Functional magnetic resonance imaging

Magnetic resonance imaging is a means of non-invasively imaging volumes in three dimensions and is commonly used to examine the structure of the brain for evidence of pathology. However, the brain's structure (at least from such a macroscopic perspective) does not necessarily determine function.

Since the development of phrenology by Franz Joseph Gall in the late 18th century, identifying brain regions with specific functions has become a central theme in neuroscience, and has arguably guided neuroimaging and 'brain mapping' over the past few decades. Centuries of scientific exploration have concluded that the functions cannot simply be localized to parts of the brain. Rather, the brain adheres to two fundamental and mutually dependable properties of functional organization:

1. Functional segregation, which states that the cortical infrastructure supporting a behaviour or function involves multiple components that are specialized in processing task-relevant information, and are often spatially distributed
2. Functional integration, which states that proper function is dependent upon optimal integration, or communication, of specialized components (Friston, 1994, 2011).

The specialization of a given cortical region is determined by its functional connections, and vice versa, thus any disturbance in either the viability of a specialized brain region or the brain's ability to integrate information will lead to a disturbance of proper function. From a clinical perspective, this could produce signs, symptoms or a clinical phenotype. Early lesion studies and invasive electrical stimulation observations were once the only way to ascertain cortical function. In the early 1990s, it was discovered that magnetic resonance imaging scanners could be used to provide an alternate, non-invasive means to map cerebral function. So-called 'functional magnetic resonance imaging' has the added benefit of being able to sample function across the whole brain (almost) simultaneously. This article discusses the basics of functional magnetic resonance imaging, and how it is used in neuroscientific research.

Dr Joshua Kahan* *is MB PhD student in the Sobell Department of Motor Neuroscience and Movement Disorders and* **Dr Stephen Auger*** *is MB PhD student in the Wellcome Trust Centre for Neuroimaging, UCL Institute of Neurology, London WC1N 3BG*

Correspondence to: Dr J Kahan (joshua.kahan@ucl.ac.uk)
**Both authors contributed equally to this article*

The BOLD contrast mechanism

Active neurons consume both glucose and oxygen but have no internal stores of either substance. When metabolic activity within a brain region increases, it therefore receives an increase in blood flow to match the demand. This 'neurovascular coupling' is tightly regulated in space and time.

Oxygenated and deoxygenated haemoglobin in the blood have different magnetic properties that can be detected with a magnetic resonance imaging scanner. Functional magnetic resonance imaging measures this blood oxygen level dependent (BOLD) signal as a means of tracking the haemodynamic response accompanying neural activity in the brain (Ogawa et al, 1990, 1992). Using functional magnetic resonance imaging, it is therefore possible to infer which areas in the brain are more active when people perform specific tasks.

The brain is separated into three-dimensional units called voxels, representing small cubes of data (usually ≈ 3x3x3 mm, corresponding to approximately a million neurons) collected from a specific location within the brain. This 'spatial resolution' of functional magnetic resonance imaging (how precisely it indicates where activity is occurring) is better than that of other comparable neuroimaging methods, such as electroencephalography or positron emission tomography.

Numerous scans are acquired to track changes in the BOLD signal over time and each scan usually takes a few seconds to acquire. This 'temporal resolution' of functional magnetic resonance imaging (how precisely it indicates when activity is occurring) compares favourably to positron emission tomography, but is inferior to that of electroencephalography which operates over a period of milliseconds.

Functional magnetic resonance imaging, therefore, provides a non-invasive, radiation-free technique to evaluate the function of the whole brain with a relatively high degree of spatial precision over a period of seconds.

Analysis

Locating function in the brain

After collecting the data from a magnetic resonance imaging scanner, it is first 'preprocessed'. Preprocessing is a series of transformations applied to the raw data in order to ensure that the serial images are:

1. Spatially aligned with each other
2. Aligned to a common template (used throughout the imaging literature to allow comparisons between studies)
3. Optimized with regard to signal-to-noise for statistical analysis.

Standard functional magnetic resonance imaging statistical analyses traditionally consider activity in every individual voxel separately using a general linear model.

At each voxel, the general linear model is used to test the hypothesis that the observed data (the BOLD signal at various time points – y in *Figure 1*) can be explained by a linear combination of explanatory variables (e.g. the task being performed – x1 in *Figure 1*), with some residual error (c in *Figure 1*). The onsets and durations of stimuli presentation are specified, creating a 'neural model', which essentially tells the analysis when the stimuli are present. The neural model is then combined with (convolved with) the so-called 'haemodynamic response function' to produce a model of the haemodynamic effect that stimuli elicit. The resulting 'haemodynamic model' attempts to predict what a BOLD response to the stimuli onsets would look like.

The haemodynamic response function takes into account the time delay between neural activity occurring and the associated

increase in supply of oxygenated blood. Other covariates of no experimental interest can also be included in the model, to account for potential sources of noise such as head movement.

Model parameters are then estimated for every voxel to provide the best fitting model of the observed functional magnetic resonance imaging responses. The output of this analysis is a set of parameter estimates (β) for each explanatory variable at every voxel. The β values correspond to how much the explanatory variable predicts the observed BOLD signal, taking into account the noise and covariates of no interest (*Figure 1*).

The β values from each voxel for a given explanatory variable can be interrogated using classical statistics to produce an 'image' of statistics, in other words identifying which voxels in the brain are most significantly predicted by the stimulus presentation. In crude terms, it shows 'blobs' on brains, highlighting which areas are particularly active for certain experimental conditions – see example in *Figure 2*.

The β values from each voxel can be combined with comparable values from other subjects to look for group effects. In turn, groups of individuals (e.g. patients *vs* controls) can similarly be compared to identify potential biomarkers of disease.

Beyond 'blobology'

The standard analysis method described above, known as the 'mass-univariate' approach (or its more derisory alternative 'blobology'), considers voxels in isolation, looking for those which show a linear increase in activation in response to certain experimental conditions. However, there are numerous powerful new techniques which offer additional opportunities for interrogating functional magnetic resonance imaging data.

More subtle differences in neural representations related to the patterns of activity across multiple voxels can be examined using machine learning algorithms. One such technique, multi-voxel pattern analysis, can be used to decode patterns of activation within a brain region to make accurate predictions about what a subject was doing at the time; for example, which of two different memories he/she was recalling (Chadwick et al, 2010), or what type of items he/she was viewing (Auger and Maguire, 2013).

Multi-voxel pattern analysis and related techniques provide a more sensitive means to examine subject-specific patterns of activity and how they relate to specific experimental conditions. This contrasts with mass-univariate analyses which reveal information about the variability and mean levels of activity across numerous participants. Both these techniques are useful for investigating how responses within a brain region relate to the demands of different tasks. However, given the widespread structural and functional associations between different brain regions, there is also a need to consider these interactions. This can be done with functional magnetic resonance imaging using various types of connectivity analysis.

Figure 1. Locating function using the general linear model. A 'neural model' is specified – the blue time series – that shows activity increases during presentation of the stimulus (the happy face). The haemodynamic model – the red time series – is simply the neural model combined with the haemodynamic response function (see glossary), forming an explanatory variable (x_1) in the general linear model. For each voxel, a β value is estimated, which is converted to a T statistic to determine the strength of the effect of x_1 on y.

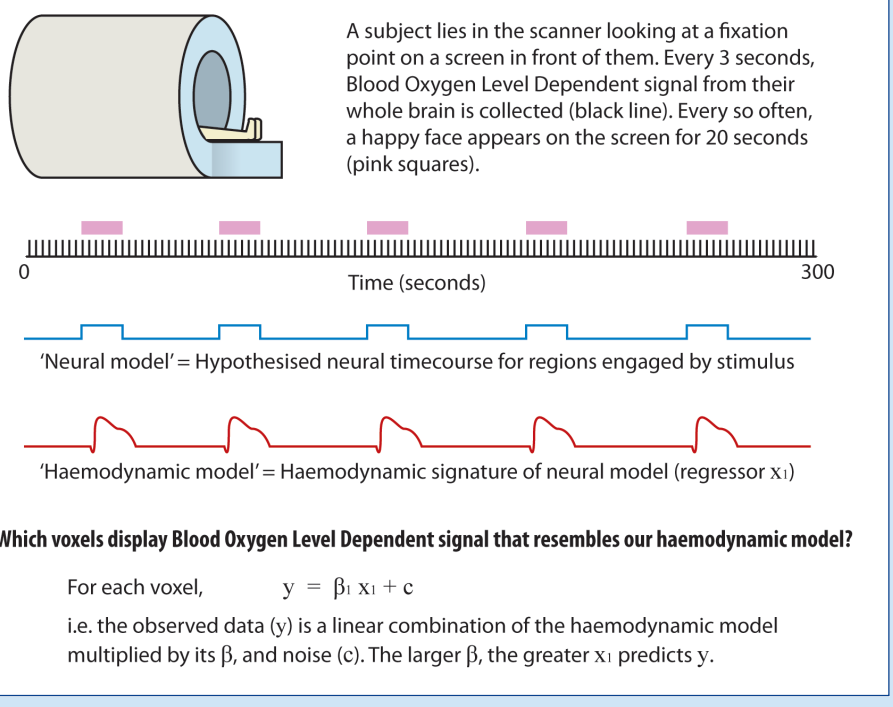

'Connectivity' and functional magnetic resonance imaging

The 100 billion neurons that comprise the human brain do not work in isolation. Anatomically, these neurons form a relatively invariant structural architecture, upon which rests a dynamic, modular and hierarchical functional network, capable of the rich computations that underpin cognition. The brain can be considered a large and densely connected network, composed of sub-networks, which are themselves comprised of many sub-modules that compute lower-level operations (Buckholtz and Meyer-Lindenberg, 2012; Park and Friston, 2013). Defining a connection in vivo has traditionally been more complicated than defining a component (or node) of a network; definition of nodes has been discussed above. At the macro-level, the literature generally examines three types of connectivity:

Structural connectivity

Structural connectivity refers to the abundance of white matter projections between nodes, usually inferred using diffusion tensor imaging – not functional magnetic resonance imaging. These measures are unable to determine direction or differentiate excitatory or inhibitory connections. For this, some tracing techniques can be used, but obviously are limited to post-mortem examination.

Functional connectivity

Functional connectivity refers to any statistical dependency between functional time series of distant nodes. This is most commonly a correlation between the BOLD signal of two regions. As before, this form of connectivity is symmetrical, and is unable to identify excitatory or inhibitory connections, simply that the functional signature of two regions is correlated.

One feature of functional organization that became apparent from the early days of functional magnetic resonance imaging is that regions with similar functional specialization display significant functional connectivity both during task performance, and while subjects lie at rest in the scanner (Biswal et al, 1995). Given the relative ease of collecting data in the 'resting state' in patient cohorts, and the existence of endogenous functional connectivity, these forms of analyses have become increasingly popular in biomarker development. For example, resting state functional magnetic resonance imaging biomarkers with sensitivities and specificities of >80% have been developed for Alzheimer's disease (Chen et al, 2011), and machine learning techniques have been able to classify schizophrenic patients from healthy controls at accuracies varying from 75% to 92% (Zarogianni et al, 2013).

Effective connectivity

Effective connectivity is defined as the influence that one neural population exerts on another under a particular network model. In contrast to functional connectivity, effective connectivity describes directed connections, which may be context-sensitive, and can be either excitatory or inhibitory. The advantage of such analyses is that they provide a simplified model of information transfer in a given circuit, permitting mechanistic hypotheses to be tested (Kahan and Foltynie, 2013; Kahan et al, 2014).

One such published example comes from patients with primary progressive aphasia. While patients with primary progressive aphasia tend to engage an identical language network during semantic and phonological tasks, Sonty et al (2007) demonstrated reduced language-specific effective connectivity between Wernicke's and Broca's areas that was predictive of semantic task accuracy, highlighting the role of dysfunctional network interactions in primary progressive aphasia.

Common pitfalls and the dead fish

As with all real life data collection, functional magnetic resonance imaging suffers a number of limitations that readers should acknowledge before drawing their conclusions. First, as discussed above, the BOLD signal is an indirect measure of neural activity, specifically the haemodynamic response to neural dynamics. That is not to say that collected data are purely a haemodynamic phenomena; the haemodynamic response to discrete neural events has been extensively studied, and most contemporary analysis packages model the data in terms of a convolution of unobserved (hidden) neural events with a canonical haemodynamic response function. In other words, signals that are neural ultimately cause the observed data. However, such haemodynamic repercussions typically occur over a much slower timescale (4–15 seconds) than their neural causes (milliseconds), somewhat limiting the temporal resolution.

Furthermore, the BOLD signal appears to be most closely predicted by incoming synaptic activity arriving at a region, as opposed to its output (Logothetis and Wandell, 2004), and appears indiscriminate of excitatory or inhibitory drive. Further work into the origins of the signal, including simultaneous functional magnetic resonance imaging and invasive electrophysiology, is ongoing (Logothetis, 2008).

One confused, but often cited criticism of functional magnetic resonance imaging

Figure 2. An example analysis using a motor task. The participant was equipped with a pair of headphones and a joystick in his left hand and was asked to lie in the scanner with his eyes closed. Every 1–3 seconds, an audio tone was heard, and every 30 seconds a voice would say 'rest' or 'move'. The subject was instructed to keep his hand on the joystick and ignore the tones during the 'rest' blocks. During 'move' blocks, the subject was asked to move the joystick in a direction of his choice with each tone, and then return the joystick handle back to the centre point. A neural model was specified as an alternating 'boxcar function', demarcating the 'rest' and 'move' periods. The neural model was convolved with the haemodynamic response function, and a general linear model analysis was performed asking which voxels were significantly predicted by left hand movement. Put simply, the general linear model was built to identify voxels with greater activity during 'move' periods than 'rest' periods. The resulting heat map (statistical parametric map) was thresholded at P<0.05 (family-wise error corrected for multiple comparisons) with a cluster size threshold of 100 voxels. Two axial slices are shown at different levels. The left slice at z=-26 mm demonstrates bilateral cerebellar activity, greater in the left cerebellar hemisphere. The right slide at z=60 mm demonstrates increased activity in the primary motor and sensory cortices (precentral and postcentral gyri respectively), as well as lateral premotor area and midline supplementary motor area activation in the frontal lobe.

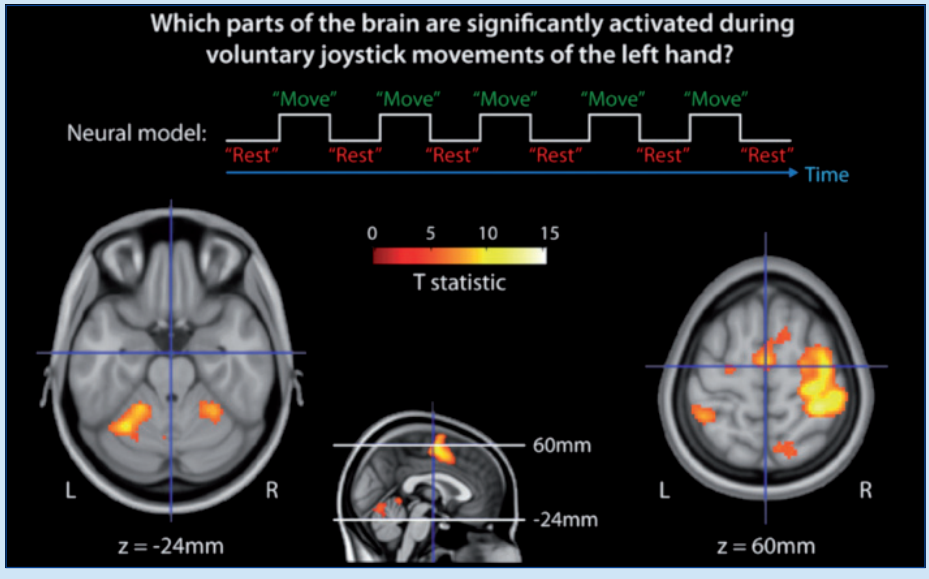

relates to spurious activations, and the findings of Bennett et al (2010). To summarize, Bennett et al whimsically but importantly demonstrated the multiple comparisons problem by detecting voxels 'significantly' activated by a cognitive task 'performed' by a dead salmon. The multiple comparisons problem is common to all classical statistics and simply states that if your significance threshold is 5% (i.e. $P<0.05$), then if you performed 100 independent tests, you would expect 5 to show an effect as a result of chance alone. The mass univariate approach discussed above essentially performs thousands of such tests, thus is prone to false positive errors. However, common practice is to correct for this post-hoc using methods that are robust to such errors (for example by using family-wise error corrected statistics). Indeed, when Bennett et al (2010) repeated their analysis using corrected statistics, no spurious activity was detected.

Conclusions

Functional magnetic resonance imaging offers a non-invasive means of collecting brain activity data. This article has discussed common analysis pipelines and its use in inferring brain connectivity, as well as its common pitfalls. **BJHM**

Conflict of interest: none.

GLOSSARY

- **Blood oxygen dependent (BOLD) signal** — This is the signal collected by the magnetic resonance imaging scanner that is sensitive to the proportions of oxygenated and deoxygenated haemoglobin. When neural tissue becomes more metabolically active, a functional hyperaemia occurs, producing an excess of oxygenated haemoglobin in the vicinity of the active neurons. As a result, the BOLD signal collected by the scanner increases.
- **General linear modelling (GLM)** — This is a statistical method that is used in a number of different fields. The essence of GLM analyses is to be able to identify (a combination of) factors that predict an outcome measure. In the context of functional magnetic resonance imaging, the GLM is used to identify voxels that are predicted by experimental factors such as viewing a smiling face – see *Figure 1*.
- **Haemodynamic response function (HRF)** — This takes into account the time delay between neural activity occurring and the associated increase in supply of oxygenated blood. The peak in the BOLD signal occurs about 6 seconds after the onset of neural activity and reflects the time taken for an increased metabolic demand to be detected and bring about dilation of local vasculature. The function of the HRF is essentially to make the explanatory variable as realistic as possible, such that as much variance in the BOLD signal could be explained by the explanatory variable.
- **Multi-voxel pattern analysis (MVPA)** — Instead of considering each voxel in isolation, MVPA examines how patterns of activation across numerous different voxels relate to experimental factors.

Auger SD, Maguire EA (2013) Assessing the mechanism of response in the retrosplenial cortex of good and poor navigators. *Cortex* **49**: 2904–13 (doi: 10.1016/j.cortex.2013.08.002)

Bennett CM, Baird AA, Miller MB, George L (2010) Neural correlates of interspecies perspective taking in post-mortem Atlantic salmon. *J Serendipitous Unexpected Results* **1**: 1–5

Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar *MRI Magn Reson Med* **34**: 537–41 (doi: 10.1002/mrm.1910340409)

Buckholtz JWW, Meyer-Lindenberg A (2012) Psychopathology and the human connectome: toward a transdiagnostic model of risk for mental illness. *Neuron* **74**: 990–1004 (doi: 10.1016/j.neuron.2012.06.002)

Chadwick MJ, Hassabis D, Weiskopf N, Maguire EA (2010) Decoding individual episodic memory traces in the human hippocampus. *Curr Biol* **20**: 544–7 (doi: 10.1016/j.cub.2010.01.053)

Chen G, Ward BD, Xie C et al (2011) Classification of Alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional MR imaging. *Radiology* **259**: 213–21 (doi: 10.1148/radiol.10100734)

Friston KJ (1994) Functional and effective connectivity in neuroimaging: A synthesis. *Hum Brain Mapp* **2**: 56–78 (doi: 10.1002/hbm.460020107)

Friston KJ (2011) Functional and Effective Connectivity: A Review. *Brain Connect* **1**: 13–36 (doi: 10.1089/brain.2011.0008)

Kahan J, Foltynie T (2013) Understanding DCM: ten simple rules for the clinician. *Neuroimage* **83**: 542–9 (doi: 10.1016/j.neuroimage.2013.07.008)

Kahan J, Urner M, Moran R et al (2014) Resting state functional MRI in Parkinson's disease: the impact of deep brain stimulation on 'effective' connectivity. *Brain* **137**: 1130–44 (doi: 10.1093/brain/awu027)

Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. *Annu Rev Physiol* **66**: 735–69 (doi: 10.1146/annurev.physiol.66.082602.092845)

Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* **453**: 869–78 (doi: 10.1038/nature06976)

Ogawa S, Lee TM, Kay AR, Tank DW (1990) Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci USA* **87**: 9868–72 (doi: 10.1002/mrm.1910140108)

Ogawa S, Tank DW, Menon R et al (1992) Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci USA* **89**: 5951–5 (doi: 10.1073/pnas.89.13.5951)

Park H-J, Friston K (2013) Structural and functional brain networks: from connections to cognition. *Science* **342**: 1238411 (doi: 10.1126/science.1238411)

Sonty SP, Mesulam M-M, Weintraub S, Johnson NA, Parrish TB, Gitelman DR (2007) Altered effective connectivity within the language network in primary progressive aphasia. *J Neurosci* **27**: 1334–45 (doi: 10.1523/JNEUROSCI.4127-06.2007)

Zarogianni E, Moorhead TWJ, Lawrie SM (2013) Towards the identification of imaging biomarkers in schizophrenia, using multivariate pattern classification at a single-subject level. *NeuroImage Clin* **3**: 279–89 (doi: 10.1016/j.nicl.2013.09.003)

KEY POINTS

- Functional magnetic resonance imaging offers a non-invasive means of measuring whole-brain activity at a relatively high degree of spatial precision, at temporal resolutions in the order of seconds.
- Inferences are made upon changes in an indirect measure of blood oxygenation, which is an indirect measure of regional metabolism, thus neuronal activity.
- General linear modelling is commonly used to identify voxels displaying significant effects of experimental manipulations.
- Correlation analyses as well as biophysical modelling are used to quantify connectivity among regions of the brain, both during tasks and while at rest.