

REVIEW ARTICLE

Experimental Design and Interpretation of Functional Neuroimaging Studies of Cognitive Processes

David Caplan*

*Neuropsychology Laboratory, Department of Neurology, Massachusetts General Hospital,
Boston, Massachusetts*



Abstract: This article discusses how the relation between experimental and baseline conditions in functional neuroimaging studies affects the conclusions that can be drawn from a study about the neural correlates of components of the cognitive system and about the nature and organization of those components. I argue that certain designs in common use—in particular the contrast of qualitatively different representations that are processed at parallel stages of a functional architecture—can never identify the neural basis of a cognitive operation and have limited use in providing information about the nature of cognitive systems. Other types of designs—such as ones that contrast representations that are computed in immediately sequential processing steps and ones that contrast qualitatively similar representations that are parametrically related within a single processing stage—are more easily interpreted. *Hum Brain Mapp* 30:59–77, 2009. © 2007 Wiley-Liss, Inc.

Key words: functional neuroimaging inferences; experimental design; validity and soundness



INTRODUCTION

The ability to design experiments that isolate cognitive operations is critical to the use of functional neuroimaging to provide evidence regarding the neural basis of cognitive processes. Much of the discussion of whether psychological methods isolate cognitive operations in neuroimaging experiments deals with the choice of materials and tasks in

particular experiments, asking questions such as whether these choices adequately control for “nuisance” variables that are confounded with the variables that have been manipulated, whether these choices reflect a theoretically justified analysis of the cognitive processes under study, and other similar questions that pertain to a particular area of study [see, e.g., Caplan, 2006a,b]. Only a few articles discuss the design of functional neuroimaging experiments from a more general perspective [for examples, see, Caplan and Moo, 2004; Coltheart, 2006a,b; Friston et al., 1996; Henson, 2005, 2006a,b; Jack et al., 2006; Jennings et al., 1997; Newman et al., 2001; Page, 2006; Poldrack, 2006; Price and Friston, 1997; Price et al., 1997; Schutter et al., 2006; Shallice, 2003; Sidtis et al., 1999; Umiltà, 2006]. This article continues the discussion of general features of experimental designs that affect their interpretability by examining the conditions under which the contrast of conditions in a functional neuroimaging study

Contract grant sponsor: NIDCD; Contract grant number: DC02146.

*Correspondence to: David Caplan, M.D., Ph.D., Neuropsychology Laboratory, 175 Cambridge Street, Suite 340, Boston, MA, 02114.

E-mail: dcaplan@partners.org

Received for publication 13 December 2006; Revised 21 August 2007; Accepted 22 August 2007

DOI: 10.1002/hbm.20489

Published online 2 November 2007 in Wiley InterScience (www.interscience.wiley.com).

could possibly isolate an operation or a component of the cognitive processing system.

The discussion here assumes that models of cognitive functions consist of two basic elements: a specification of the computations that occur at various stages of processing and a “functional architecture” that specifies how these computations are ordered relative to one another. Computations consist of operations on representations. A representation is a basic entity in a theory of a cognitive function, such as a structural description of an object or the form of a word. Some representations are stored in long term memory; others are constructed by perceptual and other processes and exist transiently in short-duration memory systems (sometimes called “working memory” systems). An “operation” changes the state of a representation. Operations may change the level of activation of an existing representation, such as when the presentation of a written stimulus leads to increases in the level of activation of the form of that word in a long term memory store, or may create a new representation such as when operations transiently construct a syntactic representation on the basis of a series of words.

The specifics of cognitive representations and operations differ greatly from model to model. For instance, models of assigning syntactic structure (parsing) and using it to constrain sentence meaning (sentence interpretation) range from ones that postulate highly abstract representations and operations that create them [e.g., Frazier and Clifton, 1996] to connectionist (neural net) models that deny the existence of any syntactic representations at all, even syntactic categories such as nouns and verbs [McDonald and Christiansen, 2002]. A great deal of theorizing in cognitive psychology focuses on the appropriate way to represent information and the nature of the operations that apply to these representations (see, e.g., the debates about “connectionist” and “procedural” models of processing in many areas of cognition).

Whether two representations or operations are qualitatively different or the same is determined by the theory of cognition that is under investigation. For instance, while most contemporary models consider memory representations of personal events and factual knowledge to be different and to be maintained in different memory stores (episodic and semantic memory), some models question this distinction [Howard et al., 2007]. In a similar way, whether two representations or operations are considered to be qualitatively different or the same is in part a matter of the grain size of the theory under investigation. A theory may consider animals and tools to be qualitatively different types of representations within a model of visual perception, but all objects may be treated as being qualitatively similar for the purpose of a given study (e.g., if a researcher is interested in differences in semantic information that is activated by words and by pictures of objects).

“Components” of the cognitive processing system (sometimes called “modules”) are sets of representations that are qualitatively similar and/or sets of similar operations that

apply to the same type of representation. For instance, some theories group the ensemble of representations of the phonological forms of words into a component of the cognitive processing system known as a “phonological lexicon” and some group a set of operations that activate the individual sounds of words from acoustic waveforms into a component known as “phoneme recognition.”

Operations and components may affect representations in series and in parallel, and may be subject to feedback and interactions; these relationships are specified in a “functional architecture.” The grain size of a functional architecture can vary from large components to individual operations. For instance, a model of object naming might postulate a functional architecture that involves three major cognitive components—object recognition, word form activation, phonological planning—or it might postulate a functional architecture that involves specification of the relation between operations within each of these components (e.g., activation of representations of visual features, of groupings of features, and of object properties may be operations within a model of object recognition). Claims regarding the neural basis of cognitive functions have been made at all levels of detail, from specific operations through components.

The study of the neural basis of cognitive operations involves contrasting tasks that are assumed to engage operations or components that are related to one another in ways that are specified in a functional architecture. This article considers the conclusions that can be drawn from particular contrasts. It reviews seven functional architectures and argues that contrasts of operations that are related in particular ways can, in principle, isolate cognitive operations or components, and that contrasts of operations that are related in other ways cannot, in principle, do so.¹ Though the examples of studies are drawn from functional magnetic resonance imaging, the conclusions of this article apply to all forms of neuroimaging that use the

¹An analogy to syllogisms may in some ways capture the difference between the issues discussed here and those that arise in connection with individual studies. Syllogisms can be valid (i.e., the conclusion follows from the premises) and/or sound (i.e., the premises are true). Some syllogistic figures cannot yield a valid conclusion no matter what the truth of their premises; e.g., the premises *Some carpenters are plumbers* and *some plumbers are electricians* cannot yield a valid conclusion relating carpenters and electricians even if the premises are true. Syllogisms with false premises can yield valid conclusions, but the conclusions are likely to be false; e.g., the premises *All carpenters are plumbers* and *all plumbers are electricians* yields the valid conclusion that *all carpenters are electricians*, but this conclusion is likely to be false because both of the premises are false. Like syllogisms, studies whose designs could, in principle, yield interpretable results can be considered valid. However, such designs can still fail to provide data about the operations they purport to study because of flaws in the choice of materials and tasks (on analogy to syllogisms, they are not sound).

designs discussed here, including event related potentials, magnetoencephalography, optical imaging, and others.²

PRELIMINARY ISSUES

A few preliminary issues pertaining to how cognitive processes are related to neural observations are necessary before turning to the discussion of experimental designs. Cognitive operations require “processing resources” [Shallice, 1988, 2003]. Different operations and sets of operations make greater or lesser demands on processing resources, leading to changes in responses such as accuracy, reaction time, BOLD signal level, amplitude and latency of electrophysiological signals, etc. Two fundamental assumptions about neural activity are that (1) measures of neural activity such as BOLD signal, elec-

²The issues discussed in this article are relevant to studies that use functional neuroimaging to study the neural basis of cognition and the organization of cognitive processes. They are not directly relevant to studies of the neural correlates of cognitive functions that are not directed at the question of how the brain is organized to perform a cognitive operation. For instance, a study could identify different neurovascular responses to a cognitive task in teenagers with and without ADD. This might be useful in many ways (e.g., in making a diagnosis in unclear cases) even if the task did not provide information about how the brain supported a cognitive function. However, if a researcher wants to know whether there is a difference in the way the brains of teenagers with and without ADD perform a particular cognitive function, such as dividing attention, it becomes necessary to use experimental designs that could identify brain regions that support those functions. In that case, the issues discussed here become important to consider.

³These assumptions have been questioned; this note is designed to answer some of these concerns.

Page [2006] has argued that Logothetis et al. [2001] showed that BOLD signal is not a measure of axonal spike activity, and thus not a measure of neuronal activity relevant to information processing since spike activity is the efferent signal of a brain area. This, however, misconstrues Logothetis et al. [2001]. Logothetis et al. [2001] found that BOLD signal correlated more highly with measures of dendritic activity (local field potentials) than with axonal spike frequencies, not that the correlation of BOLD signal and axonal spike frequency was not significant. It also bears noting that dendritic activity is an electrochemical state of a brain region that is likely to be informationally-relevant.

Page [2006] has also argued that increases in neurovascular responses might reflect increased inhibitory functioning of an area of the brain. This is unlikely: the percentage of GABAergic inhibitory neurons cortical neurons is estimated at between 15 and 25% [Kisvarday et al., 1990]. It is not now known whether there are areas of cortex with higher percentages of inhibitory neurons. If there are, BOLD signal increases in these areas can be interpreted accordingly.

Finally, there are several reasons to believe that increased BOLD signal reflects increased cognitive demand [see Shallice, 2003, for discussion]: linear models are good fits to BOLD signal [Boynton et al., 1996]; stimuli that affect behavioral measures in ways that indicate that they are more demanding lead to increases in BOLD signal; as tasks become easier, more practiced, and more automatic, neurovascular activity decreases [Raichle et al., 1994]. It bears mention that the relation between behavioral DVs and IVs is not always monotonic [the Yerkes-Dodson Law; Kahnemann, 1973] or linear [McClelland, 1979] but it is possible to use stimuli whose behavioral effects are known to be linear in the range of the IV that is being manipulated.

trophysiological signals measured at the scalp, etc., are related to cellular and subcellular neural events that are the basis for cognitive functions; and (2) as a cognitive task requires more operations, or more “effort” or “processing resources” to apply an operation, the measured neural activity increases.³

Much of neuroimaging deals with the question of whether a cognitive operation is supported by spatially contiguous neural tissue and whether any spatially contiguous neural tissue that does support an operation is located within recognizable macroscopically or microscopically defined brain areas. Models range from invariant localization in a small area through distribution over a large area, and include variable localization [e.g., Caplan et al., 2007], uneven distribution [e.g., Mesulam, 1990], and degeneracy [a variant of localization in which more than one structure independently supports a function—Noppeney et al., 2004].⁴

The areas of the brain that are referred to in these models are hard to define and identify. The usual areas in which a function is said to be localized or distributed are gyri and cytoarchitectonically defined brain regions. Gyri and other macroscopically defined areas are identifiable radiologically in individual brains [Caviness et al., 1996] but are only approximated in normalized images because of their variability across individuals [Geschwind and Galaburda, 1987]. Macroscopically defined areas of the brain are only shorthand for brain areas that contain cellular and subcellular elements that support particular cognitive operations, and the boundaries of such microscopically defined areas do not align with visible macroscopic boundaries [Caspers et al., 2006; Mazziotta et al., 2001] and are not themselves visualizable by current imaging methods. Moreover, the physiological basis for localization or distribution of a function is dynamic on a very short temporal scale [Recanzone et al., 1993]. Despite this, approximating the neural areas in which activity increases in response to a cognitive demand is a first step towards developing deeper understanding of how cognitive operations are related to the brain.

The neural areas that are activated by a cognitive contrast are a subset of those that are sufficient to support the operations that the contrast isolates. It cannot be assumed that the neural activity that is observed in an experiment is the only neural response to the experimental variables. This “completeness assumption” is obviously false due to sensitivity limits of existing methods. However, if the completeness assumption is made (perhaps on the meta-theoretical grounds that science is always constrained by the sensitivity of its measurements and that the neural measures we use are the best we have to identify neural responses to cognitive contrasts), neural areas that are activated by a cognitive contrast are those that are sufficient to support those operations. These areas may not be necessary for those

⁴A last model would be one that postulates variable uneven distribution, in which a function is unevenly distributed in a brain region, and the parts of the entire area that are most critical to it differ in different individuals. To my knowledge, no such models have been suggested.

operations; they could be damaged and the operations yet be performed if other areas of the brain can assume their functions. As in much of the literature, these areas will be referred to as “being involved in” or “supporting” these operations. These terms are used for ease of exposition and style only; the relationship they intend to convey is that of sufficiency, either total or partial depending upon whether the completeness assumption is made or not.

MODELS UNDERLYING EXPERIMENTAL DESIGNS

Functional neuroimaging studies involve a contrast between one or more experimental conditions, the experimental conditions, and one or more baseline conditions, the baseline conditions. The experimental condition(s) and baseline condition(s) are chosen in relationship to a cognitive model in such a way that the contrast of the two conditions will reveal something about the neural basis for a cognitive operation specified in the model, or about an aspect of the cognitive model itself. As indicated above, I shall argue that whether the experimental condition/baseline condition contrast can accomplish these goals depends upon the relationship between the experimental condition(s) and baseline condition(s) that is specified in the model that underlies the study. The article is thus organized along lines corresponding to the relationships between the experimental condition(s) and the baseline condition(s) specified in the underlying models. Seven major types of relations, each of which has been proposed in some part of a cognitive functional architecture, are reviewed.⁵ Figure 1 shows these seven functional architectures.

⁵It has sometimes been argued that it is unnecessary to consider all these models, because facts about the brain show that some cannot be correct. For instance, Friston et al. [1996] argue that “even if, from a functionalist perspective, a cognitive component can be added without interacting with preexisting components, the brain’s implementation of these processes is almost certainly going to show profound interactions. This follows from the observation that neural dynamics are nonlinear (p 98)” Similarly, in a review of a previous version of this article, Rik Henson suggested that “the brain is a highly interactive system with massive feedforward and feedback connectivity. Thus a “cascade” model would seem to be more appropriate, in which the parallel/serial distinction vanishes.” These conclusions about how parts of neural systems operate and how they support cognitive functions are premature. Whatever the dynamics of neural activation may be at the cellular and subcellular level, and however densely connected parts of the brain may be, the temporal scheduling of neural events in particular parts of the brain that support cognitive operations remains a question that is open to study. For instance, to take a trivial example, neither the details of the response properties of cells in lateral geniculate to retinal input nor their connectivity to V1 completely determines the temporal relation between the operations carried out in LGN and V1. The neuroanatomy establishes that LGN provides input to V1, but it does not establish whether that input occurs in serial or in cascade; similarly, projections from V1 to LGN strongly point to feedback from V1 to LGN but give no information about what operations in LGN that feedback affects, or whether the feedback occurs in a serial or cascade manner.

In all models, the diagrams indicate that the components of the cognitive processing system that are under study operate after other processing has occurred (this is indicated by the stage of “antecedent processing” in the diagrams). In all models except Functional Architecture 7, the model specifies a relationship between two or more components of the cognitive processing system; in Functional Architecture 7, the model considers the operations under study to be part of a single component of the cognitive processing system.

In the first model (Functional Architecture 1: “serial” models), the operations involved in performing the baseline condition (Processing Component 1) are required for the performance of the experimental condition (Processing Component 2) and are complete before the operations that are unique to the experimental condition are engaged. For instance, many models maintain that processes such as identification of acoustic energy at different frequencies over time precede and are the input to processes that categorize sounds phonologically [e.g., Klatt, 1979; Stevens, 2000].

In the second model (Functional Architecture 2: “parallel” models), the operations of the experimental condition (Processing Component 2) and the baseline condition (Processing Component 1) are not required for the other condition, and are postulated to require a common input. In most cases, they are assumed to occur at the same time. For instance, most models that postulate different operations in perceptual identification of different types of visual stimuli, such as faces, animals, tools, words, etc., maintain that these different operations occur in parallel after operations that provide the input to these processes (e.g., line detection; identification of geons, etc.) have occurred [e.g., Dehaene, 2001; Kanwisher, 1997; Martin et al., 1996, 1999].

In the third model (Functional Architecture 3: “two serial components operating in parallel”), two experimental condition/baseline condition contrasts are used, each of which involves a serial relation between the experimental condition and the baseline condition (in Fig. 1, Processing Component 1 is the baseline and Processing Component 2 is the experimental condition for one contrast, and Processing Component 3 is the baseline and Processing Component 4 is the experimental condition in a second).

In the fourth model (Functional Architecture 4: “cascade” models), the operations involved in performing the baseline condition (Processing Component 1) are required for the performance of the experimental condition (Processing Component 2) but are not complete before the operations that are unique to the experimental condition are engaged [McClelland, 1979]. For instance, Humphreys et al. [1999] have developed a model of object naming in which later stages of processing (e.g., activation of semantic knowledge about objects; word retrieval) occur while processing at earlier stages (e.g., activation of the representation of an object in long term memory; i.e., of the “structural description” of an object) is still taking place.

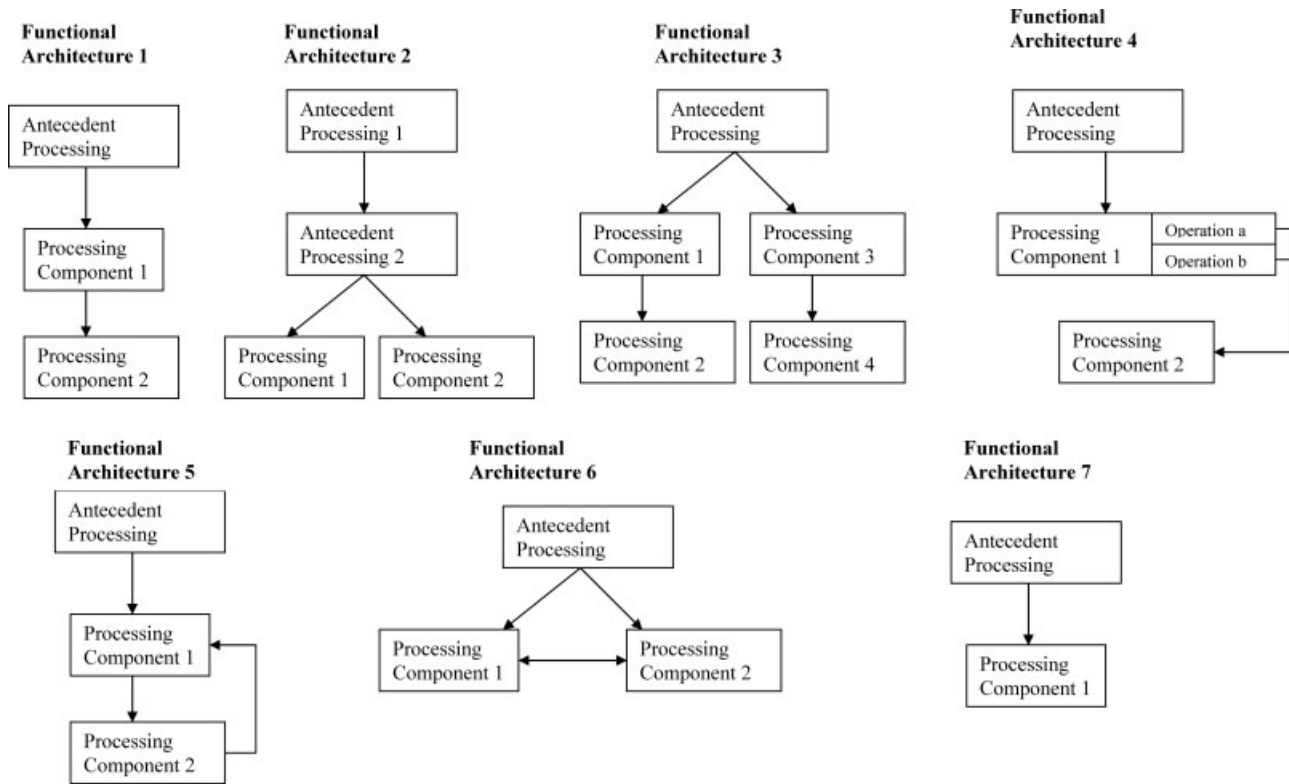


Figure 1.

Functional Architectures. Figure 1 shows diagrams of seven relations between cognitive processing components (functional architectures) that are tested in neuroimaging experiments. In all cases, the diagrams indicate that stimuli are processed in some way before the operations that are under investigation (indicated as Antecedent Processing).

The next two models, (Functional Architectures 5 and 6), are versions of serial and parallel models in which there is some sort of interaction between the operations in the experimental condition and the baseline condition—either one-way feedback from the experimental condition (Processing Component 2) to the baseline condition (Processing Component 1) (Functional Architecture 5: “serial model with feedback”) or mutual influence (Functional Architecture 6: “interactive parallel model”).

In the last model (Functional Architecture 7: model with only one processing component), the experimental condition(s) and baseline condition(s) contain operations that compute the same types of representations at the same stage of processing; the effort is to see how changes in the processing demand made upon the operations in a single component affect neural responses.

The basis for arguing that some part of the cognitive system makes use of components that are arranged in series, parallel or cascade, and that are subject to feedback or interaction, has been extensively discussed, dating at least from

the work of Donders [1868; see Sternberg 1969, 1998, 2001, for important contributions]. Both formal mathematical treatments [e.g., Townsend and Ashby, 1987] and simulations [e.g., McClelland, 1979] that have shown that many models are hard to distinguish.⁶ For instance, a widely held view is that distinguishing serial and cascade models is possible depending upon whether factors that are thought to affect a temporally early stage of processing interact with those that are thought to affect a later stage [Humphreys et al., 1988]. McClelland [1979], however, showed that a sim-

⁶A great deal of modeling of this sort has been directed at the question of how patterns of behavioral results are related to serial or parallel models of operations of the same type that apply at the same stage of processing, such as the question of whether search through perceptual arrays or memory stores occurs in serial or in parallel [e.g., Townsend and Ashby, 1987]. Most of the discussion in this article deals with more than one processing stage.

ulation of results by Meyer et al. [1975] showing a super-additive interaction of visual masking and context relatedness in lexical decision in a cascade model resulted in an interaction between these factors only if these factors affected the asymptotic level of accuracy at the stages of letter and word activation, not the temporal dynamics of those processing stages. To know whether this simulation accurately captures the functional architecture of the processes involved in word recognition, it is necessary to undertake speed-accuracy trade-off studies of the effects of these factors on measures that reflect letter and word activation. Studies at this level of detail have rarely been undertaken, leaving uncertainties about most of the models that have been proposed for cognitive functional architectures.

These considerations are important to cognitive functional neuroimaging. If the behavioral data underlying a neuroimaging study are ambiguous between several models of the cognitive function under study, neural data must be evaluated for whether they choose between these models, not just for whether they are consistent with one. However, the discussion to follow accepts the analysis of cognitive function that is assumed in the studies that are presented as examples of particular designs. While it would be valuable to consider all the possible models that are compatible with the behavioral or neural data in these studies, this is well beyond the scope of this article.

GENERAL CONDITIONS ON EXPERIMENTAL DESIGNS

One last issue needs to be discussed before considering different designs. The discussion to follow proposes conditions that have to be met regarding the relationship of the baseline condition and the experimental condition for an experimental condition/baseline condition contrast to isolate a cognitive operation. In addition to these conditions on how the baseline condition and experimental condition are related to each other, there are also two conditions that must apply to the baseline condition and the experimental condition themselves for a study to be interpretable. As these conditions always must be met for a study to be interpretable, they are presented once here.

First, it must be possible to identify the operations that are necessary for the baseline condition and the experimental condition to be performed (this may be called the “necessary operation requirement”). The “necessary operation requirement” requires using a task that can only be accomplished if a particular type of representation is processed at a particular cognitive stage. In that case, successful performance implies that that representation has been processed at that stage of processing.

Second, the operations in the baseline condition or the experimental condition must be restricted to those that are sufficient for the task to be performed (this may be called the “exclusive operations assumption.”) Without this assumption, it is always possible that what Page [2006] calls “epiphenomenal” activity—operations that are associ-

ated with those under study—is responsible for neurovascular effects. The “exclusive operations assumption” is obviously not met in many functional neuroimaging (or behavioral) studies.

Many studies use designs that logically permit inferences about the localization of a cognitive operation and/or about the nature of a cognitive functional architecture but are empirically weak because the materials and tasks are not created in ways that meet the “necessary operation requirement” and the “exclusive operations assumption.” As noted in the introduction, one source of progress in the field is due to researchers recognizing potential problems with existing studies and undertaking new studies to come closer to meeting these conditions. The challenge of meeting these conditions is not restricted to functional neuroimaging studies but applies to studies using behavioral measures as well. The discussion of designs in the text does not focus on whether the studies used as illustrations meet these conditions (brief comments on this are added as footnotes) but on the relation of the experimental condition and baseline condition.

We now turn to different designs. Each is related to a functional architecture. For each design, two issues are considered: (a) whether the choice of experimental condition(s) and the baseline condition(s) allows a cognitive operation or component to be related to a neural structure, and (b) how the neuroimaging results might bear on the model that underlies the study. Each design is illustrated with an example drawn from the literature.

EXPERIMENTAL DESIGNS

i. Designs Related to Functional Architecture I: Serial Cognitive Operations or Components

The first type of design consists of a contrast in which the operations of the experimental condition and baseline condition are thought to occur in serial. This is the simplest case, and is relevant to other contrasts to be reviewed below.

Studies of this sort are subject to the Immediate Sequentiality Principle (recall that the “necessary operation requirement” and the “exclusive operations assumption” have to be met).

Immediate sequentiality Principle

If the relation between the experimental condition and baseline condition that underlies the experimental condition/baseline condition contrast is such that

- a. a representation is computed in the experimental condition by an operation that requires and immediately follows the operations required in the baseline condition; and
- b. the experimental condition does not require operations later in the cognitive architecture, then
- c. differences in neural parameters between the experimental condition and baseline condition can be inter-

preted as reflections of brain processes that support this additional operation.

Requirement (a)—that the experimental condition and the baseline condition differ in one immediately sequential operation—is critical: if the operation unique to the experimental condition is separated from the operations of the baseline condition by many intervening cognitive processes, the difference cannot be clearly interpreted because it could be due to any of several operations. For instance, comparing viewing written words (experimental condition) against viewing randomly arranged parts of letter shapes (baseline condition) cannot isolate word recognition because, among other things, a number of cognitive operations occur in the experimental condition and not the baseline condition (e.g., letter identification).

An obvious point, but one that needs to be made explicit, is that a design that meets these requirements provides no information about the neural basis of the operations required in the baseline condition; in particular, it does not show that they do not take place in the area that was activated in the experimental condition/baseline condition contrast. What is shown is that the operation(s) unique to the experimental condition is supported by that area; if the completeness assumption is made, what is shown is that the operation(s) unique to the experimental condition is supported only by that area.

Zatorre et al. [1992] provides an example of a study that was based on a serial model. The goal of the study was to identify areas of the brain that supported phonemic processing and acoustic processing. These authors had participants listen to pairs of syllables and perform one of two tasks as experimental conditions: indicate either whether the two items ended with the same speech sound (a phonemic task) or whether the pitch of the second was higher than that of the first (an acoustic task). The baseline condition for both experimental conditions consisted of presenting the same stimuli as in the experimental condition and having participants press a response key when a stimulus was heard (“passive listening”). The status of the second syllable (same/different) was orthogonally varied with respect to phoneme and pitch in a Latin Square design.

Zatorre et al. found that the phonological judgment experimental condition minus the passive listening baseline condition activated left inferior frontal gyrus, and that the pitch judgment experimental condition minus the passive listening baseline condition activated the right inferior frontal gyrus. The authors say “Our results, taken together, support a model whereby auditory information undergoes discrete processing stages, each of which depends upon separate neural substrates (p 848).” They conclude that the lateralized frontal activations reflected processing of the auditory signal for speech or nonspeech properties. The operation that was isolated was described in somewhat different terms in the text and the Abstract of the article. In the text, the authors say “when a phonetic decision is required, there is a large focus of activity in Broca’s area . . .

right hemisphere mechanisms appear to be crucial in making judgments related to pitch (p 848).” In the Abstract, they claim that the activation in the areas seen in the experimental condition/baseline condition contrast reflects “discrimination of phonemes” and “pitch perception.” The description in the text suggests that the contrast isolates the *comparison* of phonemes and pitches that have been identified (in which case, the baseline conditions must require operations that identify phonemes and pitches). The description in the Abstract suggests that it isolates the perceptual identification of phonemes and pitches (in which case, the baseline condition must not require these operations). This ambiguity is due to the violation of the “necessary operation requirement” that results from using a baseline condition that does not require a response, which makes it impossible to know how participants processed the stimuli in the baseline condition.⁷ However, both of these conclusions are consistent with one or another serial model in which the operations in the baseline conditions precede an operation in the experimental conditions.

Turning to the second goal of functional neuroimaging of cognition—to provide data regarding cognitive processes—the demonstration that the experimental condition/baseline condition contrast activates a brain region is consistent with the serial model that maintains that the operations of the experimental condition consist of those of the baseline condition plus one more operation. It does not prove that the model is correct. It is possible that the oper-

⁷To clarify this issue, Zatorre et al. ran a fourth condition in which participants heard noise bursts with spectral qualities similar to the syllable stimuli and pressed alternate keys on alternate trials. They used this as the baseline condition for the previous baseline condition; that is, they subtracted this “noise” condition (the new baseline condition) from the “passive listening” condition (now the experimental condition in this new contrast). Again, they assumed a serial model of processing, in which the experimental condition is assumed to require all the operations required by the baseline condition (perception of properties of spectral noise that mimic those of speech) and one additional operation (the use of such spectral properties to activate speech sounds or pitch levels). This new contrast resulted in BOLD signal in left and right superior temporal gyri and left inferior frontal gyrus. The authors suggest that the activity in STG was due to “higher order auditory processing of complex signals (p 847)” and that in L IFG due to semantic processing and therefore concluded that phonemic and pitch perception depended on STG. Unfortunately, there is no way to know if these interpretations are correct, because there is no way to know whether the participants processed the speech stimuli semantically, and also because it is not clear that the “higher order auditory processing of complex signals” that putatively occurred in the STG in the baseline condition consists of the same phonemic and pitch processing as in the experimental conditions. The example is chosen because, despite the ambiguity of the study, Zatorre et al. do endorse a serial processing model and particular operations that their study may localize, and because other studies I know of that use this design are also subject to some degree of interpretive ambiguity [e.g., Scott et al., 2000].

ations occur in cascade (see below), or that the operations in the experimental condition and the baseline condition do not differ qualitatively but that those of the experimental condition require more processing than those of the baseline condition. Designs that contrast what are thought to be sequential operations can, however, provide strong disconfirming evidence for the serial model in the form of greater activation for the baseline condition than the experimental condition (sometimes called “relative deactivation”).

There are two types of relative deactivation. In the first, relative deactivation occurs with a decrease in BOLD signal from a baseline that can be interpreted as reflecting overall BOLD signal activity not affected by a specific task (such as prestimulus BOLD signal level, mean BOLD signal, or some other level). In this case, it can be due to inhibition of the area in which it occurs by the activity in the area(s) in which the experimental condition/baseline condition contrast produces a positive effect [Alison et al., 2000]. Second, relative deactivation can occur in association with an increase in BOLD signal from a neutral baseline. In this case, no such explanation is possible, and the model that underlies the study is contradicted. If the experimental condition involves all the operations required by the baseline condition plus one subsequent operation, there should be no areas of the brain in which there is an increase in BOLD signal relative to a neutral baseline and greater activity in the baseline condition than in the experimental condition. Finding that there is more activity somewhere in the brain for the baseline condition than the experimental condition disconfirms the model that underlies the design. This is true even if the completeness assumption is not made: if BOLD signal increases reflect increased processing load and the model is correct, the additional operation in the experimental condition compared to the baseline condition will exert a demand and no area should be more active in the baseline condition than in the experimental condition.

A relative deactivation (associated with an increase in BOLD signal from a baseline) does not, in and of itself, reveal the locus of the error in the model. The error may lie either in the model of the processes required in the two tasks or in the model of their relation. This is similar to disconfirming behavioral data, which require additional studies to determine which aspects of a model of cognitive function need to be changed.

A very large number of studies do not report relative deactivations, and do not indicate that they were not present in the BOLD signal analysis. This lack of attention may reflect a greater interest on the part of most researchers in localization of components of cognitive function than in the use of functional neuroimaging data to develop models of cognition. If so, ignoring relative deactivations is ill-placed, because such results provide evidence that the cognitive functions under study are not as well understood as was thought: either they are not used in tasks as envisaged, or the relationship between them is not that specified in the functional architecture. Note that relative deac-

tivations cannot be simply dismissed and ignored without denying the interpretability of the activations in the expected direction.

ii. Designs Related to Functional Architecture 2: Parallel Cognitive Operations or Components

In these designs, qualitatively different representations that are modeled as being computed in parallel are contrasted directly. The representations in question may be computed at a single stage of processing or at different stages of processing. The critical feature of parallel models is that the output of neither condition is the input to the other; labeling one condition as the experimental condition and one as the baseline condition is a statement about how a contrast is set up in an experiment, not a statement of how the conditions are related in the model of cognitive function that is the basis for the experiment.

This design is subject to severe interpretive limitations. The comparison of an experimental condition in which one type of representation is processed compared to a baseline condition in which a second type of representation is processed in parallel can never identify the neural location of an operation or a component of the cognitive processing system, because it can never isolate an operation or a component of the system.

Figure 2 shows why this is the case. In a parallel architecture, the experimental condition and baseline condition both require some antecedent processes (A–E), and the experimental condition then requires one immediately succeeding process F and the baseline condition requires the immediately succeeding process G. In terms of cognitive operations, the subtraction of the baseline condition from the experimental condition is the difference between the operations F and G. But this difference is not a cognitive operation. By definition, F is the operation in the experimental condition that occurs immediately following and that requires the processing operations in A–E; thus F cannot be isolated by subtracting processing in A–E plus processing in G from the experimental condition. Including G in the baseline condition guarantees that F cannot be isolated. If F cannot be isolated this way, it follows that F plus any other cognitive operations cannot be isolated this way. The identical argument applies in reverse if G is the experimental condition and F the baseline condition.

Although designs of this type are invalid as a means of identifying the neural basis of components of the cognitive processing system, they are ubiquitous in the functional neuroimaging literature. Martin et al. [1996] is a typical study of this sort. The authors compared saying the names of pictures of animals to oneself (“silent naming”) against doing the same for tools, and *vice versa*. They found that the calcarine cortex was activated in the comparison of animals (experimental condition) against tools (baseline condition) and the left inferior frontal and the left middle temporal gyri in the comparison of tools (experimental condition) against animals (baseline condition). It is not

clear whether the authors believe they localized the neural substrate for perceptual identification (“the brain regions active during object identification” p 649) or the neural basis for processing “semantic representations of objects” (p 652). However, they believe that, whichever of these processes are responsible for the neurovascular differences across the conditions, these processes differ for the categories of tools and animals, and these different processes or representations are supported by the areas of the brain identified in the contrasts. They say “the brain regions active during object identification are dependent, in part, on the intrinsic properties of the objects presented” p 649 and “areas recruited as part of this [semantic] network depends on the intrinsic properties of the object to be identified” p 652).⁸

Applying the considerations discussed in general above to perceptual identification of tools and animals makes it clear why these conclusions are invalid. Consider the suggestion that the experimental condition/baseline condition contrast isolated perceptual identification of tools and animals, and let us adopt a simple model in which visual perceptual identification is seen as a process in which geons are activated from features and mapped onto structural descriptions. The features and geons for tools and animals

⁸Many issues arise regarding the “soundness” of this study. Differences in silent naming of pictures of animals and tools could have arisen at many stages of processing—early visual processing, perceptual identification, feed-forward/feed-back loops between perceptual identification and semantic activation, phonological activation, or others. The pictures were taken from Snodgrass and Vanderwart [1980]. The Snodgrass and Vanderwart pictures of tools tend to contain straight lines and discrete angles; those of animals have curves and irregularities. Making the reasonable assumption that calcarine activation in part reflects early visual processes, these differences could have been the source of the category differences in this area. This possibility is not addressed by Martin et al.’s finding that the differences persisted in silent naming of silhouettes, since silhouettes of items in the two categories differ in the same way as pictures. The categories also differ in the types of geons they activate. The pictures were not matched for their likelihood to be categorized at the basic or subordinate level of their categories. If animals were categorized more often at the subordinate level than tools, this could have led to increased visual processing [Jolicoeur et al., 1984] and increased calcarine activity. Stimuli in the two categories were matched for word length of responses (on a pretest), but not for word frequency or name agreement (the number of different first responses). Silent naming of low compared to high frequency words leads to increased activity in L IFG [Kuo et al., 2003]. Differences in name agreement would affect selection of a single name if participants interpreted the instructions as requiring them to find not just any name but the most appropriate name, and this affects activation of L IFG [Kan and Thompson-Schill, 2004]. Even if the items in the categories had been matched on these—and all other conceivably relevant dimensions - in a pretest, the absence of overt responses would lead to uncertainty about whether the responses of the participants in the scanning environment were balanced along these dimensions. See also Devlin et al., 2002, for comments on this study.

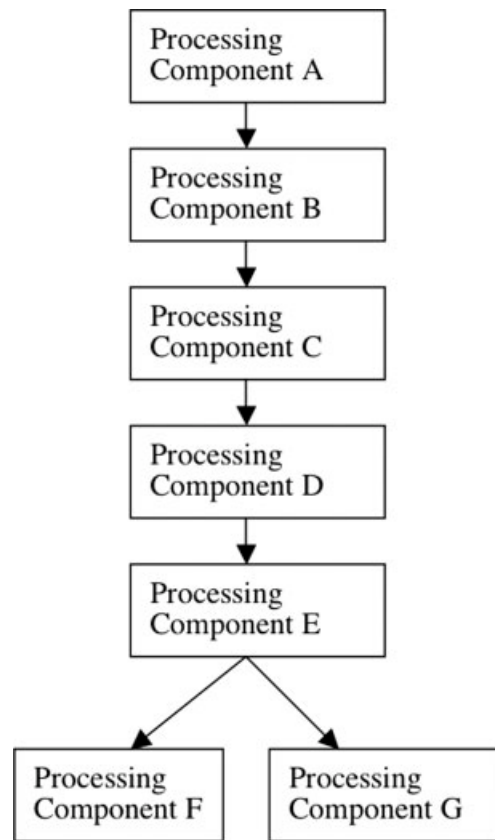


Figure 2.

Details of Figure 1, Functional Architecture 2. More detailed diagram of a functional architecture in which two components operate in parallel.

(tend to) differ, and therefore the operations that apply in perceptual identification of tools and animals (tend to) differ. Higher-order perceptual components may develop in which features are mapped onto geons or geons onto structural descriptions separately for animals and for tools. Nonetheless, though the separate representations, operations, and components involved in perceptual identification of tools and animals are each real entities in a cognitive model, the *difference* between any two representations, operations, or components is not an entity in the functional architecture underlying perception of either animals or tools in any model. It is not even clear what the difference between a representation activated in the course of perceiving an animal and one activated in the course of perceiving a tool would be: what is the difference between two features, geons, or structural descriptions? This type of consideration applies to any operations or components that are modeled as occurring in parallel in any cognitive functional architecture.

Because an experimental condition/baseline condition contrast of qualitatively different representations that are assumed to be computed in parallel cannot isolate a cogni-

tive operation, an increase in BOLD signal in such a contrast does not imply both that any operations unique to the experimental condition are supported in the area(s) that was activated in the contrast and that operations unique to the baseline condition are not supported by that area(s) (where “and” is interpreted inclusively; that is, it is not the case that these statements are both true). In fact, because the difference between the operations in the experimental condition and the baseline condition is not an entity in any cognitive theory, all of the operations involved in both conditions might bear *any* relation to the brain: they might be localized in one area, multifocally instantiated, or distributed throughout a wide area of the brain. This is the case even if we make the completeness assumption, because operations unique to the baseline condition may exert equivalent demand to those exerted by operations unique to the experimental condition in any subset of the areas in which operations unique to the experimental condition exert demand.

Despite the fact that the contrast of two conditions in which representations are processed at a parallel stage of a functional architecture can neither isolate a component of the cognitive processing system nor identify the neural locus of the operations of a component of the system, it still may be the case that the pattern of neural activation in such a contrast can provide data relevant to models of cognitive processing. In a review of an earlier version of this article, Rik Henson expressed such a view, suggesting that “Martin et al.’s findings [are] still relevant to cognitive theorizing. The fact is that they find a dissociation between brain activity relating to animals and tools (no matter how the activations are interpreted) that is predicted by some theories and not others . . . In other words, we still have another piece of evidence that is relevant for distinguishing competing cognitive theories.”

There is one aspect of models of cognitive function for which this is the case. On the assumption that increased BOLD activity reflects increased processing load, a “one way” effect of a contrast of two operations that occur in parallel—finding that there is an increase in BOLD signal in some brain region when one condition is used as the experimental condition and the other as the baseline condition but not in the reverse subtraction—rules out the possibility that the operations that are unique to the baseline condition are more demanding than those unique to the experimental condition, and would choose between theories that differed on this point. In this respect, it has the same value as finding that reaction time is longer or accuracy lower for responses to one set of stimuli than to another; if neurovascular effects (or neural effects in general) are more sensitive than behavioral measures to differences in demand between (at least some) cognitive operations, this may be valuable in developing models of cognitive function. In addition, such a “one-way” effect permits the conclusion that the increased processing demand in the experimental condition is supported by the area of the brain in which the BOLD signal increase occurred.

Whether this is informative about the psychological reason for the increased processing demand depends upon the ability to “reason backwards” from the location of an activation to what caused it. This will be discussed below, in connection with designs based on contrasts of two serial operations that occur in parallel.

Neuroimaging findings in studies with this type of design cannot, however, provide evidence relevant to deciding between two models, one of which claims that the operations that occur in parallel are qualitatively different (i.e., that the operations unique to the experimental and baseline conditions are part of separate components), and one that claims that these operations are qualitatively similar (i.e., that the operations unique to the experimental and baseline conditions are part of the same component). To be concrete, the data in Martin et al. cannot distinguish between a theory that claims that perceptual identification of animals is qualitatively different—i.e., is a separate component—from perceptual recognition of tools and a competing theory that maintains that perceptual identification of animals is accomplished by the same types of operations as perceptual identification of tools.

Clearly, a “one way” effect of a contrast of two operations that occur in parallel does not distinguish between two theories that contrast in this manner. Such a result could occur if one set of operations was more demanding than another regardless of whether the sets were qualitatively different or not. A “two-way” effect of these contrasts—as in Martin et al., where both contrasts lead to increases in BOLD signal, in different brain areas—also does not distinguish between these two theories. This pattern is the counterpart in functional neuroimaging of double dissociations of deficits in two patients [see Henson, 2005; Shallice, 2003, for discussion]. In discussing this type of result, it is preferable to refer to the two conditions as C_1 and C_2 , since each is used as the experimental condition and as the baseline condition, in different contrasts. The finding that BOLD signal for C_1 is greater than for C_2 (“ $C_1 > C_2$ ”) in area X and that BOLD signal for C_2 is greater than for C_1 ($C_2 > C_1$) in area Y is logically compatible with four possibilities:

1. C_1 includes an operation not found in C_2 and the application of the set of operations common to C_1 and C_2 is more demanding for C_2 than for C_1 because of how these operations are applied to C_1 and C_2 . The increase in BOLD signal in area X associated with $C_1 > C_2$ is due to the unique operation in C_1 and the increase in BOLD signal in area Y associated with $C_2 > C_1$ is due to the increased demand on a common set of operations for C_2 compared to C_1 .
2. The opposite of (1): C_2 includes an operation not found in C_1 and the application of the set of operations common to C_1 and C_2 is more demanding for C_1 than for C_2 because of how these operations are applied to C_1 and C_2 . The increase in BOLD signal in area Y associated with $C_2 > C_1$ is due to the unique

operation in C_2 and the increase in BOLD signal in area X associated with $C_1 > C_2$ is due to the increased demand on a common set of operations for C_1 compared to C_2 .

3. Both C_1 and C_2 include unique operations. The increase in BOLD signal in area X associated with $C_1 > C_2$ is due to the unique operation in C_1 and the increase in BOLD signal in area Y associated with $C_2 > C_1$ is due to the unique operation in C_2 .
4. Neither C_1 nor C_2 have unique operations. The increases in BOLD signal in area X associated with $C_1 > C_2$ and in area Y associated with $C_2 > C_1$ are due to the differential partitioning of activation due to a single operation over a single area by C_1 and C_2 .

The fourth combination of possibilities was dismissed by Shallice [2003] as being internally contradictory, and it is if we consider a behavioral variable such as accuracy or RT. It is impossible for the same set of operations to apply to two different categories in a way such that their application to each category leads to longer RTs or lower accuracy than to the other. But this is not the case when a measure that reflects the location of neural tissue that supports that operation is considered. The distribution of activity in a neural net can be uneven following presentation of a stimulus, such that different parts of the net become more active when one type of stimulus is presented than another and other areas show the opposite pattern of activation, even when the operations that apply to the two sets of stimuli are the same and the entire net is in fact activated by the presentation of both types of stimuli [Cangelosi and Parisi, 2004; Devlin et al., 1998; French, 1998; Plaut and Shallice, 1993; Plaut, 2002]. Therefore, a two-way pattern of activation of contrasts of operations that occur in parallel is compatible with all possible functional architectures. Specifically, the pattern found in Martin et al. could result from perceptual identification of animals and tools sharing a common set of operations that are more difficult to apply to animals than to tools and there being an operation unique to perceptual identification of tools, or *vice versa*, or both categories having unique operations, or both categories having the same operations that are distributed in different ways over a single area of the brain depending upon whether an animal or a tool is being perceived.

iii. Designs Related to Functional Architecture 3: Two Serial Components Operating in Parallel

Questions of the sort posed by Martin et al. [1996]—whether the neural basis of two operations that are postulated to occur in parallel are supported by distinct brain areas—can be answered by an extension the first type of design considered above—the contrast of sequential operations—to two sequential contrasts. Since each contrast can reveal the neural areas associated with an operation or a representation, distinct activations associated with each

contrast provide evidence of separate neural substrates for each.

Martin et al. [1996] utilized this approach as well as the direct comparison of representations computed in parallel we have just reviewed. The two silent naming experimental conditions were compared to passive viewing of nonsense drawings. Silent naming of animals and tools both activated unique brain regions in these contrasts (these areas included the areas that had been activated when silent naming of animals was contrasted with that of tools, described above). We can conclude that the operations isolated in each contrast are supported by the areas that are activated in that contrast.⁹

Results of this sort have implications about cognitive functional architectures. They provide evidence that the operations that are identified in at least one of contrasts are, in fact, different from those in the second (i.e., that one of possibilities 1, 2, or 3 listed above is correct; possibility 4 is ruled out), and thus can add significantly to behavioral measures [such as different patterns of interference in dual task studies, e.g., Shah and Miyake, 1996]. If we assume that the two contrasts in Martin et al. [1996]—the experimental conditions of silent naming of animals and tools compared to the baseline conditions of passive viewing of nonsense drawings—differed only at a single stage of processing—say, perceptual identification—this result would imply that that stage of processing differs for animals and tools.

The use of results of this sort to draw these conclusions about functional architecture rests on a claim about the relation of neural responses to cognitive operations that Henson [2005] calls the “function to structure deduction” principle: “if conditions C_1 and C_2 produce qualitatively different patterns of activity in the brain, then conditions

⁹However, in Martin et al., 1996, we cannot tell what these operations are because of violations of the necessary operation requirement. The stimuli in the two experimental conditions differed from the stimuli in the baseline condition in different ways. The nonsense object foils were taken from Kroll and Potter [1984], and are far more similar to the Snodgrass and Vanderwart pictures of tools than to the pictures of animals in terms of types of visual features (surfaces, straight lines, angles, geons). The absence of calcarine activation in the subtraction of nonsense figures from tools may thus reflect the fact that these stimuli are similar in terms of visual features. Second, the tasks—if carried out as the investigators intended - do not require immediately sequential stages of processing. If participants generated names for all the presented pictures, which of the multiple stages of processing that differ between each experimental condition and the baseline condition generated the increased neurovascular signal cannot be determined. The fact that the experimental conditions may have differed in different ways at different stages (cf, the multiple differences in linguistic and selection stages of processing between the animal and tool stimuli cited in footnote 8) is as detrimental to the interpretation of the contrast of these two experimental conditions against the nonsense object baseline condition as it was to the interpretation of the contrasts of the two experimental conditions against each other.

C_1 and C_2 differ in at least one function, F_2 . (p 197)" Page [2006], however, has presented three reasons to reject the "function to structure deduction" principle, which need to be examined.

Page's first objection is that the identification of the functions that drive an activation is often affected by what he calls "epiphenomenal activity," cognitive activity correlated with the function thought to be under study. This objection deals with how well the choice of experimental materials and tasks accomplish their goals, which is beyond the scope of this article. Page's two remaining objections are, however, relevant to the more general issues under consideration here. His second objection is that the principle is too strong. He illustrates with the example of high and low tones activating different parts of auditory cortex, which does not imply that perception (or any processing) of high and low tones involves different mechanisms. Third, he argues that the principle is either redundant with respect to its implications for cognitive function or its application is unwarranted. Page divides the cases of different activations in different conditions into two sets: one in which the conditions are associated with different behaviors, in which case the neural findings are redundant, and one in which they are not, which he takes as evidence that the same function can be performed in two different regions [cf, Noppeney et al., 2004, re: "degeneracy"].

There is an answer to Page's second objection: it is that the strength of the inference that separate activations imply separate functions depends upon the strength of the bridging inference that ties the functions to the brain [see Henson, 2005, for discussion along lines similar to what follows]. Consider the tonotopy example. The reason we reject the conclusion that different cortical areas of responsiveness to high and low tones implies different perceptual functions is that we prefer an alternative triplet of theories: the psychological theory that tone constitutes a dimension along which a unitary function—tonal perception—operates; the neural theory that the contiguous and cytoarchitectonically homogenous section of cortex along which responses to tones are spatially aligned is a unit of brain that might support a single function; and the bridging theory that relates elements of the two theories. The second theory maintains that the responses to high and low tones are not in separate brain areas but systematically arranged in a single area; the case is logically identical to the finding that responses to point changes in luminance are retinotopically mapped in V1, which would surely not lead anyone to suggest that there are multiple perceptual functions for detecting luminance changes in different parts of the visual field.

Page points out that testing this model requires cognitive operations be related to brain in spatially necessary ways. Page's point is correct, and, in the case of tonotopy, the requirement is met because we have an idea of why the bridging theory holds: it is because of the spatial arrangement of cortical projections of a series of neurons

whose first members are activated by hair cells in particular parts of the basilar membrane of the cochlea, a structure in which mechanical factors can be related to tonotopy.¹⁰

The issue of the strength of bridging inferences is also critical to answering Page's third argument—that separate activations for putatively separate operations is either redundant as evidence for separate processes (because behavioral data already establish the separation of the operations) or unwarranted (because behavioral data do not establish the separation of the operations). Page does acknowledge that there is one circumstance under which separate activations for putatively separate operations can be nonredundantly informative about cognitive architecture: "[An] imaging study [would show that] two types of process were operative in a task for which the behavioral evidence suggested no such decomposition ... [if] a function ... is shown to engage qualitatively different brain areas under circumstances in which both epiphenomenal activity and a one-to-many function-to-structure mapping can be ruled out." As noted above, controlling for epiphenomenal activity is a matter of choice of experimental materials and tasks and is not discussed here. This leaves the question of how to distinguish between two activations being a manifestation of two functional processes, each

¹⁰Page [2006] argues against the use of bridging arguments to inform models of cognitive function on the grounds that cognitive theory is not based on neural correlates of cognitive operations. Page points out that "as scanners become more sensitive to activation differences at smaller scales, then the number of hypothesized functions will proliferate. (p 430)" But, Page says, this "is not the way a cognitive psychologist attributes functions. Functions are hypothesized on the basis of a (potentially incorrect) theoretical analysis of the task at hand (p 430)." Page is right that it is a mistake to attribute the status of a function to every cognitive condition that has a different neural correlate. Doing so leads to an obvious *reductio*: assuming a materialist view of mind, a sufficiently sensitive way of identifying brain states will associate every individuated cognitive state with an individual brain state, so every mental state will be a function in some theory of cognition. But this is incorrect, because almost none of those different brain states will correspond to functions in cognitive theories; most will just be cognitive states of individuals. Cognitive states may be individuated neurally, but this will not immediately help develop models of cognitive functions because cognitive functions are *generalizations* over cognitive states. The most fundamental question that cognitive neuroscience needs to address is how cognitive functions correspond to generalizations over neural states. In relation to the issue of whether separate activations for high and low tones in auditory cortex imply separate processing operations for high and low tones, the question is whether the standard analysis of tonotopy and localization of cortical activity is an example of a mutually reinforcing confluence of models in psychology and neuroscience (whether the neural divisions map onto the psychological constructs—or, in this case, whether the *absence* of psychological divisions maps onto a *unitary* view of auditory cortex with respect to tonal perception), or whether it is an example of circular reasoning. The answer is that what we accept depends upon the strength of the bridging inference.

with a separate localization, and their being an instance of one-to-many function-to-structure mapping.

To make the question concrete, consider the Martin et al. [1996] result. Let us accept the following interpretation of Martin et al.'s finding: when effects of all preceding operations are removed, the perceptual identification of animals activates occipital cortex and that of tools L IFG and L MTG; that is, matching a visual stimulus to the long term representation of the physical properties of animals takes place, at least in part, in occipital cortex, and matching a visual stimulus to the long term representation of the physical properties of tools takes place, at least in part, in L IFG and L MTG. At the same time, there is no behavioral evidence for differences in the way stimuli belonging to these two categories are identified. Martin et al. take this as evidence for separate processes underlying perceptual identification of tools and animals; Page argues that there could be only a single mechanism that applies to both categories that is supported by several brain areas.

How can we decide between these alternatives? Given the premise that the evidence for two processes is solely neurological, an argument that is related to the neural findings—i.e., a bridging argument—is necessary. The neural finding is the difference in localization of perceptual identification of these two categories. What is therefore needed is an explanation of the correspondence between the processes underlying perceptual identification of the categories and the localizations. If such an explanation is forthcoming, the two-process model is supported. If not, the claim that this neural pattern results from a one-to-many-function-to-structure mapping is viable, perhaps preferred (due to the application of Occam's Razor).

Martin et al. proposed an account of the correspondence between the categories and the localizations—that the categories of animals and tools activate the regions they do because the operations involved in perceptual identification differ as a function of the sensory and motor features that individuate items within a category. Tool identification evokes nonbiological motion and affordances, which activate MTG and IFG respectively; identification of animals evokes biological movement, which activates occipital structures. The argument regarding the use of affordances in the perception of tools is quoted below (similar arguments are made regarding the mechanisms underlying location of activation in L MTG for tools and in occipital lobe for animals):

"The region of the left premotor cortex activated when subjects named tools, but not animals, was nearly identical to the area activated when subjects imagined gasping objects with their right hand ... Thus identifying tools may be mediated, in part, by areas that mediate knowledge ... of object use ... The capacity to identify a tremendous number of objects ... is dependent in part upon the automatic activation of previously acquired knowledge of the physical and functional attributes that define these objects." (pp 641–642)

This argument applies and extends a second principle enunciated by Henson [2005]: the "structure-to-function

induction" principle: "if condition C_2 elicits responses in brain region R_1 relative to some baseline condition C_0 and region R_1 has been associated with function F_1 in a different context (e.g., in a comparison of condition C_1 versus C_0 in a previous experiment), then function F_1 is also implicated in condition C_2 ", [Henson, 2005, p 198].¹¹ Martin et al. extend this principle to argue that adjacent regions can serve related functions. This sort of "sensory/motor extrapolation" on the basis of proximity [and, at times, connectivity: cf., Geschwind, 1965, re: the role of the inferior parietal cortex in lexical semantic representations] is a bridging theory.¹²

Once again, however, a principle enunciated by Henson has been called into question. Page [2006] objects to Henson's principle, and presumably would to the extension of the principle to contiguity and connectivity, for several reasons: (1) the basis upon which F_1 was originally associated with R_1 is often itself suspect; (2) the possibility that R_1 accomplishes two unrelated functions; (3) the possibility that C_2 requires the suppression of F_1 ; (4) the possibility of epiphenomenal activity in R_1 . These objections require answers if Henson's "structure-to-function induction" principle and Martin et al.'s model are to be accepted.

Objection (1) is obviously a potential problem but, hopefully, some studies have some merit. Objection (3) was dealt with briefly in footnote 3. As noted several times, objection (4) is important in practice but beyond the scope of this article.

Objection (2) is the inverse of Page's concern reviewed above that separate activations for separate contrasts may reflect one-to-many function-to-structure mapping. Objection (2) argues, correctly, that a single region may support more than one function and that something other than mere overlap (or adjacency) of activations is needed to show that the function that produced one activation in an area in one task is the same as the function that produced activation in a second task. Applied to Martin et al.'s argument regarding the role of affordances in perception of tools, the objection requires that we have some reason to

¹¹Poldrack [2006] refers to this as a "reverse inference." He discusses the strength of reverse inferences in relation to the probability that an area has been activated by a contrast that requires a particular type of processing, which can be related to complex probability maps of the involvement of a cortical area in a function. In my view, Poldrack's approach is valuable, but only if the studies upon which the estimation of the probability that an area supports a function are valid and sound; if not, the probability function will be heavily affected by results that are ambiguous and subject to misinterpretation.

¹²There is a long history of theory construction in neuropsychology in which adjacency arguments of this sort play an important role. The classic examples are ones in which association cortex adjacent to a koniocortical region is claimed to construct more elaborate modality-specific representations of the stimuli processed in that koniocortex; this view dates to the nineteenth century [Wernicke, 1874; see Geschwind, 1965, for a seminal modern statement] and thus antedates physiological evidence in its favor.

believe that the proximity of the regions activated when tools are identified and when subjects imagine tool use implies that representation of tool use is involved in tool perception. More generally, it requires something other than identity, overlap, proximity, or connectedness to accept that an attested function of one area is part of a second function. One way to appreciate the need for such evidence is to consider that other “proximity” arguments could be made that connect the activation of L IFG in the task involving tools to other functions. Perhaps L IFG was activated not because subjects activated representations of tool use when they perceptually identify tools but because of a role of L IFG in response selection, which may differ for the words for tools and animals.

There are possible approaches to this question. One that has never been explored, to my knowledge, is the following. Henson’s “structure-to-function induction” principle predicts that changes in the stimuli that produce an activation along the dimension that is hypothesized to explain the activation will produce correlated changes in neural responses in the area of activation, and changes in the stimuli along other dimensions will not. In terms of Martin et al.’s hypothesis, increases in the complexity of the affordances of tools should lead to increased BOLD signal in L IFG and increases in the complexity of response selection for tools should not (see below for discussion of the use of parametric designs to investigate a single processing component). Thus, Page’s concern can be addressed.

To summarize the main points of this section, contrasts of two sequential operations that operate in parallel can be informative about the neural location of the operations that differ in the two experimental condition/baseline condition contrasts. Accordingly, such designs can provide information relevant to cognitive architectures. As I have emphasized, and as Page [2006] recognizes, results of this type can provide information about the separation of processes that may not be as easily available from behavioral data. Assuming the subtractions are themselves unambiguously interpretable, the degree to which the implications for both functional localization and functional organization are convincing depends on the strength of the bridging inference connecting the functions and aspects of the neural activity, in most cases its location.

iv. Designs Related to Functional Architecture 4: Cascade Models

Cascade models (Fig. 1, Functional Architecture 4) are similar to serial models in that they postulate that there are operations in the experimental condition that follow those in the baseline condition. They appear to be similar to parallel models in that they postulate that operations in the experimental condition go on at the same time as operations in the baseline condition, but the relation between the operations in the experimental condition and the baseline condition is different in the case of cascade and parallel models. In the parallel models, the output of neither condition

is the input to the other; in cascade models, the operations in the baseline condition provide input for those of the experimental condition—they do so continuously and the operations of the experimental condition can begin before the operations of the baseline condition are complete.

The effects of experimental manipulation of factors that affect processing stages in a cascade model on behavior are complex. As noted, the hallmark of a cascade system is taken to be an interaction of factors that affect the hypothesized processing stages on a task that requires the late processing stage. These interactions can either be super-additive [see McClelland, 1979, pp 314–316, for an example in which a cascade model simulates a super-additive interaction of visual masking and context relatedness in lexical decision] or under-additive [see Humphreys et al., 1999, for an example of a cascade model that explains an under-additive interaction of perceptual similarity and word frequency in picture naming]. In addition, as noted above, not all cascade models generate interactions [McClelland, 1979]. Thus, the predictions of a cascade model for BOLD signal response must be made on a model-by-model basis.

At one level, however, these uncertainties do not affect the implications of a design in which the baseline condition and the experimental condition are modeled as operating in cascade. If the operations of a baseline condition and an experimental condition occur in cascade, the experimental condition/baseline condition contrast can isolate the processing operation or stage of processing unique to the experimental condition, just as is the case in serial models.

In addition, cascade models can generate predictions about interactions of factors that affect the experimental condition and the baseline condition, and can therefore be supported by finding such interactions. If operations occur in cascade, and the cascade model predicts an interaction of stimuli that are varied orthogonally along dimensions that are hypothesized to affect the baseline condition and the experimental condition, activity in the area that supports the operation unique to the experimental condition will show an interaction of these variables. We can capture this as follows:

Cascade principle

If the relation between the experimental condition and baseline condition that underlies the experimental condition/baseline condition contrast is such that

- a. a representation is computed in the experimental condition by an operation that operates in cascade with those of the baseline condition; and
- b. the cascade model generates an interaction of the factors that affect the experimental condition and the baseline condition in tasks requiring the experimental condition, and
- c. the experimental condition does not require operations later in the cognitive architecture, then

- d. an area that supports the operations unique to the experimental condition will show a main effect of the experimental condition/baseline condition contrast and an interaction of factors that affect the baseline condition and the experimental condition.

Unfortunately, to my knowledge, there are no neuroimaging studies that are based on, or that are interpreted as providing evidence for, cascade models. However, it may be useful to describe what such a study might look like.

Consider a study by Humphreys et al. [1999]. These authors reported that the perceptual similarity of objects in a category interacted with name frequency in determining reaction times in a naming test. They simulated this interaction in a cascade model in which perceptual identification operated in cascade with lexical retrieval. One could study the neural basis for this cascade model by orthogonally varying the perceptual similarity of objects in a category and the frequency of the names of objects in two tasks: one that require perceptual identification (Humphreys et al. recommend a properly constructed object decision task) and one that requires lexical retrieval (e.g., picture homophone matching). Areas activated by the contrast of perceptually similar and dissimilar stimuli in the object decision task would be interpreted as ones that support perceptual identification, and areas activated by the contrast of low/high name frequency in the homophone matching task would be interpreted as ones that support lexical retrieval (see discussion of studies of a single processing stage, below). If the model is correct, the areas activated by the contrast of low/high name frequency in the homophone matching task should also be activated by the contrast of the homophone matching task used as an experimental condition against the object decision task used as a baseline condition. The critical prediction of the cascade model is that there will be an interaction of perceptual similarity and name frequency on BOLD signal in these areas in the contrast of the homophone matching task used as an experimental condition against the object decision task used as a baseline condition. If the behavioral data are replicated, the effect of frequency will be less in categories with perceptually similar items than in categories with perceptually dissimilar items. Such results would support a cascade model.

v. Designs Related to Functional Architecture 5: Serial Models With Feedback

Feedback from a later to an earlier component of a cognitive architecture (Fig. 1: Functional Architecture 5) would lead to an increase in the BOLD signal in the area of the brain that supports the earlier component. An example is found in Price et al. [1996] in a study of the effects of naming on early visual processing. Price et al. [1996] had participants (a) name objects; (b) say “yes” when shown the same objects; (c) name the color of a nonobject;

and (d) say “yes” when shown the same nonobjects. Price et al. [1996] argued that subtracting the “say yes” conditions ((b) and (d)) from the “name” conditions ((a) and (c)) “reveals the areas common to object and color naming” (p 1505). These included left occipital, lingual, and fusiform gyri (among other areas). According to the authors, “the enhanced activation in these visual regions when subjects name coloured objects or the color of shapes (relative to viewing the same stimuli) implies that there is modulation of visual form and color processing when identification processes are required (p 1505)”. The [“name” – “say yes”] effect was greater for objects than for colors in bilateral medial anterior temporal lobe, left superior temporal sulcus, left posterior temporal lobe, left anterior insula, and right cerebellum. In discussing these results, Humphreys et al. [1999] say that the greater effect of naming compared to recognizing objects than of naming compared to recognizing colors is due to “the greater visual differentiation needed for naming compared to recognition . . . and to the greater visual differentiation required when natural objects are named . . . it may well be that these relative posterior areas are activated in a top-down fashion to enable the visual differentiation necessary for naming to take place” (p 127). To make this argument, it is necessary to show that the areas that increased their activation in the naming—recognition tasks were activated by the stimuli in the recognition tasks.¹³

vi. Designs Related to Functional Architecture 6: Interactive Parallel Operations

Just as cascade models are more complex versions of serial models, models that postulate parallel operations that interact are complex variants of models that postulate parallel operations. In these models (Fig. 1, Functional Architecture 6), the stages that occur in parallel also interact with each other. The addition of this feature does not change the conclusion that the use of a contrast of conditions that require parallel operations cannot isolate a component of the cognitive processing system. The implications of results of such subtractions for the nature of cognitive operations are also the same as for noninteractive parallel models, discussed above. We note that this model does lead to the empirical prediction that there will be interactions of the parameters that affect the “experimental condition” and “baseline condition” when the two conditions are contrasted.

vii. Designs Investigating a Single Processing Component (Functional Architecture 7)

The final set of designs I shall consider are ones in which a single processing component is studied. In these

¹³These conclusions are weakened by the fact that the “say yes” task is not an adequate test of recognition, because it violates the necessary operation condition; responses could have been made on the basis of detection of luminance changes.

designs, unlike all those considered thus far, the functional architecture underlying the study models the experimental condition and baseline condition as both containing representations of the same type and requiring the same types of operations. Designs of this sort are common in experimental psychological studies using behavioral measures, in which independent variables are often varied in ways that either increase the demand made by the application of a type of operation (e.g., increasing the number of distracters in a visual search task; presenting lower frequency words in a lexical decision task) or that require additional operations of a given type (e.g., adding a syntactic element such as a complement clause to stimulus sentences in a sentence comprehension task).

There are two types of designs of this sort. In the first, the representations differ qualitatively, though both belong to the same representational class on some theory. The validity of these designs is identical to that of designs in which qualitatively different representations belong to different representational classes. In the second, the experimental condition and baseline condition differ parametrically along a dimension that is relevant to how a single type of operation affects the processing of a single class of representations.

a. Qualitative contrasts

In this design, the experimental condition contains a representation, and requires an operation, that is not found in the baseline condition. In one version of this design, the model underlying the study postulates that the experimental condition requires the operations of the baseline condition and an additional operation. Such designs are similar to those for serial models involving immediately sequential different representations discussed above insofar as the model of the tasks holds that the experimental condition requires the operations of the baseline condition and one additional operation.

An example of this type of is found in Ben-Shachar et al. [2004], who contrasted each of two types of sentences that contained embedded questions (the experimental conditions) with a third that did not (the baseline condition) in a verification task. For both contrasts, the authors reported increased BOLD signal in the left inferior frontal gyrus, the left ventral precentral sulcus, and bilaterally in superior temporal gyrus (marginally on the right). They concluded that the L IFG, the left ventral precentral sulcus, and bilateral STG support the operation found in the sentences in the experimental and not the baseline conditions (in their view, processing certain noun phrases that have been moved from underlying to surface positions in sentences).¹⁴

In the second version of this type of design, the operations in the experimental condition and baseline condition

differ, and do not require each other. An example from syntax might be comparing BOLD signal responses to sentences with reflexive and nonreflexive pronouns (e.g., *The daughter of the queen washed herself/her*). These designs contrast parallel operations and are not valid as a means of isolating a cognitive operation, as discussed above.

b. Parametric contrasts

In this type of design, the experimental condition and baseline condition differ along a dimension that can be parameterized in some fashion. There are many ways in which this can occur. The experimental condition can differ from the baseline condition in requiring more memory or more computations, in requiring an additional operation of the same type as the baseline condition, or in other ways. In such designs, any difference in BOLD signal is attributed to the complexity difference between the conditions. We have encountered an example of this type of design in the gedanken experiment outlined above in connection with cascade models, in which perceptual similarity of items in a category was a parameter that determined the processing demand exerted at the stage of perceptual identification and name frequency was a parameter that determined the processing demand exerted at the stage of lexical retrieval. A real example consists of a study of syntactic processing by Just et al. [1996]. These authors used stimuli of the following type in a verification task:

4. Subject–object sentences.
The policeman who the guard admired congratulated the officer.
5. Subject–subject sentences.
The policeman who congratulated the officer admired the guard.
6. Conjoined sentences.
The policeman congratulated the officer and admired the guard.

They found progressively larger increases in BOLD signal for these sentence types—(4) > (5) > (6)—bilaterally in the posterior inferior frontal gyri and posterior superior temporal gyri, with greater BOLD signal on the left than on the right in both areas. They concluded that these areas supported the working memory demands of parsing and interpreting these sentences [Just and Carpenter, 1992].

The design of this study is parametric because the operations required in sentences (4)–(6) are viewed as identical (on some models of parsing and sentence interpretation) and the differences between how these sentences are processed are a result of how these operations are applied. In Gibson's [1998] model of parsing, for instance, (4) requires both more "storage costs" and more "integration costs" than (5) or (6), and (5) requires more "storage costs" than (6), and the combination of "storage" and "integration" costs both add to the "working memory" requirements of parsing and interpretation. The design is parametric:

¹⁴The design is unsound because the sentences differ semantically as well syntactically [see Caplan, 2006b].

increases in working memory load are varied in ways that could be used to systematically generate additional stimuli that varied along the same lines.

Regions of the brain in which this type of parametric variation leads to increases in BOLD signal can be interpreted as ones that support the demands made by the different values of the parameters that are used in a study. If these parameters affect a single processing component, this area of the brain supports that component. The neural responses can also provide disconfirming evidence regarding the model—no brain areas should show noninhibitory relative deactivations if the model is correct. The problems with these designs center on the extent to which the psychological process affected by varying the parameters can be isolated. In particular, issues regarding epiphenomenal activity are a concern (see Caplan et al., in press, for discussion).

CONCLUDING COMMENTS

I have argued that two types of simple designs can identify areas of the brain that support a function—designs based upon a cognitive functional architecture in which the experimental condition contains a single operation not found in the baseline condition and designs that parametrically vary representations of the same type within a task. The first of these designs can be applied both when the operation unique to the experimental condition is thought to be part of the same processing component as the operations in the baseline condition and when it is thought to be part of an immediately sequential processing component. It is possible to apply these designs in complex experiments that examine more than a single contrast to study the neural correlates of operations and components of functional architectures in which two or more sequences of operations are modeled as applying in parallel, to cascade models, and to models that include feedback from later to earlier components. If the conclusions in the article are correct, a simple first step in evaluating a functional neuroimaging study is to specify the functional architecture that underlies the study and to ask what operations in that functional architecture are contrasted in the experimental and baseline conditions. Studies that contrast conditions that cannot isolate a cognitive operation cannot yield information about the neural basis of cognitive operations.

If designs incorporating interpretable contrasts are used, functional neuroimaging can provide unique information about functional neuroanatomy. Neuroimaging in intact individuals is the only method that can provide information about the neural areas and processes that are sufficient to support a function. Deficit-lesion correlations provide information about the areas that are necessary to support a function, not areas that are sufficient, and functional neuroimaging after lesions (which has to deal with the same design issues as functional neuroimaging in intact

individuals) can provide information about areas of the brain and neural processes that are sufficient to support a function after the areas that normally do so are damaged.

I have also argued that functional neuroimaging can contribute to modeling cognitive processes. Page [2006] and others have argued that the investment in this approach is counterproductive; that researchers whose interests are in cognitive models would be better served by using behavioral measures. In the fields I am familiar with, models of cognitive processes have been formulated and tested at a much greater level of detail in experiments using behavioral measures than in those using neuroimaging techniques—raising concerns about a “dumbing down” effect of functional neuroimaging on experimental psychology (Shallice, personal communication). In part, this has occurred because many details of many theories can only be tested through a series of experiments, which the expense associated with neuroimaging makes very hard to undertake. In addition, behavioral measures are more sensitive in some ways than neuroimaging measures. The biggest difference between neurovascular and behavioral measures is the temporal window that can be isolated. Measurements of eye movements [Traxler et al., 2002], hand movements [Spivey et al., 2005], and even RTs have shown that the time frame that is relevant to capturing effects of cognitive processes is on the order of deciseconds. Measures that reflect neural activity that occurs over longer time periods risk being influenced by cognitive activity that occurs before and after the cognitive operations of interest, no matter how well designed a study. However, this limitation does not apply to measures of physiological activity—EEG and MEG offer millisecond level resolution. This property is rarely exploited, as most analyses select large temporal windows for analysis, but it can be, and the integration of MEG, EEG, and fMRI may allow the temporal resolution of neural measures to surpass that of behavioral ones. As these and other barriers to more sophisticated use of functional neuroimaging fall, neural measurements may match or exceed behavioral measurements with respect to their sensitivity to cognitive events. One area that neural measures may have particular applicability is the ability to provide evidence that two operations are separate and operate in parallel, based on different areas of activation.

I will conclude with the observation that there is one important reason that cognitive theorists may be interested in functional neuroimaging: neural events cause cognitive events. A truly developed model of cognition would be able to derive the nature of cognitive functions in part from an understanding of neural events. This will require much greater understanding of bridging relations between cognition and neural events than we have now. As Page [2006] notes, even the weaker capacity for spatial prediction is not available for the interpretation of almost all current functional neuroimaging data. As noted above, to my knowledge, all bridging inferences now developed are variants on the idea that cognitive operations are localized

where they are because of their relations to sensory and motor functions. This clearly cannot be the only principle relating brain and mind: cognitive operations must be related to brain areas because of the cellular and molecular contents of those areas, but we know very little about how these features of the brain support and constrain cognitive operations. Developing principles that govern the relation of brain events to cognitive operations is admittedly a distant goal, but identifying the areas in which cognitive processes take place is one important step towards identifying features of cellular identity and organization that might support particular operations. This is one reason why it is worthwhile to persist with the functional neuroimaging enterprise, despite the many challenges it poses.

ACKNOWLEDGMENTS

I am greatly indebted to Gayle DeDe, David Gow, Gina Kuperberg, Lauren Moo, Jeremy Schmahman, Louise Stanczak, Gloria Waters, Rik Henson (a reviewer), and two anonymous reviewers for discussion of the issues raised here and comments on earlier drafts of this article, and to the editors of *Human Brain Mapping* for suggestions regarding exposition and style. Remaining errors are solely my responsibility.

REFERENCES

- Allison JD, Meador KJ, Loring DW, Figueroa RE, Wright JC (2000): Functional MRI cerebral activation and deactivation during finger movement. *Neurology* 54:135.
- Ben-Shachar M, Palti D, Grodzinsky Y (2004): The neural correlates of syntactic movement: Converging evidence from two fMRI experiments. *Neuroimage* 21:1320–1336.
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996): Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207–4221.
- Cangelosi A, Parisi D (2004): The processing of verbs and nouns in neural networks: Insights from synthetic brain imaging. *Brain Lang* 89:401–408.
- Caplan D (2006a): Functional neuroimaging of syntactic processing: New claims and methodological issues. *Curr Rev Med Imag* 2:443–451.
- Caplan D (2006b): Functional neuroimaging studies of syntactic processing in sentence comprehension: A critical selective review. *Ling Lang Compass* 1:32–47.
- Caplan D, Moo L (2004): Cognitive conjunction and cognitive analysis. *Neuroimage* 21(2):751–756.
- Caplan D, Waters G, Kennedy D, Alpert N, Makris N, DeDe G, Michaud J, Reddy A (2007): A study of syntactic processing in Aphasia II: Neurological aspects. *Brain Lang* 101:151–177.
- Caplan D, Chen E, Waters G (in press). Task-dependent and task-independent neurovascular responses to syntactic processing. *Cortex*.
- Caspers S, Geyer S, Schleicher A, Mohlberg H, Amunts K, Zilles K (2006): The human inferior parietal cortex: Cytoarchitectonic parcellation and interindividual variability. *Neuroimage* 33: 430–448.
- Caviness VS, Jr., Meyer J, Makris N, Kennedy DN (1996): MRI-based topographic parcellation of human neocortex: An anatomically specified method with estimate of reliability. *J Cogn Neurosci* 8:566–587.
- Coltheart M (2006a): What has functional neuroimaging told us about the mind (so far)? *Cortex* 42:323–331.
- Coltheart M (2006b): Perhaps cognitive neuroimaging has not told us anything about the mind (so far). *Cortex* 42:422–427.
- Dehaene S, Naccache L, Cohen L, Bihan DL, Mangin J-F, Poline J-B, Rivière D (2001): Cerebral mechanisms of word masking and unconscious repetition priming. *Nat Neurosci* 4:752–758.
- Devlin JT, Gonnerman LM, Andersen ES, Seidenberg MS (1998): Category-specific semantic deficits in focal and widespread brain damage: A computational account. *J Cogn Neurosci* 10: 77–94.
- Devlin JT, Russell RP, Davis MH, Price CJ, Moss HE, Fadili MJ, Tyler LK (2002): Is there an anatomical basis for category-specificity? Semantic memory studies in PET and fMRI. *Neuropsychologia* 40:54–75.
- Donders FC (1868): On the speed of mental processes, translation in Koster WG (ed) *Attention and Performance, II*. Amsterdam: North-Holland, 276–315.
- Frazier L, Clifton C (1996): *Construal*. Cambridge, MA: MIT Press.
- French RM (1998): A simple recurrent network model of bilingual memory. In: Gernsbacher MA, Derry SA (eds.) *Proceedings of the Twentieth Annual Cognitive Science Society Conference*, Lawrence Erlbaum NJ: LEA, pp 368–373.
- Friston K, Price C, Fletcher P, Moore C, Frackowiak R, Dolan R (1996): The trouble with cognitive subtraction. *NeuroImage* 4: 97–104.
- Geschwind N (1965): Disconnection syndromes in animals and man. *Brain* 88:237–294, 585–644.
- Geschwind N, Galaburda A (1987): *Cerebral lateralization: Biological Mechanisms, Association and Pathology*. Cambridge, MA: MIT press.
- Gibson E (1998): Syntactic complexity: Locality of syntactic dependencies. *Cognition* 68:1–76.
- Henson RNA (2005): What can functional imaging tell the experimental psychologist? *Quart J Exper Psychol A* 58:193–233.
- Henson R (2006a): Forward inference using functional neuroimaging: Dissociations versus associations. *Trends Cogn Sci* 10:64–69.
- Henson R. (2006b): What has (neuro)psychology told us about the mind (so far)? A reply to Coltheart (2006). *Cortex* 42:387–392.
- Howard MW, Addis KM, Jing B, Kahana MJ (2007): Semantic structure and episodic memory. In: McNamara D, Landauer T, Dennis S, Kintsch W, editors. *LSA: A Road to Meaning*. Mahwah NJ: Earlbaum, pp. 121–141.
- Humphreys GW, Riddoch MJ, Quinlan PT (1988): Cascade processes in picture identification. *Cogn Neuropsychol* 5:67–103.
- Humphreys GW, Price CJ, Riddoch MJ (1999): From objects to names: A cognitive neuroscience approach. *Psychol Res* 62: 118–130.
- Jack AI, Sylvester C, Corbetta M (2006): Losing our brainless minds: How neuroimaging informs cognition. *Cortex* 42:418–421.
- Jennings J, McIntosh A, Kapur S, Tulving E, Houle S (1997): Cognitive subtractions may not add up: The interaction between semantic processing and response mode. *NeuroImage* 5:229–239.
- Jolicoeur P, Gluck MA, Kosslyn SM (1984): Pictures and names: Making the connection. *Cognitive Psychology* 16:243–275.
- Just MA, Carpenter PA (1992): A capacity theory of comprehension: Individual differences in working memory. *Psychol Rev* 99(1):122–149.

- Just MA, Carpenter PA, Keller TA, Eddy WF, Thulborn KR (1996): Brain activation modulated by sentence comprehension. *Science* 274:114–116.
- Kahnemann D (1973): *Attention and effort*. New York, Englewood Cliffs: Prentice-Hall.
- Kan IP, Thompson-Schill SL (2004): Selection from perceptual and conceptual representations. *Cogn Affect Behav Neurosci* 4:466–482.
- Kanwisher N, McDermott J, Chun MM (1997): The fusiform face area: A module in human extrastriate cortex specialized for the perception of faces. *J Neurosci* 17:4302–4311.
- Kisvarday G, Beroukas D, North JB, Chubb IW, Somogyi P (1990): Synapses, axonal and dendritic patterns of GABA-immunoreactive neurons in human cerebral cortex. *Brain* 113:793–812.
- Klatt D (1979): Speech perception: A model of acoustic phonetic analysis and lexical access. *J Phonetics* 7:279–342.
- Kroll JF, Potter MC (1984): Recognizing words, pictures and concepts: A comparison of lexical, object, and reality decisions. *J Verbal Learn Verbal Behav* 23:39–66.
- Kuo WJ, Yeh T-C, Lee C-Y, Wu YT, Chou CC, Ho L-T, Hung DL, Tzeng OJL, Hsieh J-C (2003): Frequency effects of Chinese character processing in the brain: An event-related fMRI study. *Neuroimage* 18:720–730.
- Logothetis NK, Pauls J, Augath MA, Trinath T, Oeltermann A (2001): Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
- MacDonald MC, Christiansen MH (2002): Reassessing working memory: A reply to Just & Carpenter and Waters & Caplan. *Psychol Rev* 109(1):35–54.
- Martin A, Wiggs CL, Ungerleider LG, Haxby JV (1996): Neural correlates of category-specific knowledge. *Nature* 379:649–652.
- Meyer DE, Schvaneveldt RW, Ruddy MG (1975): Loci of contextual effects on visual word-recognition. In: Rabbitt PMA (Ed.), *Attention and Performance V*, 98–118.
- Mazziotta J, Toga A, Evans A, et al. (2001): A probabilistic atlas and reference system for the human brain: International consortium for brain Mapping (ICBM). *Phil Trans R Soc London, Biol Sci* 356:1293–1322.
- McClelland JL (1979): On the time relations of mental processes: An examination of systems of processes in cascade. *Psychol Rev* 86:287–330.
- Mesulam M-M (1990): Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Ann Neurol* 28(5):597–613.
- Newman S, Twieg D, Carpenter P (2001): Baseline conditions and subtractive logic in neuroimaging. *Hum Brain Mapp* 14:228–23.
- Noppeney U, Friston KJ, Price C (2004): Degenerate neuronal systems sustaining cognitive functions. *J Anat* 205:433.
- Page MP (2006): What can functional imaging tell the experimental psychologist? *Cortex* 42:428–443.
- Plaut DC (2002): Graded modality-specific specialization in semantics: A computational account of optic aphasia. *Cogn Neuro-psychol* 19:603–639.
- Plaut DC, Shallice T (1993): Deep dyslexia: A case study of connectionist neuropsychology. *Cogn Neuropsychol* 10:377–500.
- Price CJ, Friston KJ (1997): Cognitive Conjunction: A new approach to brain activation experiments. *Neuroimage* 5:261–270.
- Price C, Moore CJ, Humphreys GW, Frackowiak RSJ, Friston KJ (1996): The neural regions sustaining object recognition and naming. *Proc Royal Soc B: Biol Sci* 263:1501–1507.
- Price CJ, Moore CJ, Friston KJ (1997): Subtractions, conjunctions, and interactions in experimental design of activation studies. *Hum Brain Mapp* 5:264–272.
- Poldrack RA (2006): Can cognitive processes be inferred from neuroimaging data? *Trends Cogn Sci* 10:59–63.
- Raichle ME, Fiez JA, Videen TO, MacLeod AMK, Pardo JV, Petersen SE (1994): Practice-related changes in human brain functional anatomy during non-motor learning. *Cerebral Cortex* 4:8–26.
- Recanzone GH, Merzenich MM, Dinse HR (1993): Expansion of the cortical representation of a specific skin field in primary somatosensory cortex by intracortical microstimulation. *Cerebral Cortex* 2:181–96.
- Schutter DJLG, de Haan EHF, van Honk J (2006): Scaling problems in the mind-brain conundrum. *Cortex* 42:411–413.
- Scott SK, Blank SC, Rosen S, Wise RJS (2000): Identification of a pathway for intelligible speech in the left temporal lobe. *Brain* 123:2400–2406.
- Shah P, Miyake A (1996): The separability of working memory resources for spatial thinking and language processing: An individual differences approach. *J Exper Psychol: General* 125:4–27.
- Shallice T (1988): *From Neuropsychology to Mental Structure*. Cambridge, UK: Cambridge University Press.
- Shallice T (2003): Functional neuroimaging and neuropsychology findings: how can they be linked? *NeuroImage* 20:S146–S154.
- Snodgrass JG, Vanderwart M (1980): A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *J Exper Psychol: Learn, Mem Cogn* 6:174–215.
- Sidtis J, Strother S, Anderson J, Rottenberg D (1999): Are brain functions really additive? *NeuroImage* 9:490–496.
- Spivey M, Grosjean M, Knoblich G (2005): Continuous attraction toward phonological competitors. *Proc Natl Acad Sci* 102(29): 10393–10398.
- Sternberg S (1969): The discovery of processing stages. *Acta Psychol* 30:34–78.
- Sternberg S (1998): Discovering mental processing stages: The method of additive factors. In: Scarborough D, Sternberg S, editors. *Invitation to Cognitive Science, Vol. 4: Methods, Models, and Conceptual Issues*. Cambridge, MA: MIT Press.
- Sternberg S (2001): Separate modifiability, mental modules, and the use of pure and composite measures to reveal them. *Acta Psychol* 106:147–246.
- Stevens K (2000): *Acoustic phonetics*. The MIT Press.
- Townsend JT, Ashby FG (1987): *The stochastic modeling of elementary psychological processes*. Cambridge, UK: Cambridge University Press.
- Traxler M, Morris R, Seely R (2002): Processing subject and object relative clauses: Evidence from eye movements. *J Mem Lang* 47:69–90.
- Umiltà C (2006): Localization of cognitive functions in the brain does allow one to distinguish between psychological theories. *Cortex* 42:399–401.
- Wernicke K (1874): *The aphasic symptom complex: A psychological study on a neurological basis*. Kohn and Wegert, Breslau; reprinted In: Cohen RS, Wartofsky MW, editors. *Boston Studies in the Philosophy of Science, Vol. 4*. Boston, MA: Reidel.
- Zatorre RJ, Evans AC, Meyer E, Gjedde A (1992): Lateralization of phonetic and pitch processing in speech perception. *Science* 256:846–849.