JOHNS HOPKINS BLOOMBERG SCHOOL of PUBLIC HEALTH

SMART JHSPH Biostatistics

Statistical Methods in functional MRI

Martin Lindquist
Department of Biostatistics
Johns Hopkins University

Functional MRI

- Functional magnetic resonance imaging (fMRI) is a non-invasive technique for studying brain activity.
- During the course of an fMRI experiment, a series of brain images are acquired while the subject performs a set of tasks.
- Changes in the measured signal between individual images are used to make inferences regarding task-related activations in the brain.

Advantages of fMRI

- In the past decade fMRI has become the dominant tool for functional imaging due in large to its balance of spatial and temporal resolution, ease of access, and collaborative community.

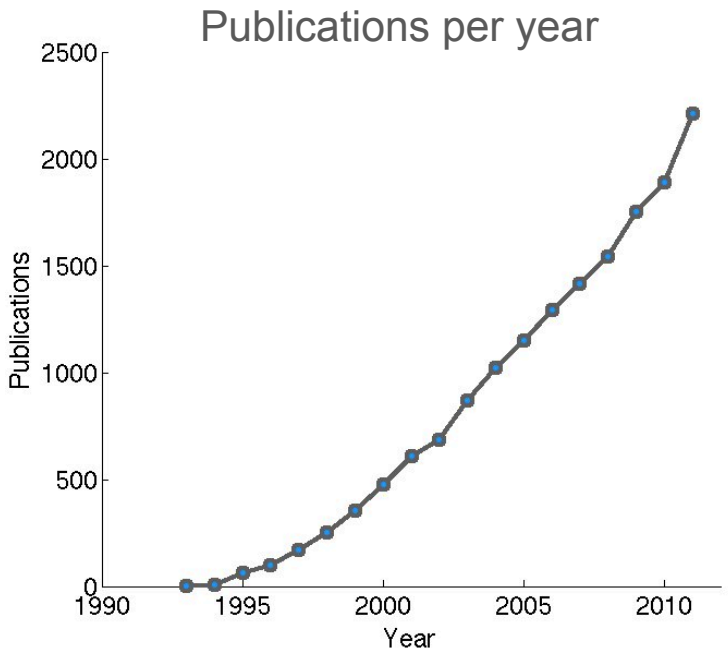

fMRI Data

- Each fMRI image consists of ~100,000 brain 'voxels' (cubic volumes that span the 3D space of the brain).

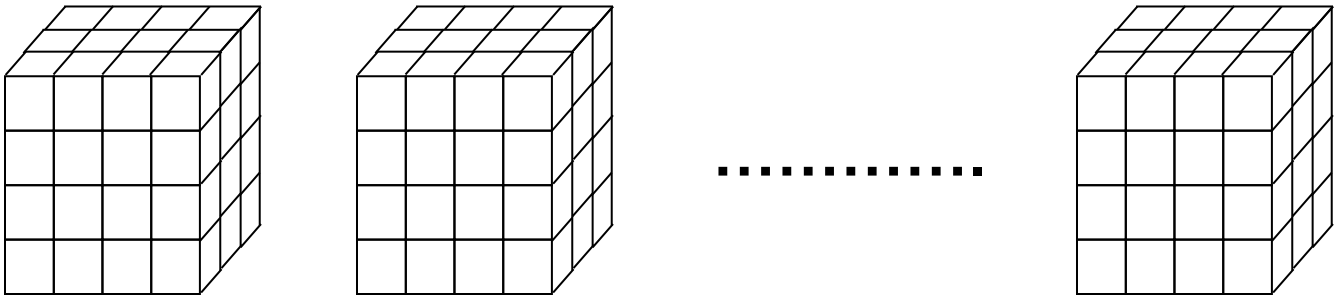

- During the course of an experiment several hundred images are acquired (~ one every 2*s*).

fMRI Data

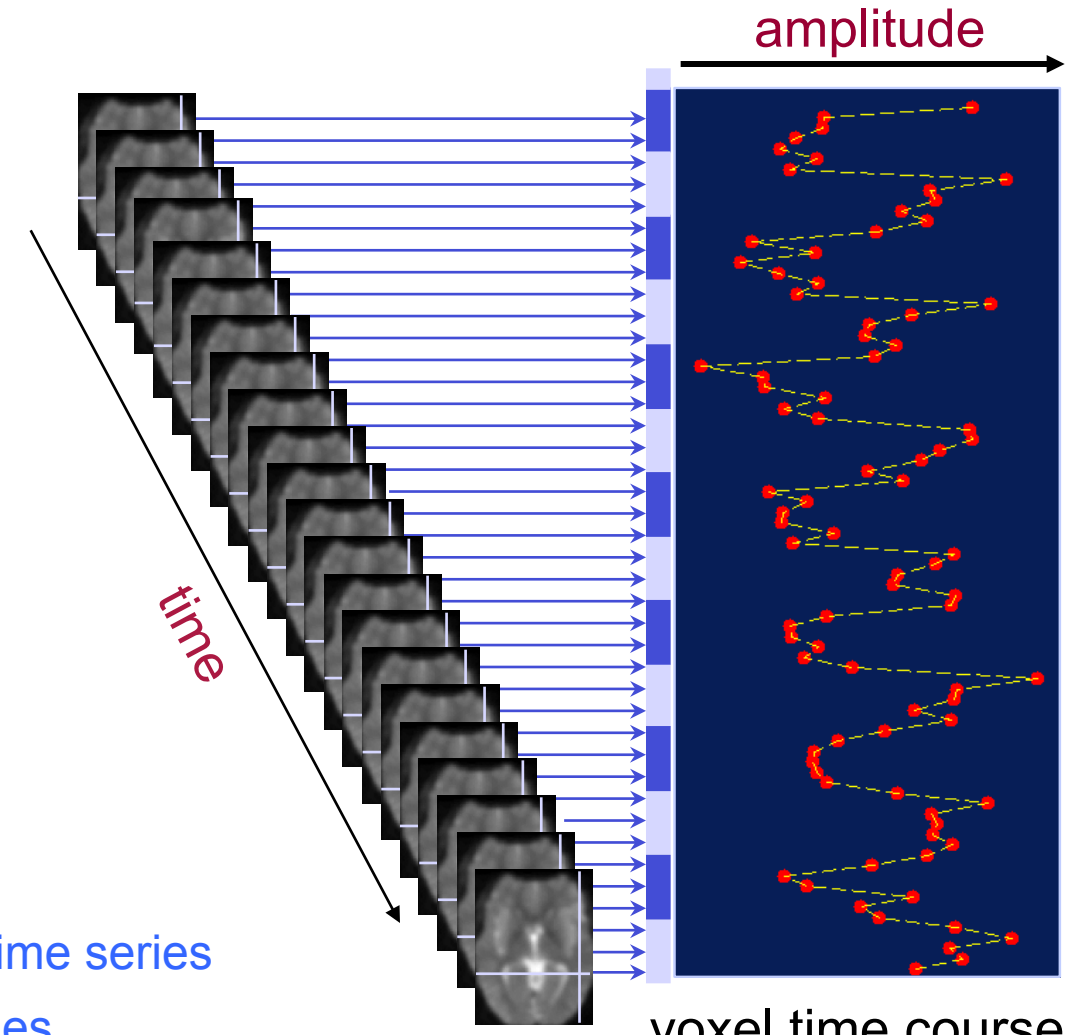

One voxel = One time series

~100,000 time series

fMRI Data

- In addition, the experiment may be repeated for multiple subjects (e.g.,10–40) to facilitate population inference.

- The total amount of data that needs to be analyzed is staggering.

Statistical Analysis

- There are multiple goals in the statistical analysis of fMRI data.

- They include:
  - localizing brain areas activated by the task;
  - determining networks corresponding to brain function; and
  - making predictions about psychological or disease states.

Statistical Analysis

- The statistical analysis of fMRI data is challenging.
  - It is a massive data problem.
  - The signal of interest is relatively weak.
  - The data exhibits a complicated temporal and spatial noise structure.

Major Analysis Packages

SPM	**Open source, but requires Matlab, which is expensive and comparatively slow. Widespread. Large variety of add-ons. Very active development community. Runs on all platforms. Emphasis on fast processing. Free.**
Voxbo	**Open source, written in C++. Has a focus on scheduling and network batch processing. Linux, Mac OS X, Windows using Cygwin. Small community. Currently being developed? Free.**
FSL	**Open source, written in C++. Many unique tools. Command-line and graphical (GUI) interface. Linux, OS X, Windows using Cygwin. Emphasis on timeseries modeling and cool new techniques. Actively developed. Free.**
AFNI	**Open source, written in C. Part command-line, part GUI. Free.**
fMRIStat	**Keith Worsley's package (updated by Jonathan Taylor). Open source, Matlab/Python. Command-line, advanced statistics. Free.**
BrainVoyager	**Closed source commercial product, written in C++. Runs on all platforms. Slick GUI interface. Good surface inflation visualization. Fairly large set of tools, but smaller community. Has real-time options. Extensible via C++. $30K for 6 licenses / ~$8K each.**

Magnetic Resonance Imaging

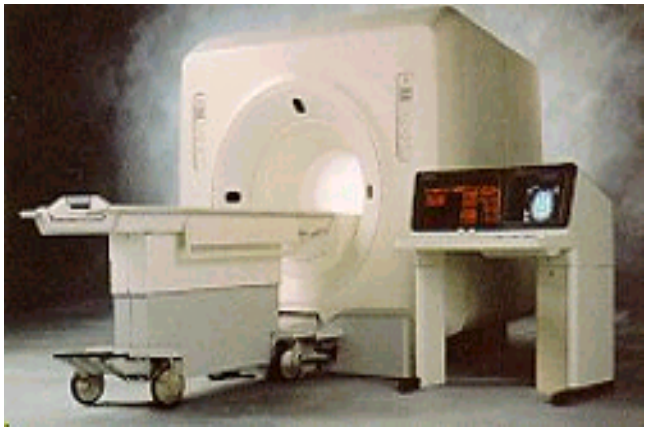

An MR scanner consists of an electromagnet with a very strong magnetic field (1.5 - 7.0 Tesla)

Earth's magnetic field = 0.00005 Tesla

3 Tesla is 60,000 times stronger than the Earth's magnetic field.

What MRI Measures

- MRI scanners are extremely versatile and can be used to study both brain structure and brain function.
- Both structural and functional MRI images are acquired using the same scanner.
- Different types of brain images can be generated to emphasize contrast related to different tissue characteristics.

Signal Formation

- The subject is placed into the MR scanner.
  - Nuclei of ^{1}H atoms align with the magnetic field.
  - The nuclei precess about the field at similar frequencies, but at a random phase.
  - Net longitudinal magnetization in the direction of field.

- Within a slice, a radio frequency (RF) pulse is used to align the phase and ‘tip over’ the nuclei.
  - Causes the longitudinal magnetization to decrease, and establishes a new transversal magnetization.

Signal Formation

- After the RF pulse is removed, the system seeks to return to equilibrium.
  - The transverse magnetization disappears (transversal relaxation), and the longitudinal magnetization grows back to its original size (longitudinal relaxation).
  - Longitudinal relaxation: exponential growth described by time constant T1.
  - Transverse relaxation: exponential decay described by time constant T2.

- During this process a signal is created that can be measured using a receiver coil.

Image Contrast

- By altering how often we excite the nuclei (TR) and how soon after excitation we begin data collection (TE) we can control key characteristics of the signal.

- The measured signal is approximately

$$M_0(1-e^{-TR/T_1})e^{-TE/T_2}$$

where M_0 depends on the proton density.

fMRI Contrast

- T_2^* is the combined effect of T_2 and local inhomogeneities in the magnetic field.
- Certain pulse sequences are able to eliminate the effects of these inhomogeneities, while others seek to emphasize them.
- The latter types of procedures form the basis of BOLD fMRI.

Image Contrast

- Different pulse sequences produce images that are sensitive primarily to T_1, T_2, or T_2^*.
- Because T_1 and T_2 vary with tissue type, they are able to represent boundaries between CSF, gray and white matter.
- Because T_2^* is sensitive to flow and oxygenation, it is can be used to image brain function.

Image Contrast

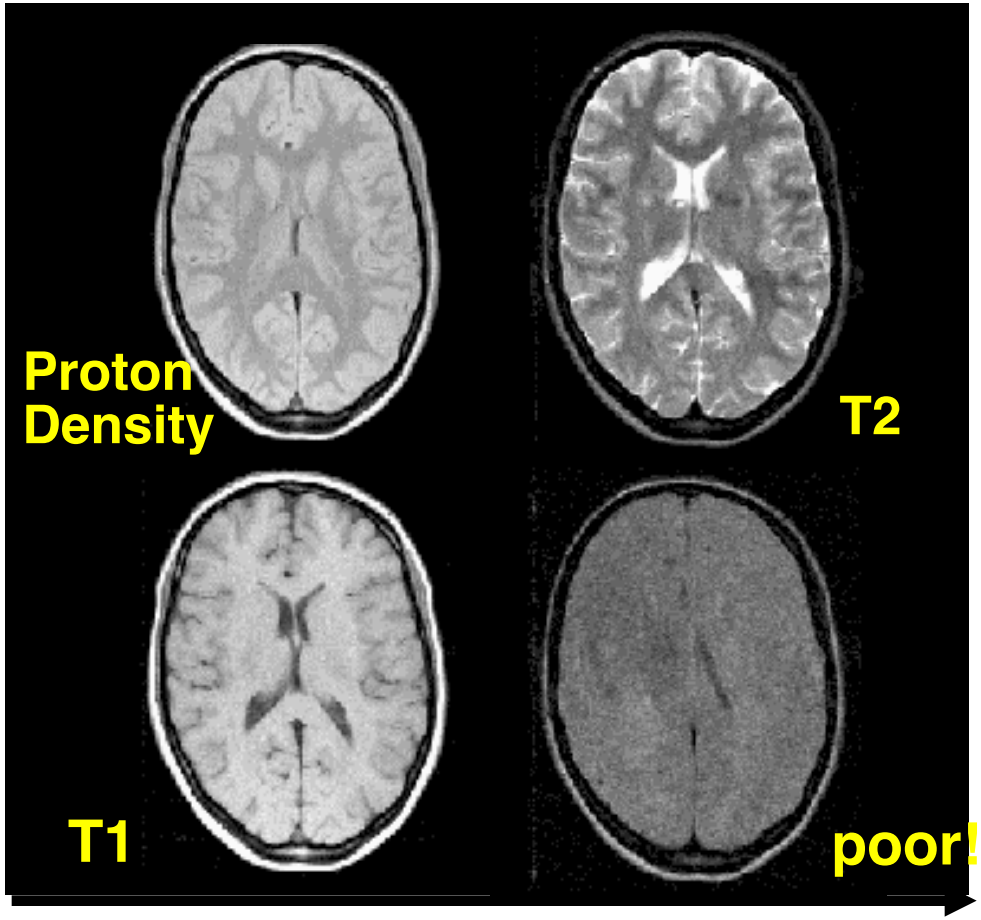

TE (echo time) - the time between excitation and data collection.

$$M_0(1-e^{-TR/T_1})e^{-TE/T_2}$$

Image Formation

- The goal of MRI is to construct an image, or a matrix of numbers that correspond to spatial locations.
- The image depicts the spatial distribution of some property of the nuclei within the sample.
- This could be the density of nuclei, their mobility, or the relaxation time of the tissues in which they reside.

Image Formation

- Imagine a brain slice split into a number of equally sized volume elements or voxels.

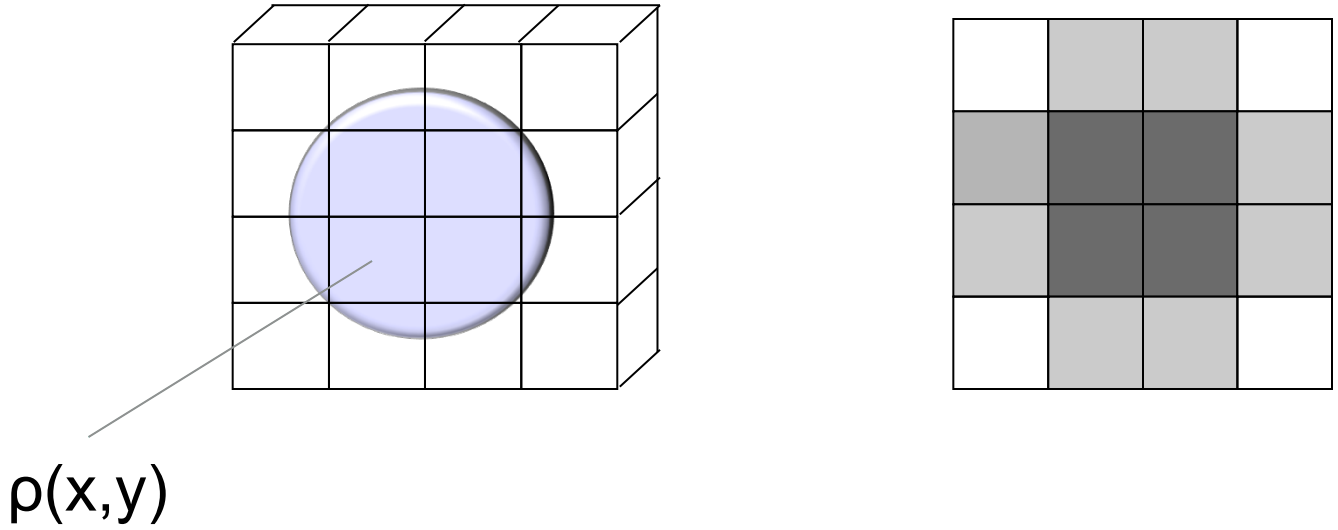

Gradients

- The measured signal combines information from the whole brain:

$$S(t) = \int\int \rho(x,y)\,dx\,dy$$

- A magnetic field gradient is used to sequentially control the spatial inhomogeneity of the magnetic field, so each measurement can be expressed:

$$S(k_x,k_y) = \iint \rho(x,y)e^{-i2\pi(k_x x+k_y y)}\,dxdy$$

K-space

- The measurements are acquired in the frequency-domain (k-space).

- By making measurements for multiple values of (k_x, k_y) we can gain enough information to solve the inverse problem and reconstruct $\rho(x, y)$.

- We can use the inverse Fourier transform (IFT):

$$\rho(x,y) = \iint S(k_x,k_y) e^{i2\pi(k_x x + k_y y)} dk_x dk_y$$

Image Formation

k-space

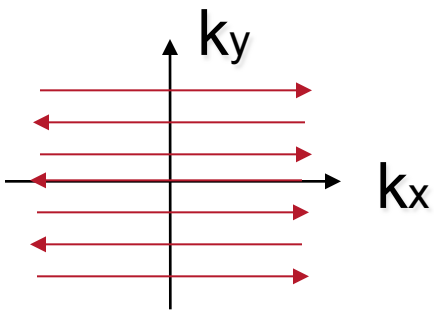

IFT

FT

Image space

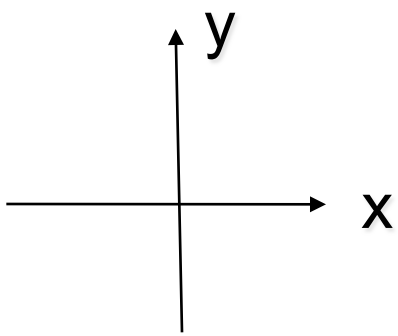

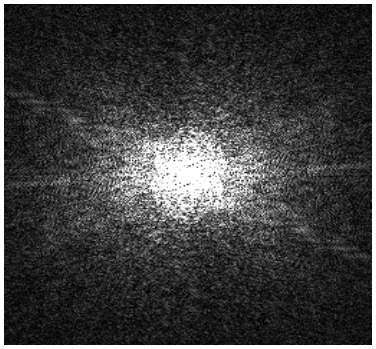

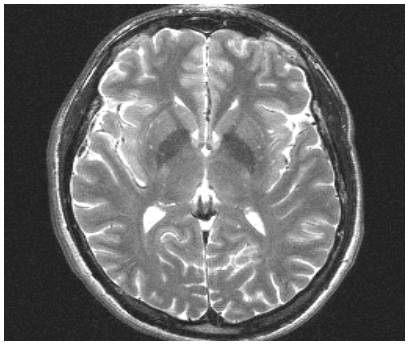

Spatial and Temporal Resolution

- When designing an fMRI experiment one must balance the need for adequate spatial resolution with that of adequate temporal resolution.

- The temporal resolution determines our ability to separate brain events in time.
  - In fMRI the temporal resolution is determined by how quickly each individual image is acquired (TR).

- The spatial resolution determines our ability to distinguish changes in an image across different spatial locations.

Terminology

Structural (T1) images:

- High spatial resolution
- No temporal information
- Can distinguish different types of tissue

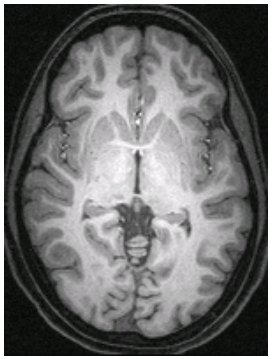

Functional (T2*) images:

- Lower spatial resolution
- Higher temporal resolution
- Can relate changes in signal to an experimental task

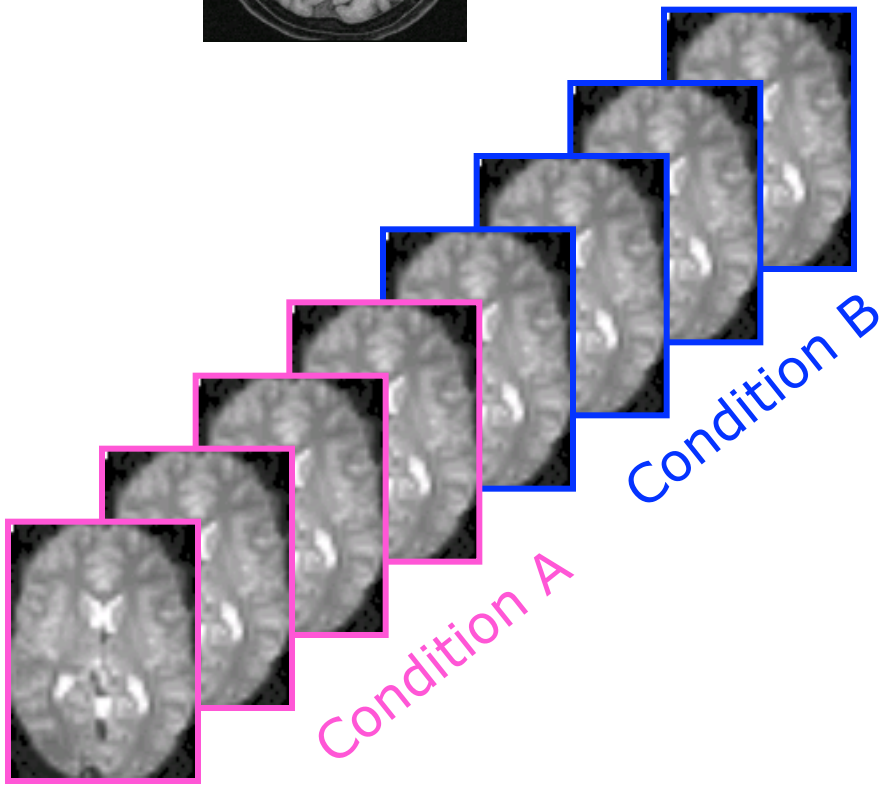

Terminology

Coronal

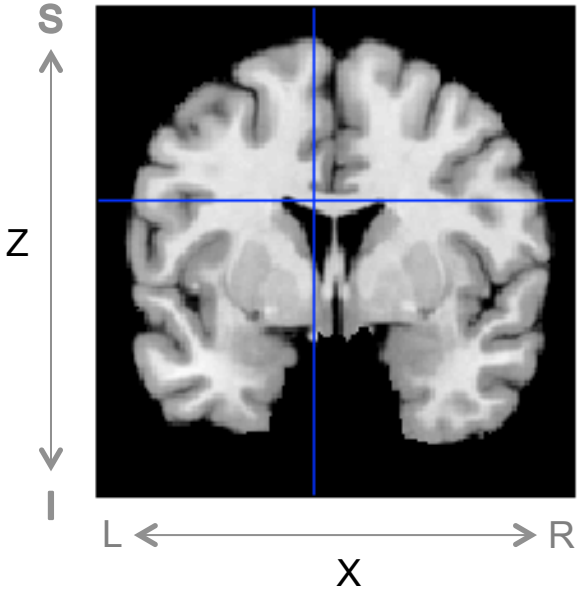

Saggital

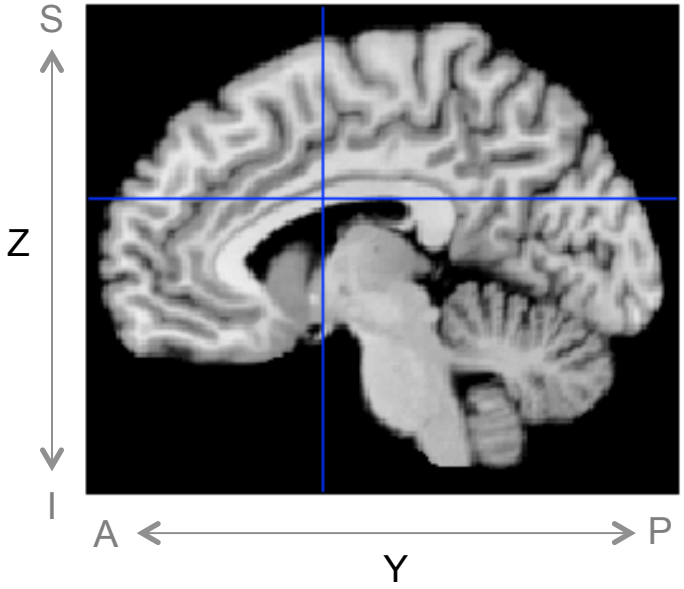

Axial

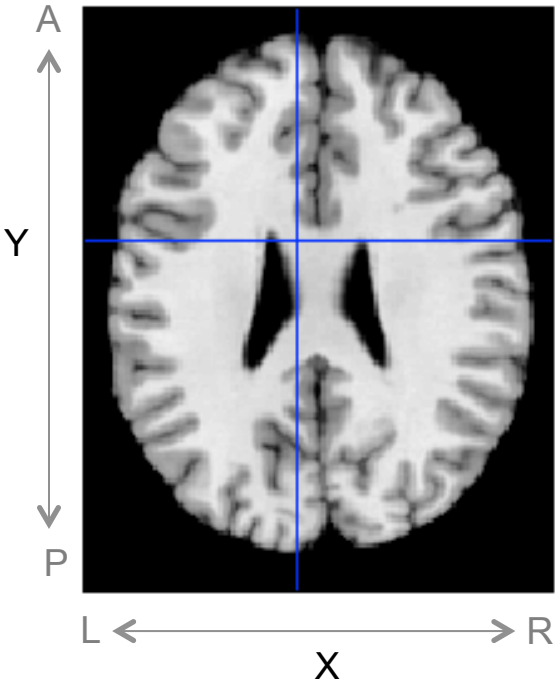

Terminology

- MRI images are typically acquired in axial slices - one at a time.
- This can be performed in either a sequential or interleaved manner.
- Together the slices make up a 3 dimensional brain volume.

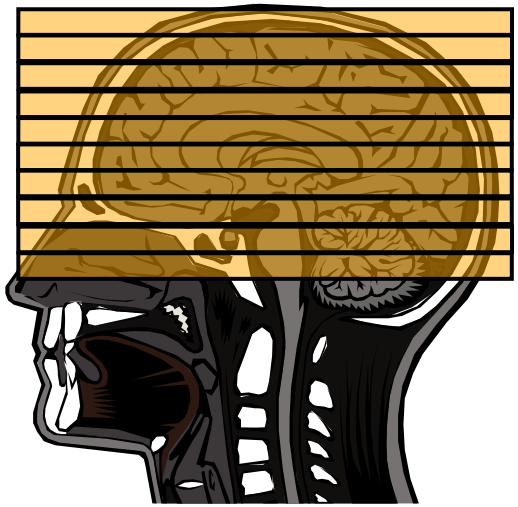

Terminology

Field of View (FOV)
e.g., 192 mm

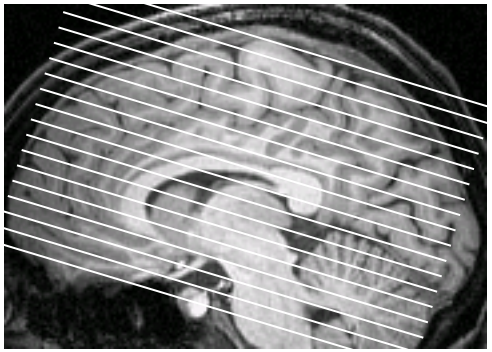

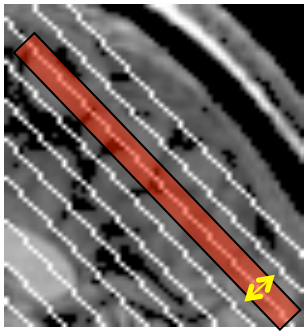

Slice thickness
e.g., 3 mm

Matrix Size
e.g., 64 x 64

In-plane resolution
192 mm / 64
= 3 mm

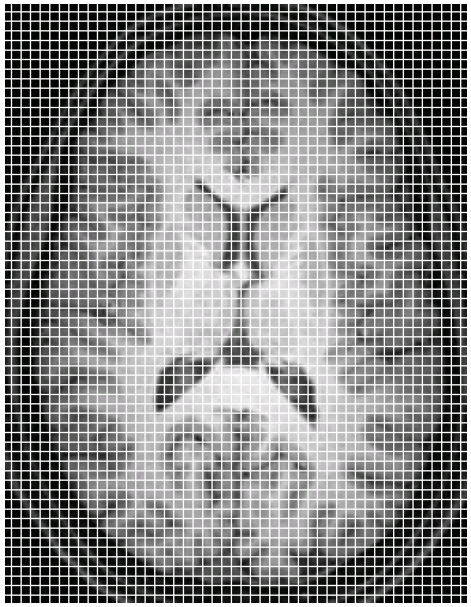

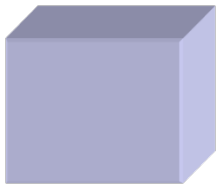

Voxel Size

Terminology

Subjects

Sessions

Runs

A single run

Volume

Slices

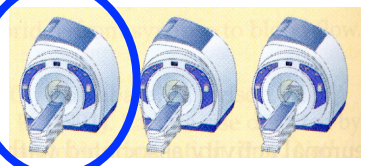

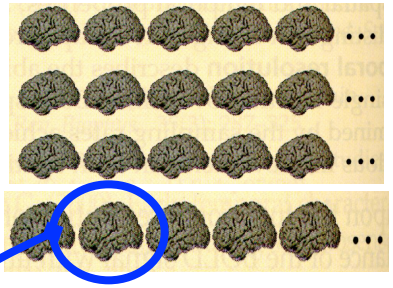

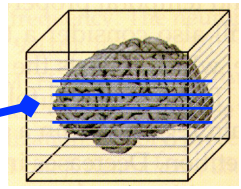

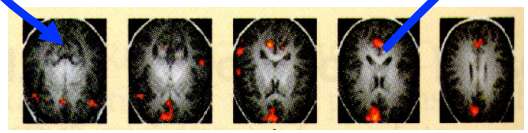

TR = repetition time
time required to scan one volume

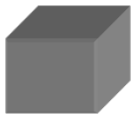

Voxel

BOLD fMRI

- The most common approach towards fMRI uses the Blood Oxygenation Level Dependent (BOLD) contrast.

- BOLD fMRI measures the ratio of oxygenated to deoxygenated hemoglobin in the blood.

- It is important to note that BOLD fMRI doesn't measure neuronal activity directly, instead it measures the metabolic demands (oxygen consumption) of active neurons.

HRF

The hemodynamic response function (HRF) represents changes in the fMRI signal triggered by neuronal activity.

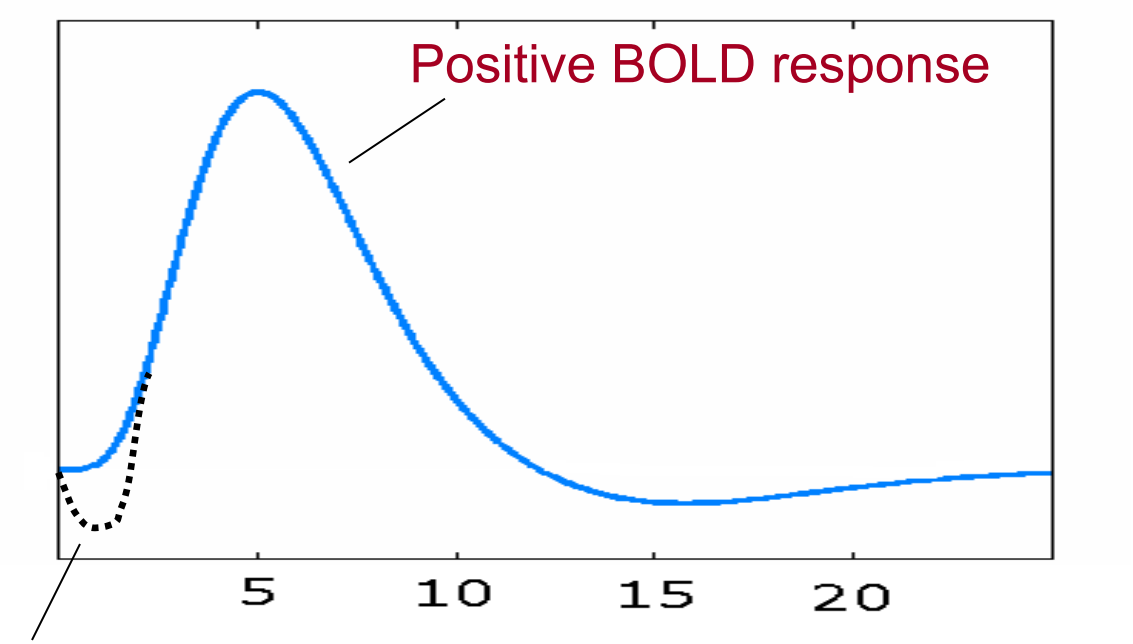

Properties of the HRF

- Magnitude of signal changes is quite small
  - 0.1 to 5%
  - Hard to see in individual images
- Response is delayed and quite slow
  - Extracting temporal information is tricky, but possible
  - Even short events have a rather long response
- Exact shape of the response has been shown to vary across subjects and regions.

BOLD Response

- The evoked BOLD response in fMRI is a complex, nonlinear function of the results of neuronal and vascular changes.
- The shape of the response depends both on the applied stimulus and the hemodynamic response to neuronal events.
- There exist a number of methods for modeling the BOLD response and the underlying HRF.

Models

- Nonlinear physiological-based models
  - Consists of a set of ordinary differential equations that model changes in blood volume, blood inflow, deoxyhemoglobin and flow inducing signal and describe how these changes impact the observed BOLD response.

- Linear time invariant (LTI) system
  - Assumes that the neuronal activity (based on task manipulations) constitutes the input, or impulse, and the HRF is the impulse response function.

LTI System

- The relationship between stimuli and the BOLD response is often modeled using a linear time invariant (LTI) system.
  - Here the neuronal activity acts as the input or impulse and the HRF acts as the impulse response function.

- In this framework the BOLD response at time *t* is modeled as the convolution of a stimulus function *v(t)* and the hemodynamic response *h(t)*, that is,

$$x(t) = (v * h)(t)$$

Convolution Examples

Block Design

Event-Related

Experimental Stimulus Function

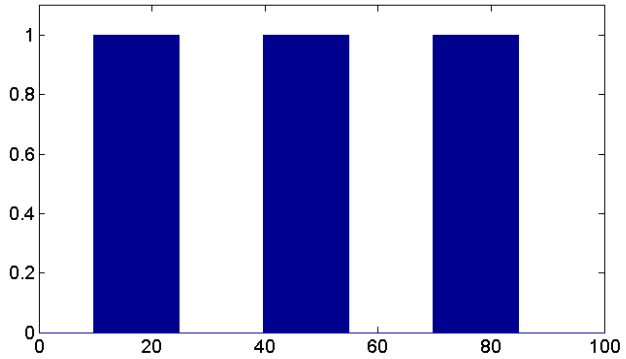

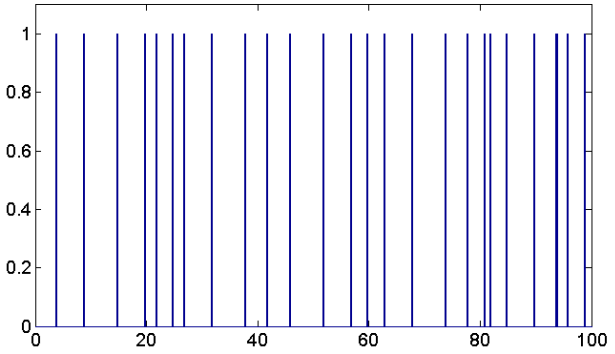

Hemodynamic Response Function

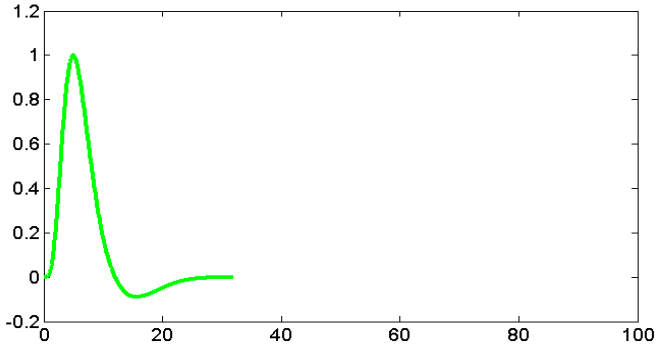

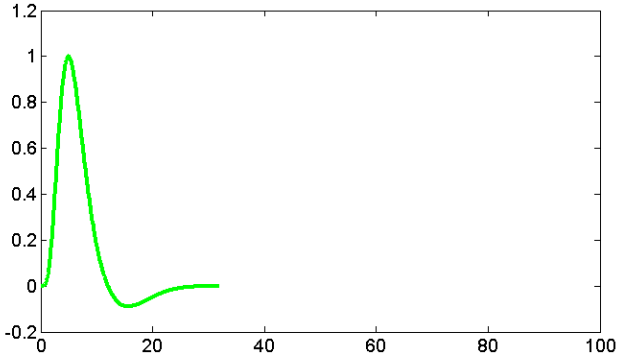

Predicted Response

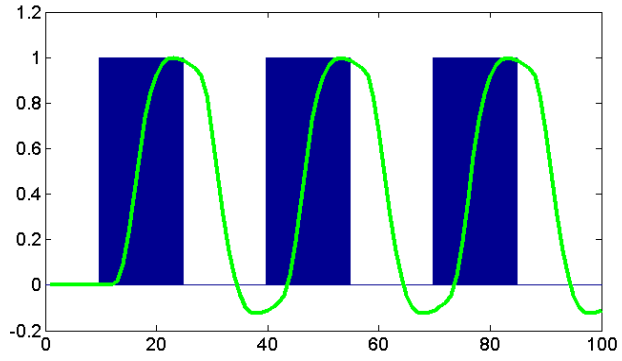

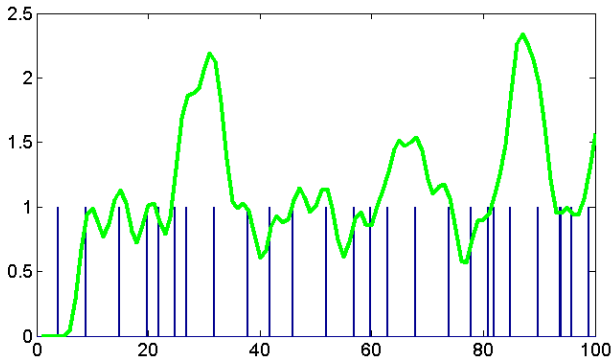

fMRI Noise

- The measured fMRI signal is corrupted by random noise and various nuisance components that arise due to hardware reasons and the subjects themselves.

- Sources of noise:
  - Thermal motion of free electrons in the system.
  - Patient movement during the experiment.
  - Physiological effects, such as the subject's heartbeat and respiration.
  - Low frequency signal drift.

fMRI Noise

- Some of these noise components can be removed prior to analysis, while others can be included as components in subsequent models.

- It is difficult to remove all sources of noise and therefore significant autocorrelation will be present in the signal.

- Characteristics of the noise:
  - “1/f” in frequency domain
  - Nearby time-points exhibit positive correlation

END OF PART I

Data Processing Pipeline

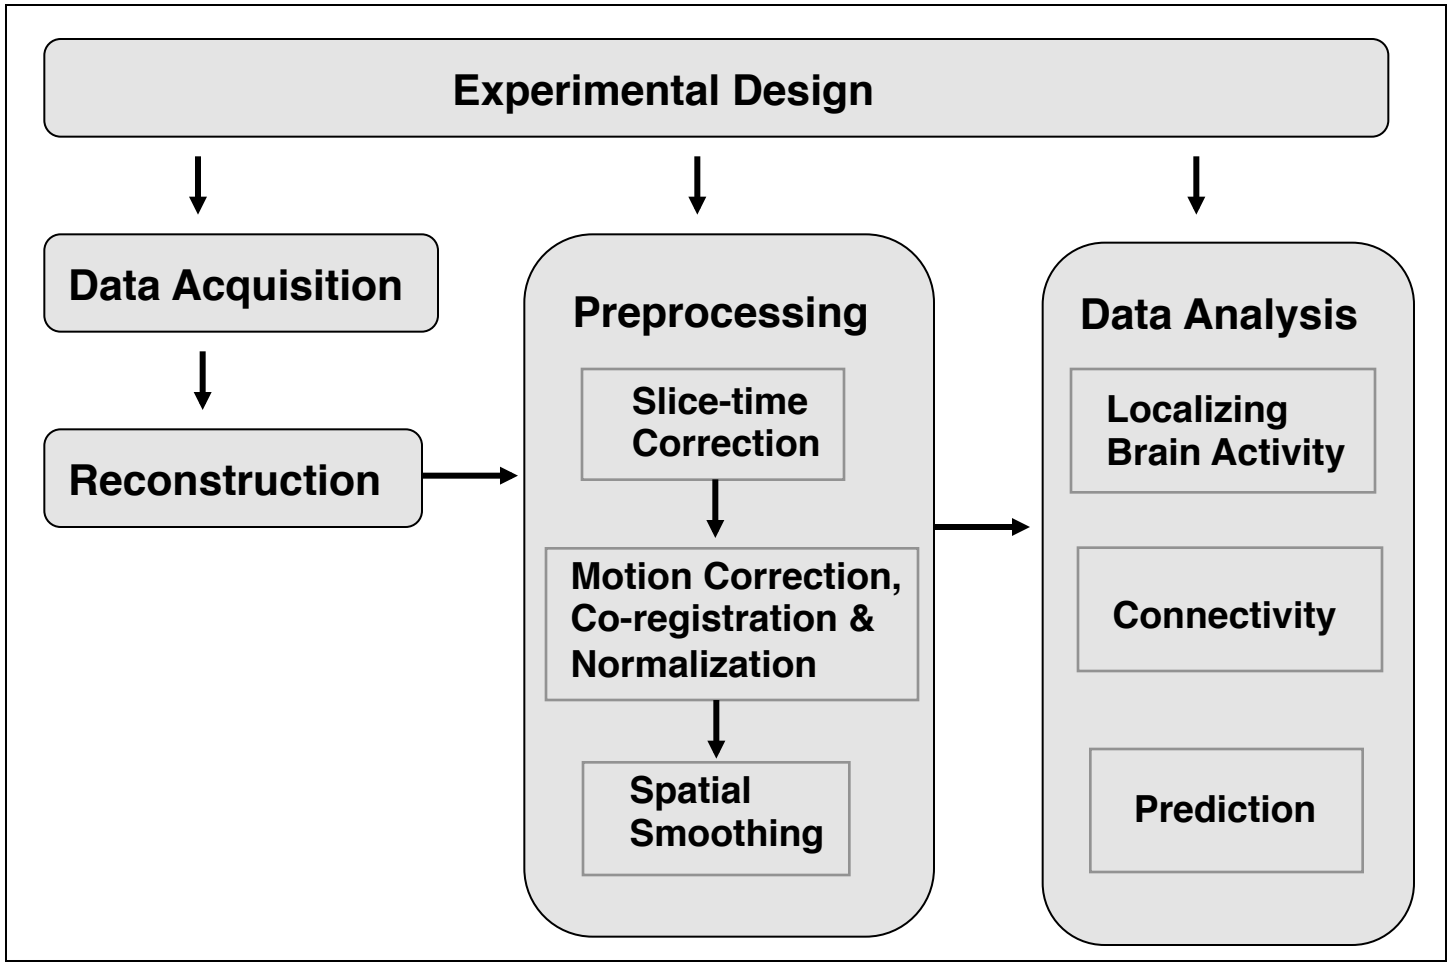

Experimental Design

- Block design: Similar events are grouped

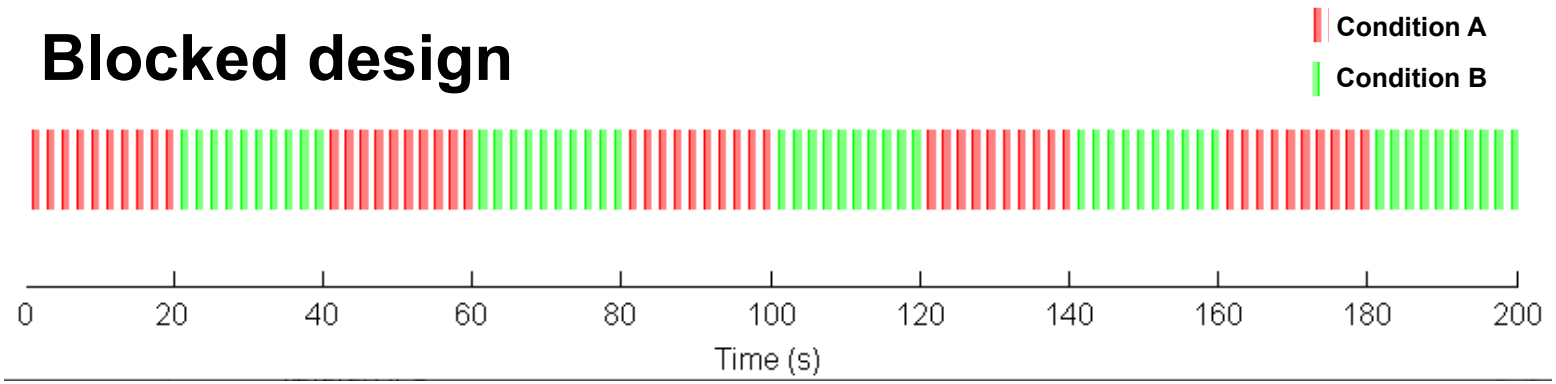

- High statistical power to detect activation and robust to uncertainties in the shape of HRF.

- Can’t directly estimate features of the HRF.

Experimental Design

- Event-related design: Events are mixed

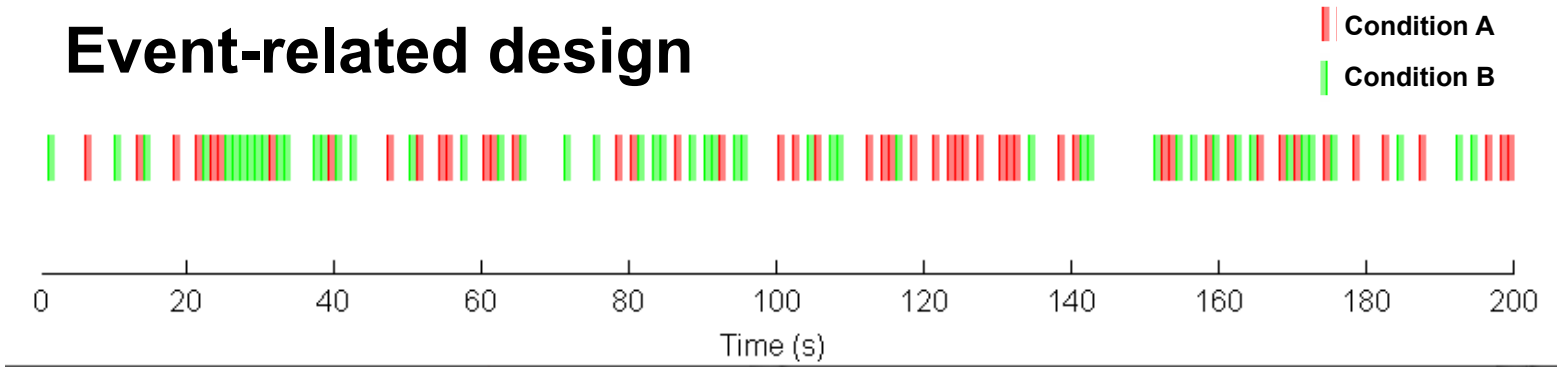

- Allows for the estimation of features of the HRF.

- Decreased power to detect activation.

Pre-processing

- Prior to analysis, fMRI data undergoes a series of preprocessing steps aimed at identifying and removing artifacts and validating model assumptions.

- The goals of preprocessing are
  - To minimize the influence of data acquisition and physiological artifacts;
  - To check statistical assumptions and transform the data to meet assumptions;
  - To standardize the locations of brain regions across subjects to achieve validity and sensitivity in group analysis.

Pre-processing Pipeline

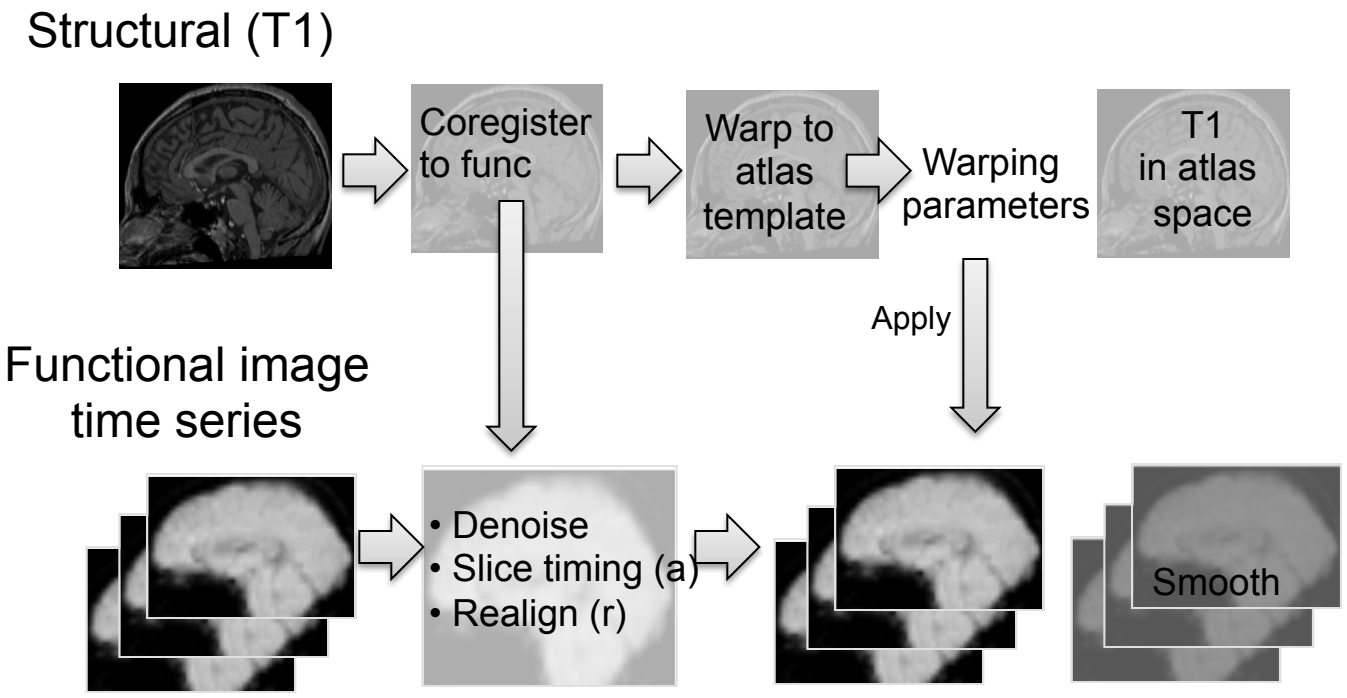

Preprocessing is performed both on the fMRI data and structural scans collected prior to the experiment.

Pre-processing Steps

- Visualization and Artifact Removal
- Slice Time Correction
- Motion Correction
- Physiological Corrections
- Co-registration
- Normalization
- Spatial Filtering
- Temporal Filtering

Slice Time Correction

- We often sample multiple slices of the brain during each individual repetition time (TR) to construct a brain volume.

- Typically each slice is sampled at a slightly different time points (i.e., 2D imaging; not 3D).

- Slice time correction shifts each voxel's time series so that they all appear to have been sampled simultaneously.

Slice Time Correction

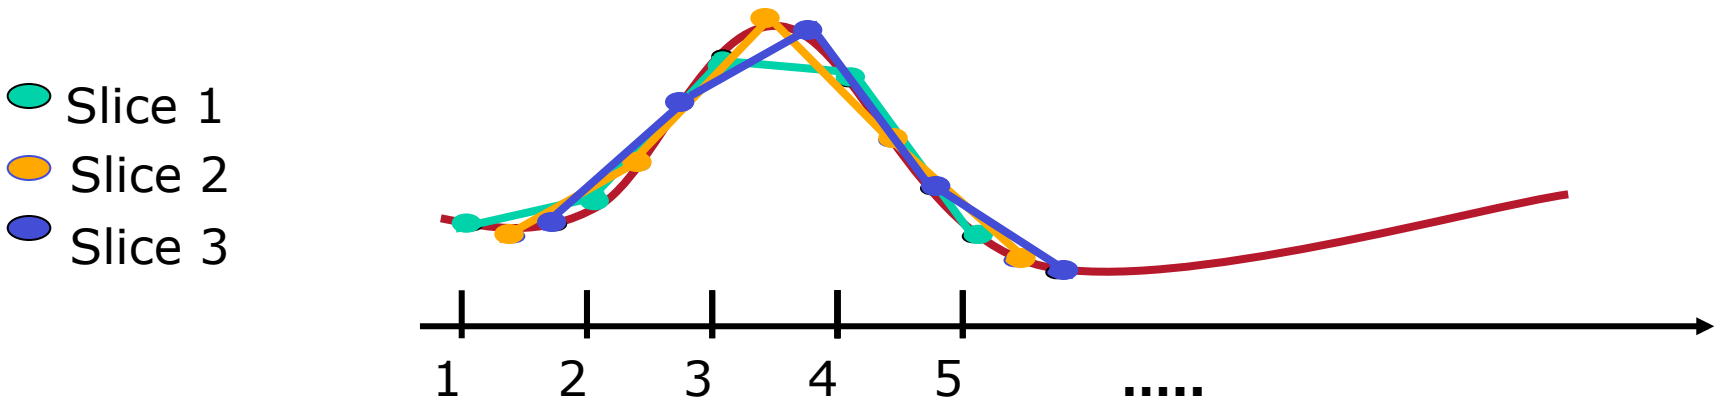

Can be corrected using temporal interpolation.

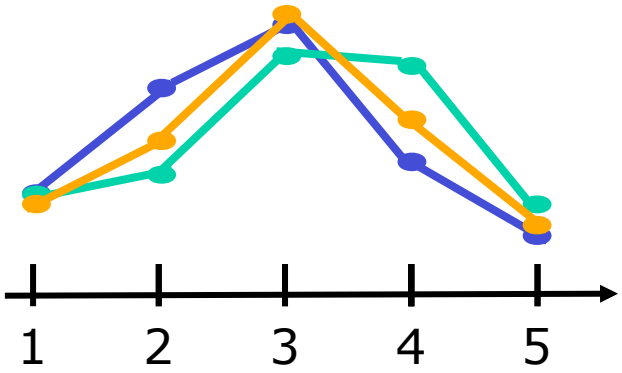

Head Motion

- Very small movements of the head during an experiment can be a major source of error if not treated correctly.
- When analyzing the time series associated with a voxel, we assume that it depicts the same region of the brain at every time point
  - Head motion may make this assumption incorrect.
- Can be corrected using a rigid body transformation.

Transformations

Linear transformations

- Rigid body (6 DOF) – translation and rotation
- Similarity (7 DOF) – translation, rotation and a single global scaling
- Affine (12 DOF) – translation, rotation, scaling and shearing.

Warping

Transformations where the equations relating the coordinates of the images are non-linear.

Illustration

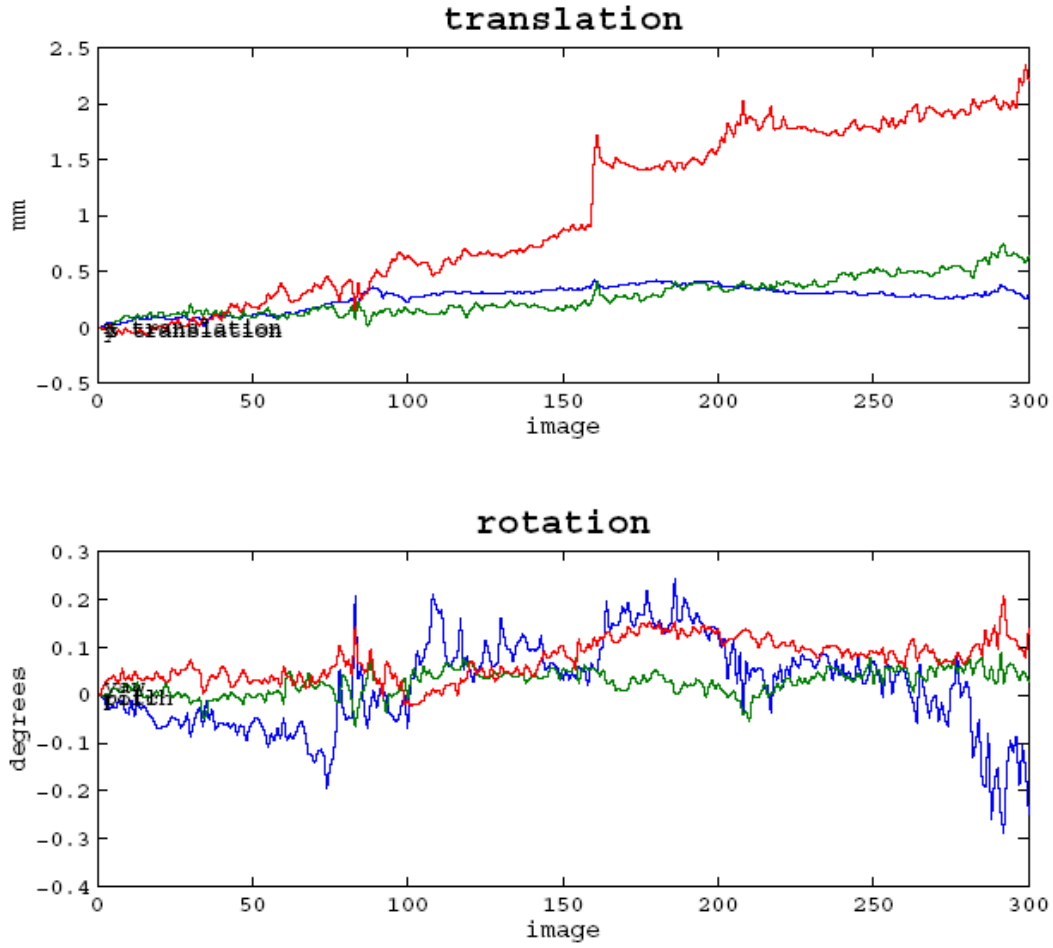

Coregistration

- Next, a structural MRI collected in the beginning of the session is registered to the fMRI images in a process referred to as coregistration.
  - Allows one to visualize single-subject task activations overlaid on the individual's anatomical information.
  - Simplifies later transformation of the fMRI images to a standard coordinate system.

- Use at least an affine transformation to perform coregistration and the mutual information cost function.

Normalization

- All brains are different. The brain size of two subjects can differ in size by up to 30%.
- There may also be substantial variation in the shapes of the brain.
- Normalization allows one to stretch, squeeze and warp each brain so that it is the same as some standard brain.
- Affine transformation (12 DOF) or warping.

Brain Atlases

Talairach space

- Talairach and Tournoux (1988)
- Based on single subject (A cadaver of a 60 year old female)

Montreal Neurological Institute

- Combination of many MRI scans on normal controls. All right-handed subjects.

The structural MR image used in the coregistration procedure is warped onto a template image.

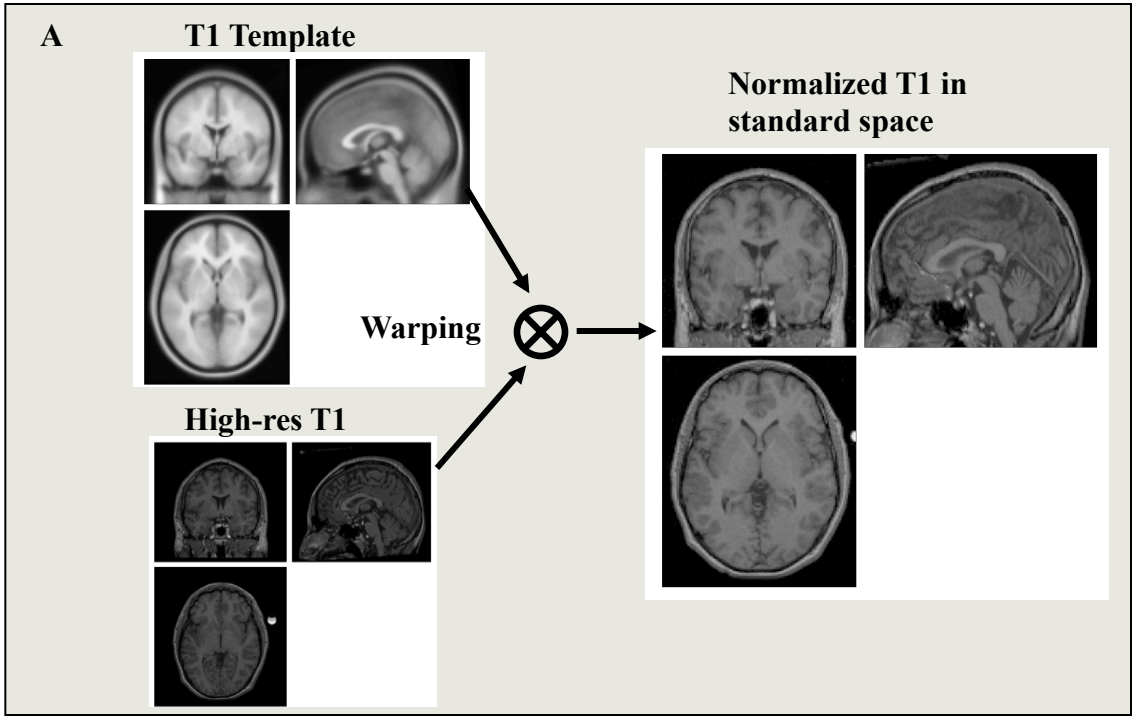

Spatial Filtering

- In fMRI it is common to spatially smooth the acquired data prior to statistical analysis.

- Can increase signal-to-noise, validate distributional assumptions and remove artifacts.

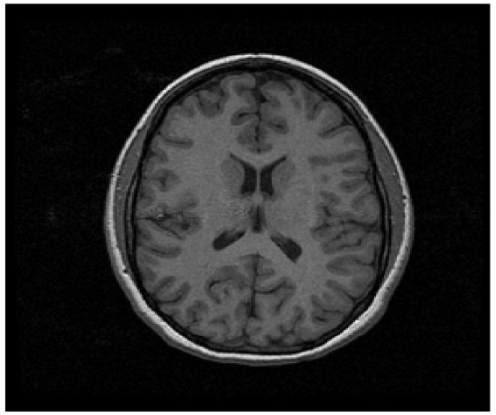

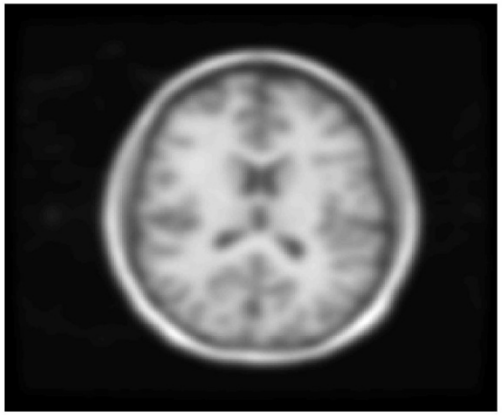

Statistical Analysis

- There are multiple goals in the statistical analysis of fMRI data.

- They include:
  - localizing brain areas activated by the task;
  - determining networks corresponding to brain function; and
  - making predictions about psychological or disease states.

Localizing Activation

1. Construct a model for each voxel of the brain.
   - "Massive univariate approach"
   - Regression models (GLM) commonly used.

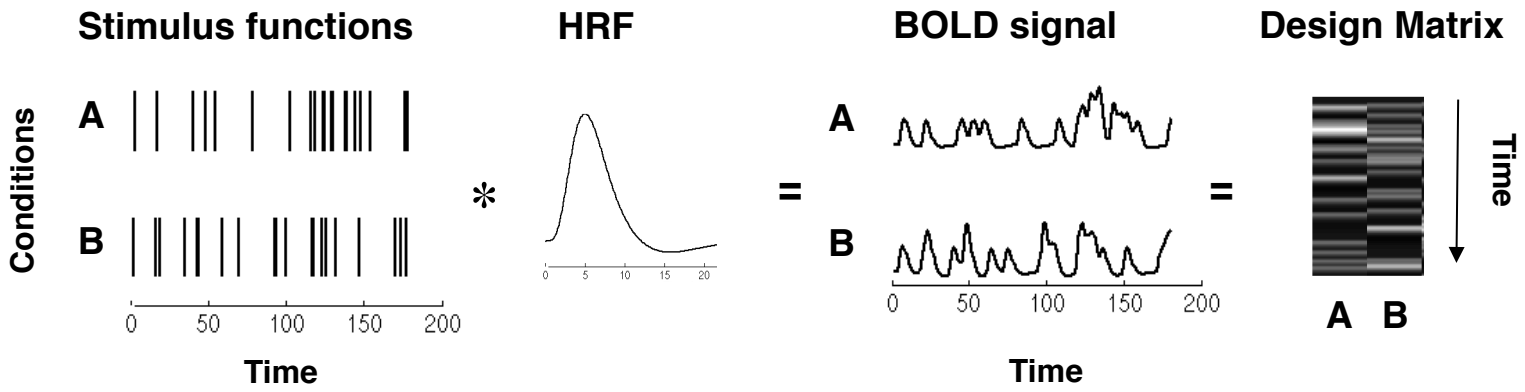

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \qquad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{V})$$

Localizing Activation

2. Perform a statistical test to determine whether task related activation is present in the voxel.

$$H_0 : \mathbf{c}^T \boldsymbol{\beta} = 0$$

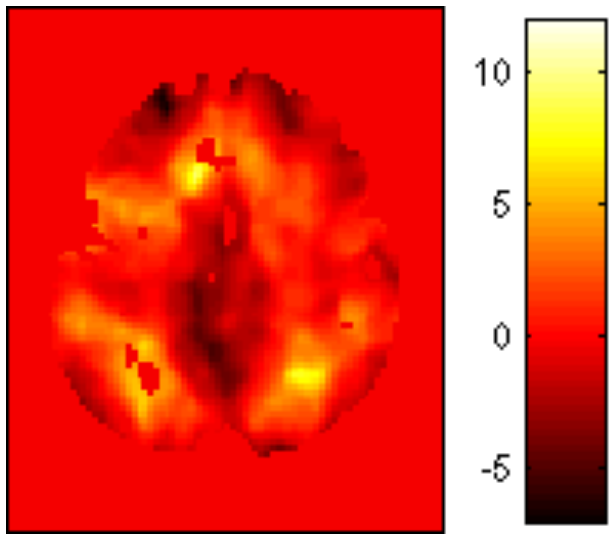

Statistical image:
Map of t-tests across all voxels (a.k.a t-map).

Localizing Activation

3. Choose an appropriate threshold for determining statistical significance.

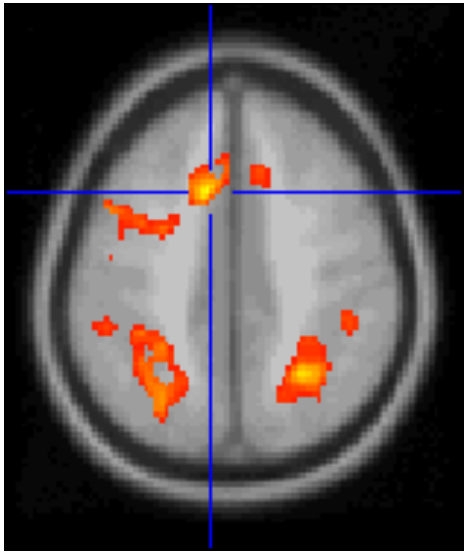

Statistical parametric map: Each significant voxel is color-coded according to the size of its p-value.

Multiple Comparisons

- Which of 100,000 voxels are significant?
  - α=0.05 $\Rightarrow$ 5,000 false positive voxels

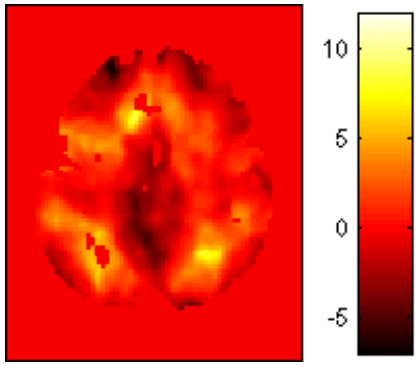

- Choosing a threshold is a balance between sensitivity (true positive rate) and specificity (true negative rate).

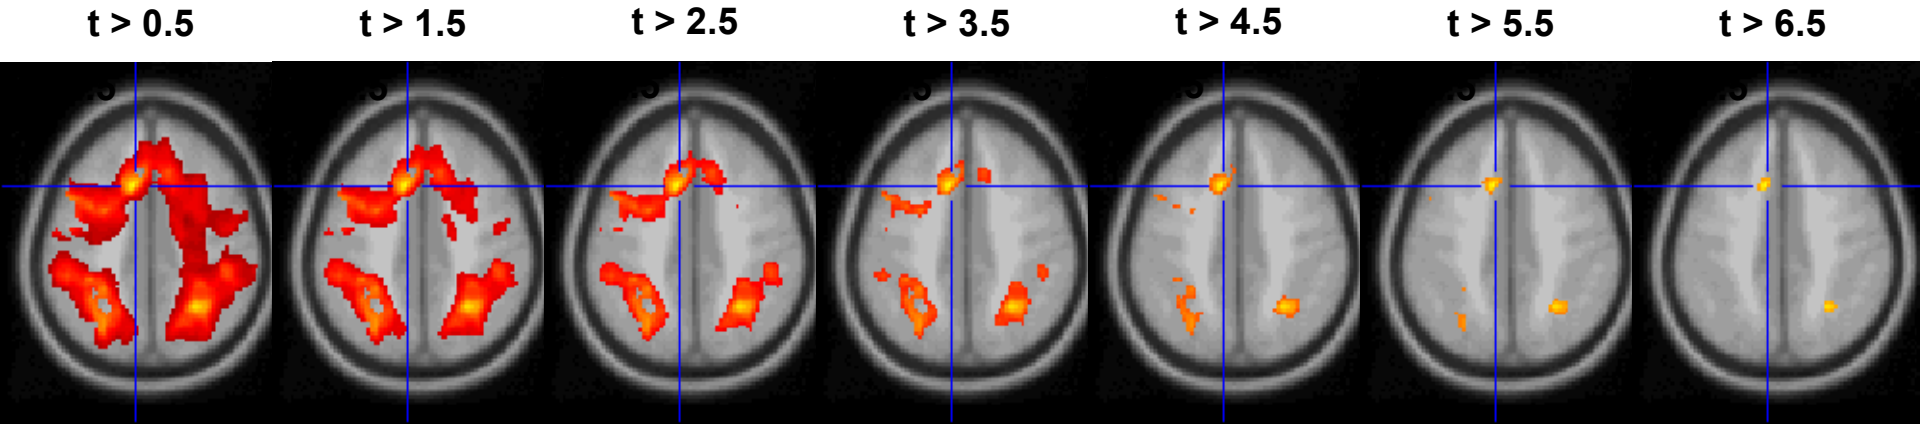

General Linear Model

A standard GLM can be written:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \qquad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{V})$$

where

$$\begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{bmatrix} = \begin{bmatrix} 1 & X_{11} & \cdots & X_{1p} \\ 1 & X_{21} & \cdots & X_{2p} \\ \vdots & \vdots & & \vdots \\ 1 & X_{np} & \cdots & X_{np} \end{bmatrix} \times \begin{bmatrix} \beta_0 \\ \beta_1 \\ \vdots \\ \beta_p \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{bmatrix}$$

fMRI Data

Design matrix

Model parameters

Noise

V is the covariance matrix whose format depends on the noise model.

The quality of the model depends on our choice of X and V.

Model Building

- Proper construction of the design matrix is critical for proper use of the GLM.

- It can be complicated by the following properties of the BOLD response:
  - It includes low-frequency noise and artifacts related to head movement and cardiopulmonary-induced brain movement.
  - The neural response shape may not be known.
  - The hemodynamic response varies in shape across the brain.

Nuisance Covariates

- Often model factors associated with known sources of variability, but that are not related to the experimental hypothesis, need to be included in the GLM.

- Examples of possible 'nuisance regressors':
  - Signal drift
  - Physiological (e.g., respiration) artifacts
  - Head motion, e.g. six regressors comprising of three translations and three rotations.
    - Sometimes transformations of the six regressors also included.

Drift

- Slow changes in voxel intensity over time (low-frequency noise) is present in the fMRI signal.
- Scanner instabilities and not motion or physiological noise may be the main cause of the drift, as drift has been seen in cadavers.
- Need to include drift parameters in our models.
  - Use polynomial basis or discrete cosine basis

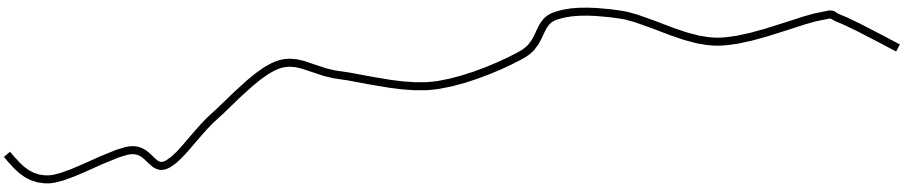

Physiological Noise

- Respiration and heart beat give rise to high-frequency noise.
- It can potentially be modeled, but if the TR is too low there will be problems with aliasing.
  - Sampling rate must be at least twice as big as the frequency of the curve you seek to model.
- Hence, this type of noise is difficult to remove and is often left in the data giving rise to temporal autocorrelations.

Motion

- Subject motion during the experiment can also give rise to serious problems.
- Typically motion correction is performed in the pre-processing stages of the analysis.
- However, ‘spin-history’ artifacts may remain that cannot be removed.
  - This is caused by through-plane motion.
  - Head motion parameters often included in GLM.

Head Motion

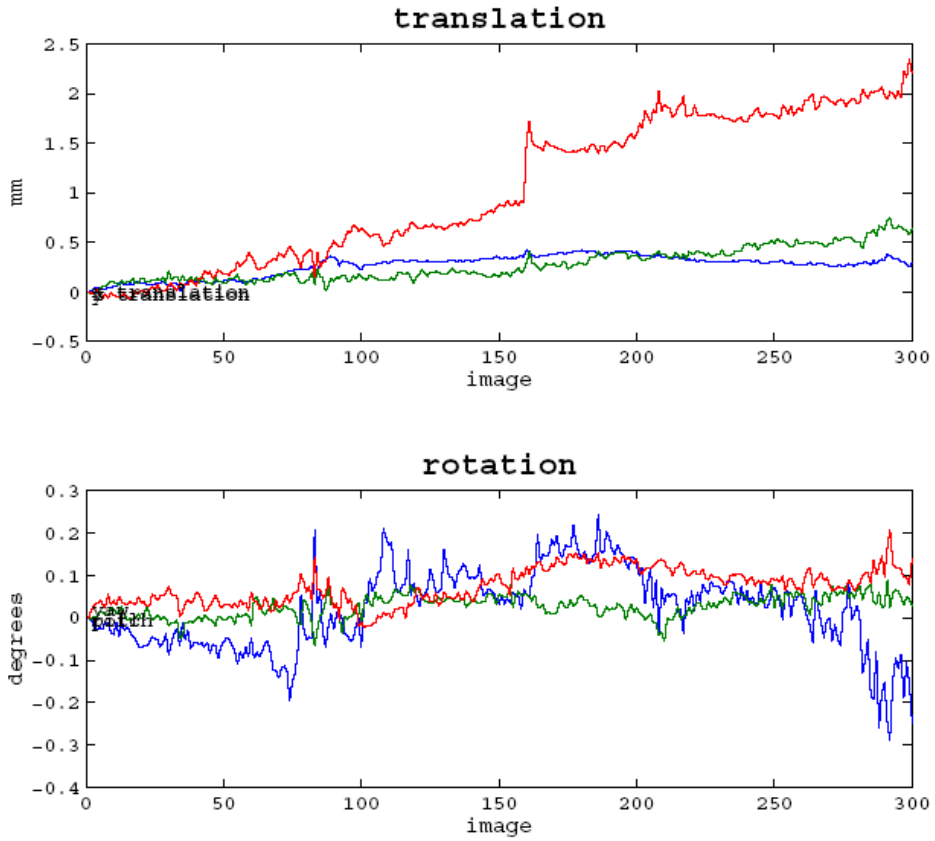
translation
mm
2.5
2
1.5
1
0.5
0
-0.5
0
50
100
150
200
250
300
image
translation
rotation
degrees
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
0
50
100
150
200
250
300
image

Head Motion Example

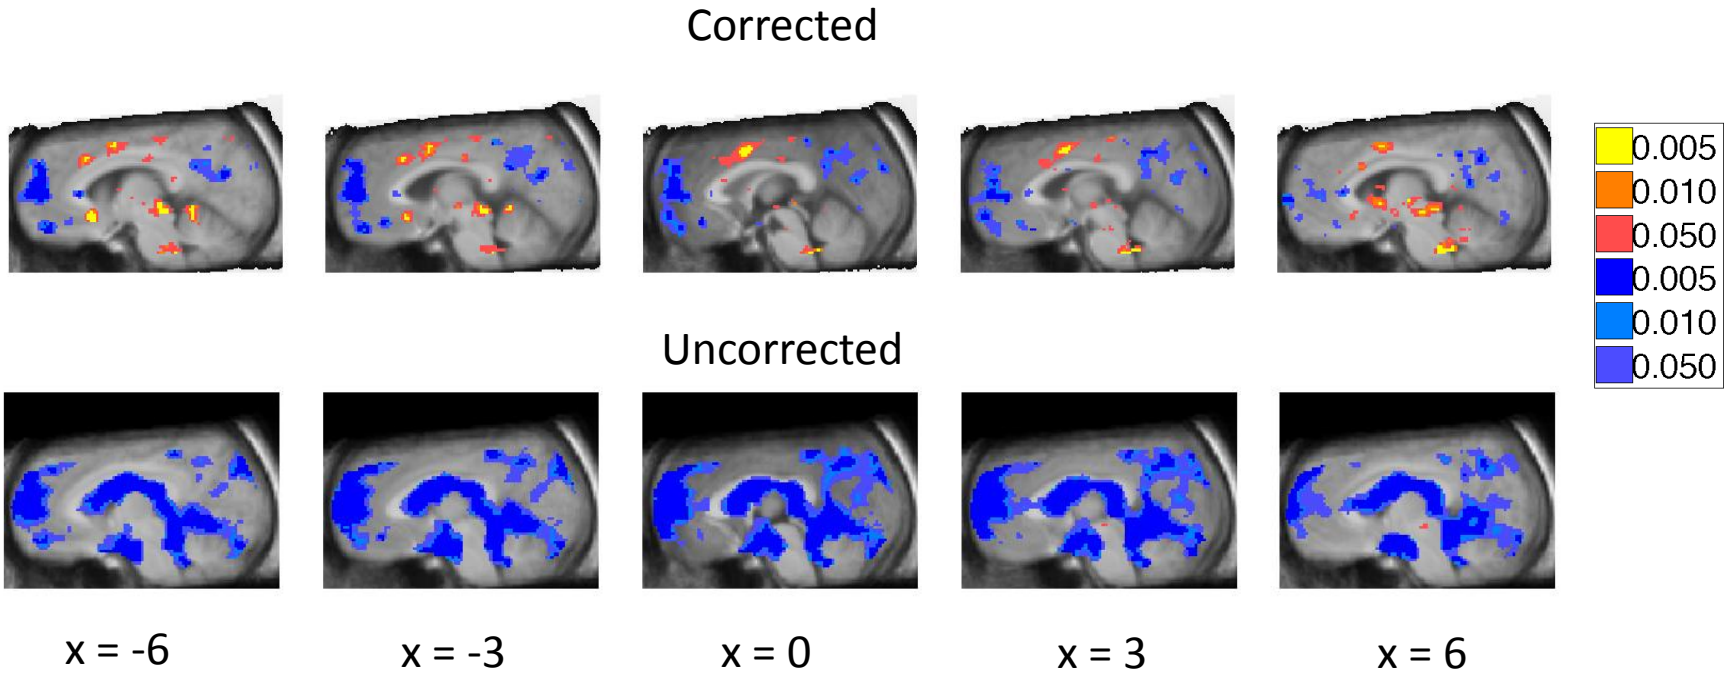

Task Related Signal

Indicator functions (Onsets) | Assumed HRF (Basis function) | Design Matrix (X^T) | Design Matrix (X)

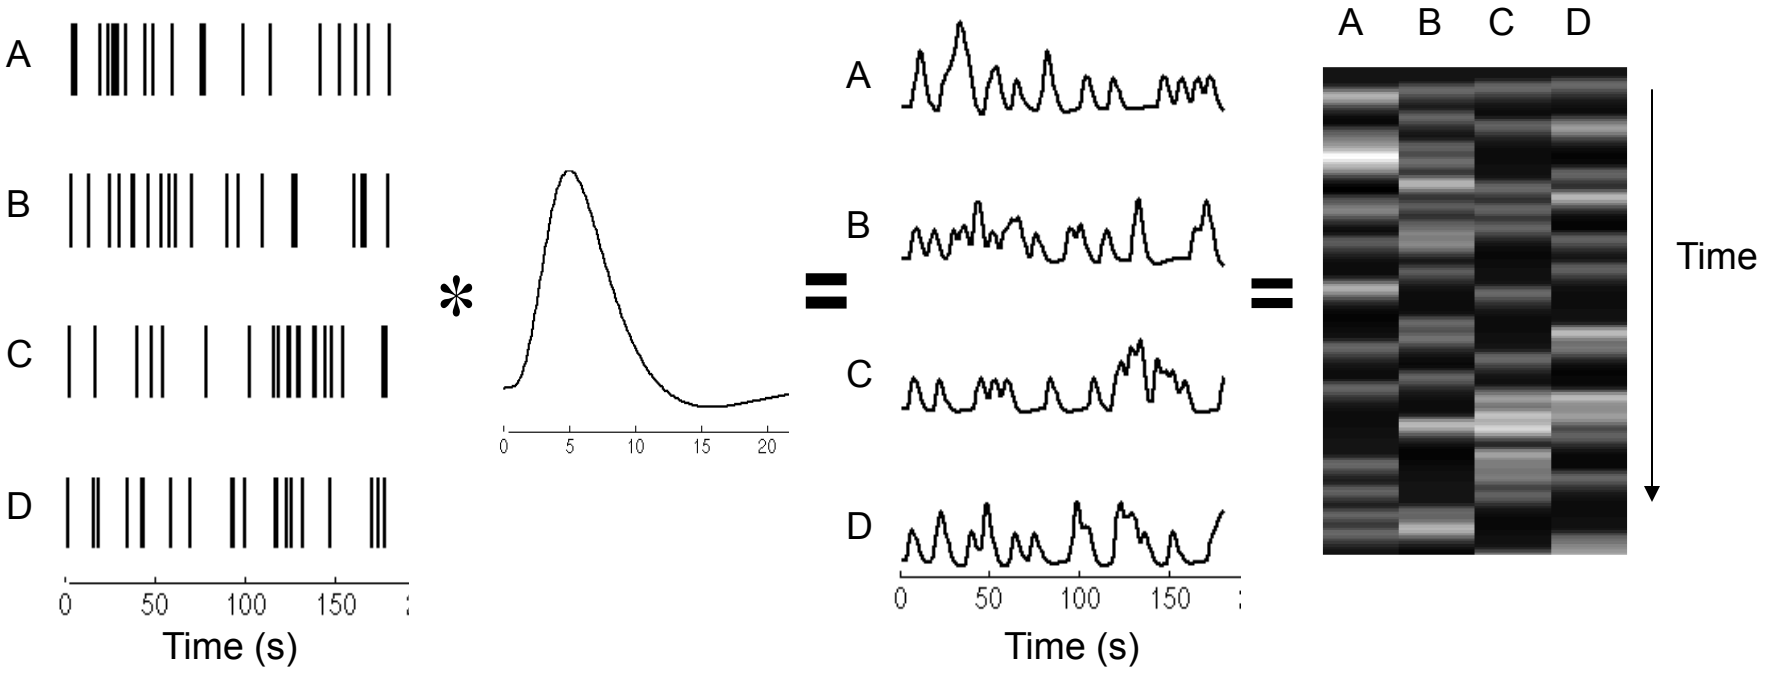

Assumptions:

Assume neural activity function is correct

Assume HRF is correct

Assume LTI system

Problems

The HRF shape depends both on the vasculature and the time course of neural activity.

Assuming a fixed HRF is usually not appropriate.

Checkerboard, n = 10

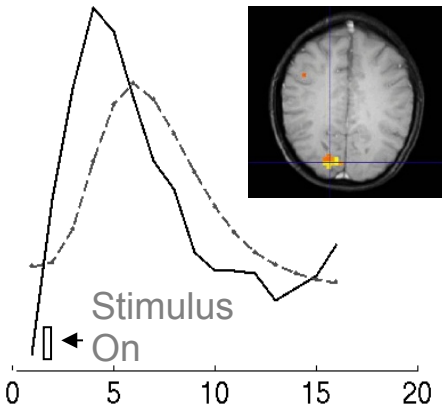

Thermal pain, n = 23

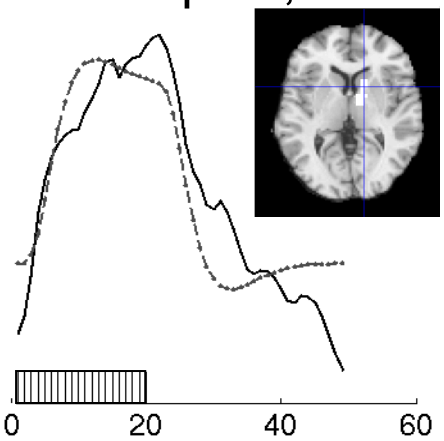

Aversive picture, n = 30

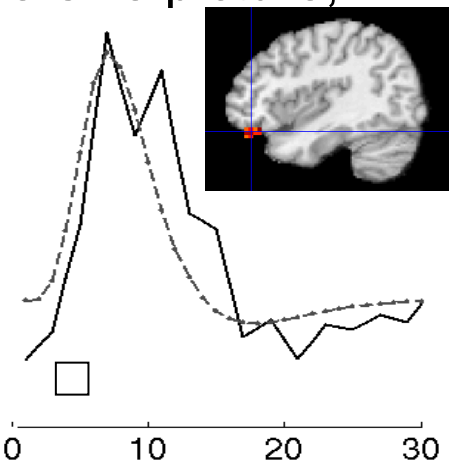

Aversive anticipation

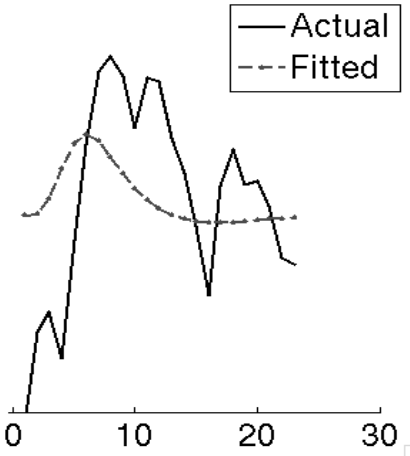

Temporal Basis Sets

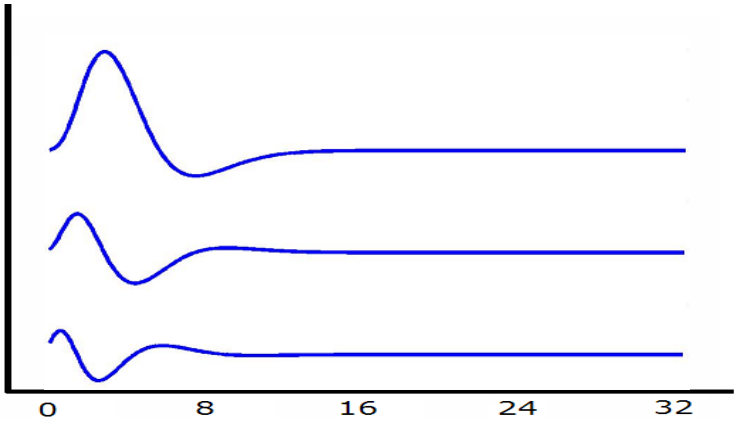

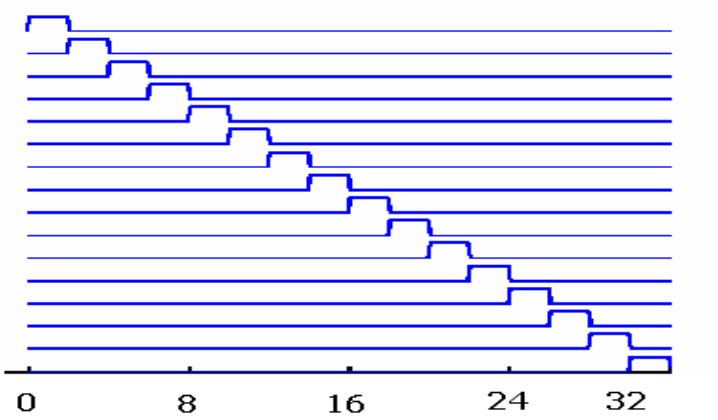

- Canonical HRF + Derivatives
  - Including the derivatives allows for a shift in delay and dispersion.
- Finite Impulse Response
  - The model estimates an HRF of arbitrary shape for each event type in each voxel

Basis sets

Model

Image of predictors

Data & Fitted

Single HRF

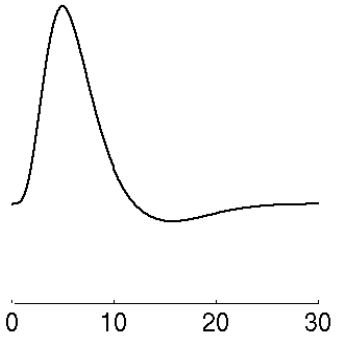

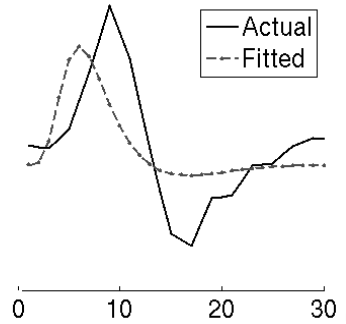

HRF + derivatives

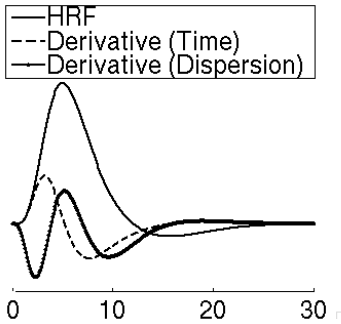

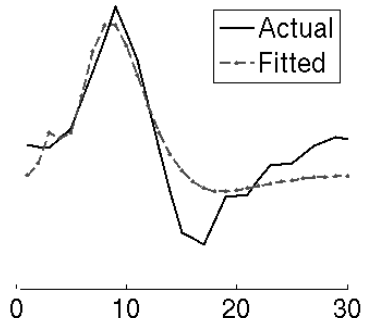

Finite Impulse Response (FIR)

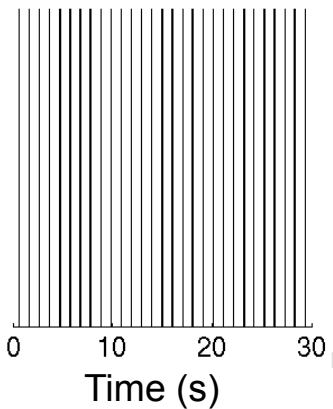

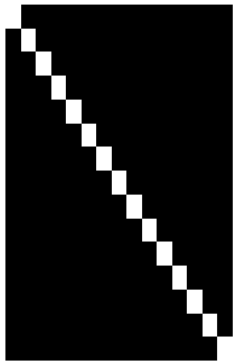

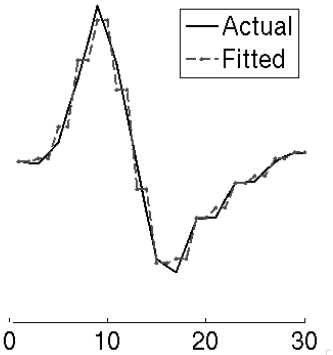

Inverse Logit Model

- Superposition of three inverse logit (sigmoid) functions.

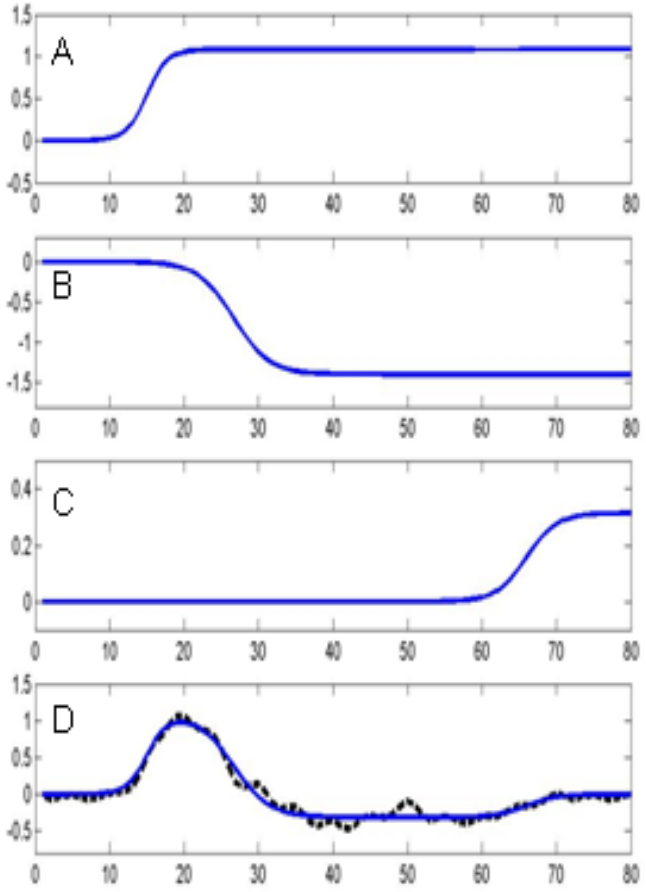

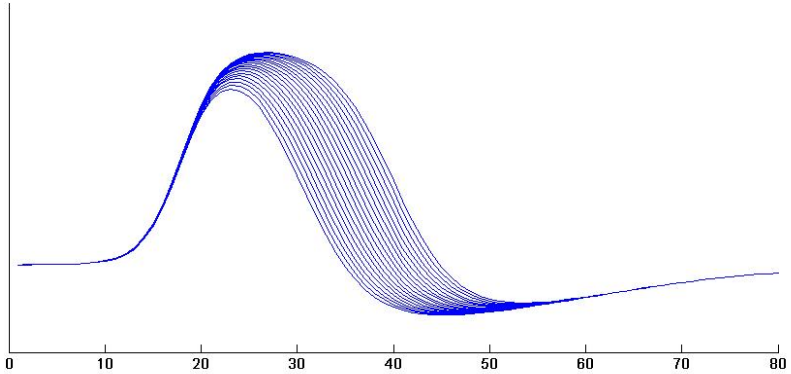

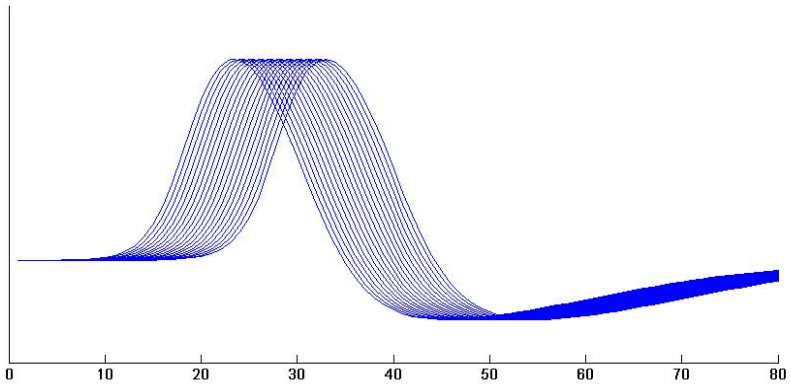

Example

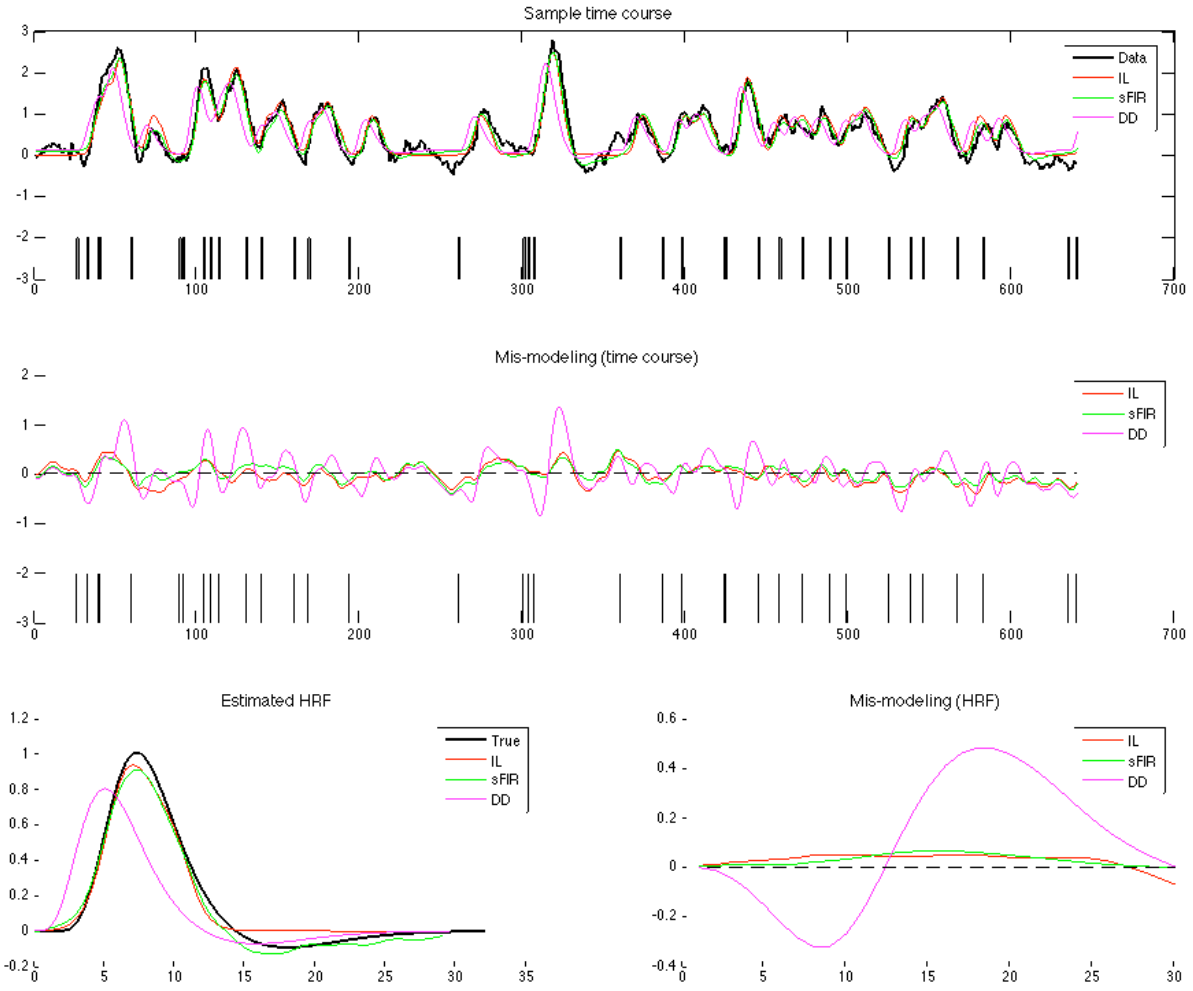
Sample time course
Data
IL
sFIR
DD
Mis-modeling (time course)
IL
sFIR
DD
Estimated HRF
True
IL
sFIR
DD
Mis-modeling (HRF)
IL
sFIR
DD

Noise Models

- What about **V**?
  - Different software packages use different models to handle fMRI noise.
  - White noise, AR or ARMA noise models.

- White noise is the easiest to work with, but is not particularly realistic.
  - Often used when performing group analysis

- The parameters of the AR model are easier to estimate than those for ARMA parameters.
  - Method of Moments vs. Maximum likelihood

Multi-level Model

- When performing group analysis we often use multi-level models. Often performed in two levels:
  - The first level deals with individual subjects.
  - The second level deals with groups of subjects.

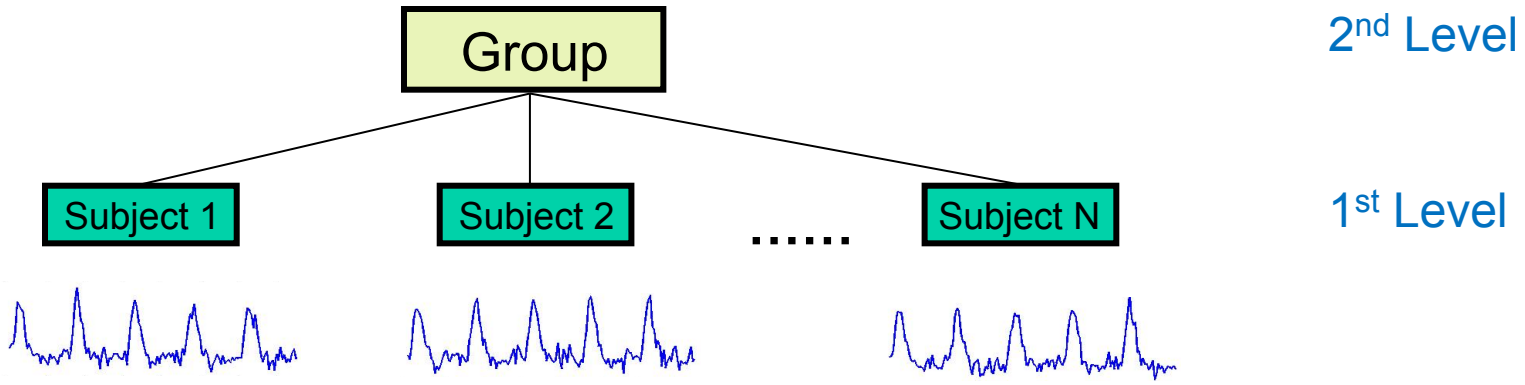

- All inference typically performed in the 'massive univariate' setting.

Group Analysis

- One can set up a multi-level GLM as follows:

$$\mathbf{Y} = \mathbf{X}\beta + \mathbf{e} \qquad \mathbf{e} \sim N(0, \mathbf{V})$$

$$\beta = \mathbf{X}_g \beta_g + \boldsymbol{\eta} \qquad \boldsymbol{\eta} \sim N(0, \mathbf{V}_g)$$

- Comments:
  - The model can be expanded to include more levels
  - The model parameters can be estimated using a mixed-effects model.
  - Inference based on group-level β.

Group Analysis

- Most software packages have implemented mixed-effects models.
  - They differ in which method and algorithm they apply.
- However, a simple non-iterative two-stage least squares approach is used in most fMRI analysis.
  - The Summary Statistics Approach.
- Results from individual subject are reused in the second level, reducing the computational burden of fitting a full model.
  - Certain assumptions are required.

Summary Statistics Approach

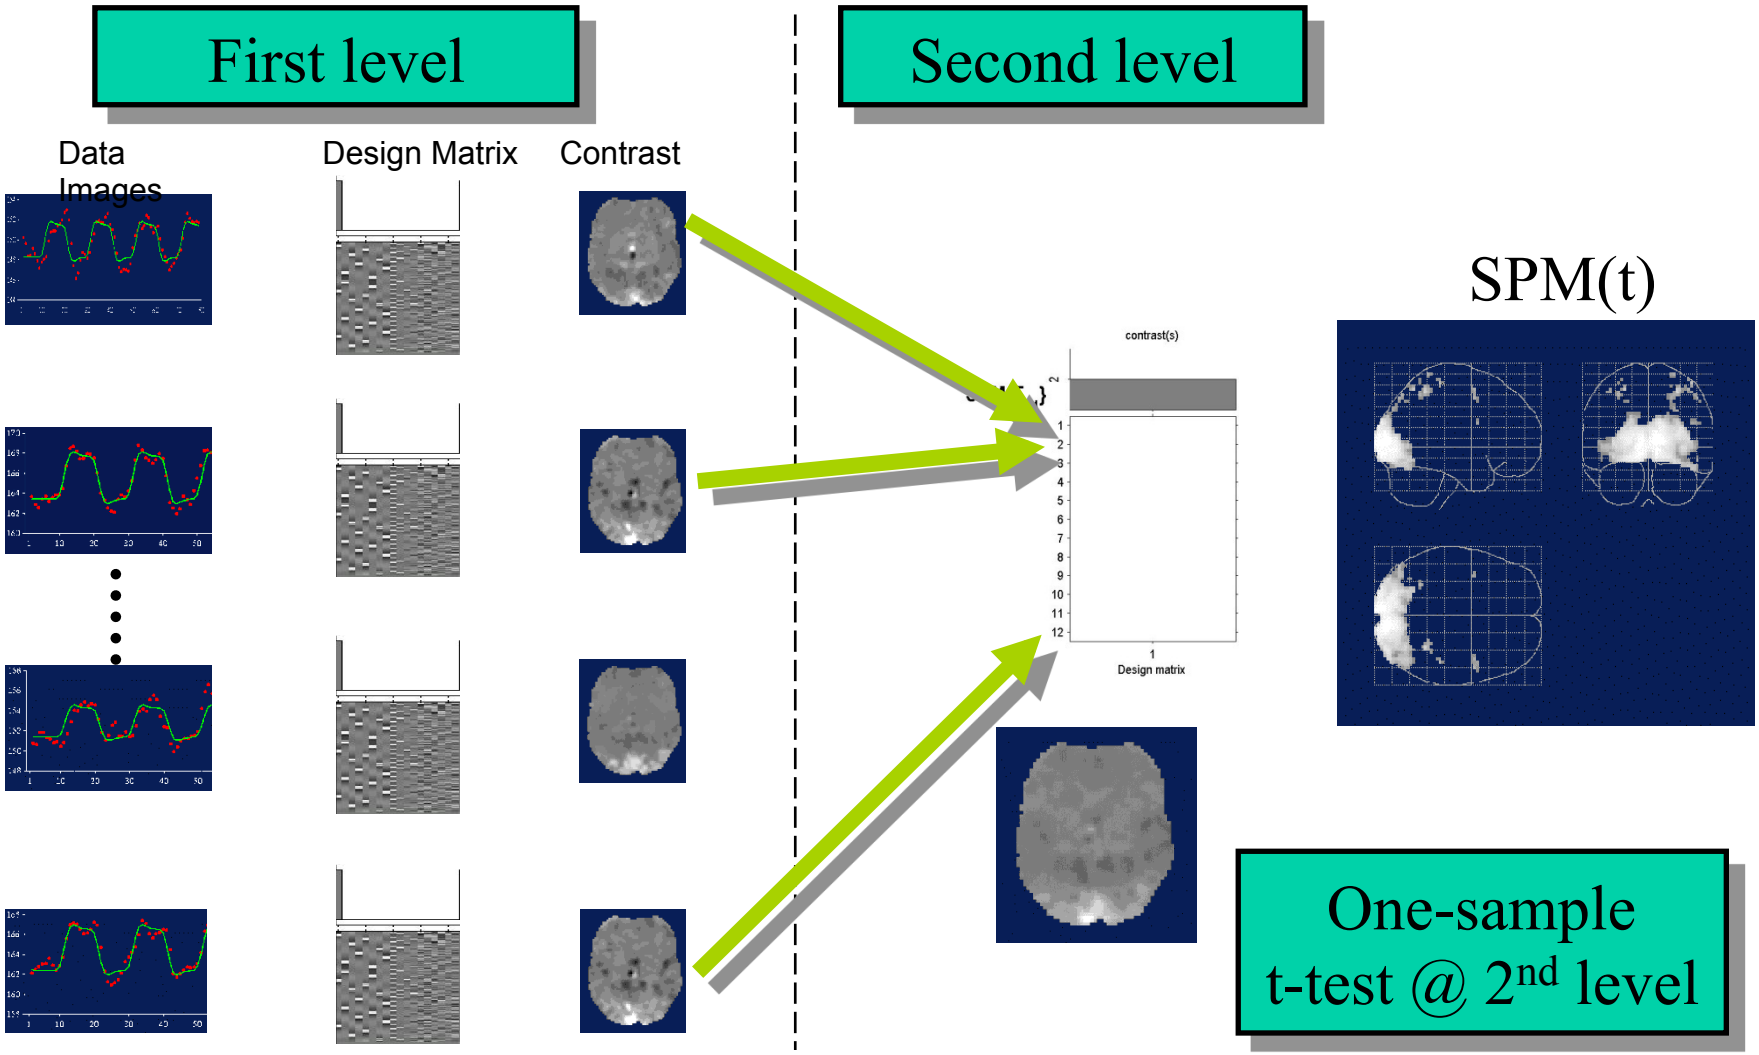

Group Analysis

- When using temporal basis sets at the first level it can be difficult to summarize the response with a single number, making group inference difficult.

- Here we can perform group analysis using
  - the “main” basis function,
  - all basis functions, or
  - re-parameterized fitted responses (Calhoun et al. (2004); Lindquist et al. (2009)).
    - Recreate the HRF and estimate the magnitude.
    - Use this information at the second level.

Interesting Questions

- Validation and greater understanding of the effects of the various pre-processing steps.
- Creation of an omnibus pre-processing method.
- Optimal experimental design
- Uncertain stimuli onset
- HRF estimation
- Model Diagnostics
- Appropriate control for multiple comparisons.
- Spatio-temporal models

END OF PART II

Data Processing Pipeline

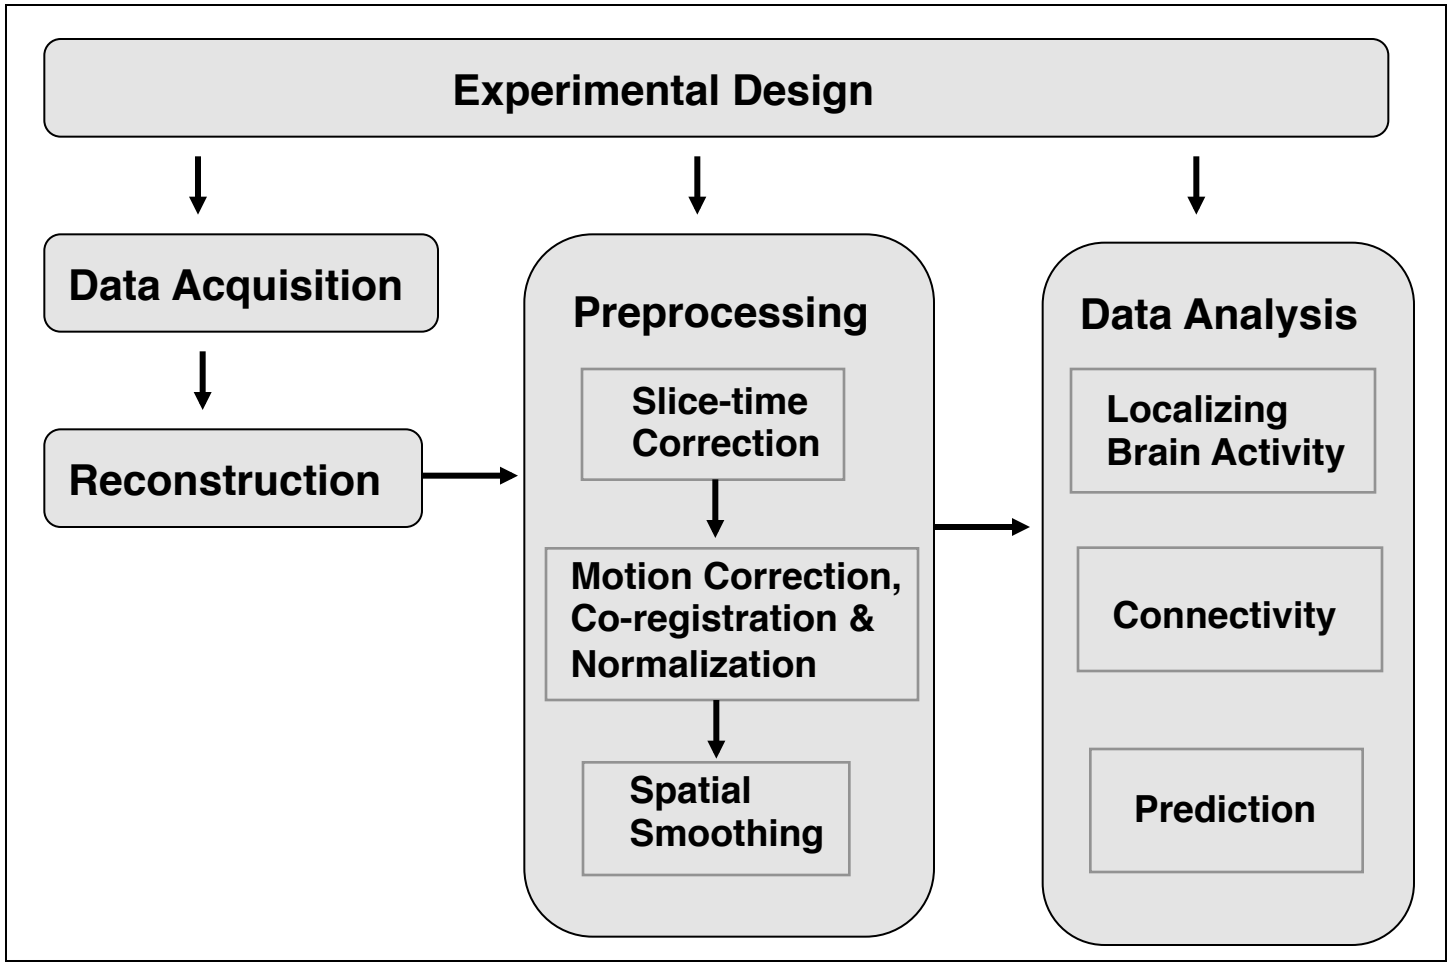

Brain Networks

- It has become common practice to talk about brain networks, i.e. sets of interconnected brain regions with information transfer among regions.

- To construct a network:
  - Define a set of nodes (e.g., ROIs)
  - Estimate the set of connections, or edges, between the nodes.

	A	B	C
A	0	1	0
B	0	0	1
C	0	1	0

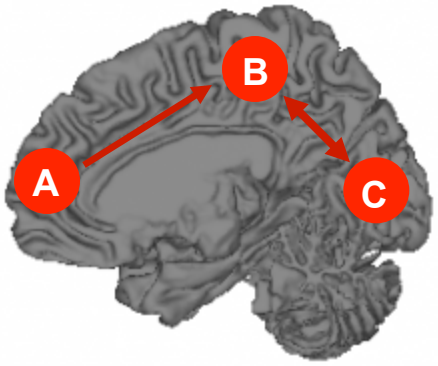

Network Methods

- A number of methods have been suggested in the neuroimaging literature to quantify the relationship between nodes/regions.

- Their appropriateness depend upon:
  - what type of conclusions one is interested in making;
  - what type of assumptions one is willing to make;
  - the level of the analysis;
  - and the modality used to obtain the data.

Brain Connectivity

- Functional Connectivity
  - Undirected association between two or more fMRI time series and/or performance and physiological variables.
  - Makes statements about the structure of relationships among brain regions.
  - Usually makes no assumptions about the underlying biology.

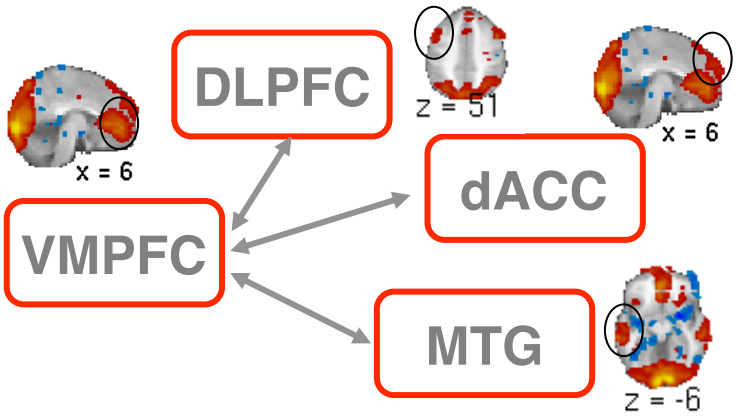

Brain Connectivity

- Effective Connectivity
  - Directed influence of one brain region on the physiological activity recorded in other brain regions.
  - Claims to make statements about causal effects among tasks and regions.
  - Usually makes anatomically motivated assumptions and restricts inference to networks comprising of a number of pre-selected regions of interest.

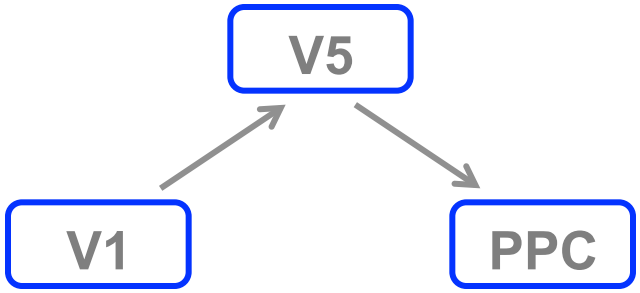

Functional Connectivity

- Methods include:
  - Seed analysis
  - Inverse covariance methods
  - Multivariate decomposition methods
    - Principle Components Analysis
    - Independent Components Analysis
    - Partial Least Squares

Seed Analysis

- In seed analysis the cross-correlation is computed between the time course from a predetermined region (seed region) and all other voxels.
- This allows researchers to find regions correlated with the activity in the seed region.
- The seed time course can also be a performance or physiological variable

Correlations between brain activity and heart-rate

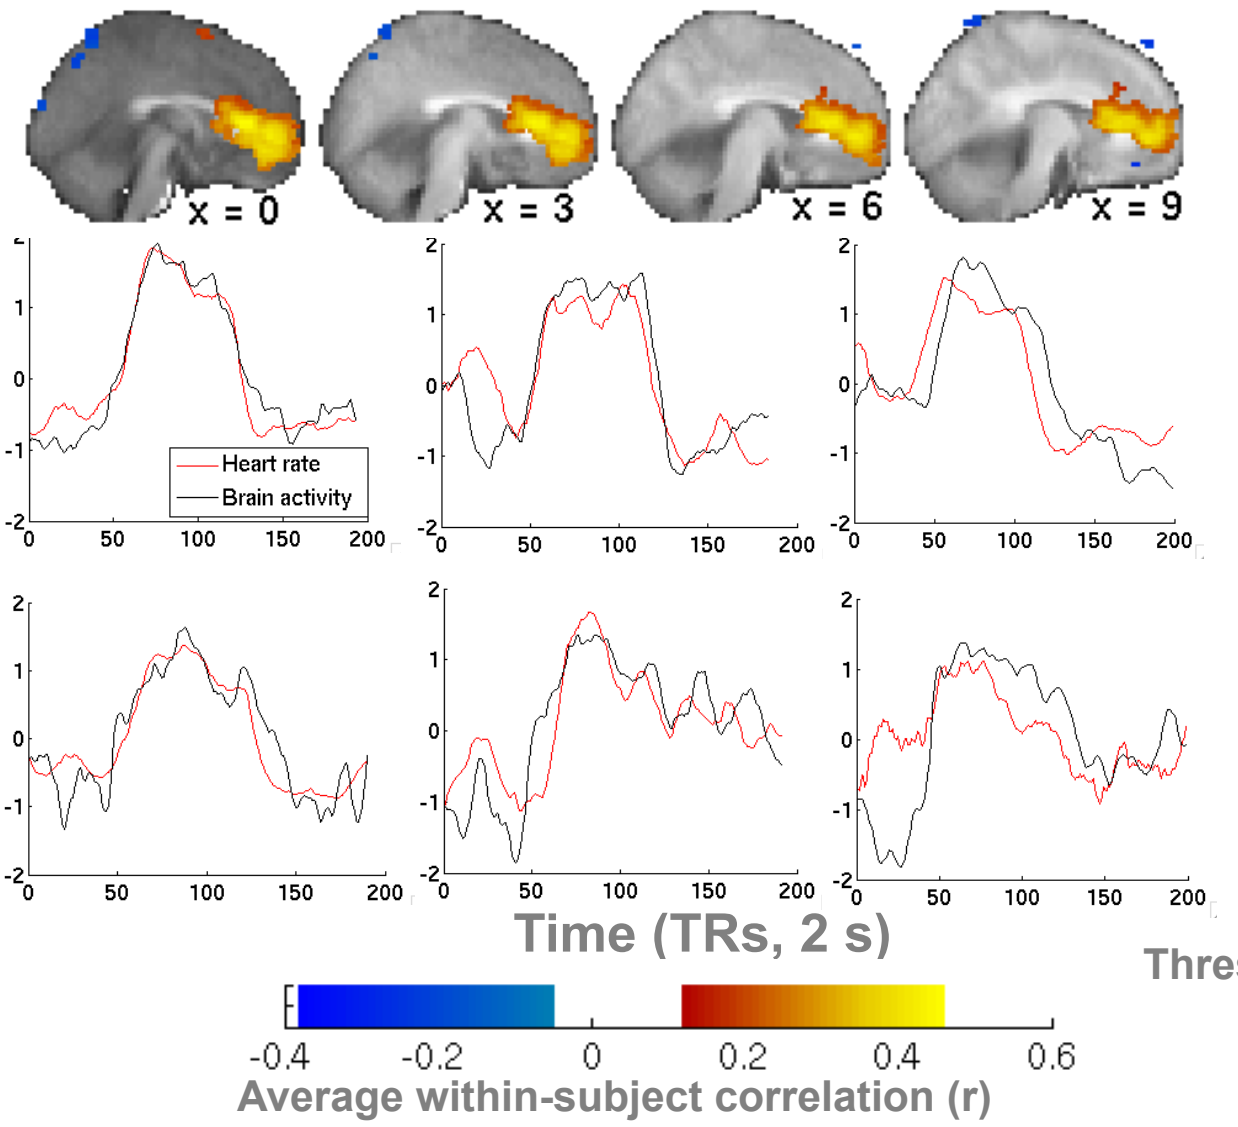

VMPFC

Threshold: p < .005

Inverse Covariance Methods

- For multivariate normal data there exists a duality between the inverse covariance (precision) matrix and the graph representing relationships between regions.
  - Conditional independence between variables (regions) corresponds to zero entries in the precision matrix.
  - Graphical lasso (GLASSO) can be used to estimate sparse precision matrices and graphs.

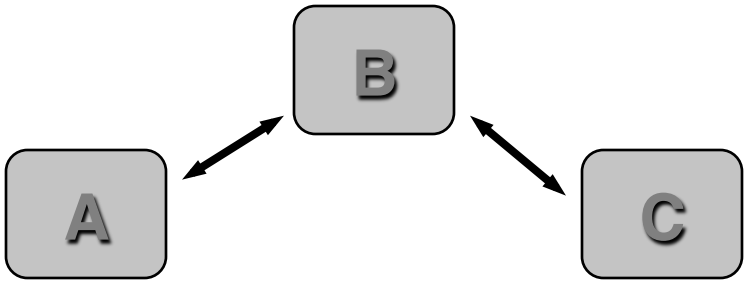

$$\Sigma^{-1} = \begin{bmatrix} & & 0 \\ & & \\ 0 & & \end{bmatrix}$$

Principal Components Analysis

- Principal components analysis involves finding spatial modes, or eigenimages, in the data.
  - These are the patterns that account for most of the variance-covariance structure in the data.
  - They are ranked in order of the amount of variation they explain.

- The eigenimages can be obtained using singular value decomposition (SVD), which decomposes the data into two sets of orthogonal vectors that correspond to patterns in space and time.

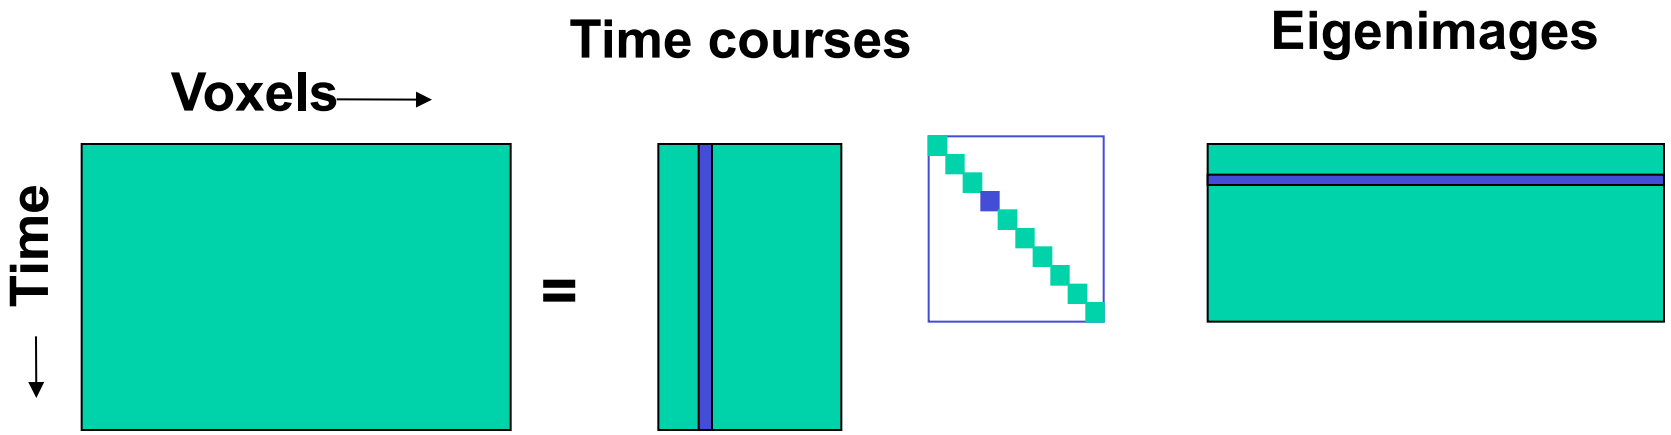

$$\mathbf{X} = \mathbf{USV}^T$$

$$\mathbf{X} = s_1\mathbf{u}_1\mathbf{v}_1^T + s_2\mathbf{u}_2\mathbf{v}_2^T + \ldots + s_N\mathbf{u}_N\mathbf{v}_N^T$$

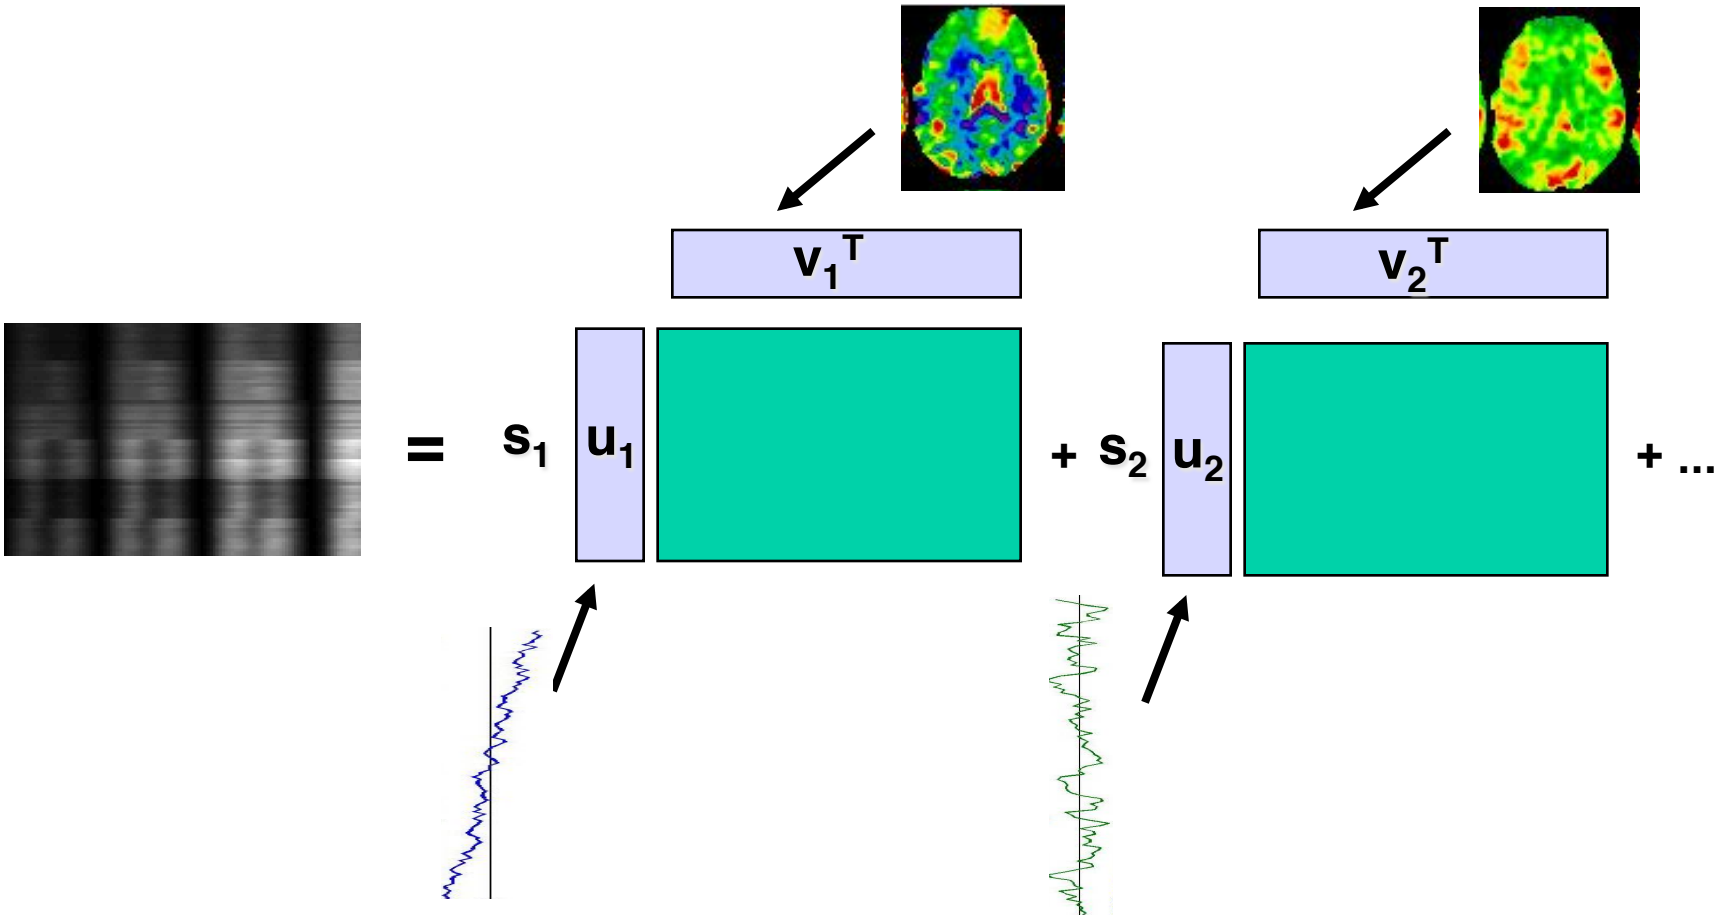
v_1^T
v_2^T
=
s_1
u_1
+ s_2
u_2
+ ...

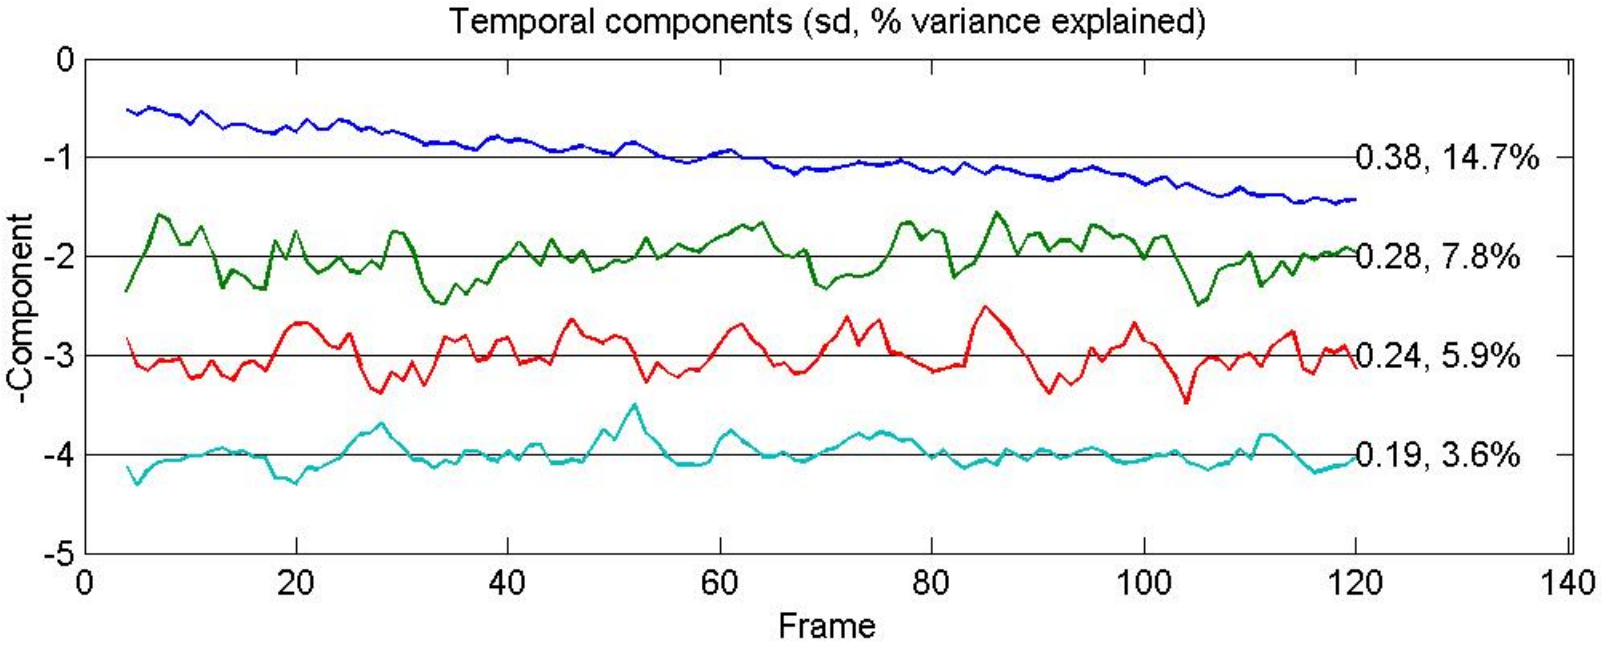
Temporal components (sd, % variance explained)
0
-1
-2
-3
-4
-5
-Component
0.38, 14.7%
0.28, 7.8%
0.24, 5.9%
0.19, 3.6%
0
20
40
60
80
100
120
140
Frame

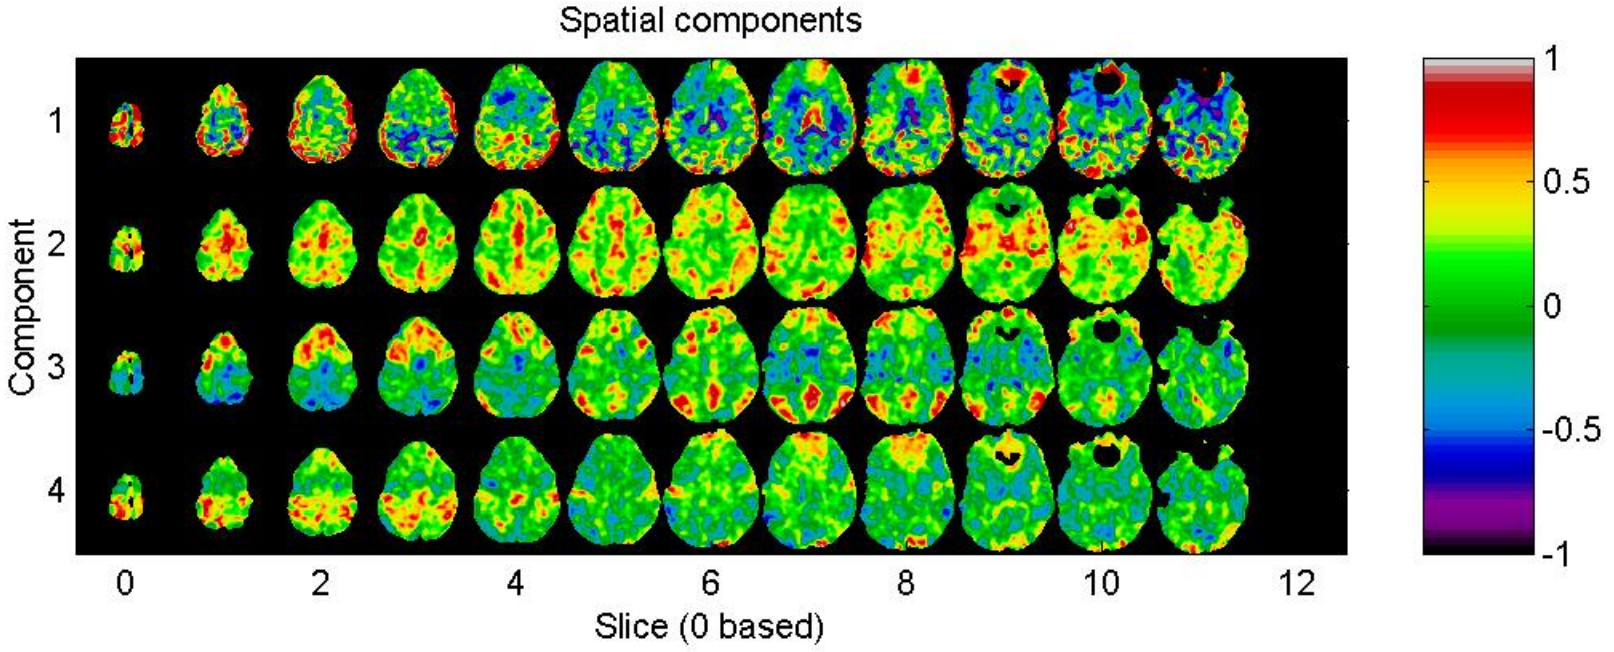
Spatial components
1
2
3
4
Component
0
2
4
6
8
10
12
Slice (0 based)
1
0.5
0
-0.5
-1

Worsley

ICA for fMRI

- Independent Components Analysis (ICA) is a family of techniques used to extract independent signals from some source signal.

- In ICA we seek to decompose $\mathbf{X}$ as follows:

$$\mathbf{X} = \mathbf{AS}$$

where the matrix $\mathbf{S}$ contains statistically independent maps in its rows each with a related time-course contained in the associated column of the mixing matrix $\mathbf{A}$.

Overview

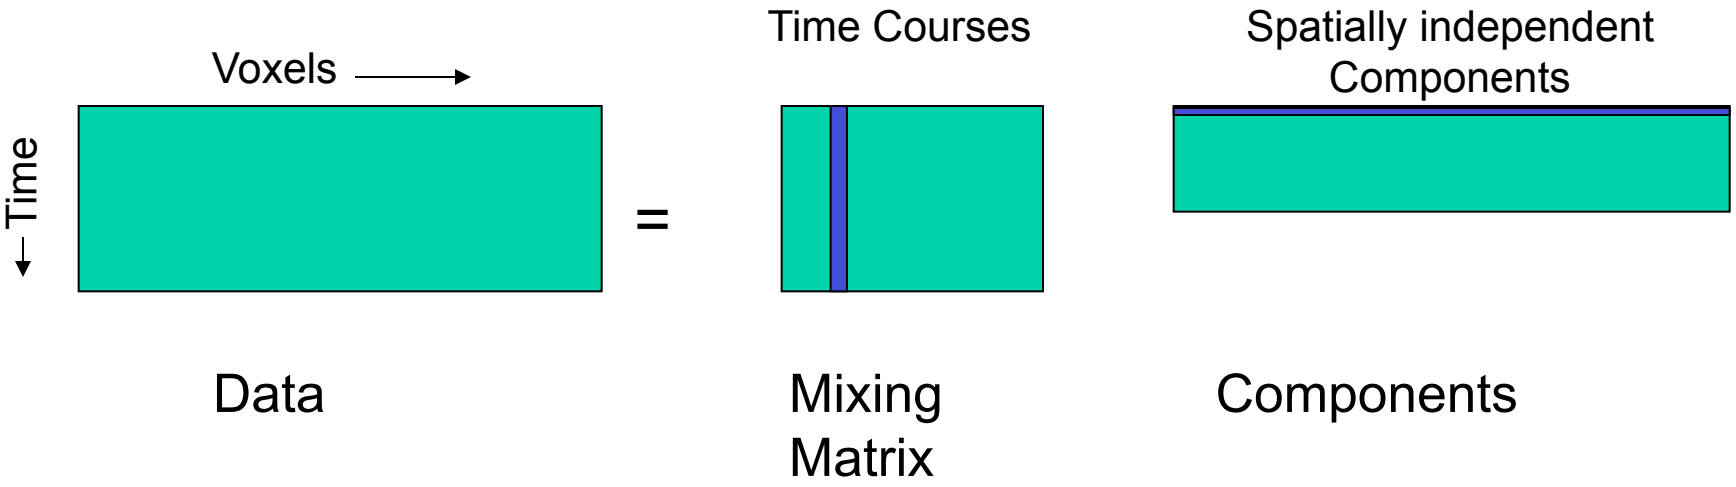

$$\mathbf{X} = \mathbf{AS}$$

Use an ICA algorithm to find A and S.

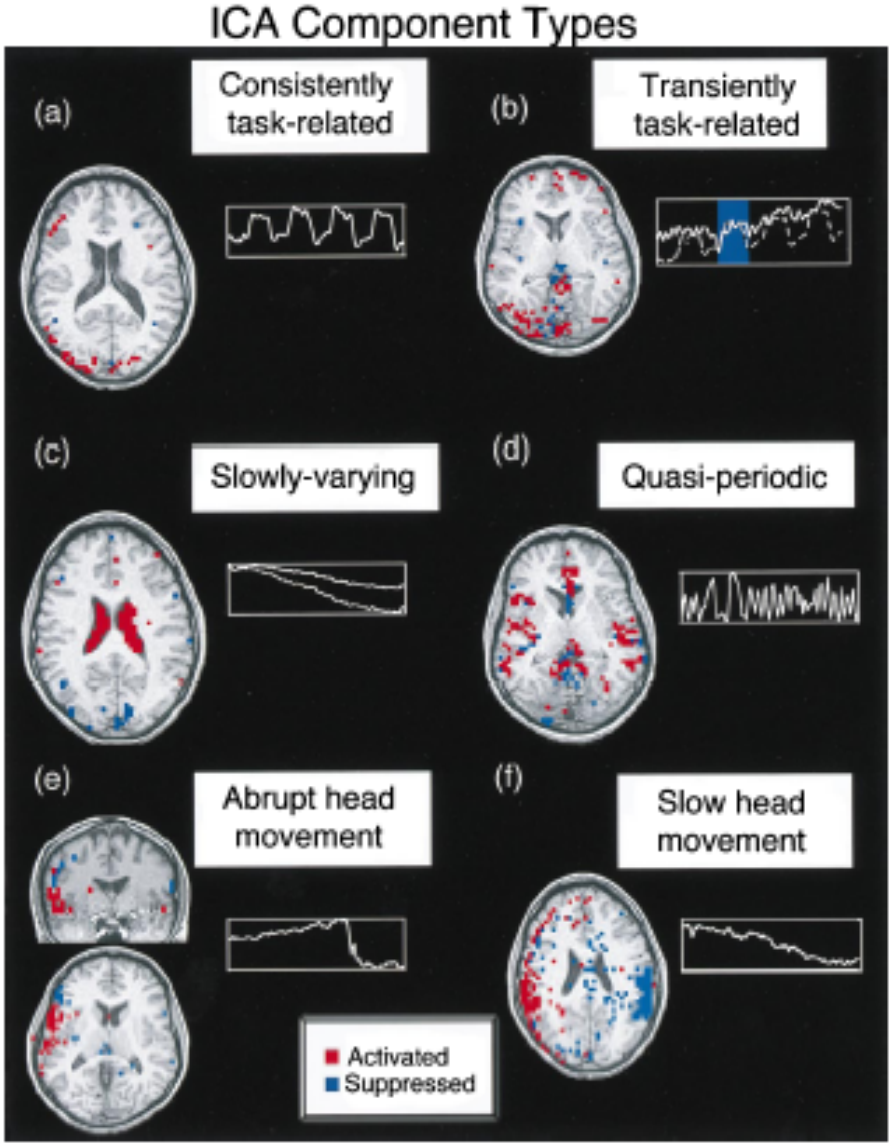

FIG. 1. Different classes of components detected by ICA decomposition of Stroop task fMRI data. (red, $r \geq 2.8$; blue, $r \leq 2.0$). Negative r values mean those voxels are activated opposite to the plotted time course. (*a*) Consistently task-related (CTR) component. (*b*) Transiently task-related (TTR) component. The dotted line shows the time course of the consistently task-related component for comparison. (*c*) Slowly varying, non-task-related component. The active region for this component was mostly localized to the ventricular system. The lower line shows the mean time course of the active voxels for this component. (*d*) Quasi-periodic component. This component was largely active in a single slice and had a dominant period of about 12 sec. The spatial distributions of such components were highly reproducible between trials. (*e*) Suspected abrupt head movement. Note abrupt change in time course, suggesting an abrupt head movement. (*f*) Component with a "ring-like" spatial structure suggestive of a head movement.

McKeown, et. al.

Comments

- Unlike PCA which assumes an orthonormality constraint, ICA assumes statistical independence among a collection of spatial patterns.
  - Independence is a stronger requirement than orthonormality.

- However, in ICA the spatially independent components are not ranked in order of importance as they are when performing PCA.

Effective Connectivity

- Methods include:
  - Structural Equation Modeling
  - Granger Causality
  - Dynamic Causal Modeling
  - Bayes Net

SEM

- Structural Equation Models comprise a set of regions and a set of directed connections.

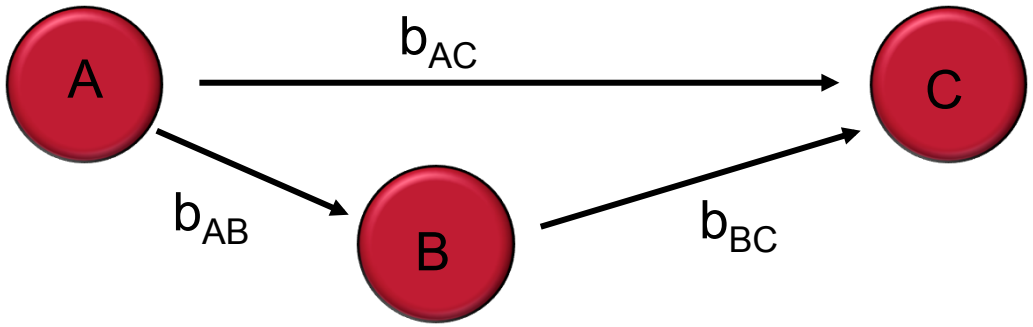

- Path coefficients defined between pairs of nodes.

- Directional relationships are assumed *a priori*.
  - Often given a causal interpretation.

Example

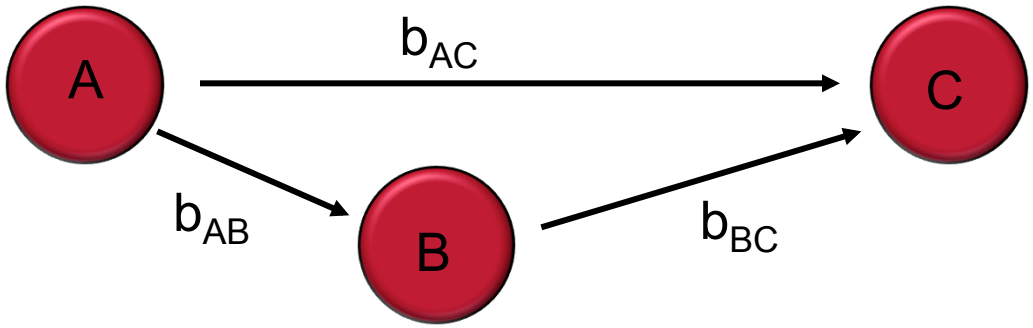

$$\begin{bmatrix} A \\ B \\ C \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ b_{AB} & 0 & 0 \\ b_{AC} & b_{BC} & 0 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} + \begin{bmatrix} e(1) \\ e(2) \\ e(3) \end{bmatrix}$$

$$y_t = My_t + e_t \qquad t = 1,\ldots T$$

Set-Up

- We can rewrite:

$$y_t = My_t + e_t$$

as

$$y_t = (I - M)^{-1} e_t$$

- Hence, we can write the covariance matrix of y_t as

$$\Sigma(\theta) = (I - M)^{-1} R((I - M)^{-1})^T$$

- The parameters θ are the unknown elements of the matrices M and R.

Estimation

- The covariance of the data represents how activities in two or more regions are related.

- In SEM we seek to minimize the difference between the observed covariance matrix and the one implied by the structure of the model.
  - The parameters of the model are adjusted to minimize this difference.
  - Typically maximum likelihood estimation is used to estimate the parameters.

Dynamic Casual Modeling

- DCM attempts to model latent neuronal interactions using hemodynamic time series.
  - Based on a neuronal model of interacting regions, supplemented with a forward model of how neuronal activity is transformed into the observed response.

- Effective connectivity is parameterized in terms of the coupling among unobserved neuronal activity in different regions.
  - We can estimate these parameters by perturbing the system and measuring the response.

Illustration

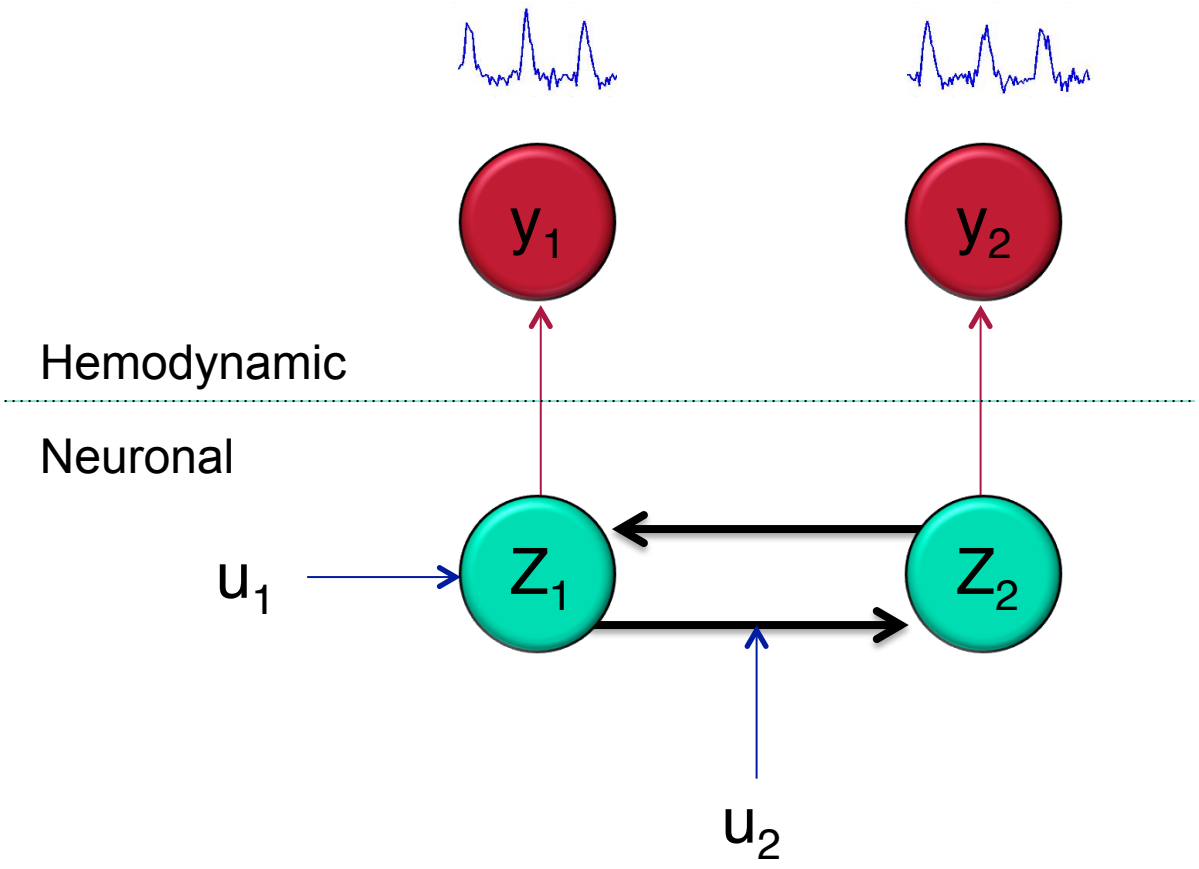

Neuronal Model

- Define the neuronal states as:

$$z = \left(z_1, \ldots z_N\right)^T$$

- The effective connectivity model is described by:

$$\dot{z}_t = \left(A + \sum_{j=1}^{J} u_t(j) B^j \right) z_t + C u_t$$

where z_t is the neuronal activity at time t (latent) and $u_t(j)$ is the j^{th} of J inputs at time t (known).

Interpretation

- The matrix A represents the first order connectivity among regions in the absence of input.
  - Specifies how regions are connected and whether these connections are uni- or bidirectional.
- The matrix C represents the extrinsic influence of inputs on neuronal activity.
  - Specifies how inputs are connected to regions.
- The matrices B_j represent the change in coupling induced by the jth input.
  - Specifies how connections are changed by inputs.

$$\dot{z}_t = \left(A + \sum_{j=1}^{J} u_t(j) B^j \right) z_t + C u_t$$

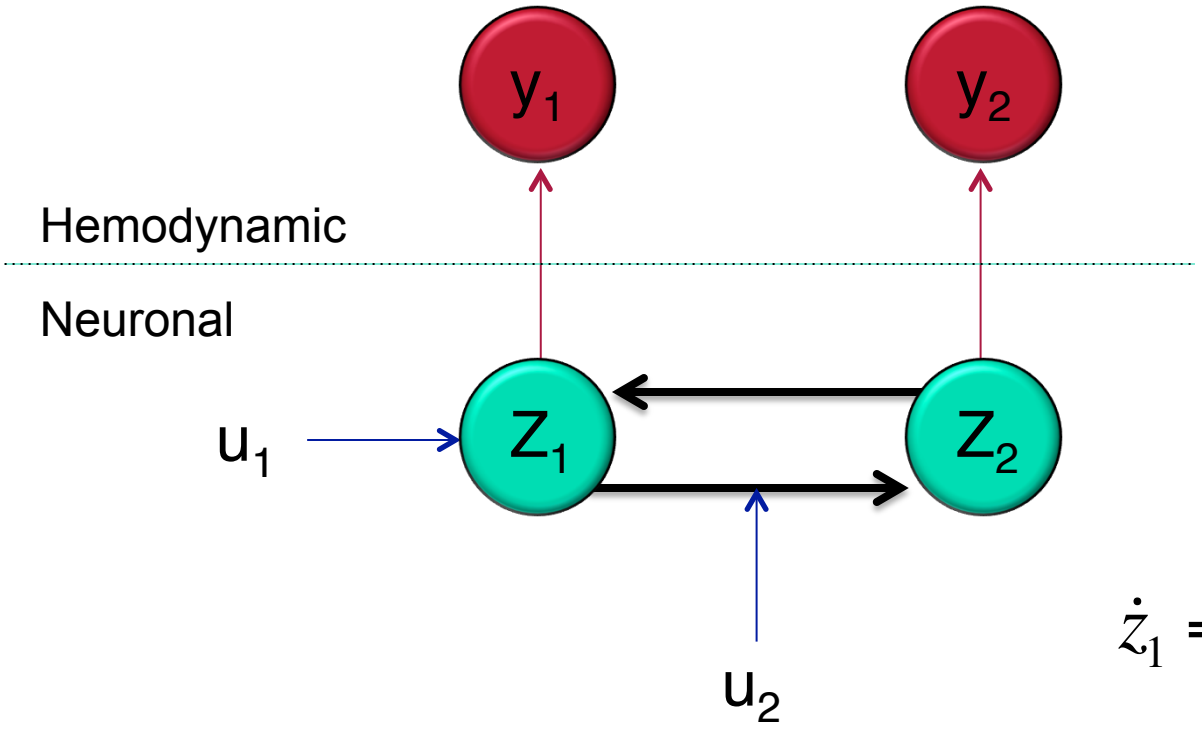

$$\dot{z}_1 = a_{11} z_1 + a_{12} z_2 + c_{11} u_1$$

$$\dot{z}_2 = a_{21} z_1 + a_{22} z_2 + b_{21}^2 u_2 z_1$$

Hemodynamic Model

- Neuronal activity causes changes in blood volume and deoxyhmoglobin that cause changes in the observed BOLD response.

- The hemodynamics are described using an extended Balloon model, which involves a set of hemodynamic state variables, state equations and hemodynamic parameters θ^h.

Extended Balloon Model

Activity-dependent signal: $\dot{s} = z - \kappa s - \gamma(f - 1)$

Flow induction: $\dot{f} = s$

Changes in volume: $\tau\dot{v} = f - v^{1/\alpha}$

Changes in dHb: $\tau\dot{q} = fE(f,\rho)/\rho - v^{1/\alpha}\, q/v$

Hemodynamic response $y = \lambda(v, q)$

State Equations

Neuronal state:

Neuronal activity - z_t with parameters θ^c.

Hemodynamic states:

Vasodilatory signal - s_t

Inflow - f_t

Blood volume - v_t

Deoxygenation content - q_t

The observed data: $y_t=\lambda(q_t,v_t)$ with parameters θ^h.

Bayesian Analysis

- Combining the neuronal and hemodynamic states $x=\{z, s, f, v, q\}$ gives us the following state-space model:

$$\dot{x} = f(x, u, \theta)$$

$$y = \lambda(x, \theta)$$

- Analysis performed using Bayesian methods
  - Normal priors are placed on θ.
  - The posterior density is used to make inferences about the connections.

Comments

- Effective connectivity is thought to provide more powerful conclusions than functional connectivity, but makes much stronger assumptions.
  - Sometimes the difference is not so obvious.
  - The validity of the conclusions depend strongly on the assumptions being correct.
  - The necessary assumptions often poorly specified and difficult to check; this is a major shortcoming of the field.
- More care needs to be taken in discussing these concepts in connectivity studies.
  - Connectivity should be based on carefully defined estimands, not on the applied estimation algorithm.

Prediction

- There is a growing interest in using fMRI data for classification of mental disorders and predicting the early onset of disease.
  - Requires the application of novel statistical and machine learning techniques.

- The application of machine learning methods to fMRI data is often referred to as multi-voxel pattern analysis (MVPA)
  - A classifier is trained to discriminate between different brain states and used to predict states in a new data.

Machine Learning

- When applied to fMRI data the result is often a pattern of weights across brain regions that can be applied to new brain activation maps to quantify the degree to which the pattern responds to a particular type of event.

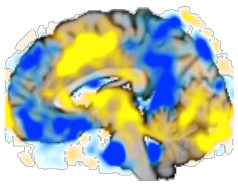

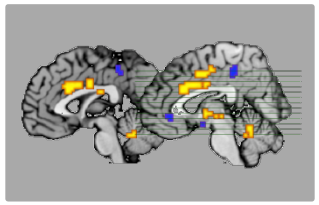

$\underline{w}^T \underline{x} > 0$ Group A

$\underline{w}^T \underline{x} < 0$ Group B

$\underline{x} = (x_1, \ldots x_V)$ $\underline{w} = (w_1, \ldots w_V)$

Performing MVPA

- The process of performing MVPA follows a series of steps:
  - Defining features and classes
    - Raw data, averages, beta-maps
  - Feature selection
    - ROI, PCA, meta-analysis
  - Choosing a classifier
    - SVD, LDA, logistic regression
  - Training and testing the classifer
    - Cross-validation
  - Examining results
    - Prediction accuracy, brain weights

Example

- Participants (n=41; Kross et al., 2011) had all recently been rejected in a romantic relationship.
  - They were asked to bring pictures of the former partner and a non-rejecting friend.

- Experiment consisted of four randomly intermixed conditions: (1) viewing the rejector; (2) viewing the friend; (3) receiving high thermal pain; (4) receiving low thermal pain.

- Subjects reported perceived pain after each trial.

Figure 3. Study 2: Prediction of physical pain vs. emotional distress

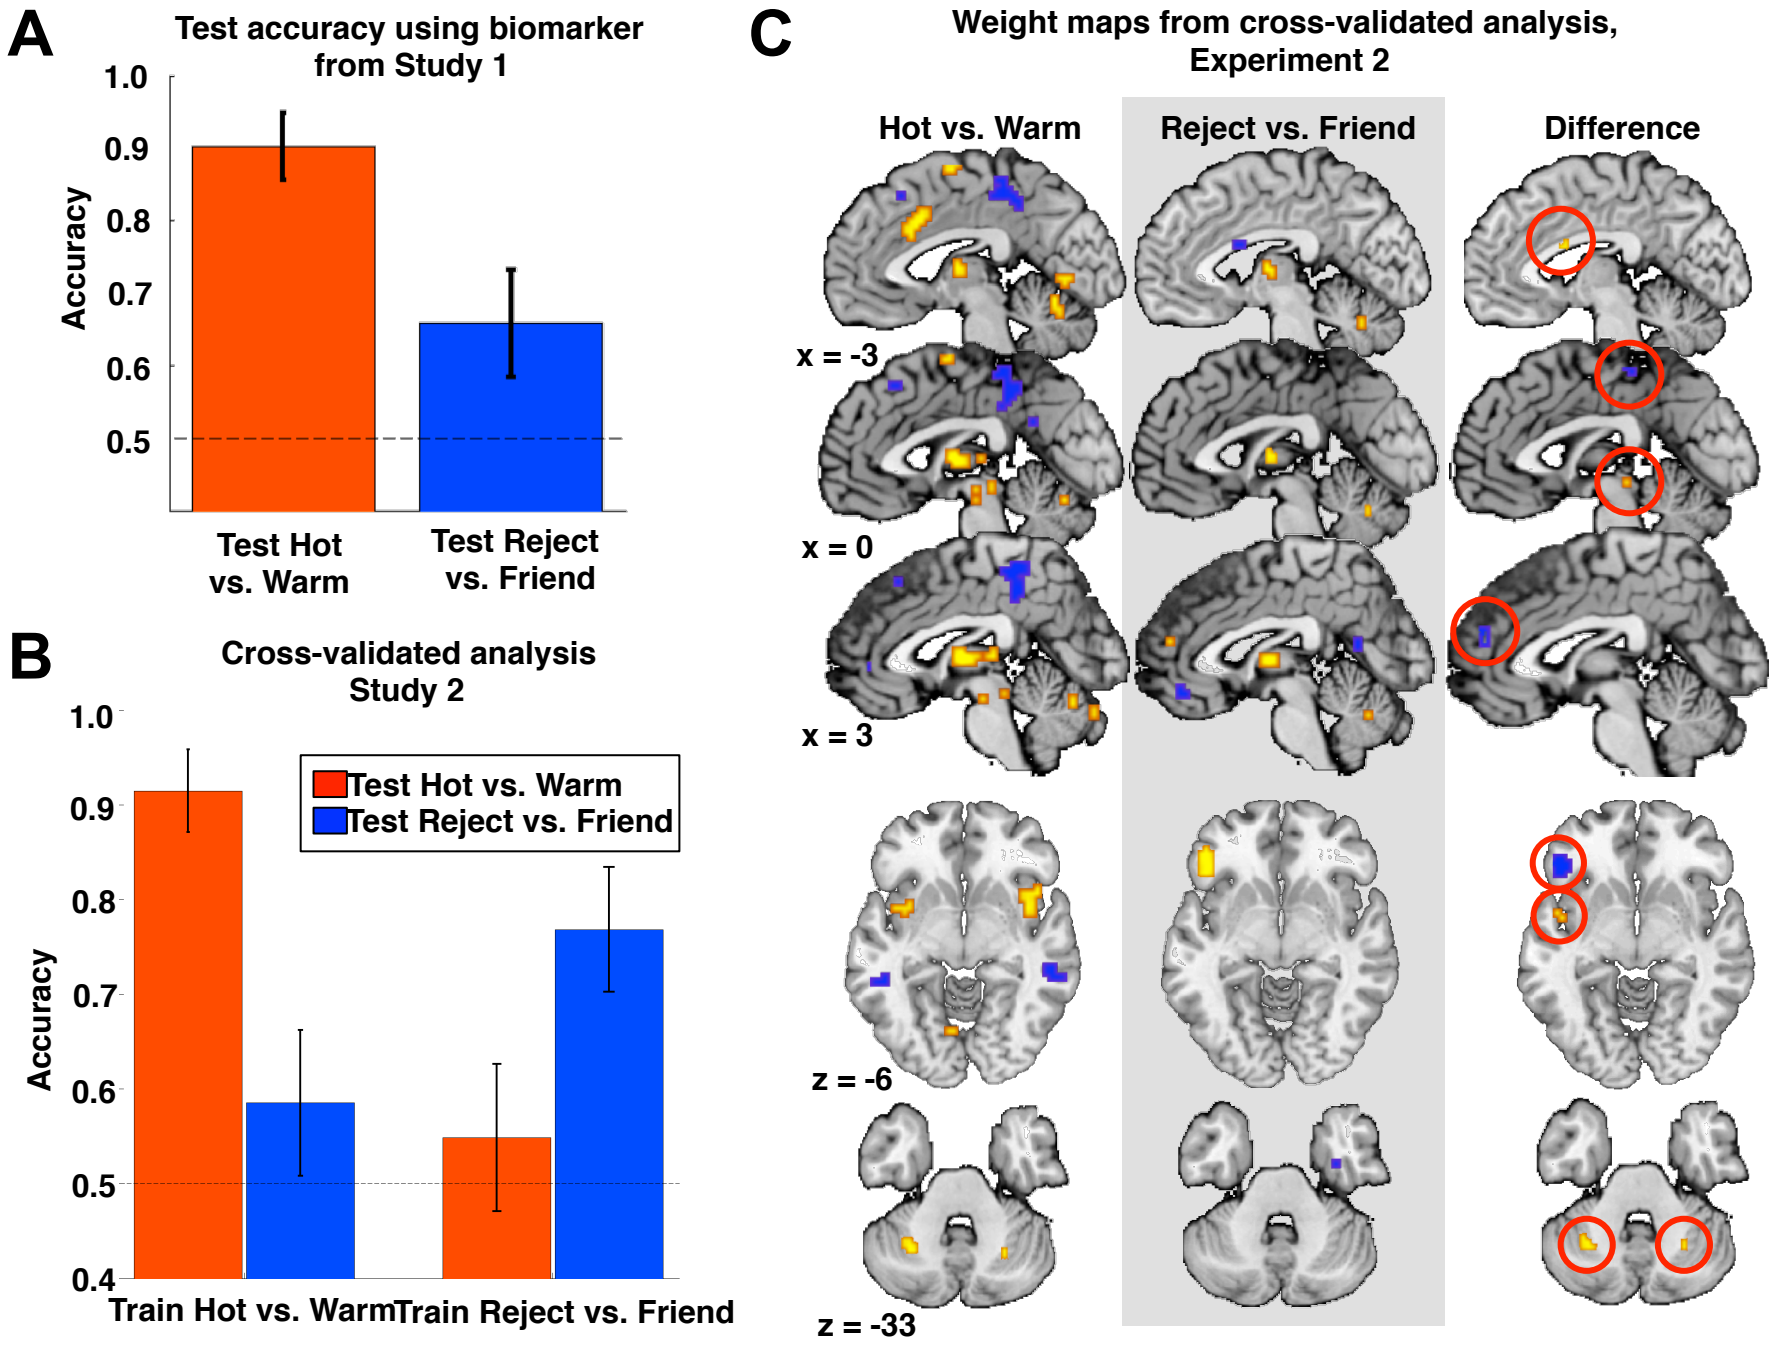

Multi-modal Experiments

- There is a trend towards using multiple imaging modalities to overcome some of the limitations of each method used in isolation.
  - Combined EEG and fMRI
  - Combined DTI and fMRI
  - Combined TMS and fMRI
  - Imaging genetics (combined fMRI and genetics)
- All very promising, but even more data intensive than fMRI alone.

Interesting Questions

- Dynamic connectivity
- Casual inference for brain connectivity
- Graph network analysis methods
- Small n, large p type questions
- Biomarker development
- Efficiently combine data across modalities

MOOC coming this Fall

- New massive open online course (MOOC) on

  'The Statistical Analysis of fMRI Data'

- Coming this Fall to Coursera.