Causal Inference for fMRI Time Series Data with Systematic Errors of Measurement in a Balanced On/Off Study of Social Evaluative Threat

Michael E. Sobel, Columbia University, Martin Lindquist, Johns Hopkins University

August 21, 2013

TREMENDOUS INTEREST IN CAUSAL RELATIONS IN NEUROSCIENCE

Small Scale—what is the effect of this stimulus on activity in particular brain locations

Large Scale—EFFECTIVE CONNECTIVITY—-causal relations among brain areas in producing outcomes of interest

Methods used—Granger causality, dynamic causal models, structural equation models, directed graphical models, network analysis methods

Some writing on causality—-all over the map.

No principled work on causal inference in this area, almost nothing informed by the modern statistical literature on causal inference.

DATA—for each subject, an image collected every 2 seconds for several minutes, each image consisting of measurements taken on (typically) the BOLD response at 100,000 voxels. Typically, the number of subjects is not large, e.g., ¡ 30.
APPROACH HERE—-BACK TO BASICS—START AT THE MOST ELEMENTAL LEVEL (BOLD response of a subject at a single voxel at a point in time) AND BUILD FROM THERE ENDS UP WITH SOME STRONG IMPLICATIONS EVEN AT THE SMALL SCALE—how neuroscientists should analyze data, for causal inference, defining effects in the presence of systematic error, alternative identification strategies for longitudinal data.

OUTLINE

I. WHAT fMRI MEASURES
A. The connection between Blood Flow and Neural Activity
B. BOLD—A noisy measurement of Blood Flow
II. CAUSAL INFERENCE WITH SYSTEMATIC ERROR RELATED TO INTERVENTION
A. Defining Unit and Average Causal Effects
B. Identification in "Balanced Designs"
C. A MODEL FOR THE BOLD RESPONSE
IV. EXAMPLE—A SOCIAL EVALUATIVE THREAT STUDY
V. DISCUSSION

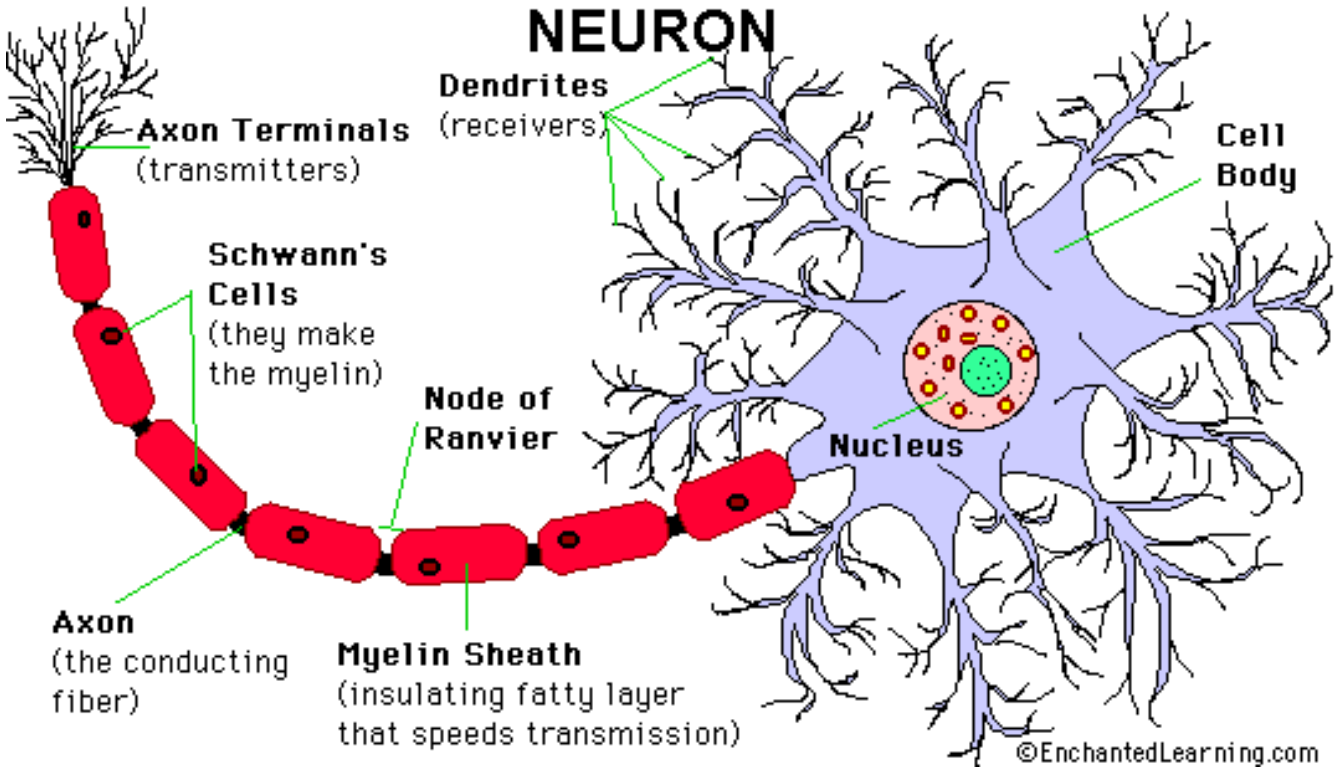
NEURON
Dendrites
(receivers)
Axon Terminals
(transmitters)
Cell
Body
Schwann's
Cells
(they make
the myelin)
Node of
Ranvier
Nucleus
Axon
(the conducting
fiber)
Myelin Sheath
(insulating fatty layer
that speeds transmission)
©EnchantedLearning.com

WHAT YOU WOULD LIKE TO MEASURE—-NEURAL ACTIVITY

NEURON—-approximately 100 billion of these in the human brain—-these are the basic units in information processing in the central nervous system.

MAIN COMPONENTS

Nucleus

Axon and axon ending—send signals to other neurons through axon terminals via small gaps (pre-synaptic process)—-when the axon hillock gets sufficiently depolarized, it sets up a chain (action potential) that propagates down the axon opening channels that allow the release of a neurotransmitter into the presynaptic terminal that forms a synapse with another neuron

Dendrites—integrate signals from other neurons via small gaps (post synaptic process)

WHAT fMRI MEASURES

Cannot measure neural activity directly except on a very limited basis—-invasive, might involve radioactive dyes, e.g., PET, single unit recording— cannot measure over whole brain.
fMRI is a technique from the 1990's that measures "blood flow" with good spatial resolution, adequate temporal resolution.

WHAT fMRI MEASURES

SO WHAT DOES BLOOD FLOW HAVE TO DO WITH NEURAL ACTIVITY?

KEY points:

1. The brain needs energy to restore membrane potential after action potential has occurred (like a pump)—uses ATP (adenosine triphosphate) made from glucose

2. The brain doesn't store energy, so need to import it—-so blood supplies oxygen and glycogen to the brain, which are important to synthesizing ATP.

3. When neurons become active, there is an increase in the amount of oxygenated hemoglobin which changes the balance of oxygenated to deoxygenated hemoglobin, i.e., decreases the proportion of deoxygenated hemoglobin. The deoxygenated hemoglobin is paramagnetic (attracted to a magnetic field) whereas the oxygenated blood is diagmagnetic.
4. Changes in the proportion of deoxygenated to oxygenated hemoglobin can be detected using strong magnetic fields (which is the MRI scanner) and certain types of pulse sequences. That is what the fMRI signal measures—-more signal when the blood is highly oxygenated.

WHAT fMRI MEASURES

The data from a fMRI experiment are (for each subject i) a multivariate (about 100,000 voxels)) time series (usually a sequence of several hundred equally spaced (TR =2) time points) of BOLD response values.
The time points are not really equally spaced, as the data are collected in slices throughout the interval (TR). Because of the nature of blood flow (the hemodynamic response function to be discussed later) a SLICE-TIMING adjustment is applied.
MOTION (especially head motion) is a real problem——(one wants the BOLD response series for voxel v in subject i to be voxel v at each time point)–the images are realigned—MORE on this later.
To enable analyses across subjects, the data are then normalized to an anatomical template (here the 152 brain template of the Montreal Neurological Institute). The normalized images are then smoothed using a Gaussian kernel.

CAUSAL INFERENCE

Subjects, i = 1,...,n
Times, $k_i + t$, t = 1,...,T. k_i is subject i's starting time.
In period t, either no stimulus is applied or a stimulus $j \in \{1, ..., J\}$, represented as a $J \times 1$ unit vector $\mathbf{z}_t$, with 1 in position j, is to be administered.
A treatment regimen, denoted $\bar{\mathbf{z}}_T \equiv (\mathbf{z_1}, ..., \mathbf{z_T})$ is a sequence of T stimuli.
The notation $\bar{\mathbf{z}}_t \equiv (\mathbf{z_1}, ..., \mathbf{z_t})$ is used to denote a sub-regimen of $\bar{\mathbf{z}}_T$ through period t.
For each subject i and voxel $v \in \{1, ..., V\}$, we consider the potential BOLD series $Y_{iv,k_i+1}(\bar{\mathbf{z}}_T), ..., Y_{iv,k_i+T}(\bar{\mathbf{z}}_T)$ for each regimen $\bar{\mathbf{z}}_T \in \Omega$,
Notice we assume SUTVA—quite reasonable here.

CAUSAL INFERENCE

We described preprocessing steps taken to render valid assumptions about the spatial and temporal alignment of the BOLD responses. However, even after such steps, these contain both random errors and systematic errors due to scanner drift and head motion not corrected for during preprocessing. We decompose the potential responses as:
(A1) BOLD Response Decomposition — For all $\bar{\mathbf{z}}_T \in \Omega$,

$$Y_{iv,k_i+t}(\bar{\mathbf{z}}_T) = \mathcal{T}_{iv,k_i+t}(\bar{\mathbf{z}}_T) + S_{iv,k_i+t}(\bar{\mathbf{z}}_T) + \varepsilon_{iv,k_i+t}(\bar{\mathbf{z}}_T), \qquad (1)$$

where $\mathcal{T}_{iv,k_i+t}(\bar{\mathbf{z}}_T)$ is the subject's "true" BOLD response, $\varepsilon_{iv,k_i+t}(\bar{\mathbf{z}}_T)$ is a random error with mean 0, and $S_{iv,k_i+t}(\bar{\mathbf{z}}_T)$ is the systematic error or measurement bias for subject i at voxel v at time $k_i + t$.

UNIT EFFECTS

For subject i, the effect of treatment regimen $\bar{\mathbf{z}}_T$ vs. $\bar{\mathbf{z}}_T^*$ at period $t = 1, ..., T$ would normally be defined as $Y_{iv,k_i+t}(\bar{\mathbf{z}}_T) - Y_{iv,k_i+t}(\bar{\mathbf{z}}_T^*)$.

If systematic error is task related (for example, task related head-motion not removed during pre-processing), such a definition incorporates errors $S_{iv,k_i+t}(\bar{\mathbf{z}}_T) - S_{iv,k_i+t}(\bar{\mathbf{z}}_T^*)$ as a component of the unit effect. Second, even in the absence of systematic error, the observed BOLD response is subject to random errors of measurement produced by cardiac and respiratory activity, for example.

DEFINE UNIT EFFECTS as:

$$\tau_{iv,k_i+t}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}^*_T) \equiv \mathcal{T}_{iv,k_i+t}(\bar{\mathbf{z}}_T) - \mathcal{T}_{iv,k_i+t}(\bar{\mathbf{z}}^*_T). \tag{2}$$

(A2) Invariance of True Outcomes to Start Times– $\forall i$, $\bar{\mathbf{z}}_T \in \Omega$, and (k_i, k_i'), $\mathcal{T}_{iv,k_i+t}(\bar{\mathbf{z}}_T) = \mathcal{T}_{iv,k_i'+t}(\bar{\mathbf{z}}_T) \equiv \mathcal{T}_{ivt}(\bar{\mathbf{z}}_T)$. So we can write $\tau_{iv,k_i+t}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}^*_T) = \tau_{ivt}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}^*_T)$. Note however, the BOLD responses $Y_{iv,k_i+t}(\bar{\mathbf{z}}_T)$ may depend on k_i.

CAUSAL INFERENCE

Average Treatment Effect

$$E(\tau_{ivt}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}_T^*)) \tag{3}$$

Variance of Subject Effects

$$V(\tau_{ivt}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}_T^*)) = E(\tau_{ivt}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}_T^*) - E(\tau_{ivt}(\bar{\mathbf{z}}_T, \bar{\mathbf{z}}_T^*)))^2, \tag{4}$$

where the expectations are taken over subjects.

CAUSAL INFERENCE

Fundamental Problem of Causal Inference
Usually INDIVIDUAL effects are treated as unidentified and AVERAGE effects identified using sequential ignorability and positivity assumptions, in conjunction with modeling assumptions.

In "BALANCED" designs, where each subject receives the same treatment regimen, comparison of subjects under different regimens cannot be used to identify average treatment effects. Thus, we make alternative assumptions to identify the individual effects (2). The effects (3) and (4) are then obtained from these.

(A3) Temporal Consistency — For all $\bar{\mathbf{z}}_T \in \Omega$, the BOLD response $Y_{iv,k_i+t}(\bar{\mathbf{z}}_T)$ in period t and its components $\mathcal{T}_{ivt}(\bar{\mathbf{z}}_T)$, $S_{iv,k_i+t}(\bar{\mathbf{z}}_T)$, $\varepsilon_{iv,k_i+t}(\bar{\mathbf{z}}_T)$ do not depend on stimuli administered after period t. So we may write $Y_{iv,k_i+t}(\bar{\mathbf{z}}_T) = Y_{iv,k_i+t}(\bar{\mathbf{z}}_t)$, $\mathcal{T}_{ivt}(\bar{\mathbf{z}}_T) = \mathcal{T}_{ivt}(\bar{\mathbf{z}}_t)$, etc.

(A4) P Period Carry-Over — Let $0 \leq P \leq T-1$ denote the smallest integer such that $\mathbf{z}_{t-P} = \mathbf{z}^*_{t-P}, ..., \mathbf{z}_t = \mathbf{z}^*_t$ implies $\mathcal{T}_{ivt}(\bar{\mathbf{z}}_t) = \mathcal{T}_{ivt}(\bar{\mathbf{z}}^*_t)$ for all $t \geq P+1$.

(A5) No Treatment by Period Interaction for True Outcomes — For P in (A4) and $P+1 \leq t < t'$, $\mathbf{z}_{t-P} = \mathbf{z}_{t'-P}, ..., \mathbf{z}_t = \mathbf{z}_{t'}$ implies $\mathcal{T}_{ivt'}(\bar{\mathbf{z}}_{t'}) = \mathcal{T}_{ivt}(\bar{\mathbf{z}}_t) + c(t, t')$.

CAUSAL INFERENCE

Now we construct a model for the dependence of BOLD responses on the treatment sequence that implies the assumptions above: 1. Key things
HEMODYNAMIC RESPONSE FUNCTION

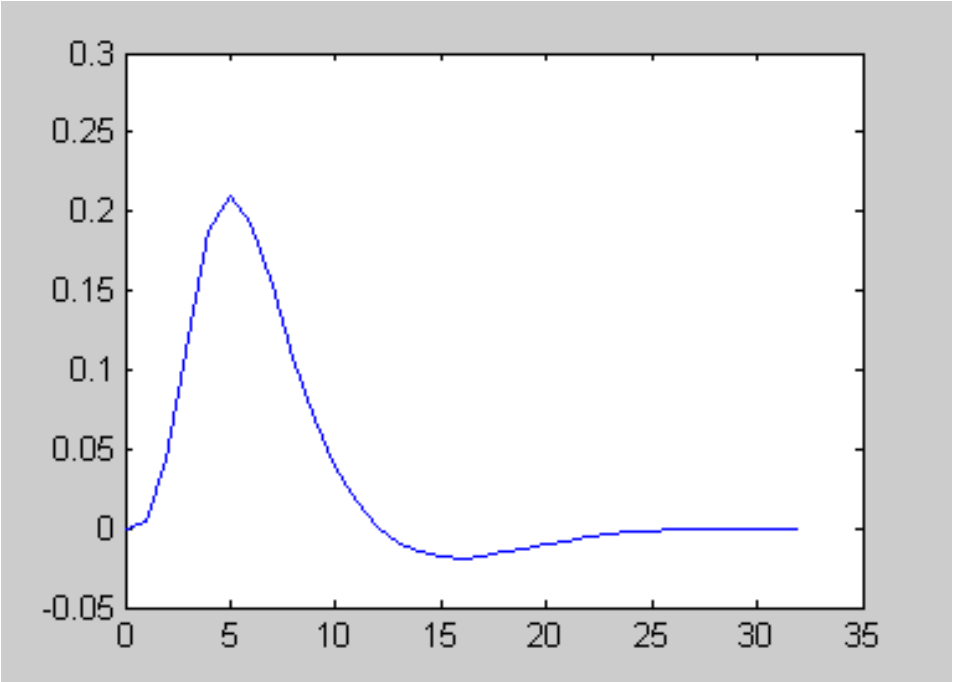
0.3
0.25
0.2
0.15
0.1
0.05
0
-0.05
0
5
10
15
20
25
30
35

SUPERPOSITION If more than one stimulus is applied, the individual hemodynamic response sequences are added. We describe this using convolution. For stimulus j, with $(z_{j,t-P}, ..., z_{jt})$, $f_j(\bar{\mathbf{z}}_t) = \sum_{p=0}^{P} z_{j,t-p} h_p$.

CAUSAL INFERENCE

fMRI researchers are typically interested in the average amplitude (and occasionally variance) in amplitude of a response to treatment. We model these quantities and relate these to the effects defined in equations (2)-(4) using a linear mixed model for the potential outcomes in Ω that implies assumptions (A1)-(A5):

CAUSAL INFERENCE

$$Y_{iv,k_i+t}(\bar{\mathbf{z}}_t) = (\beta_v + \mathbf{b}_{iv})'\mathbf{f}_t(\bar{\mathbf{z}}_t) + (\gamma_{1v} + \mathbf{g}_{i1v})'\mathbf{N}_t + (\gamma_{2v} + \mathbf{g}_{i2v})'\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t) + \varepsilon_{iv,k_i+t}(\bar{\mathbf{z}}_t), \quad (5)$$

where $\mathbf{f}_t(\bar{\mathbf{z}}_t) = (f_1(\bar{\mathbf{z}}_t), ..., f_J(\bar{\mathbf{z}}_t))'$, $(\beta_v + \mathbf{b_{iv}})'\mathbf{f_t}(\bar{\mathbf{z}}_t) = \mathcal{T}_{ivt}(\bar{\mathbf{z}}_t)$, $f_j(\bar{\mathbf{z}}_t) = \sum_{p=0}^{P} z_{j,t-p}h_p$ is the convolution of the treatment subsequence $(z_{j,t-P}, ..., z_{jt})$ with the hemodynamic response function $(\gamma_{1v} + \mathbf{g}_{i1v})'\mathbf{N}_t + (\gamma_{2v} + \mathbf{g}_{i2v})'\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)$ is the systematic error, $(\mathbf{N}_t', \mathbf{N}_{i,k_i+t}'(\bar{\mathbf{z}}_t))'$ is a vector of measured "nuisance covariates" and $\mathrm{E}(\varepsilon_{iv,k_i+t}(\bar{\mathbf{z}}_t) \mid \theta_{iv}, \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}) = 0$, where $\theta_{iv} = (\mathbf{b}_{iv}', \mathbf{g}_{i1v}', \mathbf{g}_{i2v}')'$.

Thus, the subject and average causal effects (2) and (3) are, respectively:

$$\tau_{ivt}(\bar{\mathbf{z}}_t, \bar{\mathbf{z}}_t^*) = (\boldsymbol{\beta}_v + \mathbf{b}_{iv})'(\mathbf{f}_t(\bar{\mathbf{z}}_t) - \mathbf{f}_t(\bar{\mathbf{z}}_t^*)), \tag{6}$$

$$E(\tau_{ivt}(\bar{\mathbf{z}}_t, \bar{\mathbf{z}}_t^*)) = \boldsymbol{\beta}_v'(\mathbf{f}_t(\bar{\mathbf{z}}_t) - \mathbf{f}_t(\bar{\mathbf{z}}_t^*)), \tag{7}$$

amplitude β_{jv} ($\beta_{jv} + b_{ijv}$), the j^{th} component of $\boldsymbol{\beta}_v$ $\boldsymbol{\beta}_v + \mathbf{b}_{iv}$) is the average effect (effect for subject i) of a 1 unit increase in the j^{th} component of $\mathbf{f}_t(\bar{\mathbf{z}}_t) - \mathbf{f}_t(\bar{\mathbf{z}}_t^*)$, $f_{jt}(\bar{\mathbf{z}}_t) - f_{jt}(\bar{\mathbf{z}}_t^*)$, with the other components $\mathbf{f}_{j't}(\bar{\mathbf{z}}_t) - \mathbf{f}_{j't}(\bar{\mathbf{z}}_t^*) = 0$, $j' \neq j$.

CAUSAL INFERENCE

The sources of systematic bias that fMRI researchers sometimes adjust for include 1) scanner drift and 2) head motion related artifacts not accounted for during preprocessing. Scanner drift is a low frequency change in the MR signal due to the imaging hardware, which creates slow changes in voxel intensities, and thus slow changes in the BOLD response. This source of bias does not depend on the treatment regimen $\bar{\mathbf{z}}_T$ and depends on the subject only through $k_i + t,\ t = 1, ..., T$.

Drift may be modeled as

$$D_{iv,k_i+t} = (\gamma_{1v} + \mathbf{g}_{i1v})'\mathbf{N}_t + \delta_{iv,k_i+t}, \tag{8}$$

where $E(\delta_{iv,k_i+t} \mid \boldsymbol{\theta}_{iv}, \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}) = 0$ and the ℓ^{th} coordinate of the basis vector $\mathbf{N}_t$, $\ell = 1, ..., L_1$, is, typically, $N_{\ell t} = t^{\ell-1}$.

Model head motion not corrected for using preprocessing:

$$H_{iv,k_i+t}(\bar{\mathbf{z}}_t) = (\gamma_{2v} + \mathbf{g}_{i2v})'\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t) + \kappa_{iv,k_i+t}(\bar{\mathbf{z}}_t), \qquad (9)$$

where $E(\kappa_{iv,k_i+t}(\bar{\mathbf{z}}_t) \mid \boldsymbol{\theta}_{iv}, \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}) = 0$, and $N_{i\ell,k_i+t}(\bar{\mathbf{z}}_t)$, $\ell = L_1 + 1, ..., L$, is the ℓ^{th} "nuisance covariate" for head motion for subject i at voxel v at time $k_i + t$. Typically the motion regressors are six time courses corresponding to three translations and three rotations of the brain computed using a rigid-body transformation for each functional scan thus, the values do not depend on v.

CAUSAL INFERENCE

Head motion may depend on $\bar{\mathbf{z}}_T$. In randomized experiments with $T = 1$, it is well known that adjustment for intermediate outcomes affected by treatment leads to biased estimates of causal effects. Here, however, if motion regressors are not included and they are correlated with $\mathbf{f}_t(\bar{\mathbf{z}}_t)$, estimates of the amplitudes $\boldsymbol{\beta}_v$ will be biased, hence estimates of the unit and average effects (2) and (3) will also be biased.

CAUSAL INFERENCE

To complete the model, we specify the distribution of the random components and the relationship among these. Let $\varepsilon_{iv}(\bar{\mathbf{z}}_T) = (\varepsilon_{iv,k_i+1}(\bar{\mathbf{z}}_1), ..., \varepsilon_{iv,k_i+T}(\bar{\mathbf{z}}_T))'$. We assume $\varepsilon_{iv}(\bar{\mathbf{z}}_T) \mid \boldsymbol{\theta}_{iv}, \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T} \sim N(\mathbf{0}, \Sigma_{v\varepsilon\varepsilon}(\bar{\mathbf{z}}_t))$, where N denotes the normal distribution, $(\boldsymbol{\theta}_{iv} \mid \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}) \sim N(\mathbf{0}, \Sigma_{v\theta\theta})$, with $\mathrm{E}(\mathbf{g}_{1iv}\mathbf{b}'_{iv} \mid \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}) = 0$, $E(\mathbf{g}_{1iv}\mathbf{g}'_{2iv} \mid \{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}) = 0$. Given $\{\mathbf{N}_{i,k_i+t}(\bar{\mathbf{z}}_t)\}_{t=1}^{T}$, $\boldsymbol{\theta}_{iv}$ and $\varepsilon_{iv}(\bar{\mathbf{z}}_T)$ are assumed independent. The random vectors $\boldsymbol{\theta}_{1v}, ...\boldsymbol{\theta}_{nv}, \varepsilon_{1v}(\bar{\mathbf{z}}_T), ..., \varepsilon_{nv}(\bar{\mathbf{z}}_T)$ are assumed to be independent. Estimation by ML is straightorward.

CAUSAL INFERENCE

An important caveat is in order. In a balanced design, where all subjects are assigned the same regimen, limited departures from the model (5) can be assessed by adding additional terms. But some types of departures cannot be assessed. In particular, the parameters in (5) are assumed invariant over treatment regimens, but if only one regimen is examined, it is not possible to evaluate this assumption. In a randomized experiment, where subjects are assigned to different regimens, this assumption can be tested by comparing parameter estimates from different regimens.

EXAMPLE: The SET Task Data Set

The data, a time series of 215 images acquired every two seconds for each of 25 subjects, were collected to study the effects on brain activity and heart rate of an anxiety producing experiment threatening the subjects self-evaluation. An off/on/off balanced design was used. Subjects were initially told they would be asked to prepare a 7 minute speech for possible presentation to their peers. At the start of scanning, they viewed a fixation cross for 2 minutes (resting baseline). They then viewed a slide for 15 seconds describing the topic (why subject is a good friend). The slide instructed subjects to prepare enough material for the entire 7 minutes. After 2 minutes of silent preparation, a 15 second relief instruction, telling subjects they would not have to give the speech, was given. Subjects were then scanned for an additional 2.5 minutes.

EXAMPLE

To illustrate the impact of adjusting for systematic error in the SET task study, we fit two models. Both include (a) two activation regressors: 1) "EarlyPrep" for activation from the first visual cue to the middle of speech preparation ($z_{1t} = 1$ for $t \in [61, 90]$, 0 otherwise) and 2) "LatePrep" for activation from the middle of preparation until the presentation of the relief cue ($z_{2t} = 1$ for $t \in [91, 120]$, 0 otherwise), and (b) three drift regressors: $N_{\ell t} = t^{\ell - 1}$, $\ell = 1, ..., 3$. Model 1 also includes motion regressors $N_{\ell t}$, $\ell = 4, ..., 9$, specifically the 6 head movement regressors from realignment depicting how far the brain lies from a reference image at each time point. These are included to correct for motion related 'spin-history' artifacts present in the data. Thresholded t-maps for the parameter associated with "EarlyPrep" obtained using Model 1 are presented for a single slice (number 18 out of 46) in the figure below.

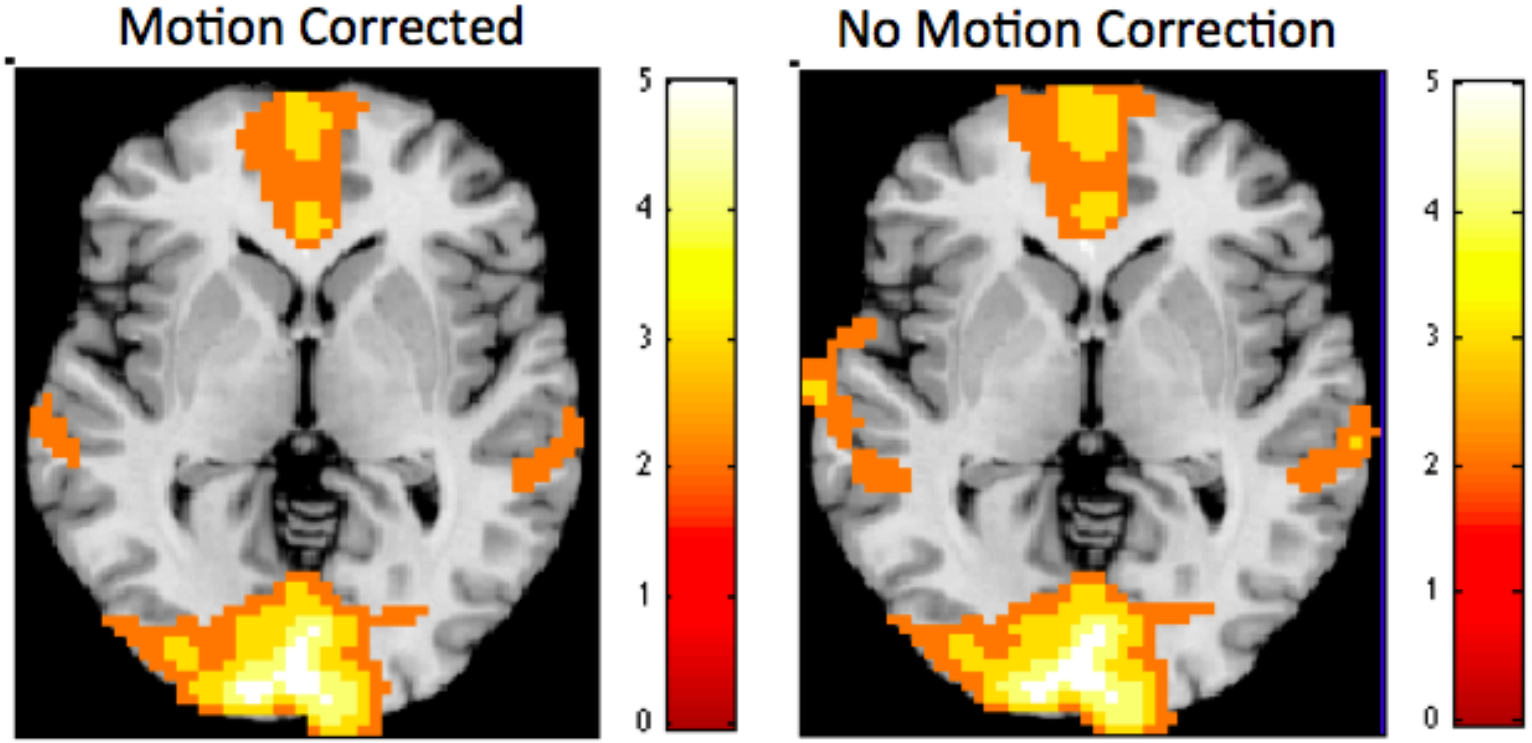

Left: Activation in visual cortex, superior temporal cortices (associated with agency) and ventromedial prefrontal cortex (associated with visceromotor control, self related attention, generation and regulation of emotion).
Right: Also ventricles and edges

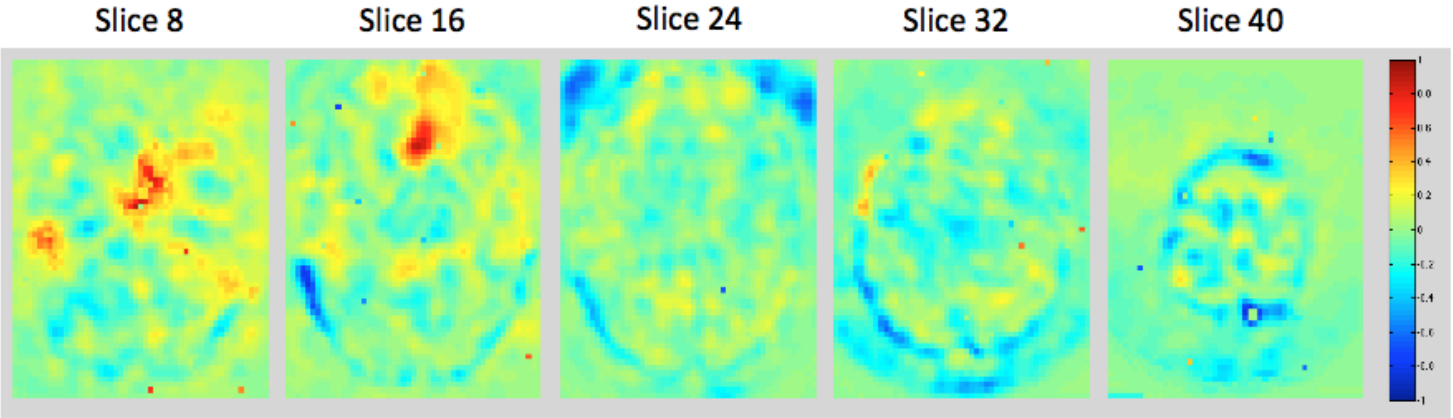

Early Prep, differences in motion corrected and uncorrected coefficients, slices 8,...40; blue edges.

DISCUSSION

1. DEFINITION OF TREATMENT EFFECTS SUBJECT TO SYSTEMATIC ERROR
2. IDENTIFICATION

DISCUSSION

FINDINGS

Visual cortex and superior temporal cortices are activated during the early part of speech preparation and the ventromedial prefrontal cortex is activated in both early and later phases of the task. In addition, we find little between subject variation in activation in the visual cortex, greater variation in the ventromedial prefrontal cortex.

DISCUSSION

IMPLICATIONS AND NEXT STEPS Our approach and results have implications more generally for the manner in which neuroscientists make causal inferences using fMRI data. Currently, less than a third of analyses in the literature include motion regressors in models for the BOLD responses But an immediate consequence of our analysis is that motion regressors should be included in models for the BOLD response to adjust for task related systematic errors.

DISCUSSION

REGION OF INTEREST ANALYSES–typically performed by averaging a subject's BOLD responses over the voxels comprising the region and modelling the aggregated responses. However, unless the systematic errors average out to 0 over the region, and there is no reason to think that will be so, the averaged responses will also contain task related error, necessitating adjustment, as before. In addition, such an analysis will obscure any variation in signal across the region. Such variation, if present, may be of interest in localizing subregions within ROIs that account for how resources are allocated and processed within the ROI. A better approach, conceptually straightforward, but computationally demanding, is to extend the model of Section 3 over the voxels in the predefined region, and treat the amplitudes as random effects over voxels and subjects: if the signal is homogeneous throughout the region, the null hypothesis that the voxel variance is 0 will not be rejected.

DISCUSSION

FUNCTIONAL MEDIATION—As another example, Lindquist modeled the relationship between a thermal stimulus and reported pain, as mediated by the time course of activity, treated as a functional mediator, in various brain regions. The BOLD responses $Y_{iv,k_i+t}(\mathbf{z}_t)$ in the region were averaged and the aggregated responses were treated as the components of the functional mediator. As above, these averaged responses will generally reflect task related effects and systematic error. In future work, we shall separate these components, defining the components of the functional mediator using the "true" hemodynamic responses $\tau_{iv,k_i+t}(\mathbf{z}_t)$.