Brain-Mind and Structure-Function Relationships: A Methodological Response to Coltheart

Adina L. Roskies†‡

In some recent papers, Max Coltheart has questioned the ability of neuroimaging techniques to tell us anything interesting about the mind and has thrown down the gauntlet before neuroimagers, challenging them to prove he is mistaken. Here I analyze Coltheart's challenge, show that as posed its terms are unfair, and reconstruct it so that it is addressable. I argue that, so modified, Coltheart's challenge is able to be met and indeed has been met. In an effort to delineate the extent of neuroimaging's ability to address Coltheart's concerns, I explore how different brain structure-function relationships would constrain the ability of neuroimaging to provide insight about psychological questions.

1. Introduction. While many people err in thinking that functional magnetic resonance imaging is a key for unlocking every scientific door guarding the secret mysteries of human brain function, others err in claiming not only that neuroimaging can't tell us *everything* we wish to know about the mind, but (much more implausibly) that neuroimaging is unsuited to telling us *anything* about the mind. Max Coltheart is one of these skeptics. His interest is in learning about "cognition itself" (Coltheart 2006b). By this, I believe that he means learning about cognition independently of implementation details, such as where cognitive processes are instantiated in the brain. Coltheart claims that neuroimaging has not yet told us anything about cognition itself and doubts that it can. He is mistaken to think this. In understanding where he goes wrong, we gain a deeper understanding of the nature of mind-brain relations.

†To contact the author, please write to: Department of Philosophy, Dartmouth College, Hanover, NH 03755; e-mail: adina.roskies@dartmouth.edu.

‡This work was supported in part by an Australian Research Council fellowship at the University of Sydney. I would like to thank Max Coltheart for his correspondence and comments on an earlier version of this manuscript.

Philosophy of Science, 76 (December 2009) pp. 927–939. 0031-8248/2009/7605-0028$10.00
Copyright 2009 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/605815 Published online by Cambridge University Press

2. Coltheart's Challenge. Several critiques of neuroimaging are extremely deflationary (Poeppel 1996; Van Orden and Paap 1997; Uttal 2001), and others have expressed their skepticism (e.g., Fodor 1999), but Max Coltheart has provided the most focused and thought-provoking challenge to neuroimagers yet. While he refrains from making any in-principle claim about brain imaging, acknowledging that future studies could prove him wrong, he does make the following rather strong claim: "I'll claim that no functional neuroimaging research to date has yielded data that can be used to distinguish between competing psychological theories" (Coltheart 2006b, 323). Coltheart's provocative challenge to those who disagree is to provide examples of studies that disprove his claim. In a special issue of *Cortex* (2006, vol. 42), a number of neuroimagers take on this challenge, addressing the question of what, if anything, we have learned from neuroimaging. In Coltheart's rejoinder he attempts to refute the candidates on a case-by-case basis (Coltheart 2006a).

Coltheart's question is whether neuroimaging has ever given us reason to distinguish between two candidate psychological theories. Coltheart makes an effort to be quite clear about the terms of his challenge. First, he says, this is not an in-principle claim, but merely (at this stage) an in-practice one, concerning what we have learned from neuroimaging *so far*. Second, rather than providing an exhaustive review of the literature or arguing on a purely theoretical basis, his strategy is to refute purported counterexamples on a case-by-case basis. Finally, he tells us exactly what we would need to show to successfully respond to the challenge. We need to have two competing psychological theories, where these are theories "expressed solely at the psychological level." As he puts it,

> I must always begin by stating two or more psychological theories $\mathbf{T_a}$ $\mathbf{T_b}$ . . . concerning that domain. Then I can consider whether there has been any neuroimaging work that has yielded data which provides good reason to favour one of these theories over the others. So what's needed is to show that $\mathbf{T_a}$ predicts X whilst $\mathbf{T_b}$ predicts $\sim X$, where X is some pattern of neuroimaging data, and then to show that there exists functional neuroimaging work which demonstrates that X is the case or demonstrates that $\sim X$ is the case. (Coltheart 2006b, 324–325)

Coltheart also tells us what he means by "distinguish[ing] between theories." Recognizing that definitive evidence in science is hard to come by, he doesn't require that any result provide certainty about the truth of one theory and the falsity of another. Rather, his criteria are quite moderate: he is looking for data that make one of two theories a sounder bet. In other words, the challenge is to identify neuroimaging data that provide evidence in favor of one psychological theory rather than another.

https://doi.org/10.1086/605815 Published online by Cambridge University Press

The challenge appears to be stated with admirable clarity. However, I contend it is just clear enough to seem straightforward and just obscure enough to hide the fact that, as stated, the challenge is framed so that it can't be met. This becomes clearer when we consider what he means by a psychological theory.

As Coltheart says, a psychological theory is a theory of cognition "expressed solely at the psychological level." By this he means that the theory can refer only to psychological functions and processes, as well as inputs to, and outputs from, these functions. Any references to implementation details, such as where some process goes on or what mechanisms realize such functions, are not part of the psychological theory, for they depart from the psychological level. The reason why one might think that location or mechanism can't be a part of theories "expressed solely at the psychological level" is that implementation details are presumably irrelevant to function, at least as conceived according to classical and functionalist views of mind. Any theory, therefore, that includes references to location or mechanism would thereby fall outside the class of psychological theories. Thus, the predicates of $\mathbf{T_a}$ and $\mathbf{T_b}$ must be functional predicates. Let us call the functional predicates $F_1, \ldots, F_n$, the inputs (e.g., stimuli) $I_1, \ldots, I_n$, and outputs (i.e., behavior) $O_1, \ldots, O_n$. Thus, any theory that is expressed solely at the psychological level will be a theory expressed only in terms of *F*'s, *I*'s, and *O*'s.

Predictions that can be empirically tested must be consequences of $\mathbf{T_a}$ and $\mathbf{T_b}$: they must be derived from these theories. But if this is so, then the X and $\sim X$ of $\mathbf{T_a}$'s and $\mathbf{T_b}$'s predictions must be consequences of statements containing only *F*'s, *I*'s, and *O*'s as predicates. That is, the predictions of $\mathbf{T_a}$ and $\mathbf{T_b}$ must likewise be a function of only *F*'s, *I*'s, and *O*'s, referring only to inputs, outputs, and functions; no purely psychological theory could lead to predictions about implementation details. However, for Coltheart's challenge to be met, the psychological theories in question must predict an X, where X is "some pattern of neuroimaging data." Patterns of neuroimaging data are locations of regions of activation in the brain, by hypothesis those regions in which the mechanisms realizing the functions reside. But no pattern of location data is identical to some function. Therefore, there is no X such that X could be predicted by both a theory expressed purely at the psychological level and a pattern of brain activations. The way Coltheart's challenge is framed simply rules out the possibility that any neuroimaging experiment will satisfy his criteria.

Corroborating evidence that Coltheart relies on this way of framing the challenge can be found in his responses to a few of the neuroscientists that take on his challenge. For example, in his response to Umiltà (2006), Coltheart accedes that Umiltà has put forth two theories about the nature of visual attention, described solely at the psychological level:

https://doi.org/10.1086/605815 Published online by Cambridge University Press

> $\mathbf{T_a}$: Endogenous and exogenous attention are governed by a single cognitive system.
>
> $\mathbf{T_b}$: Endogenous and exogenous attention are governed by separate cognitive systems.

Umiltà argues that imaging has shown that endogenous attention activates two different brain networks and takes this to be evidence that supports $\mathbf{T_b}$ and not $\mathbf{T_a}$.

Coltheart argues that Umiltà has failed to show that imaging is relevant to cognition for two reasons. First, he claims that Umiltà "says nothing about the nature of these two theories at the psychological or cognitive level: nothing for example about the different functional architectures that the two theories propose" (Coltheart 2006a, 423). While perhaps this is strictly true, it is surely not true in spirit: while Umiltà's theories make no explicit predictions about the functional architecture of attention, it is clear that $\mathbf{T_a}$ suggests that the cognitive systems will share functional components, while $\mathbf{T_b}$ does not.

Coltheart argues that

> one can show that the two theories he [Umiltà] considers are not psychological because nothing in his paper would be changed if he had stated the two theories thus:
>
> $\mathbf{T_a}$: endogenous and exogenous attention are governed by a single brain system.
>
> $\mathbf{T_b}$: endogenous and exogenous attention are governed by separate brain systems. (2006a, 423)

He concludes that "These are theories about the brain, not about the cognitive level" (423). This argument may look plausible, but it is faulty. It is a little like complaining that the statement "Whales eat fish and live in the ocean" is really a statement about octopi, because it would have been about octopi if "whales" was replaced by "octopi." The original theories as stated by Umiltà are about cognition, and the reworded statements are about the brain. Coltheart wants us to conclude that studies like this one can be dismissed as irrelevant, because they fail to meet his criteria for what counts as a psychological theory.

Coltheart uses a similar strategy to reply to several others. He argues that Jonides' (Smith and Jonides 1977; Jonides, Dee, and Berman 2006) results don't support his theory $\mathbf{T_a}$, because his theory does not predict the results:

> The finding that different parts of the brain are associated with visual and verbal working memory is *not* support for $\mathbf{T_a}$ because $\mathbf{T_a}$ does not predict this result. . . . Suppose the imaging study had found

https://doi.org/10.1086/605815 Published online by Cambridge University Press

> that the same brain regions were involved in visual and verbal working memory: would that be inconsistent with $\mathbf{T_a}$? No. So all possible results of this study are compatible with $\mathbf{T_a}$ and so no result could have contradicted this theory. I think the same is true of the example offered by Vallar. (Coltheart 2006a, 423)

Coltheart continues:

> However, there is a pair of theories between which the results of Smith and Jonides (1997) *did* adjudicate, namely
>
> $\mathbf{T_a}$: working memory for verbal information is mediated by a different brain system than working memory for spatial information.
>
> $\mathbf{T_b}$: there is a single brain system of storage and rehearsal processes that works on both verbal and spatial information.
>
> The functional neuroimaging results supported $\mathbf{T_a}$ and conflicted with $\mathbf{T_b}$. But of course $\mathbf{T_a}$ and $\mathbf{T_b}$ here are not psychological theories: they are theories about the brain, and so not relevant. (423–424)

Thus, Coltheart rules out results as being relevant if brain information is taken to bear on psychological theory. This is an extreme position for anyone who takes psychological function to be realized in neural architecture in any systematic way.

I do not think Coltheart misses the point here. In another article, he states his radical position more straightforwardly. He calls this "the ultra-cognitive-neuropsychological position":

> The other possible aim of cognitive neuroimaging is to use imaging data for testing or adjudicating between cognitive models. Here the ultra-cognitive-neuropsychological position is a particularly extreme one: The assertion is that this aim is impossible to achieve in principle, because facts about the brain do not constrain the possible natures of mental information-processing systems. No amount of knowledge about the hardware of a computer will tell you anything serious about the nature of the software that the computer runs. In the same way, no facts about the activity of the brain could be used to confirm or refute some information-processing model of cognition. This is why the ultra-cognitive-neuropsychologist's answer to the question "Should there be any 'neuro' in cognitive neuropsychology?" is "Certainly not; what would be the point?" (Coltheart 2004, 22)

As it happens, I think Coltheart's claim about computers is mistaken; the same claim about the brain is surely mistaken. In summary, Coltheart admits that neuroimaging can distinguish between theories about the brain but denies that they can distinguish theories about the mind. The reason

https://doi.org/10.1086/605815 Published online by Cambridge University Press

that this is so is that no theory about the mind can mention the brain, so no data about the brain can bear on theories about the mind. It should now be apparent that the task Coltheart has set for the relevance of imaging is an impossible one. Because of his stipulation that the theories of interest must be formulated solely at the psychological level, no psychological theory can possibly predict any pattern of brain activation.

3. Coltheart's Challenge Revived. Although his challenge is, strictly speaking, an impossible one, it is unlikely that Coltheart means to pose a challenge to neuroimaging that is impossible by definition. Surely to make the above theoretical argument demonstrating the impossibility of meeting the challenge would have been easier than to refute various suggestions one by one.

Since the in-principle difficulty with the challenge is due to its ruling out statements that make location information relevant to functional information, the path for modifying the challenge to render it viable is clear. Coltheart's challenge can be revived by allowing bridge principles or auxiliary assumptions that enable one to infer function from location. Thus, if the prediction derived from $\mathbf{T_a}$ in purely functional terms (such as "If $\mathbf{T_a}$ then F; if $\mathbf{T_b}$ then $\sim F$") can be combined with statements $S_1 \ldots S_n$ such as "F corresponds to activation in region(s) $R_{i \ldots n}$," then $\mathbf{T_a}$ conjoined with S can yield a prediction that could be corroborated or disconfirmed by patterns of neuroimaging data. Thus, Coltheart's challenge requires one either to accept a body of statements $S_1 \ldots S_n$ that specify functional-anatomical mappings or else to accept a general principle S that there is a systematic relation between function and location in the brain. The importance of this principle to the field of cognitive neuroimaging was emphasized by Henson (2005). Importantly, and somewhat puzzlingly, Coltheart claims to accept this for the purposes of his argument: "I fully accept Henson's assumption that there is some systematic mapping from psychological function to brain structure" (2006b, 323). He does not, however, seem to have accepted its implications.

The revised challenge is thus to show that some neuroimaging study has enabled us to distinguish between competing psychological theories by providing evidence that weighs in favor of one over the other, where such theories are expressed at the psychological level. However, we must also be able to appeal to specific or general claims about functional-anatomical mappings in evaluating those theories. These claims, of necessity, cannot be part of the psychological theories because of the requirement that the theories themselves be expressed purely at the psychological level, but now we recognize that some appeal to bridging claims is essential. To see how this changes the game, let us return to the two studies discussed above. Consider again Umiltà's study. Coltheart

https://doi.org/10.1086/605815 Published online by Cambridge University Press

complained that his competing theories make no predictions about the functional architecture of attention. However, I think we can reword them so as to make some, albeit rather vague, predictions. We can charitably read Umiltà's theoretical hypotheses thus:

> $\mathbf{T}'_{\mathbf{a}}$. Endogenous and exogenous attention are functionally similar.
>
> $\mathbf{T}'_{\mathbf{b}}$. Endogenous and exogenous attention are functionally diverse.

Just how to word these theories is debatable. Perhaps we should say that they share/do not share functional components and so forth. However, though such wording does not put very distinct constraints on the predicted functional architecture, even broad claims like this put some constraints on subsequent theorizing.

Now, suppose that Umiltà's data show that attention is governed by different brain systems (Coltheart's construal of his $\mathbf{T}_{\mathbf{b}}$; Coltheart does not say whether or not the data convince him of this). If one accepts the idea that there is a systematic structure-function mapping in the brain, the following claim S deserves some credence:

> **S**. Activation of the same structures is evidence for operation of the same functions, and activation of different structures is evidence for operation of different functions.

Here I use "is evidence for" to suggest that it should raise our credence in the subsequent claim, but not that it necessitates it or that the proposed relationship cannot fail to hold in particular instances. Claim S seems to follow naturally from the claim that Coltheart accepts, namely "that there is some 'systematic' mapping from psychological function to brain structure" (2006b, 323). Given the data, plus S, one is licensed to infer that there is more support for $\mathbf{T}'_{\mathbf{b}}$ over $\mathbf{T}'_{\mathbf{a}}$. That is, the studies cited by Umiltà provide reason to prefer $\mathbf{T}'_{\mathbf{b}}$ over $\mathbf{T}'_{\mathbf{a}}$. And that is precisely the challenge that Coltheart has claimed has not been met.

Similarly, with respect to the Jonides et al. study, Coltheart has agreed that the results support the claim that working memory for verbal information is mediated by a different brain system than working memory for spatial information. Call this $\mathbf{T}'_{\mathbf{a}}$. The claims $\mathbf{T}'_{\mathbf{a}}$ and S provide reason to believe $\mathbf{T}_{\mathbf{a}}$ (working memory for verbal information is mediated by a different cognitive system than working memory for spatial information) over $\mathbf{T}_{\mathbf{b}}$, where $\mathbf{T}_{\mathbf{a}}$ and $\mathbf{T}_{\mathbf{b}}$ are the psychological theories of interest to Coltheart. Coltheart's challenge still stands, for he contends that no one actually has advocated the opposing theory, $\mathbf{T}_{\mathbf{b}}$ (that a single cognitive system supports both verbal and spatial working memory). But one can see, in principle, that neuroimaging data can lend differential support to one theory over another.

https://doi.org/10.1086/605815 Published online by Cambridge University Press

4. Possible Objections and the Nature of Structure-Function Mappings. The general form of neuroimaging studies that purport to provide some insight into the psychological basis of cognition therefore involve some theory stated at the psychological level, plus auxiliary hypotheses $S_1 \ldots S_n$ regarding brain structure–cognitive function relations. When $\mathbf{T_a} + S$ leads to predictions that are borne out by data that are contrary to those predicted by $\mathbf{T_b} + S$, any credence one places in S provides some reason to favor $\mathbf{T_a}$ over $\mathbf{T_b}$ and thus to distinguish between theories.

One way to maintain that neuroimaging cannot be of use to cognitive psychology is to demonstrate that there is no systematic structure-function mapping to be had. As I have noted, Coltheart agrees that there is such a mapping, so that avenue is not open to him. Alternatively, he could argue that the particular form of S used in various particular inferences is either false or unwarranted, so that there is no reason to believe that version of S. Here I consider a few potential arguments for discounting various versions of S, including the most strong one that Coltheart rejects, but that others, such as Van Orden and Paap (1997), seem to accept.[1]

4.1. There Is No Systematic Structure-Function Mapping. This argument would clinch the ultra-cognitive neuropsychologist's in-principle argument against the possibility of brain imaging data informing psychology, but it is unsound. Brains are not all-purpose computing devices like von Neumann computers: they have highly constrained architectural and functional properties. The fact that focal brain lesions result in disruption of specific functions, leaving others unimpaired, provides very strong evidence for systematic structure-function mappings in the brain. Decades of neurophysiological research likewise support the existence of a mapping between specific brain regions and neural function. It is overwhelmingly likely that developmental, energetic, and evolutionary constraints have ensured local structure-function relationships. While no neuroscientist disputes the main claim, many disagreements in the field concern the nature and extent of the mapping.

4.2. There Is No Systematic Structure-Function Mapping at the Scale to Which Functional Imaging Is Sensitive. The resolution of neuroimaging is currently on the order of a cubic millimeter. There are on the order of 100,000 neurons in a cubic millimeter of neural tissue, and it is certainly conceivable (and probably undeniable) that there are multiple functionally specialized networks operating within areas of the brain with these dimensions. If this is so, we can't expect neuroimaging to enable us to

1. I do not intend this list of objections to be exhaustive, merely illustrative.

https://doi.org/10.1086/605815 Published online by Cambridge University Press

distinguish these functional components. However, although neuroimaging cannot tell us everything about brain function, it would be shortsighted to claim that it therefore can't tell us anything. There is ample evidence that certain brain regions respond reliably to some tasks rather than others, more strongly for some than for others, and so on. Recently, using statistical techniques and the anisotropies in brain organization, some neuroimaging studies have managed to make use of information from functional units in visual cortex that are below the spatial resolution of imaging. For example, the orientation of a visual stimulus could be predicted by the pattern of activation in visual cortex despite the fact that the spatial organization of orientation columns in cortex is at the submillimeter level (Kamitani and Tong 2005; Yacoub, Harel, and Uğurbil 2008). Admittedly, these have relied on a prior understanding of the basic functional architecture of this area of cortex and thus are not pertinent to the question at issue, but they do illustrate that with the right techniques and auxiliary hypotheses, information about brain function can be gleaned at relatively fine spatial scales.

In order for Coltheart or someone like him to help himself to this claim, he would have to argue that there is no functional-structural mapping at spatial scales above a millimeter. But, although many physiological studies may focus on functional properties of very small regions of the brain, a vast body of data, from lesions to electrophysiology, demonstrate that there is structure-function specificity at much larger scales as well. The brain demonstrates functional organization at many scales. The burden of proof lies heavily on the side of the ultra-cognitivist if this is his claim.

4.3. There Is a Systematic Structure-Function Mapping, but Multiple Functions Are Colocalized. This argument is similar to the above, although here the focus is on colocalization of function at spatial scales above the level of resolution of the technique. We know that functions are colocalized, especially for functions individuated finely. What this means is that inferences from region to function are defeasible: Suppose that region R performs function F_1 in task A and function F_2 in task B. One cannot be sure whether it is performing F_1 or F_2 (or a novel function F_3) in task C. But perhaps you have independent reason to think that task C involves F_1. Then seeing signal R in task C should increase your credence in F_1 playing a role in task A. A Bayesian analysis can tell you how much it should increase your credence (see Poldrack 2006).

It is conceivable that relations of colocalized functions could be clarified if one could differentially involve F_1 or F_2 or include both in a single task. (Although such analyses involve further assumptions, they are ones amenable to independent empirical substantiation.) One technique that helps distinguish colocalized functions is called fMRI adaptation (Grill-Spector

https://doi.org/10.1086/605815 Published online by Cambridge University Press

and Malach 2001; Grill-Spector 2006). In fMRI adaptation studies, one task is used to "fatigue" a group of neurons during a habituation phase (task A); then one determines whether there is reduced response in the same area during a different task (task B). If there is reduction during performance of task B, it is concluded that the same group of neurons is involved in both tasks, and one infers from the same neurons to the same computational process or function. If there is no decrement in that region during performance of task B, it is inferred that a different group of neurons in the same region is involved in that task, so that two separate functions are colocalized. The limits of this technique are bound to limitations in designing tasks with separate functional components. If no task allows one to dissociate functions F_1 and F_2, one may infer a single function. The limits of imaging for inferring function are reached in this case.

No one thinks that all functions are colocalized everywhere. In fact, the more similar functions are to each other, the more independent reason there is to expect colocalization. If so, then even given the possibility of colocalized functions, the similarity itself may increase our credence in some inferences over others. Again, it does seem that the burden of proof is on the ultra-cognitivist to argue that functions are colocalized in such a way that the evidence from any given experiment provides no reason to prefer one hypothesis over another.

4.4. There Is a Systematic Structure-Function Mapping at Larger Scales, but It Is Redundant, So Different Structures Can Implement the Same Function. This is also surely true to some extent. It is clear that at some level redundancy is an important feature of brain organization. For example, in primary visual cortex an array of (hypothetically) functionally equivalent neural modules processes information from different regions of the visual field. However, though the computations may be redundant, their spatial receptive fields differ, so there are dimensions along which they can be distinguished. Interestingly, this massive parallelism (at least with respect to stimulus sensitivity or receptive field) diminishes as one progresses up the cortical hierarchy, so that cells later in the visual stream respond in a much more stimulus-specific manner. It remains an open question, however, how much redundancy there is in processing at the levels of interest to cognitive psychology and whether there are methods that can individuate redundant functions.

The problem of identifying function when responses depend on both functional role and stimulus (input) specificity is a conceptually difficult one, but nonetheless one that can be informed by differential and converging patterns of activity. It is true that for any two separate regions of activity it is *possible* that both perform the same function. For instance, Coltheart could claim that Umiltà's example fails to meet the challenge

https://doi.org/10.1086/605815 Published online by Cambridge University Press

because massive redundancy could make it the case that endogenous and exogenous attention both involve the very same cognitive operations; they are merely realized in different brain circuits. However, the relevant question is not whether this is *possible* but whether it is *likely*. Given architectural and evolutionary constraints as well as current understanding of functional specialization in the brain, it is arguably much more likely that different brain regions are involved in performing different rather than the same computational or cognitive functions. For instance, it is unlikely that brain circuits mediating initial performance of a motor task and its performance after learning would differ if both circuits implemented the very same computations. The fact that large-scale differences in networks are evident with imaging is suggestive that different functional components are involved in the two tasks (Petersen et al. 1998). It is possible that the degree to which separate activations should increase credence in functional differences could be formalized by using data from functional, lesion, and physiological studies to provide us with prior probabilities of shared function across different brain areas (Poldrack 2006).

Again, given that one accepts a general principle of structure-function mapping in the brain, it seems that the burden of proof is on the ultracognitivist to explain why reliable differences in activation should provide *no* evidence regarding the likelihood of difference in function.

4.5. There Is a Systematic Structure-Function Mapping, but the Mapping at the Scales Amenable to Functional Imaging Is Cognitively Uninteresting. This is a stronger version of several of the above, for it suggests that (A) all the functions we are interested in discriminating between are either (1) colocalized in the sense outlined above, so that they are spatially indistinguishable and functionally indistinguishable, so that no pattern of data could help distinguish between models; or (2) multiply instantiated so that whatever the pattern we see, it gives us no reason to prefer one model over another; in addition (B) no visible differences in patterns of activity that we do see that are reliably correlated with various tasks have bearing on interesting cognitive questions.

I suppose that Coltheart would endorse this option, but this is quite a strong claim. It relies heavily on a view of what is "cognitively interesting" for one thing. Cognitively interesting just can't be "whatever is not amenable to functional imaging," although the dismissal of studies that show dissociation of regions during different processes has that kind of ring to it. Functions are presumably hierarchically structured and can be identified at various levels of grain (from performance of entire tasks, to various task components, and probably ultimately to operations of small local networks of neurons), and there is no single level that is the purview of cognitive psychology. Neuroimaging doesn't purport to provide infor-

https://doi.org/10.1086/605815 Published online by Cambridge University Press

mation about all conceivable levels, but it requires arguments to claim that it is relevant to none of those of interest to cognitive psychology. Moreover, both A and B are substantive claims that seem to require argument: they are not a default. One might argue that any version of S that licenses the inferences that defeat B is more likely than A and B are to be true, so again the burden of proof is on the ultra-cognitivist. But once we see the burden of proof in this way, Coltheart's strategy of dismissing candidate studies as failing his challenge by denial of S appears ill motivated.

5. Conclusion. Coltheart's challenge makes plain the importance of functional-anatomical mapping to neuroimaging. Some form of functional-anatomical specification is absolutely essential to imaging (Henson 2005). However, there has been remarkably little discussion of what sort of specification is needed and little discussion of what form(s) are likely to be true and of what implications the possibilities might have for constraining the ability of imaging techniques to illuminate cognitive questions. My impression is that many (ultra-cognitivists and others) suppose that a much stronger form of mapping is needed than is really needed in order for imaging to play a role in understanding cognition. Only a weak form is needed for imaging to be relevant, but how informative it can ultimately be will depend on how strong the mapping is. On the other hand, my impression is that many cognitive neuroimagers are surprisingly oblivious to the extent to which the relevance of their methods to their object of interest depends on the nature of this mapping: The imaging literature is rife with sloppy inferences and lack of clarity about implicit assumptions. Although it would be mistaken to suppose that neuroimagers are unconcerned or unaware of these issues, their papers rarely discuss the assumptions that underlie the interpretation of their data and often seem to take as an unwritten and unproblematic premise a strong functional-anatomical mapping. More discussion about the commitments underlying the interpretation of individual studies, and relevant evidence that can underwrite those commitments, would help make the arguments from brain activity to functional conclusion much more transparent and would be a step toward maturity for imaging as a scientific discipline.

Coltheart's challenge, suitably modified, has been met. However, the deeper point is that any challenge of the form of Coltheart's is going to be hostage to the degree and nature of functional localization. To deny that neuroimaging data can constrain our psychological reasoning is to deny that there is any structure-function mapping, and this position is plainly false. What form the true proposition takes is an important and difficult empirical question. It is virtually certain also that it is not the case that for all functions there is a one-to-one structure-function mapping

https://doi.org/10.1086/605815 Published online by Cambridge University Press

at scales amenable to imaging techniques. There are limits to what imaging can tell us about psychology, and we have yet to determine what they are. One can acknowledge this while also accepting that neuroimaging can bear on questions of mind.

REFERENCES

Coltheart, Max (2004), "Brain Imaging, Connectionism, and Cognitive Neuropsychology", *Cognitive Neuropsychology* 21 (1): 21–25.

——— (2006a), "Perhaps Functional Neuroimaging Has Not Told Us Anything about the Mind (So Far)", *Cortex* 42: 422–427.

——— (2006b), "What Has Functional Neuroimaging Told Us about the Mind (So Far)?", *Cortex* 42: 323–331.

Fodor, Jerry A. (1999), "Let Your Brain Alone", *London Review of Books* 21 (September 30).

Grill-Spector, Kalanit (2006), "Selectivity of Adaptation in Single Units: Implications for FMRI Experiments", *Neuron* 49 (2): 170–171.

Grill-Spector, Kalanit, and Rafael Malach (2001), "fMR-Adaptation: A Tool for Studying the Functional Properties of Human Cortical Neurons", *Acta Psychologica* 107 (1–3): 293–321.

Henson, Richard (2005), "What Can Functional Neuroimaging Tell the Experimental Psychologist?", *Quarterly Journal of Experimental Psychology* 58A (2): 193–233.

Jonides, John, Derek E. Nee, and Mark G. Berman (2006), "What Has Functional Neuroimaging Told Us about the Mind? So Many Examples, So Little Space", *Cortex* 42: 414–417.

Kamitani, Yukiyasu, and Frank Tong (2005), "Decoding the Visual and Subjective Contents of the Human Brain", *Nature Neuroscience* 8 (5): 679–685.

Petersen, Steven E., Hanneke van Mier, Julie A. Fiez, and Marcus E. Raichle (1998), "The Effects of Practice on the Functional Anatomy of Task Performance", *Proceedings of the National Academy of Sciences of the United States of America* 95: 853–860.

Poeppel, David (1996), "A Critical Review of PET Studies of Phonological Processing", *Brain and Language* 55: 317–351.

Poldrack, Russell A. (2006), "Can Cognitive Processes Be Inferred from Neuroimaging Data?", *Trends in Cognitive Sciences* 10 (2): 59–63.

Smith, Edward E., and John Jonides (1997), "Working Memory: A View from Neuroimaging", *Cognitive Psychology* 33: 5–42.

Umiltà, Carlo (2006), "Localization of Cognitive Functions in the Brain Does Allow One to Distinguish between Psychological Theories", *Cortex* 42: 399–401.

Uttal, William (2001), *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*. Cambridge, MA: MIT Press.

Van Orden, Guy C., and Kenneth R. Paap (1997), "Functional Neuroimages Fail to Discover Pieces of Mind in Parts of the Brain", *Philosophy of Science* 64 (Proceedings): S85–S94.

Yacoub, Essa, Noam Harel, and Kâmil Uğurbil (2008), "High-Field fMRI Unveils Orientation Columns in Humans", *Proceedings of the National Academy of Sciences of the United States of America* 105: 10607–10612.

https://doi.org/10.1086/605815 Published online by Cambridge University Press