Fitting computational models to fMRI data

F. GREGORY ASHBY AND JENNIFER G. WALDSCHMIDT

University of California, Santa Barbara, California

Many computational models in psychology predict how neural activation in specific brain regions should change during certain cognitive tasks. The emergence of fMRI as a research tool provides an ideal vehicle to test these predictions. Before such tests are possible, however, significant methodological problems must be solved. These problems include transforming the neural activations predicted by the model into predicted BOLD responses, identifying the voxels within each region of interest against which to test the model, and comparing the observed and predicted BOLD responses in each of these regions. In the present article, methods are described for solving each of these problems.

fMRI has revolutionized psychology. The effects have probably been greatest within cognitive psychology and perception, but the influence of fMRI has spread to almost every area of psychology. For example, there are now emerging new fields of social neuroscience, developmental neuroscience, and neuroeconomics, all largely due to fMRI. There is even a new field of neuromarketing. By far, the great majority of fMRI research in these fields uses fMRI as a tool for exploratory research. The most common question in this approach is: Which brain areas are activated by the task I am studying? This is clearly an important question that has had enormous benefit, and will continue to motivate future research.

On the other hand, we believe fMRI has significant untapped potential as a tool for confirmatory research. In such an approach, the research question instead would be: Does a model or theory correctly predict the results of an fMRI experiment? Or, alternatively, which of two competing models gives the best account of fMRI data? Requiring that a satisfactory model is consistent not only with available behavioral data but also with fMRI data greatly constrains the class of acceptable models. In most cognitive behavioral experiments, at best only two pieces of behavioral data are available on each trial: which response was made and the response time (RT). Many models that postulate qualitatively different underlying psychological processes might nevertheless give equally good accounts of such sparse data sets. In contrast, fMRI data is incredibly rich, and the class of models that are consistent with both the behavioral and fMRI data is likely to be quite small.

It is true that models potentially capable of making predictions about the outcome of fMRI experiments must make neurobiological assumptions that may not be necessary if the goal is to account simply for behavioral data. Even so, there are now many neurobiologically detailed computational models in psychology that make predictions (or that could make predictions) about how neural

activity changes in specific brain regions as participants perform certain cognitive tasks, and fMRI seems ideally suited to testing such predictions (see, e.g., Anderson & Qin, in press; Ashby, Ell, Valentin, & Casale, 2005; Brown & Braver, 2005; Chadderdon & Sporns, 2006; Frank & Claus, 2006; Mason & Just, 2006; Monchi, Taylor, & Dagher, 2000; O'Reilly, 2006).

Yet, in order to test computational models in this way, several problems must be overcome. First, psychological models might make predictions about neural activation, but they generally do not make predictions about the blood-oxygen-level-dependent (BOLD) signal that is most commonly measured in fMRI experiments. Thus, the first problem is to generate predicted BOLD responses from the model's predicted neural activations.

Second, most models in psychology that make neuroanatomical predictions typically do not make predictions that are specific enough to identify a small set of voxels in the region of interest (ROI) that could be used to test the model. For example, a model might specify that during a certain period of a working memory task, a specific type of activation should occur in dorsolateral prefrontal cortex (dlPFC). However, such a model would generally not predict that every voxel in dlPFC would show this activation pattern—only that some would. So, a second significant problem that must be solved is to identify exactly which voxels within dlPFC should be used to test the model. Finally, the third problem is to compare the observed and predicted BOLD responses in the selected voxels and to decide on the basis of this comparison whether the model succeeds or fails at accounting for the results of the experiment.

This article proposes solutions to these three problems. The solutions are straightforward, and standard fMRI statistical methods are used. There have been a few previous attempts to fit computational models to fMRI data (see, e.g., Danker & Anderson, 2007), and these have used some of the same methods proposed in the present article

(e.g., convolution with a hemodynamic response function). However, we know of no attempts that used all of the methods described below, nor do we know of any other articles that are devoted solely to this topic.

An Example

As an example of a model that could be tested against fMRI data, we will consider a recent computational model called SPEED (“subcortical pathways enable expertise development”), which describes how automaticity develops in perceptual categorization tasks (Ashby, Ennis, & Spiering, 2007). Briefly, the model assumes that there are two neural pathways from the sensory association area that mediate perception of the stimulus to the premotor region of cortex that mediates response selection. A longer and slower path projects to the premotor area via the striatum, globus pallidus, and thalamus. A faster, purely cortical path projects directly from the sensory association area to the premotor area. SPEED assumes that the subcortical path—although slower—has greater neural plasticity because of a dopamine-mediated learning signal from the substantia nigra. In contrast, the faster cortical–cortical path learns more slowly via (dopamine independent) classical two-factor Hebbian learning. Because of its greater plasticity, early performance is dominated by the subcortical path, but the development of automaticity is characterized by a transfer of control to the faster cortical–cortical projection. The model includes differential equations that describe activation in each of the relevant brain areas, as well as a set of difference equations that describe the relevant two- and three-factor learning.

Ashby et al. (2007) showed that SPEED accounts for some classic single-cell recording and behavioral results, but the model has not yet been tested against fMRI data. Even so, it makes detailed predictions about many of the brain regions that should show task-related activation during category learning. For example, SPEED predicts how neural activation in a variety of brain regions should change over the course of a single trial and how these activations will change after many days of practice on the task. Thus, fMRI seems like an ideal tool for rigorously testing the model.

Problem 1: Generate BOLD Predictions From the Model

The fMRI BOLD response increases with the amount of oxygenated hemoglobin in a voxel relative to the amount of deoxygenated hemoglobin (Ogawa, Lee, Kay, & Tank, 1990). In comparison with the neural activation that presumably drives it, the BOLD response is highly sluggish. It reaches a peak around 6 sec after the neural activation that induced it, and it slowly decays back to baseline 20–25 sec later. The first step in fitting a neurocomputational model to fMRI data, therefore, is to convert the predicted neural activations into predicted BOLD responses.

Almost all current applications of fMRI assume that the transformation from neural activation to BOLD response can be modeled as a linear, time-invariant system (e.g., Boynton, Engel, Glover, & Heeger, 1996). In the linear systems approach, one can conceive of the vascular

system that responds to a sudden oxygen debt as a black box in which the input is neural activation, and the output is the BOLD response. Suppose we present a stimulus event, E_i , to a subject at time 0. Let $N_i(t)$ denote the neural activation induced by this event at time t , and let $B(t)$ denote the resulting BOLD response. A system of this type is linear if and only if

$$B[N_1(t) + N_2(t)] = B[N_1(t)] + B[N_2(t)].$$

If the system is linear, then it is well known that the BOLD response to any neural activation $N_i(t)$ can be computed from

$$B(t) = \int_0^t N(\tau)h(t - \tau)d\tau. \quad (1)$$

Equation 1 is the well-known *convolution integral* that completely characterizes the behavior of any linear, time-invariant system (see, e.g., Chen, 1970). The function $h(t)$ is traditionally called the *impulse response function* because it describes the response of the system to an input that is a perfect impulse. In the fMRI literature, however, $h(t)$ is known as the *hemodynamic response function*, often abbreviated as hrf. The hrf is the hypothetical BOLD response to an idealized impulse of neural activation, often peaking at 6 sec and lasting for 30 sec or so.

To generate a predicted BOLD response from our model, we first use the model to generate a predicted neural response, $N(t)$. For example, the top panel of Figure 1 shows the neural activation predicted by SPEED (Ashby et al., 2007) in premotor cortex (i.e., either supplementary motor area or presupplementary motor area) during one categorization trial and after 50 trials of practice. In this example, the trial begins with a 1-sec fixation period.

The next step is to select a mathematical form for the hrf. These two functions are then numerically convolved (via Equation 1) to produce a predicted BOLD response. There are many alternative models of the hrf. A common approach is to select a canonical form for the hrf that is characterized by several free parameters that can be adjusted to account for differences in the hrf across brain regions or subjects. Many alternative models of this type have been proposed (Boynton et al., 1996; Clark, Maisog, & Haxby, 1998; Cohen, 1997; Dale & Buckner, 1997; Friston, Holmes, & Ashburner, 1999; Friston, Josephs, Rees, & Turner, 1998; Zarahn, Aguirre, & D’Esposito, 1997). For example, Boynton et al. proposed a gamma function of the type

$$h(t) = \begin{cases} \frac{(t - T_0)^{n-1}}{\lambda^n (n-1)!} e^{-\frac{t-T_0}{\lambda}} & \text{for } t > T_0, \\ 0 & \text{for } t < T_0, \end{cases} \quad (2)$$

where λ , T_0 , and n (an integer) are free parameters. The parameter T_0 is the lag between stimulus presentation and the initial rise in the BOLD response, whereas λ and n determine the shape of the hrf. For example, the peak of the hrf occurs at time $T_0 + (n-1)\lambda$. The parameters can be either estimated from the data or just set to some typical values. Figure 2 shows an example of this model with the parameters set to $T_0 = 0$, $n = 4$, and $\lambda = 2$. The bottom panel of

Figure 1. The top panel shows activation predicted by the computational model SPEED (Ashby et al., 2007) in premotor cortex during a single categorization trial. The bottom panel shows the predicted BOLD response under these same conditions.

Figure 1 shows the result of convolving the Figure 1 neural activation predicted by SPEED (top panel) with the hrf shown in Figure 2. This is the BOLD response predicted by SPEED in this particular ROI for one categorization trial.

All models of the hrf have the same basic shape, in the sense that they all rise to a peak at around 6 sec and then slowly decay back to baseline. Some include a late negative dip because empirical estimates of the BOLD response often show this property (see, e.g., Fransson, Krüger, Merboldt, & Frahm, 1999; Glover, 1999). However, all models that use an hrf to predict a BOLD response via the Equation 1 convolution integral depend critically on the linearity assumption. There is good evidence that linearity is approximately satisfied if the time between stimulus events exceeds several seconds and if brief stimulus exposure durations are avoided (Vazquez & Noll, 1998). However, if stimulus events quickly follow one another, or if brief stimulus ex-

posure durations are used, then it is well documented that the BOLD response exhibits significant nonlinearities (Hinrichs et al., 2000; Huettel & McCarthy, 2000; Ogawa et al., 2000; Pfeuffer, McCullough, Van de Moortele, Ugurbil, & Hu, 2003). In this case, Equation 1 is no longer valid. Some complex approaches directly model the physical processes that cause these nonlinearities (e.g., the balloon model of Buxton, Wong, & Frank, 1998). Alternatively, Friston et al. (1998) proposed a computationally simpler method that uses a second-order Volterra series for modeling nonlinearities that simply adds a correction term to Equation 1.

In summary, there is a large literature on alternative methods for modeling the BOLD response (for a review, see, e.g., Henson & Friston, 2007). The simplest solution is to use a model of the hrf with no free parameters. An intermediate choice is to use an hrf model that has free parameters that must be estimated from the data. A third, more sophisticated choice is to replace the hrf with a second-order Volterra series that models nonlinearities in the transformation from neural activation to BOLD response. The decision about which of these solutions to choose should depend on the experimental design, the amount of detail specified by the computational model to be tested, and the state of the field within which that model resides. For example, the decision about whether to use linear (i.e., hrf) or nonlinear (e.g., Volterra series) methods should depend on whether the experiment used brief stimulus exposures and short interstimulus intervals (ISIs). As another example, a parameter-free model of the hrf might suffice in research areas where there have been no previous attempts to fit computational models to fMRI data.

Problem 2: Identify Voxels Against Which to Test the Model

Once predicted BOLD responses have been generated from the model, the second problem is to identify exactly which voxels should be used to provide the data that the model will be tested against. Our solution to this problem includes three steps. The first is to apply the general linear model (GLM) to each voxel using the method of Ollinger,

Figure 2. A typical hrf used in fMRI data analysis. This is the gamma function model of Equation 2 that was proposed by Boynton et al. (1996) with $T_0 = 0$, $n = 4$, and $\lambda = 2$.

Corbetta, and Shulman (2001) and Ollinger, Shulman, and Corbetta (2001). The second step is to form a statistical parametric t -map (t -SPM) from these GLM results, and the third step is to identify the set of candidate voxels within each ROI from this SPM.

Step 1: Apply the GLM to each voxel using the method of Ollinger and colleagues. Consider an fMRI experiment in which the time between onsets of successive whole brain scans (i.e., the repetition time, or TR) is 3 sec, and the BOLD response to a stimulus onset lasts for 12 sec, or equivalently five TRs. This is an unrealistic value, since in most cases we would expect the BOLD response to last at least twice this long. However, this shorter value is sufficient to illustrate the computational method. Ollinger, Corbetta, and Shulman (2001) and Ollinger, Shulman, and Corbetta (2001) proposed that, in this case, we define five different parameters that describe the BOLD response to the stimulus presentation at each of the ensuing five TRs. Call these five parameters $\beta_1, \beta_2, \beta_3, \beta_4$, and β_5 . Assume also that every time the same stimulus is presented, the exact same BOLD response occurs.

Now consider an event-related experimental design in which a stimulus event occurs on TRs 1, 3, 4, and 7. Then, the stimulus events and their BOLD responses at each TR will be

TR:	1	2	3	4	5	6	7	8	...
Events:	E		E	E			E		
Response to 1st event:	β_1	β_2	β_3	β_4	β_5				
Response to 2nd event:			β_1	β_2	β_3	β_4	β_5		
Response to 3rd event:				β_1	β_2	β_3	β_4	β_5	
Response to 4th event:							β_1	β_2	

Therefore, if the system is linear, then the observed BOLD response at each TR will equal

TR	BOLD
1	β_1
2	β_2
3	$\beta_1 + \beta_3$
4	$\beta_1 + \beta_2 + \beta_4$
5	$\beta_2 + \beta_3 + \beta_5$
6	$\beta_3 + \beta_4$
7	$\beta_1 + \beta_4 + \beta_5$
8	$\beta_2 + \beta_5$

These equations can be expressed in terms of the GLM as follows:

$$\begin{bmatrix} B(1) \\ B(2) \\ B(3) \\ B(4) \\ B(5) \\ B(6) \\ B(7) \\ B(8) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 2 \\ 1 & 0 & 1 & 0 & 0 & 1 & 3 \\ 1 & 1 & 0 & 1 & 0 & 1 & 4 \\ 0 & 1 & 1 & 0 & 1 & 1 & 5 \\ 0 & 0 & 1 & 1 & 0 & 1 & 6 \\ 1 & 0 & 0 & 1 & 1 & 1 & 7 \\ 0 & 1 & 0 & 0 & 1 & 1 & 8 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \\ \beta_5 \\ B_0 \\ \Delta \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \\ \varepsilon_5 \\ \varepsilon_6 \\ \varepsilon_7 \\ \varepsilon_8 \end{bmatrix}, \quad (3)$$

where B_0 estimates the baseline activation, and Δ measures the amount of linear drift in the magnetic field strength. Jittering the stimulus events guarantees that $X'X$ is nonsingular. Therefore, uniformly minimum variance unbiased estimators of $\underline{\beta}$ (i.e., the gold standard of parameter estimation) are computed from the solution of the GLM normal equations

$$\hat{\underline{\beta}} = (X'X)^{-1}X'y. \quad (4)$$

This GLM analysis is performed separately in every voxel. The end result is a $\hat{\underline{\beta}}$ vector for every voxel. In the next step, we will reduce each of these vectors to a single t statistic.

Step 2: Form the SPM. The SPM we will form compares the central $\hat{\beta}_i$ with the early and late $\hat{\beta}_i$. For example, if the BOLD response to an event is assumed to persist for 25 sec and the TR is 2.5 sec, then there will be 11 $\hat{\beta}_i$. In this case, we will construct the SPM that tests the null hypothesis

$$H_0: (\beta_2 + \beta_3 + \beta_4) - (\beta_1 + \beta_{10} + \beta_{11}) = 0 \quad (5)$$

against the alternative

$$H_1: (\beta_2 + \beta_3 + \beta_4) - (\beta_1 + \beta_{10} + \beta_{11}) > 0.$$

Although SPMs based on other, simpler hypotheses could be produced, this particular SPM has several attractive advantages. Most importantly, it uses more data than do SPMs that are based on simpler hypotheses, and it makes stronger structural assumptions. For example, the t statistic produced from the Equation 5 null hypothesis will be large only if the BOLD response starts small, grows to a large value between 2.5 and 7.5 sec after the stimulus event, and then decays back to a small value by 22.5 sec after the event.

Note that the null hypothesis specified by Equation 5 can be written in matrix form as

$$H_0: \underline{w}'\underline{\beta} = 0,$$

where $\underline{w}'$ is the row vector

$$\underline{w}' = [-1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ -1 \ -1 \ 0 \ 0].$$

Note that the last two 0s in this vector remove the influence of B_0 and Δ on the SPM. If, as is common, we assume homogeneity of variance and independence across TRs, then under this null hypothesis, the statistic

$$t = \frac{\underline{w}'\hat{\underline{\beta}}}{\sqrt{\hat{\sigma}^2 \underline{w}'(X'X)^{-1}\underline{w}}} \quad (6)$$

has a t distribution with degrees of freedom equal to

$$df_E = (\# \text{ of TRs}) - 1 - (\# \text{ of parameters in } \underline{\beta}).$$

The maximum likelihood estimator of σ^2 is known to equal

$$\hat{\sigma}^2 = \frac{1}{df_E} (\underline{y} - X\hat{\underline{\beta}})' (\underline{y} - X\hat{\underline{\beta}}). \quad (7)$$

We now have a single t statistic in every voxel. Large values will come from voxels that appear to be task responsive.

Step 3: Select voxels for model testing. The next step is to use the t -SPM to select the specific voxels within each

ROI that will be used in the model testing. There are a number of plausible methods for achieving this goal. One possibility is to select the n voxels within each ROI that have the largest t values. The value of n should probably increase with the size of the ROI, but for purposes of discussion, a value of $n = 10$ is reasonable. One advantage of this method is that no statistical threshold is used to identify voxels. Thus, the 10 largest t statistics in an area might all fall below the statistical threshold for significance. Presumably, the model predicts that these t values should all be significant; thus, by excluding a significance criterion from our selection procedure, we are minimizing one important type of selection bias.

A disadvantage of this method is that the 10 voxels selected might not be contiguous. So, another approach might be to use a cluster-based selection procedure, where a cluster is defined as a set of contiguous (suprathreshold) voxels. In this method, an intermediate threshold is set—for example, at $T = 2.5$ —after the t -map is formed in Step 2. Next, voxels with t values greater than this threshold are used to identify all clusters in the ROI. Finally, within each ROI specified by the model, one might identify the largest cluster and use data from these voxels to test the model.

At this point, there is no empirical rationale for choosing one of these methods over the other. Thus, the choice of which method is used should probably depend on the theoretical predictions of the model. For example, SPEED predicts that in a task with two response alternatives, there should be two separate task-sensitive striatal units. The model assumes that if the stimuli associated with each response are similar and the motor actions required to produce each response are similar, then the two striatal units should be near each other in the striatum, but SPEED does not assume that they must be contiguous. For this reason, the method that chooses the n voxels with the largest t values is probably preferable to the method that chooses the largest above-threshold cluster.

Problem 3: Compare Observed and Predicted BOLD Responses in Selected Voxels

We are now in a position to address the last of our three problems. There are four steps to complete.

Step 1: Prepare the data. In each ROI specified by the model, we have identified 10 voxels that we will use for model testing. In each voxel, a BOLD response was recorded at every TR of the experiment. Suppose that the experiment included a total of N TRs. Then, the data from each of the selected voxels is an $N \times 1$ vector. The first step is to compute the mean of the 10 selected data vectors in each ROI. After this averaging process, there will be one $N \times 1$ data vector in each ROI. Our task is to generate a predicted vector from the model in each ROI and to compare the predicted vector with this observed data vector.

Step 2: Prepare the model. The goal of this step is to produce an $N \times 1$ vector containing predicted BOLD responses at each TR of the experiment. To begin, as was described in the solution to Problem 1, the equations of the model are used to generate predicted neural activation in each ROI during every TR of the experiment. For example, the top panel of Figure 3 shows the neural activations predicted by SPEED in premotor cortex during the

first 30 sec of a hypothetical event-related fMRI experiment. In this example, the TR is 3 sec, and a categorization stimulus is presented on TRs 1, 4, 5, 7, and 10. Each categorization TR begins with a 2-sec pause and ends with a 1-sec stimulus presentation.

Next, these predicted neural activations are numerically convolved with an hrf to produce the predicted BOLD response at each ROI throughout the experiment. The bottom panel of Figure 3 shows the results of convolving the Figure 2 hrf with the neural activations shown in the top panel of Figure 3. This process is completed by filling an $N \times 1$ vector (for every ROI) with the values of this predicted BOLD response at each of the N TRs of the experiment. Let $x_i(1), x_i(2), \dots, x_i(N)$ denote these predicted values in the i th ROI. The bottom panel of Figure 3 shows how the value $x_i(5)$ is defined.

Finally, we build the following GLM at each ROI:

$$\begin{bmatrix} B_i(1) \\ B_i(2) \\ \vdots \\ B_i(N) \end{bmatrix} = \begin{bmatrix} x_i(1) & 1 & 1 \\ x_i(2) & 1 & 2 \\ \vdots & \vdots & \vdots \\ x_i(N) & 1 & N \end{bmatrix} \begin{bmatrix} \theta_M \\ B_0 \\ \Delta \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \vdots \\ \epsilon_N \end{bmatrix}, \quad (8)$$

Figure 3. The top panel shows activation predicted by the computational model SPEED (Ashby et al., 2007) in premotor cortex during 30 sec of an event-related fMRI experiment. The bottom panel shows the predicted BOLD response under these same conditions.

where θ_M measures the correlation between the predictions of the model and the data. As before, B_0 is the baseline activation, and Δ measures the linear drift in the magnetic field strength.

If the model is correct, then, of course, θ_M should be large. Equation 6 can be used to test the null and alternative hypotheses $H_0: \theta_M = 0$ and $H_1: \theta_M > 0$, with $\underline{w}' = [1 \ 0 \ 0]$. Under this null hypothesis, the Equation 6 t statistic has a t distribution with $df_E = N - 4$. Thus, rejecting this null hypothesis in each ROI specified by the model is evidence that the model provides a good fit to the data from this experiment.

The use of the GLM to evaluate the fit of the model requires that no free parameters are estimated during the fitting process. One possibility is to estimate the free parameters in the model by first fitting to single-unit recording data (as in Ashby et al., 2005) or behavioral data. If free parameters in the model must be estimated from the fMRI data, then a traditional iterative minimization routine must be used instead of the GLM during this step (i.e., for parameter estimation and model fitting).

In general, the same is true if a model of the hrf is chosen that has free parameters that must be estimated from the data. Obviously, if a poor model of the hrf is chosen, then the correlation between predicted and observed BOLD responses will be low, even if the model accurately predicts the neural activation. If an iterative estimation procedure is used to estimate parameters of the hrf, then the convolution integral of Equation 1 must be computed numerically at each iteration.

Friston et al. (1998) had a clever idea that eliminates most of this numerical integration. Their idea was to model the hrf as a weighted linear combination of basis functions:

$$h(t) = \theta_1 b_1(t) + \theta_2 b_2(t) + \theta_3 b_3(t), \quad (9)$$

where $b_i(t)$ is the i th basis function, and θ_i is its weight. Friston et al. (1998) suggested gamma probability density functions for the $b_i(t)$ with means and variances equal to 4, 8, and 16, respectively. The key, however, is that none of the $b_i(t)$ have any free parameters. The only free parameters are the weights θ_1 , θ_2 , and θ_3 .

The advantage of defining the hrf in terms of basis functions can be seen when Equation 9 is substituted into the convolution integral of Equation 1:

$$\begin{aligned} B(t) &= \int_0^t N(\tau) h(t - \tau) d\tau \\ &= \int_0^t N(\tau) [\theta_1 b_1(t - \tau) + \theta_2 b_2(t - \tau) \\ &\quad + \theta_3 b_3(t - \tau)] d\tau \\ &= \theta_1 \left[\int_0^t N(\tau) b_1(t - \tau) d\tau \right] + \theta_2 \left[\int_0^t N(\tau) b_2(t - \tau) d\tau \right] \\ &\quad + \theta_3 \left[\int_0^t N(\tau) b_3(t - \tau) d\tau \right]. \end{aligned}$$

Note that none of these integrals include any free parameters. As a result, if we define

$$y_i(t) = \int_0^t N(\tau) b_i(t - \tau) d\tau, \quad (10)$$

then

$$B(t) = \theta_1 y_1(t) + \theta_2 y_2(t) + \theta_3 y_3(t),$$

where the only free parameters are the θ_i .

With this reformulation of $B(t)$, the iterative and time-consuming parameter estimation process can be avoided. Instead, optimal parameter estimates can be computed in one step. First, the three integrals specified in Equation 10 are computed numerically. Each of these computations produces a vector of numerical values: $y_1(t)$, $y_2(t)$, and $y_3(t)$. The parameters θ_i can now be determined using standard linear regression techniques. First, the vector $x(t)$ in column 1 of the Equation 8 design matrix is replaced by the three columns $y_1(t)$, $y_2(t)$, and $y_3(t)$, and the $\underline{\beta}$ vector becomes

$$\underline{\beta} = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ B_0 \\ \Delta \end{bmatrix}.$$

Given this formulation, the θ_i can be estimated using the standard GLM estimator of Equation 4.

Given this basis function approach to modeling the hrf, a good fit of our model to the data occurs if any of the θ_i are large. This can be tested statistically via the null hypothesis

$$H_0 : \sum_{i=1}^3 \theta_i^2 = 0.$$

If this hypothesis is true, then the statistic

$$\chi^2 = \frac{\hat{\theta}' [A(X'X)^{-1}A']^{-1} \hat{\theta}}{\hat{\sigma}^2}$$

has an asymptotic χ^2 distribution with three degrees of freedom. The matrix A is given by

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}.$$

Step 3: Construct a comparison model. We would have much more confidence in the model, however, if we could verify that it fit better than some other commonly applied model. By far, the most widely applied fMRI model is the so-called *boxcar* model, which assumes that neural activation equals 1 when the stimulus is present, and that it equals 0 otherwise. Thus, to construct a comparison model, we begin by constructing a boxcar function for the entire experiment. An example is shown in the top panel of Figure 4 for the same 30 sec of the event-related design that was used to produce the model predictions shown in Figure 3. Next—as before—this neural activation func-

Figure 4. The top panel shows activation predicted by the boxcar model during 30 sec of an event-related fMRI experiment. The bottom panel shows the predicted BOLD response under these same conditions.

tion is numerically convolved with an hrf to produce the predicted BOLD response shown in the bottom panel of Figure 4. The results are used to fill an $N \times 1$ vector with the values of this predicted BOLD response at each of the N TRs of the experiment. Let $w(1), w(2), \dots, w(N)$ denote these values. Note that the BOLD response predicted by the boxcar model is the same in every ROI.

Finally—as before—at each ROI, we build the GLM:

$$\begin{bmatrix} B_i(1) \\ B_i(2) \\ \vdots \\ B_i(N) \end{bmatrix} = \begin{bmatrix} w(1) & 1 & 1 \\ w(2) & 1 & 2 \\ \vdots & \vdots & \vdots \\ w(N) & 1 & N \end{bmatrix} \begin{bmatrix} \theta_C \\ B_0 \\ \Delta \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \vdots \\ \epsilon_N \end{bmatrix}, \quad (11)$$

where θ_C measures the correlation between the predictions of the boxcar model and the data.

Step 4: Compare the two models. A strong test of our model is to ask whether it provides a better fit than the boxcar model. Both models are the same except for the first column of the design matrix X . A poor fit of a model will mean larger residuals at each TR, and, consequently,

a larger estimate of the error variance σ_ϵ^2 . Therefore, the magnitude of the error variance can be used as a measure of goodness of fit. In fact, at each ROI, the goodness-of-fit measure BIC for model i (i = our model, or comparison model) is equal to

$$\text{BIC}(\text{Model } i) = (N - 3)\ln\hat{\sigma}_i^2 + 3\ln(N - 3). \quad (12)$$

Since N is the same for the two models, the model with the smaller error variance will have the smaller BIC, and, when comparing BIC scores for various models, the model with the smallest BIC score is chosen as the best. Therefore, we can conclude that our model gives a better account of the data than the boxcar model if its BIC score is lower in each ROI. Note that such a result is more impressive than it might first seem, because our model specifies the ROIs a priori. In contrast, the boxcar model makes no predictions about which brain regions should show task-related activation. In fact, in traditional exploratory fMRI data analyses, this boxcar model is applied to the data from every voxel in the brain, and those voxels for which the null hypothesis

$$H_0: \theta_C = 0$$

is rejected are classified as task sensitive.

Although there is no statistical test to determine whether the BIC score of one model is significantly better than the BIC score of another model, quantitative conclusions can nevertheless be drawn. One possibility is to use the difference in BIC scores to compute the Bayes factor, which estimates the probability that the better fitting model is the correct model of the two (see, e.g., Raftery, 1995). For example, the approximate probability that our model is correct (if one assumes that either our model or the boxcar model is correct) equals

$$P(\text{our model is correct} | \text{Data}) \doteq$$

$$\frac{1}{1 + \exp\left\{-\frac{1}{2}[\text{BIC}(\text{our model}) - \text{BIC}(\text{boxcar model})]\right\}}.$$

Thus, for example, if our model fits better by a BIC difference of 2, then the probability that our model is correct is approximately .73, whereas if the BIC difference is 10, then this probability is approximately .99.

Qualitative Tests

Another way to avoid problems caused by inaccurate models of the hrf is to test predictions of the model in a more qualitative fashion. One attractive possibility in this domain is to compute the coherence between the predictions of the model and the observed BOLD response in each ROI (Sun, Miller, & D'Esposito, 2004). Coherence is a measure of correlation, but in the frequency domain rather than the time domain. Whereas the standard Pearson correlation would be reduced if an inaccurate hrf was used to generate the model's predictions, the coherence would be unaffected. If superposition holds, then the coherence between two BOLD responses equals the coherence between the neural activations that elicited those BOLD responses. In addition, coherence at low frequencies is un-

affected by the presence of (high frequency) independent noise. Thus, coherence analysis avoids some of the most significant problems that plague standard correlational techniques.

Coherence analysis can also be used to test quantitative predictions of our model about the temporal delay between the onset of neural activations in two brain regions. This analysis uses the phase spectrum of the cross-correlation function between the BOLD responses from the two regions. Assuming that certain conditions are met, the temporal delay between the onset of the neural activations in the two regions is proportional to the slope of this phase spectrum (Sun, Miller, & D'Esposito, 2005).

Qualitative predictions about causality can also be tested by computing the Granger causality between BOLD responses in pairs of ROIs. For example, if a model predicts that activation in region A causes activation in region B, then the observed BOLD responses from regions A and B can be used in a Granger causality analysis to test this prediction statistically (Geweke, 1982).

Conclusions

fMRI has revolutionized psychology, but to date, almost all applications have been exploratory in nature. We believe that fMRI has tremendous untapped potential to provide rigorous and unique tests of computational models. As just one example, many computational models of working memory assign a key role to the head of the caudate nucleus. However, these models disagree as to whether the primary role of the caudate during working memory tasks is to provide a phasic, gating input to the thalamus (i.e., via the globus pallidus) (see, e.g., Beiser & Houk, 1998; Frank, Loughry, & O'Reilly, 2001) or a tonic, sustained input (e.g., Ashby et al., 2005; Monchi et al., 2000). Testing between these alternatives with behavioral or lesion data would be extremely difficult, but the fMRI predictions of these two classes of models are quite different. As a result, we believe the techniques proposed in this article could be used to resolve this debate.

Some new problems must be solved when fMRI is used in this confirmatory manner, but for the most part, the statistical methods required are straightforward extensions of currently popular techniques. It is important that we acknowledge that other solutions are possible. However, since there is currently no accepted method for fitting models to fMRI data, we believe that the methods described in this article represent a reasonable first solution. Most importantly, we hope that this article encourages researchers to expand their use of fMRI into the confirmatory domain.

AUTHOR NOTE

This research was supported in part by National Institutes of Health Grant R01 MH3760-2. Correspondence concerning this article should be addressed to F. G. Ashby, Department of Psychology, University of California, Santa Barbara, CA 93106 (e-mail: ashby@psych.ucsb.edu).

REFERENCES

- ANDERSON, J. R., & QIN, Y. (in press). Using brain imaging to extract the structure of complex events at the rational time band. *Journal of Cognitive Neuroscience*.
- ASHBY, F. G., ELL, S. W., VALENTIN, V. V., & CASALE, M. B. (2005). FROST: A distributed neurocomputational model of working memory maintenance. *Journal of Cognitive Neuroscience*, *17*, 1728-1743.
- ASHBY, F. G., ENNIS, J. M., & SPIERING, B. J. (2007). A neurobiological theory of automaticity in perceptual categorization. *Psychological Review*, *114*, 632-656.
- BEISER, D. G., & HOUK, J. C. (1998). Model of cortical-basal ganglionic processing: Encoding the serial order of sensory events. *Journal of Neurophysiology*, *79*, 3168-3188.
- BOYNTON, G. M., ENGEL, S. A., GLOVER, G. H., & HEEGER, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, *16*, 4207-4221.
- BROWN, J. W., & BRAVER, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, *307*, 1118-1121.
- BUXTON, R. B., WONG, E. C., & FRANK, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. *Magnetic Resonance in Medicine*, *39*, 855-864.
- CHADDERDON, G., & SPORNS, O. (2006). A large-scale neurocomputational model of task-oriented behavior selection and working memory in prefrontal cortex. *Journal of Cognitive Neuroscience*, *18*, 242-257.
- CHEN, C.-T. (1970). *Introduction to linear system theory*. New York: Holt, Rinehart and Winston.
- CLARK, V. P., MAISOG, J. M., & HAXBY, J. V. (1998). fMRI study of face perception and memory using random stimulus sequences. *Journal of Neurophysiology*, *79*, 3257-3265.
- COHEN, M. S. (1997). Parametric analysis of fMRI data using linear systems methods. *NeuroImage*, *6*, 93-103.
- DALE, A. M., & BUCKNER, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping*, *5*, 329-340.
- DANKER, J. F., & ANDERSON, J. R. (2007). The roles of prefrontal and posterior parietal cortex in algebra problem solving: A case of using cognitive modeling to inform neuroimaging data. *NeuroImage*, *35*, 1365-1377.
- FRANK, M. J., & CLAUS, E. D. (2006). Anatomy of a decision: Striato-orbitofrontal interactions in reinforcement learning, decision making and reversal. *Psychological Review*, *113*, 300-326.
- FRANK, M. J., LOUGHRY, B., & O'REILLY, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, & Behavioral Neuroscience*, *1*, 137-160.
- FRANSSON, P., KRÜGER, G., MERBOLDT, K.-D., & FRAHM, J. (1999). MRI of functional deactivation: Temporal and spatial characteristics of oxygenation-sensitive responses in human visual cortex. *NeuroImage*, *9*, 611-618.
- FRISTON, K. J., HOLMES, A. P., & ASHBURNER, J. (1999). *Statistical Parametric Mapping (SPM)*. Retrieved October 2003 from www.fil.ion.ucl.ac.uk/spm/.
- FRISTON, K. J., JOSEPHS, O., REES, G., & TURNER, R. (1998). Nonlinear event-related responses in fMRI. *Magnetic Resonance in Medicine*, *39*, 41-52.
- GEWEKE, J. (1982). Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association*, *77*, 304-313.
- GLOVER, G. H. (1999). Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage*, *9*, 416-429.
- HENSON, R., & FRISTON, K. (2007). Convolution models for fMRI. In K. J. Friston, J. T. Ashburner, S. J. Kiebel, T. E. Nichols, & W. D. Penny (Eds.), *Statistical parametric mapping: The analysis of functional brain images* (pp. 178-192). London: Academic Press.
- HINRICHS, H., SCHOLZ, M., TEMPELMANN, C., WOLDORFF, M. G., DALE, A. M., & HEINZE, H.-J. (2000). Deconvolution of event-related fMRI responses in fast-rate experimental designs: Tracking amplitude variations. *Journal of Cognitive Neuroscience*, *12*, 76S-89S.
- HUETTEL, S. A., & MCCARTHY, G. (2000). Evidence for a refractory period in the hemodynamic response to visual stimuli as measured by MRI. *NeuroImage*, *11*, 547-553.
- MASON, R. A., & JUST, M. A. (2006). Neuroimaging contributions to the understanding of discourse processes. In M. Traxler & M. A. Gerns-

- bacher (Eds.), *Handbook of psycholinguistics* (pp. 765-799). Amsterdam: Elsevier.
- MONCHI, O., TAYLOR, J. G., & DAGHER, A. (2000). A neural model of working memory processes in normal subjects, Parkinson's disease, and schizophrenia for fMRI design and predictions. *Neural Networks*, **13**, 953-973.
- OGAWA, S., LEE, T.-M., KAY, A. R., & TANK, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, **87**, 9868-9872.
- OGAWA, S., LEE, T.-M., STEPNIOSKI, R., CHEN, W., ZHU, X.-H., & UGURBIL, K. (2000). An approach to probe some neural systems interaction by functional MRI at neural time scale down to milliseconds. *Proceedings of the National Academy of Sciences*, **97**, 11026-11031.
- OLLINGER, J. M., CORBETTA, M., & SHULMAN, G. L. (2001). Separating processes within a trial in event-related functional MRI: II. Analysis. *NeuroImage*, **13**, 218-229.
- OLLINGER, J. M., SHULMAN, G. L., & CORBETTA, M. (2001). Separating processes within a trial in event-related functional MRI: I. The method. *NeuroImage*, **13**, 210-217.
- O'REILLY, R. C. (2006). Biologically based computational models of high-level cognition. *Science*, **314**, 91-94.
- PFEUFFER, J., MCCULLOUGH, J. C., VAN DE MOORTELE, P.-F., UGURBIL, K., & HU, X. (2003). Spatial dependence of the nonlinear BOLD response at short stimulus duration. *NeuroImage*, **18**, 990-1000.
- RAFTERY, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology*, **25**, 111-163.
- SUN, F. T., MILLER, L. M., & D'ESPOSITO, M. (2004). Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. *NeuroImage*, **21**, 647-658.
- SUN, F. T., MILLER, L. M., & D'ESPOSITO, M. (2005). Measuring temporal dynamics of functional networks using phase spectrum of fMRI data. *NeuroImage*, **28**, 227-237.
- VAZQUEZ, A. L., & NOLL, D. C. (1998). Nonlinear aspects of the BOLD response in functional MRI. *NeuroImage*, **7**, 108-118.
- ZARAHN, E., AGUIRRE, G., & D'ESPOSITO, M. (1997). A trial-based experimental design for fMRI. *NeuroImage*, **6**, 122-138.

(Manuscript received December 4, 2007;
revision accepted for publication March 19, 2008.)