

A. Specific Aims

Recent technological improvements in functional Magnetic Resonance Imaging (fMRI) are making it possible to study the brain as more than a collection of volume elements (voxels) but rather as a *system* of interacting components. Instead of individual regions, we can study functional networks. Instead of computing voxels' individual response curves, we can estimate their collective response to a stimulus. Instead of settling for responses averaged over brain regions, we can image fine spatial structure. Such a system-oriented approach requires advances in both imaging and statistical methodology.

This project consists of two intertwined components. First, we will perform fMRI experiments to address three questions about the representation of space in the human brain. Second, we will develop and validate new statistical techniques for the systems-level inferences needed to answer the neuroscientific questions. These techniques are motivated by and developed for the proposed experimental studies, but with minor adaptation, they will be broadly applicable to other neuroimaging studies.

Question 1. What is the functional network for visual remapping?

Aim 1. Develop methods to identify and characterize distributed functional networks.

Our remapping experiments to date have identified a number of cortical areas in the human brain that exhibit remapping. The proposed experiments will test functional interactions between these areas. The hypothesis is that topographic subregions in each area will become tightly correlated during remapping. We predict that the particular subregions forming the network will depend on (a) the location of the stimulus, and (b) the trajectory of the saccade. We will identify functionally coherent regions and estimate the dependence among these regions by combining neighborhood-discovery methods with a new approach to causal-state reconstruction for discrete time series. We will begin by discretizing the voxel time series, estimate the causal states of the resulting process and in turn the mutual information between each pair of voxel time series. We will adapt recent neighborhood discovery algorithms using the mutual informations to identify functional networks and to characterize how these networks evolve with time.

Question 2. Does visual cortex contain a representation of eye position?

Aim 2. Develop methods for simultaneously estimating fMRI gain fields.

Neurons in visual cortex respond to stimuli at a particular location on the retina, but their response is modulated by non-visual signals, such as the position of the eyes in the orbit. Single-unit studies in monkeys have characterized the influence of eye position as *gain fields* in which the amplitude of visual activity varies smoothly as a function of eye position. Following on these studies in monkeys, we will test the influence of eye position on visual activity in the human brain. We will use a novel experimental approach that presents deterministic “noise” at every location in the visual field. Using sparse regression and stochastic search techniques, we will estimate the voxels' functional receptive field – the part of the visual field to which the voxel responds – simultaneously and then estimate how these fields change in shape, location, and amplitude of the receptive fields – as a function of eye position and other covariates.

Question 3. Do eye position signals have a fine-scale structure?

Aim 3. Develop adaptive spatial smoothing techniques for high-resolution fMRI.

Columnar organization is a fundamental principal in the mammalian cerebral cortex. The proposed experiments will use very high resolution fMRI to test the fine-scale organization of eye position modulation in visual cortex. High resolution data combines low signal-to-noise with idiosyncratic noise distributions, making it especially valuable to borrow inferential strength across voxels with some sort of smoothing. But critically, our objective is to recover *fine-scale* spatial structure which tends to be washed out by standard smoothing techniques because they average over locations with disparate signals. We propose instead to develop adaptive smoothing techniques that use the data to determine how voxels are to be combined. We predict substantially improved performance with high-resolution data, and a slightly modified approach should improve performance for standard fMRI data sets as well.

B. Background and Significance

The new statistical methods in this proposal address basic, recurrent problems in fMRI data analysis, motivated by issues we have faced in previous fMRI studies. Our three neuroscience questions address fundamental issues regarding neural mechanisms that underlie sensorimotor transformations in humans. The questions stem directly from our previous fMRI studies conducted during the previous grant period. These studies have informed and guided the design of the proposed experiments.

The proposed statistical methods and neuroscience experiments are strongly linked. The methods will be developed in parallel with the experiments guided by our results at each stage, and our plan includes data collection explicitly for testing and validating the methods. The experiments draw heavily upon the new statistical methods both for their design and analysis.

Before giving specific background on our statistical and scientific approaches, we provide some basic terminology for fMRI and visual neuroscience. An fMRI data set consists of a sequence of three-dimensional images obtained at regular intervals in an MR scanner while the subject performs a carefully arranged sequence of behavioral tasks. Each image is a three-dimensional array of volume elements or *voxels*, usually arranged in two-dimensional *slices*. We refer to the time series of measurements at each voxel as the voxel’s *time course* or *time series*. Methods exist to map the data from the voxel coordinates to the surface of the cortex, where all the action is. This mapping is visualized by inflated to the cortical surface into a sphere. We call these coordinate systems *voxel* and *cortical surface* coordinates, respectively. Concentrated neural activity gives rise to a blood-flow (hemodynamic) response that can be detected as small, systematic changes in the time course. We refer to the underlying changes as the *response* or the *signal* and the remaining part of the time course as the *noise*, which comprises several sources of variation. A *saccade* is a rapid eye movement. Saccades are made with high frequency during visual exploration to bring the high acuity region of the retina, the *fovea*, onto objects of interest. A visual *receptive field* (RF) is the part of space which causes a neuron to fire. The visual system is composed of a number of different *regions* or *areas* that are organized hierarchically. The visual system is crossed so that neurons in the left hemisphere have receptive fields in the right side of space, referred to as the *right visual field*, and vice versa.

B.1. COMMUNITY DISCOVERY METHODS FOR FMRI

Neuroscience distinguishes two types of connectivity between neural units, be they neurons, microcircuits, or larger aggregations: *anatomical* connectivity, defined by physical connections such as synapses and fibers, and *functional* connectivity, defined through coordinated behavior, and the associated sharing of dynamical information (Sporns et al., 2002; Friston, 2002). These need not map neatly onto each other, and the same anatomical system can support many different patterns of functional organization in different dynamical or behavioral regimes. This picture, of specialized anatomical regions supporting multiple patterns of functional connectivity, originates in experimental work on neuronal networks (Selverston and Moulins, 1987) and in lesion studies (Luria, 1973; Cytowic, 1996), and qualitatively agrees with brain imaging results (Friston, 2002).

Quantitatively, functional connectivity has usually been operationalized as the correlation between simultaneous or slightly lagged activity (David et al., 2004); this is inadequate, because it misses nonlinear relationships and dependence between temporally-extended patterns of activity. Using mutual information (Sporns et al., 2002) handles nonlinearities, but still misses temporal extension. Moreover, little attention has been given to the systemic pattern of functional connectivity, despite its acknowledged importance for cognitive architecture.

We can solve these problems by using adaptive state-space modeling to capture extended patterns of activity, using mutual information estimation among states to measure functional connections, and giving system-level summaries of functional connectivity by adapting new techniques for decomposing networks into nearly-decoupled sub-networks. The first phase, of state-space modeling, draws on methods of nonlinear prediction (Shalizi and Crutchfield, 2001; Shalizi and Klinkner, 2004); the second on work on estimating mutual information and functional connectivity measurement (Paninski, 2003; Klinkner et al., 2006); and the third, on new approaches from the statistical mechanics of complex networks, known as “community discovery” or “module discovery” (Newman and Girvan, 2003; Newman, 2006; Reichardt and Bornholdt, 2004; Reichardt and Bornholdt, 2006a; Shalizi et al., 2007).

The basic idea is as follows (Shalizi et al., 2007). (For concreteness, we speak of voxels, but the method applies to any sort of dynamical network or spatio-temporal data.) The starting point is the strong notion of “state” used

in physics and dynamics: the state of the system is a variable which determines the distribution of all present and future observables. In statistical terms, the state is a minimal sufficient statistic for predicting future observations (Shalizi and Crutchfield, 2001), and it always evolves according to a homogeneous Markov process (Knight, 1975; Shalizi and Crutchfield, 2001). The observations, in general non-Markovian, are noisy, nonlinear functions of the state. In symbols, each voxel i has an associated time-series of observations, $X_i(t)$. This is generated as a noisy, non-linear function of a Markovian state process $S_i(t)$ for that voxel, which in turn is coupled to the observations:

$$\begin{aligned} X_i(t) &= Q_i(S_i(t), \epsilon_i(t)) \\ S_i(t+1) &= T_i(S_i(t), X_i(t)) \end{aligned} \quad (1)$$

where the innovation process $\epsilon_i(t)$ is a suitable source of noise, unpredictable from the pasts of $X_i(t)$ and $S_i(t)$. The function Q is the *measurement function* or *emission function*, effectively giving the distribution of observations conditional on the state. The latent states $S_i(t)$ control the future distribution of observations via Q , and so the states summarize temporally extended, statistically-reproducible patterns of behavior. The function T is the *transition function*, which simultaneously represents the dynamics in the latent state space, and shows how observing a particular outcome, $X_i(t)$, leads to an updating of the optimal prediction for the next time step, $S_i(t+1)$. Models of the form (1) differ from conventional hidden Markov models (HMMs), which would take the form

$$\begin{aligned} X_i(t) &= Q_i(S_i(t), \epsilon_i(t)) \\ S_i(t+1) &= T_i(S_i(t), \eta_i(t)) \end{aligned}$$

where η and ϵ are *independent* noise sources. This means that the latent state in a conventional HMM is *not* a sufficient statistic, since it is not a function of the observable history at all. Except in the special case of linear-Gaussian HMMs, this leads to great difficulties in estimating the state. Partially-observable Markov chain models of the form (1) avoid this, with no loss of predictive or representational power (Shalizi and Crutchfield, 2001; Knight, 1975).

The model (1) allows for dependence between voxels i and j , through the innovations ϵ_i and ϵ_j . We measure this dependence by the **informational coherence** (Klinkner et al., 2006) between the voxels,

$$IC_{ij} \equiv \frac{I(S_i, S_j)}{\min H(S_i), H(S_j)} \quad (2)$$

where $I(S_i, S_j)$ is the mutual information shared by S_i and S_j , and $H(S_i)$ is the self-information (Shannon entropy) of S_i . Specifically, $I(X, Y) = \sum_{x,y} P(X=x, Y=y) \log \frac{P(X=x, Y=y)}{P(X=x)P(Y=y)}$, the Kullback divergence between the joint distribution and the product of the marginals, and that $H(X) = -\sum_x P(X=x) \log P(X=x)$ (Cover and Thomas, 1991). Since $I(S_i, S_j) \leq \min H(S_i), H(S_j)$, this is a symmetric quantity, normalized to lie between 0 and 1 inclusive. The IC_{ij} values for each pair of voxels i, j are then inputs to a module-discovery procedure (Reichardt and Bornholdt, 2004; Reichardt and Bornholdt, 2006a) which identifies clusters of voxels whose interconnections are stronger than would be expected in a randomized system without modular structure, and in fact delivers nested hierarchies of such clusters. (We explain this in more detail below.) Comparison of discovered modules and their hierarchies across different behavioral regimes allows us to identify changes in functional connectivity.

A key step is to find and estimate a model of the form (1) for each voxel. When the observations X_i are discretized, and the number of latent, predictive states is known (or postulated) to be finite, this can be done through an adaptive, non-parametric model discovery procedure, specifically the CSSR algorithm introduced in (Shalizi and Klinkner, 2004). This is related to the “state space reconstruction” technique for nonlinear dynamical systems (Kantz and Schreiber, 1997), but unlike conventional reconstruction, it applies to stochastic state processes. So far as we know, CSSR is currently the only stochastic state reconstruction algorithm which has been proved statistically consistent (for conditionally stationary discrete sequences, a restriction we discuss later).

CSSR rests on the strong notion of “state” introduced above, which implies testable conditional independence properties. The past of a voxel’s time course up to time t , $X_i(-\infty : t)$, is independent of its future, $X_i(t+1 : \infty)$, conditional on the state at t , $S_i(t)$: in symbols, $X_i(-\infty : t) \perp X_i(t+1 : \infty) | S_i(t)$. Given a candidate set of states, then, CSSR can check its adequacy by seeing whether all the requisite conditional independence relations hold in the training data, i.e., cannot be rejected at a level α . When one or more screening-off relations fail, the corresponding states are sub-divided, and their parts assigned to new states, or merged with existing states if

they have the same predictive distributions. This procedure starts with a single state, corresponding to an IID model, and continues until screening-off has been checked for all history segments up to some pre-set length L . The state-reconstruction process is a kind of smoothing of the mapping from time-series histories to predictive, conditional distributions, starting from a maximally smooth (constant) function, and adaptively allowing more or less roughness in different parts of the history space.

This sounds similar to the idea of fitting finite-order Markov models up to order L by likelihood-ratio tests (Billingsley, 1961). A more exact analogy is fitting a variable-length Markov chain (VLMC; (Bühlmann and Wyner, 1999)), which uses a CART-like procedure to grow a prediction tree with a maximum depth L . (VLMCs with depth L are formally equivalent to order- L Markov chains, but more efficiently parameterized.) This impression is misleading. The merging operation in CSSR means that the prediction structure it learns can contain loops, rather than having to be a tree, and so it can include infinite-order Markov chains. The class of models which can be learned exactly by CSSR strictly includes all VLMCs, as well as processes which have only infinite VLMC representations (Shalizi and Klinkner, 2004), laying on them a curse of dimensionality which CSSR escapes.

If the data come from a process of this form with finitely many states, L is large enough to allow for a transition from every state to every other state, and $\alpha \rightarrow 0$ sufficiently slowly, then the probability that CSSR reconstructs the wrong state model goes to zero. (For a sketch proof, see (Shalizi and Klinkner, 2004); full details are forthcoming in (Shalizi et al., 2008).) Under these conditions, CSSR thus converges to the time series' optimal nonlinear predictor. Work is underway on establishing CSSR's properties when these assumptions are violated, particularly when the true number of states is infinite and L is allowed to grow with the data; simulation studies suggest that, if L grows temperately, CSSR finds a sequence of increasingly close finite-state approximations to the infinite-state optimal nonlinear predictor. (See D.1.1.)

Dependence between the states of two voxels' optimal predictors means that their time series share information which is relevant to their future, so that their behavior is coordinated. Mutual information quantifies this dependence. Accurate estimation of mutual information, particularly from fairly short time series, is non-trivial. It is well-known (Victor, 2000) that the standard, plug-in estimate is both biased and noisy. We currently use the Paninski estimator (Paninski, 2003), which optimizes a bound on the total estimation error.

Informational coherence is not the only possible way of measuring behavioral coordination, or functional connectivity. However, it has a number of advantages over others (Klinkner et al., 2006). Unlike measures of strict synchronization, which insist on voxels doing exactly the same thing at exactly the same time, it accommodates phase lags, phase locking, chaotic synchronization, etc., in a straightforward and uniform way. Unlike cross-covariance or spectral coherence, it easily handles nonlinear dependencies, and does not require the choice of a particular lag (or frequency, for coherence), because the predictive states summarize the entire relevant portion of the history. Generalized synchrony measures (Quian Quiroga et al., 2002) can handle nonlinear relationships among states, but inappropriately assume determinism. Finally, mutual information among the observables, $I[X_i; X_j]$, can handle nonlinear, stochastic dependencies, but suffers, especially in neural systems, because what is really of interest are coordinated *patterns* of behavior over time, rather than coordinated instantaneous actions. Because each predictive state corresponds to a unique statistical pattern of behavior, mutual information among these states is the most natural way to capture functional connectivity. One could as an alternative employ spatio-temporal autoregressive models and impose a sparsity constraint to obtain a network structure (cf. D.2.1); however, it is not at all clear that sparsity is an appropriate constraint, or that it can accommodate long-range functional connectivity.

Having obtained the IC_{ij} matrix, the final step is the use of the Reichardt-Bornholdt community discovery algorithm (Reichardt and Bornholdt, 2004; Reichardt and Bornholdt, 2006a), which assigns voxels to communities (clusters) in such a way that within-community links are stronger than would be expected under randomization, and between-community links are weaker. Specifically, the objective function to be maximized is $H_\gamma(\vec{\sigma}) = \sum_{i \neq j} (IC_{ij} - \gamma K_i K_j / M) \delta(\sigma_i, \sigma_j)$ where $\vec{\sigma}$ is the discrete vector giving the community assignments of all voxels, IC_{ij} is the informational coherence between voxels i and j , K_i is the average coherence of voxel i with all other voxels, M is a normalizing factor, and $\gamma \geq 0$ is a scale factor. (The optimization is currently done by simulated annealing.) When $\gamma = 1$, intra-community links must be just stronger than expected under randomization. As γ shrinks, communities merge into each other, and the coherence of each community weakens. As γ grows, communities split, and each community becomes more coherent — more functionally connected — internally.

The results of our analyses, then, are not just the matrix of informational coherences IC_{ij} measuring functional connectivity between voxels, but also a division of the brain into a hierarchically-organized collection of

functionally-coherent communities, which need not be geometrically contiguous. This over-all pattern of functional organization can be compared across experimental conditions.

B.2. VISUAL REFERENCE FRAMES IN HUMANS

We move our eyes nearly three times every second. Each eye movement causes the image projected onto the retina to shift. While we perceive a stable image of the world, the retina receives a chaotic barrage of ever-changing snapshots. How do our brains create from this a stable perception? More than a century ago, Helmholtz (1866) proposed that the world appears to stay still when the eyes move because the “effort of will” involved in generating a saccade simultaneously adjusts perception to account for the eye movement. A simple experiment supports the essence of Helmholtz’s explanation: We see the world move when the retina is displaced by pressing on the eye but are generally oblivious to retinal shifts from saccades. This perceptual stability has long been understood to imply that we see not a direct impression of the external world but a construction derived from an internal representation. The experiments in this proposal address the neural mechanisms by which the brain constructs this internal representation of space.

Neurophysiological studies in monkey parietal cortex have described two mechanisms by which the brain may compute the coordinates of objects in the world when the eyes move: visual remapping and eye position gain fields (Colby and Goldberg, 1999; Andersen and Buneo, 2002). While these two mechanisms have been studied extensively at the level of single neurons in monkeys, little is known about the either mechanism in humans. We have studied both of these mechanisms in the current grant period. Our initial work focused on human parietal cortex. We investigated what happens to visual information in the parietal cortex when subjects move their eyes. Parietal cortex is a logical place to investigate perceptual stability because it receives information from both the visual system and the motor systems that control eye movements. Neurons in parietal cortex encode visual information, and visual stimuli that appear within their receptive fields elicit strong responses. Parietal RFs are defined with respect to the location of gaze: they are retinotopic and move with the eyes. When the eyes move so that the RF of a neuron lands on a previously stimulated location, the neuron fires as though a stimulus were still present, even though the screen is blank (Duhamel et al., 1992). This response to the trace of a previous stimulus indicates that parietal neurons update the internal representation of visual space.

C. Progress Report and Preliminary Results

C.1. PROGRESS REPORT

Work during the previous grant period produced 12 invited presentations and 10 refereed publications in top journals in both statistics and neuroscience. Here we review a selected set of accomplishments from this work.

Adaptive Confidence Bands (Genovese and Wasserman 2007). In this paper, we show that for a variety of commonly used nonparametric regression models, it is impossible to construct confidence bands that adapt to the data. That is, we would like a confidence band procedure to produce narrower bands automatically when the truth is smooth and wider bands when it is rough, analogously to adaptive estimators (e.g., wavelet shrinkage). But this turns out to be impossible; in fact, the width of the bands is driven completely by the roughest functions allowed by the assumptions. We show that it is possible, however, to construct adaptive confidence bands if we extend the notion of coverage to include selected surrogates for the unknown function. The paper develops a procedure that is optimal to within a small constant factor. The result for smoothing is that we can make rigorous adaptive inferences for functions with a small dilation of the bands from standard techniques. Put another way: by slightly dilating the bands, we can ignore bias and still get correct inferences. The implication for fMRI is that we can construct sharp bands for the hemodynamic response that enable us to make a variety of useful inferences.

Exceedance Control for False Discovery Control (Genovese and Wasserman 2006b). Many methods for multiple testing problems have been developed to control the False Discovery Rate (FDR), the expected proportion of falsely rejected null hypotheses among all rejections. For much the same reason that one does not often use 50% confidence intervals, it can be valuable to control instead the probability that the FDP exceeds a specified bound, which we call False Discovery Exceedance (FDX) control. This paper constructs a general class of methods for FDX control based on inverting tests of uniformity. The new method, which we call adaptive FDX control, maintains the same uniform control on FDX while attaining substantially higher power. It also remains valid under arbitrary

dependence among test statistics, which is a significant advance. In comparisons with competing methods for handling dependence while controlling FDX, our method performed best.

FDR Control with P-value Weighting (Genovese, Roeder, and Wasserman 2006). This paper shows how prior information can be used to improve the power of False Discovery controlling procedures in multiple testing problems. We show how to use p-value weighting to improve the power of FDR and False Discovery Exceedance controlling methods. The power improvement is substantial when the weights are associated with the correct pattern of alternatives but power does not reduce much even when the weights are poorly chosen.

Scan Clustering: A False Discovery Approach (Perone Pacifico, Genovese et al. 2006). This paper puts the problem of testing in spatial data such as fMRI into a scan-clustering framework. For example, fMRI data tend to be spatially smoothed at fixed bandwidth, with little attention to the bias-variance tradeoff this produces. Building on our earlier work (Perone Pacifico et al. 2004), consider smoothing at multiple scales simultaneously and construct statistically efficient procedures for detecting activity based on combining information across scales. The procedure has strong theoretical performance and can be efficiently computed. In our studies, the technique successfully improved resolution of the active sets.

A Stochastic Process Approach to FDR Control (Genovese and Wasserman 2006a). This paper extends the theory of false discovery rates (FDR) pioneered by Benjamini and Hochberg (1995). We develop a framework in which the False Discovery Proportion (FDP) – the number of false rejections divided by the number of rejections – is treated as a stochastic process. After obtaining the limiting distribution of the process, we demonstrate the validity of a class of procedures for controlling the False Discovery Rate (the expected FDP) and construct a confidence envelope for the whole FDP process. From these envelopes we derive confidence thresholds, for controlling the quantiles of the distribution of the FDP as well as controlling the number of false discoveries.

False Discovery Control for Random Fields (Perone Pacifico, Genovese et al. 2004). This paper extends False Discovery Rate control to random fields. false discovery criteria and construct confidence envelopes for the proportion of false discoveries as a function of threshold. Of particular note, we develop a method for finding regions in the field's domain where there is a significant signal while controlling either the proportion of area or the proportion of clusters in which false rejections occur. This is the first work to extend the notion of error control to *regions* rather than locations (i.e., voxels).

Remapping in Human Parietal Cortex (Merriam, Genovese, and Colby, 2003). Our first question was whether we would observe activity in human parietal cortex during the same task that elicits remapping in monkey parietal cortex. We addressed this question with a modified version of the classic single-step task used in neurophysiological studies (Fig. 1A). In this task, an intervening eye movement brings the recently-stimulated screen location into the opposite hemifield. We found that an area in the medial portion of the intraparietal sulcus exhibits strong responses to the remapped stimulus trace (Fig. 1B). The time course of activity from this region indicates that the response occurs in conjunction with the saccade, 1,000 ms after the onset of the stimulus, as we would expect during remapping (Fig. 1C). Finally, equivalent activity was not observed in two conditions that controlled for purely visual and motor components of the task, indicating that both a stimulus and a saccade are required to produce remapping.

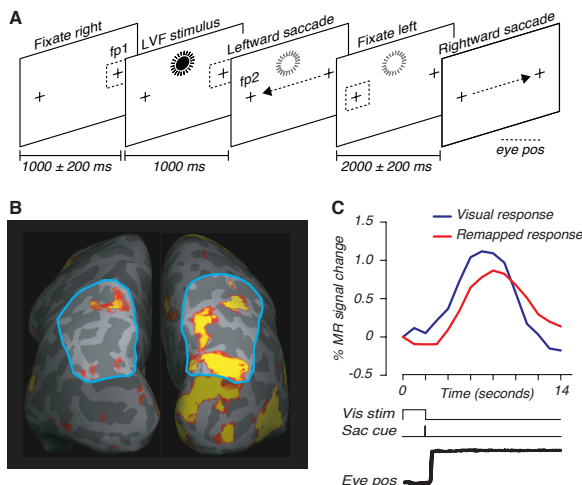


Figure 1. Remapping in human parietal cortex. **A**, Single-step task. Subjects fixate FP1, located 8° . After $1,000 \pm 200$ ms, a stimulus (full circle) appears in the left visual field for 1,000 ms, activating contralateral (right) hemisphere. The stimulus is then extinguished and a tone cues the subject to make a leftward eye movement to FP2, remapping the trace of the now-extinguished stimulus (empty circle) into the right visual field. **B**, Activation from an example subject in the single-step task on trials in which a LVF stimulus was followed by a leftward saccade, resulting in widespread activation in contralateral (right) hemisphere. Activation was also observed in ipsilateral (left) parietal lobe, indicating that the visually-evoked activation was remapped in conjunction with the eye movement. **C**, Time course of activation evoked by visual (blue) and remapped responses (red). From Merriam, Genovese, and Colby, 2003.

Strength of Remapping is Graded in Visual Cortex (Merriam, Genovese, and Colby, 2007). If remapping is important for perceptual constancy, remapping should not be limited to brain regions with attentional and oculomotor functions. Rather, updated spatial information should reach cortical areas involved in visual perception. To test the role of human extrastriate visual cortex in remapping, we first identified the boundaries of several occipital lobe visual areas using standard retinotopic mapping techniques. We then tested subjects while they performed the single-step remapping task. We found that the strength of remapping was roughly monotonic with position in the visual hierarchy: remapped responses were largest in areas V3A and hV4 and smallest in V1 and V2 (Fig. 2). This demonstrates that remapping is widely distributed in the human brain. Furthermore, updated visual representations are present in primary sensory areas that are directly linked to visual perception.




Figure 2. Proportion of in which responses reached a posterior probability threshold of $q < 0.95$. Light gray bars, responses to ipsilateral visual stimuli during stimulus-only fixation task (stim). Medium gray bars, responses to saccades in the absence of salient visual stimuli (sac). Dark gray bars, responses in the single-step task when a visual stimulus appears in the ipsilateral visual field and is followed by an ipsiversive saccade (sstep). The prevalence of voxels activated by the single-step task increases with position in the visual hierarchy. From Merriam, Genovese, and Colby, 2007.

Visual Cortex is Retinotopic, Not Spatiotopic (Gardner et al., 2008). With each eye movement, objects in the world change position on the retina, but almost all visual areas encode space in a retinotopic reference frame—the cortical representation of a point in the visual world is different across eye positions. Remapping provides one mechanism by which the visual system may compute the location of stimuli from an ever-changing retinal image (Merriam and Colby, 2005). A second mechanism would be to construct an explicit representation in external, *spatiotopic* coordinates that is invariant to changes in eye position (Colby, 1998; Snyder et al., 1998; Duhamel et al., 1998; Bremmer et al., 2001; Avillac et al., 2005). It was recently reported that human cortical area MT represents stimuli in a spatiotopic reference frame (d’Avossa et al., 2007). We used visuotopic mapping with fMRI to define 12 human visual cortical areas, and then determined whether the reference frame in each area was spatiotopic or retinotopic. We scanned subjects on three different fixation conditions ($-10^\circ, 0^\circ, +10^\circ$, relative to center). Visual stimuli, consisting of black and white moving dots, appeared for 3 s, in a pseudorandomized order, at each of four screen locations. ($-15^\circ, -5^\circ, +5^\circ, +15^\circ$, from the center of the screen). Between each stimulus presentation the screen was gray for 3–9 s.




Figure 3. Retinotopic responses from left hemisphere MT. Mean responses were computed for all voxels in which there was a significant visual response in at least one of three fixation conditions ($p < 0.01$, permutation test). Columns 1–3 plot the responses for fixation conditions at $-10^\circ, 0^\circ$, or $+10^\circ$, relative to the center, respectively. Purple, red, orange, and blue traces show the responses to visual stimuli presented at $-15^\circ, -5^\circ, +5^\circ$, and $+15^\circ$ from the center of the screen, respectively. From Gardner, Merriam, Movshon, and Heeger, 2008.

We found that all 12 areas, including MT, represented stimuli in a retinotopic reference frame. Example responses are shown from left hemisphere MT from a single subject (Fig. 3). When the eyes were fixated to the left of the screen, the three stimuli located to the right of fixation elicited strong responses (left panel); during central fixation, the two stimuli to the right of fixation elicited strong responses (center panel). When the eyes were to the right of the screen, only the one stimulus to the right of fixation elicited a response (right panel). Although there were patches of cortex in and around these visual areas that were ostensibly spatiotopic, responding to the same set of stimulus locations regardless of eye position, none of these patches exhibited reliable stimulus-evoked responses. We concluded that the early, visuotopically organized visual cortical areas in the human brain represent stimuli in a retinotopic reference frame.

C.2. PRELIMINARY RESULTS

C.2.1. Remapping Increases Inter-Area Interactions (Aim 1)

The spatial transformation that underlies remapping depends on information about (1) the size and direction of the eye movement, and (2) the retinal coordinates of the visual stimuli that will be remapped (Quaia et al., 1998). How are these two sources of information integrated? One possibility is that information is transmitted from visual and oculomotor areas to parietal cortex via long-range cortico-cortical projections. This implies that communication between cortical regions increases during remapping. fMRI provides an ideal methodology for testing this prediction because it is possible to measure brain activity from the entire cortex simultaneously.




Fig. 4. Inter-area correlations increase during remapping. Correlation between IPS and V3A in the single-step task on trials that involve remapping (**A**), and trials that do not (**B**). Each dot is a time point taken from the fMRI time series after removing activity driven by the stimulus and the saccade. **C**, Matrix of correlation differences between trials that involve remapping and those that do not.

We have begun to investigate functional interactions between cortical areas during remapping (Fig. 4). We divided trials of the single-step task into (1) trials in which the stimulus was remapped into the contralateral visual field (where we expect remapping), and (2) trials in which the stimulus was remapped away from the contralateral visual field (where we do not expect remapping). We then computed a series of correlations between every pair of ROIs. Finally, for each ROI pair, we took the difference in correlation between the two trial types. We found that remapping enhances correlation between IPS and each of the other visual areas tested. Remapping does not, however, have a significant effect on correlations between the other visual areas. This observation suggests that parietal cortex is central to the computations necessary for remapping.

C.2.2. Measuring Functional Connectivity with State Reconstruction and Mutual Information

To assess the feasibility of measuring functional connectivity via mutual information and state-space reconstruction, we took a pre-existing fMRI data set of a subject performing a remapping task, of the kind described in Section D.1.2. We used a single run of the experiment, of total duration 500 s, with data acquired at 1 ms intervals. Space was discretized by the fMRI reconstruction in a $64 \times 64 \times 16$ grid, including a substantial number of voxels containing no brain tissue.

We first discretized the time-course each voxel by estimating its time derivative with a local polynomial fit, and binning the derivatives by quintiles. CSSR then reconstructed a state-space model for each voxel's discrete time series. (Control settings for CSSR were picked, heuristically, with BIC; see D.1.1 for directions for more rigorous future methods.) We then, for each pair of voxels, calculated the mutual information between the reconstructed states. We display a sub-set of these results in Fig. 5, namely the functional connectivity between the voxel at coordinates (27,34,13) and every other voxel in the brain. We did not use any form of masking. Rather, CSSR reconstructed one-state, IID models for all voxels outside the brain, though it had no *a priori* reason to treat them any differently than any other voxels; in effect it discovered that there was no signal worth predicting there. This in turn automatically lead to an estimate of zero functional connectivity.




Figure 5. Measuring functional connectivity with mutual information. Color values indicate strength of functional connectivity between a right parietal lobe seed voxel and all other voxels in the brain. Widespread connectivity is evident in the 2D slice (**A**), on an inflated representation of the cortical surface (**B**), and on a computationally flattened representation of the right hemisphere (**C**).

C.2.3. Eye Position Modulates the Gain of Visual Activity (Aim 2)

Our studies of remapping highlight one mechanism for the dynamic integration of eye movement and visual signals. Computational models have described an alternative mechanism based on multiplicative interactions between visual and oculomotor signals (Avillac et al., 2005; Pouget and Sejnowski, 1997). According to this account, static eye position modulates the response gain of visual neurons, without affecting their spatial selectivity. While there is substantial evidence in support of eye position gain modulation in monkeys (Boussaoud and Bremmer, 1999; Andersen et al., 1990a; Snyder, 2000), little is known about this mechanism in humans. We tested for gain modulation in humans using a simple rotating wedge stimulus, based on our previous work in human area MT (Huk et al., 2002). On each run, subjects viewed the stimulus while fixating one of eight possible screen locations arranged around the perimeter of the screen ($-11^\circ, 0^\circ, +11^\circ$ horizontal, and $-5^\circ, 0^\circ, +5^\circ$ vertical, excluding screen center). A large proportion of voxels ($> 75\%$) in area V3A exhibited systematic changes in response amplitude (Fig. 6). In most cases, the effect of eye position was well fit a planer function (Fig. 6, center) that peaked in one of the four quadrants of visual space (up and to the right in the example shown).

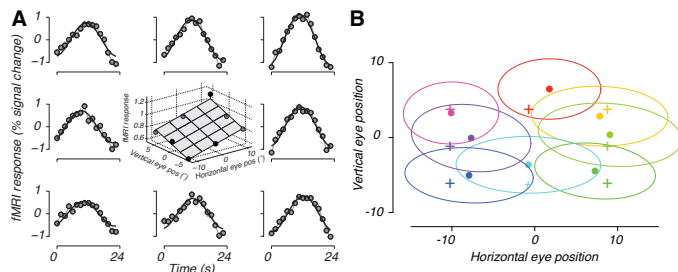


Figure 6. Eye position modulates the amplitude of visual responses. **A**, Time series plots show the cortical response in a single V3A voxel for each of the eight eye positions tested. Thin gray lines and dots indicate the fMRI data averaged over 20 trials. Thick black lines indicate the best fitting sinusoidal model (amplitude and phase) for that response. For this voxel, the response was largest when the subject fixated up and to the right and was smallest when the subject fixated down and to the left. Center plot illustrates the plane

that best fits the eight response amplitudes. **B**, Reconstruction of eye position from the planner gain field fits from the entire population of voxels in V3A. Reconstruction was performed for each trial's work of data using a leave-one-out procedure. Mean and 95% confidence interval for each reconstructed eye position is indicated by bullet (•) symbols and ellipses. Actual eye position is indicated by crosses (+). From Merriam, Gardner, Movshon, and Heeger (2008).

C.2.4. High-Resolution fMRI Reveals Fine Spatial Structure in Cortical Activity (Aim 3)

In Aim 3, we propose new methods for spatially smoothing high-resolution fMRI data. We have developed a zoomed, 3D, multi-shot EPI pulse sequence that is capable of achieving a voxel resolution of $700 \times 700 \times 700 \mu\text{m}$, representing an almost 80-fold increase in spatial resolution over conventional methods (Merriam et al., in preparation). We scanned subjects while they viewed a set of concentric rings that alternated with a set complimentary “anti-rings” (Fig. 7A). The spacing between the rings compensated for the cortical magnification factor, which we measured in another experiment (Larsson and Heeger, 2006), in order to produce bands of activation of a constant thickness (Fig. 7, B & C). The activation bands are 1.5 mm thick in the example shown, representing a cortical spatial frequency that cannot be resolved using conventional imaging methods. Our preliminary results demonstrate the feasibility of collecting data at a resolution that is much higher than is conventionally used.




Figure 7. Cortical responses to checkerboard stimuli measured with high-resolution fMRI. **A**, Flickering checkerboard stimulus consisted of a set of rings and anti-rings, alternating according to an ABAB block design at a rate of $1/24$ s. **B**, Responses to the ring stimuli overlaid on a standard anatomical scan. Inset indicates the field of view, restricted to the calcarine sulcus. **C**, Statistical map overlaid on a computationally flattened representation of the cortical surface. **D**, Amplitude spectrum of the time series indicates the signal to noise ratio of the measurement.

D. Research Design and Methods

In this section, we describe in detail our plan for research, both methodological and empirical. In each subsection, we begin by describing the new statistical methodology and follow by describing the proposed experiments. Our aims and questions, methods and experiments are closely linked, and we highlight those links throughout, showing how the methods are useful for analyzing the experiments and how the experiments motivate the methods.

A theme connecting the statistical methods we propose is identifying regions within which the signal exhibits a shared structure. Whereas the voxel grid is an artificial framework imposed by data collection and of little intrinsic interest, it is such regions that are the true target of inference in fMRI. The problem of identifying such regions is an important one in many applications, in neuroimaging and beyond. In Aim 1, we identify a hierarchy of regions where the signals exhibit functional dependence and show how to statistically detect changes in those regions across time or condition. In Aim 2, we identify regions called response fields that show a similar response to a distributed stimulus. In Aim 3, we identify regions over which the signals can be combined to denoise the data while reducing the structural bias of current smoothing methods.

Neurophysiological studies in monkeys over the last two decades have supported the idea that the brain locates visual stimuli by combining information about the position of stimulus excitation on the retina and the position of the eyes relative to the head. But relatively little is known about how these two signals are combined in the human brain. For Question 1, we will extend on our previous studies of visual remapping by studying how functional networks in the brain change when subjects make eye movements. For Question 2, we will investigate cortical representations of visual stimuli when the eyes are held fixed at different locations in the visual field. And for Question 3, we will exploit a new technique for very high resolution fMRI to test for fine-scale structure of eye position signals in early visual cortex.

D.1. QUESTION 1 / SPECIFIC AIM 1

One of the chief advantages of fMRI over single-unit physiology is the ability to sample activity in multiple regions simultaneously, but the field currently lacks adequate statistical tools for making inferences about coordinated and widespread patterns of brain activity. In this Aim, we propose rigorous and flexible statistical methods that can characterize regions in which the time courses exhibit related patterns activity and can detect changes in these regions over time or across conditions. The result is a new way to assess functional connectivity that is more powerful than existing techniques. Remapping is a good test bed for these developments because computational models of remapping posit interactions among regions in a broad, interconnected network and yet existing studies of visual remapping in both monkeys and humans have focused on studying individual brain regions. The goal of our experiment is to characterize interactions between cortical areas during remapping. Our overarching hypothesis is that a functional network involving parietal cortex and topographically aligned visual areas becomes activated around the time of an eye movement. We will use the new methods and the data from our experiments to test specific predictions about activity in this network.

D.1.1. Characterizing Distributed Functional Networks (Aim 1)

As described in B.1, our proposed methodology for quantifying functional connectivity and discovering functional modules has three stages (Shalizi et al., 2007):

1. Discretize the time course for each voxel, $X_i(t)$, and use CSSR to reconstruct a predictive state-space model for that time series, producing a series of states $S_i(t)$.
2. Estimate the mutual information $I(S_i, S_j)$ between each pair of voxels, normalizing to give information coherence IC_{ij} .
3. Maximize the objective function $H_\gamma(\vec{\sigma}) = \sum_{i \neq j} (IC_{ij} - \gamma K_i K_j / M) \delta(\sigma_i, \sigma_j)$, where $\vec{\sigma}$ assigns voxel i to community σ_i , K_i is the averaged coherence of voxel i with all other voxels, M is a normalization factor, and $\gamma \in [0, \infty)$ controls the hierarchical scale of the division into communities. As $\gamma \rightarrow 0$, we obtain fewer, large, less coherent modules, rooted in the limit $\gamma = 0$ in a single module embracing all voxels. As $\gamma \rightarrow \infty$, these modules sub-divide into smaller, more numerous, and more internally coherent sub-modules. Both the organization at a particular γ level and the hierarchy of levels are of interest.

State-space Reconstruction: Consistency and Convergence. As mentioned in B.1, CSSR is pointwise consistent when the data-generating process has a finite number of predictive states (and some weak regularity conditions

hold) (Shalizi and Klinkner, 2004). Work is already under way, with L. Kontorovich, on extending this to uniform consistency over classes of processes parameterized by a measure of the amount of memory (in bits) required for prediction, using his recent results on measure-concentration inequalities for partially-observable Markov processes (Kontorovich, 2007; Kontorovich and Ramanan, 2007). A second pressing issue is the behavior of CSSR when the true number of predictive states is infinite; here simulations show the algorithm building a series of increasing complex and accurate finite-state approximations. We believe this observation can be formalized using techniques developed in ergodic theory for studying the weak convergence of mixing processes, and their representation as transformations of IID noise sources (Ornstein and Weiss, 1990; Gray, 1988).

State-Space Reconstruction: Multivariate Methods for Multiple Voxels. Having used informational coherence to establish that voxels’ dynamics are coupled, it is natural to want to jointly model the linked processes. Symbolically, the model would partition the whole set of voxels into m non-overlapping sets I_1, I_2, \dots, I_m , and then

$$\begin{aligned}\forall i \in I_a, \quad X_i(t) &= Q_i(S_a(t), \epsilon_a(t)) \\ S_a(t+1) &= T_a(S_a(t), X_{I_a}(t))\end{aligned}$$

with innovations $\epsilon_a(t)$ coming from a noise process unpredictable from the histories of $S_a(t)$ or $X_{I_a}(t)$. That is, each group of voxels has its own Markovian state process, which generates a multivariate time series for the corresponding voxels. Using models of this form, rather than that of (1), has important benefits: more accurate estimates of dependence and information flow, improved prediction and filtering, and a better representation of the underlying information-processing taking place in the brain (Shalizi and Crutchfield, 2001; Haslinger et al., 2008). If the interaction topology of a spatio-temporal system is known, then there are techniques for “composing” the predictive state models of multiple voxels into one joint predictive-state model (Shalizi, 2003; Shalizi et al., 2006). However, we are interested in, precisely, cases where the interaction topology is unknown. Alternatively, we could fit a single global predictive-state model to all the voxels at once. Non-negative factorizations, or near-factorizations, of its state space would then indicate the functional modules of the system. Unfortunately, reconstructing a global state space from a high-dimensional multivariate time series faces a serious curse of dimensionality.

To get the advantages of a multivariate model in a tractable form, we propose a **one-step refinement** procedure, where in the first stage a separate state space is reconstructed for each voxel independently, informational coherence is assessed pairwise, and clusters of strongly-dependent voxels are identified. We will then reconstruct a joint state-space for each cluster. This should improve prediction and filtering for the cluster, as well as giving a refined estimate of the total functional connectivity within the cluster, namely the difference between the sum of the complexities for the marginal state spaces and the complexity of the joint state space. We will investigate the sensitivity of this procedure to different clustering methods (giving special emphasis to clusters found via community discovery), and establish its reliability via simulation studies.

State-space Reconstruction: Discretization. A final issue for state-space reconstruction is that the current approach requires discrete-valued time series. Currently, we discretize by taking a local polynomial fit to the time course for each voxel, estimating the derivative, and binning the derivatives by quintiles. On the practical side, we will examine the sensitivity of our results on functional connectivity to different discretizations (including the number of bins, and whether to discretize by level or by derivative) in both real data and realistic simulations. Theoretically, we will investigate whether adaptive-discretization schemes developed for deterministic dynamical systems (Kennel and Buhl, 2003; Hirata et al., 2004; Hirata et al., 2005) can be consistently carried over to stochastic processes.

Mutual Information Estimation. Accurate and precise mutual information estimates are crucial to our approach. The naive approach plugs the the empirical joint histogram into Shannon’s formula. This is consistent and quite accurate for large samples and small numbers of values for each variable. However, it has long been known to be biased (Miller, 1955), with the magnitude of the bias growing with the size of the sample space. While analytical expressions exist for the bias of this estimator under IID sampling (Victor, 2000), they are difficult to work with and even more difficult to extend to serially-dependent data.

Currently we use the estimator proposed by Paninski (Paninski, 2003), which optimizes a variational bound on the estimation error. However, we need to investigate the finite-sample biases properties of the combination of this estimator with CSSR reconstruction, in particular the use of a model-based bootstrap to obtain the error distribution and rates of convergence. We will also investigate the theoretical properties of alternative estimators.

The joint-state reconstruction approach, mentioned above, will give us model-based estimates of the mutual information between voxels. Experience with estimates of sequence entropy and entropy rate shows that such estimates often converge much faster than estimates based directly on the sequence statistics. We expect that something similar will be true here. However, if we try to jointly reconstruct a single state-space for too many voxels, we run into a curse of dimensionality and excessive bias. We therefore need to determine the risk curve for these joint-model estimates of mutual information.

Module Discovery: Consistency. The final part of our method is to build modules from the global pattern of functional connectivity, which requires us to address the statistical properties of the Reichardt-Bornholdt module-discovery procedure. The only prior work on this topic known to us (or to Drs. Reichardt and Bornholdt; personal communication), is an investigation of the sampling distribution of the objective function under the null hypothesis of the absence of modularity (Reichardt and Bornholdt, 2006b). The issues involve both the use of estimated rather than true values of the mutual information, and the choice of hierarchical scale γ .

In our preliminary work, we have established pointwise consistency, in the following senses. Holding fixed the matrix of true mutual information values a_{ij} , and assuming that our estimates of those values are themselves consistent, $IC_{ij} \xrightarrow{P} a_{ij}$, at each scale γ of the hierarchy, the optimal partition $\hat{\sigma}_\gamma^*$ estimated from data converges on the population-optimal partition σ_γ^* at that scale:

$$d(\hat{\sigma}_\gamma^*, \sigma_\gamma^*) \xrightarrow{P} 0 \quad (3)$$

(The result is independent of the metric $d(\cdot, \cdot)$ on partitions.) This essentially follows from the convergence of the objective function $H_\gamma(\vec{\sigma})$ for each partition (since it depends continuously on IC_{ij}), and the finite, though large, number of possible partitions $\vec{\sigma}$. We have also extended this result to one which holds uniformly over ranges of scales, $\gamma \in [0, \gamma_1]$,

$$\sup_{\gamma \in [0, \gamma_1]} d(\hat{\sigma}_\gamma^*, \sigma_\gamma^*) \xrightarrow{P} 0 \quad (4)$$

The extension works because the objective function is affine in γ , which reduces the problem to the union of the fixed-scale problem at finitely many scales.

The necessary next step is to extend these fixed-network results to ones which hold uniformly over networks of suitably-bounded complexity. We believe this can be accomplished by measuring “complexity” in terms of the growth of the number of cells in the optimal partition with γ . Controlling this complexity should allow us to build combinatorial proofs for uniform consistency along VC or Rademacher lines. Once rates of convergence of mutual information estimates become available, we will be able to propagate them through to rates for module discovery.

A recent extension of the Reichardt-Bornholdt procedure (Reichardt and White, 2007) allows it to recover “generalized block models” or “role models”, in which units are assigned to distinct functional roles based on their pattern of connectivity with other units in their own roles. This includes the module-finding procedure as a special case, where each module is a distinct role, and preferentially connects only to units in the same role. We will adapt this to identifying the roles of voxels from time series. Consistency of role-discovery can be tackled in the same way as consistency of module-discovery.

Module Discovery: Significance of Differences. Finally, since we are most interested in comparing patterns of functional connectivity across experimental conditions, we need methods to assess the significance of differences in such patterns. Asymptotically, consistency will take care of this, but the limited duration of our experiments means we need something more. A simple way of obtaining the sampling distribution of distances between estimated partitions, under the hypothesis that the generating distribution remains fixed, is to randomly exchange entries between the two mutual information matrices from two experimental conditions and re-estimate the partitions. Sampling distributions under weaker null hypotheses which require fewer invariants can be obtained by model-based bootstrapping, simulating from models fitted to pooled data from both experimental conditions. This will be considerably slower than the simple exchange test, but will extend the scope to a wider variety of questions.

CAVEATS AND PITFALLS. The biggest drawback of the procedure is that it contains multiple steps at which potentially high-dimensional entities (predictive state-space models, mutual informations, functional modules) must be estimated, with weak prior constraints. The amount of data required for such estimates to be reliable can be quite considerable. When sample sizes are too small to obtain precise estimates, we would expect to find that differences in functional connectivity were almost never statistically significant, not because nothing has really

changed but because our models are under-constrained. Assessing rates of convergence is thus of more than merely theoretical interest. Our preliminary work with existing data (cf. C.2.2) does suggest that we can get meaningful and reliable estimates with reasonable sample sizes.

DATA ANALYSIS. We will carry out the following steps:

1. Implement a fast, fully-automated system for processing data from voxel time courses all the way to hierarchies of functional modules.
2. Compare the results of using CSSR for state-space reconstruction to using linear-Gaussian HMMs for each voxel, estimating states with a Kalman filter, and then calculating mutual information from the estimated states. We will do this both for simulated systems, with known functional connectivity, and real data.
3. Assess the impact of different discretization schemes on real and simulated data.
4. Develop bootstrapping methods, as described above, for assessing the significance of differences in functional-connectivity patterns, and use simulation studies to determine how large a change must be before it can be reliably detected in data sets of an experimentally-reasonable size.
5. Implement the one-step refinement procedure and use simulation to assess the risk curve for using it in a model-based estimate of mutual information between voxels.

D.1.2. What is the functional network for visual remapping? (Question 1)

Remapping reflects the transfer of activation from one group of neurons (those that encoded a stimulus location before the eye movement) to a second group (those that will represent the stimulated location after the eye movement). The full network should, in theory, be able to transmit spatial information from neurons representing any part of the visual field, regardless of the trajectory of the saccade or the location of the stimulus. By systematically varying the geometry of the remapping task, we will test the hypothesis that remapping is universal. We predict that remapping will occur in all conditions, but that the community of voxels detected by the procedures in Aim 1 will vary according to the retinotopic location of the stimulus.

EXPERIMENT 1. We will scan subjects during performance of the single-step remapping task. The sequence of events will be identical to our previous studies of remapping (Fig. 1A): The trial begins when subjects fixate FP1. A stimulus appears and flickers on the screen for 1,000 ms, activating the contralateral hemisphere. As the stimulus disappears, a tone cues the subject to make an eye movement to FP2. As we have shown previously, the representation of the stimulus will be remapped in conjunction with the eye movement. In separate blocks of trials, we will vary the direction of the eye movement and the location of the visual stimulus (Fig. 8). For each configuration of stimulus and saccade trajectory, we will also test subjects on the stimulus-only and saccade-only control tasks needed to subtract activity related to purely visual or motor aspects of the main remapping task. The two control tasks will be run in a separate block of trials; previous work as shown that the single-step task can induce a long-term memory response that can influence activity in the control task (Umeno and Goldberg, 2001).




Figure 8. Single-step task for testing functional networks for remapping. **A**, Saccade vectors all terminate at the same position (center). In each case, the stimuli (filled circle) are flashed at the *identical retinotopic locations* in upper left visual field. Dotted circles indicates stimulus trace location *after* the eye movement. **B**, Single-step task with four stimulus locations. In each case, retinotopic location of the stimulus is varied (upper left, upper right, lower left, lower right). Saccade vector is held constant. These two manipulations are expected to have different effects on the functional network for remapping.

Potential Outcomes. This experiment is designed to address three fundamental questions about the neural basis of remapping. First, do functional networks in the brain change during remapping? The expected result is that remapping will enhance the mutual information between pairs of voxels in a number of cortical regions. The particular functional communities detected by our analysis will yield insights into the underlying neural mechanism that gives rise to remapping. One possible outcome is that changes in mutual information will be highly focused,

involving just a few sets of areas. As our preliminary data indicate, changes in mutual information may be limited to pairs of voxels in early visual areas and the IPS. This result would suggest that the IPS is directly involved in the computations underlying remapping. A second possibility is that changes in mutual information will be diffuse throughout the network. This result would suggest that remapping is carried out in parallel at multiple stages of the visual area hierarchy.

Second, do functional communities for remapping depend on the geometry of the task? Varying the stimulus location should shift remapping activity to the sub-population of neurons that encode that particular location. We expect that the community discovery algorithms will detect separate sets of voxels for each stimulus location. Moreover, we expect that these communities will be topographically aligned. For example, voxels in V3A that represent the upper left visual field may form a coherent network with a subset of voxels in IPS that represent the same location in space.

Third, is remapping is universal? The expected outcome is that parietal cortex will exhibit significant remapping in conjunction with saccades of all directions, regardless of stimulus location. With respect to the strength of remapping in different directions, there are two possible outcomes. These expected outcome is that the stimulus trace response will be equally robust for all saccade directions. The alternative is there will be systematic differences in the amplitude of responses. For example, we may observe differences across vs. within-hemifield remapping or for remapping in conjunction with horizontal vs. vertical saccades or for remapping in conjunction with contraversive vs. ipsiversive saccades. If we do find systematic differences, we will follow up with a new set of experiments to specify the conditions under which remapping occurs.

Caveats and Pitfalls. This experiment has a substantially larger number of conditions relative to our previous studies of remapping, which will place limitations on the number of trials per condition. We will compensate for this limitation by scanning subject in up to five separate sessions. We have used this approach—many repeated measurements in each of a number of experienced, well-practiced, motivated subjects with success in our previous work (Ress and Heeger, 2003; Lee et al., 2005, 2007).

DATA ANALYSIS. The goal of this analysis is to infer changes in the intrinsic interactions between cortical areas as a result of remapping. In the single-step remapping task, the driving effects of the stimulus and the saccade induce correlations that do not reflect intrinsic inter-area interactions within the brain (Gerstein and Perkel, 1969). It is thus critical to remove both the stimulus and saccade signals from the time series prior to computing inter-area interactions. As a preprocessing step, we will calculate a residual time series by subtracting a predicted mean time series associated with the two control conditions. Specifically, after computing the parameter estimates of the fMRI response in the stimulus-only and saccade-only control conditions for each saccade direction and stimulus location, a predicted mean time series \hat{y} will be computed by multiplying the design matrix by the parameter estimates, that is, $\hat{y} = A\hat{x}$. The residual time series will then be computed by subtraction, yielding $r = y - \hat{y}$, where r is the residual vector.

We will compare the functional networks during remapping and nonremapping trials by applying the community discovery procedure separately to the corresponding data. We will generate neighborhoods at all scales γ but focus our initial comparison on a particular γ that was selected from pilot data to maximize predictive power. The comparison of the networks will provide a direct test of our predictions. We will use the exchange and bootstrap methods described above to test the difference between these networks, initially under the the null hypothesis that the generating distributions are the same in the two conditions. Based on the theoretical results so far, we can also look at multiple scales simultaneously without any loss of power, and we can compare the time evolution of the neighborhoods (at coarse resolution).

In addition to the community discovery analysis, we will perform traditional analysis designed to test the effects of stimulus location and saccade vector on remapping. The saccade directions manipulation will allow us to compare the amplitude of the response for: (1) remapping across hemifields (contraversive horizontal saccade); (2) remapping within a hemifield (ipsiversive horizontal saccade); (3) remapping in conjunction with vertical saccades. The stimulus location manipulation will allow us to compare the amplitude of the response for four different stimulus locations (upper, lower, right, left). For each manipulation, the complete data set consists of a 3 tasks \times 4 directions/locations. A simple analysis of variance will be used to make these comparisons.

D.2. SPECIFIC AIM 2 / QUESTION 2

Our studies of remapping highlight one neural mechanism by which the brain may create a stable perception of the world during eye movements. Computational models have described a related mechanism based on multiplicative interactions between visual and oculomotor signals. According to this model, the location of eye fixation modulates the gain of visual activity. While such effects have been widely reported in the monkey neurophysiology literature (Boussaoud and Bremmer, 1999; Andersen et al., 1990a; Snyder, 2000), little is known about eye position modulation in humans. In this section, we introduce the notion of an fMRI response field, which is a measure of a voxels selectivity for spatial locations. We describe a new method for measuring response fields in humans using spatially complex, time-varying white noise stimuli and novel data analysis techniques. We then propose an experiment that uses the new method to test for eye position modulation in humans.

D.2.1. Simultaneous Estimation of Visual Response Fields (Aim 2)

Estimating response fields is a high-dimensional spatio-temporal problem with strong a priori constraints that are not easy to account for in statistical terms. To set up the problem, we need to distinguish coordinates in the visual field and coordinates in the brain. The white-noise paradigm decomposes the visual field into a collection of p *tiles*, and data acquisition decomposes the brain into a three-dimensional grid of *voxels*. Typically, we are interested in a particular set of voxels \mathcal{V} of size m that lies in the brain or in a particular region of interest. Of course, the voxel coordinates are artificial, and ultimately what we care about are locations on the cortical surface.

In the white-noise paradigm, the stimulus produced at the j th tile can taken as a sequence s_t^j of ± 1 s for all $t = 0, \Delta, \dots, (n-1)\Delta$. The TR parameter Δ is defined by the experimental design, but for convenience here, we will take it as 1 without loss of generality. As described above, we use M -sequences, which have the properties that they mimic white noise, have mean zero, and for each j , s^j is a circular shift of s^0 by δj , for some integer δ . The *response field* is described by the spatial distribution of responses to the stimulus in each tile. Specifically, we assume parameters β_j^v describing the response in voxel v to the stimulus from tile j , so that the total response at voxel v is just the linear combination $\sum_{j=0}^{p-1} \beta_j^v s_t^j$. We assume further that the hemodynamic response has a convolution structure with impulse response function α_t^v . Thus, if Y_t^v denotes the centered and detrended data at time t and voxel v , we assume that

$$Y_t^v = \left(\alpha^v * \left(\sum_j \beta_j^v s^j \right) \right)_t + \epsilon_t^v, \quad (5)$$

for a Gaussian, mean zero noise process ϵ and where $*$ denotes (circular) convolution. We can write this equivalently as

$$Y_t^v = \sum_j \beta_j^v \left(\alpha^v * s^j \right)_t + \epsilon_t^v. \quad (6)$$

In general, both matrices β and α are unknown. To maintain identifiability in this model we need to impose constraints that prevent ambiguous scalings of the parameters. To this end, we enforce the constraint that for every v , $\sum_t (\alpha_t^v)^2 = 1$. Although our primary interest is in making inferences about β , the impulse response functions are also of scientific interest.

By imposing other constraints, derived from prior information, we can improve estimation of β and a . For example, typically α is effectively zero after a short time span, so we assume that $\alpha_t^v = 0$ for $t > t_0$ where t_0 is a small integer on the order of 16–20. Similarly, α_t^v will tend to vary smoothly across time. Similarly, β will tend to be sparse, meaning that most β_j^v s equal 0, because of the relatively small size of receptive fields in the cortical areas that we propose to study. We will also usually require that the β 's be non-negative because they represent a positive response to the stimulus. Indeed, we can typically assume that there is only one (or at most a small number) focal response field for each tile j . We also expect α^v and β^v to vary smoothly across voxels, at least over neighborhoods where the voxels have similar anatomical and response profiles. (See also D.3.1.) The goal of our analysis is to construct good estimators of β and α , along with useful assessments of uncertainty if possible.

We will investigate three approaches: (1) Alternating Sparse Regression; (2) Sparse Bayesian Sampling, and (2) Iterative Clustering. using data in voxel coordinates and in cortical surface coordinates on a computationally inflated brain. The latter reduces the problem from three-dimensions to two, although the coordinate positions no longer remain on a grid.

Alternating Sparse Regression. If we are to estimate α and β simultaneously, the problem is nonlinear. But equations (5) and (6) reveal that it is composed of two linear problems: given α we have a linear model for β and vice versa. This suggests iteratively alternating between solving the two problems.

We begin each iteration by taking α as known, equal to the previously fitted value and initialized to the fitted value from a simple parametric family. Then, stacking entries over time and voxel with time varying fastest, equation (6) becomes a linear model,

$$Y = K\beta + \epsilon, \quad (7)$$

where K is a $nm \times pm$ matrix with $K_{(v,t),j} = (\alpha^v * s^j)_t$ and (v, t) represents the index in the stacked vector. This K can be quite large, but we can exploit the convolutional structure by Fourier-transforming both sides of equation (6). This yields a linear model in β analogous to equation (7) with a design matrix $\tilde{K} = D \cdot S$ where D is diagonal and S is block diagonal with identical blocks, and thus \tilde{K} is block diagonal and can be computed quickly.

Although β is high-dimensional, it is very sparse, so we will use a sparse regression technique to estimate it at each stage. The simplest is the lasso (Tibshirani, 1996), which minimizes $\|Y - K\beta\|_2^2 + \lambda\|\beta\|_1$, or equivalently minimizes the residual sum of squares subject to a constraint on the absolute sum of the β s. The lasso produces sparse solutions that can be computed efficiently even in very large problems and that have good performance in both prediction and inference problems. With the LARS algorithm (Efron et al., 2004), solutions can be computed for every λ simultaneously, and a value of λ can be selected with cross-validation or by minimizing an empirical risk (e.g., SURE), depending on the context. Moreover, in very large problems, the lasso can be computed relatively quickly by iterated marginal regressions (Tibshirani, personal communication; Genovese, Jin, Wasserman, in prep).

In our situation, the standard lasso needs several modifications. First, we constrain β to have only non-negative entries. For this, we will make straightforward modifications to the the LARS algorithm to only bring in variables in a positive direction and will show that this solves the positivity constrained optimization problem. For large data sets, we will use the special structure of the problem in the Fourier domain to solve the β -fitting stage iteratively over sub-blocks of \tilde{K} . We will use pilot data to estimate the temporal and spatial correlation function in the data, thus approximating the covariance structure in the regression problem.

The lasso will provide sparse solutions but by itself will not take into account the geometric regularity of the response field. Thus, the second modification of the lasso is to add spatial smoothness penalties for the β 's to the objective function. The simplest such penalty function is the sum of squares of first-differences $\beta_t^{v+1} - \beta_t^v$. We will also consider sums of squares of the discrete Laplacian of β_t , which is analogous to the typical Sobolev penalty in smoothing splines (Wahba, 1990). And we will use the adaptive smoothing neighborhoods from Specific Aim 3 (see D.3.1), penalizing first differences or Laplacians over those neighborhoods only. These 2-norm penalties are convex and differentiable in β . They change the optimum solution and add an extra tuning parameter that must be selected, but it should be straightforward to adapt existing algorithms to efficiently find solutions.

Finally, we want to use the fact that there are only a few contiguous regions, often just one, which constitute the response field. To this end, we will employ the modified sparse regression techniques in a greedy iterative algorithm. Initially, the algorithm will find the location of the strongest signal, as measured by the sum of the β 's over a small contiguous neighborhood (e.g., 1,4, or 9 voxels). Then, it will successively add the strongest contiguous locations to the region – with the marginal value of the parameter – until the objective function does not decrease by at least a specified threshold. While this “seed growing” method is not guaranteed to find the minimum objective value, it has performed very well in preliminary work.

In the second stage of each iteration, we fix β at its previously fitted value. Re-writing equation (5) reveals a linear model in α :

$$Y_t^v = \sum_u \alpha_u^v c_{t-u}^v + \epsilon_t^v, \quad (8)$$

where $c_r^v = \sum_j \beta_j^v s_r^j$. Stacking the vectors gives $Y = C\alpha + \epsilon$, where again C is block diagonal. We use the constraint that α is zero beyond a fixed number of components to keep the size of C manageable (i.e., most columns can be eliminated) and use its block-diagonal structure to solve the regression problem efficiently (e.g., iterating over the blockwise problems). For each v , we require that $\sum_t (\alpha_t^v)^2 = 1$, which adds $\#(\mathcal{V})$ constraints to complicate the optimization. We also want to impose smoothness constraints on α over time and space, using sums of squares of first or second differences as above, across t and v separately. Again, the adaptive smoothing developed under Aim 3 will be helpful here. For larger problems and to handle the number of constraints, we will use iterative backfitting (Hastie et al., 2001), cycling over the separate voxelwise problems.

Because the objective function (for both stages of the iteration) is convex, the alternating solutions should converge to the constrained global minimum for both β and α . In principle we can handle arbitrary spatio-temporal correlations in ϵ , but in practice, when the size of the problem gets large, we may need to approximate the correlation structure to avoid representing the entire correlation matrix. In such cases, we will likely restrict our attention to white noise and low order autoregressive models, in space and time, which appear to fit most fMRI data quite well. Notice also that there is nothing special about voxel-coordinates in this formulation of the problem and that the same structure works in cortical surface coordinates as well. The spatial smoothness penalties will need to be altered slightly on the cortical surface; instead of differences over a grid, we will use nearby points in the cortical projection. Our preliminary work in this direction has been successful so far, comparing favorably with commonly used methods (Hansen et al., 2004) in realistic simulations.

One of the goals of our experiments is to detect *changes* in response fields as a function of covariates such as eye position. This is complicated by the difficulty of assessing the uncertainty of lasso-type estimates. One reason for this difficulty is that the null distribution of the lasso coefficient estimators is skewed with an atom at zero. We will begin adapting the standard error estimates of Osborne et al. (2000) to compute standard errors of summary statistics of the response fields such as center of mass. We will investigate other approximations to the standard errors and evaluate the accuracy of these tests in simulations.

Sparse Bayesian Sampling. As an alternative to the heavily iterative method just described, we will implement a Bayesian version of the model that uses priors to enforce sparsity. The primary advantage of a Bayesian model here is that it is easier to handle qualitative constraints, in particular maintaining one contiguous region in the response field. The latter constraint is likely to have a quite powerful effect on performance.

The likelihood for the hierarchy will be based on the Normal distribution in equations (5) and (6). But the full likelihood need not be computed, because at each stage of the chain we can make the reductions given in equations (7) and (8), possibly using the Fourier representation for the former. Our initial prior will be relatively simple, designed to produce sparsity in β and smoothness in α . For β , we will start with a logistic prior conditioned on non-negativity of the components, which is analogous to the non-negative lasso described above and for which sampling is fairly straightforward. For α , we will use a Markov Random Field (MRF) prior (cf. section D.3.1 and (Geman and Geman, 1984)), making the conditional distribution of α_t^v a (truncated) Normal mean parameter equal to a weighted average of α_s^w for neighboring values of s and w subject to the normalization constraint. This will be sampled with a Metropolis-Hastings chain that updates one coefficient at a time, and we will explore ways to make more efficient moves. We will also consider a prior that allows only one contiguous group of β 's to be non-zero. This requires some care in choosing sampling moves that can cross low-posterior barriers. For example, moves that slightly shift or distort the region as a whole will lead to a chain with better mixing than if it is restricted to single parameter moves. We will also consider adapting the techniques under Specific Aim 3, a Bayesian hierarchical model can be extended to impose better spatial smoothness constraints on α and β .

We will compare performance of this approach, both in computing requirements and accuracy, to the iterative scheme. One additional benefit of the Bayesian method is that it makes it easy to assess uncertainties. We will use the posterior to quantify inferences about how the response fields change with covariates such as eye position and attention. Evidence for these changes include changes in the magnitudes of the β 's and changes in the shapes and sizes of the regions. We will begin by considering the posterior distributions of summary statistics such as the center of mass of the response field (weighting by the β magnitudes) and principle axes of the regions. This will be useful for detecting shifts and changes in orientation. More generally, we will extend the techniques for assessing posterior changes in contiguous regions developed in (Genovese, 2000).

Iterative Clustering. Finally, we can adapt the information neighborhood method from Specific Aim 3 (see D.3.1 below) to find a set of locations within which the voxels have high mutual information with each other and low mutual information with voxels outside the set. We will assess whether this set by itself is a good estimate of the response field and compare this method with two sparse regression techniques.

CAVEATS AND PITFALLS. There are three significant challenges here. The first is the proper spatial smoothing of α and β . This requires both some care and experimentation to avoid undue bias, and we will benefit as well from our work on Specific Aim 3. The second major challenge here is computational. Both the iterative and Bayesian techniques are computation-intensive, but we can take advantage of special structure in the problems, as described above, and modifications of existing algorithms to make the inferences efficient. Sampling from the

posterior in the Bayesian version of the models requires careful construction of moves. We will diagnose mixing and validate performance of the MCMC simulations. The third major challenge is inferential. Statistically comparing response fields when a covariate changes can lead to quite complex procedures. Initially, we will compare summary statistics related to location and orientation, but more general methods will be useful. The Bayesian formulation of the problem has an advantage in this regard.

We will also test formulations of the problem that estimate the gain fields directly by augmenting the model. Specifically, we will represent multiple response fields by a common β parameter and additional parameters for each eye position that represent deviations from β . Putting prior constraints on the deviations (in either the sparse regression or Bayesian setting) will allow us to estimate differences across eye positions. This will be computationally intensive but may be more statistically efficient because the response parameter (α) will be common to all eye positions.

DATA ANALYSIS. We will carry out the following steps:

1. Implement the alternating sparse regression method including positivity constraints and using the special structure in the Fourier representations. Implement the basic Bayesian model using a logistic prior on β to enforce sparsity.
2. We will compare these methods to each other and to the commonly used reverse correlation method (Marmarelis and Marmarelis, 1978; Sakai, 1992; Hansen et al. 2004).
3. Incorporate spatial smoothing constraints into the alternating sparse regression techniques and evaluate performance based on simulations.
4. Incorporate the single-region constraint into the Bayesian model and evaluate performance based on simulations.
5. Compare the best formulation of the sparse regression and Bayesian model with the clustering technique using real data and (additional) simulations.
6. Estimate the response fields from the experimental data using the best method from our assessment.
7. Develop tests that can detect changes in response fields as a function of covariates such as eye position and attention. Using eye position initially, we will begin looking for changes in location and orientation of response fields and extend to other functionals that describe higher-order features of the response fields' shapes. We will also attempt to estimate the gain fields directly with an augmented model, as described above.

D.2.2. Does visual cortex contain a representation of eye position? (Question 2)

The aim of this experiment is to test the interaction eye position signals and visual representations. We will use the methods developed in Aim 2 to measure visual response fields as subjects fixate different locations on the screen. The hypothesis is that static eye position will have an impact on the amplitude of the visually-evoked activation.

EXPERIMENT 2. The stimulus will consist of an 8×8 matrix of squares, or *tiles*, 1.5° on a side (Fig. 9). The contrast of each tile will be modulated in time by a pseudorandom M -sequence of -1 's and $+1$'s. When the sequence for a given pixel has a value of $+1$, the tile will flicker between black and white at 8 Hz, and when it is -1 , the tile will remain uniform gray. The sequence will update at 1 Hz. Each pixel will be governed by the same M -sequence shifted in time in steps of 64 s. On each scanning run, we will vary the screen location at which subjects fixates. Critically, the white noise tiles will always be located at the same retinotopic positions relative to fixation.




Figure 9. Stimulus for eye position experiment. In each run, the subject fixates a different location on the screen while the white noise stimulus flickers in the background. The question is whether eye position modulates the amplitude of the visually-evoked responses.

Potential Outcomes. The expected outcome is that we will observe systematic changes in response field characteristics with eye position. Based on single-unity physiology (Andersen et al., 1990b), and our preliminary results (Fig. 6), we predict that the influence of eye position on visual response amplitude will be linear with regard to horizontal and vertical eye position. The alternative outcome is that eye position does not modulate visual responses. This result would suggest that the eye-position-related baseline shifts reported in previous imaging studies (DeSouza et al., 2002) are due to signals arising in higher cortical areas (e.g., LIP) and are not related to visual processing *per se*.

Caveats and Pitfalls. A recent single unit recording study found that eye position modulation of visual responses was itself modulated by other cognitive factors (Sharma et al., 2003). In that study, the degree to which eye position influenced visual activity was dependent on the monkey’s expectation of the eye position on an upcoming trial. The experiment in Aim 3 could easily be extended to test the effects of prior expectation by manipulating the predictability that the fixation spot will appear at any given location.

DATA ANALYSIS. Data analysis for this experiment will involve two stages. The first stage will involve determining which tiles evoke a response in a given voxel (i.e., to determine each voxel’s response field). We will apply the methods described above to estimate the response fields. The second stage will involve comparing response fields across the different eye positions. By sorting the estimates by measured eye position and taking differences, we obtain estimates of the corresponding gain fields. We will use either the posterior distributions or the approximate standard errors to test for differences with eye position. If it proves computationally feasible, we will also apply the augmented model described above to estimate the gain fields direction. Posterior distributions (or approximate standard errors) would then give a useful picture of gain field uncertainty.

D.3. SPECIFIC AIM 3 / QUESTION 3

The question here is whether eye position signals in early visual cortex have a fine-scale spatial structure. Our preliminary results at low-resolution indicate that the influence of eye position on visual activity varies from voxel to voxel, with no global pattern across voxels. However, optical imaging studies in monkeys, which have a much higher resolution than fMRI, do indicate a fine topography for eye position gain fields (Siegel et al., 2003). We will take advantage of new scanning methods developed at NYU to collect data at a voxel resolution of $700\ \mu\text{m}$ isotropic, an almost 80-fold decrease in the total volume of each voxel. Signal-to-noise ratio is inversely proportional to voxel size; accordingly, our preliminary data at high resolution document a dramatic decrease in the signal-to-noise ratio. The new statistical methods in Aim 3 will help recover some of the signal lost by moving to higher spatial resolution.

D.3.1. Adaptive Smoothing of High-resolution Data (Aim 3)

In fMRI, it is common prior to statistical analysis to smooth the data across voxels, typically using a fixed bandwidth kernel with bandwidth set a priori. The idea is to increase signal-to-noise, neighboring voxels having similar signals but different noises. As usual, smoothing faces a bias-variance trade-off: larger bandwidths include more disparate voxels but borrow strength across a wider base while smaller bandwidths average over a more homogeneous but smaller set. But the trade-off is more complicated than usual here, because neighboring voxels can exhibit completely dissimilar signals when separated by anatomical or functional boundaries. The resulting bias degrades the estimated signal in voxels near the boundary and leads to imprecise estimates of region shape because fine-scale and edge features are especially degraded. For large, focal active regions with a big area-to-perimeter ratio, the added bias from these boundary effects will not unduly change results, but for smaller regions and fine-scale features (e.g., along narrow sulci) where most voxels are near boundaries, the signal can be significantly degraded and the estimated active region oversmoothed. Such degradation affects smoothing in either the original voxel coordinates or when the data are projected onto the surface of an inflated brain. Worse for us, the problem is especially severe and important for high-resolution data where fine-scale structure will be the norm and noise levels will be high.

Here we propose two methods for *adaptive* smoothing: using the data to determine the neighborhoods over which to smooth. The first method maintains smoothing as a pre-processing step. It solves a sequence of two-component clustering problems to identify neighborhoods of each voxel that have high mutual information with that voxel and constructs a weighted average of the corresponding time series. The second method merges smoothing with statistical inference. It posits a Bayesian hierarchical model for the data that includes smoothing neighborhoods as parameters. We discuss these two methods in detail in what follows.

Denote the detrended data by a collection of random variables $\mathbf{Y} = (Y_{vt} : v \in \mathcal{V}, t \in \mathcal{T})$, where v varies over voxels in a finite set \mathcal{V} (corresponding to voxels in the brain) and t varies over times in the set \mathcal{T} specified by the experimental design. We write $Y_{vt} = \mu_v + F_v(t) + \epsilon_v(t)$, where μ_v is the baseline signal (in the absence of the BOLD effect), F_v is a possibly random BOLD response to the experimental manipulation, and ϵ_v is a stationary, mean zero noise process. While in general, F_v and ϵ_v may be dependent because of movement- or physiological-related

artifacts, we will initially assume that such artifacts have been measured and corrected and take F_v and ϵ_v as independent for each v . We will follow typical practice and use a parametric model for the distributions of the response and noise processes. (One usually ignores the randomness in the response and models F_v as a parametric family of functions $F_v(t; \theta)$. Part of the observed variation in θ would thus reflect the ignored randomness.)

Information Neighborhoods. The idea behind the first method is to identify a neighborhood of each voxel v within which all the time series contain as much information as possible about the response F_v and then to smooth by averaging (or otherwise combining) the times series in this neighborhood. We frame this as a two-component clustering problem for each voxel with the clusters being a set $\mathcal{A}_v \subset \mathcal{V}$ containing v and its complement $\mathcal{V} - \mathcal{A}_v$. We will construct \mathcal{A}_v by solving an optimization problem of the form

$$\mathcal{A}_v = \arg \min_{\mathcal{A}: v \in \mathcal{A}} U_v(\mathcal{A}), \quad (9)$$

for a *submodular* set function U_v defined on subsets of the voxel set \mathcal{V} .

A set function U on the subsets of a finite set \mathcal{V} is submodular if it gives diminishing returns to adding an element to the set. Specifically, U is submodular if $\mathcal{A} \subset \mathcal{B}$ and $w \in \mathcal{V}$ implies that

$$U(\mathcal{A} \cup \{w\}) - U(\mathcal{A}) \geq U(\mathcal{B} \cup \{w\}) - U(\mathcal{B}). \quad (10)$$

(We take $U(\emptyset) = 0$ without loss of generality.) A variety of useful objective functions are submodular, including mutual information between variables and, under weak conditions, information gain and variance reduction for predicting one variable in terms of several other variables. Positive linear combinations of submodular functions are also submodular. The key feature of submodular functions is that, analogously to convex functions, they have unique minima and can be minimized efficiently (Edmonds, 1971; Lovasz, 1983). Polynomial time algorithms exist (Iwata et al., 2001) though they are not necessarily practical. Algorithms used in practice have demonstrated (Krause and Guestrin, 2007; Queyranne, 1995) good performance in a wide range of problems, suggesting that the collection of \mathcal{A}_v s should be obtainable in minutes for a single data set.

Initially, we will define U_v in equation (9) in terms of the estimated pairwise mutual information between voxels $I_{uw} = \hat{I}(F_u, F_w)$ for $u, w \in \mathcal{V}$, as follows:

$$U_v(\mathcal{A}) = \left[\sum_{w' \notin \mathcal{A}} I_{vw'} - \sum_{w \in \mathcal{A}} I_{vw} \right] + \lambda \left[\sum_{\substack{u \in \mathcal{A} \\ w' \notin \mathcal{A}}} I_{uw'} - \sum_{u, w \in \mathcal{A}} I_{uw} \right]. \quad (11)$$

where $\lambda \geq 0$ is a pre-specified constant. This is a submodular function that is small when the voxels in \mathcal{A} have responses with high mutual information with each other and with v and small mutual information with time series outside \mathcal{A} . The value of λ reflects a tradeoff between homogeneity of the time series and similarity with voxel v . We expect to have $\lambda < 1$ and that for reasonably small λ , the results will be insensitive to the specific value. We will also investigate the alternative of using only the first term ($\lambda = 0$) and will consider replacing I_{uw} by a measure of shape similarity, such as the rank correlation of (Heckman and Zamar, 2000).

Because we do not know the mutual informations $I(F_u, F_w)$, we need to estimate them. We will do this using the technique developed for Aim 1 (see D.1.1) and demonstrated in section C.1. Specifically, we will use local polynomial regression to estimate the derivative of F_v and discretize that derivative into a small number of bins. By fitting the resulting discrete time series using CSSR, we obtain estimated Markov structures from which we estimate mutual information. The work under Aim 1 will refine this technique further and develop theoretical and empirical assessments of performance.

Once the smoothing neighborhoods are computed, we will combine the corresponding time series to get a smoothed version of Y_v . Initially, we will consider distance-weighted ($d(v, u) / \sum_w d(v, w)$) and information-weighted ($I_{vu} / \sum_w I_{vw}$) averages, where distances will be measured in inflated-cortex coordinates. But we will also consider other ways to choose the coefficients such as maximizing the $I(F_v, \sum \alpha_u F_u)$ over α_u [ATTN] and allowing shifting and scaling to better match shapes.

When keeping smoothing as a pre-processing step, a natural alternative to the information neighborhoods approach is to cluster the time series so that those with similar temporal “shape” will be clustered together. While we prefer the information neighborhoods because they are local and do not require solving the difficult problem

of selecting the number of clusters, clustering has two advantages. First, the clusters need to be found only once rather than separately for each voxel. Second, clustering is more flexible, allowing a wide range of criteria, many developed in the literature on clustering for functional data (Cadez et al., 2000). Here, we will consider clustering based on the shape of the estimated responses, the estimated parameter vectors in response model, and mutual information.

Note that both variants of this method can be used for data in voxel coordinates or in cortical surface coordinates on an inflated cortex. We expect the gains to locality from using surface coordinates to outweigh the induced smoothing and correlation from doing the inflating, but we will compare results and simulations to determine which is the better domain for adaptive smoothing.

Bayesian Hierarchical Model with Collage Prior. Our second approach to adaptive smoothing is to build the smoothing into the inference through a Bayesian hierarchical model. If neighborhood dependencies among the responses are included in the model, the posterior distribution effectively smooths the data. A Bayesian formulation has several advantages over a classical approach here. In particular, discrete, qualitative parameters – like the smoothing regions – can be included easily in the model. Moreover, by building the smoothing directly into the model, downstream inferences are no longer distorted by the preprocessing steps.

We will begin with a pre-existing Bayesian model for the voxelwise data; several have been proposed (Genovese, 2000) and one can easily be constructed from the common General Linear Model formulation. Assume for now that β_v represents the d -dimensional vector of response parameters (e.g., log percent signal change for every experimental condition) for the voxelwise model at voxel v with all components on a common scale. We turn this voxelwise model into an adaptive smoothing model by adding layers to the hierarchy that make β dependent across voxels.

The most basic way to do this is to put a Markov Random Field (MRF) prior on the collection of β vectors (Cressie, 1993). The MRF decomposes the joint (prior) distribution of the β_v 's into lower-dimensional distributions on a collection of voxel subsets called *cliques*. That is, $f(\beta_V) \propto \prod_{\mathcal{C}} f(\beta_{\mathcal{C}})$, where the product is over all cliques \mathcal{C} and $\beta_{\mathcal{A}}$ denotes the collection of β_v 's for which $v \in \mathcal{A}$. A Markov property follows wherein the conditional distribution $f(\beta_v | \beta_{-v}) \propto \prod_{\mathcal{C}: v \in \mathcal{C}} f(\beta_{\mathcal{C}})$. This conditional distribution can thus be written solely in terms of the neighborhood of v : $\mathcal{N}_v = \bigcup_{\mathcal{C}: v \in \mathcal{C}} \mathcal{C}$. A common MRF model makes this conditional distribution a Gaussian whose mean is a weighted average of the neighbors:

$$\beta_v | \beta_{-v} \sim N \left(a_v + \frac{1}{\#(\mathcal{N}_v)} \sum_{u \in \mathcal{N}_v} w_{vu} \beta_u, \tau_v^2 I_d \right), \quad (12)$$

for a $d \times d$ identity I_d and for constants $a_v, w_{vu}, \tau_v^2 \in \mathbb{R}$. When β represents the log percent signal change, for example, we might take $a_v = -7$ (i.e., effectively zero) so the percent signal changes are a priori centered on the geometric mean of the neighboring voxels. Using a MRF prior for β requires specifying the cliques \mathcal{C} which results in a particular neighborhood structure (which must satisfy $u \in \mathcal{N}_v$ if and only if $v \in \mathcal{N}_u$). Given a voxelwise model, we can extend the hierarchy by including the MRF as the prior for the responses. A posterior sample can be efficiently constructed using MCMC (specifically Gibbs' sampling in the Gaussian case) as described in (Geman and Geman, 1984; Geman and McClure, 1987).

While simple, the MRF prior has two significant drawbacks for adaptive smoothing: the neighborhood structure is pre-specified, limiting adaptivity; and the MRF prior forbids sharp boundaries, oversmoothing fine structure. Methods to overcome these problems include adding additional layers of the hierarchy to allow breaking the dependence along boundaries (Geman and McClure, 1987) and carefully selecting cliques to conform to existing (e.g., anatomical) boundaries. While we will implement both approaches for adaptive smoothing as a testbed for comparison, here we propose a related but distinct alternative model.

As with the MRF, we will specify a prior on the response parameters β , which we call a *collage prior*. Rather than using a fixed neighborhood structure, the collage prior gives a distribution on the smoothing neighborhoods. The random process underlying the prior can be decomposed into the following steps:

1. Draw a set of voxels, the *centers*, from a (potentially inhomogenous) Poisson point process on the voxel grid. The intensity function of the point process is g and is fixed a priori to match anatomical and other prior information.

2. At each center location, place a random ellipsoid \mathcal{E} . Each ellipsoid is specified by its orientation ($\theta \in [0, \pi)$), radius ($r > 0$), and aspect ratio (ratio of semi-minor to semi-major radius, $0 \leq a \leq 1$), drawn independently from a Uniform, a fixed density h_1 , and a fixed density h_2 , respectively. We will choose h_1 and h_2 so that the typical active set under the collage prior is consistent with our empirical studies of the signal for high-resolution data.
3. At each voxel v and conditional on all the other voxels' responses, draw a value of β , independent across components, from a Normal distribution whose mean is the average of the average β 's over all ellipsoids to which v belongs.

Symbolically, the prior hierarchy is as follows:

$$\begin{aligned}
N, c_1, \dots, c_N &\sim \text{PoissonPointProcess}(g) \\
(\theta_1, r_1, a_1), \dots, (\theta_N, r_N, a_N) \mid N, \mathbf{c} &\stackrel{\text{iid}}{\sim} \text{Uniform}(0, \pi) \times h_1(r) \times h_2(a) \\
\beta_v \mid \beta_{-v}, \theta_1, r_1, a_1, \dots &\sim N \left(a_v + \frac{1}{\#\{i : v \in \mathcal{E}_i\}} \sum_{i: v \in \mathcal{E}_i} \left[b + \frac{1}{\#\mathcal{E}_i} \sum_{\substack{u \in \mathcal{E}_i \\ u \neq v}} \beta_u \right], \tau_v^2 \right),
\end{aligned}$$

where the constants a_v and b are such that voxels outside any ellipsoids have very negative values (i.e., essentially zero signal change) and where the outer average is taken to be zero if v does not lie in any ellipsoid. An alternative we have considered for the last stage of the hierarchy is to choose β_v from an even mixture of normals, each centered around the average of the other β 's in one ellipsoid containing v . (In this case, a voxel not in any ellipsoids would have β drawn from a $N(a_v, \tau_v^2)$ distribution, meaning that the response is effectively zero.)

Combining the collage prior with the likelihood (and some priors) from a voxelwise model for the data is straightforward. Sampling from the posterior requires a Metropolis-Hastings chain (Green, 1995a) that can, as ellipsoids are added and deleted, move across parameter spaces of different dimension. The reversible jump framework (Green, 1995b) works well in such problems, needing only selection and calibration of different move types to ensure detailed balance. In our preliminary implementations, though quite a bit slower than typical voxelwise fits, the technique has been successful.

The biggest advantage of the collage model over the MRF is adaptivity. The collage smoothing neighborhoods are not fixed a priori and can be combined to fit a wide variety of shapes, including those with complicated or fine-scale structure. While the MRF prior can be calibrated to favor certain qualitative features, its flexibility is limited. Moreover, the collage prior can produce sharp boundaries. Modifying an MRF to have similar flexibility with sharp boundaries is challenging and requires greater complexity and more parameters.

CAVEATS AND PITFALLS. There are several issues to be considered for both methods. For information neighborhoods, is mutual information a better choice of criterion than are shape-measures? The latter focus on what we most care about – a similar pattern of activation – but provides at most a summary of some features; the former gives a more comprehensive measure of dependence but may not distinguish subtle differences in shape. This is an empirical question that we will address, and in our preliminary analysis, mutual information does a good job of capturing the similarities we care about. Another question is how much the first step – the construction of the smoothing neighborhoods – affects downstream inferences. Adaptive smoothing is highly nonlinear and affects the distribution of the smoothed data. We will begin with data splitting, using two independent co-registered data sets with the same paradigm and subject, defining the neighborhoods on the training data, and using those neighborhoods to smooth. We will compare this with data and simulations to the results from using the same data for training and analysis. We expect that this will not be a large effect.

One issue for the Bayesian model is the sensitivity of the results to the structure of the voxelwise model. When the model fits well, the method will be efficient but it may not be as robust as finding information neighborhoods. Still the general success of the voxelwise models in the literature suggests that the fit will be sufficient. A related issue is the sensitivity of the results to the prior specification. We will study and quantitatively assess this sensitivity and develop tools for matching the prior to an expected spatial structure of the signal. Another issue is computation. The full MCMC is significantly more computationally costly than the typical analysis, but it remains completely feasible with currently available platforms and data sets.

DATA ANALYSIS.. An immediate goal in our development of adaptive smoothing methods is analysis of the high-resolution data sets described above. To that end, we will take the following steps:

1. Study and quantify the structure of the noise and signal, looking in particular at correlation structure, and typical size and shape of active regions. This effort will inform our model building and let us build effective simulations for evaluation of methods.
2. Implement the information neighborhoods method using the objective in equation (11). We will compare the results with the experimental data using standard smoothing and adaptive smoothing, both in voxel coordinates and in cortical surface coordinates on the inflated brain. We will also compare the methods using simulated data under a range of assumptions consistent with our study in step 1.
3. Using both experimental data and simulations, compare neighborhoods based on mutual information with those based on shape, and compare local two-component clusters with a full agglomerative clustering.
4. Implement the Bayesian hierarchical model with the collage priors. Tune and validate the MCMC simulation and determine mixing times using simulated data. Compare results using the Bayesian model with the adaptive smoothing on information neighborhoods. This comparison is complicated because the Bayesian method gives a full analysis, so we will use the Bayesian voxelwise model on the adaptively smoothed data. We will also compare the collage priors to an extended MRF prior that includes parameters to break dependency links (Geman and Geman, 1984) and allow sharp boundaries.
5. Finally, we will use the method that performs best in our assessments going forward for analysis of the remaining experimental data.

D.3.2. Do eye position signals have a fine-scale structure? (Question 3)

EXPERIMENT 3. The visual stimuli, illustrated in Figure. 10, will consist of a high-contrast checkerboard patterns flickering at 4 Hz. The checkerboard stimuli will form a set of concentric rings which will alternate with a complimentary set of anti-rings according to a simple block alternation protocol (period frequency of 1/24 s). In each functional run, 88 volumes (3 s / volume) will be collected corresponding to 11 cycles of the stimuli. Ten functional runs will be acquired in each scanning session. Stimulus presentations will interrupted by 500 ms blank stimuli every 3 sec to avoid perceptual filling in. The subjected will simply be instructed to fixate a high-contrast square at the center of the display.




Figure 10. Stimulus for testing effective spatial resolution. Block alternation of annulus / anti-annulus checkerboard stimuli. The annuli produce bands of activity across the cortical surface. Ring spacing is scaled with visual eccentricity to produce bands of different cortical spatial frequency. Cortical spatial frequency of 10 mm (A), 5 mm (B), and 2.5 mm (C).

Potential Outcomes. ATTN

Caveats and Pitfalls. ATTN

DATA ANALYSIS. The data will be analyzed to estimate the stimulus-evoked modulation of the fMRI signal (Bandettini et al., 1993; Engel et al., 1994). Prior to smoothing, the following preprocessing steps will be performed. 1) Data from the first cycle of block alternation will be discarded to allow the hemodynamic response and tissue magnetization to reach steady state. 2) Motion within each run will be corrected by aligning each volume to the first volume of the functional run. Motion between runs will be corrected by aligning the mean of each run to the mean of the first run (Nestares and Heeger, 2000a). Head motion parameters will be estimated after cropping each volume slightly to account for the strong intensity roll-off at the edges of the field of view due to the slab selection and outer-volume suppression. 3) The time series at each voxel was divided by its mean to convert the data from arbitrary image intensity units to percent signal modulation and to compensate for the distance from the receive coil. 4) The time series at each voxel will be high-pass filtered to compensate for slow signal drifts. Next, we will **apply the adaptive smoothing method** described above, using the variant that performs best with our validation data. On the smoothed data, we will carry out the following analysis. 1) The time series at each voxel will be fit with a sinusoid (24 s period). 2) The correlation between the time series and the corresponding

best-fitting sinusoid and the phase of the best-fitting sinusoid will be computed. The correlation value is a measure of the stimulus-evoked contrast to noise ratio, taking a value near 1 when the fMRI signal modulation at the block-alternation period is large relative to the noise (at the other frequency components) (Engel et al., 1997). The phase of the best fit sinusoid measured the temporal delay of the fMRI signal relative to the beginning of the experimental cycle, and consequently labeled cortical regions that responded to the rings versus the anti-rings.

D.4. TIMELINE FOR THE PROPOSED RESEARCH

Year 1. Under Aim 1, we will extend our preliminary work on community discovery methods, developing, and implementing, and beginning evaluation of the neighborhood construction and the empirical selection of γ . We will also begin implementation and testing of the one-step refinement method and extend our theoretical results on consistence and convergence of state-space reconstruction and module discovery. We will generate results using this method with our preliminary data. Under Aim 2, we will build on our preliminary work on sparse regression techniques for estimating functional response fields. We will incorporate smoothness constraints on the α and β parameters and test the implementation on very large problems.

Year 2. Under Aim 1, we will continue development of the one-step refinement method and refine our discretization scheme to improve performance. We will continue work on the theory underlying the method and investigate alternative estimators of mutual information. And we will apply and validate the method on data from Experiment 1. Under Aim 2, we will develop and test the Bayesian hierarchical model and compare its performance on simulation and data to the sparse regression technique. Under Aim 3, we will implement the submodular function minimization algorithm to compute information neighborhoods. We will also implement the simulation algorithm for the collage prior and perform initial comparisons with standard smoothing techniques. We will test and compare these two methods on new high-resolution data collected under Experiment ATTN.

Year 3. Under Aim 1, we will refine and validate our methods and continue analyzing experimental data. Under Aim 2, we will do the same. We will also apply our adaptive smoothing techniques to improve smoothing of the α and β parameters. We will adapt the information neighborhoods of Aim 3 as an alternative method for response field estimation. Under Aim 3, we will finish testing and validation of the information neighborhood and collage-prior techniques and perform a comprehensive comparison on data and simulation. We will also investigate the utility of adaptive smoothing to lower resolution data.

Year 4. By year 4, data collection, preliminary analysis, and basic methods development will be complete. During this year, we will focus on finalizing the data analysis of all three experiments, refining methods and extending them to more general settings, and disseminating the work.

Team Interaction. Our team has a good track record of interaction, even across the distance between our institutions. We will be in regular contact throughout the award period and will travel to collaborate with each other several times per year.

D.5. GENERAL METHODS

Subjects and scanning sessions. MR imaging will be performed in the NYU Center for Brain Imaging (see Resources). Each experiment will typically include 6 subjects. Each subject will participate in a large number scanning sessions: one session to obtain a high-resolution ($1 \times 1 \times 1$ mm) anatomical volume (using an MPRAGE pulse sequence); one session to measure the retinotopic maps in visual cortex; one session to measure hemodynamic impulse response functions; and several to measure fMRI responses in the experimental conditions. Subjects will practice the tasks extensively until they have achieved asymptotic performance, before scanning. Scanning sessions will be ~ 1.5 hours, providing sufficient time for 10–12 fMRI scans, each ~ 4 min. Each experimental condition will be repeated several times in each subject within and across several scanning sessions. We will first identify each of several visual cortical areas in each subject’s brain, analyze the data separately for each subject, and then compute the mean responses across subjects. If we note substantial individual differences in the fMRI data, then instead of averaging across subjects we will attempt to correlate those individual differences in cortical activity with individual differences in behavioral performance. We have used this approach—many repeated measurements in each of a number of experienced, well-practiced, motivated subjects (as is typical in psychophysics and awake-monkey electrophysiology)—with success in our previous work (Boynton et al., 1999; Gandhi et al., 1999; Ress et al., 2000; Zenger-Landolt and Heeger, 2003; Montaser-Kouhsari et al., 2007; Larsson et al., 2006; Ress and Heeger, 2003; Lee et al., 2007; Lee et al., 2005).

Visual stimuli. Stimuli will be presented either via a calibrated LCD flat panel display or an LCD video projector (see *Equipment*). Subjects will lie on their backs in the bore of the MR scanner and view the display through an angled mirror. Stimulus conditions will otherwise be matched to those in the psychophysics experiments.

Registration and flat maps. Each session will begin by acquiring a set of anatomical images using an MPRAGE pulse sequence in the same slices as the functional images. An image registration algorithm²⁵⁸, accurate to within ~ 1 mm, will be used to align these inplane anatomical images to a high-resolution anatomical volume of the subject's brain (acquired in a separate scanning session), so that we will be able to localize the same cortical areas across scanning sessions. Custom software (Teo et al., 1997; Wandell et al., 2000; Dougherty et al., 2003; Smith et al., 1999) will be applied to the high-resolution anatomical images of each subjects brain to restrict the functional data analyses to gray matter voxels, and to computationally flatten the gray matter to create flattened visualizations of the cortical activity.

fMRI acquisition. During each fMRI scan, a time series of volumes will be acquired using a T2* sensitive EPI pulse sequence. We have considerable flexibility in the acquisition parameters, between covering more brain and acquiring higher quality (sensitivity, reliability, spatial and temporal resolution) data. We can, for example, acquire either 27 slices ($3 \times 3 \times 3$ mm) every 1.5 s, 24 slices ($2 \times 2 \times 2$ mm) every 1.5 s, 14 slices ($3 \times 3 \times 3$ mm) every 0.8 s, or 28 slices ($700 \times 700 \times 700$ μm) over a restricted field of view ($2.1 \times 4.8 \times 19.2$ cm) every 3 s.

fMRI data preprocessing. Data will be preprocessed as follows: 1) discarding the first 9-30 s of each fMRI scan to minimize transient effects of magnetic saturation and to allow subjects to practice the behavioral tasks; 2) correcting for any residual head movements (although subjects will be stabilized on a bite bar) within each scan and across scans using custom software (Nestares and Heeger, 2000b); 3) interpolating in time to compensate for the differential slice acquisition times; 4) high-pass filtering with a cutoff frequency of 0.01 Hz to compensate for the slow signal drift (Smith et al., 1999); and 5) dividing each voxel's time course by its mean intensity to convert the data from arbitrary image intensity units to percent signal modulation and to compensate for distance from the receive coil.

Mapping retinotopic visual areas. We will define a total of 16 retinotopic visual areas (V1, V2, V3, V4, VO1, VO2, LO1, LO2, MT, MST, V3A, V3B, IPS1, IPS2, IPS3, IPS4) by measuring the radial and polar angle components of the cortical retinotopic map. Details of the experimental protocols and data analysis are provided in our previous publications (Silver et al., 2005; Huk et al., 2002; Schluppeck et al., 2005; Larsson and Heeger, 2006). We will use the Wandell et al. definition of V4 (Wandell et al., 2007), but we note that this is controversial (Tootell and Hadjikhani, 2001; Brewer et al., 2005; Hansen et al., 2007), so we will confirm that our conclusions will not differ qualitatively if we adopt any of the other proposed definitions. Visual areas will be further restricted to correspond to the portion of visual field where the stimuli are presented. In a separate block-alternation experiment, subjects will maintain fixation while a stimulus (identical to that in the main experiment) alternates with a uniform gray screen (12 cycles of alternation, 24 s each). Data will be analyzed by averaging across the repeated scans, fitting a sinusoid to the time series separately for each voxel, and then choosing a correlation threshold and phase window to select contiguous regions of gray matter that respond strongly to, and in phase with, the block alternation. This procedure for restricting the visual areas will be done separately for each experiment if the stimuli differ across experiments. However, for conditions in which we compare the responses across experimental conditions (e.g., while manipulating stimulus size), we will of course use the same restriction for the analysis of both conditions (e.g., corresponding to the smaller stimulus size).

Event-related fMRI protocols and data analysis. In each experiment, different trial types will be presented in random order and typically with randomized inter-trial intervals (3-9 s). Data will be analyzed separately for each of the predefined retinotopic cortical areas, by averaging the measured response time courses across voxels within each visual area ROI. The mean fMRI response time course, evoked by each trial type, will be estimated using deconvolution (Dale, 1999; Buracas and Boynton, 2002); which is equivalent to selective averaging with correction for overlap between temporally-adjacent responses, based on the assumption that hemodynamic responses superimpose linearly over time (Boynton et al., 1996) (for details see Donner et al., 2008). The randomized inter-trial intervals improve the efficiency of the experiment so that the mean response time courses can be estimated more accurately with fewer trials. In a complementary analysis of the data, response amplitudes will be estimated, again using linear regression, by adopting a model of the hemodynamic impulse-response function (HIRF). The HIRF will be measured separately for each visual area in each subject (see below). For both the mean response time course and response amplitude analyses, the amount of variance accounted for by the linear regression, r^2 , will be

used to quantify the goodness of fit (Gardner et al., 2005; Gardner et al., 2008). Statistical significance (p-values) for the r^2 values will be computed using a permutation analysis, randomizing the trial times and recalculating the r^2 values a large number of times to estimate of the distribution of r^2 values expected by chance. Only voxels that have reliable visual responses, as assessed by this randomization test, will be included in the analyses. Our published results demonstrate that fMRI data can be misinterpreted unless care is taken to ensure that only such reliable measurements are included in the analysis (Gardner et al., 2005). Randomization tests will be used to determine the statistical significance of the estimated response amplitudes (whether or not they are different from zero and/or whether or not they are statistically different from one another for the different trial types).

Hemodynamic impulse-response functions. Our published results demonstrate the importance of estimating the HIRF separately for each subject (Offen et al., 2008). HIRFs will be measured for each subject in a separate experiment. Target stimuli (identical to those presented in each of the main experiments) will be presented briefly (200 ms) at randomized intervals (3-9 s). Subjects will perform a demanding task at fixation to control attention. Data will be preprocessed as described above, and then fit (using nonlinear least-squares) with a difference of two gamma functions (Glover, 1999).

Eye-movement recording during MRI scanning. Trials in which subjects break fixation will be discarded from further analysis. Eye position will be monitored using an MRI-compatible eye tracker (see Resources). Eye position will be calibrated at the beginning of the session and between scans (i.e., every ~ 4 min), and transformed to visual space. Fixation will be verified by identifying the trials in which the eye deviates from a pre-determined fixation window, and measuring the SD of eye position during fixation. Our published results demonstrate that we can routinely monitor fixation with an accuracy of $\pm 0.5^\circ$ (Silver et al., 2005).

F. Progress Report Publication List

F.0.1. Research papers

- Genovese, C.R. and Wasserman, L. (to appear). Adaptive Confidence Bands. *The Annals of Statistics*.
- Genovese, C.R. and Wasserman, L. (2006). Exceedance Control for the False Discovery Proportion, *Annals of Statistics*, **101**, No. 476, 1408–1417.
- Genovese, C.R., Roeder, K., and Wasserman, L. (2006). False Discovery Control With P-Value Weighting, *Biometrika* 93(3), 509–524.
- Merriam, E.P., Genovese, C.R., and Colby, C.L. (2007) Remapping in human visual cortex. *Journal of Neurophysiology* (with commentary).
- Merriam, E.P., Genovese, C.R., Lasson, M., and Colby, C.L. Transhemispheric remapping in the absence of the corpus callosum.
- Merriam, E.P. and Colby, C.L. (2005). Active vision in parietal and extrastriate cortex. *The Neuroscientist*. 11(5): 484-93
- Perone Pacifico, M., Genovese, C. R., Verdinelli, I., and Wasserman, L. (2004). False Discover Control for Random Fields, to appear *Journal of the American Statistical Association*.
- Perone Pacifico, M., Genovese, C. R., Verdinelli, I., and Wasserman, L. (2005). Scan Clustering: A False Discovery Approach, under review.

F.0.2. Other Refereed Publications

- Genovese, C. R. and Wasserman, L. (2004). Bayesian and Frequentist Multiple Testing, *Bayesian Statistics 7*, eds. Bernardo, J.M., Bayarri, M.J., Berger, J.O., Dawid, A.P., Heckerman, D., West, M., p 145-162

F.0.3. Abstracts

- Merriam, E.P., Genovese, C.R., Lassonde, M., and Colby, C.L. (2006). Spatial updating in a subject with a complete commisurotomy. *Annual meeting of the Society for Neuroscience*, Atlanta, GA.
- Merriam, E.P., Genovese, C.R., Lassonde, M., and Colby, C.L. (2006). Spatial updating in the absence of the corpus callosum. *Cognitive Neurosciences*. New York, NY.
- Merriam, E.P., Genovese, C.R., and Colby, C.L. (2005). Spatial updating in human cortex: a Bayesian analysis. *NSF IGERT project meeting*. Washington, DC
- Merriam, E.P., Genovese, C.R., and Colby, C.L. (2004a). Modulation of visual activity in striate, extrastriate, and parietal cortex. *Annual meeting of the Society for Neuroscience*, San Diego, CA.
- Merriam, E.P., Genovese, C.R., and Colby, C.L. (2004b). Spatial updating in human extrastriate visual cortex. Center for Neuroscience University of Pittsburgh (CNUP) annual retreat.
- Merriam, E.P., Genovese, C.R., and Colby, C.L. (2004c). Dynamic receptive fields in human visual cortex. Science2004: No Boundaries, University of Pittsburgh.

F.0.4. Manuscripts in Progress

- Genovese, C.R., Shalizi, C., and Merriam, E.P. Community Discovery Methods for Estimating Functional Connectivity.
- Genovese, C.R., Huang, E., and Merriam, E.P. Analysis of White Noise Models in fMRI.
- Genovese, C.R., Merriam, E. P., and Freeman, P.F. (2006a). False Discovery Control for fMRI: An Assessment
- Genovese, C.R., Merriam, E.P., and Freeman, P.F. (2006b). Bayesian Analysis of fMRI Data: A Review

G. Protection of Human Subjects

G.1. RISKS TO HUMAN SUBJECTS

G.1.1. Human subjects involvement and characteristics

Behavioral tasks. The tasks will involve looking at visually presented patterns on a computer monitor, listening to auditorally presented tones or words; performing cognitive activity (e.g., remembering a visual pattern or making a decision about a visual stimulus); making a response (e.g., a button press or an eye movement) regarding the stimuli perceived; making hand movements.

Psychophysics experiments. Subjects will be instructed on how to perform the various visual and cognitive tasks. They will perform these tasks in blocks of several minutes, with rest periods as needed. Multiple sessions will typically take place over a period of several weeks and each such session will last about 1 hour.

fMRI experiments. Subjects will be screened ahead of time to determine eligibility. They will be given the NYU MRI screening questionnaire prior to participation. Subjects will be excluded if they indicate any risk factor on these questionnaires, which include pregnancy. Women who are or who suspect they may be pregnant will be excluded from these studies. Once a subject passes the screening, s/he will be instructed to remove all metal objects before entering the magnet room (see below for details). After entering the magnet room, subjects will be given earplugs to wear and instructed to lie on the MR table. Foam pads will be placed around their heads to limit head movement during the scan. The table will then be slid into the scanner so that the head and upper body are inside the magnet tube. They will view standard projection of visual stimuli onto a screen in the back of the bore via a mirror.

Conventional anatomical scans will be taken during a period of about 15 minutes. High-speed functional MRI images will be obtained during the rest of the session. Each functional MRI scan will last between 3-10 minutes. During the functional scans, subjects will be asked to hold still. Following each scan, the experimenter/technician operating the MRI scanner will ask the subject how they are doing and if they want to continue. If they like, subjects can take a brief rest in the scanner and then continue. They will remain in the scanner during breaks, but they may close their eyes and rest. Subjects can decide to stop at any time and will be taken them out of the scanner immediately. During the session they will hold a panic squeeze ball that will enable them to notify the experimenter if they need immediate attention. Each scanning session will typically last 1.5 hours.

Subject population. Our subject population will include approximately 50% men and 50% women (age range: 18–65), and will reflect the demographics at New York University as follows (see <http://www.nyu.edu/ir/pdf/demographics/2007/Enrollment.pdf>):

Hispanic: 6%
Asian/Pacific Islander: 15%
American Indian/Alaskan Native: 0%
Black (non-hispanic): 5%
White (non-hispanic): 45%
International: 12%
Unknown/other: 17%

These subjects will be found by advertising on campus electronic bulletin boards and by posting announcements on old-fashioned bulletin boards. These subjects will be paid \$25/hour for fMRI experiments and \$10/hour for psychophysics experiments. Some of the subjects will be laboratory personnel who volunteer (i.e., they are not required) to participate in the proposed experiments. Lab personnel will be instructed that their participation is entirely voluntarily. Some lab members will choose to serve as subjects in each others' experiments (i.e., payment "in kind" for the time). Others, who do not wish to participate as subjects, will be asked to "pay in kind" by volunteering to perform other services such as operating the scanner for someone else's experiment. Subjects who are unable to attend to or respond to stimuli as required by the experimental protocol will be excluded. There will be no cost to the subjects for participation in this study.

G.1.2. Sources of materials

N/A

G.1.3. Potential risks

The risks from this study to the subject will be minimal. Subjects will be informed that they should contact the PIs if they have experienced a research related injury.

Psychophysics experiments. The likelihood and severity of physical or psychological harm does not exceed that in ordinary life or during routine examinations or tests—risks to individuals is minimal (with the most likely negative consequence being boredom).

fMRI experiments. There are no known significant risks or side effects associated with MRI procedures. The magnetic fields, at the strengths used are not harmful and our MRI scanning procedures fall within the FDA guidelines for radiofrequency electromagnetic field exposure. These field strengths are at safe levels and less hazardous than a comparable x-ray computed tomography examination. Exceptions include if a person has electrically, magnetically activated implants (such as cardiac pacemakers), or to those who have clips on blood vessels in their brain, or other metallic objects in their body such as shrapnel, bullets, buckshot, or metal fragments. Therefore, subjects will be carefully screened for previous exposure to metallic fragments or to implanted devices. The most serious potential risks are related to the possibility of ferromagnetic objects in the vicinity of the high-field magnet in the scanner. These objects could conceivably become projectiles due to the powerful magnetic field. Therefore, subjects will be asked to place all metallic and magnetic objects in their possession (e.g. keys, jewelry, credit cards) in a locker available outside the magnet room. Although there are no known risks of an MRI scan to the unborn fetus, we do not permit participation by any woman who is or suspects she may be pregnant.

Most people do not find a magnetic resonance scan uncomfortable. However, on occasion some subjects have reported mild discomfort. The following are some types of discomfort that have been reported with a magnetic resonance scan. Some subjects have experienced claustrophobia (fear of enclosed spaces). Subjects will be asked to lie on a table that slides into a horizontal cylinder only slightly wider in all directions than their body and the head is secured to help minimize extra movement. If they are prone to claustrophobia they should notify the researcher in charge of the scan. The MR scanner makes loud knocking or beeping sounds during imaging; earplugs will be provided to help reduce this noise. Due to the rapid rate of change of the magnetic gradients during imaging, the possibility exists for peripheral nerve stimulation. If this happens, subjects may feel creeping or tingling sensations, typically along their arms or lower back. Dizziness and nausea may occur if the subjects move their head in the bore of the magnet. Some subjects find it uncomfortable to have a dental impression for more than a brief period and have found it can interfere with swallowing, so bite bars will be used only with experienced subjects (typically laboratory personnel). Finally, there may be some heating from the radio frequency coils, the cables to the coils, response and physiological monitoring devices. The machine is calibrated so that this heating will be no more than one degree of body temperature.

Subjects will be instructed to notify the investigators, as soon as possible, if at any time they feel uncomfortable, no matter what the reason. Subjects will be in contact with the research staff throughout the study through a microphone mounted on the MRI scanner. They will also be instructed in how to use an emergency handheld squeeze ball to inform the operator if they wish to immediately stop scanning and be removed from the magnet. The MRI can be stopped at any time at their request.

G.2. ADEQUACY OF PROTECTION AGAINST RISKS

G.2.1. Recruitment and informed consent

The experiments will be undertaken in compliance with the safety guidelines for MRI research. Prior to participating in the study, subjects will be informed of the risks and benefits of the study, and will fill out and sign the screening questionnaire and informed consent. Experimenters will verbally question the subjects to make sure that the screening questionnaire and informed consent have been understood. Subjects will be able to withdraw from the study at any time with no penalty (to date, we have not had any withdrawals). The human subject protocols will be reviewed annually by the University Committee on Activities Involving Human Subjects (UCAIHS).

Confidentiality of research records will be strictly maintained. All data will be stored on the PIs computer file server, under password protection. Only the PIs and their designated research associates will have access to the data. The results of these studies may be published in a book or journal or used for teaching purposes. However, subjects names or other identifiers will not be used in any publication or teaching material without specific permission. Data will be backed up on to CD, DVD, or computer tape. These backup media will be

stored in a cabinet in the PIss laboratory. Subject identity will be coded on all documents so as not to breach confidentiality. After publication, the data will be removed from the PIss computer file server. The backup media (CDs, DVDs, computer tapes) and other documents will be kept indefinitely.

G.2.2. Protection against risk

Anyone involved in data collection must have successfully completed the UCAIHS training course at NYU. Only certified investigators may be involved in subject recruitment and administration of informed consent.

To ensure safety during MRI experiments, every researcher at NYU will complete a safety training course, every year. These safety courses are offered by the staff of the NYU Center for Brain Imaging. Subjects will be carefully screened to make sure they do not have any metal before being taken into the room with the MRI scanner. Subjects will wear earplugs to protect their hearing while in the MRI scanner.

Volunteers will be excluded for the usual contraindications to MRI exams, including pacemakers, surgical aneurysm clips, known metal fragments embedded in the body including eyes. Women of childbearing age will not be included in this study if they are or think they might be pregnant. Subjects with suspected cerebrovascular or pulmonary disease or a history of such will be excluded. Subjects with a history of migraine, arterial hypertension, coronary heart disease, asthma, anemia, or epilepsy will be excluded. Subjects who have experienced claustrophobia will be excluded. Subjects will be carefully screened for these conditions prior to scanning, using a standardized form/checklist.

G.3. POTENTIAL BENEFITS OF THE PROPOSED RESEARCH TO THE HUMAN SUBJECTS AND OTHERS

There will be no direct benefits to subjects for their participation in the proposed experiments. The physiological and psychological knowledge to be gained from these studies justifies the use of human subjects.

G.3.1. Incidental findings

On occasion the brain images of a subject may reveal a potential brain abnormality. NYU has adopted a policy concerning incidental findings that is in accordance with the recommendations of an NIH-sponsored workshop on MRI Research Safety and Ethics. The summary report from that workshop is available online: <http://www.nimh.nih.gov/about/advisory-boards-and-groups/namhc/reports/mri-research-safety-ethics.pdf>

G.4. IMPORTANCE OF THE KNOWLEDGE TO BE GAINED

There are a number of ways in which the proposed research is aligned with the research objectives of the National Eye Institute. The following are excerpts from the National Plan for Eye and Vision Research (2003) stating goals and objectives for the Strabismus, Amblyopia, and Visual Processing Program:

- Understand how the brain processes visual information, how neural activity is related to visual perception, and how visual processing interacts with other brain systems underlying cognition.
- Study the function, circuitry, and development of higher order visual areas to determine the effect of attention and top-down influences on visual processing.
- Develop a mechanistic understanding of the origin of the signals that control attention and how they alter the responses of neurons in visual processing and sensorimotor transformations.
- Develop imaging technologies to further understand the neural basis of amblyopia and related visual deficits.
- Bridge the knowledge of what happens at the cellular level in the visual system with the knowledge of visual psychophysics.
- Exploit the knowledge of the functional organization of the visual system in animal models with noninvasive imaging studies in humans. The proposed research will contribute to each of these.

The experimental protocols and theoretical principles that we develop for studying normal vision will be readily applicable to patient populations. A better understanding of visual attention will lead to a better understanding of factors limiting peripheral vision, that are critical when central vision is compromised due to macular degeneration and related visual deficits. Basic knowledge of visual attention also has implications for our understanding of several neuropsychological conditions, including unilateral neglect, schizophrenia and ADHD, and for informing the development of diagnostic tests. We have an established track-record of pursuing opportunities for translational research (see Specific Aims).

H. Inclusion of Women and Minorities

Inclusion of women. Our subject population will include approximately 50% men and 50% women, reflecting the population as a whole. We will ensure very nearly 50% women by drawing actively from the undergraduate students, graduate students, and postdocs at NYU.

Inclusion of minorities. Our subject population will reflect the demographics at New York University (see above). Should we find ourselves falling short of the targeted number of minority subjects, we will actively draw from the NYU undergraduate population. The NYU undergraduate population is sufficiently diverse that it should be possible to achieve this target.

I. Inclusion of Children

NYU undergraduates under the age of 21 will represent children in our sample of subjects. We believe that it is necessary to exclude children under the age of 18 for the following reasons. 1) fMRI experiments of any type require a great deal of patience and concentration. Subjects must lie very still for long periods of time in an extremely small space without moving their heads. 2) The behavioral tasks, while they are not conceptually complicated, are nonetheless demanding and require long periods of sustained attention. 3) The fact that we use relatively small sample sizes increases the importance of having subjects that are extremely reliable and capable of performing the behavioral tasks at a high level of accuracy. 4) Our stated goal is to study cortical function in the mature brain. This is not a developmental study. Including subjects of various different ages ranging from children to adults could introduce a confound in the interpretation of the results.

J. Vertebrate Animals

N/A

K. Select Agent Research

N/A

L. Multiple PI Leadership Plan

The proposed research is largely based on Genovese's past statistical work on methods for fMRI (ATTN refs) and Heeger's past fMRI work on visual reference frames (ATTN refs). Because the proposed work relies equally on our combined expertise, it is our understanding that we should both be listed as Principal Investigators. We have a strong established track record of collaboration.

Following the instructions for grant applications with multiple PIs, we have developed a leadership plan including governance and organizational structure of the leadership team, communication plans, a process for making decisions on scientific direction, and procedures for resolving conflicts. Genovese will be the lead PI overall. He will be primarily responsible for directing the statistical work. Heeger will be primarily responsible for directing the fMRI experiments. It is expected, however, that two PIs will contribute roughly equally to all of the proposed research. Funding will be shared as outlined in the Personnel Justification and Additional Narrative Justification, i.e., no budget allocation is planned for the specific components of the project nor to the individual PIs. The postdoctoral research fellow will be jointly recruited and co-advised by the two PIs.

M. Consortium/Contractual Arrangements

N. Letters of Support

N/A

O. Resource Sharing Plan

N/A

P. Literature Cited

- Andersen, R., Bracewell, R., Barash, S., and Gnadt, J. (1990a). Eye position effects on visual, memory, and saccade-related activity in areas lip and 7a of macaque. *Journal of Neuroscience*.
- Andersen, R. A., Bracewell, R. M., Barash, S., Gnadt, J. W., and Fogassi, L. (1990b). Eye position effects on visual, memory, and saccade-related activity in areas lip and 7a of macaque. *J Neurosci*, 10(4):1176–96.
- Andersen, R. A. and Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annu Rev Neurosci*, 25:189–220.
- Avillac, M., Denève, S., Olivier, E., Pouget, A., and Duhamel, J.-R. (2005). Reference frames for representing visual and tactile locations in parietal cortex. *Nat Neurosci*, 8(7):941–9.
- Bandettini, P., Jesmanowicz, A., Wong, E., and Hyde, J. (1993). Processing strategies for time-course data sets in functional MRI of the human brain. *Magn Reson Med*, 30(2):161–73.
- Billingsley, P. (1961). *Statistical Inference for Markov Processes*. University of Chicago Press, Chicago.
- Boussaoud, D. and Bremmer, F. (1999). Gaze effects in the cerebral cortex: reference frames for space coding and action. *Experimental brain research Experimentelle Hirnforschung Expérimentation cérébrale*, 128(1-2):170–80.
- Boynton, G. M., Demb, J. B., Glover, G. H., and Heeger, D. J. (1999). Neuronal basis of contrast discrimination. *Vision Res*, 39(2):257–69.
- Boynton, G. M., Engel, S. A., Glover, G. H., and Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human v1. *J Neurosci*, 16(13):4207–21.
- Bremmer, F., Schlack, A., Duhamel, J. R., Graf, W., and Fink, G. R. (2001). Space coding in primate posterior parietal cortex. *Neuroimage*, 14(1 Pt 2):S46–51.
- Brewer, A. A., Liu, J., Wade, A. R., and Wandell, B. A. (2005). Visual field maps and stimulus selectivity in human ventral occipital cortex. *Nat Neurosci*, 8(8):1102–9.
- Bühlmann, P. and Wyner, A. J. (1999). Variable length Markov chains. *The Annals of Statistics*, 27:480–513.
- Buracas, G. T. and Boynton, G. M. (2002). Efficient design of event-related fmri experiments using m-sequences. *Neuroimage*, 16(3 Pt 1):801–13.
- Cadez, I. V., Gaffney, S., and Smyth, P. (2000). A general probabilistic framework for clustering individuals and objects. In Ramakrishnan, R., Stolfo, S., Bayardo, R., and Parsa, I., editors, *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 140–149, New York. ACM Press.
- Colby, C. L. (1998). Action-oriented spatial reference frames in cortex. *Neuron*, 20(1):15–24.
- Colby, C. L. and Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annu Rev Neurosci*, 22:319–49.
- Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley, New York.
- Cressie, N. A. C. (1993). *Statistics for Spatial Data*. Wiley, New York, revised edition.
- Cytowic, R. E. (1996). *The Neurological Side of Neuropsychology*. MIT Press, Cambridge, Massachusetts.
- Dale, A. M. (1999). Optimal experimental design for event-related fmri. *Human brain mapping*, 8(2-3):109–14.
- David, O., Cosmelli, D., and Friston, K. J. (2004). Evaluation of different measures of functional connectivity using a neural mass model. *NeuroImage*, 21:659–673.
- d’Avossa, G., Tosetti, M., Crespi, S., Biagi, L., Burr, D. C., and Morrone, M. C. (2007). Spatiotopic selectivity of bold responses to visual motion in human area mt. *Nat Neurosci*, 10(2):249–55.

- DeSouza, J. F. X., Dukelow, S. P., and Vilis, T. (2002). Eye position signals modulate early dorsal and ventral visual areas. *Cereb Cortex*, 12(9):991–997.
- Donner, T. H., Sagi, D., Bonneh, Y. S., and Heeger, D. J. (2008). Opposite neural signatures of motion-induced blindness in human dorsal and ventral visual cortex. *J Neurosci*, 28(41):10298–310.
- Dougherty, R. F., Koch, V. M., Brewer, A. A., Fischer, B., Modersitzki, J., and Wandell, B. A. (2003). Visual field representations and locations of visual areas v1/2/3 in human visual cortex. *Journal of Vision*, 3(10):586–98.
- Duhamel, J. R., Colby, C. L., and Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255(5040):90–2.
- Duhamel, J. R., Colby, C. L., and Goldberg, M. E. (1998). Ventral intraparietal area of the macaque: congruent visual and somatic response properties. *J Neurophysiol*, 79(1):126–36.
- Edmonds, J. (1971). Matroids and the greedy algorithm. *Mathematical Programming*, 1:127–136.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, 32:407–499.
- Engel, S., Glover, G., and Wandell, B. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex*, 7(2):181–92.
- Engel, S., Rumelhart, D., Wandell, B., Lee, A., Glover, G., Chichilnisky, E., and Shadlen, M. (1994). fMRI of human visual cortex. *Nature*, 369(6481):525.
- Friston, K. (2002). Beyond phrenology: What can neuroimaging tell us about distributed circuitry? *Annual Review of Neuroscience*, 25:221–250.
- Gandhi, S. P., Heeger, D. J., and Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci USA*, 96(6):3314–9.
- Gardner, J. L., Merriam, E. P., Movshon, J. A., and Heeger, D. J. (2008). Maps of visual space in human occipital cortex are retinotopic, not spatiotopic. *J Neurosci*, 28(15):3988–99.
- Gardner, J. L., Sun, P., Waggoner, R. A., Ueno, K., Tanaka, K., and Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. *Neuron*, 47(4):607–20.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741.
- Geman, S. and McClure, D. E. (1987). Statistical methods for tomographic image reconstruction. *Bulletin of the International Statistical Institute*, LII-4:5–21.
- Genovese, C. R. (2000). A bayesian time-course model for functional magnetic resonance imaging data (with discussion). *Journal of the American Statistical Association*, 95(451):691–719.
- Gerstein, G. L. and Perkel, D. H. (1969). Simultaneously recorded trains of action potentials: analysis and functional interpretation. *Science*, 164(881):828–30.
- Glover, G. H. (1999). Deconvolution of impulse response in event-related bold fmri. *Neuroimage*, 9(4):416–29.
- Gray, R. M. (1988). *Probability, Random Processes, and Ergodic Properties*. Springer-Verlag, New York.
- Green, P. J. (1995a). Markov chain Monte Carlo in image analysis. In Gilks, W., Richardson, S., and Spiegelhalter, D., editors, *Markov Chain Monte Carlo in Practice*. Chapman and Hall, London.
- Green, P. J. (1995b). Reversible jump markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82:711–732.
- Hansen, K. A., David, S. V., and Gallant, J. L. (2004). Parametric reverse correlation reveals spatial linearity of retinotopic human v1 bold response. *Neuroimage*, 23(1):233–41.

- Hansen, K. A., Kay, K. N., and Gallant, J. L. (2007). Topographic organization in and near human visual area v4. *J Neurosci*, 27(44):11896–911.
- Haslinger, R., Klinkner, K. L., and Shalizi, C. R. (2008). The computational structure of spike trains. *Neural Computation*, forthcoming.
- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, New York.
- Heckman, N. and Zamar, R. (2000). Comparing the shapes of regression functions. *Biometrika*, 87:135–144.
- Hirata, Y., Judd, K., and Aihara, K. (2005). Characterizing chaotic response of a squid axon through generating partitions. *Physics Letters A*, 346:141–147.
- Hirata, Y., Judd, K., and Kilminster, D. (2004). Estimating a generating partition from observed time series: Symbolic shadowing. *Physical Review E*, 70:016215.
- Huk, A. C., Dougherty, R. F., and Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas mt and mst. *J Neurosci*, 22(16):7195–205.
- Iwata, S., Fleischer, L., and Fujishige, S. (2001). A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM*, 48:761–777.
- Kantz, H. and Schreiber, T. (1997). *Nonlinear Time Series Analysis*. Cambridge University Press, Cambridge, England.
- Kennel, M. B. and Buhl, M. (2003). Estimating good discrete partitions from observed data: symbolic false nearest neighbors. *Physical Review Letters*, 91:084102.
- Klinkner, K. L., Shalizi, C. R., and Camperi, M. F. (2006). Measuring shared information and coordinated activity in neuronal networks. In Weiss, Y., Schölkopf, B., and Platt, J. C., editors, *Advances in Neural Information Processing Systems 18 (NIPS 2005)*, pages 667–674, Cambridge, Massachusetts. MIT Press.
- Knight, F. B. (1975). A predictive view of continuous time processes. *Annals of Probability*, 3:573–596.
- Kontorovich, L. (2007). *Measure Concentration of Strongly Mixing Processes with Applications*. PhD thesis, Carnegie Mellon University.
- Kontorovich, L. and Ramanan, K. (2007). Concentration inequalities for dependent random variables via the martingale method. *Annals of Probability*, forthcoming.
- Krause, A. and Guestrin, C. (2007). Near-optimal observation selection using submodular functions. In *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI)*, Cambridge, Massachusetts. MIT Press.
- Larsson, J. and Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *J Neurosci*, 26(51):13128–42.
- Larsson, J., Landy, M. S., and Heeger, D. J. (2006). Orientation-selective adaptation to first- and second-order patterns in human visual cortex. *Journal of Neurophysiology*, 95(2):862–81.
- Lee, S.-H., Blake, R., and Heeger, D. J. (2005). Traveling waves of activity in primary visual cortex during binocular rivalry. *Nat Neurosci*, 8(1):22–3.
- Lee, S.-H., Blake, R., and Heeger, D. J. (2007). Hierarchy of cortical responses underlying binocular rivalry. *Nat Neurosci*, 10(8):1048–54.
- Lovasz, L. (1983). Submodular functions and convexity. In Bachem, A., Grotschel, M., and Korte, B., editors, *Mathematical Programming: The State of the Art*, pages 235–257. Springer-Verlag, New York.
- Luria, A. R. (1973). *The Working Brain: An Introduction to Neuropsychology*. Basic Books, New York.
- Merriam, E. P. and Colby, C. L. (2005). Active vision in parietal and extrastriate cortex. *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry*, 11(5):484–93.

- Miller, G. A. (1955). Note on the bias on information estimates. In Quastler, H., editor, *Information Theory in Psychology: Problems and Methods II-B*, pages 95–100.
- Montaser-Kouhsari, L., Landy, M. S., Heeger, D. J., and Larsson, J. (2007). Orientation-selective adaptation to illusory contours in human visual cortex. *J Neurosci*, 27(9):2186–95.
- Nestares, O. and Heeger, D. (2000a). Robust multiresolution alignment of MRI brain volumes. *Magn Reson Med*, 43(5):705–15.
- Nestares, O. and Heeger, D. J. (2000b). Robust multiresolution alignment of mri brain volumes. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 43(5):705–15.
- Newman, M. E. J. (2006). Finding community structure in networks using the eigenvectors of matrices. E-print, arxiv.org, physics/0605087.
- Newman, M. E. J. and Girvan, M. (2003). Finding and evaluating community structure in networks. *Physical Review E*, 69:026113.
- Offen, S., Schluppeck, D., and Heeger, D. (2008). The role of early visual cortex in visual short-term memory and visual attention. *Vision Res*.
- Ornstein, D. S. and Weiss, B. (1990). How sampling reveals a process. *The Annals of Probability*, 18:905–930.
- Paninski, L. (2003). Estimation of entropy and mutual information. *Neural Computation*, 15:1191–1254.
- Pouget, A. and Sejnowski, T. J. (1997). A new view of hemineglect based on the response properties of parietal neurones. *Philos Trans R Soc Lond, B, Biol Sci*, 352(1360):1449–59.
- Quaia, C., Optican, L., and Goldberg, M. (1998). The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural networks : the official journal of the International Neural Network Society*, 11(7-8):1229–1240.
- Queyranne, M. (1995). A combinatorial algorithm for minimizing symmetric submodular functions. In *SODA: ACM-SIAM Symposium on Discrete Algorithms (A Conference on Theoretical and Experimental Analysis of Discrete Algorithms)*.
- Quian Quiroga, R., Kraskov, A., Kreuz, T., and Grassberger, P. (2002). Performance of synchronization measures in real data: case studies on electroencephalographic signals. *Physical Review E*, 65:041903.
- Reichardt, J. and Bornholdt, S. (2004). Detecting fuzzy community structures in complex networks with a Potts model. *Physical Review Letters*, 93:218701.
- Reichardt, J. and Bornholdt, S. (2006a). Statistical mechanics of community detection. *Physical Review E*, 74:016110.
- Reichardt, J. and Bornholdt, S. (2006b). When are networks truly modular? *Physica D*, 224:20–26.
- Reichardt, J. and White, D. R. (2007). Role models for complex networks. E-print, arxiv.org, 0708.0958.
- Ress, D., Backus, B. T., and Heeger, D. J. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nat Neurosci*, 3(9):940–5.
- Ress, D. and Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nat Neurosci*, 6(4):414–20.
- Sakai, H. (1992). White-noise analysis in neurophysiology. *Physiological Reviews*.
- Schluppeck, D., Glimcher, P., and Heeger, D. J. (2005). Topographic organization for delayed saccades in human posterior parietal cortex. *Journal of Neurophysiology*, 94(2):1372–84.
- Selverston, A. I. and Moulins, M., editors (1987). *The Crustacean Stomatogastric System: A Model for the Study of Central Nervous Systems*. Springer-Verlag, Berlin.

- Shalizi, C. R. (2003). Optimal nonlinear prediction of random fields on networks. *Discrete Mathematics and Theoretical Computer Science*, AB(DMCS):11–30.
- Shalizi, C. R., Camperi, M. F., and Klinkner, K. L. (2007). Discovering functional communities in dynamical networks. In Airolidi, E., Blei, D. M., Fienberg, S. E., Goldenberg, A., Xing, E. P., and Zheng, A. X., editors, *Statistical Network Analysis: Models, Issues, and New Directions*, volume 4503 of *Lecture Notes in Computer Science*, pages 140–157, New York. Springer-Verlag.
- Shalizi, C. R. and Crutchfield, J. P. (2001). Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics*, 104:817–879.
- Shalizi, C. R., Haslinger, R., Rouquier, J.-B., Klinkner, K. L., and Moore, C. (2006). Automatic filters for the detection of coherent structure in spatiotemporal systems. *Physical Review E*, 73:036104.
- Shalizi, C. R. and Klinkner, K. L. (2004). Blind construction of optimal nonlinear recursive predictors for discrete sequences.
- Shalizi, C. R., Rinaldo, A., and Klinkner, K. L. (2008). Adaptive nonparametric prediction and bootstrapping of discrete time series. Manuscript in preparation.
- Sharma, J., Dragoi, V., Tenenbaum, J. B., Miller, E. K., and Sur, M. (2003). V1 neurons signal acquisition of an internal representation of stimulus location. *Science*, 300(5626):1758–1763.
- Siegel, R. M., Raffi, M., Phinney, R. E., Turner, J. A., and Jandó, G. (2003). Functional architecture of eye position gain fields in visual association cortex of behaving monkey. *Journal of Neurophysiology*, 90(2):1279–94.
- Silver, M. A., Ress, D., and Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of Neurophysiology*, 94(2):1358–71.
- Smith, A. M., Lewis, B. K., Ruttimann, U. E., Ye, F. Q., Sinnwell, T. M., Yang, Y., Duyn, J. H., and Frank, J. A. (1999). Investigation of low frequency drift in fmri signal. *Neuroimage*, 9(5):526–33.
- Snyder, L. (2000). Coordinate transformations for eye and arm movements in the brain. *Current Opinion in Neurobiology*.
- Snyder, L. H., Grieve, K. L., Brochie, P., and Andersen, R. A. (1998). Separate body- and world-referenced representations of visual space in parietal cortex. *Nature*, 394(6696):887–91.
- Sporns, O., Tononi, G., and Edelman, G. M. (2002). Theoretical neuroanatomy and the connectivity of the cerebral cortex. *Behav Brain Res*, 135(1-2):69–74.
- Teo, P. C., Sapiro, G., and Wandell, B. A. (1997). Creating connected representations of cortical gray matter for functional mri visualization. *IEEE transactions on medical imaging*, 16(6):852–63.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society B*, 58:267–288.
- Tootell, R. B. and Hadjikhani, N. (2001). Where is ‘dorsal v4’ in human visual cortex? retinotopic, topographic and functional evidence. *Cereb Cortex*, 11(4):298–311.
- Umeno, M. M. and Goldberg, M. E. (2001). Spatial processing in the monkey frontal eye field. ii. memory responses. *Journal of Neurophysiology*, 86(5):2344–52.
- Victor, J. D. (2000). Asymptotic bias in information estimates and the exponential (Bell) polynomials. *Neural Computation*, 12:2797–2804.
- Wahba, G. (1990). *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics, Philadelphia.
- Wandell, B. A., Chial, S., and Backus, B. T. (2000). Visualization and measurement of the cortical surface. *Journal of cognitive neuroscience*, 12(5):739–52.

- Wandell, B. A., Dumoulin, S. O., and Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, 56(2):366–83.
- Zenger-Landolt, B. and Heeger, D. J. (2003). Response suppression in v1 agrees with psychophysics of surround masking. *J Neurosci*, 23(17):6884–93.