

Forward inference using functional neuroimaging: dissociations versus associations

Richard Henson

MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge CB2 2EF, UK

Many people are excited by functional neuroimaging as a new tool for cognitive science; many others are sceptical. In this opinion article, I describe a ‘forward inference’ that one can make from patterns of brain activity to distinguish between cognitive theories. I give an example of forward inferences in research on recognition memory, and outline some statistical criteria for a ‘qualitative difference’ in brain activity. Forward inferences resemble the dissociation logic long-used in behavioural studies of healthy and brain-damaged people, although I argue that dissociations in neuroimaging data can go beyond behavioural dissociations. Nonetheless, forward inferences are only as good as the cognitive theories to which they pertain, and are most valuable in conjunction with other types of inference.

Introduction

The past decade has seen an explosion of research papers and newspaper articles involving the technique of functional neuroimaging, most notably functional magnetic resonance imaging (fMRI). This technique can provide a coarse measure of activity at different locations within the human brain while people are performing various cognitive tasks. This explosion of interest has also attracted much criticism [1–3], particularly from those psychologists and cognitive scientists who maintain that the mind can be studied independently of the brain. I, like many others, believe that knowledge about the brain does constrain our understanding of the mind, and that functional neuroimaging can inform us about human cognition. However, I am also wary, like others, of many of the assumptions and inferences made with functional neuroimaging.

In a companion paper in this issue [4], Poldrack formalizes the use of ‘reverse inference’ in functional neuroimaging. I would like to describe a complementary type of inference, which might be called a ‘forward inference’. Whereas the reverse inference is a (probabilistic) assignment of a cognitive process to activation of a specific brain region [4], forward inference refers to the use of qualitatively different patterns of activity over the brain to distinguish between competing cognitive

theories. More precisely, if one can design experimental conditions that differ in the presence of a cognitive process according to one theory, but not according to another, then the observation of distinct patterns of brain activity associated with those conditions constitutes evidence in favour of the first theory.

Example of forward inference in recognition memory

This type of forward inference is perhaps best illustrated with an example. There has been much debate between so-called ‘single-process’ versus ‘dual-process’ theories of recognition memory, for which extant behavioural data would seem inconclusive (see [Box 1](#)). Given that behavioural data are likely to reflect the amalgamation of multiple cognitive processes, neuroimaging data might be more revealing of such processes. For example, multiple brain regions are likely to contribute different types of information during memory retrieval, even though the different types of information need to be collapsed (via some decision process) into a single, categorical behavioural response when making a recognition memory judgment. Henson *et al.* [5] used fMRI to compare brain activity for items that subjects said they ‘remembered’ with that for items that subjects said they just ‘knew’ ([Box 1](#)). According to single-process models, such ‘Remember’ and ‘Know’ judgments differ only in the strength of memory for an item [6] ([Figure 1a](#)). According to dual-process models however, the two judgments differ in the relative contributions of two distinct forms of memory (e.g. recollection and familiarity [7]). Henson *et al.* found that a region in posterior cingulate (among other regions) was more active for Remember than Know judgments, whereas a region in right lateral frontal cortex was more active for Know than Remember judgments ([Figure 2a](#)). Although not reported in that paper, there was a significant interaction between the two regions and the two types of judgment. Moreover, both regions were active for either Remember or Know judgments relative to the ‘baseline’ of New judgments to unstudied words. These findings fulfil the criteria for a ‘qualitative’ difference in brain activity (assuming equivalent neural–haemodynamic mappings in the two regions; see [Box 2](#)).

This observation of a qualitative, rather than simply quantitative, difference in brain activity for Remember versus Know judgments would appear to favour dual-process over single-process models of recognition memory,

Corresponding author: Henson, R. (rik.henson@mrc-cbu.cam.ac.uk).

Available online 6 January 2006

Box 1. Theories of recognition memory

Cognitive psychologists have long argued whether the behavioural data from recognition memory tasks entail the existence of two types of memory process. In such tasks, subjects study a series of items, and then at some later point in time, are presented with another series of items and asked to distinguish the studied ('old') items from randomly intermixed unstudied ('new') items. According to **single-process theories**, studied and unstudied items can be represented as overlapping distributions along a single continuum of 'memory strength', upon which subjects impose a response criterion to make a binary old/new decision (an example of 'signal-detection theory' [20]). According to **dual-process theories** however, an 'old' decision can be based on two distinct types of process, such as 'recollection' and 'familiarity' [7].

One influential variant of the recognition memory task is to ask subjects to indicate their phenomenological experience accompanying 'old' decisions [21]. They are asked to respond 'Remember' when they remember a detail about the specific episode when an item was presented at study (e.g. what occurred before or after it), or 'Know' when the item seems familiar but they cannot remember any specific detail. Several experimental variables have been shown to dissociate Remember and Know judgments [22,23], *prima facie* supporting dual-process models. However, other theorists have pointed out that these

dissociations (together with other associations) can also be explained simply by two different response criteria within a signal-detection model [6,24] (see Figure 1a in main text).

The debate between single- and dual-process theories of recognition memory extends to other types of behavioural data, such as the precise form of Receiver–Operator Curves, but these data also appear to be indeterminate (e.g. dual-process models versus unequal-variance signal-detection models) [25]. The present argument is that, whereas single-process models might predict a 'quantitative' difference in brain activity between conditions differing in memory performance, they are inconsistent with a 'qualitative' difference in brain activity (as defined in Box 2). Such a qualitative difference therefore favours dual-process theories.

Two studies are described in the main text that illustrate the nature of a forward inference, and its theory-dependence. These two studies do not resolve the debate between single- and dual-process theories. Nonetheless, further neuroimaging studies have revealed a region in anterior medial temporal cortex that seems to track familiarity independent of recollection and independent of decision processes [26], and several other studies have reported patterns of brain activity that generally appear to favour dual-process relative to single-process models of recognition memory [27–30].

and hence constitutes a forward inference. This form of inference is hardly new; it resembles the 'dissociation logic' that has been used for many years in cognitive psychology, and in particular, cognitive neuropsychology [8]. The neuroimaging data are simply being treated as an additional dependent variable with which to find functional dissociations (cf. reaction times or accuracy).

Assumptions underlying neuroimaging inferences

One feature of forward inference is that it does not require strong selectivity of brain regions (i.e. a high value of $P(COG_X|ACT_Z)$ in Poldrack's formalism [4]). In the extreme case, it does not even matter which brain regions differ in the two conditions, or even if numerous regions differ, as long as there is a qualitative difference in the overall pattern of brain activity (although see below). The only assumption made is that the same cognitive process is not supported by different brain regions during different conditions within the same experiment. This is what I

previously called a weak form of 'systematicity' in the function–structure mapping [9], as distinct from the 'one-to-one' function–structure mapping that I argued was necessary for an unambiguous (rather than probabilistic [4]) form of reverse inference. In the present example, the forward inference does not imply the association of right frontal activity uniquely with familiarity and posterior cingulate activity uniquely with recollection; this would be a reverse inference (and depends further on the relationship between Remember and Know judgments, specifically, on whether they are redundant, independent or exclusive [10]). Having said this, forward inferences are of course more valuable if the specific regions that comprise the qualitative difference in brain activity are consistent with other data: for example, those brain regions that, when disrupted following brain damage (or transcranial magnetic stimulation [11]), produce some form of behavioural deficit under the same experimental conditions [9,12].

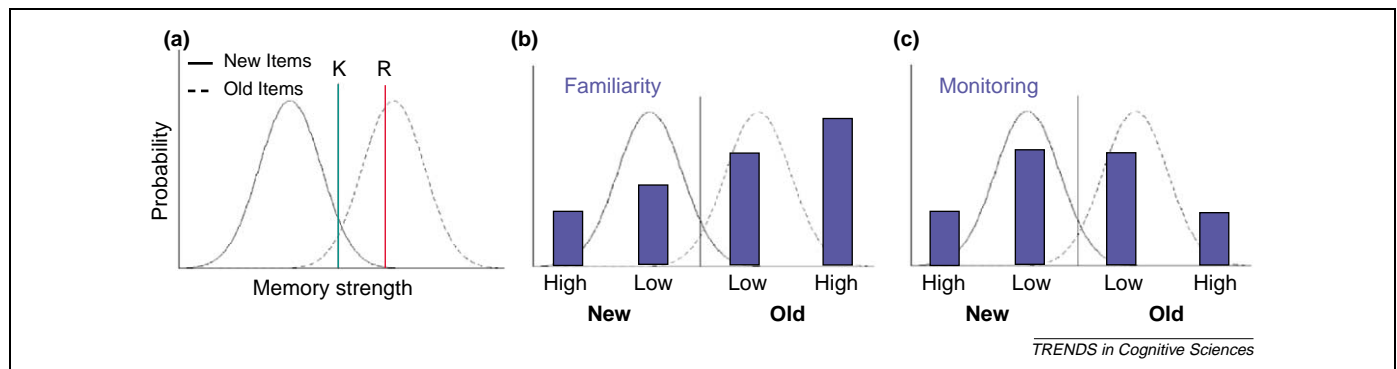


Figure 1. Schematics of a 'single-process' theory of recognition memory based on signal-detection theory [6,20], in which memory strength is a single continuum along which studied ('Old') and unstudied ('New') items are represented as Gaussian functions with different means (generally greater for 'old' items). Subjects impose a criterion for judging whether a given item is 'old' or 'new', or, as in (a), two criteria for whether an item is 'remembered' (to the right of the red vertical line labelled 'R'), 'known' (between the lines labelled 'R' and 'K') or 'new' (to the left of the green line labelled 'K'). The bars in (b) and (c) show the predicted fMRI responses for a study [14] in which old and new judgments were made with high and low confidence for (b) a brain region in which activity increases with memory strength (or 'familiarity' in some dual-process models [7]), and (c) a brain region in which activity increases as items approach the old/new response criterion, that is, for less confident responses, regardless of their studied status.

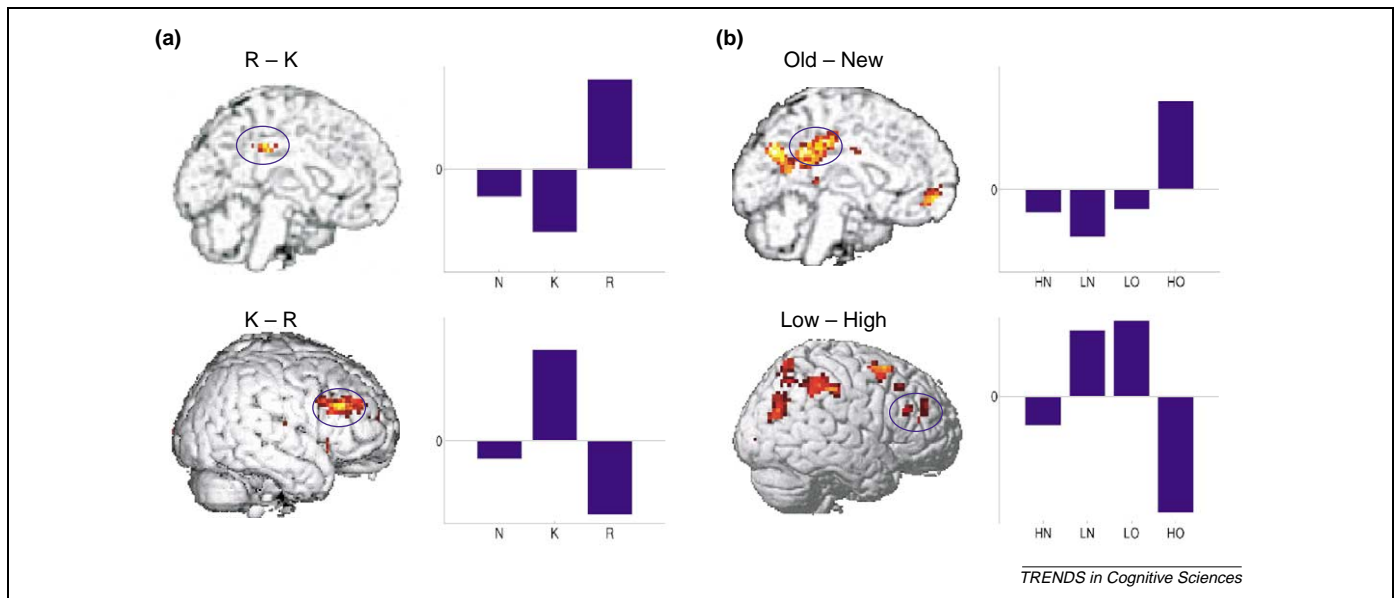


Figure 2. Results of two functional neuroimaging studies of recognition memory. (a) When contrasting old words in a recognition memory test as a function of whether they were 'remembered' (R) or 'known' (K) [5], greater responses were found in a region of posterior cingulate for R judgments (top) and in a region of right lateral frontal cortex for K judgments (bottom). The 'blobs' on the brains represent voxels surviving an *a priori* statistical threshold when tested across a sample of 11 healthy volunteers, after warping their brains to the same stereotactic space. The brain shown is an approximate rendering of a canonical brain in that space. The graphs beside each brain show the estimated fMRI response for R, K and N (correctly rejected new word) judgments from the maximum within the circled regions; the units and zero-value are arbitrary. (b) When contrasting words as a function of their correct judgment as 'old' or 'new' (O/N) and the confidence of that judgment as high or low (H/L), a region in posterior cingulate showed greatest responses for high-confidence old judgments (HO) (top), whereas a region in right lateral frontal cortex showed greatest fMRI responses for low-confidence judgments (bottom) [14]. For both (a) and (b), the different pattern of responses across the conditions for the two regions constitute a qualitative difference in brain activity, implying that the conditions differ in the presence of a cognitive process (i.e. the conditions engage at least two distinct processes).

Box 2. Qualitative differences in brain activity

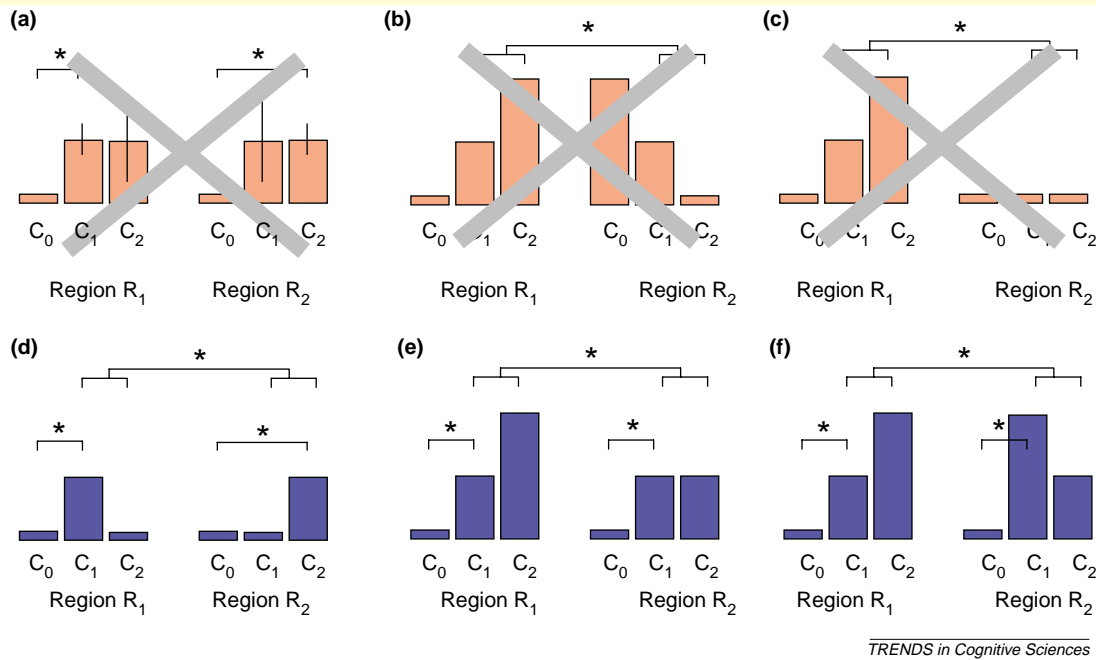
The argument concerning forward inference discussed in this article hinges on what is meant by a 'qualitative difference' in brain activity. Most analyses of neuroimaging data perform separate statistical tests for each element (voxel) within the brain image; the so-called 'mass univariate' approach. The results of these tests can be used to construct a Statistical Parametric Map (SPM) of the reliability of activity over the brain. These maps can be thresholded to show the 'significant' activations, often using sophisticated methods for correcting for the multiple statistical comparisons entailed [31]. This mass univariate approach is understandable, given an emphasis on 'localization' of cognitive processes. However, an observation of different suprathreshold activations when comparing two experimental conditions against a baseline is not sufficient for a forward inference. This is simply because the failure of voxels to survive a threshold does not imply that those voxels were not truly activated (the 'null result' problem in classical statistics; see Figure 1a, following page). The same problem arises when the two conditions are compared directly: finding that some, but not all, brain regions express significant differences might reflect only a quantitative difference in activity between the two conditions in the presence of noise levels that vary across voxels.

To claim a qualitative difference, I would argue that a minimal requirement is a significant interaction between experimental conditions and brain regions [6]. A significant interaction is evidence that the regions are *differentially* activated by the conditions. Such an interaction is not sufficient however. Even a 'crossover' interaction in which Region 1 is more active for Condition 1 than Condition 2, but Region 2 is more active for Condition 2 than Condition 1, could reflect a single cognitive process, engagement of which increases activity in Region 1 but decreases activity in Region 2 (Figure 1b). This might arise, for example, if the two regions were reciprocally interconnected. A further criterion therefore is that the regions are both significantly activated (or both deactivated) in at least one of the conditions relative to a third, baseline condition. This rules out the trivial case where one region is not differentially activated by any condition (e.g. voxels in a ventricle, which cannot contain neural activity), which might

nonetheless furnish a reliable interaction when tested against a region that is (Figure 1c). More importantly, the criterion that the regions must be activated (or deactivated) in the same direction in at least one of the conditions, relative to the baseline condition, overcomes the problem of reciprocal connections between brain regions mentioned above: if the activity in Regions 1 and 2 changes in the same way for Conditions 1 and/or 2 relative to the baseline condition, then they cannot be reciprocally connected.

The patterns in Figures 1d and e would therefore satisfy the above criteria. Nonetheless, there is an additional criterion that one might want to apply for haemodynamic measures like fMRI: given that the mapping from neural activity to haemodynamic activity can vary across brain regions, some Region-by-Condition interactions could arise simply because of different gains in this neural-haemodynamic coupling (often called 'multiplicative' or 'range effects' in behavioural data). This possibility is ruled out by the finding of a crossover interaction between regions and conditions (Figure 1f). In this final example, the combined criteria of a positive correlation across regions in one comparison (between conditions of interest and the baseline condition) and a negative correlation in another comparison (between the two conditions of interest) correspond to the 'reversed association' promoted for behavioural data [32].

Few neuroimaging studies to date report tests sufficient to meet these criteria for a 'qualitative difference'. Indeed, generalization of these criteria to multiple voxels and multiple conditions is likely to require further methodological developments. Note also that the above criteria are couched in classical statistics (rather than the Bayesian approach adopted by Poldrack [4]). Hence they still require control of false positives, in that the thresholds ('p-values') for rejecting the null hypothesis in each case must take into account how the regions were selected, for example, from searching multiple voxels, particularly when using a correlated rather than orthogonal comparison (an issue for the present examples). Moreover, the present criteria are probably not perfect, and are not meant to be proscriptive; the intention is simply to encourage consideration of the statistical issues pertinent to forward inferences.



TRENDS in Cognitive Sciences

Figure 1 (Box 2). Schematics of six possible fMRI responses across three conditions (C_0, C_1, C_2) within two regions (R_1, R_2) to illustrate the notion of a 'qualitative difference' in brain activity. C_1 and C_2 are conditions of interest, C_0 is some form of baseline. The y-axis represents level of fMRI activity. (a) shows a situation where C_1 is reliably greater than C_0 in Region R_1 but not in R_2 , whereas C_2 is reliably greater than C_0 in Region R_2 but not in R_1 ; a common result when thresholding statistical tests performed independently for each voxel. The pattern of reliability is reflected in the different error bars, which represent the estimation error in the contrast of C_1 and C_2 relative to C_0 : large error bars are the cause of 'null results' for $C_2 - C_0$ in R_1 and $C_1 - C_0$ in R_2 . (b) shows a reliable interaction between conditions C_1 and C_2 and regions R_1 and R_2 (indicated by the asterisk), which overcomes the problem of null results. However, the inverse pattern across conditions C_0 to C_2 for the two regions could reflect a quantitative difference across the conditions in a single cognitive process, which happens to activate R_1 but deactivate R_2 . (c) shows another reliable interaction, but in this case R_2 is not differentially activated by any of the conditions, preventing a clear interpretation. These three panels therefore do not constitute qualitative differences. (d) and (e) do constitute qualitative differences: both regions are differentially activated by at least one of the conditions, and show a crossover interaction with C_1 and C_2 (d), or a non-crossover interaction (e). Nonetheless, the patterns in (d) and (e) rely on equivalent mappings from neural to haemodynamic activity in regions R_1 and R_2 (otherwise differences in the shape of this mapping could cause the haemodynamic differences despite no qualitative differences in neural activity). (f) shows a crossover interaction between C_1/C_2 and R_1/R_2 , plus a difference between C_1/C_2 compared with C_0 that is in the same direction (+ve) for R_1 and R_2 . This pattern of association plus dissociation overcomes any differences in the neural-haemodynamic mapping in the two regions (assuming only that the mapping is monotonic in both).

Revisiting forward inferences in recognition memory

Another important feature of forward inferences is that they are always theory-dependent. In other words, a third theory can be created that interprets the qualitatively different pattern of activity in a different way from the original theories tested. To take the above example of Remember/Know judgments, a single-process memory theorist might argue that Remember and Know judgments are based on the same, single dimension of memory strength, but nonetheless differ in other, non-mnemonic ways. Indeed, we wondered whether the right frontal activation for Know relative to Remember judgments had more to do with decision processes, rather than reflecting memory strength *per se*. For example, when an item seems familiar, but details of its previous occurrence are not remembered, subjects might engage in greater checking or 'monitoring' of retrieved information [13]. To test this, we conducted a further neuroimaging experiment [14], in which subjects indicated the confidence of their old/new decision ('high confidence new', 'low confidence new', 'low confidence old', or 'high confidence old'). If right frontal activity reflects memory strength, it should vary monotonically across these confidence ratings (Figure 1b). Alternatively, if it reflects decision processes, it should be greatest for low- relative to high-confidence judgments, regardless of old/new status (i.e. greatest when memory

strength, according to a single-process model, is close to the old/new response criterion; Figure 1c).

Our data supported the latter outcome – that right lateral frontal activity reflects confidence of the decision, rather than memory strength (Figure 2b). The posterior cingulate region, on the other hand, continued to show greatest activity for 'high confidence old' decisions. Thus, the data again revealed a qualitative difference in brain activity as a function of experimental condition, but are now interpretable in terms of a new model in which decisions are based on a single-continuum of memory strength, and memory strengths close to a response criterion entail additional cognitive processes, such as 'monitoring' [14], before a decision is made.

Further considerations for forward inferences

The purpose of the above two examples is to illustrate the theory-dependence of forward inferences (as applies equally to behavioural data from patients or healthy controls). Indeed, one can always invent re-interpretations of data *post hoc*. To be valid however, an alternative explanation must be specified at a level that allows further experiments to be conducted that distinguish the original explanation from the new explanation [15,16]. For example, further neuroimaging experiments have found differences in brain activity associated with recollection

Box 3. Questions for further research

- Is it a reasonable assumption that the same cognitive process cannot activate different brain regions within different conditions of the same experiment?
- What exactly is meant by a cognitive process (consider, for example, the case of 'selective attention' in relation to the above question)?
- Are some cognitive processes invisible to current neuroimaging techniques (even putting aside issues related to spatial resolution), for example, by being realised by rapid changes in communication between brain regions, without changes in the mean metabolic activity within those regions?
- Are the criteria for a qualitative difference in brain activity (outlined in Box 2) necessary and sufficient?
- How are the criteria for a qualitative difference generalizable to multiple brain regions/voxels?
- Have we made any progress in cognitive theory to date using functional neuroimaging data?

and familiarity, even when decision processes are controlled by various experimental manipulations (see Box 1). A common abuse of this principle, however, is the objection that the difference in brain activity arises simply because one condition is 'more difficult' than another. Without an additional specification of how 'difficulty' is manipulated experimentally, this objection has little force. Although one might equate difficulty with poorer accuracy or longer response times, the cause of such behavioural differences is often the cognitive process of interest itself: such operationalizations of difficulty are only useful if they can be experimentally manipulated while holding the process of interest constant. Thus, although one can always list potential confounds of a specific experimental comparison, the only ones of value are those that (i) relate to existing theories (i.e. are theoretically-interesting potential confounds) and (ii) can be experimentally manipulated in a manner orthogonal to the cognitive process of interest.

Now one might argue that if two experimental conditions differ (and subjects are aware of this), then there must be a difference in activity somewhere in the brain, almost by definition. The immediate response is that it must be a qualitative rather than simply quantitative difference (see Box 2). But more importantly, the main goal of experimental design is to create conditions that differ only in the specific hypothetical process of interest. This has been called the 'pure insertion' assumption [17]; an assumption generic to psychological experiments and not specific to function neuroimaging [9]. So if, for example, Remember and Know judgments differed in the specific fingers with which subjects responded, one would design the experiment to control for this difference (e.g. by counter-balancing finger assignment across subjects). The aim is to contrast conditions that differ only in the cognitive process that distinguishes the competing theories been tested, and hence there is no reason to always expect a qualitative difference in brain activity.

Dissociation logic is not without its critics [18], one criticism being that dissociations can only result in further fractionation of cognitive processes. It is also closely tied to the 'modularity' assumption that is prevalent in cognitive science (see, for example, [8,19]), but which again has its critics [1]. This is why I believe that associations are also

necessary, as provided for example by the reverse inference discussed by Poldrack [4]. At a minimum, reverse inference is necessary to determine whether a specific pattern of activity that dissociates two conditions has been replicated in a second experiment. More importantly, reverse inference allows us to relate cognitive processes across different theories. Indeed, common sites of activation across quite different experiments might allow connections to be made between fields of theorizing that have traditionally been distinct (what I previously called 'structure-to-function induction' [9]). In other words, the brain can provide a framework within which to investigate associations as well as dissociations across different experimental tasks.

Conclusions

The increased use of functional neuroimaging to address cognitive theory has practical consequences. Given its considerable expense for example, fMRI research imposes financial burdens on funding bodies and universities (although of course there are good reasons for funding MRI scanners other than using them for making inferences about human cognition). This is why it is important to think carefully about the type of inferences that can be made from functional neuroimaging data (see also Box 3). This article has outlined one type of inference – forward inference – and some of the associated assumptions and caveats (mostly shared with conventional dissociation logic). Other inferences, such as reverse inference [4], make related, but different and sometimes stronger assumptions [9]. However, only by making these caveats and assumptions explicit, and criticizing them, will we be able to assess the real value of functional neuroimaging for cognitive science.

Acknowledgements

The author thanks Russ Poldrack and Shbana Rahman for helpful comments during this article's preparation. The author is supported by the UK Medical Research Council.

References

- 1 Uttal, W.R. (2001) *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*, MIT Press
- 2 Coltheart, M. (2002) Brain imaging, connectionism and cognitive neuropsychology. *Cogn. Neuropsychol.* 21, 21–25
- 3 Van Orden, G.C. and Paap, K.R. (1997) Functional neuroimages fail to discover pieces of mind in the parts of the brain. *Philos. Sci.* 64, S85–S94
- 4 Poldrack, R.A. (2006) Can cognitive processes be inferred from neuroimaging data? *Trends Cogn. Sci.* 10, doi:10.1016/j.tics.2005.12.004
- 5 Henson, R.N.A. et al. (1999) Recollection and familiarity in recognition memory: an event-related fMRI study. *J. Neurosci.* 19, 3962–3972
- 6 Donaldson, W. (1996) The role of decision processes in remembering and knowing. *Mem. Cogn.* 24, 523–533
- 7 Yonelinas, A.P. (2002) The nature of recollection and familiarity: a review of 30 years of research. *J. Mem. Lang.* 46, 441–517
- 8 Shallice, T. (1988) *From Neuropsychology to Mental Structure*, Cambridge University Press
- 9 Henson, R. (2005) What can functional neuroimaging tell the experimental psychologist? *Q. J. Exp. Psychol. A* 58, 193–234
- 10 Knowlton, B.J. and Squire, L.R. (1995) Remembering and knowing: two different expressions of declarative memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 21, 699–710

- 11 Cowey, A. and Walsh, V. (2001) Tickling the brain: studying visual sensation, perception and cognition by transcranial magnetic stimulation. *Prog. Brain Res.* 134, 411–425
- 12 Sarter, M. *et al.* (1996) Brain imaging and cognitive neuroscience. *Am. Psychol.* 51, 13–21
- 13 Burgess, P.W. and Shallice, T. (1996) Confabulation and the control of recollection. *Memory* 4, 359–411
- 14 Henson, R.N.A. *et al.* (2000) Confidence in recognition memory for words: dissociating right prefrontal roles in episodic retrieval. *J. Cogn. Neurosci.* 12, 913–923
- 15 Mill, J.S. (1843, reprinted 2002) *A System of Logic*, University Press of the Pacific
- 16 Boring, E.G. (1963) *History, Psychology and Science: Selected Papers* (Watson, R.I. and Campbell, D.T. eds), John Wiley & Sons
- 17 Sternberg, S. (1969) The discovery of processing stages: extensions of Donders method. *Acta Psychol. (Amst.)* 30, 276–315
- 18 Van Orden, G.C. *et al.* (2001) What do double dissociations prove? *Cogn. Sci.* 25, 111–172
- 19 Henson, R.N.A. What has (neuro)psychology told us about the mind (so far)? A reply to Coltheart (2005). *Cortex* (in press)
- 20 Green, D.M. and Swets, J.A. (1966) *Signal Detection Theory and Psychophysics*, John Wiley & Sons
- 21 Tulving, E. (1985) Memory and consciousness. *Can. Psychol.* 26, 1–12
- 22 Gardiner, J.M. and Java, R.I. (1990) Recollective experience in word and nonword recognition. *Mem. Cogn.* 18, 23–30
- 23 Rajaram, S. (1993) Remembering and knowing: two means of access to the personal past. *Mem. Cogn.* 21, 89–102
- 24 Dunn, J.C. (2004) Remember–Know: a matter of confidence. *Psychol. Rev.* 111, 524–542
- 25 Heathcote, A. (2003) Item recognition memory and the receiver operating characteristic. *J. Exp. Psychol. Learn. Mem. Cogn.* 29, 1210–1230
- 26 Henson, R.N. *et al.* (2003) A familiarity signal in human anterior medial temporal cortex. *Hippocampus* 13, 301–304
- 27 Ranganath, C. *et al.* (2004) Dissociable correlates of recollection and familiarity within the medial temporal lobes. *Neuropsychologia* 42, 2–13
- 28 Davachi, L. *et al.* (2003) Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. *Proc. Natl. Acad. Sci. U. S. A.* 100, 2157–2162
- 29 Yonelinas, A.P. *et al.* (2005) Separating the brain regions involved in recollection and familiarity in recognition memory. *J. Neurosci.* 25, 3002–3008
- 30 Rugg, M.D. and Yonelinas, A.P. (2003) Human recognition memory: a cognitive neuroscience perspective. *Trends Cogn. Sci.* 7, 313–319
- 31 Friston, K.J. *et al.* (1995) Statistical parametric maps in functional imaging; a general linear approach. *Hum. Brain Map.* 2, 189–210
- 32 Dunn, J.C. and Kirsner, K. (1988) Discovering functionally independent mental processes: the principle of reversed association. *Psychol. Rev.* 95, 91–101

OXFORD SUMMER SCHOOL ON CONNECTIONIST MODELLING
Department of Experimental Psychology, University of Oxford
23 July–4 August, 2006

Applications are invited for participation in a two-week residential Summer School on techniques in connectionist modelling. The course is aimed primarily at researchers who wish to exploit neural network models in their teaching and/or research, and will provide a general introduction to connectionist modelling, biologically plausible neural networks and brain function through lectures and exercises on PCs. The instructors with primary responsibility for teaching the course are Kim Plunkett and Edmund Rolls.

No previous knowledge of computational modelling will be required although simple word processing skills will be assumed. Participants will be encouraged to start work on their own modelling projects during the Summer School.

The cost of participation in the Summer School is £1050. A small number of partial bursaries (£200) will be available for graduate students.

If you are interested in participating in the Summer School, please complete the application form no later than 23 January 2006, at: <http://www.psy.ox.ac.uk>