Pre-Processing of BOLD FMRI Data

Mark Jenkinson and Stephen M. Smith

FMRIB (Oxford University Centre for Functional MRI of the Brain)

1 Introduction

Pre-processing is necessary in fMRI analysis in order to take raw data from the scanner and prepare it for statistical analysis. This talk describes the various pre-processing steps which occur prior to the the statistical analysis (which estimates where significant activation occurred, and is described in the following talk). The pre-processing steps take the raw MR data and apply various image and signal processing techniques to it that reduce noise and artefacts. These steps are crucial in making the statistical analysis valid and greatly improve the power of the subsequent analyses.

The usual method of pre-processing is to apply a sequence of individual processing steps in a particular order. However, it is increasingly being recognized that many of these steps are not independent and that some power is lost in the statistical analysis by ignoring the interactions between these pre-processing steps. Consequently, there is a move towards developing more integrated approaches to pre-processing, and ultimately it is likely that more advanced analysis techniques will incorporate both these pre-processing steps and the statistical analysis within one unified framework. At present, though, it is still common to have separate, independent pre-processing steps, each of which can have considerable impact on the later statistical analysis and so need to be understood and considered carefully. The following sections will describe these individual pre-processing steps, including information on why they are important, how they are implemented and what impact they have on subsequent analysis stages.

2 Reconstruction from K-Space Data

The raw MR signal is obtained by digitizing the demodulated RF signal detected by the receiver coil. This raw data does not resemble a typical image, but instead is "k-space" data – that is, a spatial frequency transformation of the familiar image-space. In order to reconstruct the k-space data into image-space, so that the image may be viewed and analysed, an inverse transformation is required. This is usually performed using a Fast Fourier Transform if the data is sampled in a regular grid in k-space, whereas for sequences like spiral-EPI it is necessary to either re-grid the data prior to transformation, or to use a different form of transformation. In addition, fMRI data acquisitions normally occur one slice at a time and so the Fourier transform is performed in 2D on each slice independently, although other types of images, such as T1-weighted structural images, often use MR sequences that require 3D Fourier transforms.

As MR images, and particularly EPI images, are susceptible to various acquisition-related artefacts, the reconstruction stage often includes some form of artefact correction technique. For example, standard EPI data usually contains an artefact known as the *N/2 ghost* (or Nyquist ghost). This artefact manifests itself as a lower intensity replication of the main image, shifted by half the field-of-view in the phase-encode direction. By using additional phase information, acquired in a separate reference image, reconstruction techniques can perform *ghost correction* which typically reduces the intensity of the ghosted image to less than 5% of the main image intensity. For more detail on this and other artefact correction techniques in EPI data, see [18].

3 Motion Correction

If a subject moves their head during an fMRI session, the position of the brain within the functional images will vary over time. This means that any particular voxel's time series does not (over time) refer to the same point in the brain. This motion can vary from small, sub-voxel motion to large, obvious motion. However, even sub-voxel motion has a detrimental effect on the statistical analysis – especially for smaller activations and in areas near major intensity boundaries. Therefore, motion correction is almost always desirable in fMRI experiments, and is often of crucial importance.

Motion correction amounts to finding a common orientation for all images within a given session and resampling the original data to this reference orientation. This is usually achieved by performing a separate 3D image registration of each image in the series with a chosen reference image. For example, if the reference image is chosen as image 1, then every image in the series (from image 2 to image N) will be registered and resampled to be in the same orientation as image 1. As all images in an fMRI experiment are of the same object, taken with the same MR sequence, rigid-body spatial transformations (rotation and translation only) along with intra-modal voxel-similarity functions (like normalised correlation) are typically used to model the change between one image and the next, as described below.

In addition to rigid-body motion, there are also some sources of non-rigid motion that are usually present. For example, there is pulsatile motion of the soft brain tissues during the cardiac cycle. Even bulk motion of major chest organs during respiration will change the magnetic (B_0) field distribution throughout the body (including the head) and will therefore affect the geometry of the images, inducing non-rigid motion. It is possible to reduce the extent of some of these motions by using methods such as cardiac gating of the images (acquiring at the same point in the cardiac cycle each time). However, the major component of motion is due to rigid movement of the head in the scanner. For more detail on physiological noise, see [13, 14].

3.1 Using Registration for Motion Correction

The standard approach taken in motion correction is to split the problem into a series of separate 3D registration problems, as described above. These registrations usually employ a within-modality voxel similarity measure, like mean-square-difference or normalised-correlation, to quantitatively determine when the images are in good alignment. This, together with a standard local optimisation algorithm, is sufficient for most motion correction problems since the images are all taken with the same MR sequence (giving the same tissue intensities and contrast) and do not have large changes in position (compared with changes between sessions). Such registrations can be performed quite rapidly, since the resolution of the images is relatively low and only 6 degrees-of-freedom (DOF) transformations are used (i.e. rigid body). For instance, a 200-image series can typically be motion corrected in less than 15 minutes – that is, less than 5 seconds per image.

In this standard approach to motion correction, it is necessary to choose a *reference* image, with which all other images will be aligned. Normally this is simply an arbitrary single image from the original data. However, care needs to be exercised when choosing this reference, as the first image from the sequence can look quite different to all subsequent images, due to MR saturation effects (if the initial images have not been removed). Also, if within-scan motion blurring is significant in this image or large activation signals are present, then it is a poor candidate for the reference. Furthermore, to minimise the amount of interpolation required to resample the images it is desirable to choose a reference image in an "average" position, which for long, slow movements (which are relatively common) is likely to occur near the middle of the session. There is no consensus on which is the "best" choice for the reference image, although it is likely to be less important in the future as more sophisticated methods of motion correction are developed that do not separate out a single reference, but instead use all of the time series information simultaneously [3, 4, 16].

3.2 Additional Motion Artefacts

Simple alignment of the images, however, cannot remove all the effects of motion on the images, and intensity-based motion artefacts will still remain in the images [11]. These artefacts are due to several causes such as within-scan motion, between-slice movement, and induced B_0 field changes, as well as the effect of interpolation during resampling in the motion correction stage. Such artefacts introduce extra intensity variation into the time series which affects the later statistical analysis by either adding extra unwanted noise (making true activations harder to find) or, more problematically, mimicking activation signals (leading to the detection of false activations). The latter situation arises most commonly when there is *Stimulus Correlated Motion* – that is, motion which is correlated, or synchronous, with the applied stimulus.

One of the most discussed motion artefacts is the *spin-history effect* [9]. This occurs when there is between-slice movement such that certain nuclei in the object go from one slice to another, resulting in a change in the timing of the RF excitations they receive. Such timing changes are significant as they depart from the stable saturation cycle that the nuclei are in when there is no motion. Therefore, the signal will be artefactually increased or decreased (depending on the direction of motion), displaying saturation-like effects similar to those seen in the first few images of a scanning sequence.

Another artefact that has received attention recently is the interaction of motion and B_0 field inhomogeneity. Field inhomogeneities are the result of differences in tissue susceptibilities in the head and cause significant local signal loss and geometric distortion in EPI which varies with the orientation of the head. This effect can be reduced by modelling or estimating the field inhomogeneities (e.g. see [3]).

One way to deal with general motion-related intensity variations is to remove all trends from the time series that have the same form as the voxel displacement (as measured during the transformation estimation stage in motion correction). This assumes that the artefact will be proportional to the displacement of the voxel from its usual position. The removal of the trend can be achieved in either the pre-processing stage (by decorrelation of the data) or in the later statistical analysis (by using these displacements as regressors of no interest – or confounds – see [5, 9]). However, for strongly stimulus-correlated motion, there is little difference between true activation signals and the motion displacements and so it is difficult, if not impossible, to distinguish between the two. Consequently, in this situation, the removal of such trends has the undesirable effect of also removing a substantial proportion of the activation signals themselves.

Until recently, attempts at motion artefact correction have treated the various effects separately. However, as there is significant interaction amongst the various artefacts, and the artefacts have significant impact on the accuracy of the basic motion correction methods, more research is focusing on simultaneous motion correction and artefact removal methods (e.g. see [3]).

4 Slice Timing Correction

Functional volumes are normally acquired one slice at a time with the timing of the slice acquisition evenly spread over the repetition time (TR), which is typically a few seconds. Therefore it is not correct to assume that all slices were captured at the same time, or that the timing in all voxels is the same, which is a common assumption made in the later statistical analysis. If this change in timing is not correctly accounted for then the relative timing of the stimulus and response will not be matched and consequently the statistical analysis will not be able to fit the model with optimal accuracy.

Slice timing correction aims to adjust the voxel time series so that a common reference timing exists for all voxels. The reference time is often chosen as that corresponding to the first slice. The temporal adjustment is usually achieved by shifting the time series of values slightly forward or backward in time (as all corrections are less than one TR), using some form of interpolation. For signals with a limited range of frequencies (band-limited) the optimal interpolation method is sinc interpolation – that is, convolution in time with a sinc function kernel. However, not all measured signals will be within the correct frequency range for typical fMRI experiments and so other forms of interpolation may be better, and other options are available.

One complication in the above formulation is if head motion occurs such that locations in the brain move between slices. The signal from these locations will therefore be incorrectly located both in space and in time. Motion correction attempts to correct for the global spatial locations but usually assumes a simple rigid-body motion has occurred and cannot correct for non-parallel slice acquisitions. Similarly, slice timing correction methods usually assume that all voxels in a slice have the same timing and cannot correct for spatial variability within the slice. Ideally a combined approach to both slice timing and motion correction is required (e.g. see [4]), although such techniques are still very new and not in common use. In the meantime, it is necessary to correct for the two effects separately, where the most common choice is to apply slice timing correction *after* motion correction since the interpolation in time is normally a second order correction compared to the spatial interpolation of intensities near tissue boundaries.

5 Spatial Filtering

The next stage of fMRI analysis is the spatial filtering (or “blurring”) of each volume. There are two reasons for applying spatial filtering as a pre-processing step; firstly, blurring can increase signal-to-noise ratio in the data, and secondly, certain later statistical steps, in order to be valid, require the functional images to be spatially smooth.

5.1 Reasons for Spatial Filtering

The Signal-to-Noise Ratio (SNR) is a measure of the amplitude of signal of interest compared with the noise level (standard deviation). The signal of interest, in this case, is the change in image intensity resulting from the stimulus induced activation via the BOLD effect. The noise is the unavoidable random variations in image intensity which are present even when no stimulation is applied. Typically, the change in intensity due to stimulation is between 0.5% and 5% of the average intensity, and the noise level is between 0.5% and 1%.

The main reason for performing spatial filtering of the fMRI data is to reduce the noise level whilst retaining the underlying signal. Noise is reduced since the blurring function is effectively a local averaging, so that unrelated noise values in the local neighbourhood will tend to cancel each other out. However, in order for the underlying signal to not be reduced along with the noise, it is required that the extent of the blurring is no larger than the size of the activated region. Therefore, if it is likely that the activated regions will be very small then spatial filtering should not be carried out.

The secondary reason for spatial filtering is that Gaussian Random Field theory [10], which is often used for statistical thresholding, is only valid when the data is spatially smooth. However, the amount of smoothing required for this is generally quite small – a 4mm width blurring function is generally adequate.

Alternative methods to simple spatial filtering also exist. These methods not only avoid the necessity of setting an arbitrary smoothing amount (the width of the blurring function) but, by incorporating more sophisticated models, can considerably improve the sensitivity of the statistical analysis. One alternative approach [17, 20] is to use a set of differently sized spatial filters to form a scale-space which can then be analysed later to find activations of a range of sizes. This approach can be used to set an appropriate spatial scale for a given set of experiments or, along with appropriate multiple-comparison corrections, can be combined and thresholded directly. Another alternative approach is to combine spatial information in different ways, either with Markov Random Fields or with non-linear filtering [6, 12]. These methods make it possible to spread information spatially without necessarily blurring the signal of interest. In other words, voxels can *selectively* use information from their neighbours, which should result in clearer fine detail within activation clusters. As such alternative developments mature, the use of simple spatial filtering may become obsolete.

5.2 Implementation of Spatial Filtering

The most common method of carrying out spatial filtering is to convolve each volume with a Gaussian filter. The width of this filter, usually expressed in mm, determines the extent of the blurring that takes place. A width of between 3 and 10 mm full-width-half-maximum (FWHM) for fMRI images is most common, where the term FWHM refers to the full width of the kernel (the Gaussian) at the points which are half of the value of the central maximum.

Figure 1 shows an example image and the effect that using a 10mm FWHM smoothing filter has on it. Both the images themselves and the example intensity profiles through the images clearly show the effect of the blurring.

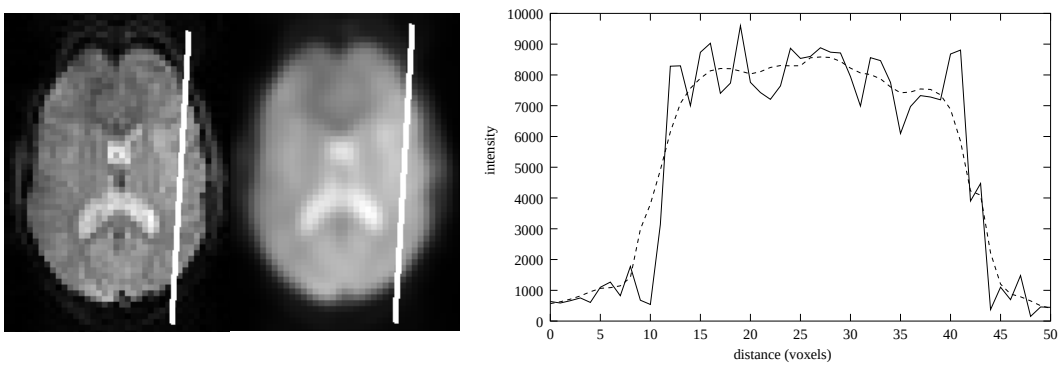

Figure 1: Effect of spatial filtering. Left: single slice from original image with source of 1D profile (plotted as solid line) shown. Middle: single slice from filtered image (FWHM=10mm) with source of 1D profile (plotted as dashed line) shown.

In three dimensions, the convolution filter can be visualised as a blurred sphere, although in practice the 3D filtering operation is usually implemented either as 3 separate 1D filters (in x, y and z) or via multiplication in Fourier space (that is, k-space).

6 Intensity Normalisation

The intensity normalisation stage of the pre-processing attempts to rescale the mean intensity of the fMRI signal in order to compensate for variations of global signal both within sessions and between sessions. In the former case the normalisation attempts to compensate for changes of intensity with time, while in the latter case the compensation is for changes from session to session. These two cases will be discussed separately here as not everyone agrees that they are both necessary in fMRI.

Within-session intensity normalisation is particularly necessary in image modalities where the global signal changes significantly, such as in PET (Positron Emission Tomography) where the decay of the radioactive tracer leads to a substantial global decrease in signal over time. However, for fMRI, the magnitude of the global change is smaller and considerably slower than the acquisition rate (TR), making it appear as a slow drift. This slow drift can be removed in several ways, such as temporal filtering (described in the next section), modelling the mean intensity as a confound in the statistical analysis or by using intensity normalisation. There is currently no consensus on whether within-session intensity normalisation is actually necessary in fMRI to compensate for such drifts, although it is still relatively common to apply it to fMRI data in practice.

In fMRI, within-session intensity normalisation is commonly carried out by finding, for each volume separately, the mean intensity across all voxels which have an intensity above a predetermined threshold[1]; then all intensity values (within the volume) are rescaled by a constant value, so that the new mean intensity becomes a preset value – for example, 10000. There is, however, a widely recognised problem with such a method of intensity normalisation. That is, if strong activation occurs, then the activation itself will increase the mean intensity, such that a volume which contains strong activation will have an artificially high mean intensity. Thus, after normalisation the "non-activated" parts of the volume will be negatively correlated with the stimulation, and will show up as "deactivation" in the final statistical image. This is clearly to be avoided – see, for example, [1]. Similarly, some experiments directly alter mean intensity level (for example, if overall oxygenation is being modulated during the session) in which case mean-based intensity normalisation is, again, problematic.

There are two possible solutions to this problem. The easy solution is simply not to carry out intensity normalisation – especially as it is arguably unnecessary for normal fMRI data. Alternatively, in cases where intensity normalisation is indeed necessary, a more robust estimation of global intensity which is not sensitive to strong local activation can be used. For example, a median intensity, instead of a mean, might be expected to be appropriate, as it is relatively insensitive to such "outliers" in data. Alternatively, it is possible to iterate between the statistical analysis and intensity normalisation [2], using the statistical results to mask out the activated areas in the intensity normalisation stage. However, such iterative schemes are very time consuming and not guaranteed to converge to the correct (globally optimal) solution. Currently, no general robust solution is in wide use.

On the other hand, between-session intensity normalisation is used to compensate for variations of the global intensity across sessions. This is almost always used in fMRI as otherwise problems can occur in second-level (multi-subject) analyses. For example, if one session had a mean intensity twice that in another session then the fitted activation parameter estimates would also be double (given the same activation). A random-effects group analysis would then see this, incorrectly, as between-session variability in response.

In the case where within-session intensity normalisation has been applied, then nothing more needs to be done, otherwise, a between-session intensity normalisation is normally performed. This operates in the same way as the within-session version but instead of a volume by volume mean intensity estimate, it makes a session by session estimate, scaling each session (4D data set) by a single value in order to have a preset mean intensity across the session. This is sometimes referred to as *grand mean* normalisation, and is less problematic than within-session normalisation since the mean is taken over many volumes and is therefore less affected by the activation signal.

[1] This could be defined, for example, as 10% of the maximum volume intensity. The reason for thresholding is that one is only interested in the mean intensity of "brain" voxels, and not background.

7 Temporal Filtering

Temporal filtering works on each voxel's time series separately, instead of working on each (spatial) volume separately, as in spatial filtering and intensity normalisation. This stage of the pre-processing chain is carried out after all the spatial processing is done since operations like high pass filtering remove the mean signal, which is often critical for stages such as motion correction or intensity normalisation. Furthermore, as most basic statistical analyses also operate directly on voxel time series this filtering stage can be implemented as a separate pre-processing stage or directly as part of the statistical model fitting.

The main point of temporal filtering is to remove unwanted components of a time series, without, of course, damaging the signal of interest. For example, if a stimulation is applied for 30 seconds, followed by 30 seconds rest, and this pattern is repeated many times, the signal of interest will be close to a square wave of period 60 seconds. Temporal filtering will normally attempt to remove components in the time series which are more slowly varying than this 60s periodic signal (high-pass filtering or drift removal) and also remove components which are more quickly varying (low-pass filtering, or noise reduction). Figure 2 shows an example time series, decomposed into the different frequency components.

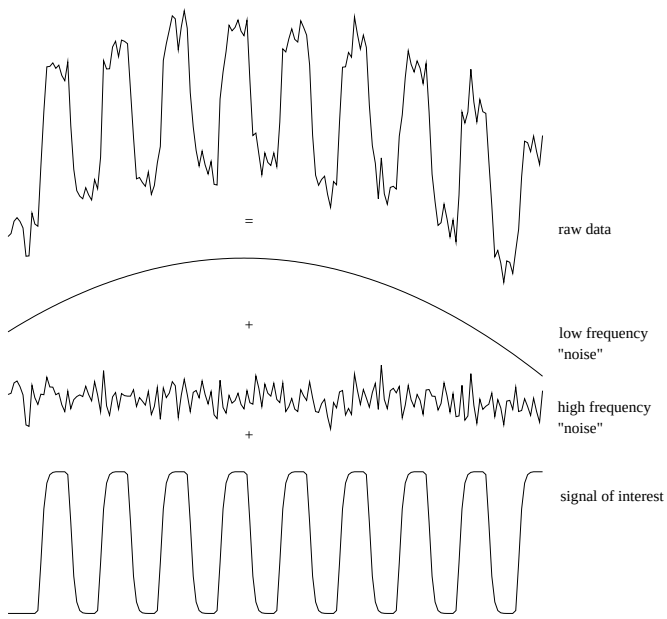

Figure 2: An example unfiltered time series from an activated voxel. The raw time series is decomposed into three components: low frequency artefact (or drift), high frequency artefact (or noise) and the expected signal of interest, caused by the stimulation.

7.1 High-pass Filtering

High-pass filtering attempts to remove all slowly varying unwanted signals in each voxel's time series. Such confounds could be physiological effects like heartbeat or breathing, or scanner-related drifts. Note that although the physiological effects may primarily occur at higher frequencies than the (stimulus-related) signal of interest, the temporal sampling (taking one volume every few seconds) interacts with the original frequency to give an apparent signal at a much lower frequency. This is known as aliasing, and is equivalent to watching propeller blades appear to turn slowly when viewed on television. (Note that some researchers – e.g. [13] – have attempted to remove physiological artefacts in a more sophisticated way than simple high-pass filtering, including feeding measurements of the cardiac and respiratory cycles into the algorithm, but this kind of approach is not yet in common use.)

If slowly varying signals are not removed from a voxel's time series then the apparent noise (or *residuals*) in the later statistical analysis will be higher than necessary. If one supposes that low frequency signals do not relate to the stimulation, it is therefore worthwhile removing them. This will result in better fitting of the stimulation-

derived model to the data, and more significant activation.

Of course, it is important that the high-pass filtering does not corrupt the stimulus-related signal. If the cutoff period of the filter is too low (i.e. cutoff frequency too high), the signal of interest will be reduced or even eliminated. This is particularly important in block-design experiments, whereas for event-related designs the danger is much less. Thus it is normal to set the cutoff period at 1.5 times the period of the stimulation. For example, if the stimulation is 10 volumes off, 10 volumes on, repeated many times, then a cutoff period of $1.5 \times (10 + 10) = 30$ volumes is safe. The selection of an appropriate cutoff period when complex block-design experiments are to be analysed requires careful thought. For example, take an ABACABAC... design, where A is the rest (or control) condition, and B and C are different types of stimulation; all blocks (A, B and C) consist of 10 volumes. If the response to B is to be compared with the response to C, the cutoff period should be at least $1.5 \times 4 \times 10$; anything smaller would result in a loss of sensitivity to this contrast.

High-pass filtering is often achieved by finite impulse response (i.e., convolution-based) linear filters, with designs such as the Butterworth filter. The design normally permits both the cutoff period and the sharpness of the cutoff to be specified. Although this kind of filtering is simple and in common use, some have an unfortunate side-effect in that they can induce negative auto-correlations into the signal (i.e., an oscillatory component), which can confound later measurements of intrinsic temporal smoothness (see the following section for more on this).

Alternatively, as mentioned previously, it is possible to model low frequencies within the later statistics instead of removing them at this stage. These trends are treated as confounds within the modelling and so effectively ignored when calculating contrasts of interest. This is commonly achieved by placing a series of low frequency cosine waves at a range of frequencies in the model – a linear combination of these then fits the actual drift in the data. This solution does not suffer from the problem of inducing negative auto-correlations in the data and correctly accounts for the change in degrees of freedom induced by the filtering.

7.2 Low-pass Filtering

Low-pass filtering attempts to reduce high frequency noise in each voxel's time series, without affecting the signal of interest. As with the high-pass filtering, it is important to choose a filter that removes noise without corrupting the underlying stimulus-related signal.

It is common to carry out low-pass filtering via a simple linear convolution with a Gaussian kernel. This can be thought of as a "blurring" function, and is identical in theory to the Gaussian spatial filtering described earlier. Because it is normal to use a fairly narrow Gaussian, i.e., carry out relatively little blurring, this does not do much more than replace each time point with the original value plus a small fraction of its immediate (temporal) neighbours'.

One danger with low-pass filtering occurs when using it with event-related experiments, since these often contain signals of interest which are rapidly changing. For example, very brief stimulation may give rise to narrow peaks in the resulting time series. Low-pass filtering may suppress these signals, thus reducing the power of later statistical analysis.

Another danger arises because of the increases in smoothness of the time series due to low-pass filtering. ("Smoothness" – sometimes referred to as positive autocorrelation – means that the intensity at any given time point is likely to be closer to its neighbours' values than values elsewhere in the time series.) Smooth data contains less effective, or truly independent, time points than the actual number of points. Later analysis needs to correct for this smoothness, otherwise reported significances will be overestimated, resulting in false positives. Related to this, certain statistical approaches *rely* on low pass filtering, in order for later calculations of smoothness to be well-conditioned [7, 8].

It has relatively recently been shown [19] that if one is able to robustly and accurately characterise the nature of the intrinsic temporal smoothness of the fMRI data, then the most efficient approach is to carry out no smoothing at all, and instead "prewhiten" the data within the statistical analysis. Thus, with this method, no low-pass filtering is applied.

Other researchers have looked at more sophisticated low-pass filtering, including wavelet filtering, Markov random fields, local enforced monotonicity and local neighbourhood smoothing [15]. It has not yet been conclusively shown that such approaches are of general benefit; however, whilst many researchers still regularly use linear low-pass filtering of fMRI data, the interest in more sophisticated non-linear methods, both in pre-processing and as part of the statistical analysis, is increasing.

References

[1] G.K. Aguirre, E. Zarahn, and M. D'Esposito. The inferential impact of global signal covariates in functional neuroimaging analyses. *NeuroImage*, 8:302–306, 1998.

[2] J.L.R. Andersson. How to estimate global activity independent of changes in local activity. *NeuroImage*, 6:237–244, 1997.

[3] J.L.R. Andersson, C Hutton, J Ashburner, R Turner, and K Friston. Modelling geometric distortions in EPI time series. *NeuroImage*, 13(5):903–919, 2001.

[4] P.R. Bannister, S.M. Smith, M. Brady, and M. Jenkinson. Spatio-temporal realignment of FMRI data. In *Eighth Int. Conf. on Functional Mapping of the Human Brain*, 2002.

[5] E.T. Bullmore, E.T. Suckling, S. Overmeyer, S. Rabe-Hesketh, E. Taylor, and M.J. Brammer. Functional magnetic resonance image analysis of a large-scale neurocognitive network. *NeuroImage*, 4:16–33, 1996.

[6] X. Descombes, F. Kruggel, and D. Y. von Cramon. FMRI signal restoration using a spatio-temporal markov random field preserving transitions. *NeuroImage*, 8:340–349, 1998.

[7] K.J. Friston, A.P. Holmes, J.-B. Poline, P.J. Grasby, S.C.R Williams, R.S.J. Frackowiak, and R. Turner. Analysis of fMRI time series revisited. *NeuroImage*, 2:45–53, 1995.

[8] K.J. Friston, P. Jezzard, and R. Turner. Analysis of Functional MRI time-series. *Human Brain Mapping*, 1:153–171, 1994.

[9] K.J. Friston, S.R. Williams, R. Howard, R.S.J. Frackowiak, and R. Turner. Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine*, 35:346–355, 1996.

[10] K.J. Friston, K.J. Worsley, R.S.J. Frackowiak, J.C. Mazziotta, and A.C. Evans. Assessing the significance of focal activations using their spatial extent. *Human Brain Mapping*, 1:214–220, 1994.

[11] J.V. Hajnal, R. Myers, A. Oatridge, J.E. Schwieso, I.R. Young, and G.M. Bydder. Artefacts due to stimulus correlated motion in functional imaging of the brain. *Magnetic Resonance in Medicine*, 31:283–291, 1994.

[12] N.V. Hartvig and J. Jensen. Spatial mixture modelling of fMRI data. *Human Brain Mapping*, 11(4):233–248, 2000.

[13] X. Hu, T.H. Le, T. Parrish, and P. Erhard. Retrospective estimation and correction of physiological fluctuation in functional MRI. *Magnetic Resonance in Medicine*, 34(2):201–212, 1995.

[14] P. Jezzard. Physiological noise: Strategies for correction. In C.T.W Moonen and P.A. Bandettini, editors, *Functional MRI*. Springer, 1999.

[15] F. Kruggel, D.Y. von Cramon, and X. Descombes. Comparison of filtering methods of fMRI datasets. *NeuroImage*, 10:530–543, 1999.

[16] J Orchard, C Greif, GH Golub, B Bjornson, and MS Atkins. Simultaneous registration and activation detection for fMRI. *IEEE Trans. on Medical Imaging*, 22(11):1427–1435, 2003.

[17] J.-B. Poline and B.J. Mazoyer. Enhanced detection in brain activation maps using a multifiltering approach. *Journal of Cerebral Blood Flow and Metabolism*, 14:639–642, 1994.

[18] F Schmitt, MK Stehling, and R Turner, editors. *Echo-Planar Imaging: Theory, Technique and Application*. Springer-Verlag, Berlin, 1998.

[19] M.W. Woolrich, B.D. Ripley, J.M. Brady, and S.M. Smith. Temporal autocorrelation in univariate linear modelling of FMRI data. *NeuroImage*, 14(6):1370–1386, 2001.

[20] K.J. Worsley, S. Marrett, P. Neelin, and A.C. Evans. Searching scale space for activation in PET images. *Human Brain Mapping*, 4:74–90, 1996.