

Imaging cognitive anatomy

Karl J. Friston

One of the major challenges in imaging neuroscience is the integration of cognitive science with the empiricism of neurophysiology. The cognitive architectures and principles offered by cognitive science have been essential in shaping experimental design and image analysis strategies from the outset. Now some of the cognitive models and their assumptions (for example, cognitive subtraction) are being re-evaluated in the light of how the brain actually implements putative components and processes. In this review we will consider experimental designs that go beyond cognitive subtraction and also consider how functional imaging can be used to assess the context-sensitivity of cognitive processing (using conjunction analyses), and the integration of different processes (in terms of interactions, using factorial designs) and how both these themes can be developed in the context of parametric designs. These new approaches reflect an ongoing discourse between cognitive science and the emerging principles of functional anatomy.

Our interpretation of functional imaging studies relies upon an ensemble of models. Some are very explicit (for example, the general linear model used in statistical analysis) and some are implicit (for example, functional specialization when inferring the nature of activations¹ or pure insertion in cognitive subtraction²). These models embody the assumptions required to make sense of the data. Models relating to what is being measured (Box 1) and how these measurements are analysed (Box 2) are biophysical, anatomical, mathematical, statistical and, usually, explicit. In this review we consider some of the more conceptual and implicit models that are commonly employed in experimental design and their relationship to cognitive science.

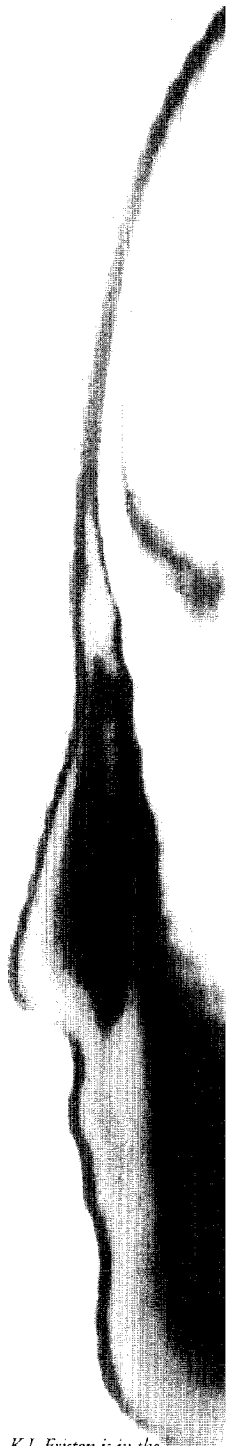
Functional specialization or integration?

Patterns of brain activity elicited by sensory, motor or cognitive tasks are usually understood by reference to 'functional specialization' or the 'functional integration' of specialized areas. Functional specialization refers to the expression of stereotyped patterns of neuronal activity in response to specific attributes of a stimulus, cognitive processing, or motor behaviour by specialized cortical areas, sub-areas or neuronal populations. The principle of functional specialization is now well established, particularly in visual neuroscience³. Functional integration refers to the interactions among specialized neuronal populations and how these interactions depend upon the sensorimotor or cognitive context. Functional specialization and integration are not exclusive; they are complementary; one only makes sense in the context of the other. From the perspective of neuroimaging, functional specialization calls for the identification of 'regionally specific effects' that can be attributed

to changing stimuli or task conditions. Functional integration, on the other hand, is usually assessed by examining the correlations between activity in different brain areas, or by trying to explain the activity in one area in relation to activities elsewhere⁴⁻⁶. These analyses are usually framed in the terms of 'effective connectivity' (the influence that one neuronal system exerts over another)⁶. There is a fundamental distinction between demonstrating effective connectivity (in relation to some model of neuronal interactions) and simply observing correlated activity: correlations can arise from many sources that do not reflect teleologically meaningful interactions (for example, stimulus-evoked transients in two neuronal populations that are not connected, or the modulation of two cortical areas by a common subcortical input). This review will focus primarily on functional specialization and concludes with an example that brings together integration and specialization in terms of the top-down modulation of stimulus-specific responses. To date, brain activation experiments have been predicated on a functional specialization view of the brain and there are three main types: categorical; factorial and parametric.

Cognitive subtraction

Categorical designs⁷⁻¹⁰, such as cognitive subtraction³, have been the mainstay of functional neuroimaging over the past decade. Cognitive subtraction involves elaborating two or more tasks that differ in a separable component. Ensuing differences in brain activity are then attributed to this component. For example, the difference between simply saying 'yes' when an object is seen, and naming that object, includes retrieval of the phonology of the object's name. Regionally-specific differences in brain activity that distinguish between



K.J. Friston is in the Wellcome Department of Cognitive Neurology, Institute of Neurology, 12 Queen Square, London, UK WC1N 3BG.

Tel: + 44 171 833 7456
 Fax: + 44 171 813 1445
 e-mail: k.friston@fil.ion.ucl.ac.uk

Box 1. What are we measuring?

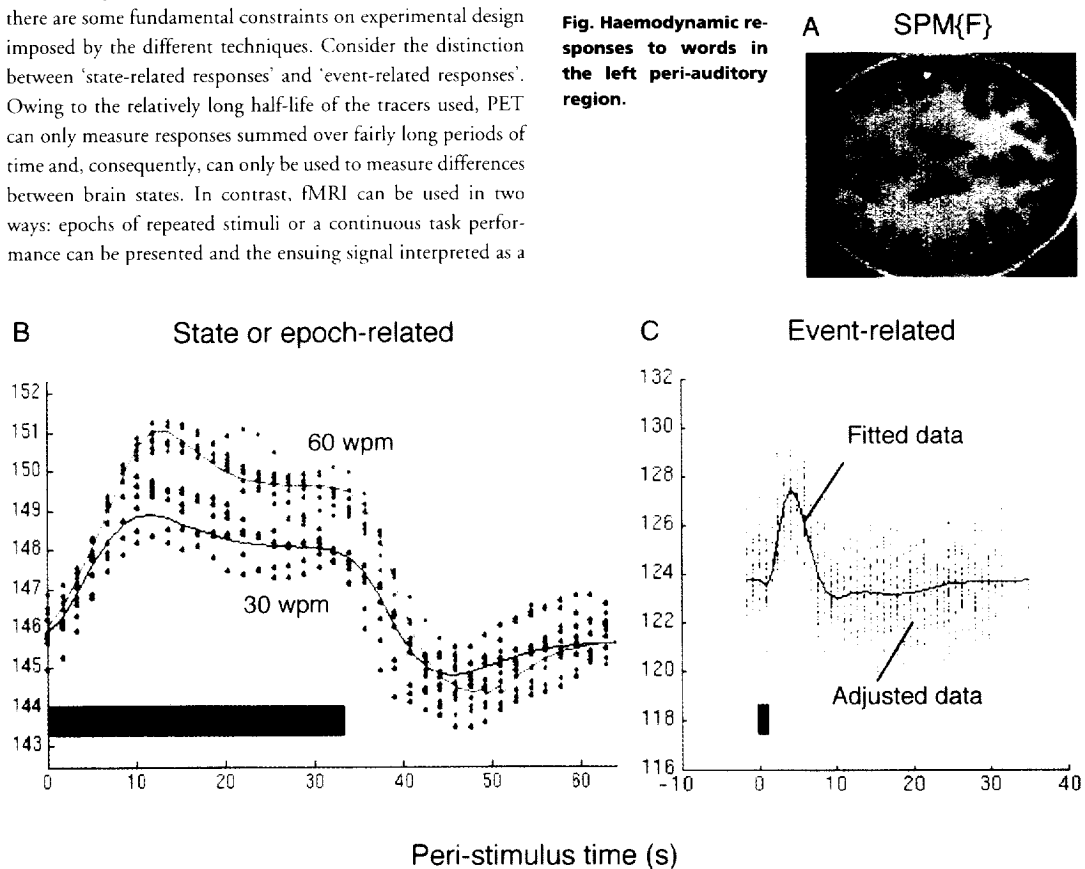
From the early 80s positron emission tomography (PET) dominated the field of functional neuroanatomy. However, in the past five years, functional magnetic resonance imaging (fMRI) has developed into an alternative and powerful technique. PET measures blood flow on a spatio-temporal scale of ~6 mm and 30 sec; fMRI is sensitive to the oxygenation of blood and has a spatio-temporal scale of ~1–3 mm and one or more seconds. The lower limits on the 'effective resolution' of fMRI are physiological and imposed by the spatio-temporal organization of evoked haemodynamic responses (2–5 mm and 5–8 sec). Local increases in neural activity cause both a relative deoxygenation of blood and an increase in perfusion that quickly reverses the deoxygenation, leading to an increase in oxygenation that endures for several seconds. The observed haemodynamic response can be thought of as a smoothed version of the underlying neural activity (see Figure). PET measures changes in blood flow or perfusion directly, in terms of the amount of radio-labelled water that accumulates locally, but this takes up to a minute or so.

Although there is a great overlap between PET and fMRI, there are some fundamental constraints on experimental design imposed by the different techniques. Consider the distinction between 'state-related responses' and 'event-related responses'. Owing to the relatively long half-life of the tracers used, PET can only measure responses summed over fairly long periods of time and, consequently, can only be used to measure differences between brain states. In contrast, fMRI can be used in two ways: epochs of repeated stimuli or a continuous task performance can be presented and the ensuing signal interpreted as a

brain-state dependent measure (see Fig. B); alternatively, event-related responses to a single stimulus can be measured by analogy to evoked potentials in electrophysiology (see Fig. C).

The data shown in the figure were acquired from a single subject using echo planar imaging (EPI) fMRI at a rate of one volume image every 1.7 s. In Fig. B, the subject listened to words in periods of 34 s at a variety of different frequencies. The fitted peri-auditory responses (lines) and adjusted data (dots) are shown for two rates (30 and 60 words per minute). The solid bar denotes the presentation of words. By removing confounds and specifying the appropriate design matrix (see Box 2) one can show that fMRI is exquisitely sensitive to single events. The data shown in Fig. C were acquired from the same subject whilst simply listening to single words presented every 34 sec. Event-related responses were modelled using a small set of temporal basis functions of the peri-stimulus time. The SPM{F} (statistical parametric map) reflecting the significance of these event-related responses (see Box 2) has been thresholded at $p = 0.001$ (uncorrected) and displayed on a T_1 -weighted structural MRI.

Fig. Haemodynamic responses to words in the left peri-auditory region.



these two tasks could, therefore, be implicated in phonological retrieval.

Although its simplicity is appealing, this approach embodies some strong assumptions about the way that the brain implements cognitive processes. A key assumption is 'pure insertion'. Pure insertion asserts that one can insert a new component into a task without affecting the implementation of pre-existing components. (For example, how do we know that object recognition, implicit in both the tasks above, is not itself affected by naming?) The fallibility

of this assumption has been acknowledged for decades, perhaps most explicitly by Sternberg's revision of Donder's subtractive method¹¹. The problem for subtraction is as follows: if one develops a task by adding a component, the new task comprises not only the previous components and the new component but the integration of the new and old components (for example, the integration of phonology and aspects of structural identification during object recognition). This integration or 'interaction' can itself be considered as a new component. Therefore, the difference

Box 2. How are the data analysed?

Overview

The Figure depicts the transformations that start with the imaging time-series and end with a statistical parametric map (SPM). SPMs can be thought of as 'X-rays' of the significance of an effect. Voxel-based analyses require the data to be in the same anatomical space: this is effected by realigning the data (in fMRI it is usually necessary to remove signal components that are correlated with displacements and remain after realignment^a). After realignment the images are subject to non-linear warping^b so that they match a template that already conforms to a standard space. After smoothing, the general linear model is employed to perform the appropriate univariate test at each and every voxel. The test statistics that ensue (usually t or F statistics) constitute the SPM. The final stage is to make statistical inferences on the basis of the SPM and characterize the responses observed using the fitted responses or parameter estimates. If one knows where to look beforehand, then this inference can be based on the value of the statistic without correction. If however an anatomical site cannot be predicted, *a priori*, then a correction for the multiple dependent comparisons performed has to be made. These corrections are usually made using distributional approximations from the theory of Gaussian fields^{c-f}.

The general linear model

The general linear model is an equation which expresses the observed response variable in terms of a linear combination of explanatory variables plus a well-behaved error term. The general linear model is variously known as 'analysis of covariance' or 'multiple regression analysis' and subsumes simpler variants, like the 't test' for a difference in means. The matrix that contains the explanatory variables (for example, designed effects or confounds) is called the design matrix. Each column of the design matrix corresponds to some effect one has built into the experiment or that may confound the results. The example below relates to an fMRI study of visual stimulation under four conditions. The effects on the response variable are modelled in terms of functions of the presence of these conditions

(smoothed with some appropriate haemodynamic response function) and constitute the first four columns of the design matrix. There then follows a series of terms that are designed to remove or model low-frequency variations in signal due to artifacts such as aliased biorhythms. The final column is whole brain activity. The relative contribution of each of these columns is assessed using standard least squares and inferences about these contributions are made using t statistics or F ratios, depending upon whether one is looking at a particular linear combination of these contributions, or all of them together. In fMRI the data are sometimes smoothed in time and one has to include the resulting serial correlations in the general linear model^d.

Statistical inference and the theory of Gaussian fields

Correcting the p values for the volume analysed is complicated as the data are not independent, owing to smoothness in the original data, and can be achieved by treating the SPMs as realizations of Gaussian random fields^e. The inferences obtained are based on the probability of obtaining c or more clusters of volume k or more voxels above a threshold of u in the volume analysed. This formulation has a number of special cases that give corrected p values pertaining to the number of clusters, the size of clusters or the heights of voxels within each cluster.

References

- a Friston, K.J. et al. (1996) Movement related effects in fMRI time series *Magn. Reson. Med.* 35, 346-355
- b Friston, K.J. et al. (1996) Spatial realignment and normalization of images *Hum. Brain Mapp.* 3, 165-189
- c Talairach, J. and Tournoux, P. (1988) *A Co-planar Stereotaxic Atlas of a Human Brain*, Thieme
- d Friston, K.J. et al. (1995) Statistical parametric maps in functional imaging: a general linear approach *Hum. Brain Mapp.* 2, 189-210
- e Worsley, K.J. (1994) Local maxima and the expected Euler characteristic of excursion sets of χ^2 , F and t fields *Adv. Appl. Prob.* 26, 13-42
- f Worsley, K.J. and Friston, K.J. (1995) Analysis of fMRI time-series revisited - again *NeuroImage* 2, 173-181

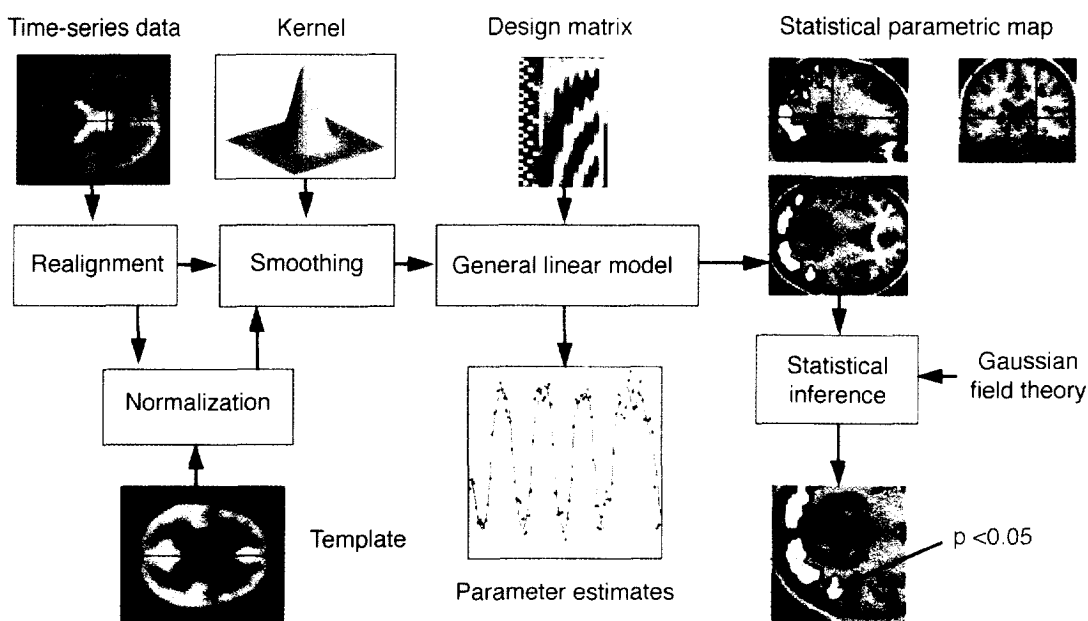
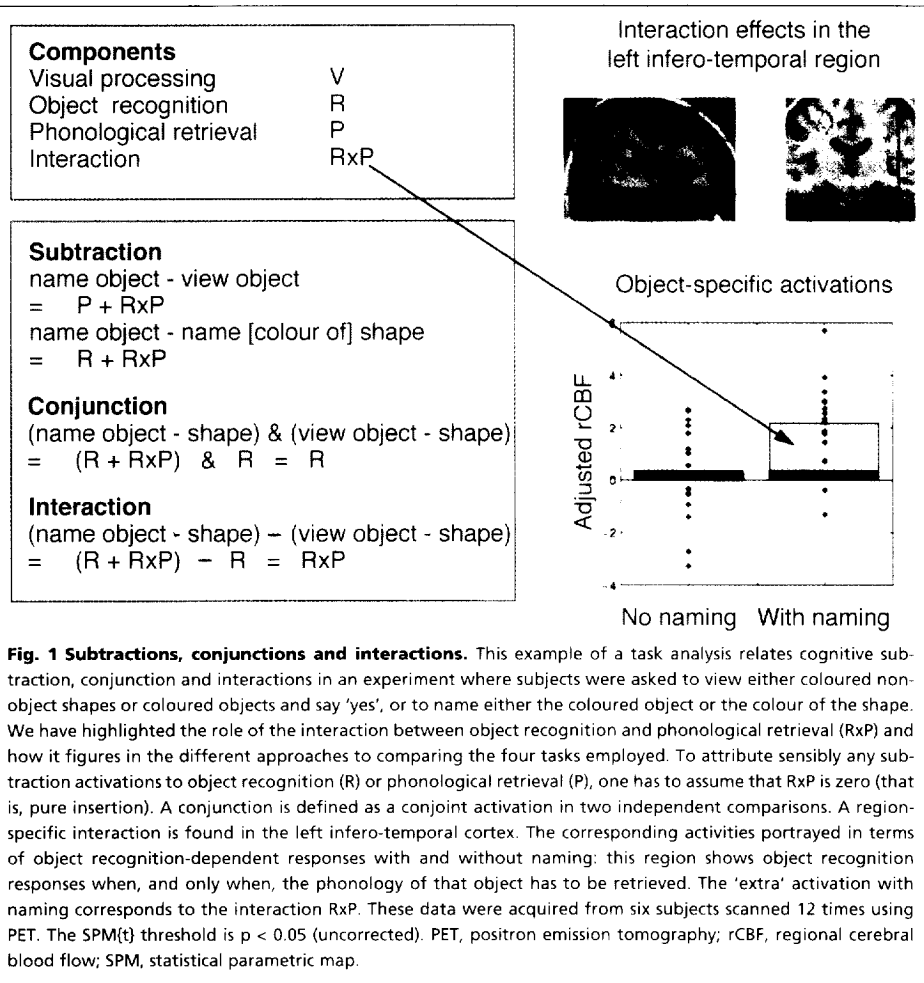
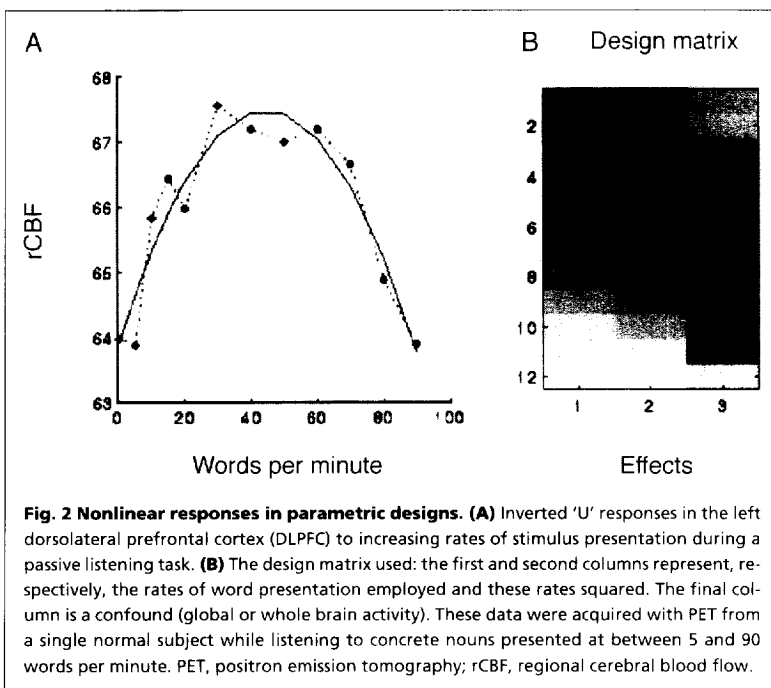


Fig. Data transformations.



between two tasks includes the new component and the interactions between the new component and those of the original task (see Fig. 1). Pure insertion requires that all these interaction terms are negligible but in many instances they obviously are not. A compelling example is modulation



of sensory processing by directed attention¹². The augmented responses to stimuli in specialized visual areas represent an interaction between attention, for a particular sensory attribute, and the processing of perceptual stimuli that have this attribute. In general, the context-sensitivity of cognitive and sensorimotor processing implies that the interaction between the component and its context (other components or processes) cannot be ignored. This renders cognitive subtraction a rather unsound conceptual basis for many lines of research and we will, therefore, consider two complementary approaches: cognitive conjunctions and factorial designs that eschew the assumption of pure insertion.

Cognitive conjunctions

Cognitive conjunctions can be thought of as an extension of the subtraction technique in the sense that they combine a series of subtractions. In subtraction we test a hypothesis pertaining to activation in one task relative to another. In cognitive conjunctions several hypotheses are tested, asking whether all the activations, in a series of task pairs, are jointly significant. Consider the problem of identifying regionally specific activations due to a particular cognitive

component (for example, object recognition). If one can identify a series of task pairs whose differences have only that particular component in common, then the region which activates in all the corresponding subtractions can be associated uniquely with the component in question. This is tenable even if interactions are prevalent because the interactions will (by design) be specific to each pair. In this way, cognitive conjunctions can be used to discount interactions and render the inference less sensitive to the context in which a particular cognitive component is expressed. To make this argument clear consider the experiment depicted in Fig. 1 (Ref. 13): subjects either viewed coloured shapes or coloured objects (and said 'yes') or they named coloured shapes or coloured objects (naming either the object or the colour). The differences between naming an object or the colour of a shape include object recognition and the interaction between recognition and explicit phonological retrieval. The difference between simply viewing an object and a coloured shape is object recognition (for simplicity we ignore interactions with visual processing). The common difference is object recognition and this conjunction discounts the interaction between recognition and phonological retrieval of the name of the recognized object. Regionally specific activations that are significant in both comparisons (and are not significantly different) will identify areas that are specialized for object recognition per se. In some situations the interaction or integration of the naming and recognition of an object may, itself, be of interest.

Factorial designs allow this interaction effect to be assessed directly.

Factorial designs

Factorial designs involve combining two or more factors within a task or tasks [the factors in the above example were objects (vs non-objects) and naming (vs saying 'yes')] and looking at the interaction between the different factors, or the effect of one factor on the responses to the other factor. The first, and perhaps the simplest, factorial design in neuroimaging involved an interaction between motor activation and time, that was interpreted in terms of physiological plasticity or adaptation¹⁴. The interaction here was simply an effect of time on activation due to motor performance. Generally, interactions can be thought of as a difference in activations brought about by another processing demand. In other words, in changing the context of a particular task one can modulate the activation and examine the interaction between the activation and the context employed. Dual task interference paradigms are a clear example of this approach¹⁵. A further example includes psychopharmacological activation studies wherein the effect of manipulating neurotransmitter systems in the brain is assessed in terms of responses elicited by cognitive challenge^{16,17}. The factorial nature of the above naming experiment can be seen by noting that object-specific responses are elicited (by asking subjects to view objects relative to meaningless shapes) with and without naming. This 'two by two' design allows one to look specifically at the interaction between phonological retrieval and object recognition. This analysis identifies, not regionally specific activations, but 'regionally specific interactions'. These interactions are evident in the left inferior temporal region (Fig. 1) and can be associated with the integration of phonology and object recognition. Alternatively, these regions can be thought of as being recognition-dependent responses that are realized in, and only in, the context of having to name the object seen.

Parametric designs

The essential difference between parametric designs and those described above is that cognitive components or sensorimotor attributes are treated as dimensions as opposed to categories. In this framework, one would expect to see systematic changes in elicited responses according to the strength of expression of a component or performance attribute¹⁸. An early example of this approach was the demonstration of increased activity in peri-auditory regions with increasing frequency of aurally presented words¹⁹, and there were some interesting differences between this response in the peri-auditory regions and Wernicke's area. In a sense, categorical designs are a limiting case of parametric designs, where only two extremes of processing are employed. Categorical approaches implicitly assume that the parametric

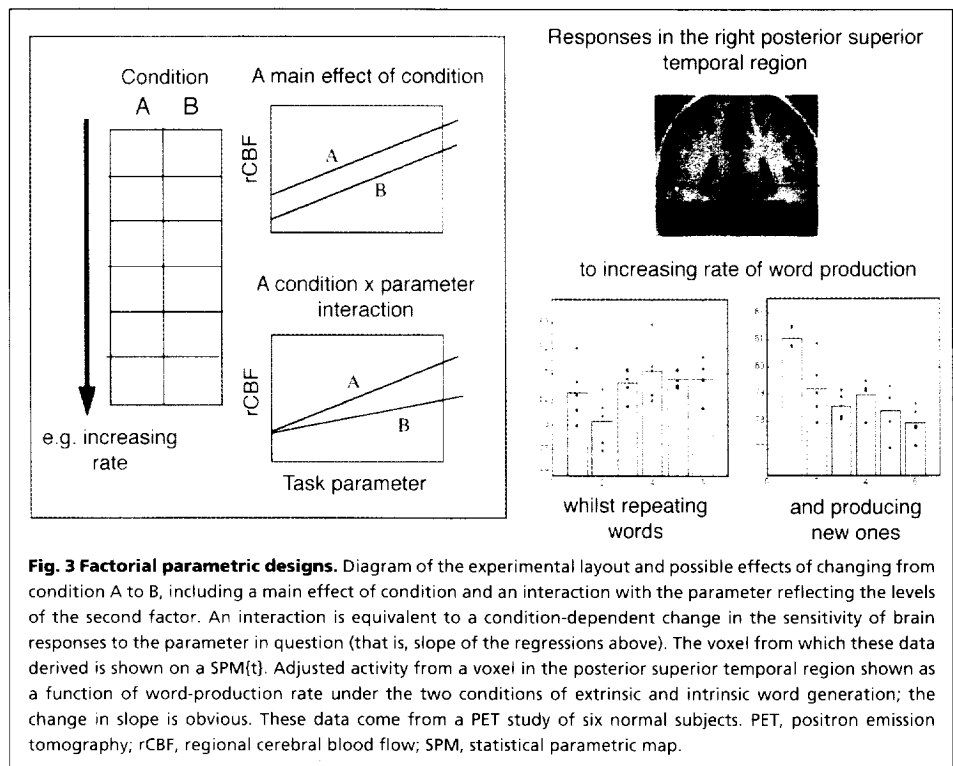


Fig. 3 Factorial parametric designs. Diagram of the experimental layout and possible effects of changing from condition A to B, including a main effect of condition and an interaction with the parameter reflecting the levels of the second factor. An interaction is equivalent to a condition-dependent change in the sensitivity of brain responses to the parameter in question (that is, slope of the regressions above). The voxel from which these data derived is shown on a SPM(t). Adjusted activity from a voxel in the posterior superior temporal region shown as a function of word-production rate under the two conditions of extrinsic and intrinsic word generation; the change in slope is obvious. These data come from a PET study of six normal subjects. PET, positron emission tomography; rCBF, regional cerebral blood flow; SPM, statistical parametric map.

response is at least monotonic. However, in some instances it may be highly non-linear²⁰; for example, an inverted U behaviour with increasing stimulus presentation rates (Fig. 2). The importance of this example for categorical designs is clear: evidence for functional specialization might be missed completely using just two rates (for example, a very low one and a very high one). Categorical approaches assume simple forms and dynamic ranges for the relationship between a parameter and response, whereas parametric designs embody no such assumptions and allow the nature of this relationship to be characterized empirically. There are many interesting aspects of parametric designs: for example, does the sensitivity of hippocampal responses to increasing word-list length, in a free-recall memory task, change with strategic differences in memory processing when the number of words exceeds the capacity of short-term systems²¹? In other words, is there a strategy-dependent change in slope, at some critical value, when we plot brain activity against the experimental parameter? How does one distinguish between the effects of increased 'effort' and processing effects specific to the parameter in question (for example, are prefrontal responses attributable to increased processing incurred by the number of distracters in a target detection task, or would these responses be seen with any task involving increases in effort?). One approach is to repeat the experiment using a different parameter, whilst matching the levels of 'effort', and looking for differences in the regression of prefrontal responses on the two parameters. This brings us back to interactions and factorial designs, but now in a parametric context. A simple example of these designs would be the examination of brain responses to increasing frequency of stimulus presentation under different (attentional) conditions²². A further example is shown in Fig. 3, where subjects were asked to produce words at different frequencies. In one context, the subjects

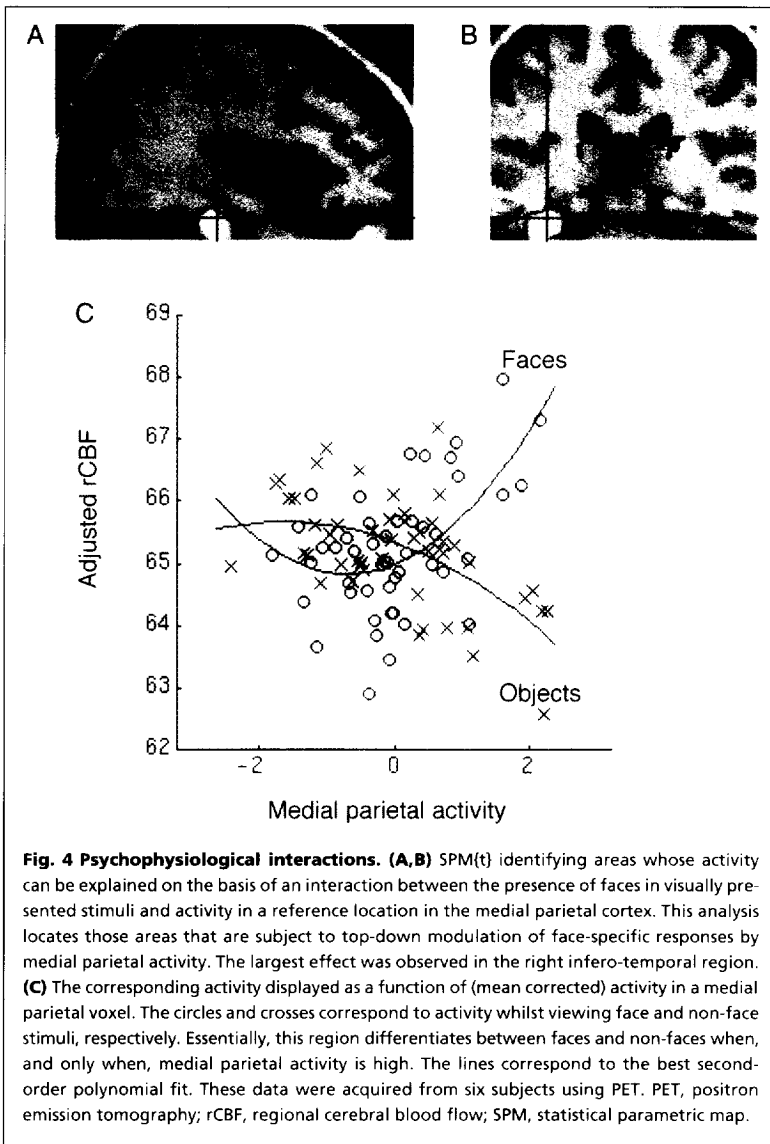


Fig. 4 Psychophysiological interactions. (A, B) SPM(t) identifying areas whose activity can be explained on the basis of an interaction between the presence of faces in visually presented stimuli and activity in a reference location in the medial parietal cortex. This analysis locates those areas that are subject to top-down modulation of face-specific responses by medial parietal activity. The largest effect was observed in the right infero-temporal region. (C) The corresponding activity displayed as a function of (mean corrected) activity in a medial parietal voxel. The circles and crosses correspond to activity whilst viewing face and non-face stimuli, respectively. Essentially, this region differentiates between faces and non-faces when, and only when, medial parietal activity is high. The lines correspond to the best second-order polynomial fit. These data were acquired from six subjects using PET. PET, positron emission tomography; rCBF, regional cerebral blood flow; SPM, statistical parametric map.

Outstanding questions

- What are we measuring? Perfusion (measured by PET) and haemodynamic responses (measured by fMRI)²³ reflect neural activity but the mechanisms mediating neural transients and the ensuing haemodynamics remain uncertain. Optical imaging, electrophysiology²⁴ and the use of biophysical models may help elucidate these mechanisms.
- What is the neural code and why is it important for neuroimaging? Functional imaging is not necessarily sensitive to the degree of synchronization among neuronal populations. If transactions among populations of neurons are mediated by temporal codes that rely on timing and synchronization then neuroimaging is blind to a whole dimension of neuronal discourse, an issue currently being addressed, using multi-unit electrode recordings and computational neurobiology²⁵.
- Is functional specialization a sufficient model?
- How do we measure functional integration? Functional integration is measured with effective connectivity; can this be measured properly? Most current approaches use simple linear models which have questionable usefulness or validity.
- Can cognitive science and neurophysiology be reconciled? Functional imaging conflates cognitive science and neurophysiology into one discipline. This could be problematic²⁶; can the cognitive architectures posited by cognitive science be endorsed by the empiricism of functional imaging and is that empiricism properly informed?

simply repeated a heard word and in the second they generated a word beginning with a heard letter. In this case the interaction corresponds to a differential sensitivity to word-production rate depending upon whether the words were extrinsically or intrinsically generated. A profound interaction is shown in the right posterior temporal region where, for extrinsically generated words, activity increases gently with word production rate, whereas for intrinsically generated words it decreases. The observation that activity systematically decreases with the increasing frequency of a neuronal event (in this case the intrinsic generation of a word) is of fundamental importance, and indicates true deactivation or reduction in activity elicited by each word.

To conclude we now look at factorial designs, of a parametric nature, from a rather different perspective allowing us to make some inferences about functional integration in the brain in terms of top-down modulatory interactions.

Psychophysiological interactions

The new concept of psychophysiological interactions can be thought of as analogous to psychopharmacological interactions: in psychopharmacology^{16,17} we are interested in the interaction between sensorimotor or cognitive evoked responses and some pharmacological or neurotransmitter manipulation; in psychophysiological interactions we are trying to explain the physiological response in one part of the brain in terms of an interaction between the presence of a sensorimotor or cognitive process and activity in another part of the brain. For example, by combining information about activity in the parietal region, mediating attention to a particular stimulus and information about the stimulus, can we identify regions that respond to that stimulus when, and only when, activity in the parietal region is high? If such an interaction exists, then one might infer that the parietal area is modulating responses to the stimulus for which the area is selective. This has clear ramifications in terms of the top-down modulation of specialized cortical areas by higher brain regions. This is an interesting analysis for two reasons: (1) the explanatory variables used to predict activity (that is, the response variable) in any brain region comprise a standard predictor variable based on the experimental design (for example, the presence or absence of a particular stimulus attribute) and a response variable from another part of the brain; (2) it uses techniques normally used to make inferences about functional specialization, instead, to make inferences about functional integration (in this instance, effective connectivity of a modulatory sort). Figure 4 illustrates a specific application of this approach: subjects were asked to view (degraded) faces and non-face (object) controls. The interaction between activity in the medial parietal region and the presence of faces was most significantly expressed in the right infero-temporal region. Changes in medial parietal activity were introduced experimentally by pre-exposure to the stimuli before some scans. These results can be interpreted as a priming-dependent instantiation of attentional, memory or learning differences in face-specific responses, in inferotemporal regions that are mediated by interactions with the medial parietal cortex. Note that we could have modelled the priming effect

explicitly in our design matrix (see Box 2) but chose to substitute medial parietal activity in its place, enabling us to make a more mechanistic inference: namely, not only do infero-temporal responses show a priming-dependent effect but this effect is mediated by modulatory influences from a higher (parietal) area.

Conclusion

It is perhaps appropriate that this review concludes with an example that depends explicitly on combining cognitive processing and physiological brain measurements when trying to understand how the brain works. In conclusion, there are many implicit assumptions that underlie experimental design in neuroimaging and ongoing re-evaluation of these assumptions is the spur to adopt new and more cognitively informed approaches (and possibly the refinement of cognitive models).

Acknowledgements

This work was supported by the Wellcome Trust. The work and ideas presented here represent the collective efforts of the Wellcome Department of Cognitive Neurology. The author thanks Nicki Roffe for preparing this manuscript, and the reviewers, Chris Frith and Cathy Price for invaluable comments and discussion.

References

- 1 Friston, K.J. et al. (1995) Statistical parametric maps in functional imaging: a general linear approach *Hum. Brain Mapp.* 2, 189–210
- 2 Zeki, S. et al. (1991) A direct demonstration of functional specialisation in human visual cortex *J. Neurosci.* 11, 641–649
- 3 Petersen, S.E. et al. (1989) Positron emission tomographic studies of the processing of single words *J. Cogn. Neurosci.* 1, 153–170
- 4 Gerstein, G.L. and Perkel, D.H. (1969) Simultaneously recorded trains of action potentials: analysis and functional interpretation *Science* 164, 828–830
- 5 McIntosh, A.R. et al. (1994) Network analysis of cortical visual pathways mapped with PET *J. Neurosci.* 14, 655–666
- 6 Friston, K.J. et al. (1995) Characterizing modulatory interactions between V1 and V2 in human cortex with fMRI *Hum. Brain Mapp.* 2, 211–224
- 7 Lueck, C.J. et al. (1989) The colour centre in the cerebral cortex of man *Nature* 340, 386–389
- 8 Pardo, J.V., Fox, P.T. and Raichle, M.E. (1991) Localization of a human system for sustained attention by positron emission tomography *Nature* 349, 61–64
- 9 Kosslyn, S.M. et al. (1995) Topographical representation of mental images in primary visual cortex *Nature* 378, 496–498
- 10 Martin, A. et al. (1995) Discrete cortical regions associated with knowledge of color and knowledge of action *Science* 270, 102–105
- 11 Sternberg, S. (1969) The discovery of processing stages: extension of Donders method *Acta Psychol.* 30, 276–315
- 12 Corbetta, M. et al. (1991) Selective and divided attention during visual discrimination of shape, color and speed. Functional anatomy by positron emission tomography *J. Neurosci.* 11, 2383–2402
- 13 Price, C.J. et al. (1996) The neural pathways sustaining object and color naming *Proc. R. Soc. London Ser. B* 263, 1501–1507
- 14 Friston, K.J. et al. (1992) Motor practice and neurophysiological adaptation in the cerebellum: a positron tomography study *Proc. R. Soc. London Ser. B* 248, 223–228
- 15 Shallice, T. et al. (1995) Brain regions associated with acquisition and retrieval of verbal episodic memory *Nature* 368, 633–635
- 16 Friston, K.J. et al. (1992) Measuring the neuromodulatory effects of drugs in man with positron tomography *Neurosci. Lett.* 141, 106–110
- 17 Dolan, R.J. et al. (1995) Dopaminergic modulation of impaired cognitive activation in the anterior cingulate cortex in schizophrenia *Nature* 378, 180–182
- 18 Grafton, S. et al. (1992) Functional anatomy of human procedural learning determined with regional cerebral blood flow and PET *J. Neurosci.* 12, 2542–2548
- 19 Price, C.J. et al. (1992) Regional response differences within the human auditory cortex when listening to words *Neurosci. Lett.* 146, 179–182
- 20 Buechel, C. et al. (1996) Nonlinear regression in parametric activation studies *NeuroImage* 4, 60–66
- 21 Grasby, P. et al. (1994) A graded task approach to the functional mapping of brain areas implicated in auditory-verbal memory *Brain* 117, 1271–1282
- 22 Frith, C.D. and Friston, K.J. (1997) The role of the thalamus in 'top-down' modulation of attention to sound *NeuroImage* 4, 210–215
- 23 Kwong, K.K. et al. (1991) Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation *Proc. Natl. Acad. Sci. U. S. A.* 88, 5675–5679
- 24 Malonek, D. and Grinvald, A. (1995) Interactions between electrical activity and cortical microcirculation revealed by imaging spectroscopy: implications for functional brain mapping *Science* 272, 551–554
- 25 Aertsen, A. and Preissl, H. (1991) Dynamics of activity and connectivity in physiological neuronal networks, in *Non-Linear Dynamics and Neuronal Networks* (Schuster, H.G., ed.), pp. 281–302, VCH Publishers
- 26 Sergent, J. (1994) Brain imaging studies of cognitive functions *Trends Neurosci.* 17, 221–227

Readers of *TICS* may be interested in the Special Issue on **Flavour perception** recently published by our sister journal, *Trends in Food Science & Technology*.

Flavour Perception: From basic research to industrial applications

Flavour perception mechanisms, by D.G. Laing and A. Jinks

Interactions among salty, sour and bitter compounds, by P.A.S. Breslin

How are bitter and sweet tastes related?, by D.E. Walters

Towards an improved understanding of sweetener synergy, by G.G. Birch

Interactions between umami and other flavor characteristics, by S. Fuke and Y. Ueda

Are there 'basic' tastes?, by J. Delwiche

Chemesthesis: Pungency as a component of flavor, by B.G. Green

Molecular mechanisms of olfactory perception: Their potential for future technologies, by G.A. Bell

Orders should be directed to the Regional Sales Offices, listed on p.iv of this issue.