Temporal preprocessing of fMRI data

Blaise Frederick, Ph.D., Yunjie Tong, Ph.D.

McLean Hospital Brain Imaging Center

Scope

- This talk will summarize the sources and characteristics of unwanted temporal features of fMRI signals, how they are recorded in MR data, and how they can be removed or attenuated prior to further processing, without compromising the signal of interest.
- I will discuss spectral filtering, removal of single point discontinuities (spikes) and slice time correction; later talks will discuss more sophisticated methods for removing specific motion- and physiology- related signals, such as GLM filters and modeling.

Temporal preprocessing - goals

- Characterize and remove noise from various sources.
- Preserve signal.
- Make sure that data are ready for analysis by ensuring that the assumptions underlying the analyses are correct.

Noise in fMRI - sources

- Subject generated
  - Physiology
    - Heartbeat, respiration, low frequency oscillations
  - Motion
- Equipment generated
  - Machine instability
- Externally generated
  - Interference from inside or outside of the scan suite

Noise in fMRI - characteristics

- Physiological
  - Cardiac signal: Pseudoperiodic, 40-120 beats/min (0.66-4.0 Hz, accounting for harmonics)
  - Respiration: Pseudoperiodic, 12-20 breaths/min (0.2-0.67Hz, accounting for harmonics)
  - Low frequency oscillations: Aperiodic, (~0.01-~0.15*Hz)
  - Motion: Aperiodic, usually impulsive discontinuities any frequency, but also pseudoperiodic at the respiratory rate.
- Equipment generated
  - Machine instability: could be anything, but in a modern scanner, this is usually constrained to very low frequencies from thermal drift in shims, gradients, and RF components, (below ~0.1Hz), or extremely high frequency (single point spikes), with a 1/f distribution.
- External signals
  - Best dealt with at the source (eliminate the interference before running the experiment).

*the exact boundaries of this range are the subject of some debate.

fMRI signal - sources

- Task response:
  - Changes in neuronal activity in response to stimuli or directed cognitive activity.
- Resting state signals:
  - Changes in neuronal activity due to correlated activity in connected brain regions.
- In both cases, the fMRI signal comes from the hemodynamic response to the neuronal activity, not the activity itself.
  - This has important implications for the temporal characteristics of the signal.

Hemodynamic response

- fMRI doesn't see neuronal activation – it sees the hemodynamic response to neuronal activation.
- As a result, any neuronal signal we want to look at is filtered by the hemodynamic response function (HRF) – this is the response to a neuronal event of the shortest possible duration (shown below on the left). No neuronal fMRI signal can be faster than this.
- This puts a hard limit on the frequency spectrum of "true" activations (see below right).

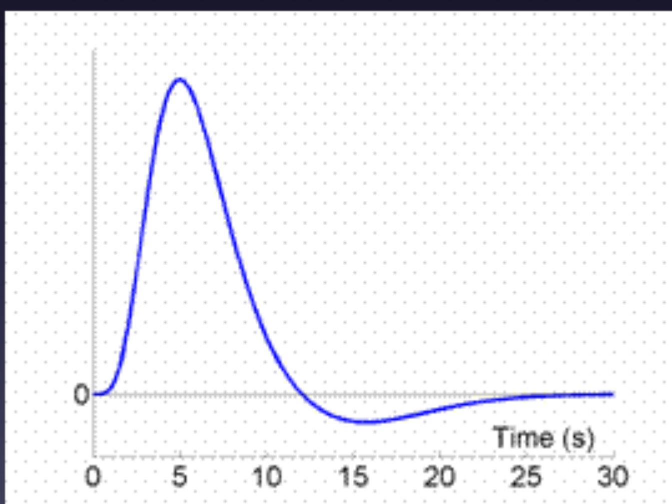

Hemodynamic response function

Fourier Transform

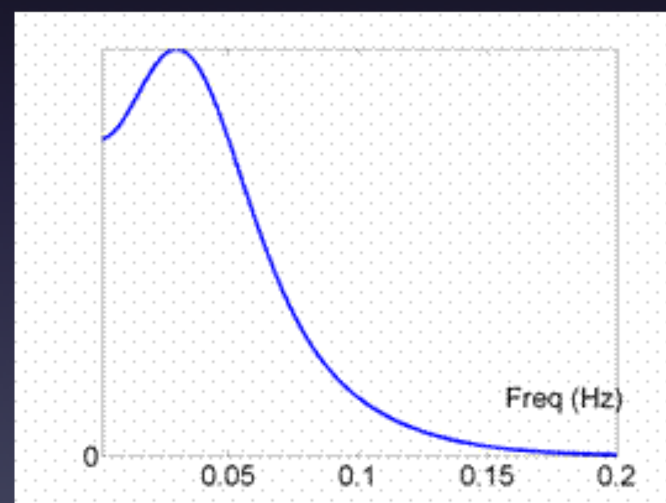

Power spectrum of HRF

Figures from: http://imaging.mrc-cbu.cam.ac.uk/imaging/DesignEfficiency

fMRI signal - characteristics

- The spectral properties of fMRI signals are determined by the underlying neuronal activation and the HRF filter function.
  - For resting state data, the spectrum is just that of the filter function.
  - For for task data, it varies (but is limited to the range of the filter function) – the HRF imposes a design constraint on tasks. For maximum design efficiency, you try to concentrate the task spectrum in the area of peak HRF response – i.e. around 0.03Hz .

Spectral filtering

- The easiest way to improve SNR is to preprocess your data with a filter matching the spectral characteristics of your signal of interest.
  - This will preserve your signal, while eliminating out of band noise.
  - In band noise you either have to live with, or model (see following talks).

fMRI spectrum (resting state)

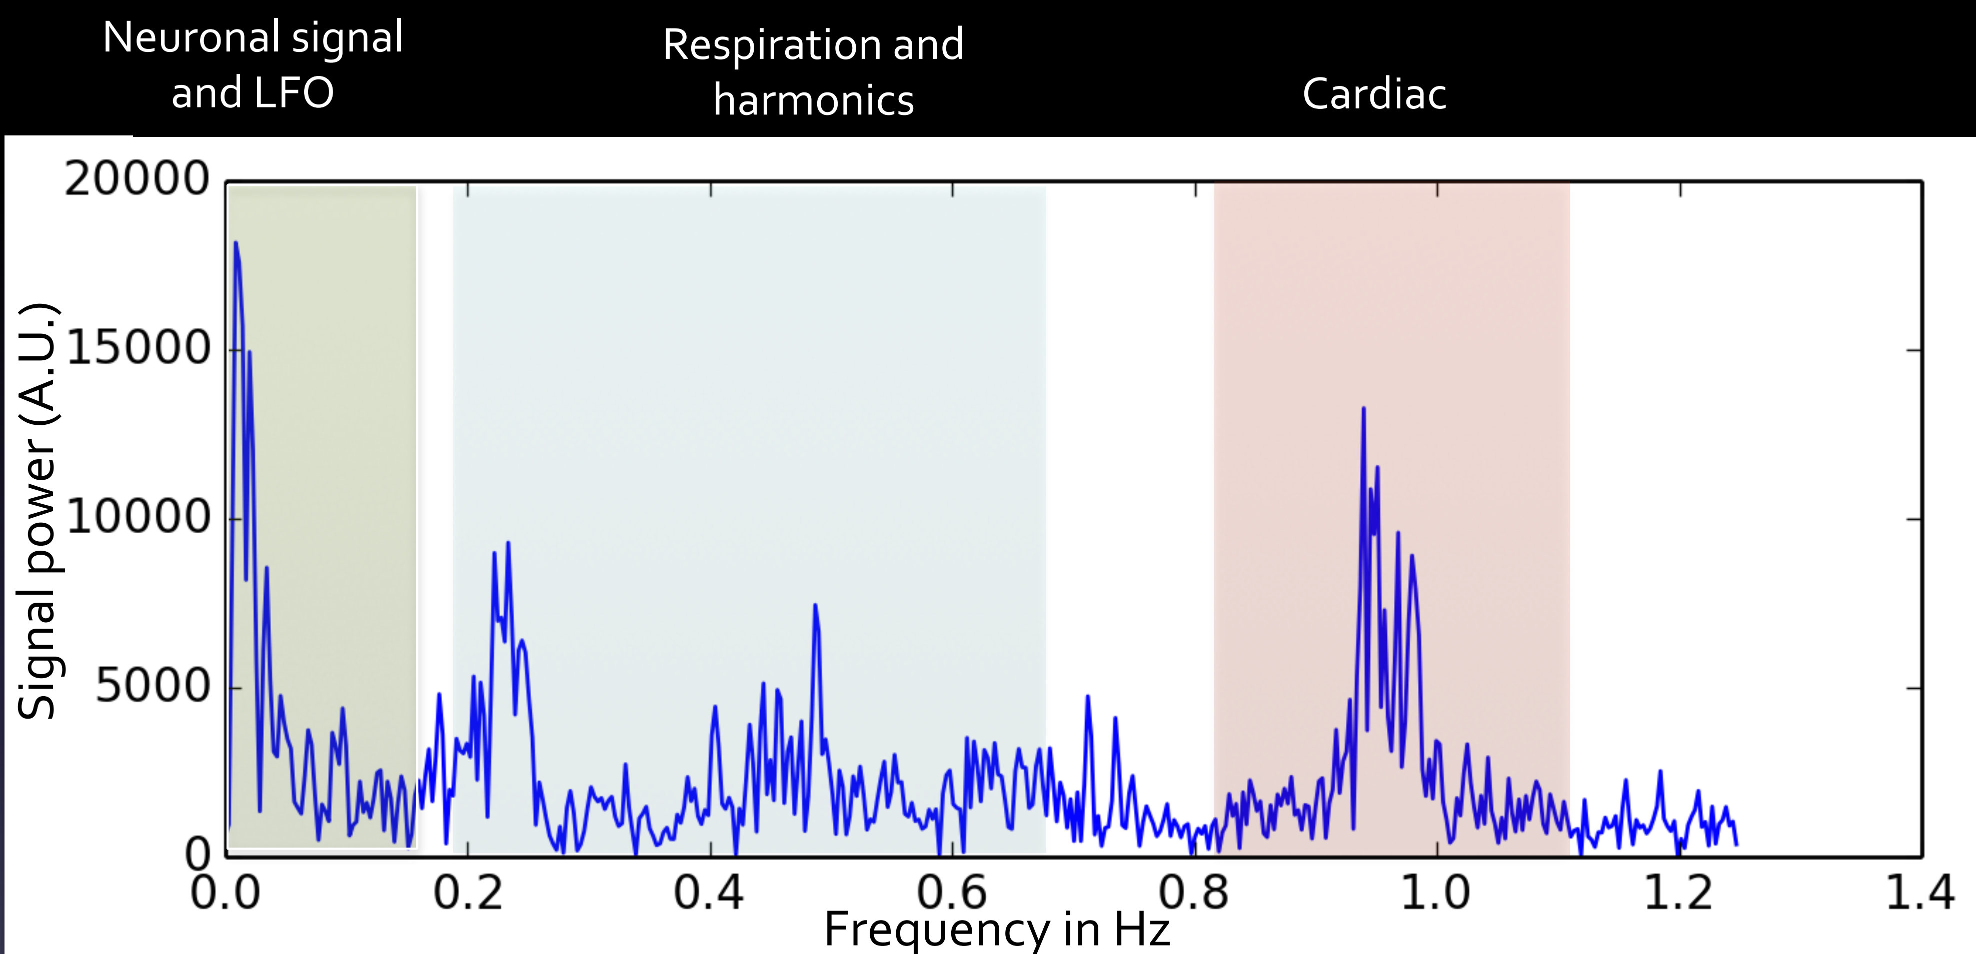

This is a voxel from the frontal lobe near a blood vessel from a 400ms TR EPI dataset. The voxel is near a blood vessel, and was chosen to show all the physiological noise components – most voxels aren't nearly this bad!

fMRI spectrum (resting state)

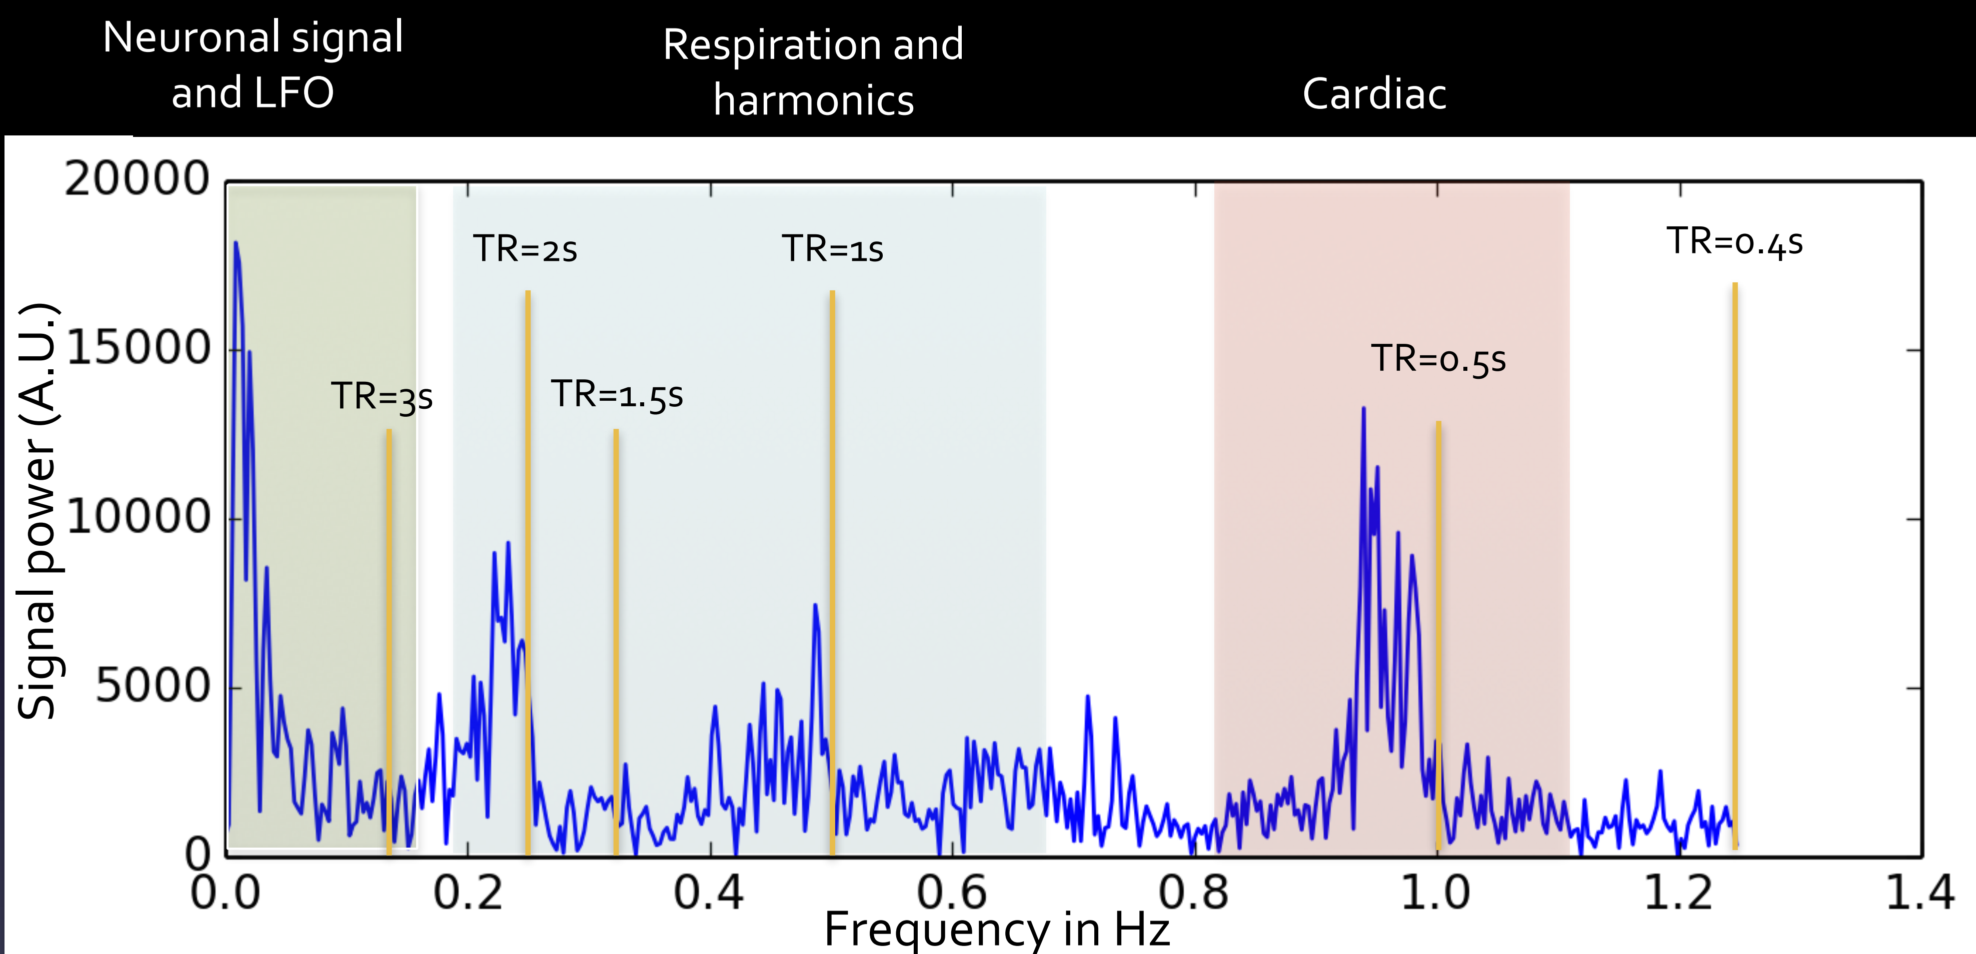

Increasing the TR will reduce the acquisition bandwidth. As a result, certain physiological processes will not be fully sampled at higher TR values, resulting in aliasing. When that happens, you can't remove that process using only a spectral filter.

What are your best choices?

- Highpass filtering – always recommended:
  - The filter cutoff value depends on your experiment.
  - For resting state data, 0.01Hz (~100s cycle time) is usually chosen to reduce 1/f noise, scanner drift, and linear trends. This is a good compromise that preserves the low frequency connectivity signals.
  - For task data, you should set your highpass filter frequency as high as possible without attenuating the task waveform. (for example, with a 20s off, 20s on block paradigm, you would eliminate everything below a 40s cycle time, or 0.025Hz.)
- **Important caveat** – you need to keep the zero frequency component (i.e. the mean) for display, masking, image interpretation, and activation scaling. So "high-pass filtered" data is actually "high-pass filtered + mean" data. All common processing packages add back the mean after filtering.

What are your best choices?

- Lowpass filtering – only recommended if certain conditions are met:
  - Lowpass filtering is typically not performed, because with standard TR values, cardiac and respiratory waveforms are severely aliased – modeling is usually a better approach to removing them – filtering the aliased signal will interfere.
  - If you do use a lowpass filter, neuronal signals will be below 0.15Hz (6.6s cycle time), so that's a good place to set the cutoff frequency.
  - If your TR is 1.5s or shorter, you can attenuate respiratory noise (if your TR is 750ms or less, you can eliminate it).
  - If your TR is 500ms or shorter, you can remove the majority of the cardiac signal as well (with a TR of ~166ms you can probably remove it all).

How do you do this?

- Highpass filtering is a standard preprocessing feature of most analysis packages (FSL, SPM, BrainVoyager, AFNI), so just make sure you set the filter cutoff appropriately and use it.
- Depending on the package, lowpass filtering may be just as easy to use as highpass filtering, or may require a little more work (command line tools, external programs) to do.
- There are some implementation differences in the packages in terms of type of filter, but they are all good.

Spectral filtering - summary

- Simple spectral filtering is good for getting rid of physiological and machine noise below your frequencies of interest, but unless you have a very short TR, it doesn't do much for things at higher frequencies – that requires more sophisticated methods, which are described in subsequent talks.

Spike removal

- Spikes are large, impulsive positive or negative going discontinuities in the timecourse, followed by a return to normal.
- Multiple sources:
  – Electrical sparks (from gradient problems, metal in bore)
  – Rapid motion (e.g. coughing)
  – External interference
- They cannot be effectively removed by spectral filtering.
  – Linear filters will "ring" when they hit a discontinuity, and contaminate adjacent points (making things worse).

Spike removal strategy

- Detect “abnormal” timepoints by examining voxel timecourses and the motion traces and flagging volumes with discontinuities (using the signal derivatives or a median filter).
- Remove the affected volumes and interpolate.
- Optionally generate a set of single point regressors to model out residual variance during analysis.

Spike removal example

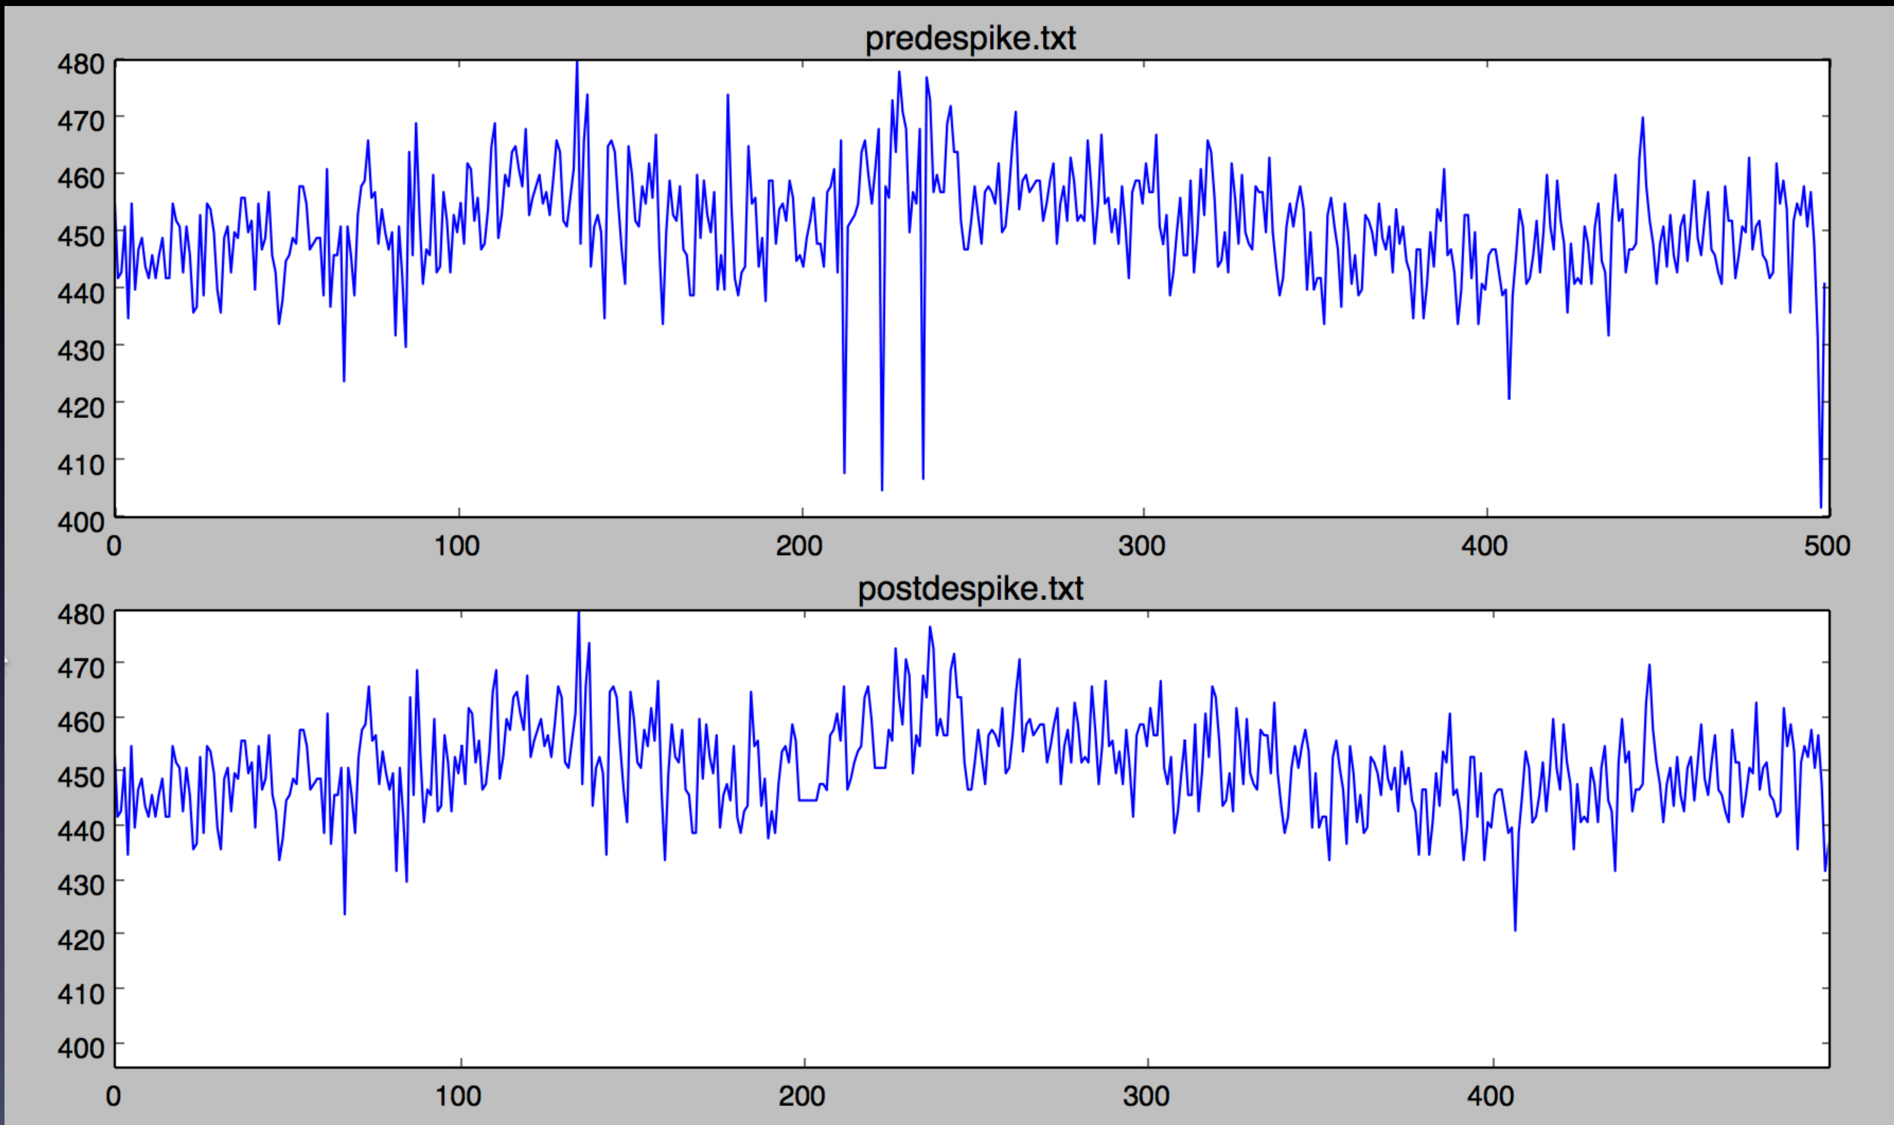

Tools for spike removal

- Unlike spectral filtering, this is not always included major fMRI processing packages, so many groups implement their own correction routines.
- There are some packages available.
  - ART (from Susan Gabrieli at MIT) (http://www.nitrc.org/projects/artifact_detect/)
  - fsl_motion_outliers (part of FSL - only does motion spikes).
  - despike (from me, bbfrederick@mclean.harvard.edu, on NITRC soon)

Artifact Detection Toolbox (ART)

- ART is a GUI tool for remedying a number of problems in datasets. Among its tools is one for detecting and removing spikes from fMRI data.
- Examines the global mean signal and motion parameters to detect "anomalous" timepoints that need to be dealt with.

Data Review

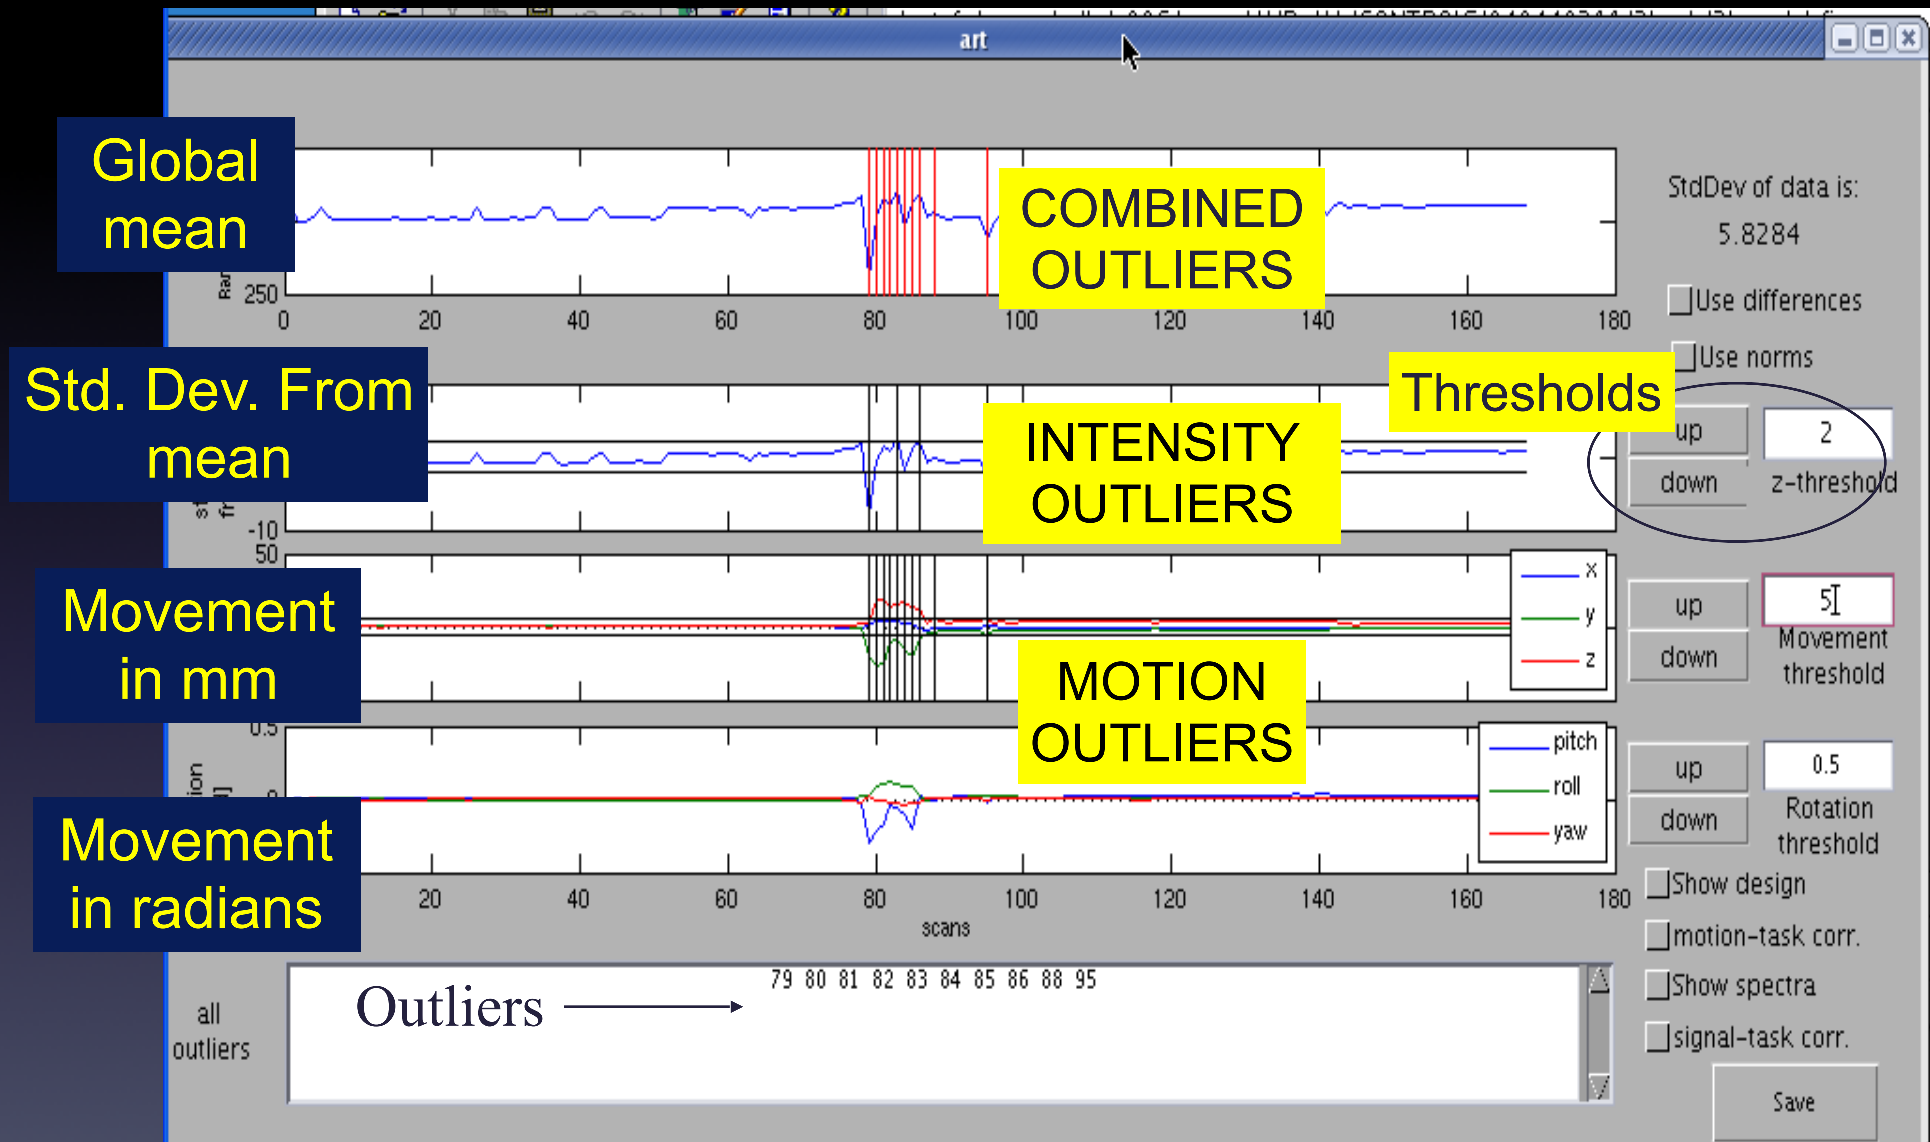

Slide courtesy of Susan Whitfield-Gabrieli

Artifact detection rs-fcMRI (10 y.o. child)

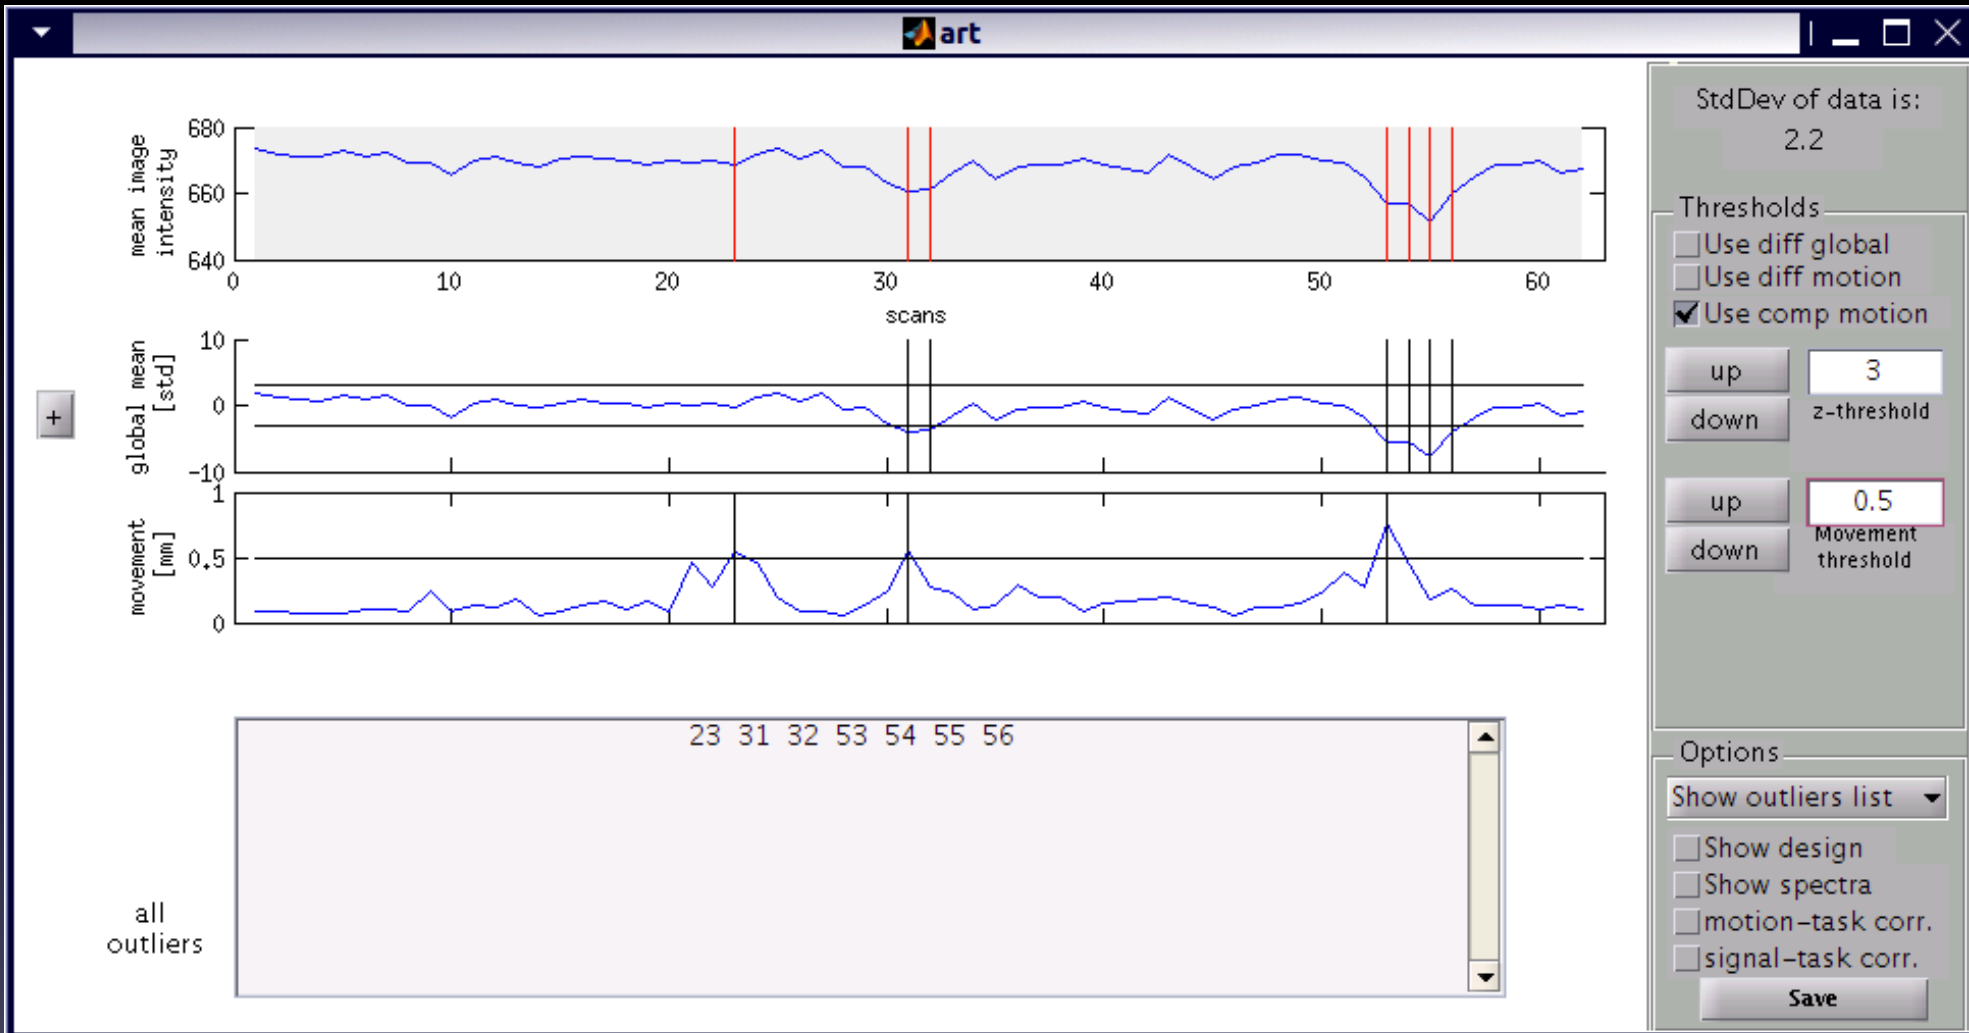

Slide courtesy of Susan Whitfield-Gabrieli

Effect of artifact rejection on first-level connectivity map – MPFC seed

Regressing out motion parameters only

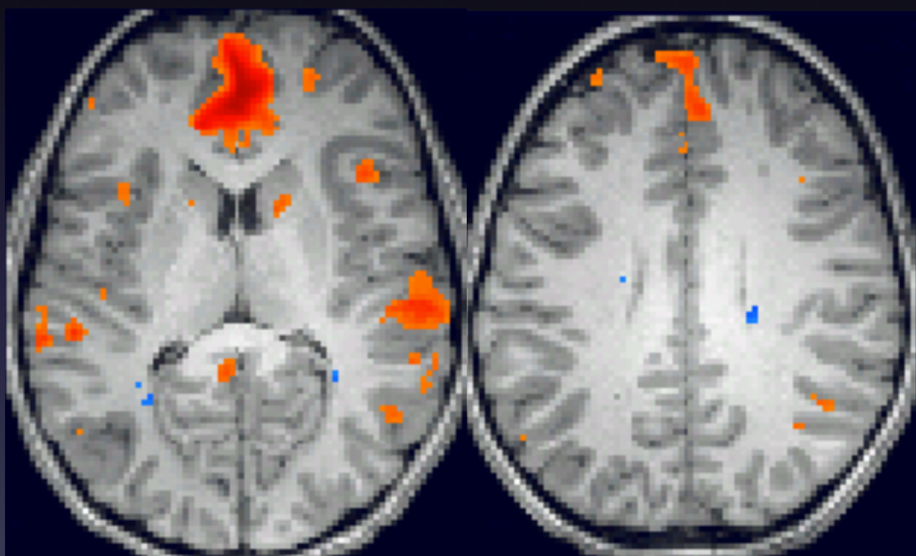

Regressing out motion parameters + outliers

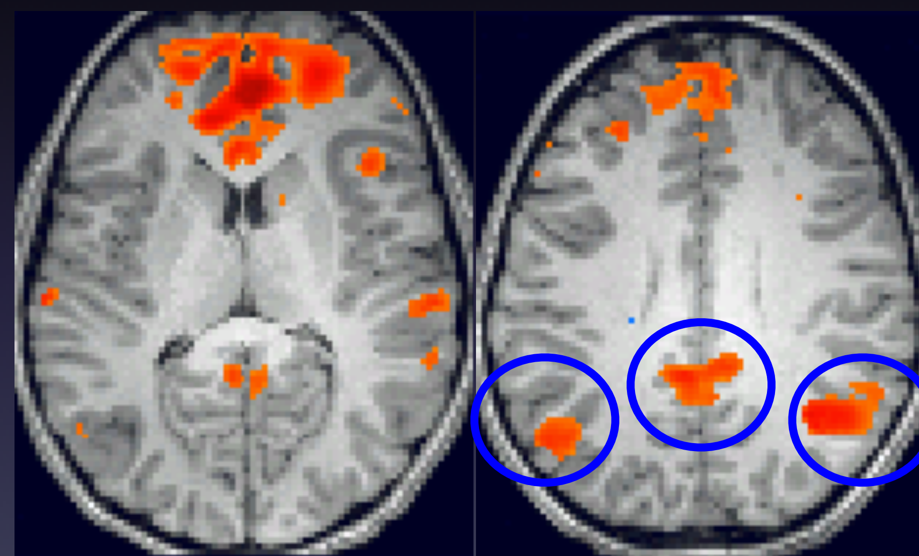

Slide courtesy of Susan Whitfield-Gabrieli

Spike removal summary

- Spikes can be very bad for your analyses, essentially destroying your ability to see activations. If you have them, you should use a tool like ART to remove them.
- The good news is that it's not common nowadays to have them unless you have a subject population that is very prone to sudden motion artifacts (e.g. children).

Slice time correction

- Slices in an fMRI volume are acquired all through the TR.
- The delay between the start of the TR and the actual slice acquisition time varies by individual slice.

Slice acquisition order - problem

- GLM analyses assume that all the slices in a volume were recorded at time zero.
- This assumption will generally be incorrect in all but one of the slices, sometimes very badly so.
- The GLM is very sensitive to time shifts
  – In event related task based studies, it will cause underestimation of the regressor betas in certain slices. (Sladsky et al found a decrease of parameter estimates of up to 63% (event-related design, average SOA=4s, TR=4s), and a reduction of 12% (block design, block length 10s, TR=4 s)).
  – In resting state analyses, slice timing differences may obscure connectivity between slices, especially with very long (6s) TR values.
- You should account for this, especially with longer TR's (>=2s or so).

Slice acquisition order - solution

- You can interpolate the data from each slice to the start of the TR (as long as you have fully sampled the signal you are interested in, this is legal).

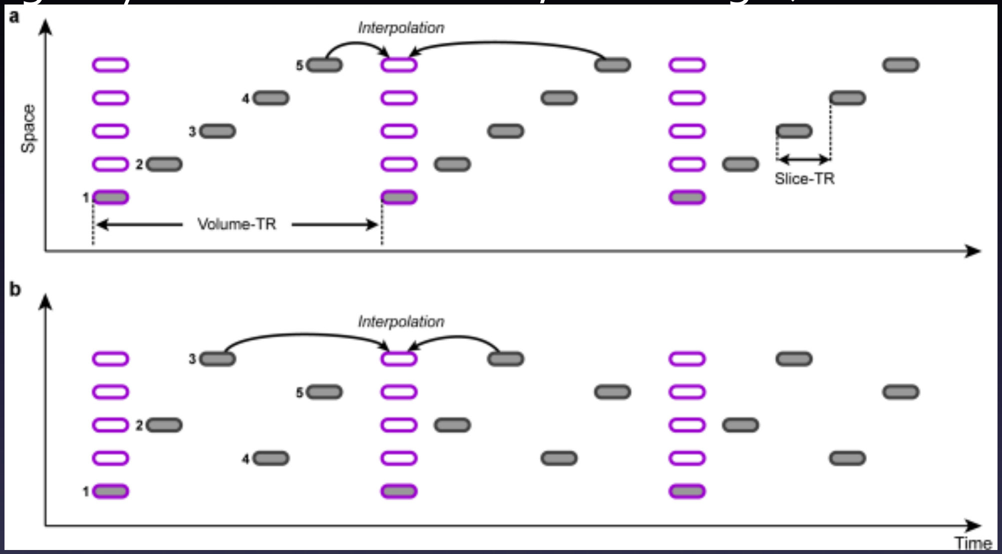

- The MOST IMPORTANT THING to remember with slice correction is to correctly specify the slice acquisition times – if you get it wrong, you can make the problem worse rather than better.

Figure courtesy of Rainer Goebel

Slice acquisition order

- There are many slice acquisition orders*.

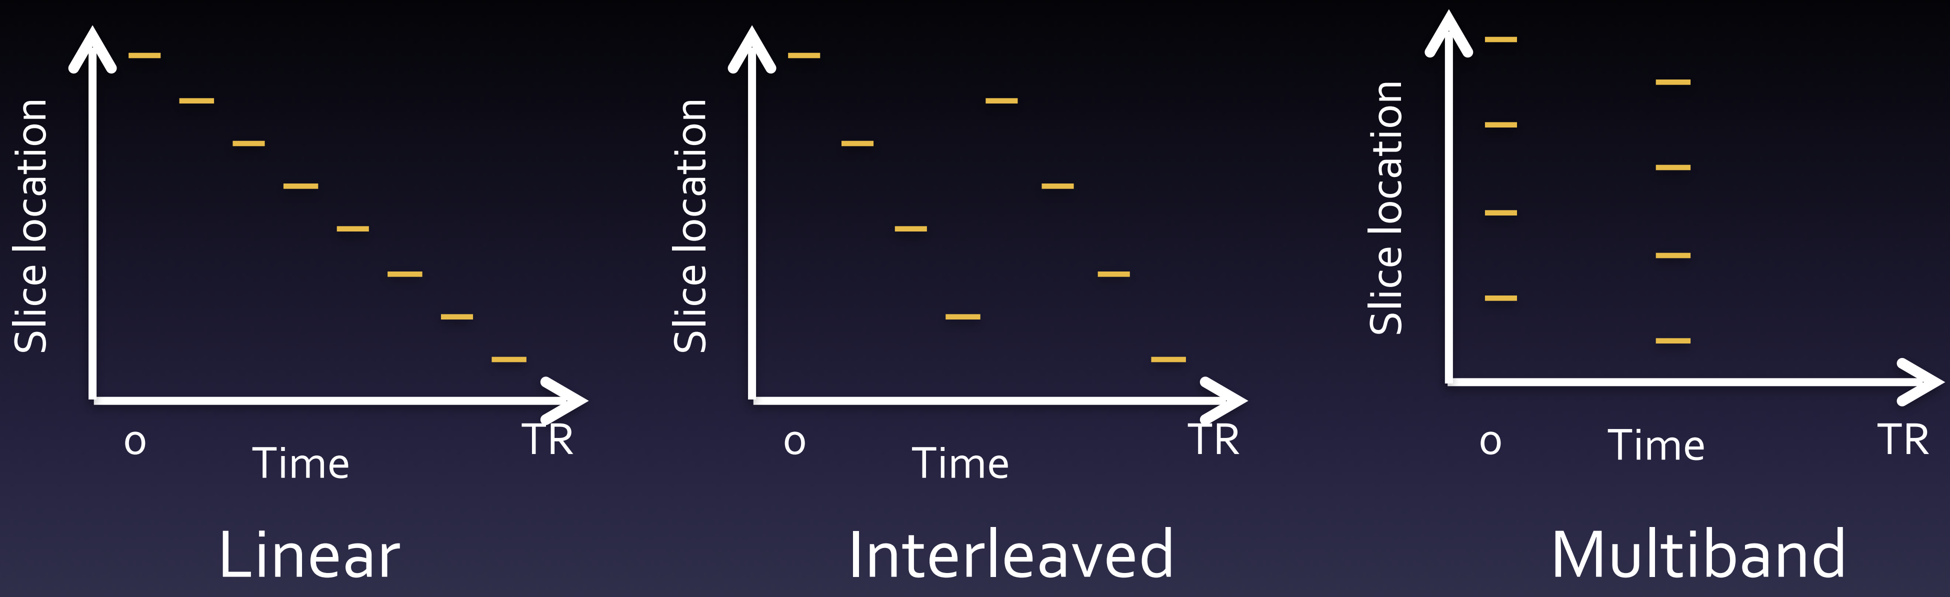

*There are also multiple variants (ascending, descending, odd or even start for interleaving, etc.). Rather than assume you have it right, it's best to check the DICOM headers.

Slice time correction - effect

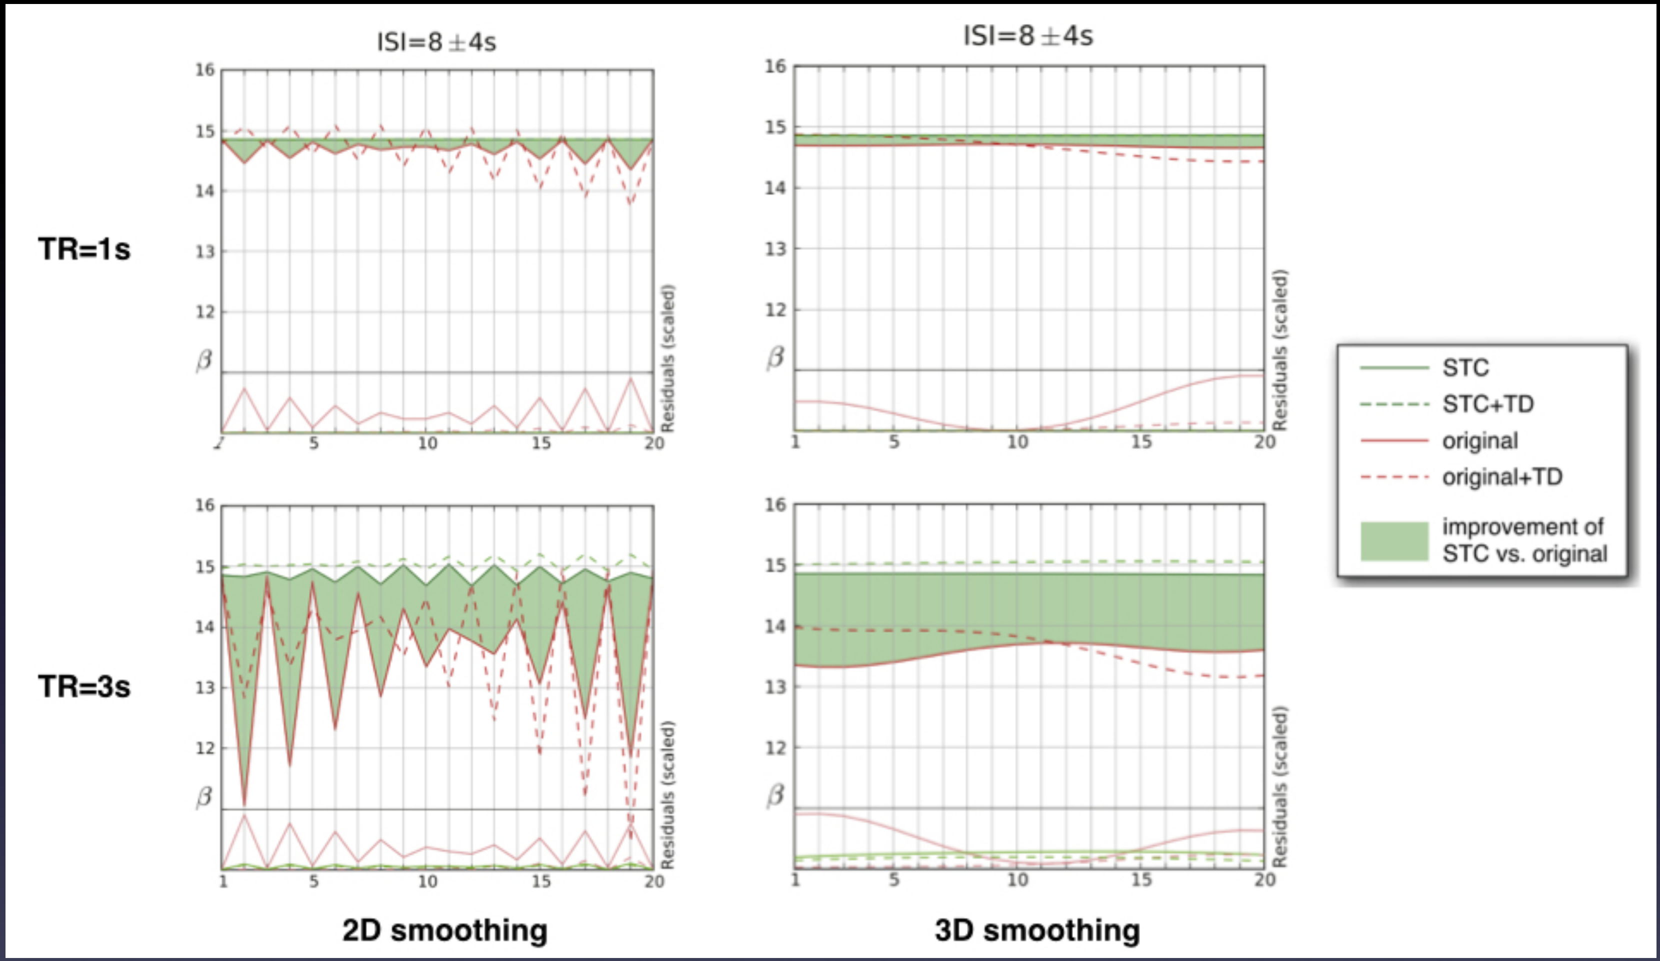

Figure adapted from Sladky, R. et al. Slice-timing effects and their correction in functional MRI. NeuroImage 58, 588–594 (2011).

Slice time correction summary

- Slice time correction prevents the underestimation of regressor betas by making sure regressors and data in all slices are properly aligned.
- While it doesn't make that big a difference for TR values of 1s or less, it can be extremely significant in event related designs with long TR values, and works better than using temporal derivatives of regressors alone.
- Since there isn't ever a downside to doing it (other than processing time)*, and sometimes a significant benefit, you should always do it.

* Or maybe there is for short responses and long TR's - opinions vary on this. Mark Jenkinson may say more…

Summary

- Temporal preprocessing is a necessary step in preparing your data for analysis.
- Filtering, spike removal, and slice time correction can each make a significant contribution to the power of your analysis (or more precisely, not doing them can significantly degrade your analysis).