Predicting Human Functional Maps With Neural Net Modeling

Barry Horwitz^{1*} and M.-A. Tagamets²

¹Language Section, Voice, Speech and Language Branch, National Institute on Deafness and Other Communication Disorders, National Institutes of Health, Bethesda, Maryland

²Georgetown Institute of Cognitive and Computational Sciences, Georgetown University School of Medicine, Washington, DC

Abstract: Formidable difficulties exist in interpreting positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) hemodynamic signals in terms of the underlying neural activity. These include issues of spatial and temporal resolution and problems relating neuronal activity (i.e., action potentials) measured in nonhuman studies by single unit electrodes to hemodynamic measurements reflecting synaptic activity. Also, regional hemodynamic measurements correspond to a mixture of local and afferent synaptic activity. To surmount these difficulties, we propose using large-scale neurobiologically realistic models in which data at various spatial and temporal levels can be simulated and cross-validated by multiple disciplines, including functional neuroimaging. A delayed match-to-sample visual task is used to illustrate this approach. *Hum. Brain Mapping* 8:137–142, 1999. © 1999 Wiley-Liss, Inc.

Key words: brain; positron emission tomography; functional magnetic resonance imaging; computational neuroscience; cognition

INTRODUCTION

The advent of functional neuroimaging has initiated a conceptual revolution in understanding the neural basis of human cognition. However, integrating results from the different imaging modalities (e.g., PET, fMRI, magnetoencephalography (MEG)) with one another, with observations based on neuropsychological investigations of brain-damaged patients, with electrophysiological recordings in primates performing similar tasks, and with pharmacological studies is not a trivial problem. Each technique results in data with features (e.g., spatial and temporal characteristics) that make direct quantitative comparison difficult. For example,

consider the problems associated with interpreting PET/fMRI measurements of hemodynamic activity in terms of the underlying neural activity during different components of a cognitive task. First, the spatial resolution of human brain imaging devices, even fMRI, is large compared with the size of neurons or cortical columns. Consequently, multiple and diverse neuronal populations often are lumped together in any resolvable PET or fMRI region of interest (even a single voxel), and local and afferent activities are combined into a single signal. Second, whereas the temporal dimension for neurons is on the order of milliseconds, it is on the order of a few seconds for the hemodynamic methods (because of the hemodynamic delay, this is the case even for fMRI). Therefore, important transient components of activity are not detectable by PET or fMRI. Third, the source of most information about neural activity comes from animal electrophysiological

*Correspondence to: Dr. Barry Horwitz, Language Section, VSL Branch, NIDCD, NIH, Bldg. 10, Rm. 6C414, MSC 1588, 9000 Rockville Pike, Bethesda, MD 20892. Email horwitz@helix.nih.gov
Received for publication 12 February 1999; accepted 16 June 1999

studies whose measurements generally correspond to action potentials, whereas the hemodynamic measurements most likely reflect synaptic activity to a greater extent than spike activity [Jueptner and Weiller, 1995]. Furthermore, because of this, both excitatory and inhibitory synaptic activity probably result in increased PET or fMRI activity [Jueptner and Weiller, 1995].

To overcome these problems, we [Horwitz and Sporns, 1994; Tagamets and Horwitz, 1998] and others [Arbib et al., 1995] proposed using neural modeling as a framework by which these multiple data sets can be quantitatively combined so that a conceptually coherent account of human cognition can be generated. Although this effort is just beginning, it is worth examining one such study [Tagamets and Horwitz, 1998] that attempts to relate functional neuroimaging data (PET and fMRI) to findings obtained by cellular level analyses to see the potential that neural modeling offers for elucidating functional neuroimaging data. We first describe the large-scale neural model and show how it was used to simulate a PET activation study. We then show how the model was modified so that it could be used to simulate fMRI experiments, in essence by reducing integration time and adding a hemodynamic delay.

SIMULATING A PET ACTIVATION STUDY

Our framework consisted of constructing a large-scale, neurobiologically realistic computational model incorporating multiple, interconnected brain regions such that neuronal activity in each brain region would be similar to that observed experimentally (generally from data obtained in nonhuman primate electrophysiological studies). Furthermore, the model was designed to simulate functional neuroimaging experiments of tasks similar to those used for the electrophysiological studies. A somewhat similar approach was used by Arbib et al. [1995] to examine a saccade generation task.

We chose a delayed match-to-sample (DMS) task for shape, since there exists much functional neuroimaging, neuroanatomical, electrophysiological, and cognitive data in human and nonhuman primates about this type of task. The DMS task involves the presentation of a shape, a delay, and the presentation of a second shape; the model has to decide if the second stimulus is the same as the first. Multiple trials (e.g., 10) are used to simulate a PET or fMRI study. The model incorporates four major brain regions (primary visual cortex = V1/V2; occipitotemporal cortex = V4; infe-

rior temporal cortex = IT; prefrontal cortex = PFC) that represents the ventral visual processing stream (the object vision pathway; see [Ungerleider and Mishkin, 1982]). Each region contains populations of neuronal assemblies of basic units, each of which is an interacting excitatory-inhibitory neuronal pair that represents a cortical column. As noted earlier, the limited spatial and temporal resolution of PET (and to a lesser degree fMRI) means that the neuroimaging signal from even a single voxel has contributions from the activity of multiple neuronal populations. In our model, each region (except IT) contains subpopulations with excitatory neuronal units having different response properties. In V1/V2 and V4, we have neuronal units with different orientation selectivities. In PFC, we have four different types of neuronal units whose response properties are based on the findings of Funahashi et al. [1990]: units that respond when a visual stimulus is present, two kinds of units that show activity during the delay interval, and units whose activities increase when a match between the second and first stimuli occurs. Feedforward and feedback connections between regions are based on primate neuroanatomical data. Parameters are chosen so that the excitatory elements have simulated neuronal activities resembling those found in electrophysiological recordings from monkeys performing similar tasks [e.g., Funahashi et al., 1990]. A functional neuroimaging study is simulated by presenting pairs of stimuli to an area of the model that represents the lateral geniculate nucleus (LGN). Like Arbib et al. [1995], we simulated rCBF data by integrating the absolute value of the synaptic activity over the time course of the study within the different areas for each task. For details about the parameters used in the model and a thorough discussion of all the assumptions employed, see Tagamets and Horwitz [1998].

The neuronal firings in each brain area at various times during presentation of a single test item are displayed in Figure 1. During the initial stimulus presentation, all brain regions show significant neural activity. During the delay interval, the period when the stimulus must be kept "in mind," there is significant activity in two prefrontal populations. When the second stimulus appears during the response portion of the task, neural activity again increases in all areas and a subpopulation in PFC responds only if the second stimulus matches the first. There is also a control task, where degraded forms of shapes are used as visual input to the model, but no representation has to be maintained in working memory. When the simulated

Figure 1.

Shown are excerpts from a “film” depicting the simulated neural activity at various times during one trial of the DMS task. **A.** Regions used in the large-scale neural simulation. Shown on the brain are approximate locations of the neural populations used in the simulation. (V1/V2 is shown on the lateral surface of the brain, although it represents a more medially placed region.) Each square represents 81 excitatory units (inhibitory units not shown). The units in V1/V2 correspond to neurons with vertical orientation selectivity (model units with horizontal orientation selectivity not shown); in V4, the units correspond to neurons that respond best to corners (not shown are neurons with vertical and horizontal orientation selectivity). Four classes of prefrontal neurons are shown whose simulated electrical activities are similar to those found experimentally [Funahashi et al., 1990] (they are displaced

spatially here, even though in real brains these neurons are intermingled); neurons in the posterior group of PFC respond when a stimulus is present; neurons in the two middle groups are active during the delay interval; neurons in the anterior group increase their activity if the stimulus presented following the delay is the same as that preceding the delay. **B–D.** Simulated neural activity in each brain region during various portions of one trial of the delayed match-to-sample task. **B.** A shape (shown behind the brain) is presented, and neural activity (high activity is in white, low activity in black) increases in most brain areas. **C.** During the delay period, activity in two of the prefrontal neural populations remains high. **D.** Because the second stimulus is the same as the first, a group of response neurons in prefrontal cortex show increased activity. Figure adapted from Horwitz [1998].

rCBF values of the two conditions are compared [Tagamets and Horwitz, 1998], the differences have values similar to those found in experimental PET studies of face working memory [Haxby et al., 1995].

An important feature of the model concerns how the “task instructions” are handled so that the model knows which task (DMS or control task) had to be

performed. This was accomplished by means of a continuous “attention” variable that modulates a subset of prefrontal units by diffuse synaptic inputs, the functional strength of which controls whether the stimuli are to be retained in working memory or not. Activity in each brain area, therefore, is some combination of feedforward activity determined in part by the

presence of an input stimulus, feedback activity determined in part by the strength of the modulatory attention signal, and local activity within each region.

The synaptic activities of all populations in a local brain area contributed to the simulated hemodynamic signal, as is the case in actual experimental data. But in our model we can keep track of the activities of each individual population (indeed, we can follow the dynamics of every individual neuronal unit) and thus determine how each affects the functional neuroimaging signal. Moreover, the synapses in a local cortical brain region, whose activities are summed to produce the hemodynamic signal recorded by PET or fMRI, are a mixture of synapses coming from neurons in other brain areas (studies suggest that these constitute ~10–20% of the total) [see Douglas et al., 1995] and those originating from local neurons [see Tagamets and Horwitz, 1998, for details]. These ratios were incorporated into the design of our model's architecture. Furthermore, the afferents coming from outside the local region themselves are a mixture of fibers from numerous brain regions. This makes it difficult in real data to ascertain neurobiologically why a given brain shows a change in PET/fMRI signal when task conditions are varied, during pharmacological stimulation, or when scanning individuals with neurological or psychiatric disorders. It is through modeling of the type shown above for the DMS task that one can start determining how critical any single neurobiological factor is in affecting performance and in contributing to the measured hemodynamic signals.

Besides its use in helping to assess the neurobiological factors underlying functional neuroimaging signals, computational neuromodeling can also aid PET and fMRI in assessing how neural activity becomes transformed by the multiple steps involved in its conversion into what is measured by PET and/or fMRI. We illustrate this in the next section.

SIMULATING AN FMRI STUDY

We altered our large-scale neural model so that fMRI experiments could be simulated [Horwitz et al., 1998]. Rather than integrating the absolute value of the summed synaptic activities in each brain region over the entire experiment as would be done to simulate a PET study, the integration time is reduced to 50 msec (the basic time step in the model is 5 msec), which compares favorably with the slice acquisition time for many fMRI scanners. However, fMRI must contend with the hemodynamic delay problem: the brain's vascular response to a transient change in neural activity is delayed and dispersed in time; the delay has

been estimated to be ~5–8 sec [see Bandettini et al., 1993; Friston et al., 1994]. We used a Poisson function to represent the hemodynamic delay [cf. Friston et al., 1994].

Thus to simulate an fMRI study, the absolute value of the synaptic activity is integrated over 50 msec (slice acquisition time). The resulting time series for each region can be thought of as the "gold standard"—what a noiseless, fast MRI scanner would show if there were no hemodynamic delay (other possible confounds such as nonlinearities affecting the relationship between blood oxygenation dependent signal and cerebral blood flow are ignored). Each regional time series is then convolved with the hemodynamic response function, i.e., a Poisson function, which is characterized by the parameter λ (its mean and standard deviation; units are in seconds) to produce a temporally smoothed time series. This smoothed time series is sampled every T_r sec (T_r is the repetition time) to generate the simulated fMRI time series, since in most fMRI studies it takes time to sample each slice in the imaged volume. No sources of noise other than neural noise are assumed, which implies that our simulated results represent the best that can be achieved with fMRI.

The recent interest in event-related fMRI [e.g., Courtney et al., 1997; Dale and Buckner, 1997] provides the setting for our simulation. The relatively high temporal resolution of fMRI allows for experimental designs that can try to capture different components of the cognitive tasks under study. For example, Dale and Buckner [1997] showed that selective averaging of mixed trial stimuli, similar to that done in evoked potential studies, was possible and that stimuli spaced even as close as 2 sec apart added in a roughly linear fashion in V1.

One problem with these approaches is that they are hard to validate; most studies have used simple sensory or motor tasks to evaluate their procedures, but their applicability to high-level cognitive tasks is difficult to ascertain. The problem with such tasks is that there can be extensive neural activity in multiple brain regions in the absence of external stimuli, or when no overt motor responses are employed. Moreover, much of this neural activity is not directly under experimental control. We use our modeling to illustrate some of these points, thus demonstrating the need for care in interpreting the results of event-related fMRI studies.

Figure 2A shows the simulated experimental design. A shape stimulus (S) is presented to the LGN for T sec; there is a delay of 3T sec during which the stimulus is to be retained in working memory. A second shape stimulus (S') is then presented for T sec; if it is the same

Simulated fMRI Delayed Match-to-Sample Task

Figure 2.

A. Shown is one trial of the simulated fMRI DMS and control task. S corresponds to a sample shape presented to the LGN for T sec. A delay of 3T sec occurs; the model is to retain the representation of S in working memory. A second shape (S') is presented for T sec during which the model decides if it is the same as S. Following an intertrial interval of 3T sec, a trial of the control task is presented, which uses degraded shapes (C and C'); during the control task no representation has to be maintained in working memory. **B.** Part of the time course of the absolute value of the summed (over 50

msec) synaptic activity in V1/V2 (black) and PFC (gray). The types of stimuli used for each temporal epoch are shown above several of the V1/V2 peaks; symbols are the same as in A. **C.** The same portion of the time course as in B after the integrated synaptic activity has been convolved with a Poisson function ($\lambda = 2T$) representing the hemodynamic delay. **D.** The time course obtained sampling the data in C every T sec; this represents the simulated fMRI activity. Units along the vertical scale are arbitrary.

shape as S, some frontal response neurons should increase their activity. Following an intertrial interval of 3T sec, a degraded shape stimulus (C) is presented for T sec, there is a delay of 3T sec, during which the degraded shape need not be retained in working memory; a second control stimulus (C') is presented for T sec, and there is an intertrial interval of 3T sec. This entire cycle is repeated six times, thus representing a simulated fMRI study. The attention level during the DMS task is set at 0.3 (arbitrary units), during the control task at 0.2, and during the intertrial interval at 0.05. Note that this design contains a mixture of

processing modes similar to those found in a typical imaging study. There are different time intervals that are distinguished by whether stimuli are present or absent, whether stimuli need to be retained in working memory or not, and whether or not the attention level is high. Neural activity, especially in anterior brain regions, can be high both when stimuli are present and when they are not. We assume for this example that the hemodynamic delay is characterized by a Poisson function with a λ of 2T and that the fMRI data is obtained by sampling the hemodynamically convolved (HC) time series every $T_r = T$ sec.

Figure 2B shows a portion of the time series for the absolute value of the integrated (every 50 msec) synaptic activities in V1/V2 (black curve) and PFC (gray curve), whereas Fig. 2C corresponds to the HC time series and Figure 2D to the simulated fMRI time series. The actual neural activities shown in Figure 2B show the features that one would like to capture by fMRI. For V1/V2, activity is high when a stimulus is present and low when a stimulus is not present. For PFC, activity during the times when stimuli are not present is high when the stimulus is being kept in working memory (e.g., around time = 50), lower during the delay period of the control task (e.g., around time = 58), and very low during the intertrial interval (e.g., around time = 63). For this particular example, the V1/V2 activity, although delayed and dispersed in time due to the hemodynamic delay, is discernible nonetheless in the fMRI time series. However, this is not the case for activity in PFC; the convolving together of various types of neural activity when stimuli are present with that when they are absent results in a time series that misrepresents the level of activity during the various components of the DMS task. This blurring together of the different cognitive epochs in our model is more pronounced in anterior areas than in posterior areas because of the higher neural activity in the former than the latter during the delay periods. Thus these simulated results in anterior brain areas, especially the prefrontal cortex, suggest caution in interpreting event-related fMRI in brain areas where there may be substantial neural activity, some not under experimental control, when stimuli are not present.

CONCLUSIONS

The centrality of functional brain imaging in linking neurobiological function to cognitive function has become increasingly evident. The spatial and temporal richness of functional brain imaging data sets, however, compels the need for neurocomputational modeling to help interpret their neuroscientific meaning. The large-scale neural model of shape working memory that we presented here provides an illustration of this approach. We think that such large-scale models offer a way by which multiple observations (e.g., results from

lesion studies, MEG data, fMRI data, nonhuman electrophysiological recordings) concerning specific cognitive tasks can be integrated into a coherent understanding of how the brain's activity relates to observed behavior.

ACKNOWLEDGMENTS

We thank Dr. Sabine Kastner for a critical reading of the manuscript.

REFERENCES

- Arbib MA, Bischoff A, Fagg AH, Grafton ST. 1995. Synthetic PET: analyzing large-scale properties of neural networks. *Hum Brain Mapp* 2:225–233.
- Bandettini PA, Jesmanowicz AJ, Wong EC, Hyde JS. 1993. Processing strategies for time-course data sets in functional MRI of the human brain. *Magn Reson Med* 30:161–173.
- Courtney SM, Ungerleider LG, Keil K, Haxby JV. 1997. Transient and sustained activity in a distributed neural system for human working memory. *Nature* 386:608–611.
- Dale AM, Buckner RL. 1997. Selective averaging of rapidly presented individual trials using fMRI. *Hum Brain Mapp* 5:329–340.
- Douglas RJ, Koch C, Mahowald M, Martin KAC, Suarez HH. 1995. Recurrent excitation in neocortical circuits. *Science* 269:981–984.
- Friston KJ, Jezzard P, Turner R. 1994. Analysis of functional MRI time-series. *Hum Brain Mapp* 1:153–171.
- Funahashi S, Bruce C, Goldman-Rakic PS. 1990. Visuospatial coding in primate prefrontal neurons revealed by oculomotor paradigms. *J Neurophysiol* 63:814–831.
- Haxby JV, Ungerleider LG, Horwitz B, Rapoport SI, Grady CL. 1995. Hemispheric differences in neural systems for face working memory: a PET-rCBF Study. *Hum Brain Mapp* 3:68–82.
- Horwitz B. 1998. Using functional brain imaging to understand human cognition. *Complexity* 3:39–52.
- Horwitz B, Bertelson J, Beauchamp M, Tagamets M-A. 1998. A large-scale neural model linking local neuronal dynamics to fMRI data. *NeuroImage* 7:S769.
- Horwitz B, Sporns O. 1994. Neural modeling and functional neuroimaging. *Hum Brain Mapp* 1:269–283.
- Jueptner M, Weiller C. 1995. Does measurement of regional cerebral blood flow reflect synaptic activity? Implications for PET and fMRI. *NeuroImage* 2:148–156.
- Tagamets M-A, Horwitz B. 1998. Integrating electrophysiological and anatomical experimental data to create a large-scale model that simulates a delayed match-to-sample human brain imaging study. *Cereb Cortex* 8:310–320.
- Ungerleider LG, Mishkin M. 1982. Two cortical visual systems. In: Ingle DJ, Goodale MA, Mansfield RJW (eds): *Analysis of Visual Behavior*. Cambridge: MIT Press, pp 549–586.