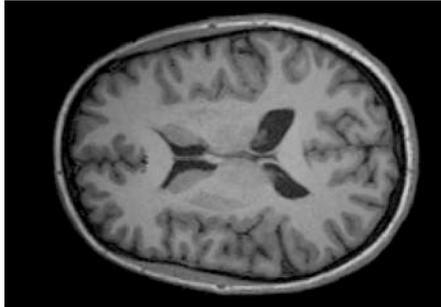


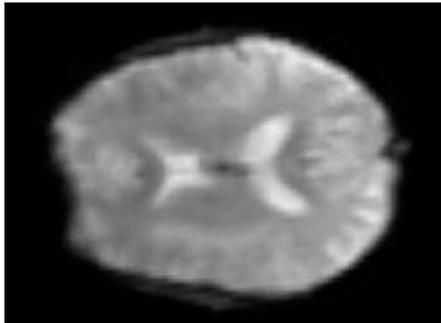
fMRI Basics: Spatial pre-processing

With thanks to Matthew Brett, Rik Henson, and the authors of Human Brain Function (2nd ed)

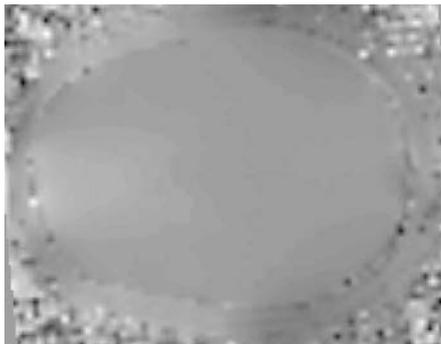
Most commonly collected types of image include:



Anatomical data: MPRAGE sequence - T1 weighted, high spatial resolution (usually 1x1x1 mm). Optimised for contrast between white and grey matter. Acquisition time ~5 minutes

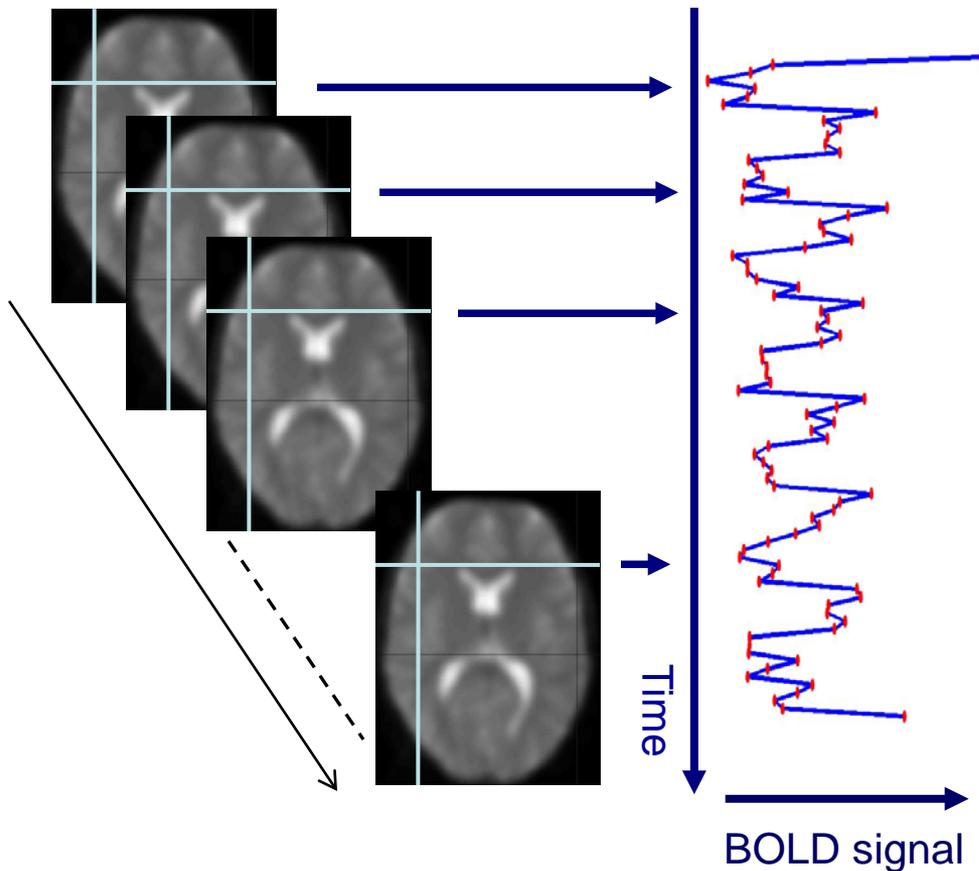


Functional data: Usually a T2* weighted gradient echo, echo planar imaging (EPI) sequence. Optimised for contrast between oxy- and deoxygenated blood. Fast acquisition (32 slices in ~2 seconds), reasonable resolution (3x3x3.75 mm)



Fieldmaps – map of magnetic field inhomogeneities within the scanner. Acquisition time ~1 minute.

Types of Data



Usually collect:

- 1 structural image

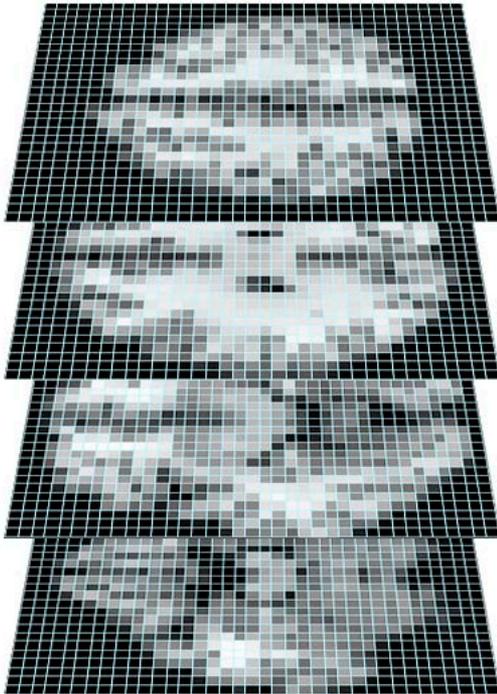
- 1 set of fieldmaps

- Many functional volumes

Repetition time (TR) for default EPI sequence = 2 seconds

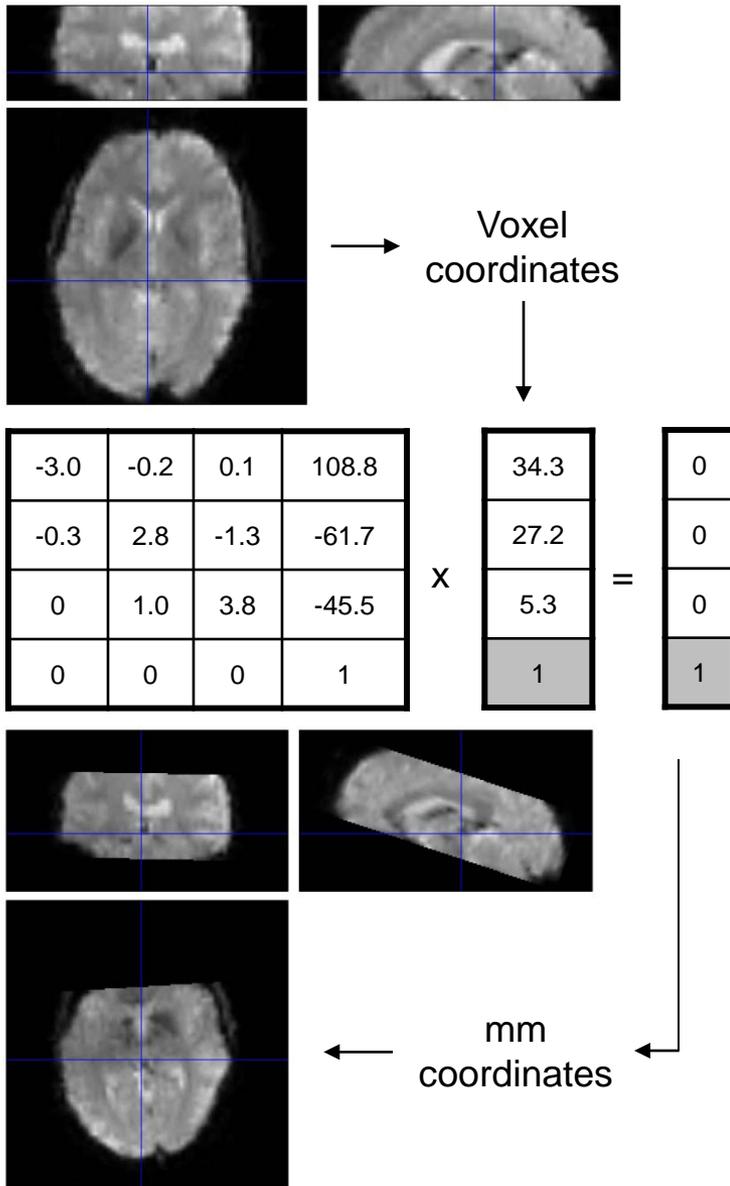
Build up a time series showing signal changes in each voxel in the brain

What's in an image?



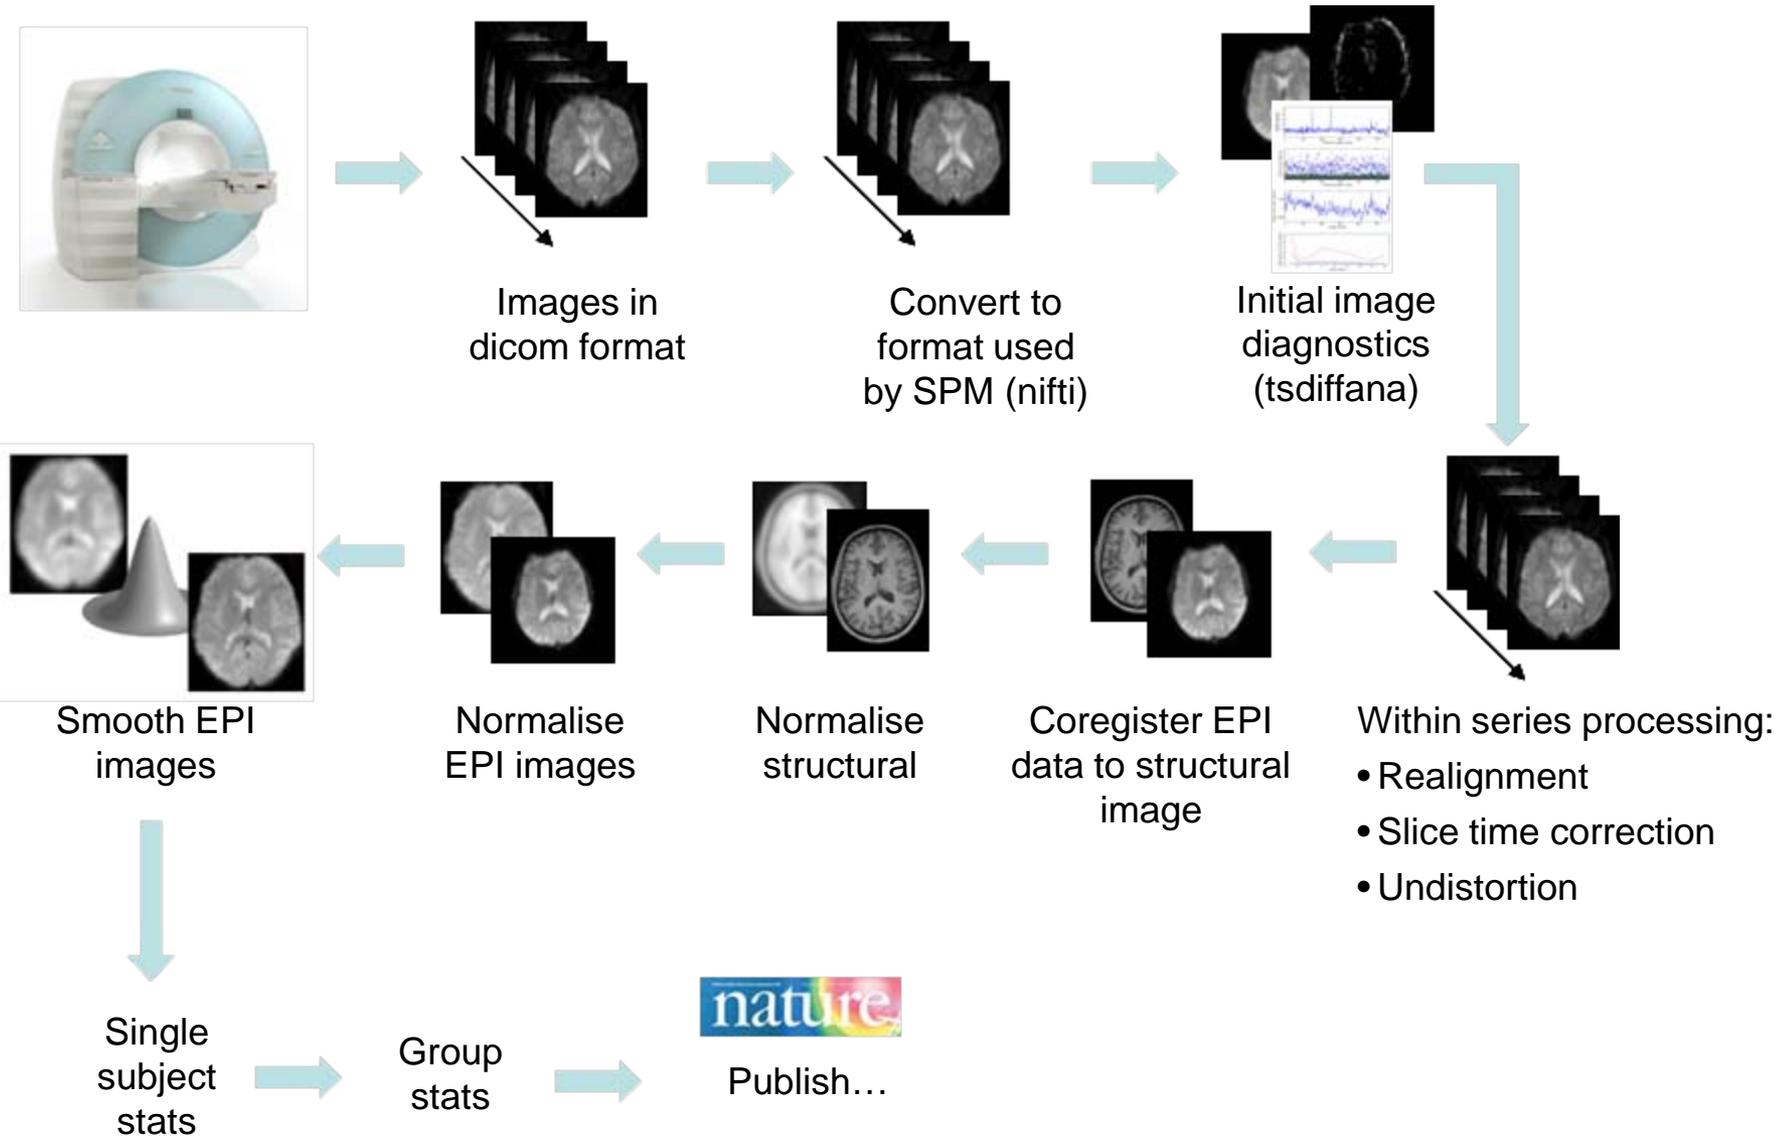
- Can think of an image as a 3D matrix of data
 - Divided into “voxels” (= volumetric pixels)
 - Each voxel contains a value showing the signal at a specific point in space and time
 - Data is accompanied by a header file that contains information about how to interpret the data:
 - Data type (integer, floating point etc)
 - Data scaling
 - Image dimensions
 - Voxel size
 - Voxel→mm transformation matrix
-
- Various different image formats - differ in the way data is stored and/or the way headers are organised
 - Most common here are dicom (.dcm, raw data from scanner), Analyze (.img / .hdr pairs), and Nifti (.nii)
 - SPM uses Analyze and Nifti formats

Voxel space vs. world space

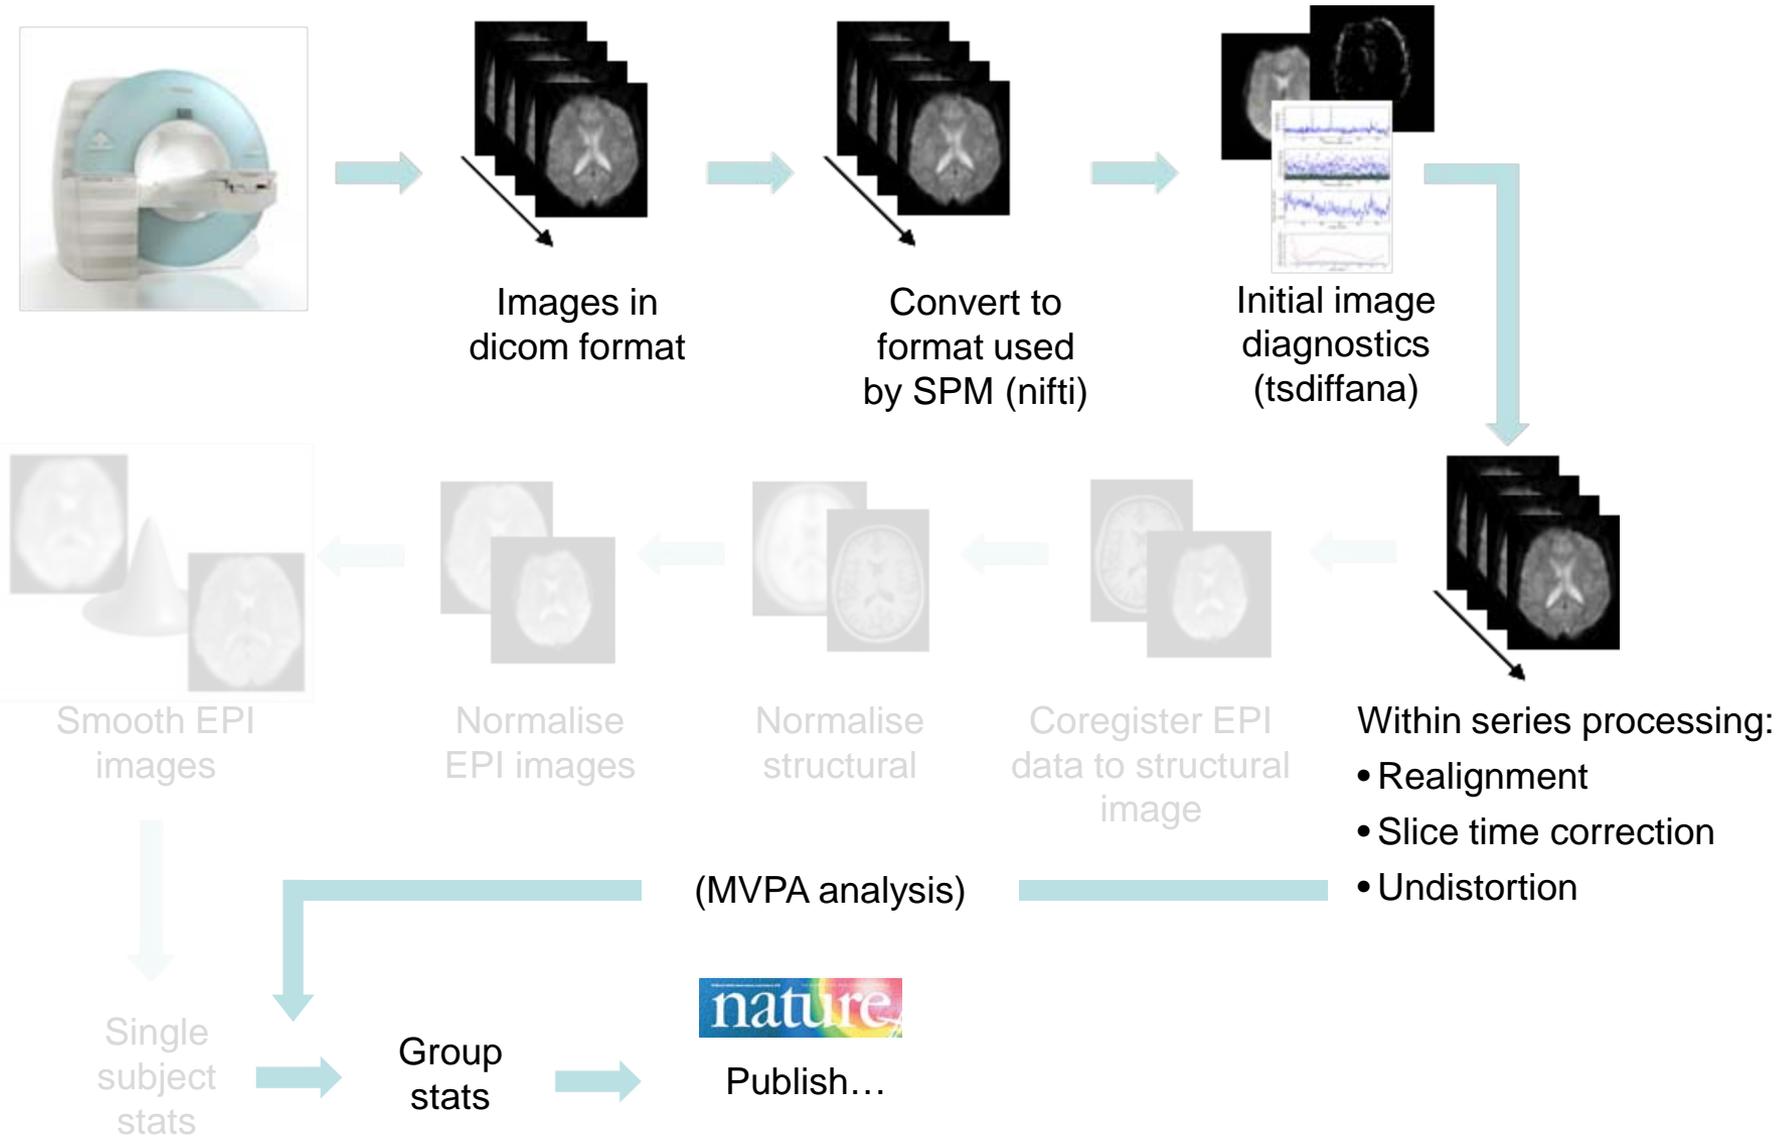


- Values at each voxel depend on where the acquisition grid is placed.
- Values in this “voxel space” are defined purely in terms of where they occur in the image
- Values in “world space” are defined in meaningful units (mm) from a point of origin
- Transformation matrix provides a way of encoding information about the spatial position of voxels
- Changing the transformation matrix changes the relationship between voxel co-ords and world co-ords
- The matrix can be used to store transformations without having to resample (reslice) the image

Overview of analysis

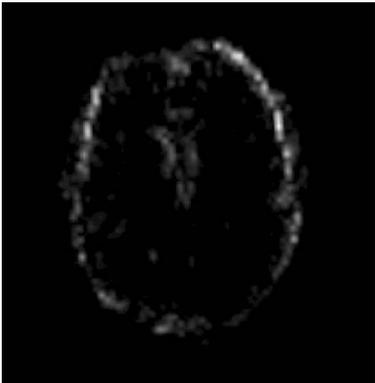
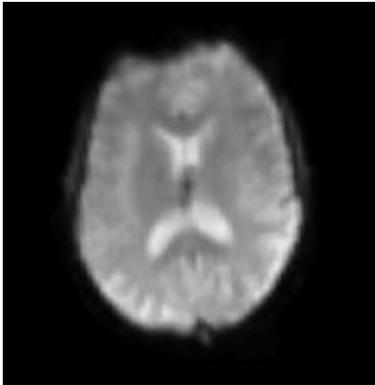


Overview of analysis



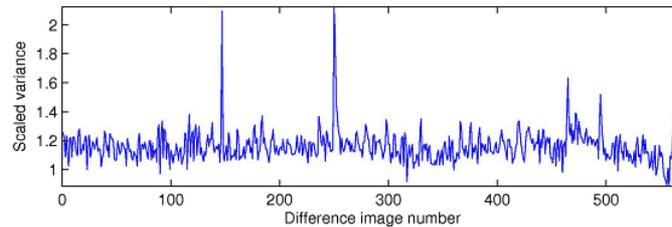
Initial diagnostics with tsdiffana

Mean and variance images:

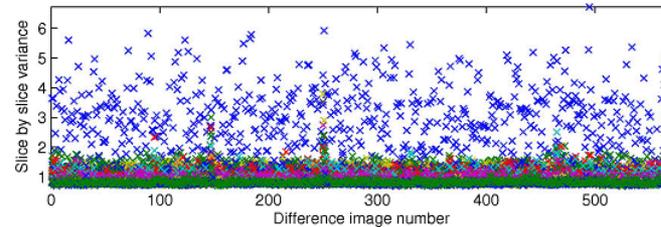


Look for obvious distortions + artefacts

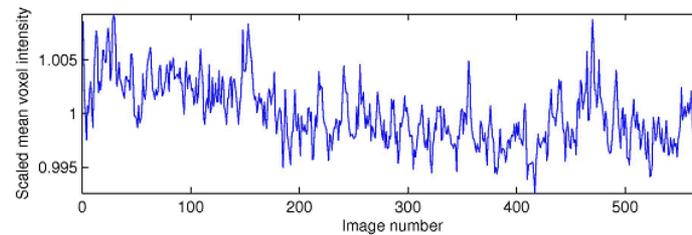
Diagnostic plots:



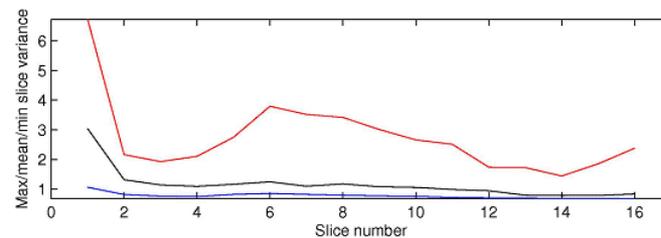
Scaled Variance



Slice by slice Variance



Scaled mean voxel intensity



Max / Min / Mean slice variance

Realignment – What & Why?

What?

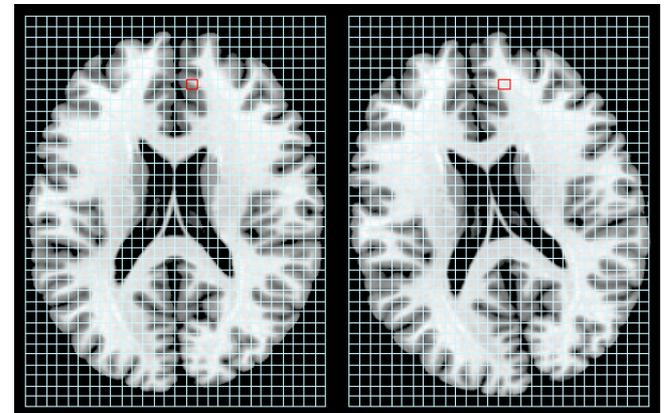
Within modality coregistration – usually this means realigning each of the images in a functional time series so that they're all in the same orientation

Why?

Because people move their heads...

This causes problems in several ways:

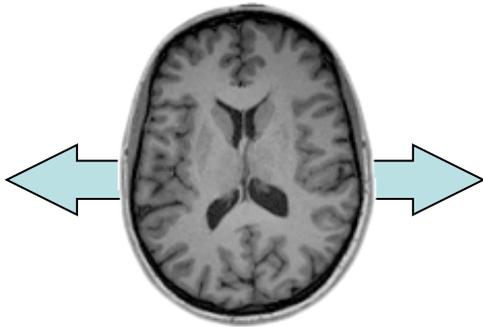
- Voxel contents change over time (e.g. from white matter to grey matter or vv), this can add considerable noise (unexplained variance) to the analysis.
- Interactions between head movements and inhomogeneities in the magnetic field – the magnetic field within the scanner isn't perfectly uniform and this can cause distortions which interact with head position.



Realignment – How?

Rigid body transformation using 6 parameters:

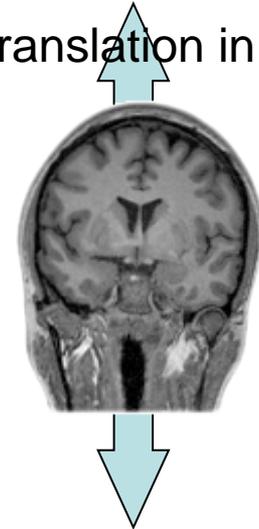
Translation in X



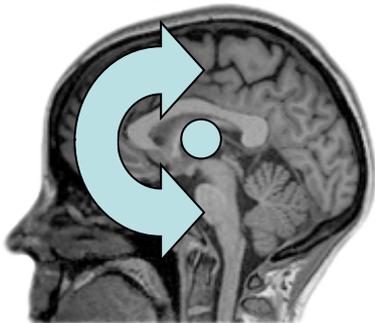
Translation in Y



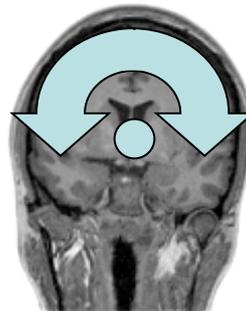
Translation in Z



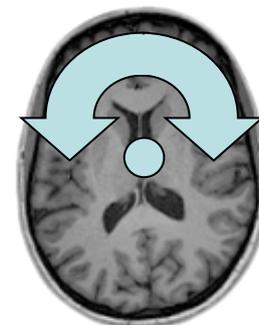
Rotation around X
(Pitch)



Rotation around Y
(Roll)



Rotation around Z
(Yaw)



Realignment

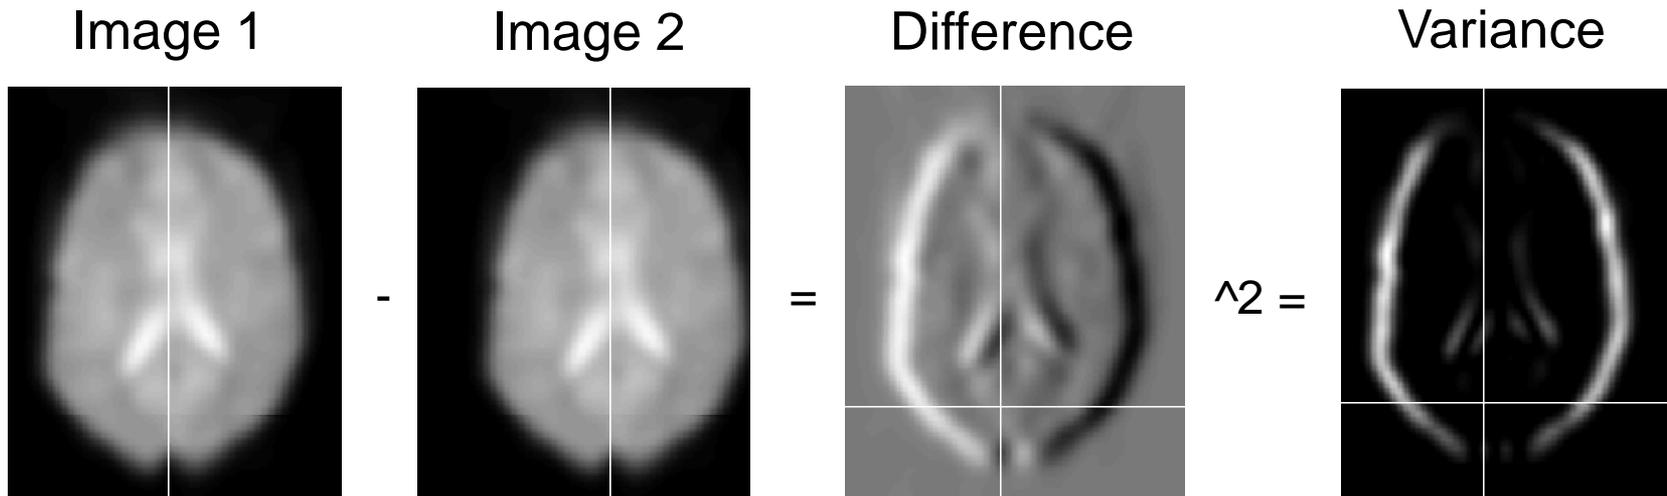
Find optimal values for these 6 parameters

Optimal = values that give the minimum value for a particular cost function

Cost function = sum of squared difference between consecutive images

Successive approximation - start with one set of parameters and iteratively try different combinations in order to find minimum sum of squared diffs

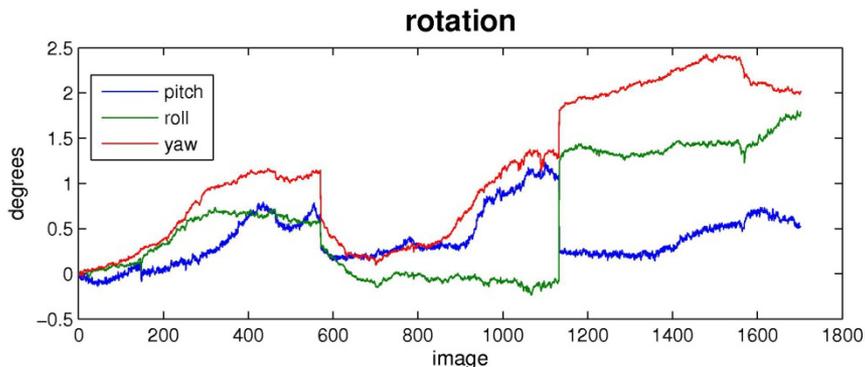
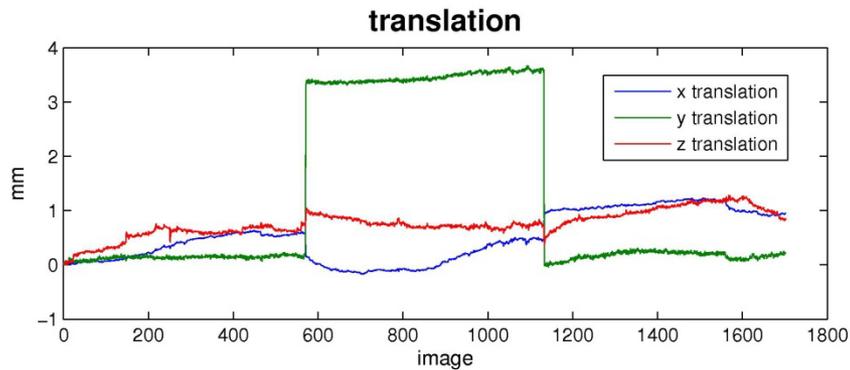
(Gauss-Newton algorithm provides a systematic way of modifying the parameters at each iteration)



Realignment – Results

Rigid body transformations parameterised by:

$$\begin{pmatrix} 1 & 0 & 0 & X_{\text{trans}} \\ 0 & 1 & 0 & Y_{\text{trans}} \\ 0 & 0 & 1 & Z_{\text{trans}} \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\Phi) & \sin(\Phi) & 0 \\ 0 & -\sin(\Phi) & \cos(\Phi) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} \cos(\Theta) & 0 & \sin(\Theta) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\Theta) & 0 & \cos(\Theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} \cos(\Omega) & \sin(\Omega) & 0 & 0 \\ -\sin(\Omega) & \cos(\Omega) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$



Transformation saved in the vox→mm matrix stored in the image header.

Image can also be “resliced”, i.e. resampled so that the transformation is permanently applied to the image

Like shifting the “voxel space” frame of reference.

Need to interpolate new voxel values

Process also displays graphs of the parameters giving an impression of movement throughout the experiment

Realignment – Possible issues

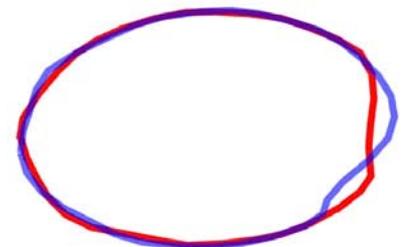
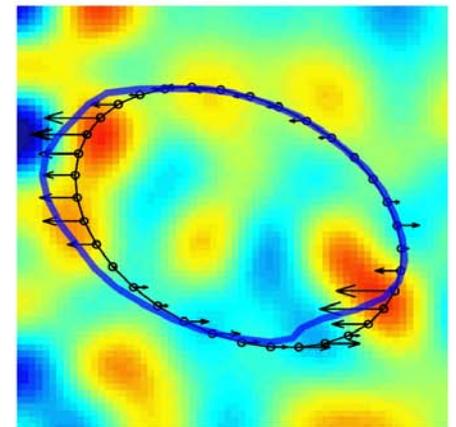
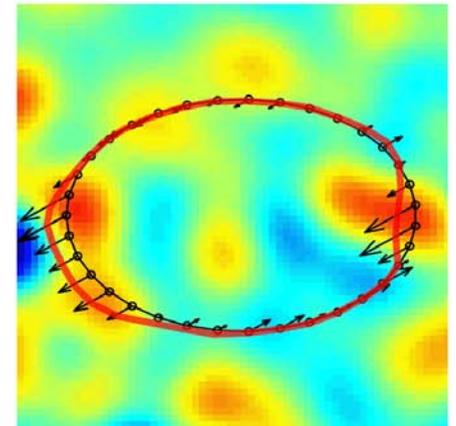
Realignment solves the voxel correspondence problem – the same voxel now contains the same bit of the brain over the entire time series

Doesn't solve all movement related problems though. In particular interactions between movement and field inhomogeneity remain.

Inhomogeneities in the magnetic field affect both signal strength and spatial encoding of signals, causing dropouts and distortions.

These effects are partially caused by the head, and are dependant on the position of the head. Distortions can be specific to a particular head position, so you can get non-rigid body movements.

The signal in each voxel could be recorded from a different position within the scanner, and potentially a different field strength, at different points in time.



Realignment – Possible issues

2 common solutions:

1. Include the realignment parameters as covariates in the statistical model
 - Idea is to capture any movement related variance in the data.
 - Can be problematic if movement is correlated with effects of interest (esp. button pushes, verbal responses etc)
 - If the movement parameters are correlated with your experimental conditions, they can remove the effects of interest.
2. Unwarping
 - Try to estimate the effects of interactions between field inhomogeneity and movement and compensate for them
 - Estimate rigid body parameters
 - Observe remaining variance
 - Estimate the “derivative fields” – how distortions at any point change with movement
 - Correct image to compensate for these
 - Re-estimate movement parameters
 - Iterate until minimal change

Realignment – Possible issues

Even after all this, movement artefacts still remain.

- There's no way of detecting rapid movements within a scan
- Spin history effects (movement may make the effective TR longer / shorter for some slices)
- Dropout by movement artefacts

The moral of the story? Stop people moving...

Make sure they're comfortable to begin with – encourage them to relax their neck and shoulders.

Discourage them from talking during breaks between sessions, be careful of using messages like “End of part 1” etc.

Reject any data with too much movement

Undistortion – what & why

What?

Deals with similar problem as unwarping – adjust images to correct for distortions caused by magnetic field inhomogeneities.

Why?

Unwarp does not actually remove the “static” distortions, it only estimates the interactions between distortions and movement (i.e. the first derivative, or the change of deformation with respect to movement). Unwarp will only undistort to some “average” distortion.

Undistortion attempts to correct for static distortions and return the image to something closer to the actual brain shape

Undistortion – how

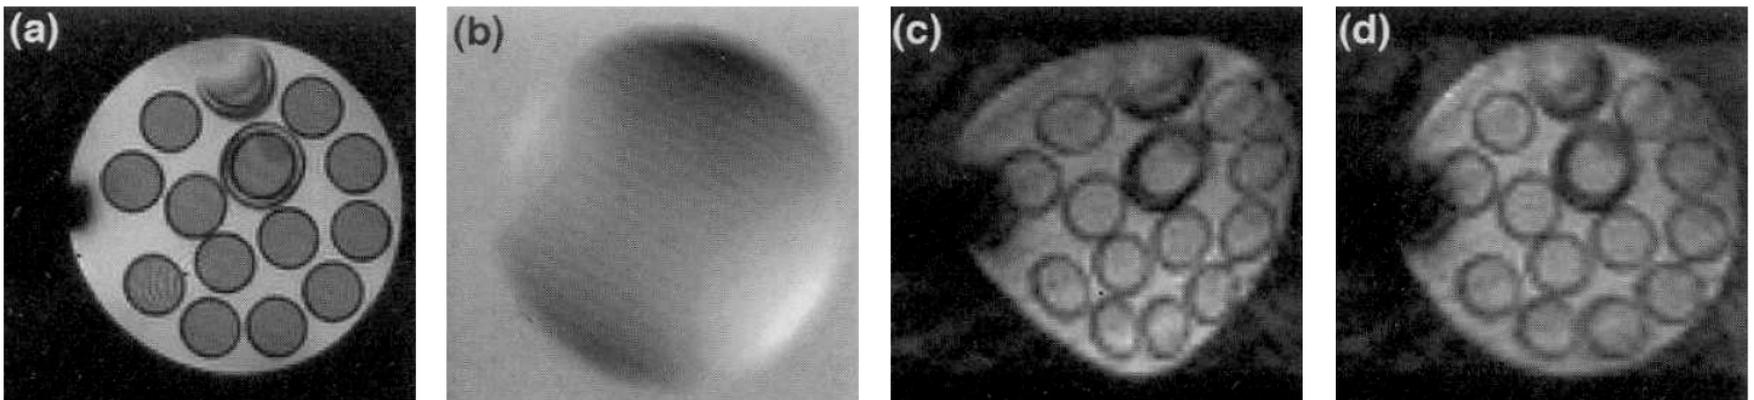
During scanning, collect a fieldmap – an image showing how the magnetic field varies within the scanner

This can be converted to a voxel displacement map – a map of how voxel values are displaced due to field inhomogeneities

This in turn can be used to calculate the original voxel values

NB – Usually only collect one set of fieldmaps which are specific to the head position at acquisition

Static undistortion can be combined with unwarping though



Slice time correction – what & why?

What?

Adjust the values in the image to make it appear that all voxels have been acquired at the same time

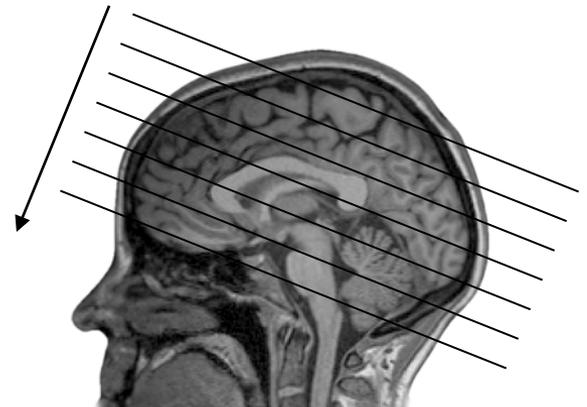
Why?

Most functional sequences collect data in discrete slices

Each slice is acquired at a different time

In an EPI sequence with 32 slices and a slice acquisition time of 62.5 ms, the signal in the last slice is acquired ~1.9 seconds after the first slice

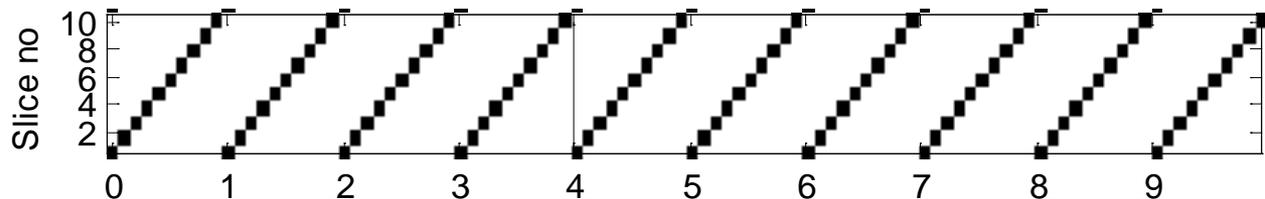
Problem if modelling rapid events (not necessarily such an issue in block designs)



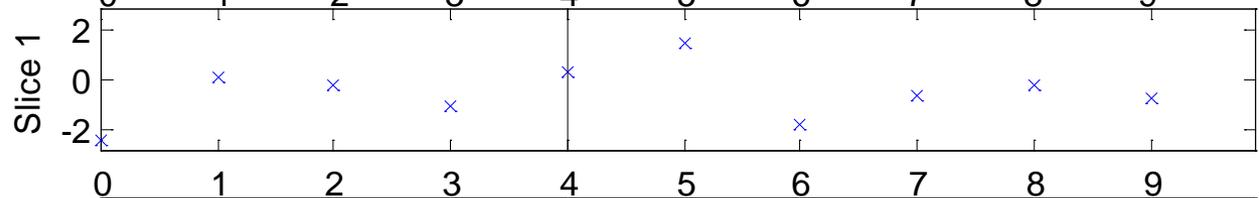
Slice time correction – how?

Create an interpolated time course for later slices
Shift each voxel's time course back in time

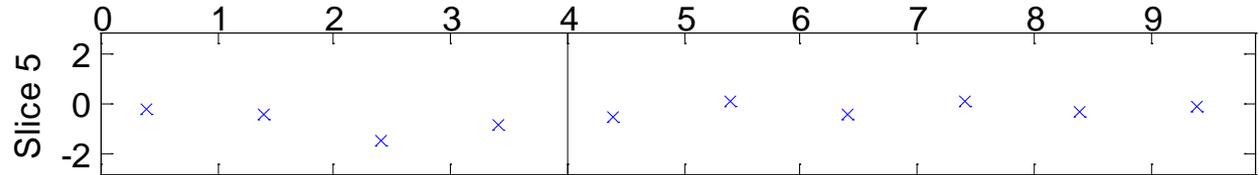
10 slice acquisition



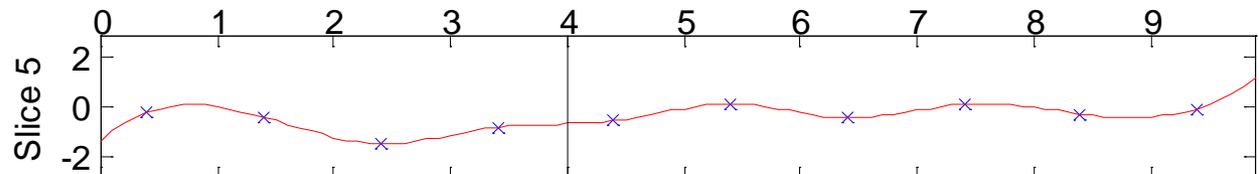
Time course of Voxel in slice 1



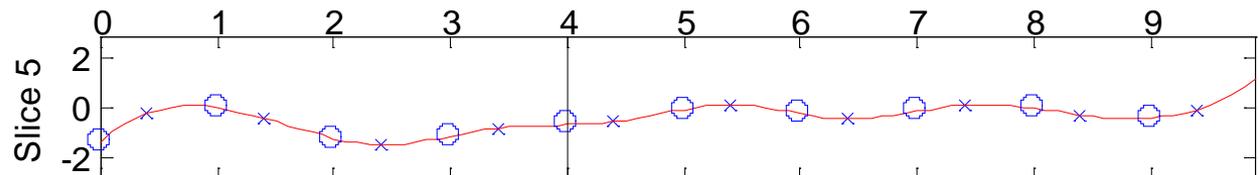
Time course of Voxel in slice 5



Interpolated time course in slice 5



Estimated value at time of first slice



Time (in TRs)

Slice time correction – possible issues

Which slice to align to?

- Using the middle rather than the first slice means the maximum interpolation necessary is reduced, which may reduce interpolation artefacts
- Possibly beneficial with long TRs
- Be careful to modify event onset times in statistical model though!

Coregistration – what & why?

What?

Cross modality registration – realigning images collected using different acquisition sequences. Most commonly registering T1 weighted structural image to T2* weighted functional images.

Why?

Head movement again...

Precursor to spatial normalisation

Often better to normalise the structural image (higher spatial resolution, fewer artefacts and distortions) and then apply the parameters to the functional data.

So, want the structural in the same space as the functional images

Coregistration – how?

- Similar to realignment - find parameters for translations in X, Y, and Z, and rotations around X, Y, and Z
- BUT - different acquisition sequences have different properties, e.g. CSF is bright in T2 functional images, dark in T1 structural images
- Can't simply subtract images and minimise the squared difference
- Have to use another cost function - “Mutual Information”
- How much does knowing the value of one variable (e.g. T1 intensity) tell us about the possible values of another variable (e.g. T2 intensity)

Joint histograms of X, Y:

Value of Y	1	1	1	1
	1	1	1	1
	1	1	1	1
	1	1	1	1

Value of X

None

0	0	0	4
0	0	4	0
0	4	0	0
4	0	0	0

Value of X

High

2	0	0	2
0	2	2	0
0	2	2	0
2	0	0	2

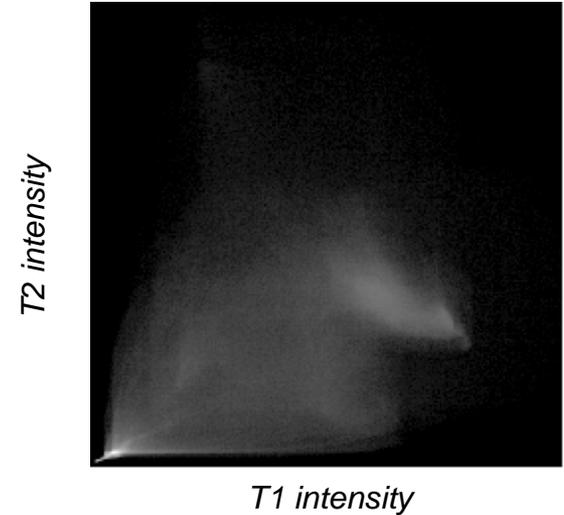
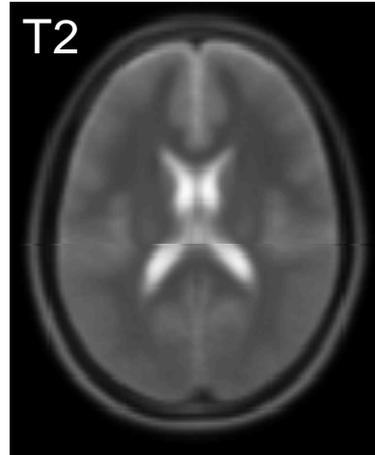
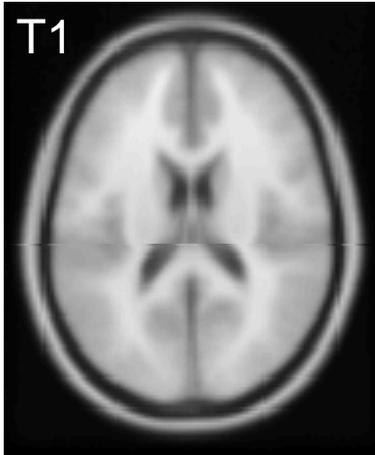
Value of X

Some

$$I(X;Y) = \sum_{y \in Y} \sum_{x \in X} p(x,y) \log \left(\frac{p(x,y)}{p_1(x)p_2(y)} \right)$$

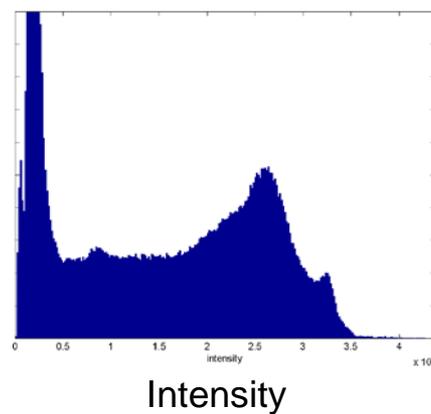
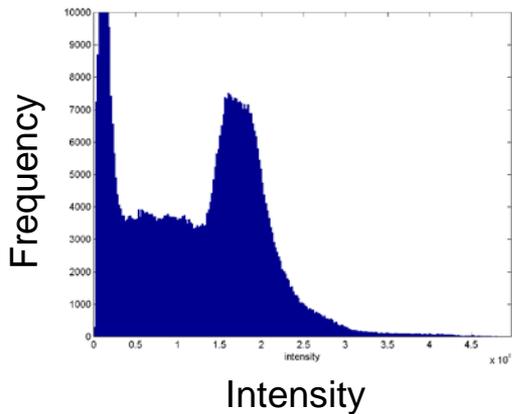
Coregistration – how?

Joint histograms pre...

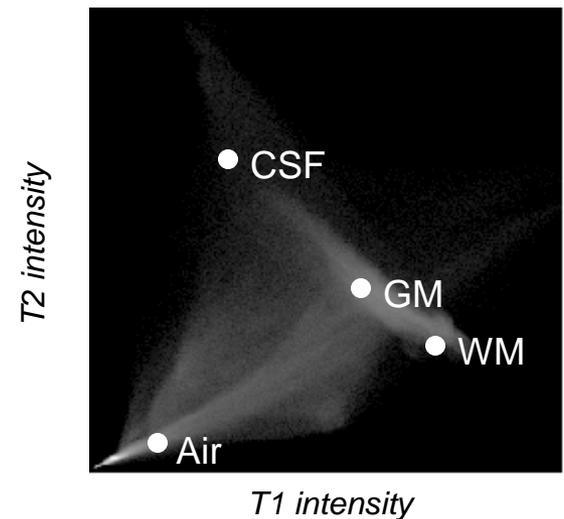


T1 histogram

T2 histogram



...and post registration



Normalisation – what & why

What?

Registration between different brains. Transforming one brain so its shape matches that of a different brain.

Why?

People have different shaped brains...

Allows group analyses since the data from multiple subjects is transformed into the same space

Facilitates cross study comparisons since activation co-ordinates can be reported in a standard space (rather than trying to identify landmarks in each individual study)

Normalisation – different approaches

Landmark matching

- try to identify, then align homologous anatomical features in different brains, e.g. major sulci.
- Time consuming and potentially subjective – manual identification of features.

Intensity matching

- Minimise differences in voxel intensity between different brains
- More easily automated – like realignment and coregistration, can assign some cost function based on differences in image intensity, then find parameters that minimise this cost function.

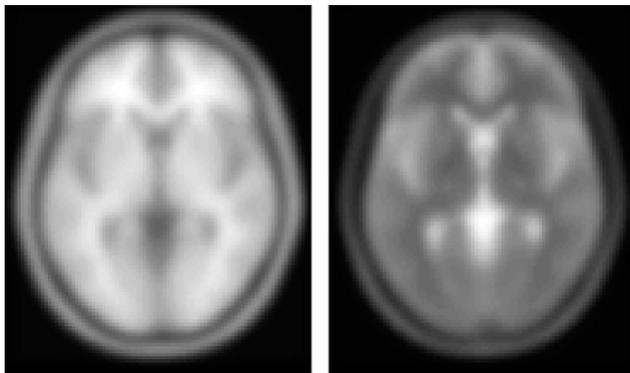
Normalisation – how

SPM uses a procedure that attempts to minimise the differences between an image and a template space

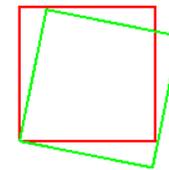
Like realignment, start with affine (linear) transformations.

As well as the 3 translations and 3 rotations, also apply 3 zooms and 3 shears.

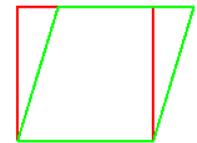
This matches the overall size and position of the images, but not necessarily differences in shape



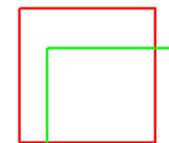
MNI T1 (left) and T2 templates



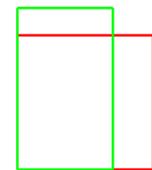
Rotation



Shear

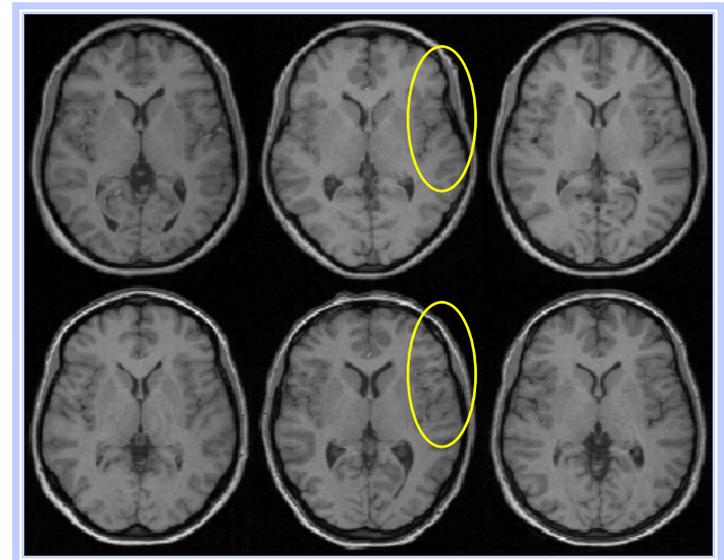


Translation



Zoom

*6 images
registered to the
MNI template
using only affine
transformations*

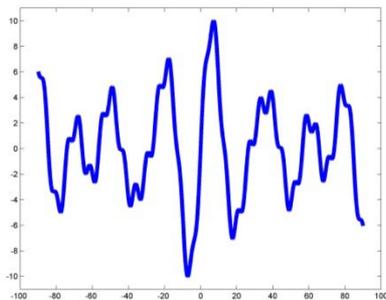


Normalisation – how

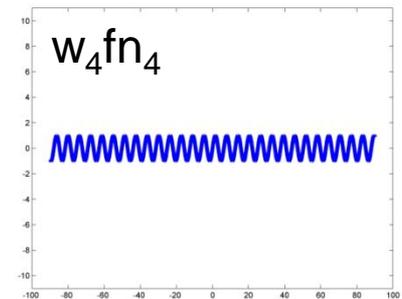
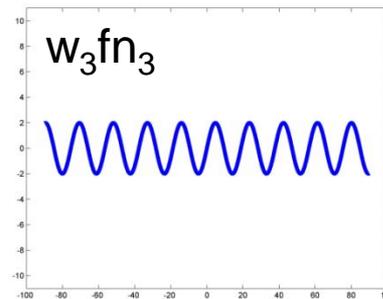
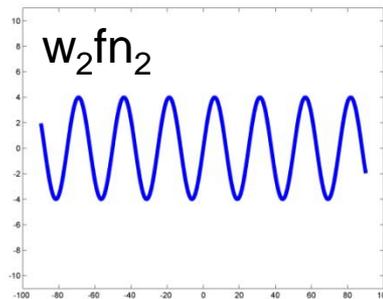
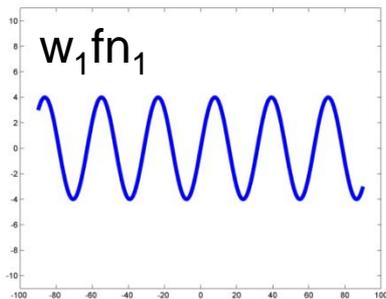
Next, apply nonlinear transformations

A quick digression into basis functions...

A complex function can be described as a linear combination of a set of simpler basis functions:



$$= \sum w_i f_n_i$$



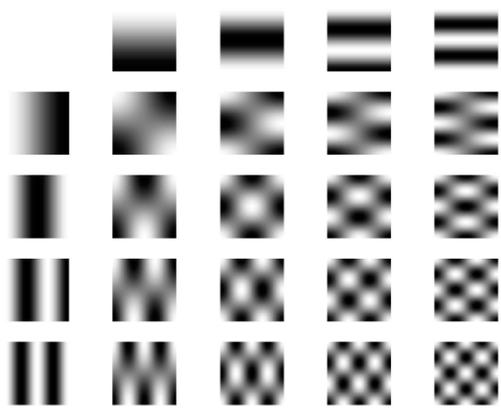
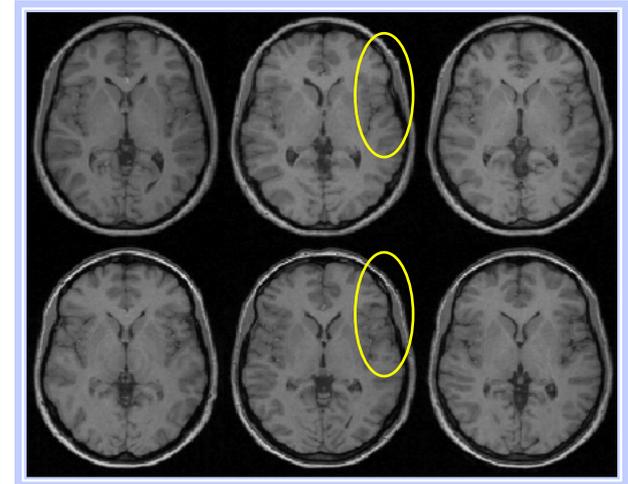
Normalisation – how

Nonlinear transformations implemented by applying deformation fields

These are modelled using a linear combination of cosine basis images

Matches size, position and global shape of template.

6 images registered to the MNI template using linear and nonlinear transformations



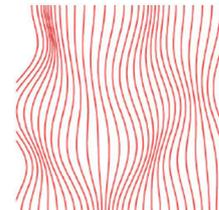
Cosine basis images

Weighted combination of basis images

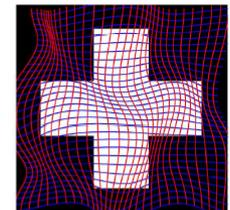
Dark – shift left, Light – shift right



Deformation Field in X



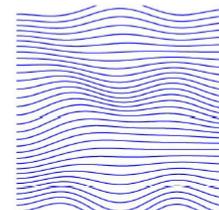
Field Applied To Image



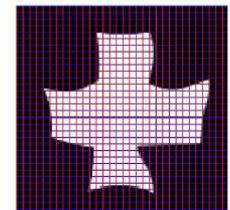
Dark – shift down, Light – shift up



Deformation Field in Y



Deformed Image



SPM Algorithm simultaneously minimises:

- Sum of squared difference between template and object image
- Squared distance between the parameters and their expected values

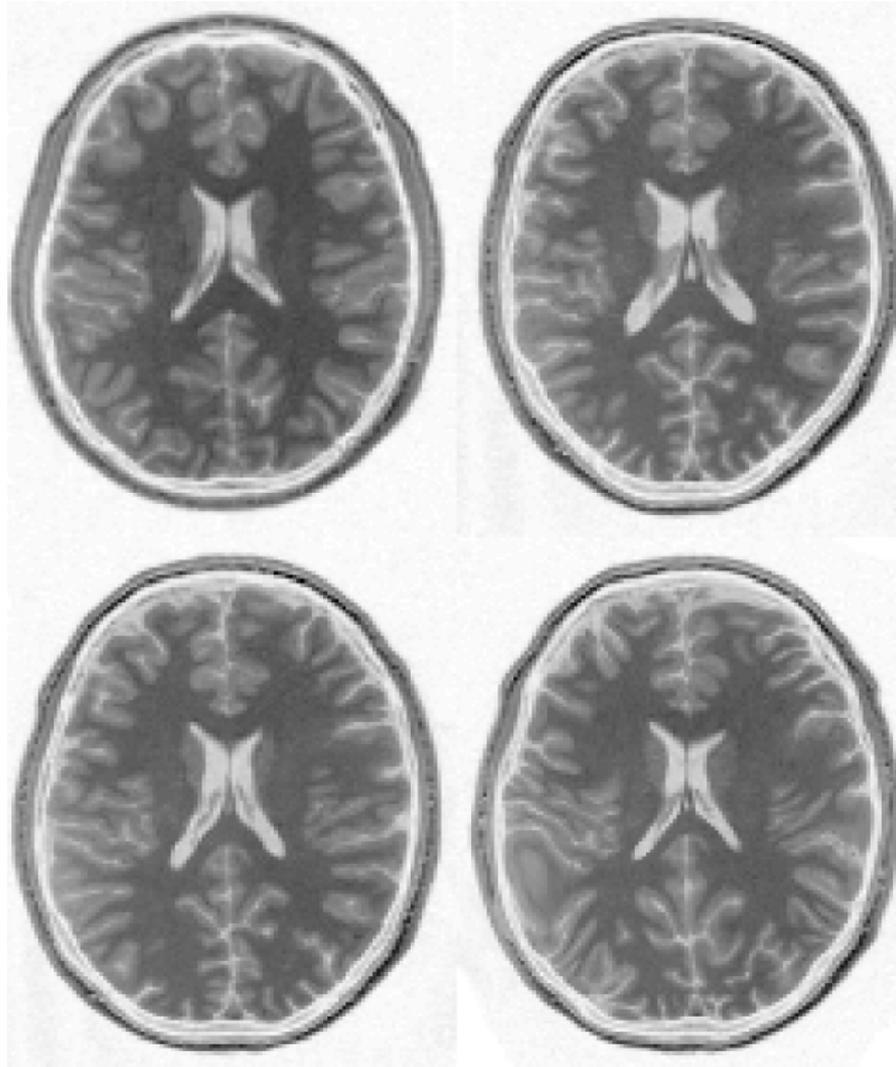
The latter condition is referred to as “regularisation”

Essentially a way of incorporating prior knowledge about the range of values that parameters can take in order to constrain current estimates.

Helps reduce unnecessary distortions (an example of overfitting due to the large number of available parameters)

Normalisation – how

Template



Affine only – still some differences in shape

Affine and nonlinear with regularisation – good match to overall shape, but some high spatial frequency differences

Affine and nonlinear without regularisation. This can introduce overfitting and unnecessary warps

Common templates:

- Talairach and Tournoux, 1988 (detailed anatomical study of a single subject...)
- Montreal Neurological Institute 152 (MNI152; averaged from T1 MRI images of 152 subjects)
- Similar, but not identical
- SPM uses MNI152 template
- To report co-ordinates in Talairach space, have to convert using something like `mni2tal.m`

Smoothing – what & why

What

Spatial averaging - replace the value at each voxel with a weighted average of the values in surrounding voxels

Why

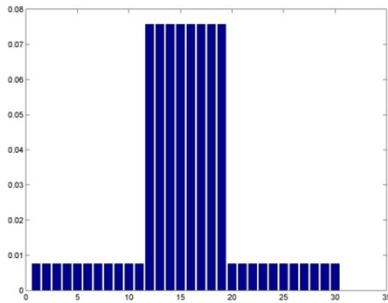
Increase signal to noise

Random noise tends to be reduced by the process of averaging (since it's a mixture of high and low values)

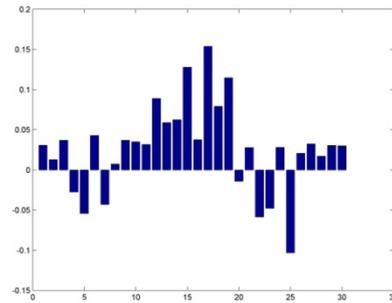
Smoothing – how

Apply a smoothing “kernel” to each voxel in turn, replacing the value in that voxel with a weighted average of the values in surrounding voxels

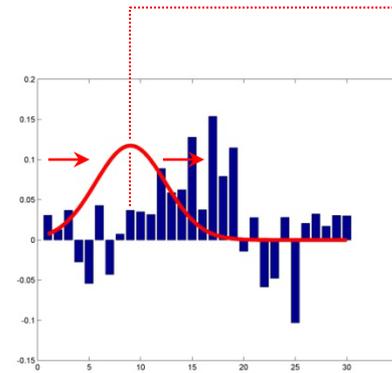
The kernel is simply a function that defines how surrounding voxels contribute to the weighted average



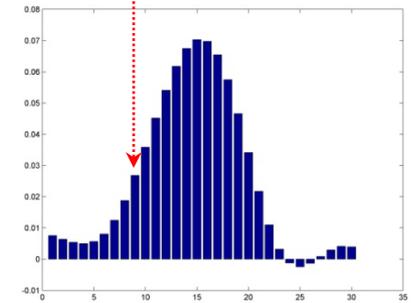
Original signal



Signal plus noise



Apply kernel to each point in turn



Recover signal

Smoothing – how

Which kernel?

Ideally, want a kernel that matches the spatial properties of the signal

“Matched filter theorem”

In practice, usually use a 3D Gaussian

Shape defined by Full Width at Half Maximum height (FWHM)

Usually don't know the spatial extent of the signal

Can make some assumptions though – e.g. if looking at specific visual areas a smaller kernel may be optimal, whereas if looking at prefrontal, a larger kernel may be best

In practice, 8-10mm is common

