and more) promises greater sensitivity to fine-grained patterns (Kriegeskorte and Bandettini, 2007). However, this sensitivity will depend on many factors, including the spatial-frequency band

of the neuronal pattern effects, the physiology and spatiotemporal dynamics of the hemodynamic response, the physics of magnetic resonance imaging, and the influence of noise and artifacts. We argue that empirical studies are needed to enable us to select appropriate acquisition schemes for pi-fMRI.

From classification to regression, and on to computational modeling: how can we best test theories about brain information processing?

In the early days of neuroimaging, a wide-spread analysis approach was subtraction of brain images acquired during two conditions to be contrasted. This approach was soon generalized to univariate multiple regression analysis (Friston et al., 1995), which can handle more sophisticated parametric designs.¹ Pattern-information analysis, similarly, started with analyses aiming to detect differences (of patterns rather than bulk activation) between two predefined classes of stimuli using pattern classifiers. Like the univariate activation approach, pattern-information methods soon branched out into parametric analyses, where brain responses are related to richer continuous descriptions of the stimulus space. These parametric pattern-information approaches also enable the testing of complex predictions of computational models of brain-information processing, where the models take the experimental stimuli as input and actually perform perceptual and cognitive feats such as object recognition.

Pattern classification

Pattern classifiers remain an important tool, and the most popular method of pi-fMRI. In this approach, a classifier learns to discriminate observed brain response patterns measured during different experimental conditions, associated with different perceptual, cognitive, or action-related mental states. The generalization performance of the classifier is then tested by examining the accuracy of predicting the experimental conditions from brain response patterns in a new data set (“decoding”). Significantly accurate decoding indicates that the response patterns convey information about the experimental condition.

Note that pattern-classifier decoding differs from classical activation analysis in three respects (Fig. 1): (a) it considers multivariate patterns, not univariate regional-average activation levels, (b) the model operates in reverse, i.e. from responses to stimuli (thus called “decoding”), and (c) independent test data are used for validation and statistical inference. The analysis of multivariate patterns (rather than univariate activation), point (a), represents the crucial innovation of the pattern-information approach. Points (b) and (c) characterize pattern-classifier decoding, but are not necessary features of pattern-information analysis in general. In contrast to the decoding approach (b), the methods for testing computational models mentioned above use *encoding* models that mimic brain information processing, proceeding from stimuli to response patterns. The use of independent test data (c) is also not a necessary feature of pattern-information approaches (see e.g. Friston et al., 2008; Kriegeskorte et al., 2006).

Pattern regression

Although multivariate classification has so far received more attention in functional neuroimaging research, multivariate regression – especially in its kernel formulations (e.g. kernel ridge regression, relevance vector machine, Gaussian processes, Bayesian linear

¹ The term parametric, in this context, refers to continuous rather than categorical variation of the independent variables, not to the use of distributional assumptions for statistical inference.

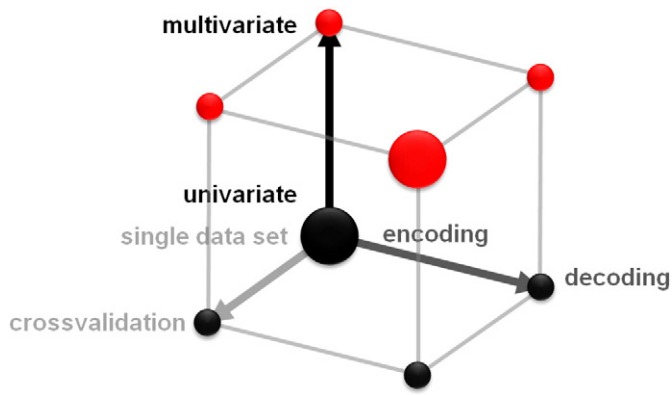


Fig. 1. Pattern-information analysis can use encoding or decoding models and perform inference using a single data set or cross validation. Classical fMRI activation analysis (large black dot) is univariate, operates from stimulus to response (encoding), and performs inference using a single data set. Pattern-classifier analysis (large red dot), by contrast, is multivariate, operates from response to stimulus (decoding), and uses an independent test set for inference (crossvalidation). Although these two combinations of the three properties are more frequent (large dots) than the other combinations, all are viable in principle. Bulk activation analysis (black dots) is thought to reflect a brain region's "involvement" in some task. Pattern-information analysis (red dots) is thought to reflect the representational content of neuronal population codes and need not rely on decoding or cross validation. Encoding models are better suited for modeling the information processing from stimulus to brain-activity pattern (perception). Decoding models are better suited for modeling the information processing from brain-activity pattern to behavioral response (readout, decision making). In bulk activation analysis, inference can usually be performed using a single data set and the assumption of univariate normality, which is reasonable in this context. In pattern-information analysis, multivariate normality does not usually hold, and inference therefore requires independent validation data or permutation testing.

regression) – is gaining popularity as a method for characterizing the relation between spatially distributed activity patterns and continuous variables (Formisano et al., 2008a, 2008b). Using multivariate linear kernel regression (MLKR), for example, it was possible to predict subject-driven behavior in a virtual reality world (Chu et al., 2011; Valente et al., 2011). Multivariate regression has also been employed to predict single-trial EEG responses from simultaneously acquired fMRI data (De Martino et al., 2011).

Similar to multivariate classification, MLKR provides a straightforward way to localize informative brain locations by ranking the model's weights associated with the voxels. This approach, however, may be problematic especially in experiments with multiple and time-continuous variables (predictors) that are not easily controllable by design. In studies of naturalistic experience (e.g. involving audiovisual movies, music, or interaction with a virtual reality) the amount of experimental control is reduced in favor of greater ecological validity. Current MLKR algorithms that are normally used in combination with large number of features deal with one predictor at a time, and thus do not explicitly account for the presence of other predictors or confounding factors which partly overlap – spatially and/or temporally – with the one considered. New fMRI applications of MLKR need to properly consider this aspect, as disregarding it may lead to the presence of 'ghost' informative patterns and complicate the interpretability of the maps. It will be also important to compare the "mapping" results obtained with these methods with those obtained using multivariate statistics with a "searchlight" approach (Kriegeskorte et al., 2006) or unsupervised algorithms, such as kernel canonical correlation analysis (Haroon et al., 2007). Finally, since kernel-based algorithms have been designed for prediction rather than for inference, it is more difficult to assess the statistical significance of the model's parameters. In fact, due to the regularization procedures employed, it is particularly problematic to use parametric models to characterize the distributions of the model parameters. Non-parametric approaches, such as permutation

testing, remain a valuable tool for statistical inference, and they are becoming increasingly feasible with rising computational power.

Computational models in pattern analysis

Relating multivariate brain-activity patterns to continuous multivariate stimulus representations is also required to test computational theories. This exciting direction incorporates computational models of brain information processing into the analysis of brain-activity data. The models, here, are no longer just statistical vehicles, but they process the experimental stimuli and are meant to mimic brain information processing. Computational modeling and brain-activity experimentation, two fields that used to be related mainly at the level of verbal theory, are thus linked in data analysis, and computational models are directly evaluated with brain data.

There are two basic approaches in the literature. One approach is voxel-receptive-field modeling (Dumoulin and Wandell, 2008; Kay et al., 2008; Mitchell et al., 2008; for a review and motivation of the approach see Gallant et al., 2011), in which the model's internal representation of the stimuli is used for linear prediction of the voxel responses. The model is then validated with independent data acquired with a different sample of stimuli. Another approach is representational similarity analysis (RSA; Kriegeskorte et al., 2008a, 2008b; Kriegeskorte, 2009). RSA relates brain and model representations at the level of the dissimilarity matrices of the representational patterns. This obviates the need for defining the mapping from model units to fMRI voxels, simplifying the analysis while addressing the same scientific questions. An independent data set with response patterns to different stimuli is only required when parameters of the computational model are fitted to the data. These two approaches have recently been compared in a review paper (Kriegeskorte, 2011).

Seeing through the hemodynamic veil: does pi-fMRI have hyperacuity?

Several studies have suggested that fMRI voxel patterns can convey information represented in columnar-scale neuronal population codes – even when the fMRI voxel resolution is insufficient to directly image the patterns of columnar selectivity. Kamitani and Tong (2005) and Haynes and Rees (2005) have shown that fMRI activity patterns acquired from V1 with 3-mm-wide isotropic voxels enable robust decoding of the orientation of a visually presented grating. Decoding worked even though the fMRI patterns were much too coarse to directly depict the fine-grained patterns of orientation columns. Does this mean that columnar-scale neuronal patterns, in general, can be investigated with standard-resolution fMRI?

Such "pi-fMRI hyperacuity" is an exciting possibility. It would enable us to reveal neuronal pattern information from spatial frequencies above the band that can be directly resolved at a given fMRI resolution. We call this spatial-frequency band of neuronal-pattern effects the "hyperband" (Kriegeskorte et al., 2010). Hyperband neuronal-pattern information in fMRI patterns constitutes a form of aliasing, by definition, because voxels reflect spatial pattern frequencies that exceed the Nyquist limit. Even subtle aliases of hyperband effects might be very useful, because pi-fMRI can combine the evidence distributed across hundreds or thousands of voxels and enable us to test neuroscientific hypotheses about columnar-scale representations. However, it is unclear whether fMRI hyperacuity is real. Several hypotheses have been proposed to explain the empirical findings.

If voxels sample neuronal patterns like compact kernels, fMRI hyperacuity is unlikely

Kamitani and Tong (2005) originally suggested that the decodability of grating orientation might result from unbalanced sampling of orientation-selective columns by the voxels. Each voxel might sample

columns selective for the different orientations in slightly different proportions (*biased sampling hypothesis*; for an intuitive illustration, see Boynton, 2005). Each voxel, then, would inherit a subtle bias for certain orientations, despite responding to all orientations. Pattern analysis would combine the evidence across voxels and enable us to robustly detect (or decode) the information. How plausible this hypothesis appears depends on our model of how an fMRI voxel samples the neuronal pattern through the hemodynamic process (Kriegeskorte et al., 2010).

The field has generally assumed that each voxel samples the neuronal activity pattern either by averaging within its boundaries, or by a local Gaussian-kernel filter that approximates the hemodynamic blur. Both of these possibilities constitute compact-kernel models of the spatial filter that characterizes a voxel. If voxels are compact-kernel filters, fMRI hyperacuity is unlikely for two reasons (Chaimow et al., 2011; Kriegeskorte et al., 2010): (1) Compact-kernel filtering minimizes aliasing. In image processing, such filters are purposely applied to prevent aliasing. For pi-fMRI, detailed simulations of the neuronal patterns and the sampling process suggest that hyperband pattern effects will not be detectable (Chaimow et al., 2011). (2) Hyperband pattern effects aliased into the fMRI patterns through any filter spatially fixed to the voxels would be highly sensitive to slight shifts of the voxel grid caused by head motion (see Fig. 3 in Kriegeskorte et al., 2010). For example, head motion on the scale of the width of an orientation column would entirely and globally change the aliased pattern and this change would not be undone by head-motion correction. When orientation columns are directly imaged with high-resolution fMRI (Yacoub et al., 2008), then head-motion correction can in principle correct the displacements. However, when orientation information only enters the fMRI patterns through aliasing with a voxel-fixed filter, head-motion correction fails to align the aliases. Although bite bars can help minimize head motion (e.g. Yacoub et al., 2008), such measures are not required for orientation decoding at standard resolution (e.g. Kamitani and Tong, 2005), which appears to be robust to small head movements.

Both of these two arguments suggest either that voxels are not compact-kernel filters or that the fMRI pattern information does not reflect hyperband neuronal-pattern information. The next two sections address these two possibilities in turn.

A large-scale organization of selectivity might underlie orientation decodability

The presence of a columnar-scale neuronal preference map does not preclude larger-scale variation of preferences. The columnar-scale preferences might be ripples on a larger wave. It has been argued that a large-scale, rather than a columnar-scale, organization of selectivity might underlie decodability of grating orientation and other stimulus properties.

Consider the case of orientation columns. We know that selectivities vary across cortex at high frequency. High-field high-resolution fMRI suggests about 2.24 pinwheels per mm² of cortex (Yacoub et al., 2008). A 3-mm-wide isotropic voxel would then sample about 20 pinwheels. However, there is evidence that a given patch of V1 will respond more strongly to a grating in its receptive field when the grating is oriented radially to the point of fixation (Sasaki et al., 2006). This suggests that the fine-scale orientation-preference map, with its many little pinwheels, rides on top of a global radial-preference map, which might be thought of as a single big pinwheel centered on fixation. A uniform grating oriented from lower left to upper right would then more strongly drive the V1 representations of the lower left and upper right quadrants of the visual field. It has recently been shown that this radial preference map contributes to the decoding of grating orientation at low resolution (Freeman et al., 2011).

Although the radial-preference map suggests a simple and appealing resolution to the puzzle of low-res orientation decodability, demonstrating the absence of any contribution from the columnar scale is difficult. Moreover, other evidence suggests that the radial-preference map is not the whole story: Mannion et al. (2009) showed that response patterns elicited by opposite logarithmic spirals can be decoded from V1 (and also from V2 and V3). The spirals are radially balanced: for each spiral, the orientation is 45° (or –45°) off the radial axis at every location of the visual field. A radial-preference map, therefore, cannot account for the decodability of these stimuli.

Another line of evidence is provided by low-pass filtering and high-pass filtering of the fMRI patterns. Op de Beeck (2010) reported that smoothing (i.e. low-pass filtering) does not hurt decoding performance. This finding suggests a stronger contribution from the low band than from the high band of the fMRI patterns. However, it does not have strong implications for the question of fMRI hyperacuity, for the following two reasons: (1) If hyperband information is aliased into the low band of the fMRI patterns, then smoothing will not interfere with it (Kriegeskorte et al., 2010). Aliasing has no preference for nearby frequency bands, so there is no reason to expect more hyperband neuronal information in the high than in the low band of the fMRI patterns. (2) Gaussian smoothing is an invertible linear operation, so it does not remove any information – unless noise is added after smoothing or numerical precision is insufficient (Kamitani and Sawahata, 2010). Even if smoothing improved decoding performance, this would not speak to the question of fMRI hyperacuity. It would not even mean that the high band (the higher portion of the directly imaged spatial-frequency range) contains no information: If the high band contained information at a lower signal-to-noise ratio (e.g. as a consequence of larger residual head motion artifacts in this band), then downscaling the high band could improve the signal-to-noise ratio despite true weak effects in the high band.

In order to address what information is present in the high band of the fMRI patterns, a more promising approach is to high-pass filter the fMRI patterns. Using this approach, Swisher et al. (2010) showed that both high and low bands of high-resolution fMRI patterns contain orientation information. For ocular dominance columns, Shmuel et al. (2010) similarly suggested that both coarse-scale and fine-scale structures reflect the stimulated eye in 7 T high-resolution pi-fMRI.

In sum, it is unlikely that neuronal pattern effects have exactly zero energy in the lower spatial frequency bands. The presence of coarse-scale pattern information has been demonstrated for visual grating orientation (Freeman et al., 2011; Sasaki et al., 2006; Swisher et al., 2010). Coarse-scale pattern information contributes to pi-fMRI. Within the frequency range directly resolved at a given fMRI resolution, both the low and the high band of the fMRI patterns contain information about the stimuli (Shmuel et al., 2010; Swisher et al., 2010). Intriguingly, visual orientations are still decodable for radially balanced stimuli with no expected radial-preference map effects (Mannion et al., 2009), suggesting that pi-fMRI might reflect hyperband pattern effects after all.

Actually, voxels sample neuronal patterns with complex spatiotemporal filters – which might afford fMRI hyperacuity

Above we considered the widely used compact-kernel model of how a voxel samples the neuronal activity pattern. We concluded that this model renders pi-fMRI hyperacuity unlikely. This left us with two possibilities: either decoding results are entirely due to neuronal pattern effects in spatial frequency bands below the Nyquist limit imposed by the size of the voxels (large-scale organization) or the compact-kernel model is incorrect.

Actually, we know that the compact-kernel model is incorrect. The processes by which neuronal activity triggers a vascular response and the spatiotemporal hemodynamics of the BOLD effect are much more

complex than a spatial compact-kernel multiplied by a temporal response function (e.g. Boynton et al., 1996). The filter by which a voxel samples the neuronal activity pattern depends on the unique vascular architecture supplying each particular voxel with its signal (Turner, 2002) and on the dynamics of blood flow (e.g. the speed of the blood in vessels of different diameters). These considerations suggest that the filter is not space–time separable (i.e. the spatial profile changes across time) and might have unexpected sensitivity (Kriegeskorte et al., 2010). If the filters have a fine-grained spatial structure (if only for some of the voxels at some of the latencies) they might conceivably lend pi-fMRI some measure of sensitivity to hyperband spatial frequencies. The complex-spatiotemporal-filter hypothesis is also compatible with the observed robustness of decoding to slight shifts of the voxel grid caused by head motion: The aliasing that projects hyperband neuronal pattern effects into the fMRI patterns would occur at the level of the vasculature, which moves with the head, remaining in a constant relationship to the neuronal patterns.

There is evidence for functional biases in large vessels (Gardner, 2010; Shmuel et al., 2010), which could be thought of as transposing high spatial-frequency neuronal pattern effects into lower spatial-frequency venous signals. Moreover, simple simulations suggest, perhaps counterintuitively, that the spatial characteristics of the filter depend on temporal parameters of the acquisition, e.g. the TR, and of the hemodynamics, e.g. the speed of blood in the capillary bed (Kriegeskorte et al., 2010). Importantly, the fact that voxels are complex spatiotemporal filters does not imply that pi-fMRI has hyperacuity. It only reminds us that there might be more fine-grained neuronal pattern information in a voxel's response than in a local average of neuronal activity.

In sum, the jury on the nature of the information in fMRI patterns is still out. To reach a verdict, we will need to understand the relationship between hemodynamic and neuronal signals at a fine-grained spatial and temporal scale. This relationship has so far been targeted mainly at the level of univariate fMRI responses, addressing the question how different aspects of neuronal activity (e.g. spiking and local field potentials) contribute to the fMRI signal (Logothetis et al., 2001). An important future direction is to characterize how spatiotemporal neuronal patterns are reflected in spatiotemporal fMRI patterns. Pattern information provides a novel, yet fundamental, perspective on the basis of the fMRI signal.

Optimal pi-fMRI acquisition: do high-resolution and high-field fMRI give more pattern information?

From spatial SNR to pattern information

The quality of an fMRI acquisition scheme can be assessed by a range of criteria at different levels of analysis (Fig. 2). MR physicists usually assess the spatial and temporal signal-to-noise ratio (SNR). However, these two do not unequivocally inform us how sensitive the acquisition will be to either regional-average activation or pattern information. It is therefore necessary to empirically compare different acquisition schemes also at the higher levels of analysis of the functional-contrast-to-noise ratio (FCNR) and pattern information (pi). This, of course, requires an experimental design with different conditions to be contrasted, and results will depend on the conditions. We will now consider each criterion in turn, and discuss the factors that come into play at each level of analysis.

Spatial SNR

The spatial SNR is the amplitude of the MRI signal relative to its variability due to instrumental noise. The spatial SNR can be assessed from a single volume, by estimating the signal (e.g. the regional-average signal amplitude in the brain) and dividing by an estimate of the noise (e.g. the standard deviation of the signal outside the head, where all changes are due to instrumental noise). The signal

amplitude, and with it the spatial SNR, is proportional to voxel volume (Edelstein et al., 1986). Lower resolutions, thus, are preferable by this criterion. The spatial SNR also depends on the strength of the magnetic field. Higher fields provide greater spatial SNR. Considering only spatial SNR would suggest measuring at the lowest possible resolution with the highest possible field strength.

Temporal SNR

The temporal SNR is the signal amplitude relative to its variability across time. The temporal SNR can be assessed from a time series, by estimating the temporal-mean signal amplitude and dividing by the standard deviation across time, which reflects both instrumental noise and physiological fluctuations. As voxels get bigger, the instrumental noise fades relative to the signal (i.e. the spatial SNR rises), but physiological fluctuations remain constant relative to the signal (Krüger and Glover, 2001). Making voxels bigger therefore yields diminishing returns in temporal SNR when the instrumental noise is already negligible relative to the physiological fluctuations. Considering only temporal SNR therefore suggests choosing a resolution that balances instrumental noise and physiological fluctuations (Bodurka et al., 2007), such that the resolution is as high as possible without the fMRI signal being dominated by instrumental noise (but see also Triantafyllou et al., 2005).

Note that the physiological fluctuations include BOLD effects related to neuronal brain activity, along with effects related to heartbeat, respiration, and subtle brain pulsations. Brain-activity effects of interest are lumped with the “noise” in the term temporal SNR. The concept of temporal SNR, thus, appears oblivious to the purpose of fMRI: to measure brain activity, which varies with time. This hints at the limits of this criterion for choosing an optimal acquisition scheme. However, temporal SNR is still a useful criterion. It can be assessed based on resting-state data and does not require or depend on an experimental design contrasting different conditions of brain activity. As long as the time to echo (TE) is optimal for BOLD, the ratio of brain-activity effects of interest and physiological fluctuations of no interest should be constant at a given resolution. As a result, an acquisition scheme with better temporal SNR is also expected to have better FCNR. As instrumental noise becomes negligible relative to physiological signals, both temporal SNR and FCNR will increase to an asymptotic level (Krüger and Glover, 2001), although this may not strictly hold in the case of multi-channel receive coils (Triantafyllou et al., 2011).

FCNR

When we are comparing fMRI acquisition schemes with different resolutions, then the local averaging within each voxel can differentially affect the brain-activity effects of interest and physiological fluctuations of no interest. Consider the scenario of a small activation blob hidden in physiological noise of a larger spatial grain. The lower-resolution sequence might have better temporal SNR, but it might be less sensitive to an activation blob that is smaller than a voxel.

Changing the resolution, the field strength, or the sequence type (e.g. gradient echo versus spin echo) can affect the relative sensitivity to brain-activity effects of interest and physiological fluctuations of no interest. Whenever such changes to the acquisition scheme are to be evaluated, temporal SNR will no longer give a good indication of FCNR. We might have greater brain-activity effects of interest in one sequence, thus greater FCNR, but also greater overall physiological fluctuations (lumping functional contrast and noise), thus lower temporal SNR. This is why we need to determine FCNR to compare acquisition schemes for a given contrast of experimental conditions.

Pattern information

At the highest level, pattern-information measures (e.g. decoding accuracy) give us an indication about the extent to which the regional

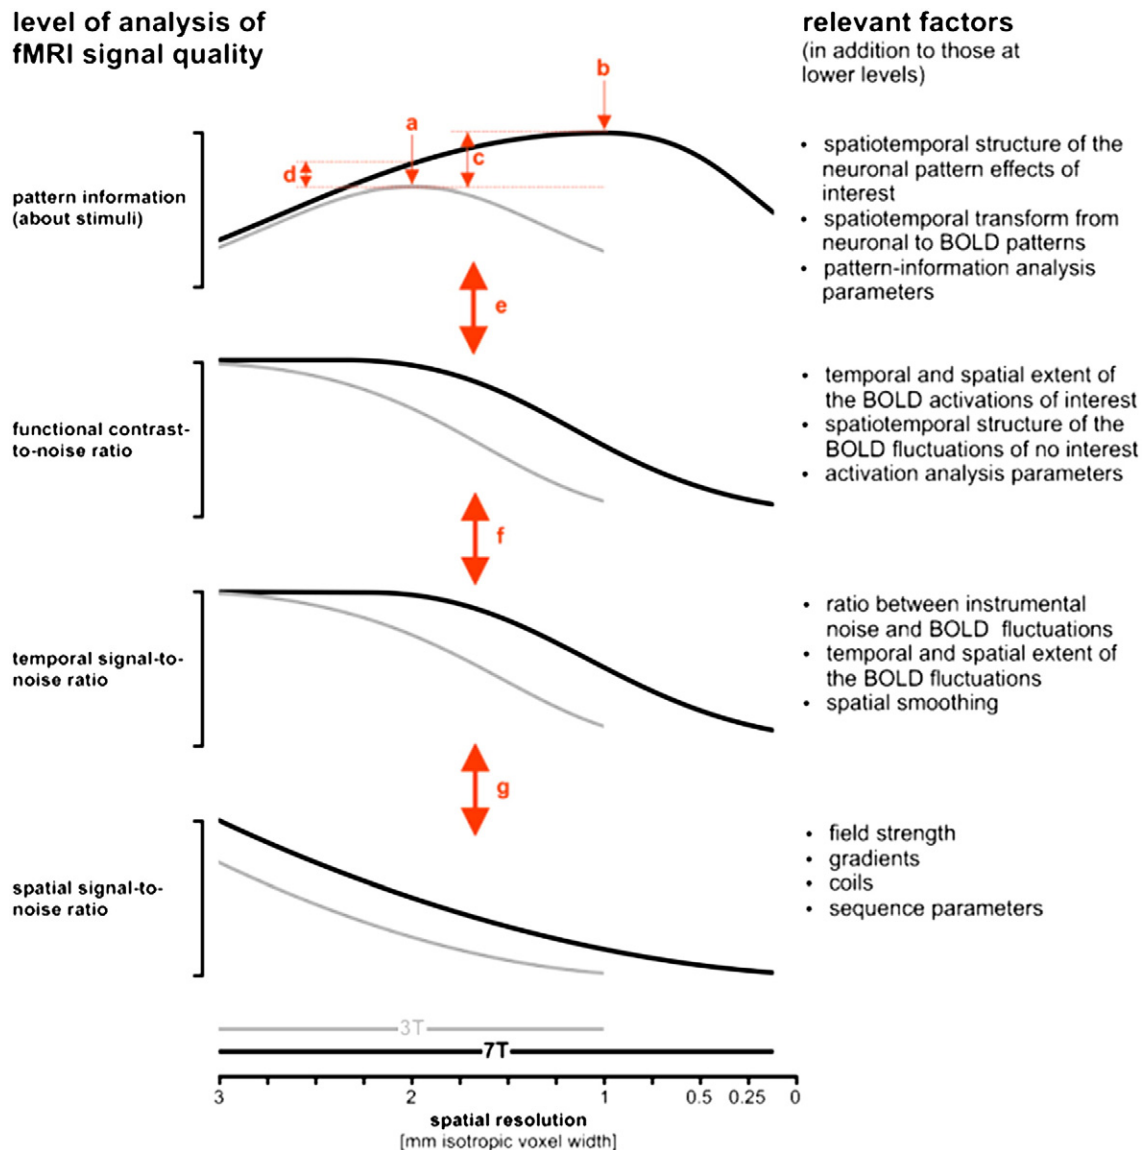


Fig. 2. Criteria of fMRI signal quality at multiple levels of analysis. This figure illustrates different criteria of the quality of 3 T (gray) and 7 T (black) fMRI acquisition schemes, as a function of spatial resolution. At the lowest level, spatial SNR rises as voxels get bigger – without saturating. At higher levels, the temporal SNR (TSNR) and the functional contrast-to-noise ratio (FCNR) also rise as voxels get bigger – but they saturate when the physiological fluctuations (for TSNR) or the physiological noise (for FCNR) dominate the fMRI signal. At the highest level, pattern-information measures (e.g. decoding accuracy) give us an indication about the extent to which the regional fMRI signal reveals the content of the neuronal population code. Pattern information relies on both good FCNR and sufficient resolution to reflect columnar neuronal pattern contrasts. It is therefore expected to peak at the optimal combination of FCNR and spatial resolution. However, several important questions remain unanswered (red letters): (a) What spatial resolution maximizes pattern information at 3 T? (b) What spatial resolution maximizes pattern information at 7 T? (c) Is 7 T better than 3 T for revealing pattern information? (d) Does the optimal field strength for pattern information studies depend on the resolution (e.g. does 3 T yield greater pattern information than 7 T at lower resolutions)? Note that the curves are not based on data but represent conventional guesses that are likely to be incorrect. Our aim here is to highlight that more research is needed to address these questions.

fMRI signal reveals the content of the neuronal population code. From a neuroscientific perspective, this might be considered the fundamental or ultimate criterion. Pattern information relies on both good FCNR and sufficient resolution to reflect neuronal pattern contrasts, which might have a lot of their energy at the columnar scale. Note that pattern-information techniques separate effects from noise in multivariate space. This is why FCNR, which considers univariate contrast only, cannot fully predict pattern information.

The promise of high-field fMRI

The recent advent of fMRI at high magnetic fields (7 T or above), opens unprecedented opportunities for unraveling neuronal coding mechanisms in humans. Compared to fMRI at standard fields (1.5 and 3 T), high-field fMRI presents several advantages. First, at high

fields spatial SNR (Vaughan et al., 2001) and BOLD FCNR (Yacoub et al., 2001, 2003) increase, which allows for proportional increases in the image resolution (i.e. smaller voxels). Crucially, higher image resolutions also result in reduced partial-volume effects, enabling us to segregate signals originating from the microvasculature from the surrounding tissue (e.g. large surface vessels, white matter), thus permitting significant improvements of fMRI spatial specificity (Polimeni et al., 2010; Yacoub et al., 2007, 2008). Further, BOLD FCNR and specificity can vary depending on the contrast (T_2^* or T_2) and the field strength used for fMRI measurements (Ugurbil et al., 2003). In particular, spin-echo (i.e. T_2 -weighted) images at high magnetic fields have been shown to be less sensitive to signals originating from large draining veins than gradient-echo (T_2^* -weighted) images, thereby increasing fMRI spatial specificity (Uludağ et al., 2009; Yacoub et al., 2005).

All the elements discussed above (increase of spatial SNR and BOLD contrast, reduced partial-volume sampling, smaller voxel size, increased sensitivity to microvasculature) indicate that high-field fMRI will enable sampling of the neuronal activity patterns with sub-millimeter specificity. It is thus expected that neuronal spatial organizations which are invisible at 3 T will become visible at higher magnetic fields. In fact, it has already been shown that cortical orientation columns can be visualized in the human brain (Yacoub et al., 2008) and that layer-specific activation and connectivity studies are possible at 7 T (Polimeni et al., 2010).

However, high-field fMRI is also associated with greater field-homogeneity challenges and higher sensitivity to head-motion artifacts. Furthermore, greater spatial resolution means more voxels, which pose a greater statistical challenge (Kriegeskorte and Bandettini, 2007). In univariate mapping, this challenge manifests in a more severe multiple-testing problem, in pattern-information analysis it manifests in a greater degree of overfitting, which can compromise generalization performance. If neuronal pattern effects are strong in low spatial frequencies or if high-band neuronal effects are projected into low spatial frequencies of the fMRI patterns through some form of aliasing (i.e. hyperacuity), then functional contrast might matter more than resolution. Fig. 2(c) therefore raises the empirical question what field strength yields the greatest pattern information. It remains an open question how the pros and cons of high-field acquisition play out in pi-fMRI.

Conclusion

Considering the relationship between neuronal activity and fMRI at the level of fine-grained spatiotemporal patterns provides a novel and important perspective on the basis of the fMRI signal. In this review, we have highlighted a range of open questions on the physiological basis of pi-fMRI and on the optimal acquisition scheme. The empirical answers will reveal the power this technique gives us in the context of current fMRI technology and will inform design, acquisition, analysis, and interpretation. But whatever the answers are, whether hyperacuity is real or pi-fMRI is limited to directly resolved pattern effects, we already know that pi-fMRI provides a unique window on neuronal representation. In fMRI and cell recording, pattern-information techniques are here to stay and will be key to understanding the combinatorial information encoded in neuronal ensembles as reflected in increasingly rich spatiotemporal measurements of brain activity.

Acknowledgments

This research was supported by Maastricht University (EF) and by the UK Medical Research Council and by a European Research Council Starting Grant (NK). We thank Federico De Martino, Kamil Uludağ, Giancarlo Valente, Rainer Goebel and Peter Bandettini for useful comments and discussions on this article.

References

- Bodurka, J., Ye, F., Petridou, N., Murphy, K., Bandettini, P.A., 2007. Mapping the MRI voxel volume in which thermal noise matches physiological noise—implications for fMRI. *Neuroimage* 34 (2), 542–549.
- Boynton, G.M., 2005. Imaging orientation selectivity: decoding conscious perception in V1. *Nat. Neurosci.* 8, 541–542.
- Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* 16, 4207–4221.
- Carlson, T.A., Schrater, P., He, S., 2003. Patterns of activity in the categorical representations of objects. *J. Cogn. Neurosci.* 15, 704–717.
- Chaimow, D., Yacoub, E., Ugurbil, K., Shmuel, A., May 15 2011. Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns. *Neuroimage* 56 (2), 627–642 Epub 2010 Sep 22.
- Chu, C., Ni, Y., Tan, G., Saunders, C.J., Ashburner, J., 2011. Kernel regression for fMRI pattern prediction. *Neuroimage* 56, 662–673.
- 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.
- De Martino, F., Valente, G., Staeren, N., Ashburner, J., Goebel, R., Formisano, E., 2008. Combining multivariate voxel selection and support vector machines for mapping and classification of fMRI spatial patterns. *Neuroimage* 43, 44–58.
- De Martino, F., de Borst, A.W., Valente, G., Goebel, R., Formisano, E., 2011. Predicting EEG single trial responses with simultaneous fMRI and relevance vector machine regression. *Neuroimage* 56, 826–836.
- Dumoulin, S.O., Wandell, B.A., 2008. Population receptive field estimates in human visual cortex. *Neuroimage* 39, 647–660.
- Edelstein, W.A., Glover, G.H., Hardy, C.J., Redington, R.W., 1986. The intrinsic signal-to-noise ratio in NMR imaging. *Magn. Reson. Med.* 3, 604–618.
- Ethofer, T., Van De Ville, D., Scherer, K., Vuilleumier, P., 2009. Decoding of emotional information in voice-sensitive cortices. *Curr. Biol.* 19, 1028–1033.
- Formisano, E., De Martino, F., Bonte, M., Goebel, R., 2008a. “Who” is saying “what”? Brain-based decoding of human voice and speech. *Science* 322, 970–973.
- Formisano, E., De Martino, F., Valente, G., 2008b. Multivariate analysis of fMRI time series: classification and regression of brain responses using machine learning. *Magn. Reson. Imaging* 26, 921–934.
- Freeman, J., Brouwer, G.J., Heeger, D.J., Merriam, E.P., 2011. Orientation decoding depends on maps, not columns. *J. Neurosci.* 31 (13), 4792–4804.
- Friston, K.J., Holmes, A.P., Poline, J.B., Grasby, P.J., Williams, S.C., Frackowiak, R.S., Turner, R., 1995. Analysis of fMRI time-series revisited. *Neuroimage* 2, 45–53.
- Friston, K., Chu, C., Mourao-Miranda, J., Hulme, O., Rees, G., Penny, W., Ashburner, J., 2008. Bayesian decoding of brain images. *Neuroimage* 39, 181–205.
- Gallant, J.L., Nishimoto, S., Naselaris, T., Wu, M.C.K., 2011. System identification, encoding models and decoding models, a powerful new approach to fMRI research. In: Kriegeskorte, N., Kreiman, G. (Eds.), *Visual Population Codes – toward a Common Multivariate Framework for Cell Recording and Functional Imaging*. MIT Press.
- Gardner, J.L., 2010. Is cortical vasculature functionally organized? *Neuroimage* 49, 1953–1956.
- Hardoon, D.R., Mourao-Miranda, J., Brammer, M., Shawe-Taylor, J., Oct 1 2007. Unsupervised analysis of fMRI data using kernel canonical correlation. *Neuroimage* 37 (4), 1250–1259.
- Harrison, S.A., Tong, F., 2009. Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458, 632–635.
- Haxby, J.V., Gobbini, M.I., Fury, M., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293, 2425–2430.
- Haynes, J.D., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* 8, 686–691.
- Kamitani, Y., Sawahata, Y., 2010. Spatial smoothing hurts localization but not information: pitfalls for brain mappers. *Neuroimage* 49 (3), 1949–1952.
- Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* 8, 679–685.
- Kay, K.N., Naselaris, T., Prenger, R.J., Gallant, J.L., 2008. Identifying natural images from human brain activity. *Nature* 452, 352–355.
- Kilian-Hutten, N., Valente, G., Vroomen, J., Formisano, E., 2011. Auditory cortex encodes the perceptual interpretation of ambiguous sound. *J. Neurosci.* 31, 1715–1720.
- Kriegeskorte, N., 2009. Relating population-code representations between man, monkey, and computational models. *Front. Neurosci.* 3 (3), 363–373.
- Kriegeskorte, N., 2011. Pattern-information analysis: from stimulus decoding to computational-model testing. *Neuroimage* 56 (2), 411–421.
- Kriegeskorte, N., Bandettini, P., 2007. Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage* 38, 649–662.
- Kriegeskorte, N., Kreiman, G., 2011. Visual population codes – toward a common multivariate framework for cell recording and functional imaging. (Editors) MIT Press.
- Kriegeskorte, N., Goebel, R., Bandettini, P., 2006. Information-based functional brain mapping. *Proc. Natl. Acad. Sci. U. S. A.* 103, 3863–3868.
- Kriegeskorte, N., Formisano, E., Singer, B., Goebel, R., 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc. Natl. Acad. Sci. U. S. A.* 104, 20600–20605.
- Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., Bandettini, P.A., 2008a. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60 (6), 1126–1141.
- Kriegeskorte, N., Mur, M., Bandettini, P.A., 2008b. Representational similarity analysis – connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* 2–4.
- 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 (3), 1965–1976.
- Krüger, G., Glover, G.H., 2001. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magn. Reson. Med.* 46 (4), 631–637.
- Li, S., Mayhew, S.D., Kourtzi, Z., 2009. Learning shapes the representation of behavioral choice in the human brain. *Neuron* 62, 441–452.
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412 (6843), 150–157.
- Mannion, D.J., McDonald, J.S., Clifford, C.W., 2009. Discrimination of the local orientation structure of spiral Glass patterns early in human visual cortex. *Neuroimage* 46 (2), 511–515.
- Mitchell, T.M., Shinkareva, S.V., Carlson, A., Chang, K.-M., Malave, V.L., Mason, R.A., Just, M.A., 2008. Predicting human brain activity associated with the meanings of nouns. *Science* 320 (5880), 1191–1195.
- Mourao-Miranda, J., Bokde, A.L., Born, C., Hampel, H., Stetter, M., 2005. Classifying brain states and determining the discriminating activation patterns: support vector machine on functional MRI data. *Neuroimage* 28, 980–995.

- Mur, M., Bandettini, P., Kriegeskorte, N., 2009. Revealing representational content with pattern-information fMRI — an introductory guide. *Soc. Cogn. Affect. Neurosci.* 4 (1), 101–109.
- Nestor, A., Plaut, D.C., Behrmann, M., 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc. Natl. Acad. Sci. U. S. A.* 108, 9998–10003.
- Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* 10, 424–430.
- Op de Beeck, H.P., 2010. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? *Neuroimage* 49 (3), 1943–1948.
- 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.
- Polyn, S.M., Natu, V.S., Cohen, J.D., Norman, K.A., 2005. Category-specific cortical activity precedes retrieval during memory search. *Science* 310, 1963–1966.
- Sasaki, Y., Rajimehr, R., Kim, B.W., Ekstrom, L.B., Vanduffel, W., Tootell, R.B., 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* 51 (5), 661–670.
- Shmuel, A., Chaimow, D., Raddatz, G., Ugurbil, K., Yacoub, E., 2010. Mechanisms underlying decoding at 7 T: ocular dominance columns, broad structures, and macroscopic blood vessels in V1 convey information on the stimulated eye. *Neuroimage* 49 (3), 1957–1964.
- Staeren, N., Renvall, H., De Martino, F., Goebel, R., Formisano, E., 2009. Sound categories are represented as distributed patterns in the human auditory cortex. *Curr. Biol.* 19, 498–502.
- 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. *J. Neurosci.* 30 (1), 325–330.
- Triantafyllou, C., Hoge, R.D., Krueger, G., Wiggins, C.J., Potthast, A., Wiggins, G.C., Wald, L.L., 2005. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage* 26, 243–250.
- Triantafyllou, C., Polimeni, J.R., Wald, L.L., 2011. Physiological noise and signal-to-noise ratio in fMRI with multi-channel array coils. *Neuroimage* 55, 597–606.
- Turner, R., 2002. How much cortex can a vein drain? Downstream dilution of activation-related cerebral blood oxygenation changes. *Neuroimage* 6, 1062–1067.
- Ugurbil, K., Adriany, G., Andersen, P., Chen, W., Garwood, M., Gruetter, R., Henry, P.G., Kim, S.G., Lieu, H., Tkac, I., Vaughan, T., Van De Moortele, P.F., Yacoub, E., Zhu, X.H., 2003. Ultrahigh field magnetic resonance imaging and spectroscopy. *Magn. Reson. Imaging* 21, 1263–1281.
- Uludağ, K., Müller-Bierl, B., Ugurbil, K., 2009. An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *Neuroimage* 48, 150–165.
- Valente, G., De Martino, F., Esposito, F., Goebel, R., Formisano, E., 2011. Predicting subject-driven actions and sensory experience in a virtual world with relevance vector machine regression of fMRI data. *Neuroimage* 56, 651–661.
- Vaughan, J.T., Garwood, M., Collins, C.M., Liu, W., DelaBarre, L., Adriany, G., Andersen, P., Merkle, H., Goebel, R., Smith, M.B., Ugurbil, K., 2001. 7T vs. 4T: RF power, homogeneity, and signal-to-noise comparison in head images. *Magn. Reson. Med.* 46, 24–30.
- Xue, G., Dong, Q., Chen, C., Lu, Z., Mumford, J.A., Poldrack, R.A., 2010. Greater neural pattern similarity across repetitions is associated with better memory. *Science* 330, 97–101.
- Yacoub, E., Shmuel, A., Pfeuffer, J., Van De Moortele, P.F., Adriany, G., Andersen, P., Vaughan, J.T., Merkle, H., Ugurbil, K., Hu, X., 2001. Imaging brain function in humans at 7 Tesla. *Magn. Reson. Med.* 45, 588–594.
- Yacoub, E., Duong, T.Q., Van De Moortele, P.F., Lindquist, M., Adriany, G., Kim, S.G., Ugurbil, K., Hu, X., 2003. Spin-echo fMRI in humans using high spatial resolutions and high magnetic fields. *Magn. Reson. Med.* 49, 655–664.
- Yacoub, E., Van De Moortele, P.F., Shmuel, A., Ugurbil, K., 2005. Signal and noise characteristics of Hahn SE and GE BOLD fMRI at 7 T in humans. *Neuroimage* 24, 738–750.
- Yacoub, E., Shmuel, A., Logothetis, N., Ugurbil, K., 2007. Robust detection of ocular dominance columns in humans using Hahn Spin Echo BOLD functional MRI at 7 Tesla. *Neuroimage* 37, 1161–1177.
- Yacoub, E., Harel, N., Ugurbil, K., 2008. High-field fMRI unveils orientation columns in humans. *Proc. Natl. Acad. Sci. U. S. A.* 105, 10607–10612.