ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 30 (2006) 1077 – 1087

Comments and Controversies

A critique of functional localisers

K.J. Friston,[a,*] P. Rotshtein,[a] J.J. Geng,[b] P. Sterzer,[a] and R.N. Henson[c]

[a]*The Wellcome Department of Imaging Neuroscience, Institute of Neurology, UCL, 12 Queen Square, London WC1N 3BG, UK*
[b]*Institute of Cognitive Neuroscience and Department of Psychology, University College London, 17, Queen Square, London WC1N 3BG, UK*
[c]*MRC Cognition and Brain Sciences Unit, Cambridge CB2 2EF, UK*

Received 4 May 2005; revised 19 July 2005; accepted 15 August 2005

In this critique, we review the usefulness of functional localising scans in functional MRI studies. We consider their conceptual motivations and the implications for experimental design and inference. Functional localisers can often be viewed as acquiring data from cells that have been removed from an implicit factorial design. This perspective reveals their potentially restrictive nature. We deconstruct two examples from the recent literature to highlight the key issues. We conclude that localiser scans can be unnecessary and, in some instances, lead to a biased and inappropriately constrained characterisation of functional anatomy.

Introduction

The use of functional localisers to constrain the analysis of fMRI data is becoming popular in neuroimaging. This approach entails a separate experiment to localise areas in the brain that serve to guide, constrain or interpret results from a main experiment. The need and motivation for functional localisers are often not stated explicitly and is sometimes unclear. Nevertheless, several colleagues have encountered reviewers who thought that omission of a functional localiser did not conform to good or standard practice. The purpose of this commentary is to provide a reference for people who do not want to use functional localisers and have to defend themselves against the contrary attitudes of reviewers (see Appendix A for some verbatim comments).

The term "functional localiser" is generally used in the context of stereotactic neurosurgery or radio-surgical treatment planning. It refers to a functional (e.g., fMRI) experiment that is used to disclose eloquent cortex (e.g., Liu et al., 2000). The term functional distinguishes this localisation from the anatomic information in structural MRI or CT scans. The human brain mapping community has adopted this term to refer to an auxiliary fMRI experiment that constrains the analysis or interpretation of a main fMRI experiment. Although every fMRI study is a study of functional localisation in the human brain, we will take functional localiser to mean a separate scanning session that has been divorced from the functional experiment proper.

Our aim is to frame some issues that may be useful when motivating and critiquing the use of localisers (or not using them). Specifically, we focus on four issues:

- Functional regions of interest (fROI), such as the Fusiform Face Area (FFA), the Lateral Occipital Complex (LOC) or Visual Word Form area (VWFA), are often viewed as useful vehicles to characterise functional anatomy. Although fROI are sufficient to establish functional selectivity, they preclude inferences about functional specialisation. This is because functional specialisation entails anatomical specificity (i.e., the specialised region exhibits more functional selectivity than another region). This anatomical specificity cannot be addressed with a single fROI.
- The validity of fROI constructs depends on their context-sensitivity. For example, the VWFA may process words in one context, but not another. Equivalently, the voxels comprising the FFA may change when processing one facial attribute relative to another. Unless an fROI is context-invariant, it may not provide the most appropriate constraint to analyse responses in a different context. Indeed, the introduction of factorial designs to neuroimaging was driven by context-sensitive specialisation implied by interactions between factors, and the empirical failure of pure insertion (Friston et al., 1996). This leads to the next point:
- Separate localiser designs often represent missed opportunities, in relation to factorial designs. Eliciting main effects in the localiser and main experiments separately precludes tests for interactions (i.e., differences in activation between the localiser and main sessions). Localiser designs could be regarded as a slightly retrograde development in experimental design, in that inferences about effects in the main experiment usually rest on simple subtraction and pure insertion. Clearly, there are many designs that cannot be made factorial. However, when they can,

* Corresponding author. Fax: +44 207 813 1445.
E-mail address: k.friston@fil.ion.ucl.ac.uk (K.J. Friston).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2005.08.012

there are compelling reasons to use factorial designs. Our main point here is "people who like localisers should like factorials even more." (Jon Driver, personal communication).

- The practice of averaging responses over voxels in an fROI has clear advantages in simplifying analyses. However, averaging entails some strong assumptions about responses that practitioners may not have considered. Some of these assumptions are untenable (e.g., homogeneity across the fROI), and historically led to the development of voxel-based analysis (i.e., Statistical Parametric Mapping). In this sense, fROI represents another retrograde development.
- An advantage of fROI averages is that they summarise subject-specific responses without assuming anatomical homology over subjects. However, there may be more sensitive and principled approaches to the problem of functional–anatomical variability.

There are several other issues, which are not the subject of this critique. These include:

- The labelling of regional responses using anatomic or functional information from another study. This study could be a localiser, a retinotopic mapping study or indeed someone else's study of a related effect. We are concerned only with functional localisers that are used to constrain analysis, not the post hoc use of localising information to label the results of an analysis.
- The use of functional [as opposed to anatomical] constraints to characterise regional responses. This is an important component of many analyses, particularly in the context of factorial designs. Functional constraints per se are essential for hypothesis-led and powerful inference. Our focus is on the use of fROI averages to summarise regional responses, not on their useful role in constraining searches to regional responses within the fROI.
- The principled use of localisers in designs that require separate sessions or are not inherently factorial.

This commentary is organised as follows. First we review the functional localiser approach to fMRI data analysis, its motivation and its relation to conventional multifactorial experimental designs. We then deconstruct two recent experiments reported in the literature, which used functional localisers, to reprise the main points of the first section. The first is an example of functional localisers in the study of object-defining properties. The second uses functional localisers to characterise evoked responses associated with attentional shifts.

Theoretical issues

Functional localisers and fROI

Functional localisers can be thought of as the splitting of a study into two parts. One part (the localiser) involves the comparison of two or more conditions (e.g., pictures of faces vs. pictures of houses) to isolate a functionally specialised region (e.g., the "fusiform face area", Kanwisher et al., 1997). The other part constitutes the main experiment and usually involves the comparison of further conditions that have not been explored previously (e.g., pictures of dogs that are named either at the basic level – "Dog" – or subordinate level – "Dalmatian"). These are used to establish the functional selectivity of the fROI. The results of the functional localiser (i.e., the fROI) are used to constrain, anatomically, the search for effects in the main experiment. When the conditions in the functional localiser and in the main experiment share a factor, there is an implicit factorial design, in which the localiser can be considered as a level in an extra factor (localiser vs. main). See Fig. 1 for two examples. Functional localisers raise two issues. First, why include an extra factor in the experimental design and second, why perform it separately? This section addresses these two questions.

Why use functional localisers?

Attribution vs. constraints

Functional localisers assume that there is some effect that has an important role in interpreting or constraining the analysis of the effects in the main experiment. One common use of localisers is to inform the labelling of region-specific effects identified in a

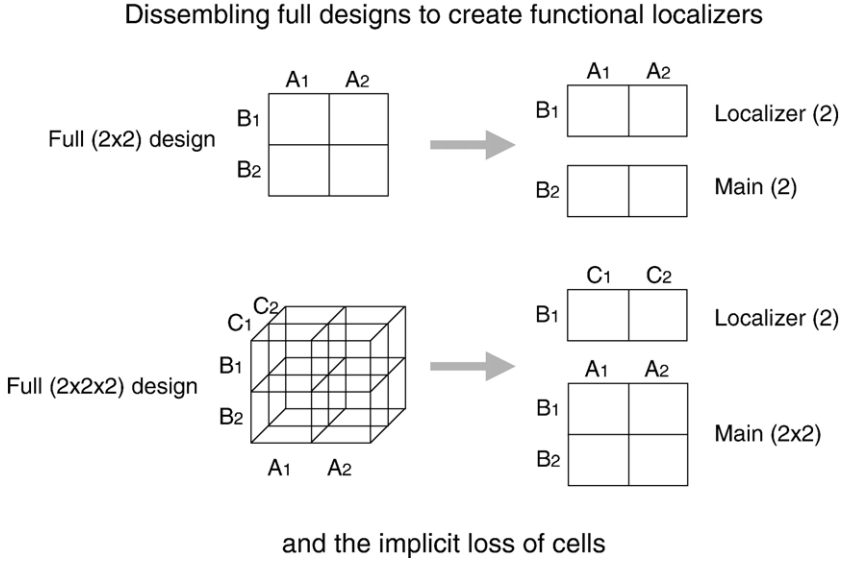

Fig. 1. A schematic showing the relationship between functional localisers and their multifactorial parents.

separate analysis of the main experiment (e.g., retinotopic mapping). Here, the results of the localiser are either used to a priori constrain the analysis of the main experiment to a subset of voxels according to some cortical parcellation scheme (e.g., retinotopy) or they can be used post hoc to assign labels to regional effects of the main experiment (e.g., V2). The usefulness of such constraints depends on their anatomic and physiological validity. The use of functional localisers is a common procedure in studies of early visual processing, because the fine-grained retinotopic organisation of early visual areas is well established and largely cue-invariant. This calls for a characterisation of structure–function relationships on an appropriate spatial scale.

This critique is not concerned with the use of localisers that inform the analysis, attribution or interpretation post hoc. It is concerned with the use of localisers as constraints on the analysis per se. These constraints may be valid or they may not be. There appear to be several ways in which such constraints operate. First, the inclusion of the localiser enables one to constrain the search for significant effects of the main factors by only looking in areas activated by the localiser. Here, the objective is to increase the sensitivity of searches for regionally specific effects in the main experiment (by reducing the problem of multiple statistical comparisons). Second, one can assess the effects of the main factors on the average response of the areas activated by the localiser. The latter is a special case of constrained inference that is not concerned with where effects are expressed (the localising fROI prescribes this). The question here is what effects the main factors have on the average response. We will deal with these two cases in turn.

Functionally constrained searches

If the localiser is chosen carefully, it can provide a tremendous increase in sensitivity because it reduces the multiple comparisons problem, entailed by searching over large volumes of brain. The anatomical constraints afforded by functional localisers usually take the form of regions of interest. These are defined operationally by reliable effects in the localiser. The search for significant effects in the main experiment can then be constrained to voxels showing maximal responses within the fROI (ideally using a Bonferroni correction for the number of such maxima). Alternatively, a random field correction, as implemented in SPM, is applied to all voxels within the fROI. This is known as a small volume correction or SVC (see Worsley et al., 1996). Constrained searches would be preferred if one wanted to search for functionally heterogeneous responses within an fROI. For example, only part of the V5 complex, defined by visual motion, may be engaged during the perception of apparent motion.

Constrained searches, of this sort, are standard practice in conventional fMRI analyses that do not use functional localisers. Perhaps the most common example are searches for differences among event types in event-related designs that are limited to regions that respond to all events, relative to the inter-event baseline. This constrained search relies on the fact that differences among event-related responses are orthogonal to their average (i.e., do not bias the inference). In balanced factorial designs, the main effects and interactions are, by design, orthogonal. This means that one can take the maxima, or fROI, of one effect and constrain the search for the other effects to the ensuing voxels. For example, two effects might define plasticity in the motor system: a main effect of movement (motor-related responses) and a movement-by-time interaction (learning-related changes in those responses). The search for interactions can be restricted to maxima exhibiting a main effect of movement to infer motor-plasticity. This example highlights an important point. Namely, main effects usually constrain the search for interactions. This is because finding a significant interaction between factors A and B in the presence of a main effect of A, entitles one to say that responses to A, in A-selective regions, depend on B. We will see below that these interactions are usually precluded in localiser designs.

This well-established approach to constrained statistical searches does not rely upon a separate localiser. The key thing to appreciate is that a contrast testing for a particular effect can be used as a localiser for the remaining [orthogonal] effects. In this sense, any factorial fMRI study has as many functional localisers, embedded within it, as there are orthogonal contrasts. A typical two-by-two design has three orthogonal contrasts. The natural conclusion is that all fMRI experiments are simply collections of functional localisers. This is quite sensible given that human brain mapping is about functional localisation. In short, the use of localising contrasts to provide constraints on the search for orthogonal effects has a long history (e.g., Friston et al., 1996), is principled and rests an explicit or implicit [localiser] factorial design that may, or may not be balanced.

It should be noted that this approach assumes a modular functional architecture. For example, functional constraints based on faces vs. houses assumes that all face-related processes occur only within regions that show a stronger response to faces compared to houses. More generally, one should be aware that constraining the search for interaction to regions showing main effects will miss crossover interactions. For example, a region that activates in one context but deactivates in another will exhibit no main effect and will be missed using a contained search procedure. As with all constraints, they should be used in an informed and careful fashion. Next, we discuss approaches that do not constrain the search for functionally selective responses but examine the functional selectivity of the fROI itself.

Functional ROI

The second reason one might want to use a localiser is to restrict the analysis to the responses of the fROI itself (i.e., responses averaged over voxels within the fROI). In this case, there is only one statistical inference and no need to adjust the P value. The motivation goes beyond simply increasing sensitivity, because the nature of the response variable is changed qualitatively, from a collection of regional responses at each voxel, to a summary of their collective response i.e., average. In this context, responses elsewhere in the brain are uninteresting; the researcher has reduced functional anatomy to a single brain region, defined operationally by the functional localiser. This is perfectly tenable, with the qualification that inferences relate to, and only to, some ad hoc fROI.

This approach has the advantage of being focussed, providing for uncomplicated accounts of responses within a pre-specified fROI. However, fROI could be regarded as colloquial in the sense that they are not derived from any formal functional ontology. The point here is that an ad hoc fROI cannot participate in structure–function ontologies unless it has some structurally invariant functional properties. In short, a useful fROI should comprise the same voxels in different contexts. There are many examples of functionally segregated regions that do this. For example, all the anatomical and physiological evidence suggests that V5 is specialised for motion processing. Similarly, the anatomical profile

of selectivity in V1 units is largely context-invariant (e.g., ocular dominance columns). However, even in V1, extra-classical receptive field effects and attentional modulation confer context-sensitivity at some level of analysis. This context-sensitivity is assessed with interactions in factorial experiments. The simple and logical critique of fROI, defined by localisers, is that the context-sensitivity of their anatomy cannot be assessed because localiser designs preclude the assessment of interactions. This means that localisers are unable to establish the structural invariance properties of the fROI they are designed to study.

Functional selectivity vs. functional specialisation

A more fundamental problem with fROIs is that they preclude any inferences about functional specialisation in the brain. There is a subtle but important distinction between functional selectivity and functional specialisation. Functional selectivity is defined operationally by demonstrating functional responses in a single unit or area that are selective or specific to some stimulus or task attribute (i.e., orientation of a visually presented bar or category of an object). Functional selectivity implies specificity in terms of what elicits a neuronal response. In contrast, functional specialisation is not an operational definition; it is an inference that a particular unit or brain area specialises in some function or computation. This inference rests on an anatomical specificity in terms of where functionally selective responses are expressed. For example, V5 or MT expresses functionally selective responses to motion and is functionally specialised for motion because motion-selective responses are restricted largely to this area. In short, functional selectivity implies responses that are specific to a domain of function or stimulus-space. Conversely, the specificity in functional specialisation refers to structural or anatomical space.

The distinction between functional selectivity and specialisation has important implications for fROIs; fROIs are entirely sufficient to establish functional selectivity because one can examine their responses in a large number of contexts. However, they cannot be used to infer functional specialisation because they are blind to responses elsewhere in the brain. Put simply, for every face-related process, there may be another area that expresses a more selective response than the FFA. The only way one would know this would be to employ a conventional SPM analysis. A simple example may help to clarify this point.

Imagine a study of implicit face recognition using an incidental same-different judgement task. The design has two factors; face vs. non-face and familiar vs. unfamiliar. A conventional analysis reveals a main effect of faces in the FFA and a main effect of familiarity higher in the ventral stream. Critically, the interaction between faces and familiarity elicits the greatest response in a face-recognition-area (FRA). One might then infer that the FRA was specialised for facial recognition and that the implicit recognition of faces involved a set of areas that included the FFA and FRA, specialised for pre-semantic processing and recognition, respectively. fROI analyses afford a much more limited inference. The FFA would show a main effect of faces and may indeed show a face by familiarity interaction that is mediated by backward connections from the FRA. The proper conclusion here is that the FFA is functionally selective for faces and face recognition. However, the inference that the FFA is specialised for facial recognition would be wrong because of the anatomical bias imposed by the fROI. This face-recognition study is not a thought experiment. It was reported in Gorno-Tempini et al. (1998): "The areas specialised for the perceptual analysis of faces (irrespective of whether they are famous or non-famous) are the right lingual and bilateral fusiform gyri, while the areas specialised for famous stimuli (irrespective of whether they are faces or names) spread from the left anterior temporal to the left temproparietal regions. One specific area, the more lateral portion of the left anterior middle temporal gyrus, showed increased activation for famous faces relative to famous proper names". Note that Gorno-Tempini et al. are entitled to talk about specialisation because they used an SPM analysis to characterise functional selectivity throughout the ventral stream.

The above example highlights the difficulties fROI-based characterisations contend with, when a cognitive process relies on several areas. This is particularly relevant to visual processing hierarchies in which the FFA and VWFA reside. The responses of these areas are subject to both bottom-up and top-down effects (Friston, 2003) and call for an analysis of selective responses that is not constrained to a single region. It had already been shown that face-selective responses were subject to top-down influences from as far away as the parietal cortex before the FFA was described (Dolan et al., 1997).

In short, to reduce functional anatomy to fROIs, and their functional selectivity, assumes we know a priori the parcellation and segregation of function within the cortex. Furthermore, it assumes that the voxels comprising the fROI do not change with neuronal context (McIntosh, 2000) or the level of task analysis. Some might argue it is pre-mature to invoke fROI to characterise functional anatomy. The equivalent agenda for "Areae anatomicae" (Brodmann, 1909), using anatomic criteria, is still incomplete after more than a century's work (Kötter and Wanke, 2005).

The background to ROI

The imaging community has already entertained the debate about ROI in the early days of brain mapping with PET. Initially, in the late 80s, people reported their results using ROIs defined using structural anatomy, perfusion or receptor binding. Note that these ROIs were based on defining characteristics of the underlying tissue and did not reflect any functional role of that region (i.e., were not fROI). These ROIs were assumed to be a useful summary of the distributed patterns of activity evoked. However, the problem was that ROIs pre-empted the questions they were supposed to answer, namely, where are region-specific responses expressed? Put simply, an ROI, was the result, not the hypothesis. One example of the pitfalls of ROI is the study of ventricular enlargement in schizophrenia. Because the ventricular ROI is easy to measure, it was the focus of imaging research in schizophrenia for almost a decade. The ventricles have no role in the pathophysiology of schizophrenia and, not surprisingly, this research went nowhere. The ad hoc and unprincipled basis for the parcellation of functional anatomy into ROI prompted the development of voxel-based approaches, namely statistical parametric mapping. The reprise of ROI in the form of fROI, many years later, raises the same issues, in a somewhat more subtle way. Now, the ad hoc nature of fROI lies in the choice of the localiser that defines the region. For example, is the only appropriate definition of the FFA the contrast of faces vs. houses? If this is so by convention, what if the convention is wrong?

On a more positive note, fROI have utility in well-established research programmes focussing on a specific part of the brain. Operational constructs like the FFA or VWFA play an important role in focussing experiments and enabling scientific exchange at a

colloquial level. For example, the adoption of a "standard" definition of the FFA (as the set of voxels more active for faces than, say, houses) allows researchers to compare directly the effects of different manipulations. This can proceed without relying on anatomical criteria, and ameliorating the effects of differences in the spatial properties of their functional images. Another 'virtue' to arise from the fROI tradition is the emphasis on replication (focus on activations that are evident in multiple separate experiments) and refinement (the use of increasingly subtle comparisons to establish selectivity). In summary, a fixed operational definition of fROI provides a strict and rigorous way of accumulating evidence across studies. However, it must be remembered that such operational definitions make assumptions about context-invariance and functional ontology that may not turn out to be true.

fROI averaging

One reason for averaging within an fROI is the assumption it increases signal-to-noise. Unfortunately, this assumption is not necessarily true. The most efficient averaging depends on the how the signal and noise are deployed within the fROI. Generally, the best (minimum variance) unbiased estimate of the fROI response would involve spatially 'whitening' the data, accounting for spatial correlations and inhomogeneity in both signal and noise. Simple averaging assumes the noise is uncorrelated and uniform. Furthermore, it assumes the signal is expressed identically at every voxel. This is a strong assumption. It is possible that functionally heterogeneous responses within the fROI cause half the fROI to activate and the other half to deactivate (this is not an uncommon architectural principle in the brain cf. surround inhibition in receptive fields or lateral interactions mediating 'winner-take-all' mechanisms). An average, in this case, will suppress signal-to-noise. We generally deal with this by taking the first eigenvariate of an fROI, which uses the temporal covariance of voxels in the fROI to find coherent spatial modes of activity (see spm_regions at http://www.fil.ion.ucl.ac.uk/spm). The principal eigenvariate is, like the average, simply a summary of the responses within an fROI. Unlike the average, it does not assume homogenous responses within the fROI. It should be noted that if one gets a significant result using the average, then it is valid. The point made here is that the average is not the most principled measure of a regional response.

Smoothing and the matched filter theorem address the equivalent issue in conventional voxel-based analysis. Under the assumption that the spatial dependencies of signal and noise are stationary, the most effective filter weights match the spatial scale of the signal. On the basis of optical imaging experiments, we know that haemodynamic response has a spatial scale around 4 mm. Usually, single-subject data are smoothed between 4 and 6 mm. This smoothing effectively transforms the data into an ensemble of fROI averages at every point in the brain. The implicit fROI corresponds to the smoothing kernel centred at each voxel. In this sense, a standard SPM analysis can be regarded as an analysis of all possible fROIs, whose spatial scale is physiologically informed and determined by smoothing. In the context of this discussion, the signal-to-noise of a single voxel (e.g., maximum of a localising contrast) can be made the same as the fROI average, provided the smoothing kernel and spatial scale of the fROI are comparable.

We have not made a distinction between analysing the average response of an fROI and averaging the estimated responses (i.e., contrast of parameter estimates, regression coefficients, etc.) it encompasses. This is because the two procedures give the same result. Averaging and response estimation with the general linear model are both linear operations and are commutative.[1] This means that the order of averaging and estimation is irrelevant. Furthermore, the same results will be obtained irrespective of whether the images are spatially normalised or not. This is because the fROI is subject to the same transformation as the underlying data. We mention this because one argument made in favour of fROI is that they can accommodate between-subject variations in functional anatomy, if the fROI response of each subject is taken to a second (between-subject) level for inference:

Inter-subject averaging

In multisubject studies, one has to account for between-subject variations in functional anatomy. The precise anatomical location of the FFA, for example, may vary over individuals. Conventionally, with voxel-based analyses, one assumes that most of this variability can be removed by spatial normalisation (Ashburner and Friston, 1999). Residual variability in functional anatomy, that persists after anatomical normalisation, is usually accommodated by further smoothing the data according to the matched filter theorem. This 'matches' the spatial dispersion induced by this residual variability (see above). In other words, smoothing is used to increase the probability that responses from different subjects overlap.

Clearly, this approach reduces spatial resolution. Furthermore, the degree of spatial dispersion of responses over subjects is unknown, and is an active area of research. This means the choice of smoothing is motivated rather anecdotally. An alternative to smoothing, in a standard anatomical space, is to pool data from different subjects using functional criteria. Specifically, the average responses of an fROI, defined for each subject, enter an analysis of variance (ANOVA), thereby discounting between-subject variations in anatomy and eschewing any need for spatial normalisation. This may be another reason why localisers have become so popular in the literature. However, even though this seems to be an important motivation for fROI, this motivation does not require fROIs to be defined from a localiser session. The same approach can be taken within a voxel-based analysis of single subject data (with or without spatial normalisation). In other words, it is very simple to perform an ANOVA on contrasts selected from the maxima of an orthogonal localising contrast in a subject-specific fashion. The advantage of this procedure over fROI averages is that the subject-specific maxima can be reported, providing a quantitative and useful characterisation of inter-subject variability in functional anatomy. Furthermore, this avoids defining fROI using ad hoc threshold criteria and avoids the assumption of functional homogeneity with the fROI. Having said this, one could argue that any attempt to define an irregular cluster of activated voxels with a few measures (e.g., size, Talairach coordinates of maximum response) is not fully adequate for informed meta-analyses.

Summary

The advantages of fROI responses include:

- They provide a simple way of summarising functional anatomy with a small number of well-defined areas that enables

[1] This point disregards any estimation of non-sphericity, which involves second-order or non-linear operations.

colloquial exchange and clarity, when addressing response properties.
- They enable a careful and comprehensive assessment of functional selectivity (of the fROI).
- They enforce reproducibility and provide a rigorous way of accumulating evidence across studies.

The disadvantages are that:

- They preclude inferences about function specialisation because of their inherent anatomical bias (i.e., failure to characterise anatomically distributed responses and, implicitly, their anatomical specificity).
- They may provide unnatural constraints on functional anatomy because they may have no structurally invariant properties (i.e., the voxels constituting the FFA under passive viewing may change under a familiarity judgement task).
- Their definition is sometimes ad hoc, both in terms of the paradigm used in their definition and the statistical criteria determining their extent. One practical caveat is the subjective component involved in specifying an fROI. Specifically, when considering the variability in anatomical location, the definition of an fROI and its borders might be observer-dependent.
- fROI averages are a poor surrogate for mass-univariate (e.g., voxel-based) or full multivariate characterisations (e.g., eigen- or canonical-variate analyses) of responses within the fROI. Eigenvariates are an important alternative to averaging and are used extensively in studies of effective connectivity because they allow for functionally heterogeneous but statistically dependent responses over the fROI.
- Their anatomy is difficult to report simply and quantitatively (i.e., for meta-analysis). This is because a meaningful meta-analysis would require a list of all the voxels comprising the fROI.
- There are no principled anatomical constraints on their inter-subject variability.

In summary, there are good reasons to use fROI or maxima to constrain statistical searches for the effects of other factors. However, the analysis of fROI responses per se has many disadvantages and is probably best motivated when other summaries of regional responses, such as eigenvariates, are inappropriate. Although fROI are sufficient to study functional selectivity, they cannot be used to comment on functional specialisation and therefore have a limited role in characterising functional anatomy. We now turn to the second question posed above and ask whether it is necessary to perform the functional localiser separately from the main experiment.

Why use separate functional localisers?

As noted above, any orthogonal contrast can be used as a functional localiser that is embedded in the main experiment. So why acquire data for a localising contrast outside the main experiment? We start by considering the advantages of factorial designs over localiser designs:

- First, localisers introduce inevitable confounds of both time and order. If the contrast of interest shows an effect of time (e.g., reduced activation in some areas and increased activation in others, due to perceptual learning), the localiser will be inappropriate because the activation pattern will have changed. This is particularly important in studies of visual categorisation, where perceptual learning may suppress activation in lower regions and increase them at higher levels (e.g., Dolan et al., 1997). More generally, one cannot look for interactions between the localiser factor and other factors because the localiser and order factors are confounded. For example, differences in response between the localiser and main experiment can be confounded by subject movement between sessions, differences in the cognitive or physiological status of the subject, difference in acquisition parameters such as temperature, etc.
- If the main experiment comprises n factors, the use of a localiser that preserves balance requires the localiser to be n-factorial. If it is not, the design is unbalanced, precluding a full analysis of the interactions (see lower panel of Fig. 1). In this example, a balanced ($2 \times 2 \times 2$) design has been replaced with a (2) localiser and a (2×2) main experiment. The ensuing loss of balance precludes the analysis of three-way interactions because a factor is missing from the localiser. The inability to test for interactions is important because it prevents inferences about functional specialisation or category specificity. For example, one can never say that a face-selective region does not respond to houses because this would be accepting the null hypothesis. However, within a balanced (2×2) design, one could use the "face vs. non-face $\times$ house vs. non-house" interaction to say that some object-responsive regions respond significantly more to faces than to houses. This would not be an option with a single face vs. house localiser because of the implicit loss of balance.
- The localiser and main experiments are often different in many aspects: scanning parameters (e.g., number of scans), design (e.g., blocked vs. event-related), task (e.g., passive viewing vs. one back) or stimuli used (e.g., expanding circles vs. moving dots). This means that the precision with which localising and experimental effects are estimated can differ profoundly. This can have a number of detrimental consequences. For example, a quick functional localiser may fail to disclose significant responses because of low sensitivity. The effects of the experimental factors, in these missed areas, may have been extremely significant. However, they precluded from analysis by the localiser, leading to biased reporting of the results.
- Factorial designs are more efficient because several orthogonal effects can be estimated using the greatest degrees of freedom. Put simply, factorial designs allow one to use the same degrees of freedom to make several inferences with no loss of statistical efficiency. Splitting the design into two sessions reduces the degrees of freedom for variance component estimation and reduces sensitivity to the effects in each session. For studies in which the localising contrast is replicated in the main experiment, it is more powerful to combine the two sessions into one long session and use a contrast testing for activation in the first half of trials as a localiser for responses in the other half. This is because the statistical model can assume that the error variances are the same for both halves and can estimate them more precisely than for the replication or split model.
- Finally, if the experiment conforms to a multifactorial design, a separate localiser is unnecessary and represents a waste of resources and unnecessary subject or patient discomfort.

So what are the potential advantages of using a localiser? We can think of the following:

- To avoid confounds that arise from interspersing the localiser contrast within the main experiment. In experiments that are perceptually or cognitively demanding, it is often important to keep the design simple to optimise performance. It may be the case that embedding a localiser in the main experiment will change cognitive set, induce task-switching costs or lead to priming of certain stimuli. There are several psychological constraints on experimental design that may be better accommodated by separate localisers. A simple example would be the use of a high-contrast stimulus to localise the processing of low-contrast stimuli. If the high-contrast stimuli are presented sporadically, they may change the context in which the low-contrast stimuli are perceived (e.g., one implicitly creates an oddball paradigm where low-contrast stimuli become standards). Separate localisers may be essential in paradigms that involve a training phase, followed by a test or probe phase. These phases cannot be inter-mixed because they entail an inherent order. In short, a localiser may be mandatory if presenting localising and main factors at the same time changes the nature of the processing under investigation.
- There may be designs that cannot be balanced. For example, imagine that the question is whether face-selective regions are also sensitive to emotional expression. The experimenter might define an "expression" factor with two levels: happy and sad. To ensure that subjects attend to the facial expression, the experimenter asks them to make an explicit expression judgement. To define "face-responsive", the experimenter has another two-level factor of faces vs. houses. Since houses do not have platonic expressions, a balanced (2×2) factorial design cannot be formed. However, even if the experimenter sticks with only three conditions (happy faces, sad faces and houses), the houses cannot be presented with the faces because subjects cannot perform an expression judgement on houses. Therefore, the experimenter might consider testing faces vs. houses in a separate localiser, using a different task (e.g., a one-back task). Another example is attentional modulation of sensory evoked responses, where a cell with attention to no stimulus is difficult to imagine.

These examples reflect experimental design issues and the problem of balance. One can imagine potential solutions, such as blocking faces and houses in the above example, and changing the task between blocks. If such solutions are inadequate (e.g., owing to task-switching or attentional confounds), and a separate localiser is performed, our arguments suggest that the experimenter needs to beware of some issues. First, they are assuming that there are no significant time or order confounds. Second, they are using their stimuli inefficiently. In the face example, faces are presented twice, once in the localiser and once in the main experiment. However, only half the face-selective responses are used to make an inference about emotional selectivity. Had all the stimuli been presented in a conventional manner, the same face-trials could have been used to test for emotional effects (happy vs. sad) and to provide the localising contrast (faces vs. non-faces). Third, the experimenter is making the important assumption that the different tasks in the localiser and main experiment do not interact with face-effects. This may not seem a big assumption to some researchers, who tend to view visual–object processing as "bottom-up" or "modular" (i.e., impenetrable by cognition). Such studies are concerned mainly with the stimulus properties (e.g., in defining fROIs like the "Lateral Occipital Complex" and its role in object processing, Malach et al., 1995). However, as we have discussed above, there is evidence that task factors can have important effects on responses in occipito-temporal cortex (Friston et al., 1996; Henson et al., 2002). This is particularly relevant to haemodynamic measures, which integrate over several seconds of synaptic activity and are likely to aggregate exogenous and endogenous processes (e.g., both "early", predominantly stimulus-driven and "late", predominantly task-related components). The balance of advantages and disadvantages would seem to suggest that functional localisers should be avoided if the question can be addressed using a factorial design.

Summary

Before turning to the case-studies, it is worth noting a few positive developments that are associated with the use of functional localisers. These developments can be viewed as going beyond simple structure–function relationships. We have already discussed the notion of pooling over subjects using functional, as opposed to anatomical criteria. Although this does not necessarily require separate localisers or fROIs, it is an important development and a challenge to the focus on anatomy as the exclusive reference for function. For example, studies of ocular dominance columns in V1 would not get very far using conventional inter-subject averaging procedures. However, these studies would be feasible if the voxels showing monocular bias were selected on a subject-by-subject basis. There have been parallel developments in analyses of functional integration (with dynamic causal modelling and structural equation modelling) where interacting regions are defined, not by their anatomical position but in terms of regions expressing the greatest functional response. This trend speaks to interesting notions; like spatially normalising with respect to a canonical localising contrast image, as opposed to a canonical anatomical template.

Another compelling trend that attends the use of localisers is a progression of questions about where a response is expressed to how functionally defined systems respond. Rather than asking where in the brain an effect is expressed, many visual scientists would ask whether (or how) an effect is expressed in a certain visual cortical region. For example, if one was interested in the role of early visual areas in perceptual awareness in a masking paradigm, it may be perfectly tenable to constrain the analysis to V1–V3 or even to the voxels within these areas that represent the stimulus retinotopically. An important advantage of these approaches is that one can increase spatial resolution and/or signal-to-noise by focussing on a specific brain region during data acquisition (e.g., using a surface coil). Again, note that these arguments do not rest on fROI, or indeed functional localisers. However, they are easily articulated in this context.

To conclude this section, we have seen that functional localisers are used to generate constraints on searches for the effects of the main factors. These constraints range from restrictions on voxel-based searches through averaging the response of an fROI. However, in many cases, a study can be designed to comprise orthogonal contrasts that can be used as mutual constraints in searching for regional effects. This means that balanced factorial designs circumvent the disadvantages of functional localisers, rendering them unnecessary. In light of this, it seems inappropriate to regard functional localisers as standard practice. In the next

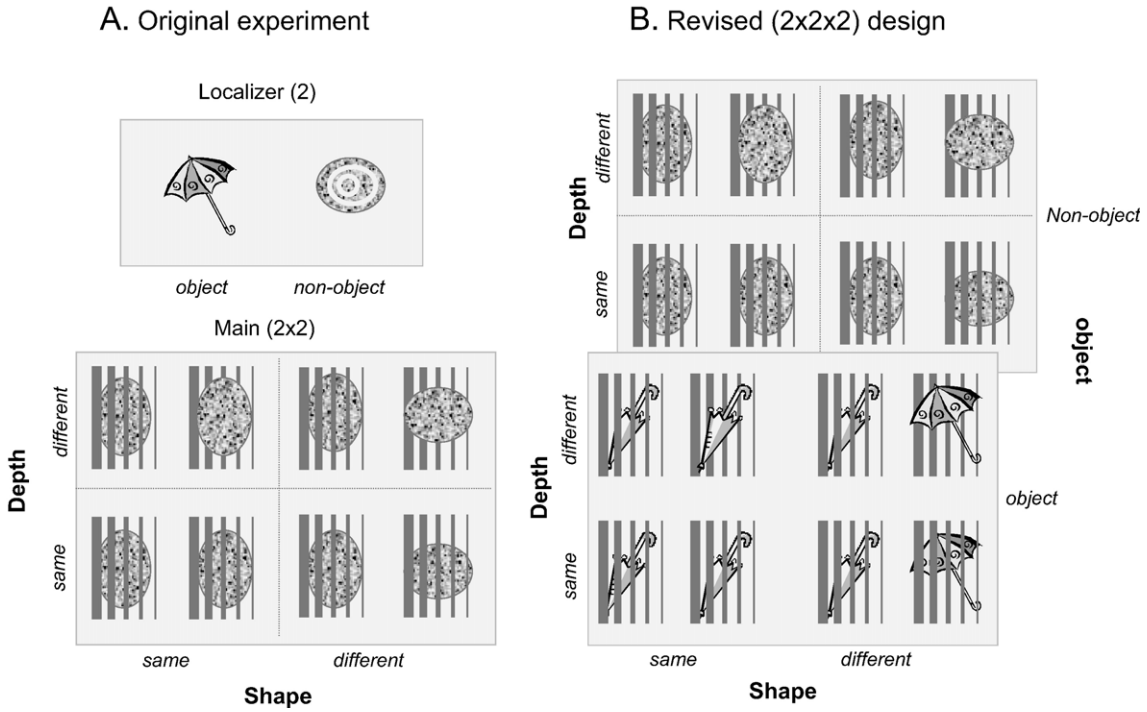

Fig. 2. A schematic illustration of the relationship between the original (left) and revised (right) designs described in the main test for the study by Kourtzi and Kanwisher (2001). The graphics are based loosely on the original paper but are used iconically in this paper. We have simplified the presentation of the design to make our points more clearly.

section, we take two recent studies and illustrate the above points in a practical setting. We want to stress that these two case-studies are used as a vehicle to make our points and do not detract from the original results reported by the authors or their significance. Furthermore, we appreciate that the original authors had specific and interesting questions in mind that our didactic deconstructions do not address. Our aim is to demonstrate how adjustments to the experimental designs enable additional questions that are precluded by functional localisers.

Case-studies

Functional regions of interest in object-selective regions

The first example comes from Kourtzi and Kanwisher (2001). In this study, subjects were shown objects (line drawings) and scrambled objects (arranged in circles) to define an fROI, called here the lateral occipital complex (LOC). The authors report the results of their main experiment purely in terms of the average response over all voxels within the LOC. The main experiment was an elegant two-by-two multifactorial design using the "adaptation" paradigm (in fact there were two experiments but we will consider only one here). The two main factors were shape with two levels (same vs. different) and depth (same vs. different); the depth manipulation entailed a change in local contours, which were different depending on whether the object appeared in front of or behind a grid (i.e., the depth manipulation was mediated by occlusion; Fig. 2A). Although the authors interpreted their results in terms of adaptation, for simplicity, we will regard haemodynamic responses as simple activations to a change in an attribute. Their key observation was that the LOC showed a main effect of shape-change but did not show a main effect of depth-change (i.e., contour-change). This was an interesting and well-received observation.

So what have we learned? This study demonstrates that the average response of object-selective voxels is sensitive to changes in shape but not in depth (i.e., occlusion or local contours).[2] What we do not know is:

- Where the main effects of shape-changes are expressed within the LOC.
- Whether a main effect of depth occurred within (subpartitions of) the LOC.
- Whether there was an interaction between shape and object: i.e., do shape-related responses depend on the stimulus being a recognisable object vs. non-object?
- Whether there was an interaction between depth and object: i.e., do depth-related responses depend on the stimulus being a recognisable object vs. non-object?
- Whether there was an interaction between depth and shape (this omission is not a reflection of original experimental design, the authors simply did not report it).
- Whether effects of depth (occlusion) or shape-changes occurred outside the LOC.

The first two limitations reflect the fact that an fROI average was analysed. The remaining questions are precluded by the localiser design. We will deal with these issues using a revision of the analysis

[2] Clearly one cannot infer that the LOC does not show depth or contour-selectivity (i.e., accept the null). One can only say there was a failure to elicit it. But this is another issue.

and experimental design respectively. The reason that localising information about the main effects is not available is that the effects of the main experiment were only evaluated for the averaged response of LOC. There is nothing wrong with this. However, we have learned nothing about the functional specialisation of shape or contour processing per se and how it may be segregated and integrated within the ventral-processing stream. It would have been possible to address this using the following analysis.

A revised analysis

The effect of visually evoked responses in the main experiment could have been used to identify maxima for further analysis. This approach invokes a localising contrast, testing for the differences between visual stimulation and inter-trial periods of no stimulation, to identify visually responsive areas. Note that this definition does not identify voxels that show larger responses to object compared to non-objects—hence, it does not identify the LOC; but it reduces the search volume. This localiser contrast is orthogonal to the main effects of shape- and depth-change and their interaction. Therefore, the two main effects and interaction can be tested in a constrained and sensitive fashion at the peaks of visual responses. One would then have presumably seen that early extrastriate cortices were more sensitive to contours (i.e., depth) and that higher visual areas preferred shapes (e.g., Murray et al., 2002). It would have also been interesting to see where changes in shape interacted with changes in contour, particularly at the maxima of the main effects of shape- and contour (i.e., depth)-change, respectively. These interactions can usually be interpreted in terms of an integration of neuronal computations (i.e., regions responsible for integrating contour information into the representations of shape). This sort of analysis presents a very different perspective on the data than that afforded by the fROI-based analysis. Note that, in the revised analysis, the separate functional localiser is completely redundant. Inferences about the functional anatomy of shape and contour processing pertain to the systems engaged by shape and contour processing, not to the processing of the unrelated objects used in the localiser.

One potential problem with this revised analysis is that the contrast of visual-evoked responses vs. baseline does not isolate object-responsive areas (such as the LOC, which is traditionally defined by comparing objects vs. textures, Malach et al., 1995). Regions activated by objects relative to inter-stimulus baseline are likely to include early visual regions that are sensitive to any transient change in luminance and not specifically the properties of objects. A better approach would involve a revised design.

A revised design

Let us assume that the authors wanted to understand object-selective processing in terms of its dependence on dynamic form (i.e., changes in shape and contour). One approach to this would be to integrate the localising and main experiments to create a conventional design with three factors; stimulus category (objects vs. non-objects), shape-change (same vs. different) and depth-change (same vs. different), see Fig. 2B. The localiser contrast, in this instance, would be the main effect of object and could be used to constrain the analysis of the shape- and contour-change effects.[3] Here, the localising contrast averages over all other factors and simply compares compound responses to paired stimuli (object vs. non-objects), irrespective of whether their shape- or depth-changed or not. Although addressing similar questions as the original design, this balanced factorial design properly controls the context in which objects are presented. In other words, object-selective regions are defined, operationally, in the context of changes in their shape and contours. The full design, in which the localiser factor is absorbed into the main experiment, has a number of advantages. For example, one can look at two- and three-way interactions that were precluded with the localiser design. The key interactions, from the point of view of the authors, are the two-way interactions involving stimulus category. For example, a significant object *x* shape-change interaction suggests that object-selective responses are sensitive to changes in object shape. This inference has a much greater focus than afforded by the original design, where the effect of shape- and contour-change could not be tested in relation to the stimulus category that was changing (object vs. non-object).

Summary

In summary, if the objective was to characterise shape and contour processing in the ventral stream, then the functional localiser was unnecessary. Indeed, using the localiser enforces a biased account of the underlying functional anatomy. Conversely, if the aim was to characterise the relative importance of shape and contour information in explaining object-selective responses, the original two-factor design could have been crossed with the localiser (object vs. non-object) factor to create a fully balanced three-factor design. By failing to integrate the localiser factor into the main experiment, certain cells are omitted and key interactions defining the context-sensitivity of object-selective responses are precluded. Fig. 2 provides a schematic illustration of the relationship between the original and revised designs. We now consider another functional localiser experiment, in this case concerning visual attention.

Functional localisers and visual attention shifts

Slotnick et al. (2003) report a clever experiment showing that spatial attention can facilitate or inhibit visually evoked responses in a retinotopically specific fashion (Fig. 3A). They presented compound stimuli with sparse flickering checkerboards at the centre, middle and periphery of the visual field. In the main experiment, these stimuli were presented continuously while subjects were cued to maintain attention or shift it to specific [middle] targets on opposite sides of the visual field. Two functional localisers were used, one for labelling early retinotopic areas and the other to identify regions responding specifically to the inner (i.e., centre), middle and outer (i.e., peripheral) stimuli by presenting them separately in blocks. The results of the second localiser fROI (stimulus vs. no stimuli) were used to constrain the analysis of responses evoked by shifts of spatial attention. Specifically, shifts to the contralateral field were compared with shifts to the ipsilateral field and vice versa.

The investigators showed that, within the fROIs, there was a significant difference between shifting to the contralateral field relative to the ipsilateral field. These differences varied in polarity and distribution over the fROI, from area to area, confirming the author's predictions. As expected, ipsilateral to contralateral shifts

[3] Note that the non-objects (scrambled blobs) in our revised design are not exactly the same as those used in the localiser of the original design (scrambled objects arranged in concentric blobs). However, our purpose is not to replicate Kourtzi and Kanwisher's comparisons precisely, but rather to illustrate the form that a more general, factorial design would take.

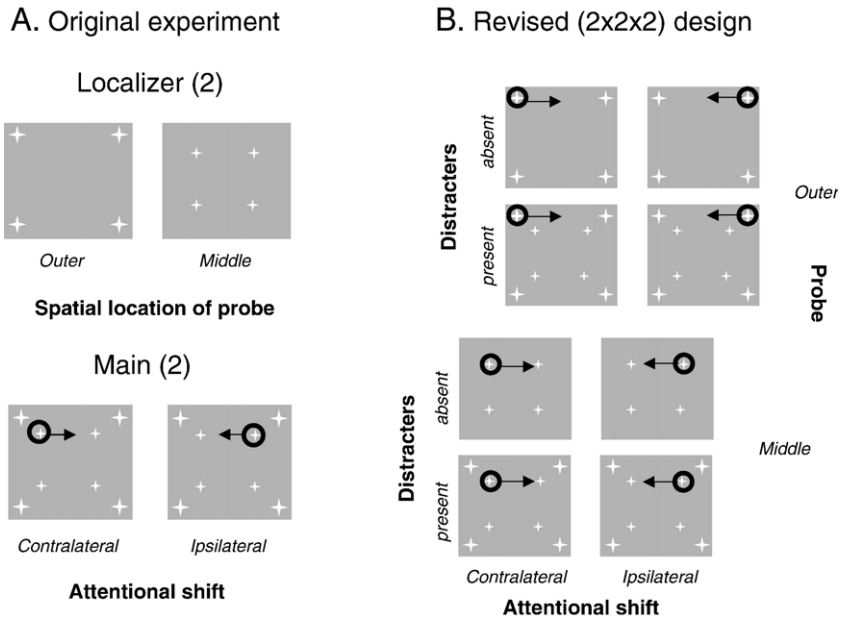

Fig. 3. A schematic illustration of the relationship between the original (left) and revised (right) designs described in the main test for the study by Slotnick et al. (2003). The graphics are based loosely on the original paper but are used iconically in this paper. We have simplified the presentation of the design to make our points more clearly.

of attention increased activity, and contralateral to ipsilateral shifts of attention decreased activity, in extrastriate representations of the upper–middle probe. Attentional facilitation extended to other stimulus probe representations in ventral cortex. Inhibitory effects were also evident. Consistent with previous findings, an inhibitory pattern was seen in the outer probe representation in ventral visual area V1v. In addition, attentional inhibition dominated many of the lower visual field probe representations in dorsal visual areas, as indicated by the activity profiles in V1d, V2d and V3.

What have we learned here? It has been shown that regions that respond to stimuli, when presented alone, show a main effect of attentional shift (contralateral vs. ipsilateral). What we do not know is where the main effect of attention was expressed (because the analysis was constrained by the a priori fROI). Furthermore, because the design was not balanced, we do not know whether the effects of attention depended on the non-attended stimuli or would have been expressed in their absence (note that the original design did not need to be balanced because this was not the authors question).

A revised design

These issues can be addressed in a revised design that brings the localiser and main experiment into the same balanced design. For simplicity, we will pretend the original experiment considered just two probes, i.e., outer (periphery) and middle (Fig. 3A). The revised design here augments the main experiment with two additional factors (Fig. 3B), one is distracter (present vs. absent) and the second is attended probe (middle vs. outer). In this design, the main experiment comprises alternating blocks of continuous stimuli during which ipsilateral and contralateral attention shifts occur.

This balanced design has all the information within it to localise retinotopically specific responses to the probes and to examine the main effects of attentional shift and how these effects depend on location (middle vs. outer) and the presence of distracters. For example, the simple main effect of probe, under the absence of distracters, acts as a functional localiser of the middle and outer probe locations (using a suitable contrast and the reverse contrast, respectively). Note that this localisation, like the original experiment, uses attended visual stimuli. However, unlike the original experiment, the baseline for the localiser is an attended stimulus elsewhere. This ensures that non-spatial attentional effects do not confound the localising contrast. The maxima of these contrasts can now be used to constrain tests for the simple main effects of attentional shift while attending to the middle or outer probes, respectively. Moreover, in the revised design, it is possible to look, not only for a main effect of attention as in the original report, but now to look for the interaction between attention and the presence of distracters. This would enable one to partition attention-related responses into components that depended upon other stimuli in the visual field and those that did not. In summary, as in the previous example, a functional localiser is both unnecessary and precludes tests of interactions among various factors that establish the context-sensitivity of the inferred effects.

Conclusion

In conclusion, the apparent advantages of the fROI approach with localiser designs apply as well or better to localisers embedded in factorial designs. The latter have many advantages, principally the ability to look at both main effects and interactions with increased statistical efficiency. If the localiser factor is an integral part of the hypothesis, then when possible, it should enter the main experiment in a balanced way. The increase in sensitivity afforded by constrained searches for region-specific effects is an important consideration but in many instances does not call for separate functional localisers. These constraints are often implicit

in multifactorial designs by virtue of the orthogonality of contrasts. The most common localising contrast is simply a difference between activation conditions and a suitable baseline (or null event).

We appreciate that in some cases there are good reasons to use separate functional localisers outside the main experiment but would like to emphasise that one should be aware of potential dangers involved in using them.

We have also addressed the use of fROIs. Although useful as constraints on statistical search spaces, their role in summarising regional responses is less compelling. A problem with the fROI idea is that it may be self-perpetuating in that fROI studies address only the behaviour of the fROI and can never ask whether the fROI is, in itself, valid. We hope this critique will help when assessing the need for functional localisers in both experimental design and peer review.

Acknowledgments

The Wellcome Trust supported this work. We thank very much Marcia Bennett for preparing the manuscript. We would like to thank Cathy Price, John Haynes and Jon Driver for invaluable contributions. We would also like to thank Zoë Kourtzi and Scott Slotnick, for ensuring that their original findings were properly represented in this discussion.

Appendix A

Two verbatim examples of anonymous reviewers' comments that refer explicitly to the lack of functional localisers in submitted scientific reports.

"It will be helpful if the authors could demonstrate the relationship between their activation maps and the FFA by comparing them to more conventional localizers—e.g. faces vs. buildings."

"I am left wondering about the anatomical relationship between the fusiform region reported here, and the well-known fusiform face area (FFA). Why didn't the authors use an independent functional localiser for the FFA? An independent localiser would also have allowed the authors to conduct a robust and independent interaction test."

References

Ashburner, J., Friston, K.J., 1999. Nonlinear spatial normalisation using basis functions. Hum. Brain Mapp. 7, 254–266.

Brodmann, K., 1909. Vergleichende Lokalisationslehre der Großhirninde in ihren Prinzipien dargestelle auf Grund des Zelenbaues. Leipzig, Barth, pp. 8.

Dolan, R.J., Fink, G.R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R.S.J., Friston, K.J., 1997. How the brain learns to see objects and faces in an impoverished context. Nature 389, 596–598.

Friston, K.J., 2003. Learning and inference in the brain. Neural Netw. 16, 1325–1352.

Friston, K.J., Price, C.J., Fletcher, P., Moore, C., Frackowiak, R.S.J., Dolan, R.J., 1996. The trouble with cognitive subtraction. NeuroImage 4, 97–104.

Gorno-Tempini, M.L., Price, C.J., Josephs, O., Vandenberghe, R., Cappa, S.F., Kapur, N., Frackowiak, R.S.J., 1998. The neural systems sustaining face and proper-name processing. Brain 121, 2103–2118.

Henson, R.N.A, Shallice, T., Gorno-Tempini, M.-L., Dolan, R.J., 2002. Face repetition effects in implicit and explicit memory tests as measured by fMRI. Cereb. Cortex 12, 178–186.

Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialised for face perception. J. Neurosci. 17, 4302–4311.

Kötter, R., Wanke, E., 2005. Mapping brain without coordinates. Philos. Trans. R. Soc. Lond. B Biol. Sci. 360, 751–766.

Kourtzi, Z., Kanwisher, Z., 2001. Representation of perceived object shape by the human lateral occipital complex. Science 293, 1506–1509.

Liu, W.C., Schulder, M., Narra, V., Kalnin, A.J., Cathcart, C., Jacobs, A., Lange, G., Holodny, A.I., 2000. Functional magnetic resonance imaging aided radiation treatment planning. Med. Phys. 27, 1563–1572.

Malach, R., Repas, J.B., Benson, R.R., Kwong, K.K., Jiang, H., Kennedy, W.A., Ledden, P.J., Brady, T.J., Rosen, B.R., Tootell, R.B.H., 1995. Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. Proc. Natl. Acad. Sci. U. S. A. 92, 8135–8139.

McIntosh, A.R., 2000. Towards a network theory of cognition. Neural Netw. 13, 861–870.

Murray, S.O., Kersten, D., Olshausen, B.A., Schrater, P., Woods, D.L., 2002. Shape perception reduces activity in human primary visual cortex. Proc. Natl. Acad. Sci. U. S. A. 99, 15164–15169.

Slotnick, S.D., Schwarzbach, J., Yantis, S., 2003. Attentional inhibition of visual processing in human striate and extrastriate cortex. NeuroImage 19, 1602–1611.

Worsley, K.J., Marrett, S., Neelin, P., Vandal, A.C., Friston, K.J., Evans, A.C., 1996. A unified statistical approach or determining significant signals in images of cerebral activation. Hum. Brain Mapp. 4, 58–73.