Simultaneous Estimation of Response Fields and Impulse Response Functions

PhD Thesis Proposal

Erich Huang
Carnegie Mellon University - Department of Statistics

February 11, 2009

Abstract

Relevant issues in using functional magnetic resonance imaging (fMRI) experiments to characterize the visual system include estimating the response fields resulting from given stimuli, namely the spatial distribution of substantial activity in the brain, as well as their temporal properties and the gain fields, the changes in the response field that variables such as attention and eye position induce. But in addition to mapping the activity patterns given fMRI data and gathering information about the time courses of the response and detecting changes that result from modulation variables, we would also like our methods to be able to identify areas of substantial activity at the level of regions and estimate the time courses of the response as smooth functions of space and time. This would produce results that are ultimately more meaningful than statistical analysis methods that proceed at the level of artificial volume elements in space and arbitrary time points.

Specifically, we frame this problem as a simultaneous estimation of the response fields in their entirety and the functions that characterize their time courses. To this end, we introduce the preliminary developments of two techniques. The first is Penalized Regression, a least squares estimation subject to penalties on the parameters that induce solutions with spatial and temporal regularity, and the second is Region Growing, a technique that maps the areas of significant activity in the brain through a greedy algorithm that gives an approximate solution to the combinatorially difficult problem of searching through all possible regions for the ones that best fit the data. We apply both techniques to simulation studies and actual fMRI data collected under the White Noise setup. The simulation studies indicate these two techniques outperform Reverse Correlation, a standard technique in the literature, and application of these methods to the actual fMRI data produce results that are sensible from a neuroscientific point of view.

1 Introduction

Our aim is the development of novel statistical methodology for characterization of the visual system given data from functional magnetic resonance imaging (fMRI) experiments. fMRI is a technique to monitor the activity of a subject's brain that capitalizes on the differences in the magnetic signature of active and non-active regions. Typically, a visual fMRI experiment proceeds as follows: during the course of the experiment, a subject observes sequences of visual stimuli and concurrently, the activity in the subject's brain is measured via fMRI at regular time intervals. fMRI has some significant advantages as a method for tracking brain activity, including the fact that it enables very rapid acquisition of brain activity measurements; some setups allow measurements to be taken as frequently as every 0.1 seconds [1]; and that fMRI is non-invasive; exposure of the subject to potentially harmful chemicals or radiation is avoided [2].

Characterizing the visual system involves gathering information about the spatiotemporal properties of the brain's response to visual stimuli. In our case, this entails three main problems. The first is finding the response fields, the spatial distribution of substantial activity across the brain and the magnitude of this activity, and how they vary with time, resulting from a sequence of external stimuli occurring in a given location in visual space. The second is the inverse problem of finding the receptive field of a given area of the brain, namely where in visual space must stimuli occur in order to produce activity in that location. The third is determining the effects of modulation, in other words, how these response fields change as a function of variables such as eye position or attention. This characterization is likely to be very useful in a number of diverse problems, from the decoding of a subject's visual experience through brain activity; in other words, reconstructing the image that the subject is viewing from fMRI readings alone [3]; to clinical applications such as the early diagnosis and monitoring of central nervous system diseases affecting the visual cortex. Being able to detect such diseases before the onset of the more obvious systems through anomalies in fMRI images may be key in preventing substantial damage to the visual cortex. Also, with this characterization, fMRI may become a non-invasive way to monitor the progression of a disease or the efficacy of a certain treatment [4].

But in addition to being able to elicit information about the spatial and temporal properties of the visual response and the effects of modulation, we also aim to develop methodology that would identify the locations of substantial activity in the brain at the level of regions and estimate the time courses of the response at the level of functions of space and time. This latter objective would ultimately produce more meaningful results than what is ordinarily done; typically, in an fMRI data set, the spatial resolution is in terms of voxels, rectangular volumes measuring millimeters on each side, while activity is usually measured at regular time intervals, and estimation of the response fields and response time courses often proceed at the level of these artificial spatial volume elements and arbitrary points in time. Instead of looking at whether or not an individual voxel shows substantial activity to given visual stimuli, we aim to look for collections of contiguous voxels that do. Similarly, instead of estimating the underlying activity at a particular point in time, we seek to examine the evolution of the response as a smooth function of time and space. So specifically, our problem is a simultaneous estimation of the response fields resulting from each of the visual stimuli and the time course functions. The ideas involved in this problem are expected to not only be a valuable contribution to neuroscience but to statistics as well; estimating objects subject to spatial and temporal regularity conditions arises in a number of situations, even outside fMRI data analysis.

Initially, we focus primarily on response field estimation. To this end, we introduce two new techniques. The first, Penalized Regression, frames this problem as a parameter estimation under a given model with penalties on the parameters to incorporate the a priori spatial and temporal constraints. The second, Region Growing, is a greedy algorithm that seeks to find an approximate solution to the combinatorially

difficult search over all possible regions on the brain for the one that best fits the data. In simulation studies we find that both techniques outperform Reverse Correlation, a standard technique in the literature, in terms of detecting regions of activity. The results of applying both of these techniques to real fMRI data agree with one another, and make sense from a neuroscientific point of view.

Section 2 gives further details on the background of fMRI, including necessary terminology and description of the operation of fMRI, as well as descriptions of the experimental paradigm under which our data were collected and the model we will use. Section 3 describes work in the literature related to our problem. Section 4 describes Penalized Regression and Region Growing, and Section 5 illustrates the use of these techniques on simulation studies and real fMRI data. Finally, in Section 6, we discuss the next steps in our research and sketch some of our plans on how to proceed with this proposed work.

2 fMRI Background, Experimental Setups, and Model

2.1 fMRI Background and the White Noise Setup

First, we give a brief overview of background on fMRI, including the mechanism by which fMRI operates as well as some necessary terminology. fMRI capitalizes on the differences in the magnetic properties of oxygenated and deoxygenated hemoglobin to measure changes in activity. Oxygenated hemoglobin is diamagnetic whereas deoxygenated hemoglobin is paramagnetic, so they will exhibit different attractions to a weak external magnetic field. Activated regions of the brain require more oxygen than inactive regions, so an increase in activity would be marked by an influx of oxygenated hemoglobin. Therefore, the ratio of oxygenated to deoxygenated hemoglobin is an indirect measure of brain activity; changes in this ratio are marked by changes in the magnetic signature of the blood in the subject's brain, which fMRI measures [2, 4].

As mentioned before, an fMRI experiment for the purpose of visual system characterization proceeds as follows: the subject observes sequences of visual stimuli at different points in space while fMRI measurements are taken at regular time intervals. Define the space that the subject observes during this experiment as the visual field. One fMRI measurement is a three-dimensional image of the brain, partitioned into voxels, rectangular volume elements typically measuring several millimeters on each side. These images are arranged into a series of slices, two-dimensional images of different horizontal cross-sections of the brain. The magnetic signature of a particular voxel is encoded into the grayscale intensity of that voxel. A full fMRI data set will typically contain N such three-dimensional images. Note that for each voxel, we have N measurements of the magnetic signature and activity, each at a different point in time; thus, each voxel has a time series of its activity, or a hemodynamic response, associated with it.

These hemodynamic responses also include non-systematic changes such as breathing and involuntary movements. Thus, hemodynamic responses consist of two components: the signal, the variation due to the activity resulting from exposure to the stimuli, and the noise, the remainder of the response. The signal-to-noise ratio is defined as the ratio of the root mean square of the signal to the root mean square of the noise.

Analysis and visualization may, and often does, proceed in terms of the voxel coordinates; the location of a site of the brain is identified by its position on the original three-dimensional images that comprise the fMRI data. These voxel coordinates, however, do not necessarily reflect the true connectivity of the brain. For example, because of the fissures and folds that the brain contains, two voxels that are adjacent in voxel coordinates may be separated by a fissure and actually not be contiguous along the surface of the brain. Alternatively, we may perform our analyses and visualizations on a flat map, a projection of the surface of

the brain upon a two-dimensional map that does reflect the true connectivity of the brain; two sites that are close on this flat map are indeed close on the surface of the brain [5].

Again, define the response field as the spatial distribution across the brain of the activity resulting from particular stimuli. The receptive field of a voxel, as described in Section 1, is the spatial distribution across the visual field of where stimuli must occur in order to evoke activity in the neurons in that voxel. Finally, we use the term gain fields to describe the changes in the response field as a result of variables such as attention and eye position.

One experimental setup that is particularly useful for our purposes is the White Noise paradigm. Here, we partition the visual space into p rectangular areas, which we refer to as tiles. Each tile takes on one of two states at any point during the experiment: on or off. When the tile is off, it stays a uniform gray color, and when the tile is on, it flickers between pure white and pure black. Each tile updates its status at regular time intervals according to a deterministic binary sequence called an M-sequence that appears completely random to the observer.

Because the subject will not be likely to identify any patterns in these stimuli, confounding variables such as attention are controlled for, thus making this setup appropriate for gathering information about modulation effects. Additionally, pseudo-random stimuli like M-sequences are also suitable for extracting details about spatial and temporal interactions between stimulus components [6].

And not only does the deterministic nature of the M-sequences allow for greater control over the stimuli for analysis purposes, their characteristics make the generation of p sequences of stimuli easy. A circular shift of one M-sequence by some positive integer Δ is another M-sequence; typically, we would set the stimulus sequence in the tile in the upper left corner equal to the original M-sequence, and set the stimuli in the other tiles to be circular shifts of the original M-sequence. Thus, we generate sequences of stimuli for all p tiles using only one M-sequence [7].

2.2 Modeling the Hemodynamic Response in Terms of Presented Stimuli

Define $R_v(t)$ to be the measured hemodynamic response at time t in voxel v and $s_j(t)$ to be the stimulus time series corresponding to tile j. The status of each tile is updated at times $t = \delta, 2\delta, \ldots, N\delta$ for some time interval δ. Set

$$s_j(t) = \begin{cases} -1 & \text{if tile } j \text{ is off at time } t \\ 1 & \text{if tile } j \text{ is on at time } t \end{cases}$$

The literature gives evidence of the mechanism by which neurons respond to stimuli as shift-invariant linear systems [1], meaning that (1) scaling the intensity of the stimuli should also scale the magnitude of the neuron's response by the same factor, (2) if the subject is exposed to two sequences of stimuli simultaneously, then the resulting response should be the sum of the responses we would have observed under exposure to each of the sequences separately, and (3) the time course resulting from a given sequence of stimuli should not change regardless of when the sequence began. Neurons, like any other linear system, are characterized by an impulse response function (IRF); an IRF is defined as the expected time course of a neuron's response to a brief stimulus that activates that neuron. However, in practice, dealing with this IRF at the level of individual neurons is not possible because of the resolution of our data; instead, we handle the IRFs in our analysis at the voxel level, and the IRF for a particular voxel is assumed to be approximately an average of the IRFs for all neurons within that voxel. These IRFs may vary across voxels, but they are expected to be smooth functions as a function of time with the same basic form; after about 5 seconds, they reach a maximum, and then return to baseline by about 16 seconds, while undershooting slightly beforehand [8].

Define $a_v(t)$ as the impulse response function for voxel v, and $\beta_j(v)s_j(t)$ as the effective sequence of stimuli from tile j for voxel v. $\beta_j(v)$ is a non-negative coefficient to correct for the fact that the neurons in voxel v will not respond the same way to stimuli in different tiles. Stimuli in some tiles may activate voxel v more strongly than others, and stimuli in some tiles will not activate voxel v at all. Naturally, $\beta_j(v)$ would be large if voxel v responds strongly to stimuli in tile j and zero if this response is nonexistent. Then it can be shown that the response in voxel v resulting from exposure to tile j is now a convolution between the impulse response function and the effective sequence of stimuli:

$$\begin{aligned}[a_v * \beta_j(v)s_j](t) &= \beta_j(v)[a_v * s_j](t) \\ &= \beta_j(v)\sum_u a_v(u)s_j(t-u)\end{aligned}$$

Note that under this model, $a_v(t)$ and $\beta_j(v)$ are not identifiable without further restrictions on either parameter. To circumvent this, we restrict $\int [a_v(t)]^2 dt = 1$.

Again, since each voxel is approximately a linear system, the response in voxel v resulting from exposure to all p tiles is a sum of the responses we would have observed had we exposed the subject to each tile individually. Since the observed hemodynamic response includes non-neural changes, we add a zero-mean noise time series $\epsilon_v(t)$. Also, these time series typically have "scanner" drift, an approximately linear trend $b_1 + b_2 t$ [9]. The hemodynamic response time series in voxel v is therefore

$$R_v(t) = b_1 + b_2 t + \sum_{j=1}^{p} \beta_j(v)\left[a_v * s_j\right](t) + \epsilon_v(t)$$

Prior to analysis, we typically remove the trend by either directly fitting a line to the time series or applying a Fourier transform to the time series and then subtracting off a low-frequency component [9] and then centering and scaling $R_v(t)$ and $s_j(t)$. This pre-processed form of the data is the one we use in this thesis proposal; we can implement straightforward changes to our techniques to simultaneously estimate the drift. The model then becomes

$$R_v(t) = \sum_{j=1}^{p} \beta_j(v)\left[a_v * s_j\right](t) + \epsilon_v(t) \tag{1}$$

To estimate the response field that results from the stimuli in tile j, we estimate $\beta_j(v)$ while estimation of $a_v(t)$ will give insight on how these response fields vary with time.

3 Literature Review

Two of the most commonly used techniques in the literature for similar problems are multiple linear regression techniques [10] and Reverse Correlation [11]. In the multiple linear regression approaches, we would assume model (1) but fix $a_v(t)$ to equal some ideal or mean impulse response function. Note that given $a_v(t)$, model (1) is a linear model; for each voxel v, we would estimate $\beta_j(v)$ through standard multiple regression techniques. We would then be able to interpret the response fields from our estimates of $\beta_j(v)$.

Reverse Correlation is another frequently used procedure that conducts, for each voxel v, p simultaneous hypothesis tests of the null hypothesis of zero cross-correlation at all time points t between the

hemodynamic response $R_v(t)$ and the stimulus sequence in each tile $s_j(t)$. Reverse Correlation has an advantage over multiple linear regression in that it does not require the convolution of $s_j(t)$ with an ideal or mean impulse response function, which makes it better suited for instances where the temporal aspect is important and unknown [11].

Some other methods involve the deconvolution of the time series $\beta_j(v)a_v(t)$ from the hemodynamic response under model (1). Glover, for example, performs a similar deconvolution by finding the least squares parameter estimates [12]. This procedure is often carried out in Fourier space, where the convolution operator becomes equivalent to point-wise multiplication; here, we would apply a Fourier transform to the hemodynamic response and sequences of stimuli, perform the deconvolution, and then apply an inverse Fourier transform to the impulse response functions to convert them back to the original coordinates.

However, these techniques traditionally proceed voxel-wise and without any further constraints on the spatial or temporal relationship between parameter estimates. Further techniques to account the spatial regularity of the response fields have also appeared. The problem of spatial regularity appears not only in our problem, but in many fMRI-related problems as well. The literature does contain some techniques to incorporate this spatial structure; many of them have been developed in the context of other fMRI problems, but may also prove very useful for our purposes as well. For example, Worsley et al suggest smoothing the hemodynamic response time series spatially with a kernel prior to analysis [13]. A related technique to incorporate spatial structure into the estimates is to perform a voxel-wise estimation of the parameters, but instead of using the response time series for that voxel, we use an average of the time series of that voxel and all voxels within a fixed distance [14].

Recently, image segmentation techniques [15], which partition a field of voxels into contiguous regions based on their intensities, have also been proposed as a technique to detect activation at the level of clusters rather than individual voxels. Various forms of these image segmentation techniques are applicable to this problem. One version involves partitioning the brain according to a summary statistic of the responses of each of the voxels in each region, and then based on these summary statistics, declaring each region as activated or not activated. Lu et al provide one notable example of this [16]; they propose finding the partitions of the brain such that the mean response of each voxel in each regions differ as little as possible, and then, given these mean responses, search for the regions of activity. Yet another image segmentation technique involves separating the field of voxels into "foreground" and "background", namely regions of substantial activity and regions of no activity respectively, through seeded region growing [17]. The idea here is to start with a seed, an initial cluster of foreground pixels, and iteratively expand the foreground region by absorbing neighboring background voxels until some stopping criterion. Our Region Growing procedure, which we describe in detail in Section 4, is based on this concept.

Work in the literature on estimating receptive fields and investigating modulation effects include Kay et al [3], who use fMRI readings to predict which image, out of a set of over one hundred, a subject was looking at. The methodology they present is likely to be widely useful in visual system characterization, though their purposes are primarily predictive whereas ours are estimative. They estimate the entire receptive field of each voxel in the brain through a Gabor wavelet pyramid model, which they later use for predicting the visual experience of the subject. By performing constrained versions of the Gabor wavelet pyramid modeling, they also investigate the effects of variables such as orientation on predictive performance.

4 Methodology

4.1 Penalized Regression

This approach is a simultaneous estimation of $\beta_j(v)$ and $a_v(t)$ under model (1) through a convex minimization problem. To incorporate the expected spatial and temporal characteristics of $\beta_j(v)$ and $a_v(t)$ into our estimates and at the same time substantially reduce the degrees of freedom of our problem, we perform a least squares minimization subject to appropriate penalty terms $P(\beta_j(v), a_v(t))$:

$$\frac{1}{2}\sum_v \sum_t \left[R_v(t) - \sum_{j=1}^{p} \beta_j(v)[a_v * s_j](t) \right]^2 + P(\beta_j(v), a_v(t))$$

To encourage estimates of $\beta_j(v)$ with spatial regularity, we include the penalty term $\mu \sum_{j=1}^{p} \int ||\nabla \beta_j(v)||^2 dv$, where $\nabla \beta_j(v)$ is the gradient of $\beta_j(v)$ with respect to location v. Penalizing first and second derivatives to enforce spatial smoothness was an idea proposed by Marx and Eilers [18]. To encourage estimates of $a_v(t)$ with spatial and temporal regularity, we can include the similar penalty term $\nu \int \int ||\nabla a_v(t)||^2 dv dt$. Finally, to encourage regularity across time in $a_v(t)$, we include the penalty term $\gamma \sum_v \int \left[\frac{d^2}{dt^2} a_v(t) \right]^2 dt$, a penalty that is similar to the usual second-derivative penalty for smoothing splines [22]. Note we may also use a similar penalty to enforce spatial smoothness, and we will consider this as a possible alternative.

Even with these regularity constraints, the estimation problem is still very high-dimensional. For a typical fMRI data set, we can have hundreds, perhaps thousands, of voxels to deal with, and a number of time points on the order of five hundred to one thousand; this translates into tens of thousands of parameters we need to estimate. To reduce the number of parameters we need to estimate, we can capitalize on the form of $a_v(t)$ and the inherent sparsity of $\beta_j(v)$. Recall the basic form of $a_v(t)$ as described in Section 2.2: $a_v(t)$ rises to a peak at around 5 seconds and then returns to baseline after about 16 to 20 seconds, while undershooting slightly beforehand. This means for t greater than about 16 seconds, $a_v(t)$ should be zero; thus, we would only need to estimate $a_v(t)$ for a small number of time points. Meanwhile, we expect $\beta_j(v)$ to be very sparse, meaning $\beta_j(v) = 0$ for most v and j, because only a few regions of the brain should show substantial activation in response to any given stimuli.

To incorporate the sparsity of $\beta_j(v)$, we include the penalty $\lambda \sum_v \sum_{j=1}^{p} |\beta_j(v)|$. Because of its geometry, this penalty will cause many $\beta_j(v)$ to be set to zero. This is the idea behind the LASSO [19] and variants such as the Relaxed LASSO [20] and the Fused LASSO [21]. However, recall from Section 2.2 that we restrict $\beta_j(v)$ to be non-negative; thus, this penalty becomes $\lambda \sum_v \sum_{j=1}^{p} \beta_j(v)$.

Therefore, the objective function to minimize is

$$\frac{1}{2}\sum_v \sum_t \left[R_v(t) - \sum_{j=1}^{p} \beta_j(v)[a_v * s_j](t) \right]^2 + \lambda \sum_v \sum_{j=1}^{p} \beta_j(v) + \mu \sum_{j=1}^{p} \int ||\nabla \beta_j(v)]||^2 dv$$

$$+ \ \nu \int \int ||\nabla a_v(t)||^2 dv dt + \gamma \sum_v \int \left[\frac{d^2}{dt^2} a_v(t) \right]^2 dt \qquad (2)$$

subject to $\beta_j(v) \geq 0$.

Note that simultaneous estimation of these parameters is a nonlinear least squares problem. But we can decompose this into two linear problems; we may rewrite model (1) as

$$\begin{aligned} R_v(t) &= \sum_{j=1}^{p} \beta_j(v) \left[\sum_u a_v(u) s_j(t-u) \right] + \epsilon_v(t) \\ &= \sum_u a_v(u) \left[\sum_{j=1}^{p} \beta_j(v) s_j(t-u) \right] + \epsilon_v(t) \end{aligned}$$

This suggests using Repeated Alternating Conditional Estimation [23], where we alternate between estimating $\beta_j(v)$ while fixing $a_v(t)$ and estimating $a_v(t)$ while fixing $\beta_j(v)$ until some stopping criterion. Equivalently, we alternate between minimizing the following two objective functions with respect to $\beta_j(v)$ and $a_v(t)$ respectively instead of directly minimizing (2):

$$\frac{1}{2} \sum_v \sum_t \left[R_v(t) - \sum_{j=1}^{p} \beta_j(v) [a_v * s_j](t) \right]^2 + \lambda \sum_v \sum_{j=1}^{p} \beta_j(v) + \mu \sum_{j=1}^{p} \int ||\nabla \beta_j(v)||^2 dv$$

$$\text{subject to } \beta_j(v) \geq 0 \tag{3}$$

$$\frac{1}{2} \sum_v \sum_t \left[R_v(t) - \sum_{j=1}^{p} \beta_j(v) [a_v * s_j](t) \right]^2 + \nu \int \int ||\nabla a_v(t)||^2 dv dt + \gamma \sum_v \int \left[\frac{d^2}{dt^2} a_v(t) \right]^2 dt \tag{4}$$

One significant advantage of rewriting the minimization this way is that we reframe it as two far easier, more familiar problems. If we approximate the derivatives in these equations as differences, which is a logical way to proceed given both the time points and locations in space, the voxel locations, are discretely sampled, minimizing (3) with respect to $\beta_j(v)$ is a quadratic programming problem [24] whereas minimizing (4) with respect to $a_v(t)$ is a minimization of a quadratic form, which has a closed-form solution. Additionally, all three objective functions; the full one, (2), and the two partial ones, (3) and (4); are convex, and so we expect this algorithm to converge to the global minimum objective value with respect to $\beta_j(v)$ and $a_v(t)$.

Our procedure is then as follows:

1. Begin with initial guesses for $a_v(t)$ and $\beta_j(v)$.
2. Fix $a_v(t)$ and update $\beta_j(v)$ via quadratic programming.
3. Fix $\beta_j(v)$ and minimize (4) with respect to $a_v(t)$.
4. Renormalize $a_v(t)$ so that $\int [a_v(t)]^2 dt = 1$.
5. Cycle through steps (2) through (4) until some stopping criterion, e. g. the norms of $\beta_j(v)$ or $a_v(t)$ between successive iterations differ by less than a prescribed ϵ.

Note that this procedure requires a choice of three parameters, λ, μ, and γ. We make a selection for the appropriate parameters at each stage of the procedure; in other words, every time we estimate $\beta_j(v)$ given $a_v(t)$ we search through different values of λ and μ and take the estimates produced by the optimal parameter values as our updated $\beta_j(v)$, and we proceed similarly for estimating $a_v(t)$ given $\beta_j(v)$. Another

advantage of this technique is that it allows the use of K-fold cross-validation [25] to select the values of these parameters, which is a logical starting point. Here, we would divide the data into K portions according to time index, and as usual, cycle through the following procedure for $k = 1, \ldots, K$: use all but the k^{th} portion for fitting, and predict using the k^{th} portion, thus obtaining an estimate of the mean squared prediction error.

4.2 Region Growing

The primary aim of Region Growing is the estimation of the support of $\beta_j(v)$ given $a_v(t)$ and the data, namely the locations on the response field where $\beta_j(v) \neq 0$. Our objective here is to find the regions on the response field, out of all possible regions, that best fit the data. However, this problem is combinatorially difficult, and so, we find an approximate solution through a greedy search algorithm.

The feasibility of this greedy search algorithm relies on two assumptions regarding the nature of the response fields. First, substantial activity, in other words, where $\beta_j(v) \neq 0$, occurs in contiguous, sizable clusters. Second, the number of such clusters is very small; for now, we assume that the response field only has one, though we are in the process of developing possible ways to allow multiple regions of substantial activity.

This technique is based on the ideas of seeded region growing [15] and coordinate descent [21], which involves iteratively cycling through the process of estimating one parameter or subset of parameters given the latest values of the others. We assume model (1); we fix $a_v(t)$ to be some mean or ideal impulse response function and then we cycle through the estimation of the response fields for each tile, namely estimating $\beta_j(v)$ simultaneously at all locations v on the response field for each j, until some stopping criterion such as the norms of successive estimates of each response field differing by less than some prescribed threshold ϵ. For each response field, we begin by finding the seed, an initial cluster of adjacent voxels where $\beta_j(v) \neq 0$, which is usually the location where the magnitudes of $\beta_j(v)$ are highest. Then, we iteratively add neighboring voxels to this "nonzero region" according to a given convex objective function; specifically, we propose the inclusion of a neighboring voxel into the nonzero region and accept this move if it results in a decrease of the objective function by more than a certain threshold λ; until decreases in this function greater than λ no longer occur.

As we add voxels to the nonzero region, we also adjust the magnitude of all $\beta_j(v)$ for v in this region according to the same objective function. Currently, Region Growing sets the $\beta_j(v)$ for the voxels in the nonzero region to all equal their mean $\beta_j(v)$, but we are implementing modifications to the technique to allow for more precise estimates. Consequently, the estimates of the magnitudes of $\beta_j(v)$ will be biased for these voxels.

Thus, our basic procedure is as follows. Given $a_v(t)$ and initial guesses for the response fields $\beta_j(v)$, we cycle through the following for $j = 1, \ldots, p$ until the stopping criterion is met:

1. Re-initialize $\beta_j(v) = 0$ for all v.

2. Search for the seed, which becomes the initial nonzero region.

3. Repeat the following until no more expansions of the nonzero region result in a decrease in the objective function of more than a prescribed threshold λ: propose the absorption of a voxel immediately adjacent to the current nonzero region, adjust $\beta_j(v)$ for all voxels in the proposed new nonzero region, and calculate the resulting change in the objective function. Accept this proposed move if it results in a decrease in the objective by more than λ.

Execution of the Technique

Assume $R_v(t)$ have been centered and scaled so that they have mean zero and standard deviation one. Define our objective function as

$$\sum_v \sum_t \left[R_v(t) - \sum_{j=1}^{p} \beta_j(v)[a_v * s_j](t) \right]^2 \tag{5}$$

where $\beta_j(v) = \beta_j I\{v \in \mathcal{A}_j\}$ for some $\beta_j \geq 0$ and some sets of contiguous voxels $\mathcal{A}_j$. Recall that we accept or reject each expansion of the nonzero region during the region growing procedure based on the resulting decrease of this objective function; this procedure is similar to performing a search for the minimum of a sums of squared residuals objective function subject to the constraints that all $\beta_j(v)$ either equal β_j or zero, and that nonzero $\beta_j(v)$ occur in contiguous regions.

The process begins by resetting the response field to equal zero everywhere, specifically $\beta_j(v) = 0$ for all v. The seed search then proceeds as follows. For each voxel v we propose the following move: form a cluster $\mathcal{V}$ including v and the immediately neighboring voxels and set $\beta_j(v)$ for these voxels equal to

$$\frac{\sum_{\mathcal{V}} \sum_t \left[R_w(t) - \sum_{k \neq j} \beta_k(w)(a_w * s_k)(t) \right] (a_w * s_j)(t)}{\sum_{\mathcal{V}} \sum_t \left[(a_w * s_j)(t) \right]^2}$$

This value of $\beta_j(v)$ minimizes the objective function (5) given $a_v(t)$, $\beta_k(v)$ for $k \neq j$, and that $\beta_j(v) \neq 0$ only for $v \in \mathcal{V}$. We calculate the decrease in the objective function resulting from each proposed move, accepting the one that results in the largest decrease.

Indeed, we may use a single voxel as the seed rather than a cluster as above. However, selection of a false single voxel seed, namely selecting a voxel v where $\beta_j(v) = 0$ as the seed, becomes more and more of a problem as the signal-to-noise ratio decreases. The problem of selecting a false cluster seed, on the other hand, is far lower.

Now suppose we currently have a response field where $\beta_j(v) = 0$ for all v except on a region $\hat{\mathcal{A}}_j$. We then select a voxel v_{new} immediately neighboring $\hat{\mathcal{A}}_j$ and propose the following moves: we absorb v_{new} into $\hat{\mathcal{A}}_j$ and set $\beta_j(v) = \hat{\beta}$ for all voxels $v \in \hat{\mathcal{A}}_j \cup \{v_{new}\}$, where $\hat{\beta}$ is the value of $\beta_j(v)$ for all $v \in \hat{\mathcal{A}}_j \cup \{v_{new}\}$ that results in the largest decrease in the objective function, given that $\beta_j(v) \neq 0$ only for $v \in \hat{\mathcal{A}}_j \cup \{v_{new}\}$:

$$\hat{\beta} = \frac{\sum_{v \in \hat{\mathcal{A}}_j \cup \{v_{new}\}} \sum_t \left[R_v(t) - \sum_{k \neq j} \beta_k(v)(a_v * s_k)(t) \right] (a_v * s_j)(t)}{\sum_{v \in \hat{\mathcal{A}}_j \cup \{v_{new}\}} \sum_t \left[(a_v * s_j)(t) \right]^2}$$

We accept this move if it decreases the objective function. We repeat this procedure until no further expansions of the nonzero region will decrease the objective function.

5 Feasibility of Our Methodology: Simulation Studies, Application to fMRI Data, and Comparison with Standard Techniques

5.1 Simulation Studies

We simulate data according to model (1) with characteristics similar to those we expect of actual fMRI data, use Penalized Regression and Seeded Region Growing to detect regions of activity, and compare their

performances with that of Reverse Correlation. Here, $s_j(t)$, for $j = 1, 2, \ldots, 10$, consist of circular shifts of the same M-sequence of length 511; we set $s_1(t)$ to be the original M-sequence and for $j = 2, 3, \ldots, 10$, set $s_j(t)$ to be a circular shift of $s_{j-1}(t)$ by 10.

For every voxel v in a 25×25 array, we set the impulse response function as

$$a_v(t) = \gamma(t; 7, 1) - 0.22\gamma(t; 14, 1)$$

where $\gamma(t; \alpha, \beta)$ is a $\Gamma(\alpha, \beta)$ density.

For each j, we generate response fields $\beta_j(v)$ that are zero everywhere except on a few contiguous regions. We accomplish this by setting $\beta_j(v)$ equal to a sum of biweight functions [26], i. e.,

$$\beta_j(v) = \sum \left[\alpha_j \left(1 - (v - c_j)^T \Omega_j^{-1} (v - c_j) \right)^2 + \kappa_j \right] I \left\{ (v - c_j)^T \Omega_j^{-1} (v - c_j) \leq 1 \right\}$$

where α_j and κ_j are scalar coefficients, Ω_j are 2×2 positive-definite matrices, and $c_j \in [1, 25] \times [1, 25]$. Each term in the above sum by itself will produce an elliptical nonzero region where the magnitude of $\beta_j(v)$ is highest in the centroid and decays smoothly toward the edges. Figure 1 depicts the impulse response function as well as a few examples of these response fields.

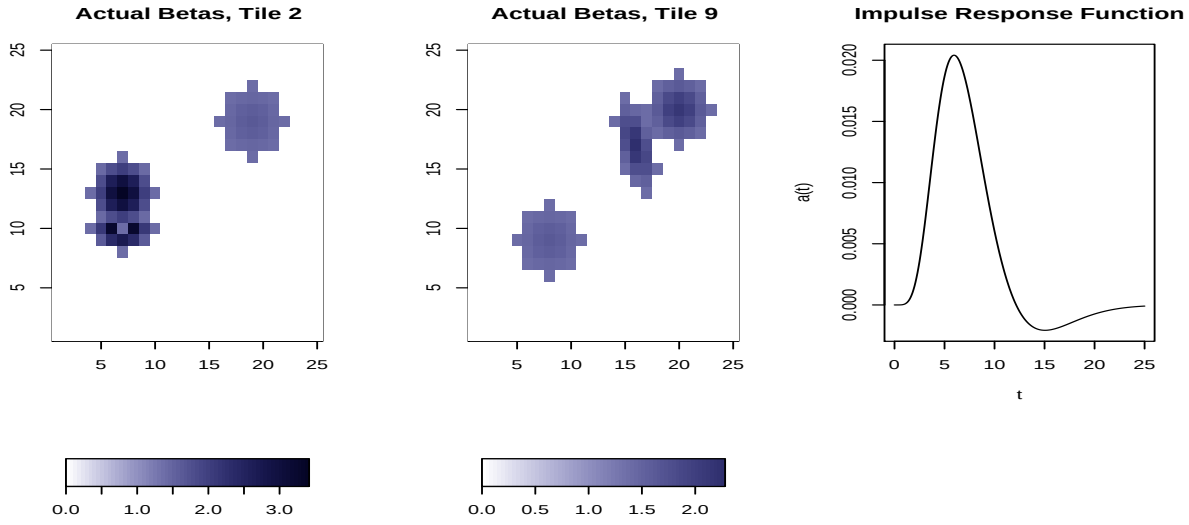

Figure 1: Left: simulated response field for $j = 2$; center: simulated response field for $j = 9$; right: impulse response function used for all voxels

We generate hemodynamic responses $R_v(t)$ for $t = 0, 1, \ldots, 511$ in each voxel v according to model (1). For simplicity, for each v we set each of the error time series $\epsilon_v(t)$ to be uncorrelated with one another and to be independent draws from a $N(0, 15^2)$ distribution.

We repeat the following procedure 50 times: we generate synthetic data according to the model above, and then we apply Penalized Regression, Region Growing, and Reverse Correlation to the simulated data. For the Seeded Region Growing procedure, we fix $a_v(t)$ to be the "canonical" impulse response function $\gamma(t; 6, 1) = \frac{1}{6}\gamma(t; 16, 1)$ [27] for all v. We thus obtain estimates for the activity detection power as well as for the squared bias and variance of the Penalized Regression estimates of $\beta_j(v)$.

Figure 2 shows the squared bias and variance of the Penalized Regression estimates of $\beta_j(v)$ for $j = 2$ and $j = 9$. The bias of the $\beta_j(v)$ estimates is low for locations where the true $\beta_j(v)$ is larger. The larger bias estimates occurs mostly along the borders of the nonzero regions, where the $\beta_j(v)$ are smaller. Smaller $\beta_j(v)$ will produce smaller signals, and recall that the variance of the noise is the same for all v; thus, the increased bias for these voxels is likely to be a result of reduced signal-to-noise ratio.

Figure 3 shows the detection proportion for $j = 2$ and $j = 9$ for each of the three methods, namely the proportion of times each method detected activity, or estimated $\beta_j(v) \neq 0$, in each location of the response fields for $j = 2$ and $j = 9$. The detection proportions for Penalized Regression indicate that this method had nearly perfect sensitivity; where $\beta_j(v) \neq 0$ truly, the Penalized Regression $\beta_j(v)$ estimates were almost always nonzero as well. Meanwhile, the sensitivity of Region Growing was lower than that of Penalized Regression, but its specificity was much higher. The fact that Region Growing can currently feasibly only detect one contiguous nonzero region contributed to this decrease in sensitivity; however, it seemed to consistently detect the region with the strongest activation, namely the largest $\beta_j(v)$, as expected. In this region, though, the proportion of times Region Growing detects this region is in general high; the detection proportions are above 0.6 in this region, and are above 0.8 in the centroids where $\beta_j(v)$ is the highest. But in locations where $\beta_j(v) = 0$, Region Growing almost never gave nonzero estimates of $\beta_j(v)$, whereas Penalized Regression sometimes estimated $\beta_j(v) \neq 0$ for these locations between 30% to 45% of the time. Both of these techniques outperformed Reverse Correlation in this simulation study. The detection proportions for Reverse Correlation were much lower in the nonzero regions; for $j = 2$, Reverse Correlation estimated $\beta_j(v) = 0$ about 60% of the time near the centroids and less than 30% everywhere else, and for $j = 9$, where the actual magnitudes of the $\beta_j(v)$ were lower, this proportion dropped to less than 25%. The specificity of Reverse Correlation was indeed higher than that of Penalized Regression; the proportion of times Reverse Correlation detected activity was in general between zero and 10% in locations where the actual $\beta_j(v)$ was zero, but this specificity was still lower than that of Region Growing.

There were a few cases in this simulation study where the detection power of Seeded Region Growing seemed relatively low, however. The results for $j = 5$ was the most notable example of this (see Figure 4); even though the specificity of this procedure is high, it seems that this technique successfully detected one of the regions half the time and the other region the other half of the time. This is likely to have been the result of the magnitudes of the $\beta_j(v)$ in both nonzero regions being similar. Modifying this technique to be able to detect activations in multiple regions is likely to remedy this problem.

5.2 fMRI Data

We apply the Penalized Regression and Seeded Region Growing techniques to data from the Center of Neural Sciences at New York University, collected under the White Noise paradigm. The visual field was decomposed into an 6 by 8 array of tiles. $s_1(t)$ corresponds to the status of the tile in the upper left corner of the array, $s_{48}(t)$ corresponds to the status of the tile in the lower right corner, and indexing proceeds downward and then to the right. Again, if the tile was on, it flickered rapidly between black and white, and if it was off, it stayed isoluminant gray in color. Each tile updated its status every 1.6 seconds according to circular shifts of an M-sequence of length 511. Hemodynamic responses were measured via fMRI every 0.8 seconds; these images consisted of 14 slices of 64 by 64 arrays of voxels. Given this data, we use Penalized Regression and Region Growing to detect the regions of the primary visual cortex (V1) that respond to the stimuli in each tile.

The spatial patterns of the activity in V1 are related to the relative position of the stimulus in the visual field [28]. Tiles 1 through 24 lie in the left side of the visual space whereas tiles 25 through 48 lie in the right; as a result, stimuli in tiles 1 through 24 should evoke activity in the right side of V1 whereas stimuli

in tiles 25 through 48 should produce a response in the left. Within each side of V1, the position of the activation also corresponds to the location of the stimulus; tiles toward the top of the array, e. g. tiles 1, 8, and 19, should produce activity primarily toward the dorsal side of V1 whereas tiles toward the bottom of the array, e. g. tiles 6, 11, and 24, should produce activity primarily toward the ventral side.

Figures 5 and 6 are plots of the response fields as estimated by Penalized Regression and Region Growing corresponding to tiles 20, 23, 26, and 29 on flat maps. Colored regions on these plots represent locations where these two techniques detected substantial activity, and these results are indeed consistent with what we expect. Activation from tiles 20 and 23 are primarily in right V1, which is sensible given these two tiles lie on the left hand side of the tile array. Meanwhile, activation from tiles 26 and 29 are primarily in left V1, which is sensible given these two tiles lie on the right hand side of the tile array. The activation from tiles 20 and 26 are on the dorsal side of V1 whereas those from tiles 23 and 29 are on the ventral side of V1. These results are expected, since tiles 20 and 26 lie toward the top of the array and tiles 23 and 29 lie toward the bottom.

The regions of substantial activation that Penalized Regression detected were in general larger than those Region Growing detected. This is likely to be due to the fact that the detection power of Penalized Regression is slightly higher and that Penalized Regression also tends to make spurious detections. Also, Penalized Regression seemed to only detect one substantial region of activity; the other smaller regions not contiguous to this region seem to be either an artifact of noise or the fact we applied our analyses in the original voxel coordinates; this indicates that the number of active regions in V1 resulting from the stimuli is indeed very small and most likely equal to one.

6 Proposed Research and Timeline

The next steps of our research consist of four main tasks:

1. Refine our current approaches, including regularization parameter selection for Penalized Regression, incorporating Region Growing into the alternating conditional estimation procedure, and changes to the techniques that make them suitable for analysis on a flat map of the brain.

2. Application of our new methods to simulation studies and fMRI data sets, including the one from Section 5.2.

3. Address the modulation aspect of the problem. This entails investigating how the response fields vary as a function of other variables such as orientation and attention, which requires methods to calculate the uncertainties in our estimates of our parameters. After developing techniques for this, we will apply them to appropriate simulation studies and actual fMRI data.

4. Derive theoretical results for our methodology wherever possible. This includes convergence and mean squared error of the Penalized Regression estimates, and convergence and sparsistency of the Region Growing technique.

We are currently implementing possible approaches to point (1) above, and expect to make substantial progress on this over the next two months. By May 2009 we aim to have results for point (3), and we hope to complete point (4) by September 2009.

6.1 Refining Our Current Methodology: Ideas on How to Proceed

Initially, the Region Growing procedure was intended as a way to estimate $\beta_j(v)$ given $a_v(t)$ in the alternating scheme described in Section 4. This procedure was designed specifically for a particular type of spatial distribution of the $\beta_j(v)$: for when $\beta_j(v)$, for each j, was equal to zero except on a few, preferably only one, sizable contiguous regions in space. We have reason to believe this is indeed true of the response fields.

The Region Growing procedure, as it stands now, is primarily useful for estimating the support of $\beta_j(v)$ rather than the values of $\beta_j(v)$ themselves. Recall that we are restricting the $\beta_j(v)$ to be piecewise constant surfaces here and assuming the $\beta_j(v)$ are nonzero on only one region. We had in fact run simulation studies where the $\beta_j(v)$ did actually have these characteristics; the Region Growing procedure within the alternating estimation scheme outperformed the other techniques in terms of bias, variance, detection specificity, and detection power. But because of these assumptions, which in general may not be true for the actual data, the current development of this technique may not be reliable in the alternating scheme, since the $\beta_j(v)$ estimates are biased, which may in turn cause divergence of the entire algorithm. Thus, incorporation of this procedure into the alternating scheme requires appropriate modifications, namely a way to allow for smooth variation of these $\beta_j(v)$ in these nonzero regions and a way to allow for multiple nonzero regions. One of our objectives here is to develop such modifications.

Allowing for variation of the $\beta_j(v)$ within the nonzero region requires a choice of a different objective function. If we use the same objective function (5) without the piecewise constant restriction, the estimate for $\beta_j(v)$ we will get will be overfit and have support extending across the entire field. To allow for smoother variation across voxels, we use a very similar objective function as the one used in Penalized Regression. We include a term to penalize the squared norm of the gradient of $\beta_j(v)$ with respect to v to incorporate smoothness in the response field across voxels. Then, to prevent the support of $\beta_j(v)$ from extending across the entire field, we add a λ term; what this will do is reject a proposed expansion of the region if the sum of the squared partial residuals and the squared norm penalty term decreases by less than λ. Thus, our objective function is

$$\sum_t \sum_v \left(R_v(t) - \sum_{k \neq j} \beta_k(v)[a_v * s_k](t) - \beta_j(v)[a_v * s_j](t) \right)^2 + \gamma \sum_{j=0}^{p-1} \int ||\nabla \beta_j(v)||^2 dv + \lambda$$

Suppose the response field $\beta_j(v)$ is zero on all except for K_j contiguous regions, and the regions are mutually non-overlapping. Then we may write the entire response field as a sum of K_j response fields $\beta_j^{(k)}(v)$ that are zero everywhere except on exactly one region; the nonzero region for each of these response fields is equal to a different nonzero region on $\beta_j(v)$. In other words,

$$\beta_j(v) = \sum_{k=1}^{K_j} \beta_j^{(k)}(v)$$

Then we may rewrite model (1) as

$$R_v(t) = \sum_{j=0}^{p-1} \sum_{k=1}^{K_j} \beta_j^{(k)}(v) \left[a_v * s_j\right](t) + \epsilon_v(t)$$

and given the K_j, the idea behind the estimation under this form of the model should be the same.

Initial developments and simulation studies indicate that using Region Growing to estimate $\beta_j(v)$ given $a_v(t)$ in the alternating scheme has the potential to perform very well.

The most significant modification that Penalized Regression requires involves the choice of the regularization parameters γ and λ. Our current implementation chooses these parameters by K-fold cross-validation. For predictive purposes, this leads to an excellent choice of γ and λ, but it has been shown that an L_1-norm penalty on the $\beta_j(v)$ is not consistent in selecting which of these coefficients are nonzero [20]. However, it turns out when we use prediction error as a criterion to select regularization parameters for an L_1-norm penalty, we tend to select spurious variables in addition to those that are truly relevant [20]. Thus, for voxels where $\beta_j(v)$ is actually nonzero, we will tend to obtain nonzero estimates for these $\beta_j(v)$; this can be seen in the simulation studies detection power plots in Section 5.1. We can then screen out the spurious nonzero $\beta_j(v)$ after the execution of Penalized Regression; we look for voxels where $\beta_j(v)$ is nonzero but small in magnitude, or small nonzero clusters, and set $\beta_j(v) = 0$ at these locations.

Finally, in Section 5.2, we applied our techniques directly to the data on the original voxel coordinates. Though this may be a good starting point, as mentioned before, the spatial constraints in either of these techniques may not reflect the true connectivity of the brain. To get improved results, therefore, we plan to perform the analysis on the flat map. However, on the flat map, the location coordinates no longer lie on a grid like they do in the voxel coordinates; this will affect the spatial penalties and the execution of the Region Growing.

6.2 Identifying Changes in the Response Induced By Modulatory Variables: Ideas on How to Proceed

To determine the effects of variables such as attention and orientation, we will need methodology for identifying substantial changes in the response field. The basic idea behind these modulation effects analyses is to find the responses under different conditions, such as different eye positions, and then detect the presence of any substantial differences between these responses. Fortunately, we expect the impulse response functions $a_v(t)$ to be common to all modulatory conditions, and so, this problem only involves detecting changes in the response fields $\beta_j(v)$.

Then one possible way to proceed is to detect changes in the center of mass of each response field. We will need to derive standard errors of the center of mass or approximations and develop tests for changes in the center of mass. Another possibility is model augmentation. We would incorporate into the model another parameter ω related to changes in the response field associated with different modulatory conditions; $\omega = 0$ if the variable is not associated with changes in the response field and is nonzero otherwise. We would then need to develop tests for the null hypothesis that $\omega = 0$.

6.3 Some Theoretical Results to Show

Finally, we plan to derive theoretical results regarding the performance of our techniques, which include

1. The probability that $\beta_j(v)$ for a given voxel or location coordinate v will be included in the nonzero region in Region Growing. This will most likely be a function of the $\beta_k(v)$ for $k \neq j$, the $a_v(t)$ estimates, and the level of noise in the data. Ultimately, these probabilities will be useful in the computation of the standard errors of the Region Growing $\beta_j(v)$ estimates, as well as proving results regarding the estimated support of $\beta_j(v)$.

2. Performance of the Penalized Regression technique in terms of correctly selecting nonzero $\beta_j(v)$. This includes the probability that if $\beta_j(v) \neq 0$ for given j and v, the LASSO penalty term in Penalized Regression causes the estimate of $\beta_j(v)$ to also be nonzero if the regularization parameter is selected

through prediction error, as well as the probability that we correctly set $\beta_j(v) = 0$ through the process of screening out spurious nonzero $\beta_j(v)$ after the execution of the technique.

3. Convergence of the Repeated Alternating Conditional Estimation procedure, and what it converges to. Because of the convexity of the overall minimization problem, it has a global minimum; we plan to show that each iteration of Repeated Alternating Conditional Estimation brings us closer to this minimum.

4. Statements about the convergence of the Region Growing procedure, specifically, what it converges to, and the rate at which the procedure converges.

A Additional Figures

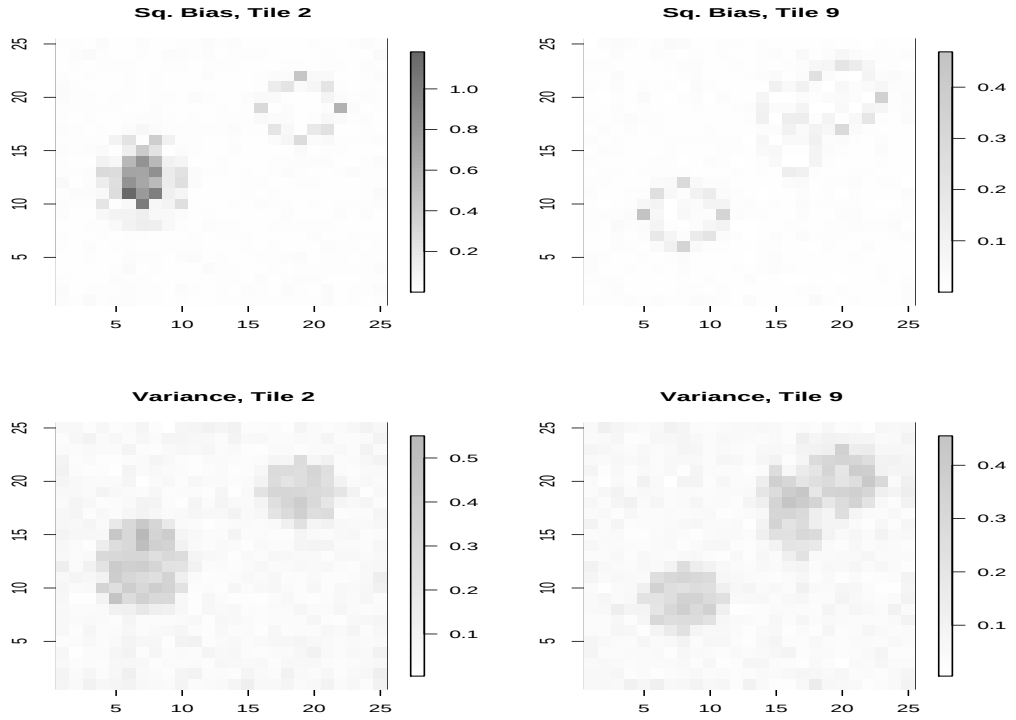

Figure 2: Squared biases and variances of the $\beta_j(v)$ Penalized Regression estimates for each voxel in the array, $j = 2$ and $j = 9$.

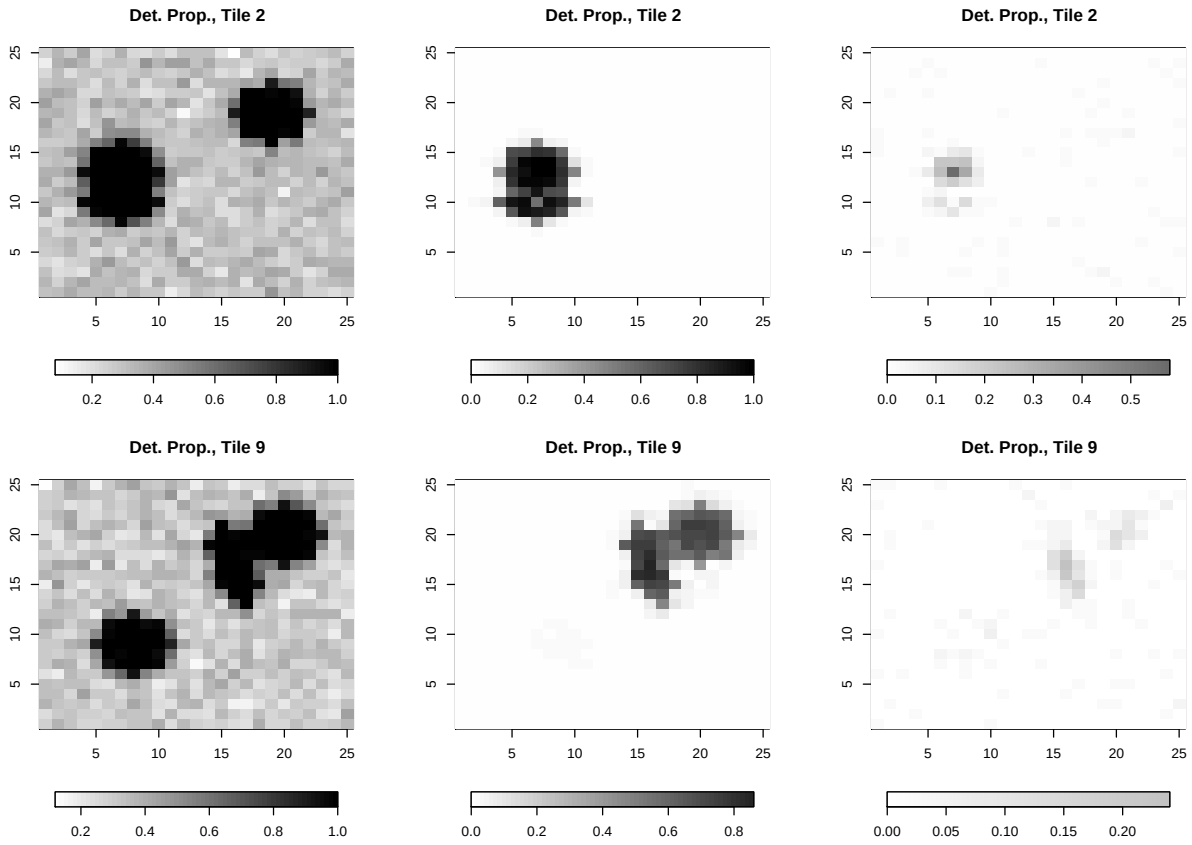

Figure 3: Detection proportions for $j = 2$ and $j = 9$ using all three methods. Left column: Penalized Regression. Center column: Region Growing. Right column: Reverse Correlation.

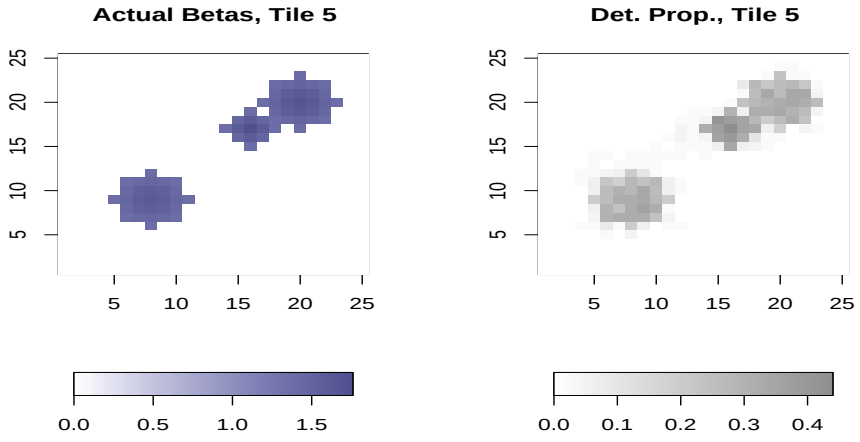

Figure 4: Region Growing results for $j = 5$. Left: actual $\beta_5(v)$. Right: detection proportions.

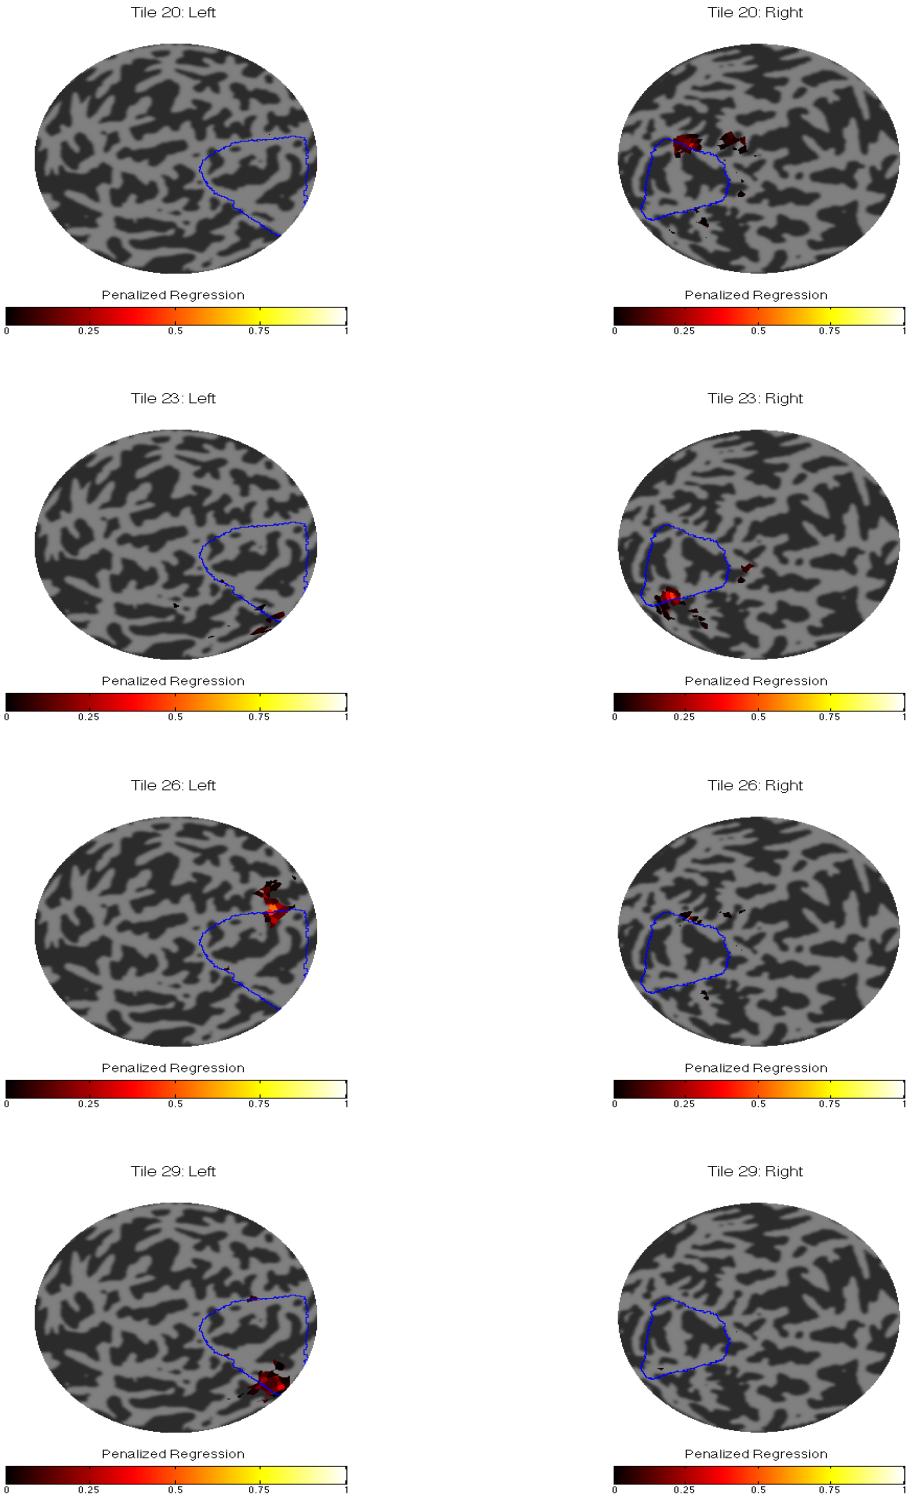

Figure 5: Penalized Regression flat maps for fMRI data analysis from Section 5.2. The area enclosed by the blue outline denotes V1 and red colors denote regions of activation.

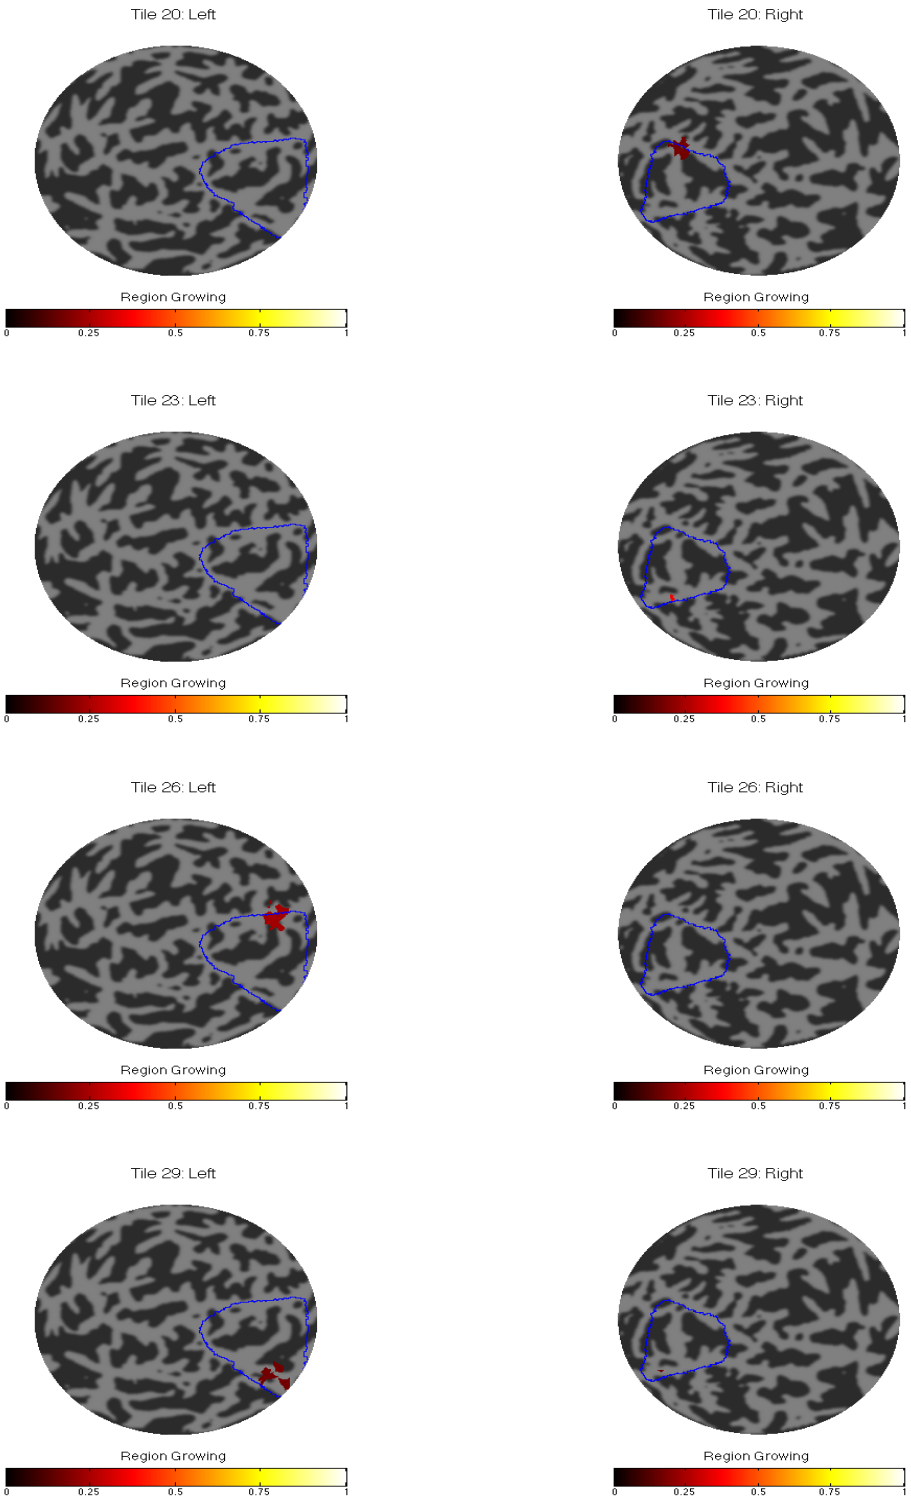

Figure 6: Region Growing flat maps for fMRI data analysis from Section 5.2. The area enclosed by the blue outline denotes V1 and red colors denote regions of activation.

References

[1] Cohen, M. S. (1997). "Parametric Analysis of fMRI Data Using Linear Systems Methods". *Neuroimage* 6, pg. 93-103.

[2] Casey, B. J., Davidson, M., Rosen, B. (2002). "Functional Magnetic Resonance Imaging: Basic Principles of and Application to Developmental Science". *Developmental Science* 5, pg. 301-309.

[3] Kay, K. N., Naselaris, T., Prenger, R. J., Gallant, J. L. (2008). "Identifying Natural Images from Human Brain Activity". To appear in *Nature*.

[4] Matthews, P. M., Honey, G. D., Bullmore, E. T. (2006). "Applications of fMRI in Translational Medicine and Clinical Practice". *Nature* 7, pg. 732-744.

[5] Fischl, B., Sereno, M. I., Dale, A. M. (1999). "Cortical Surface Based Analysis II: Inflation, Flattening, and a Surface-Based Coordinate System". *Neuroimage* 9, pg. 195-207.

[6] Reid, R. C., Victor, J. D., Shapley, R. M. (1997). "The Use of M-sequences in the Analysis of Visual Neurons: Linear Receptive Field Properties". *Visual Neuroscience* 14, pg. 1015-1027.

[7] Ringach, D., Shapley, R. (2004). "Reverse Correlation in Neurophysiology". *Cognitive Science* 28, pg. 147-166.

[8] Henson, R. N. (2001). "Event-Related fMRI: Introduction, Statistical Modeling, Design Optimization, and Examples". Paper presented at the 5^{th} Congress of the Cognitive Neuroscience Society of Japan.

[9] Buracas, G. T., Boynton, G. M. (2002). "Efficient Design of Event-Related fMRI Experiments Using M-sequences". *Neuroimage* 16, pg. 801-813.

[10] Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N. A., Dale, A. M. (1999). "Stochastic Designs in Event-Related fMRI". *Neuroimage* 10, pg. 607-619.

[11] Hansen, K. A., David, S. D., Gallant, J. L. (2004). "Parametric reverse correlation reveals spatial linearity of retinotopic human V1 BOLD response". *Neuroimage* 23, pg. 233-241.

[12] Glover, G. H. (1999). "Deconvolution of impulse response in event-related BOLD fMRI". *NeuroImage* 9, pg. 416-429.

[13] Worsley, K. J., Marrett, S., Neelin, P., Vandal, A. C., Friston, K. J., Evans, A. C. (1996). "A unified statistical approach for determining significant signals in images of cerebral activation". *Human Brain Mapping* 4, pg. 58-73.

[14] Katanoda, K., Matsuda, Y., Sugishita, M. (2002). "A spatio-temporal regression model for the analysis of functional MRI data". *Neuroimage* 17, pg. 1415-1428.

[15] Russ, J. C. (1992). *The Image Processing Handbook*. CRC Press, Inc., Boca Raton, FL.

[16] Lu, Y., Jiang, T., Zang, Y. (2003). "Region growing method for the analysis of functional MRI data". *Neuroimage* 20, pg. 455-465.

[17] Bylesjö, M., Eriksson, D., Sjödin, A., Sjöström, M., Jansson, S., Antti, H., Trygg, J. (2005). "MASQOT: a method for cDNA microarray spot quality control". *Biomed Central Bioinformatics* 6:250.

[18] Marx, B. D., Eilers, P. (1999). "Generalized Linear Regression for Sampled Signals or Curves: a P-spline Approach". *Technometrics* 41, pg. 1-13.

[19] Tibshirani, R. (1996). "Regression Shrinkage and Selection via the LASSO". *Journal of the Royal Statistical Society - Series B (Methodological)* 58, pg. 267-288.

[20] Meinshausen, N. (2007). "Relaxed LASSO". *Computational Statistics and Data Analysis* 52, pg. 374-393.

[21] Friedman, J., Hastie, T., Höfling, H., Tibshirani, R. (2007). "Pathwise Coordinate Optimization". *The Annals of Applied Statistics* 1, pg. 302-332.

[22] Wahba, G. (1990). *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.

[23] Kirchner, M. (2004). "Spatial extensions to self-modeling curve resolution". Technical Report, Multidimensional Image Processing, IWR, University of Heidelberg.

[24] Rusin, M. (1971). "A revised simplex method for quadratic programming". *SIAM: Applied Mathematics* 20, pg. 143-160.

[25] Hastie, T., Tibshirani, R., Friedman, J. (2001). *The Elements of Statistical Learning*. Springer Science, New York, NY.

[26] Fox, J. (1997). *Applied Regression Analysis, Linear Models, and Related Methods*. Sage Publications, Thousand Oaks, CA.

[27] Bai, B., Kantor, P. (2007). "A Shape-Based Finite Impulse Response Model for Functional Brain Images". *Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium.*

[28] Bear, M. F., Connors, B. W., Paradiso, M. A. (2001). *Neuroscience: Exploring the Second Brain (2nd Edition)*. Lippincott Williams and Wilkins, Philadelphia, PA.