ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 30 (2006) 1097 – 1099

Comments and Controversies

Commentary on: Divide and conquer; a defence of functional localisers

K.J. Friston* and R.N. Henson[1]

The Wellcome Department of Imaging Neuroscience, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK

Received 30 January 2006; accepted 8 February 2006

We would like to thank Rebecca and her colleagues (Saxe et al., this issue) for providing a comprehensive and engaging commentary on our target article (Friston et al. this issue). We hope this exchange will clarify the different perspectives adopted by people who do and do not use fROI. However, this debate will not be resolved here: it will only be resolved by looking at the practice of imaging neuroscientists in the years to come. We suspect that fROI will have disappeared by then, because the questions they address are fundamentally limited. fROI are already being subverted by the growing interest in high-resolution functional imaging and multivariate characterisations of fine-scale distributed responses. The current issue is more pragmatic: it is becoming more difficult for cognitive neuroscientists to publish imaging papers that involve extra-striate or inferotemporal areas, without conforming to fROI dogma (see Appendix A). Our hope is to reverse this trend.

Contentious issues sometimes arise from a misconception of the others position. It was useful to have the response of Saxe et al. because we realise now how proponents of fROI may miss the point of our critique. Our point was that, if one wishes to identify brain regions using functional criteria (i.e., a "localiser" contrast), then it is best to (1) embed the localiser within an explicit factorial design, in the same experimental session; (2) use the contrast to constrain the search for brain regions showing the effects of interest (i.e., orthogonal main effects or interactions), rather than to average data over all voxels identified by the localising contrast.

Thus, while we can imagine readers nodding thoughtfully during the first part of Saxe et al. (part 1A) – in which the authors describe a well-known issue in cognitive neuroscience that structure–function mappings may differ across individuals – this is not the point of contention. This inter-subject variability raises questions about the validity of matching brains purely on the basis of structure (e.g., by "normalising" to a template) in group analyses. This is independent of our critique of fROI. As we noted in our section on inter-subject averaging (Section 2.2.7, Friston et al., this issue):

"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) … 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."

In short, there is nothing to prevent one performing analyses in each subject's native space (i.e., without normalising). Our point was simply that such analyses should not be based on averages within fROIs or need a separate localiser session.

Having clarified the focus of the debate, we will deconstruct the key advantages of fROI as listed by Saxe et al. in their abstract. We then address an important misconception about factorial designs. The comments of Saxe et al. are in quotation marks. Appendix B details short responses to some specific points.

The key advantages

"The fROI method, which resembled long establish practice is visual neurophysiology, has methodological statistical and theoretical advantages."

This methodology was developed under the constraints of single-unit electrode recording. These constraints do not apply to imaging neuroscience. This is because we can measure evoked responses everywhere in the brain and do not have to specify where these measurements are taken from.

"Because functional properties are more consistently and robustly associated with fROIs than with locations in stereotactic space, functional hypotheses concerning fROIs are often the most straightforward to frame".

While it is true that fROI provide a straightforward solution, they only address a straightforward problem; the functional

* Corresponding author. Fax: +44 207 813 1445.
E-mail address: k.friston@fil.ion.ucl.ac.uk (K.J. Friston).
[1] MRC Cognition and Brain Sciences Unit, Cambridge CB2 2EF, UK.
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter
doi:10.1016/j.neuroimage.2006.02.007

selectivity of functionally selective voxels. The inherent tautology precludes any questions about structure–function relationships. The most important example is specificity: it is impossible to address functional segregation (i.e., the anatomical specificity of functionally selective responses), because the functionally selective responses found in a single fROI may be expressed in many other parts of the brain that were not examined. Saxe et al. note that a solution to the problem of fROI is to "combine fROI analyses with each other or with voxel-based whole-brain analysis." However, arguing for the complementary use of SPM and fROI is specious. This is because SPM (and related whole-brain analyses) is equivalent to performing all possible fROI analyses.

As noted above, structure–function relationships become an essential issue for fROI at the level of inter-subject variability, and as noted in both papers, pooling functionally defined selectivity profiles over subjects is one way to discount uninteresting subject differences in structure–function mapping. However, even in this context, dispensing with anatomy is not always appropriate. The anatomical deployment of functionally selective responses can provide important constraints on inter-subject variability (e.g., degenerate or many-to-one structure–function mappings; Price and Friston, 2002; Henson, 2005). For example, some subjects may activate one region, whereas other subjects activate an anatomically distinct region. This degeneracy would be revealed in conventional voxel-based analyses at the single-subject level but would be missed completely using fROI that, operationally, treat the two regions as the same. In short, fROI analyses are straightforward because they eschew deeper questions about structure–function relationships in the brain.

"Because hypotheses are tested in only a handful of fROIs, advanced specification of fROIs provides a massive increase in statistical power over whole brain analyses".

This is nonsense. Statistical power is determined by the search volume. The power of a whole-brain analysis can be rendered identical to fROI analyses; if the search volume is suitably constrained (e.g., using the fROI). Saxe et al. are confusing the use of fROIs with the well-established relationship between sensitivity and search volume. They note later "By contrast traditional whole-brain analyses produce an explosion of multiple comparisons requiring powerful corrections to control false-positives." This correction depends only on the volume of brain examined. It is perfectly valid to perform a search constrained to a small volume of interest, which would entail less severe corrections to P values. One can also search the whole of the remaining brain using more severe (i.e., appropriate) corrections. Whole-brain analyses enable both these extremes and intermediate searches. fROI do not. In short, one can enjoy all the advantages of a constrained search, afforded by fROI, in the context of a conventional whole-brain analysis.

"Some fROIs may serve as candidate distinct components of the mind/brain worth investigation as such."

We were not really sure what this meant. Perhaps they meant that if a particular fROI is reified sufficiently it becomes an interesting object of study. While Saxe et al. observe that fROIs do not need to be reified, this has occurred (see Appendix B). And even if reification is appropriate, it should not preclude studying the rest of the brain. In the final lines of Saxe et al., they focus on a question posed to us by a reviewer, "Why didn't the authors use an independent functional localiser for the FFA?" Saxe et al. state "we don't see an answer to this question in the commentary by Friston et al. and we still think that many studies will benefit from the use of a fROI." The answer was that the authors of the original paper were not interested in the FFA. Differences in the way we think about functional anatomy become practically important when reviewers start prescribing the research question, focus or analysis for their colleagues (see Appendix A). As peer reviewers, should we be this prescriptive? Or should we be more sensitive to the dangers of fundamentalism, be it fROI or SPM?

Factorial designs

Saxe et al. state that "We are ambivalent about factorial designs." This is significant because localisers rest upon an implicit factorial design (the demonstration of the main effect of one factor in voxels that express a significant effect of another). We wonder whether the ambivalence of Saxe et al. stems from a failure to fully understand the nature of treatment effects in factorial designs: Saxe et al. state "in a factorial design, though the test of an interaction will be independent from the ROI-definition, the test of the main effect will be biased, since the very same data used to find the region of interest is then used to estimate the magnitude of the main effect." This is a remarkable statement. First, only one of the main effects is the localising effect and clearly one would not use this effect to constrain its own search! The effects of interest comprise the other main effects and interactions. Our proposal, which is standard practice in many labs, is to test for orthogonal effects (i.e., the interesting manipulations that are combined factorially with the localising factor), at voxels that exhibit a localising response. Orthogonal effects are independent, up to second order statistics. This means the test for one main effect cannot bias the test for other main effects or interactions. This can be seen simply by noting that the sum of two independent numbers is independent of their difference, despite the fact they are mixtures of the same data. Second, it may be that Saxe et al. think that there is some advantage to replicating the localising effect with a separate localiser. There is not. Formally, localiser sessions correspond to a split-half procedure (e.g., split t test). It is well known (by the Neyman–Pearson lemma) that split-half procedures are less efficient than a single likelihood-ratio test (i.e., combining the localiser and main experimental in the same model).

Conclusion

Finally, we want to reiterate the fundamental importance of factorial designs. The use of separate localiser sessions embodies an implicit assumption that the functional selectivity of the fROI is context-independent. However, in many situations, selective responses are modulated by context (e.g., McIntosh, 2000; Mechelli et al., 2003). For example, a "standard" FFA-localiser that compares faces and objects in an N-back task may engage different functions and brain regions than those engaged by the task examined in the main experiment. Being unable to test for an interaction between stimulus and task factors means the main effect and interaction are confounded (i.e., one cannot partition the response into a face-selective component and its task-specific modulation). This issue becomes especially problematic when reviewers insist that authors add "standard" localisers that enforce an unbalanced design and this inherent confound.

There are clearly many issues to be resolved in the mapping of structure and function in the human brain and how this mapping varies from subject to subject. Saxe et al. provide a very nice treatment of this. However, the solution offered by fROI is superficial, in the sense it ignores structure–function relationships by focussing exclusively on function. Note that this is in contradistinction to retinotopic mapping that depends on the anatomical topography of functionally selective responses (see Appendix B). While fROI may remain the preferred practice for some investigators, they are not necessarily the most principled approach to functional anatomy.

Appendix A

Since writing the target article, one of us (KJF) had a paper rejected from PLoS-B: The verbatim comments of [just] one reviewer were (our italics).

"Analysis: ROIs. These data deserve to be analyzed using the ROI approach that is now standard in the field. Retinotopically defined early visual areas should be identified and MT should be delineated in an independent scan. *Response should be averaged within active portions of each area*."

Appendix B

This appendix lists some other statements by Saxe et al. and a brief comment.

"Whole brain analyses are particularly poor tools for establishing the absence of an effect."

One can never establish the absence of an effect with classical inference (i.e., accept the null hypothesis). One can only say that there was a failure to reject the null hypothesis (i.e., one was unable to find an effect). This applies to both fROI and whole-brain analyses.

"The use of fROIs is uncontroversial and indeed virtually required in any study of visual cortex involving retinotopic visual areas or MT."

Retinotopic mapping should not be confused with fROI. Retinotopic mapping entails a careful voxel-based analysis of the topography of functionally selective responses. It does not use fROI in the sense we have been discussing. There is a fundamental distinction between using phase-encode mapping to assign a regional response to V2 and using the average of all V2 voxels as the response per se.

"So the worry that the advantages of fROI analyses will lead researchers to focus exclusively on a single ROI seems unsubstantiated by the actual practice in the field."

While a matter of opinion, our perception is that an exclusive focus has occurred in the case of the "fusiform-face area" (FFA). When performing "standard" localiser contrasts of faces versus objects, researchers not only find a region within the mid-fusiform (FFA), but also regions in occipital (Gauthier et al., 2000) and superior temporal (Haxby et al., 1999) cortex, among others (as also the case in single-cell recordings from the nonhuman primate). But ask a nonexpert in this domain, and a typical answer will be "faces are processed in a part of the brain called the FFA". The FFA is often reified in this sense, at the expense of other face-selective regions. This focus may be confounded by the historical accident that many initial studies used a surface coil over the occipital lobe, rendering them less sensitive to anterior regions.

"The second basic concern about fROI analysis is that a focus on the response of a predefined region will lead researchers into the dangerous. . .assumption that the neurons that comprise the fROI are homogenous." There is no concern about assumptions. The point made in Friston et al. was that the fROI averaging procedure provides an unbiased estimate of the activation if, and only if, the response is homogenous.

"Other sophisticated designs use parametric gradations of a single factor. . ."

Of course, but there is no reason why such designs cannot be made factorial.

"Peak-smoothed averaging uses smooth data and takes the time course of the voxel showing maximum activity in the localiser."

Not quite. This time course is a weighted average of nearby voxels (c.f., fROI average) determined by the smoothing kernel.

"If the area is not homogenous, we need to treat the first eigenvariate measure with care"

While it is true that any summary measure needs to be treated with care, the reason to use an eigenvariate is precisely to deal with areas that are not homogenous.

"Finally, Friston et al. (in press) imply that there is no cost to using a full-factorial design . . . (or "no loss of statistical efficiency"). On the contrary, simply multiplying the number of conditions in the experiment may produce many conditions that are of no interest to the experimenter, and simultaneously decrease the number of observations that can be conducted for the critical conditions. The result is a loss of power where it counts, and therefore reduced statistical efficiency."

It could be argued that all cells of a factorial design are both interesting and necessary (see above). However, just absorbing the localiser cells into the main experiment increases the degrees of freedom and power: under a pooled variance assumption, the estimated variance (i.e., standard error) becomes more precise with more data, even if these data do not contain an effect of interest. In other words, pooling the data from a localiser and main experimental will always be more powerful than analysing them separately.

References

Gauthier, I., Skudlarski, P., Gore, J.C., Anderson, A.W., 2000. Expertise for cars and birds recruits brain areas involved in face recognition. Nat. Neurosci. 3, 191–197.

Haxby, J.V., Ungerleider, L.G., Clark, V.P., Schouten, J.L., Hoffman, E.A., Martin, A., 1999. The effect of face inversion on activity in human neural systems for face and object perception. Neuron 22, 189–199.

Henson, R.N.A., 2005. What can functional imaging tell the experimental psychologist? Q. J. Exp. Psychol., A 58, 193–233.

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

Mechelli, A., Price, C.J., Noppeney, U., Friston, K.J., 2003. A dynamic causal modelling study on category effects: bottom-up or top-down mediation? J. Cogn. Neurosci. 15, 925–934.

Price, C.J., Friston, K.J., 2002. Degeneracy and cognitive anatomy. Trends Cogn. Sci. 6, 416–421.