

Statistical Models for Image Sequences

Neil Crellin, Trevor Hastie and Iain Johnstone

Department of Statistics
Stanford University
Stanford, CA 94305-4065, USA

December, 1998

Abstract

Identification of brain activity using functional magnetic resonance imaging (fMRI) depends on blood flow replenishing activated neuronal sites. In this paper we describe how previous studies have sought to model the hemodynamic response function with restricted parametric models, and examine some of the inadequacies imposed by the implicit restrictions. We investigate more flexible estimation of hemodynamic response using cubic spline basis expansions, both of the time course itself and of the hemodynamic response function. In the latter, a deconvolution estimate of hemodynamic response is estimated from the observed fMRI time series at each pixel location and the designed temporal input stimulus. The estimated hemodynamic responses include both monophasic and biphasic forms, comparable with the more limited model proposed in Friston et al. [1995b]. Bootstrapping allows us to show that for our data, a Poisson-based convolution model performs no better than fitting a sinusoid to the data, but that for regions of activation, both the spline-based convolution model and periodic spline models do. A study of primary visual cortex activation is used to illustrate these findings.

1 Introduction

Functional magnetic resonance imaging (fMRI) has made it possible to conduct sophisticated human brain mapping neuroscience experiments. Such experiments commonly consist of human subjects being exposed to a designed temporal sequence of stimulus conditions while repeated MRI scans

of the brain region of interest are taken. A common experimental design alternates equal length periods of stimulus and rest for a number of cycles. The magnetic characteristics of hemoglobin in the blood are detectably changed by oxygenation. Recently activated sites of neural activity are replenished with oxygenated blood, allowing identification of brain activity using MRI scans designed to detect such changes in magnetic susceptibility [Ogawa et al., 1992, Kwong et al., 1992]. While neural activation occurs on a millisecond time-scale, the detectable blood flow replenishing the activation sites depends on local vasculature and can occur as long as several seconds after activation.

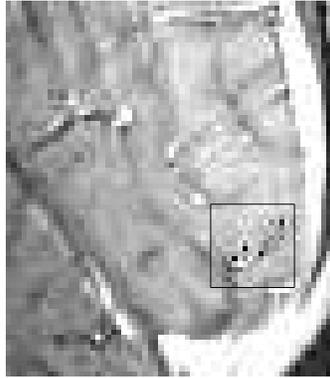


Figure 1: *An oblique anatomic MRI scan localized around the calcarine sulcus. The 16 by 16 subregion indicated contains the area of primary visual cortex being stimulated. The four darker pixels in this region are selected to demonstrate some of the methods in this paper.*

The concept of the *hemodynamic response function* introduced by Friston et al. [1994b] and their convolution model has been enthusiastically adopted in the fMRI literature. A version of this model has the form:

$$y(t) = (h * x)(t) + \epsilon(t) \quad (1)$$

where

- $y(t)$ is the observed signal at a particular time (t).
- $h * x = \int_{v>0} h(v)x(t-v)dv$ is a convolution of the neuronal response $x(t)$ with the hemodynamic response function $h(v)$.

- $\epsilon(t)$ is the residual process.

The idea is that rather than observing the neuronal activity $x(t)$ directly, MRI allows us to see a blurred and delayed version via the measured blood activity (which responds to the neuronal activity); $h(v)$ is the blurring function.

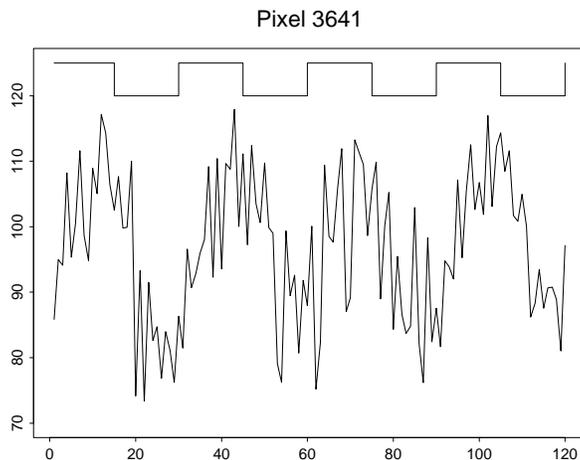


Figure 2: *The time course from a single pixel is displayed. Superimposed is the treatment profile, with the higher level indicating stimulus and the lower rest.*

Friston et al. [1994b] propose use of a Poisson form for the hemodynamic response function (to be convolved, for example, with the temporal treatment profile), with a single parameter λ globally describing delay, dispersion, and hemodynamic response function shape. Lange and Zeger [1997] and Boynton et al. [1996] use Gamma forms for the hemodynamic response. Lange and Zeger estimate parameters at each pixel location rather than globally, and use a frequency domain fitting procedure.

Friston et al. [1995b,a] observe that evoked hemodynamic responses can be biphasic or have differential ‘early’ or ‘late’ profiles. They fit a model consisting of a linear combination of two monophasic parametric curves to explore this behavior, for example. They set the parameters of their model’s monophasic curves to pre-chosen integers and do not estimate them using the data.

An acknowledged deficiency of the Poisson model for the hemodynamic response function is that it combines delay and dispersion into a single pa-

parameter. Given the possible inadequacy of the Poisson and Gamma models in describing the true nature of vascular delay and acknowledging the possibility of a biphasic response, we propose a more data-driven approach to the estimation of the hemodynamic response. In particular, by using a more flexible modeling family, such as cubic B-splines, we are able to discover more information regarding the local workings of the brain during activation such as sensory stimulation or cognitive function.

2 Data

Figure 1 shows an anatomical MRI scan of the human brain. This is an oblique slice taken as part of an investigation into activation of the primary visual cortex, which is located along the calcarine sulcus. The area of activation for this particular experiment is expected to lie within the boxed 16 by 16 pixel subregion along the sulcus. The stimulus-rest regime for this series of functional scans consists of stimulus for 15 scans then rest for 15 scans repeated for 4 cycles, a total of 120 functional scans. An image was taken every 1.5 seconds in a continuous spiral scan.

Figure 2 shows an example of a time series produced at a single pixel, with the stimulus regime superimposed at the top of the figure. Note the delay between the change in treatment and the change in signal. Observe also the periodic nature of the response in an almost sinusoidal pattern. One common form of analysis of such data consists of thresholding a map of the correlation coefficient between the data at each pixel and the best fitting sinusoid at the stimulation frequency (with respect to amplitude and phase).

Highlighted in the activation square in figure 1 are four selected pixels, whose mean-corrected time courses are shown in Figure 3. Superimposed over the data are the fits produced by the Lange-Zeger procedure using a Poisson model for the hemodynamics. Figure 4 shows the fits produced at these four pixels using periodic splines, and figure 5 shows a single cycle of the periodic splines compared to the Lange-Zeger estimates at the corresponding pixels. Precise details of the periodic spline model are described in the next section.

A large number of the periodic spline fits of the data exhibited the characteristic double-humped pattern seen in these figures. Increasing the number of knots did not dramatically alter the overall shape of the fits, suggesting this pattern is not an artifact of the fitting procedure.

The structure apparently uncovered by these exploratory fits suggested

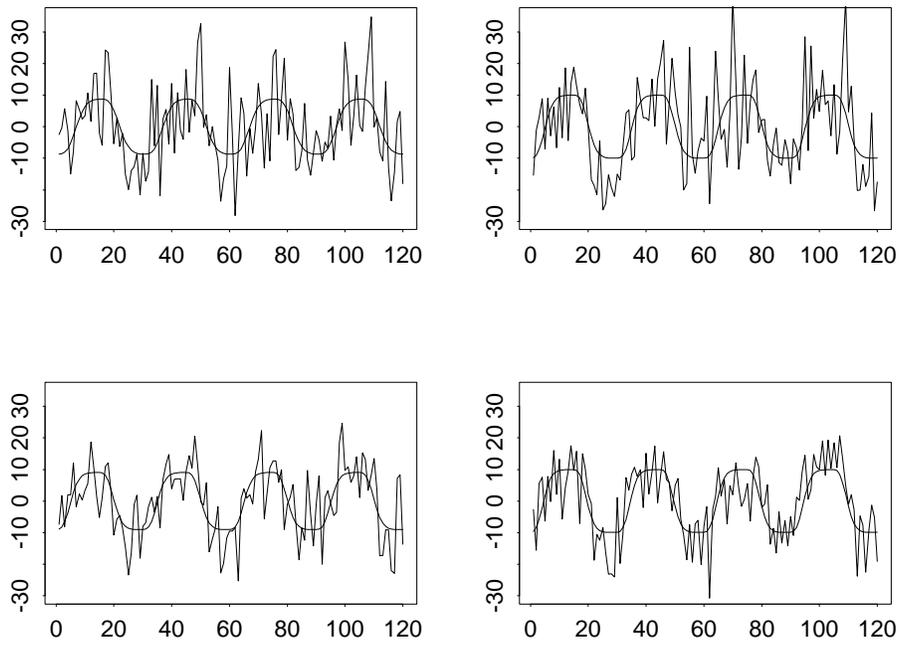


Figure 3: Mean-corrected time-courses at the four pixel locations indicated in figure 1 (in order from left to right). Superimposed over the time-courses are the parametric model fits obtained using the Lange-Zeger procedure, and a poisson model for the hemodynamic response function.

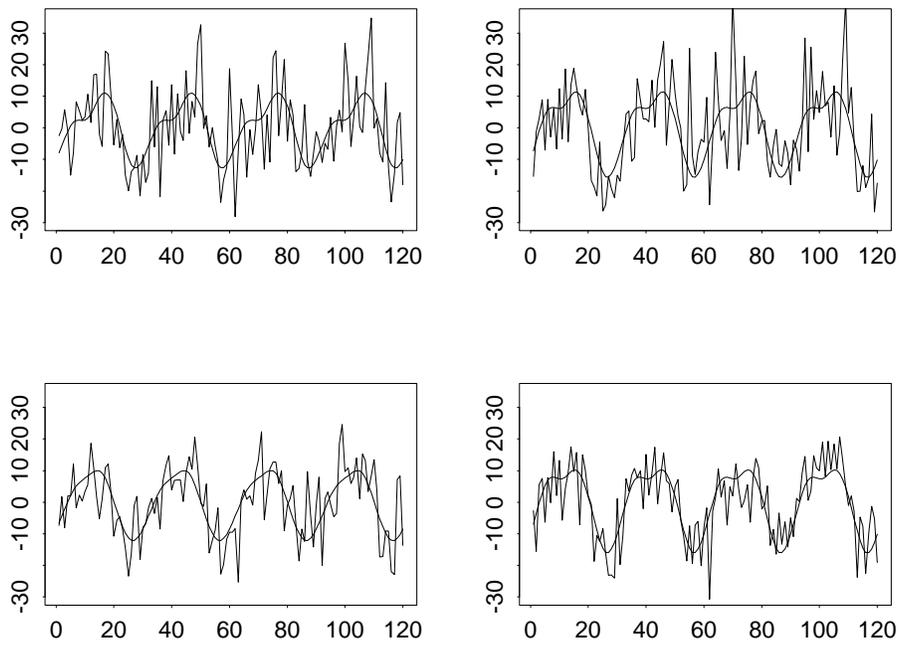


Figure 4: *Mean-corrected time-courses at the four pixel locations indicated in figure 1 (in order from left to right). Superimposed over the time-courses are the fits obtained using periodic cubic splines.*

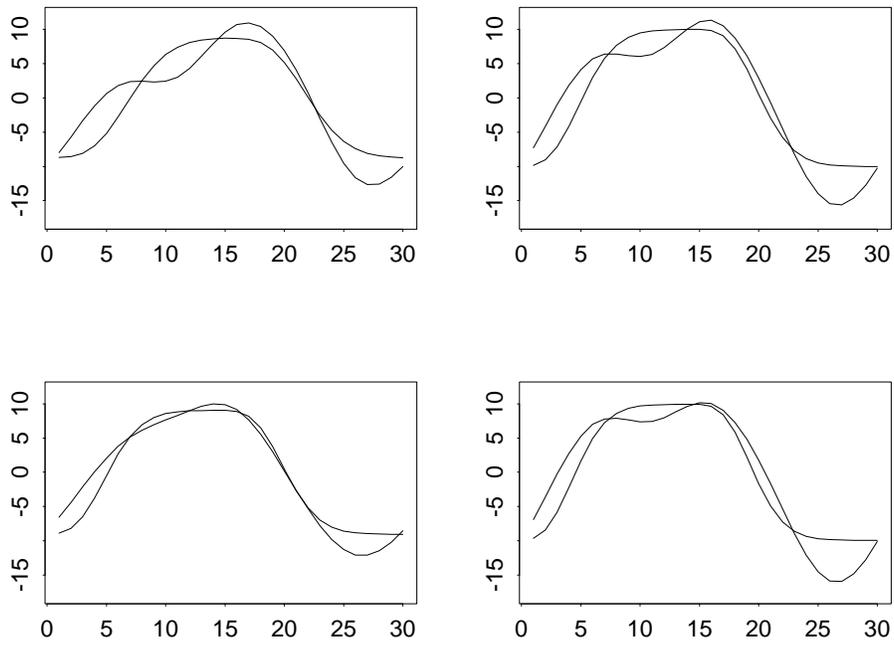


Figure 5: *The monophasic kernel fits of the Lange-Zeger procedure compared to the periodic spline fits over a single cycle at each of the 4 pixels displayed above.*

that a more flexible model might provide greater insights into the nature of the brain’s response to activation. The double-humped pattern seemed to be suggestive of a biphasic hemodynamic response. This led us to also consider cubic spline basis expansion models for the hemodynamic response function.

3 Models

The data we are considering consists of sequences or time series of image data. Let y_i denote the i^{th} image in the sequence, for $i = 1, \dots, n$, where each image is viewed as a vector of $M = m_1 \times m_2$ pixels, where m_1 and m_2 are the image dimensions. If we consider the image data as an $n \times M$ matrix Y , each row represents the pixels of an image in the sequence, and each column represents the time-series of responses at individual pixel locations. The experimental subject is scanned while being exposed to alternating periods of rest and (in our example above, visual) stimuli. Encoding whether an image is scanned during a period of rest or stimulation, we can create a covariate $x(t)$, for which we have the discretized vector with $x_i = x(t_i)$ equal to one if the i^{th} scan occurs during the stimulus condition, or zero during rest.

We now propose two different flexible approaches for modeling the series; the first is a spline model for the observed series itself, the second a spline model for the hemodynamic response function.

3.1 Periodic Spline Model for Pixel Time Series

Since the simplest commonality of the experimental regime is its periodic nature, we first explored the time-series using regression with periodic splines. Here we are modeling the observed data directly, with a model of the form

$$Y(t) = S(t) + \epsilon(t)$$

where $S(t) = \sum_{l=1}^L \theta_l b_l(t)$. The spline fits are based on equally spaced knots over a single cycle, constrained to periodically wrap at the cycle boundaries with continuous zero’th through second derivatives at the knots. The functions $b_l(t)$ for $l = 1, \dots, L$ are the cubic B-spline basis functions for L equally spaced knots over a single stimulus-rest cycle, and the values of θ_l must be such that the periodicity constraints are satisfied.

The times at which fMRI scans are taken, t_i , are converted to new times $t_i^* = t_i \bmod T$. We can then construct a basis for cubic-splines periodic on

$[0, T]$ as follows. We have available a function to generate a basis of cubic B-splines $b_l(t)$ with L equally spaced knots in $[0, T]$, and their derivatives de Boor [1978]. Using their zero'th through third derivatives at 0 and T , we can obtain a linear constraint matrix C such that $C\theta = 0$ enforces the periodic boundary conditions on the parameters θ . We now construct a basis matrix B using the periodic time points t_i^* . Using C , we can reduce B to B^* by standard linear algebra [Golub and Van Loan, 1983], where B^* has the constraints built in. We can then regress the pixel-wise series Y_{ij} at the j^{th} pixel on the columns of B^* for $j = 1, \dots, M$. Standard software such as Matlab or Splus allow these least-squares calculations to be performed simultaneously for all pixels. Examples are given in figure 4.

3.2 Hemodynamic Spline Model for response function

Alternatively, suppose the time-series $Y(t)$ arises as the convolution of an unknown hemodynamic response function $h(v)$ with the stimulus function $x(t)$

$$Y(t) = \mu + (h * x)(t) + \epsilon(t) \quad (2)$$

Let $h(k)$, $k = 1, \dots, K$ denote the discretized version of $h(v)$ evaluated at the uniform time points. Since we allow for distinct hemodynamic responses at each pixel location, we may refer to the hemodynamic response at the j^{th} pixel using $h_j(v)$ or its discretization $h_j(k)$. We can write the convolution model as

$$Y_{ij} = \mu_j + \sum_{k=1}^K x_{i-k+1} h_j(k) + \epsilon_i$$

for each pixel j and time index i . Negative subscripts are to be understood as zero value quantities. (Cyclically lagging values modulo n may be an appropriate alternative in some circumstances.) This can be rewritten in vector form as

$$Y_j = \mu_j + Xh_j + \epsilon_j \quad (3)$$

where the $n \times K$ matrix X is constructed with first column given by the vector x described above, and lagged versions of x in subsequent columns. (If h_j is periodic, cyclic lags may be quite appropriate, especially in the typical case where initial scans are not considered while the magnetic resonance relaxation effects of tissue stabilize.) For physiological interpretation, we consider μ_j as a baseline (control) response level, and values of $h_j(k)$ to be non-negative, corresponding to how much blood is delayed by that amount. We can now readily estimate h_j using least-squares, subject to a

non-negativity constraint. That is, at the j^{th} pixel, we minimize

$$Q_j = \|Y_j - \mu_j - Xh_j\|^2 \quad (4)$$

with respect to μ and $h_j = h_j(k)$, $k = 1, \dots, K$ subject to the constraints $h_j(k) \geq 0$ for $k = 1, \dots, K$.

In this description so far, we have offered a similar formulation to the convolution models proposed by Friston et al and Lange and Zeger and others. They model h using a Poisson or Gamma model, which could be fit in this framework by iterative non-linear least squares techniques in the time domain. Friston et al. [1995b] use a general linear model estimation procedure as described in Friston et al. [1994a, 1995c] and Worsley and Friston [1995]. As described above, their model assumes an uncorrelated error term as the sole source of error which is added to the response function before convolution by the hemodynamic response. Lange and Zeger transform into Fourier space and perform what is effectively iteratively re-weighted non-linear least squares in the frequency domain. Friston et al. [1994b] assume the same Poisson model across all pixels; Lange and Zeger allow for locally varying parameter values between pixels, and use Gamma models to allow for more general local variability.

Rather than presume a specific distributional form for h , we propose instead to estimate it semi-parametrically at each pixel location based on the time course data. Specifically, we propose a basis expansion model of the form $h(v) = \sum_{l=1}^L \theta_l b_l(v)$ and the values of θ are the parameters to be estimated. The studies referred to above use scaled probability distributions such as the Poisson and Gamma to estimate hemodynamic response. Given the physiological interpretation of h , we expect $h(v)$ to be “smooth” and suggest the use of cubic B-splines as basis functions to ensure smoothness without requiring a specific monophasic form. Note that the cubic B-spline basis functions $b_l(v)$ for $l = 1, \dots, L$ are themselves probability density functions. Modeling $h(v)$ as $\sum_{l=1}^L \theta_l b_l(v)$, we are estimating the hemodynamic response as a non-negative linear combination of densities. The coefficients may be negative so long as the estimate of $h(v)$ is non-negative. The non-negativity constraint on h makes this a constrained but linear least-squares problem (note that the constraints are also linear). This is easily implemented using minimization software such as CFSQP Lawrence et al. [1994]. We discretize the basis functions $b_l(v)$ for $l = 1, \dots, L$ at the uniform points at which $h(v)$ was discretized into the columns of a matrix B (which is thus $K \times L$). Criterion (4) becomes

$$Q_j = \|Y_j - \mu_j - XB\theta_j\|^2, \quad (5)$$

which we minimize with respect to μ_j and $\theta_j(l), l = 1, \dots, L (L \leq K)$ subject to the constraints $\sum_{l=1}^L \theta_j(l) b_l(k) \geq 0$ for $k = 1, \dots, K$.

4 Example

Examples of the periodic spline fits to several pixels were provided previously in figures 4 and 5. They appear to follow the observed data quite well and detect structural patterns not observed in the comparable Lange-Zeger fits. We now investigate whether the hemodynamic spline models show similar flexibility.

We use a cubic B-spline basis with eight equally spaced knots on the interval $[0, 22.5]$ seconds for the hemodynamic spline. Since a scan is taken every 1.5 seconds this corresponds to the time taken for 15 scans (see figure 2). Care needs to be taken in constructing the matrix X in (3). Since the test and stimulus phases last 15 scans each, the columns become linearly dependent if X has more than 15 columns. It follows that estimability of $h(v)$ beyond 22.5 seconds is not possible. Lange and Zeger cite Bandettini et al. [1993] and Friston et al. [1994b] as estimating hemodynamic delays in humans to be roughly between 4 to 10 seconds so this should not be restrictive. We have chosen $K = 15$ and $L = 8$ for this example. Figure 6 shows examples of the estimates obtained at two example pixels. The two pixels selected are the upper two of the four displayed in figures 3, 4, and 5. The estimated hemodynamic response function is shown on the left, and the fitted reconvolutions on the right. For each case the corresponding curves based on the best fitting Poisson model is displayed for comparison. There is little, if any, obvious improvement or difference in the reconvolution estimates. Also, neither convolution models' estimates appear to capture the double-humped pattern seen in the periodic spline fits.

The general shape of the estimates of hemodynamic response for these pixels are fairly close in gross features. In many cases the estimates obtained for the hemodynamic response closely follow the matching Poisson fit. Figure 7 below shows the corresponding estimates at two pixels selected to have more obvious biphasic dissimilarity. Observe that the estimate of the hemodynamic response in each case is notably different, apparently uncovering structure not observed in the Lange-Zeger and Friston models. The reconvolutions, however, are still quite similar. The periodic spline fits (not shown) in these cases are similar to both convolution models and are more sinusoidal, not exhibiting the double-humped pattern noted previously.

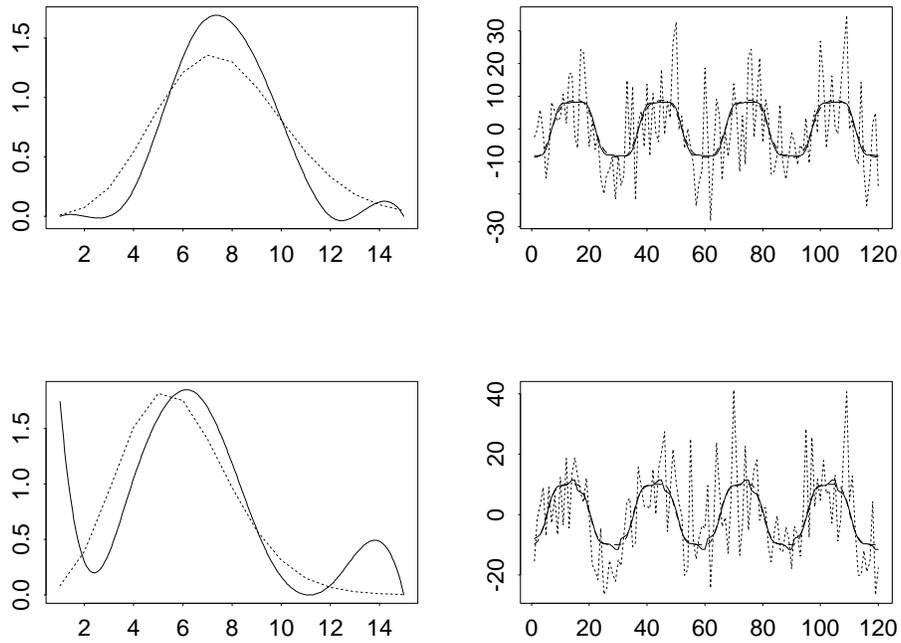


Figure 6: *On the left are estimates of the hemodynamic response function for the left-most two selected pixels in figure 1 (which are also the top two panels in figures 3,4,5). The solid curve is the periodic spline estimate, the dashed curve the Lange-Zeger poisson estimate. On the right appear the convolutions of both, superimposed over the original time series.*

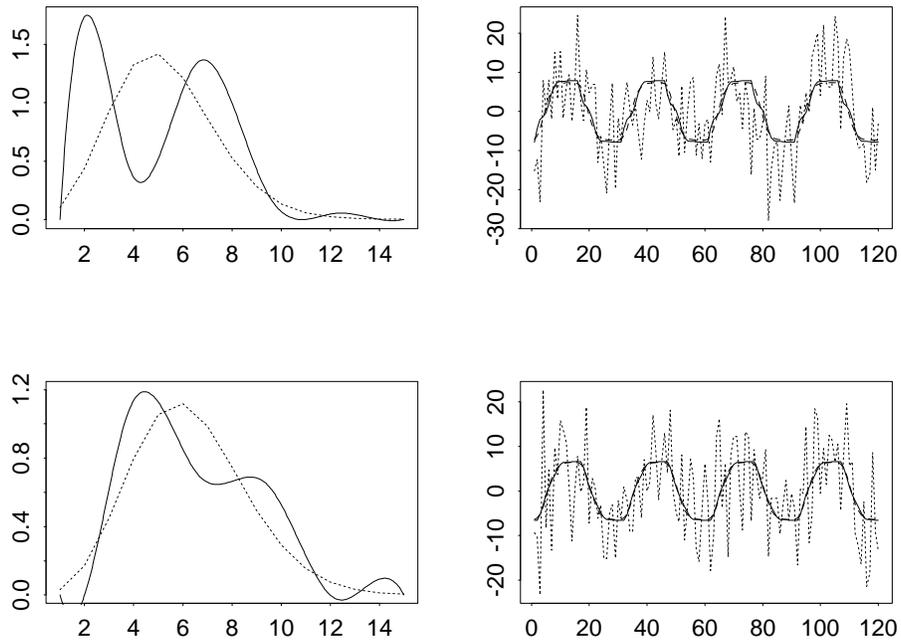


Figure 7: *On the left are estimates of the hemodynamic response function for two pixels selected to differ in their hemodynamic estimates. The solid curve is the periodic spline estimate, the dashed curve the Lange-Zeger poisson estimate. On the right appear the convolutions of both, superimposed over the original time series.*

5 Inadequacies of convolution models

To more clearly display the differences between the two convolution models and the periodic spline fits, figure 8 below shows a single cycle of each model. The fits displayed are from the first example pixel in figure 5, and the first example pixel in figure 7. The convolution models are convolving a square wave with the hemodynamic response function. As the stimulus regime switches from rest to stimulus, the convolution accumulates a sum of the values of the hemodynamic response function. As it switches back from stimulus to rest, the accumulated terms are dropped in the order they were added. This accounts for the mirror anti-symmetry at the half cycle point of the convolution model reconstructions. This is more clearly seen in figure 9, which shows one cycle of the hemodynamic spline model from the first set of panels of figure 7.

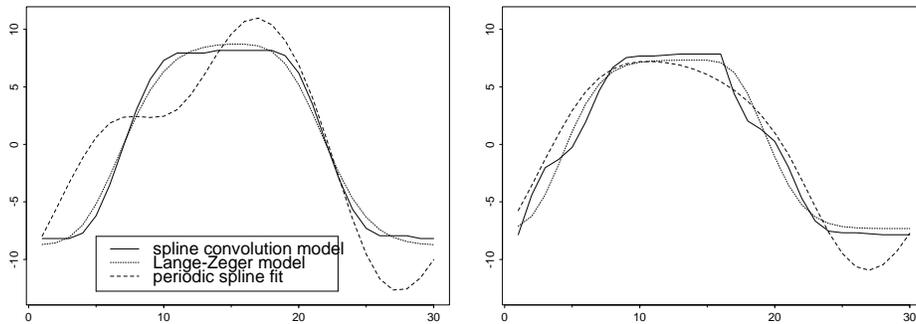


Figure 8: A single cycle of each of the convolution models together with the periodic spline fit. At left are the fits from the first example pixel above. At right are the fits from the strongly biphasic pixel seen in Figure (7).

The window smoothing nature of the convolution models for this stimulus regime also forces the estimates to be monotonically non-decreasing up to the half cycle and monotonically non-increasing in the second half cycle. By contrast, the periodic spline model is only constrained to be smooth and periodic and admits different rise and decay behaviors.

Next we attempt to model the periodic spline fitted values from the first panel of figure 8 with a convolution model. Fitting a completely general hemodynamic response function produces the deconvolution and reconvolution estimates shown in figure 10. Clearly the convolution model is unable to capture the types of patterns suggested by the spline fits to the data.

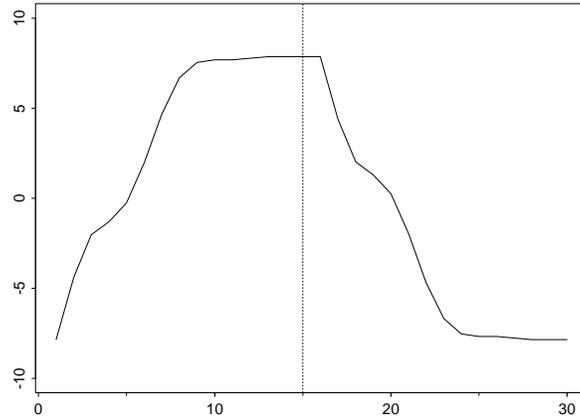


Figure 9: *A single cycle of the strikingly biphasic hemodynamic spline estimates reconvolution. Observe both the half-cycle monotonicity and mirror anti-symmetry.*

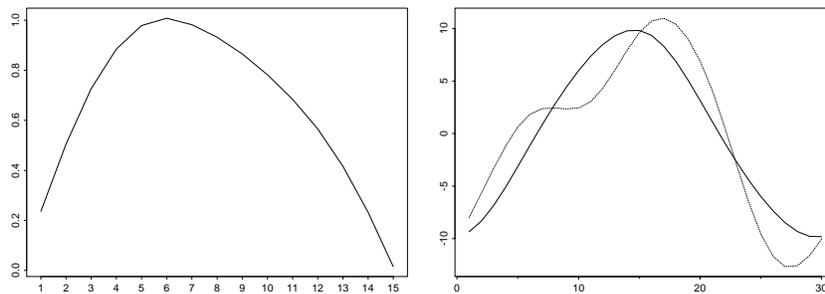


Figure 10: *On the left is the best deconvolution kernel estimate when attempting to fit the periodic spline fitted values using a convolution model. On the right are the curve being fitted and the convolution reconstruction. Note the convolution model is unable to adequately describe the shape of the spline fit.*

6 Signal or Noise?

Our investigations so far suggest that the convolution models may be unfairly restrictive and fail to represent structure present in the data. It should be stressed that thus far these are empirical results, and we need to better understand the nature of variability of the the data. The observed time series can be quite noisy, and we need to be sure we are not just fitting models to the noise. The example power spectrum in figure 11 shows the periodogram of the time series displayed in figure 2. Most of the power lies at the fundamental frequency. Implicit in the convolution model is the assumption that signal power at the off-harmonic frequencies comes from noise. The spectral power at the harmonic frequencies is of comparable magnitude to that seen at the off-harmonic frequencies. This poses the following questions: if all of the spectral power of the time series lies at the fundamental frequency, do any of these models perform significantly better than fitting a single sinusoid at that frequency, and can we statistically test if the pattern observed in the periodic spline fits is real?

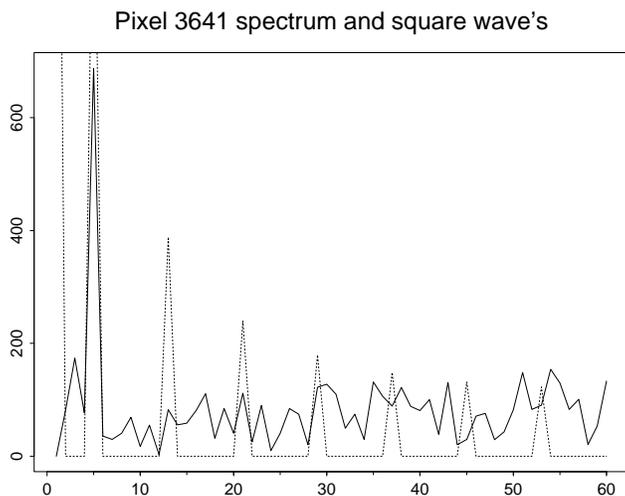


Figure 11: *The power spectrum of the first example time series above. Superimposed is a scaled periodogram of the stimulus regime square wave, whose power lies at odd harmonics of the fundamental frequency.*

6.1 Periodic spline model versus sinusoid

We wish to test whether the bumps we have found in our periodic spline models result from overfitting the data. We do this via a bootstrap hypothesis test, which attempts to preserve the autocorrelation in the data. We use the residuals from the fitted sinusoid at each pixel to form our resampling distribution.

As a test statistic, we use the sum-of-squared deviations (SS) between the periodic-spline fit and the sinusoid fit at each pixel. It can be shown that the family of sinusoids is a linear sub-family of the periodic splines, and hence the fits are nested. The bootstrap is used to estimate the null distribution of this statistic, assuming the sinusoid model is correct.

Specifically at each pixel we fit a sinusoid model, compute the residuals, and then take each cycle of the residuals as our sampling units. In the example above, each pixel time-course thus yields four cycles of residuals. We re-sample in temporal blocks of cycles in an attempt to retain the temporal autocorrelation structure of the underlying data. Under the null model that the data are ideally modeled by a sinusoid at the fundamental frequency, these can reasonably be assumed to arise from the same distribution.

At each pixel, $B - 1$ bootstrap resamples are formed by taking the original fitted sinusoid for that pixel, and adding a separate resampled residual cycle to each of its four cycles to produce a bootstrap simulated time-course. For each of these $B - 1$ simulated time series, the sinusoid and periodic-spline models are fit, and the SS statistic is computed. To test at the α level, we check whether the observed SS value lies within the largest αB values of SS.

Figure 12 shows the results for the 16×16 array of pixels in our example, and $B = 500$. Both the sinusoid and periodic-spline fit are shown in each plot, along with the realized “P-value” from the bootstrap simulations, rounded to the nearest percent. We notice that not only are several of the bumpy periodic-splines in the active area significantly different from a sinusoid, but also that there is a spatial pattern to this significance which adds further support.

6.2 Poisson convolution versus sinusoid

Comparing the poisson convolution model to the sinusoid requires some care. These are not nested fits, and so each can dominate under different circumstances. For this comparison we used a different bootstrap procedure. Let $\Delta\text{RSS} = \text{RSS}_S - \text{RSS}_P$ denote, for a given time course, the change in residual-sum-of-squares between the sinusoid fit and the poisson fit. We use

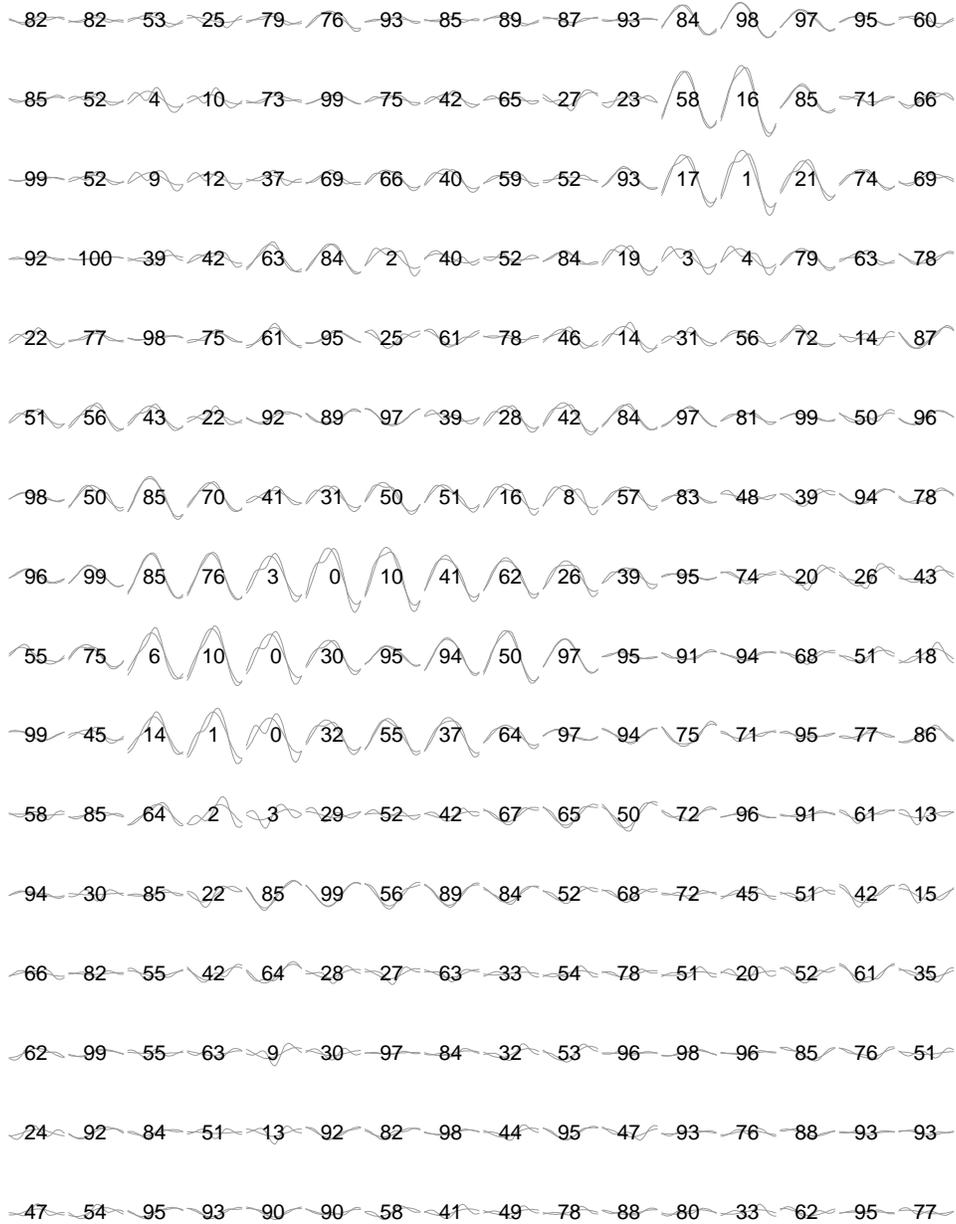


Figure 12: *Observed P-value (rounded to nearest percent) obtained in bootstrap simulation comparing periodic spline fits to sinusoid. Each sub-panel corresponds to a pixel, and includes the sinusoid and periodic spline fit for one cycle.*

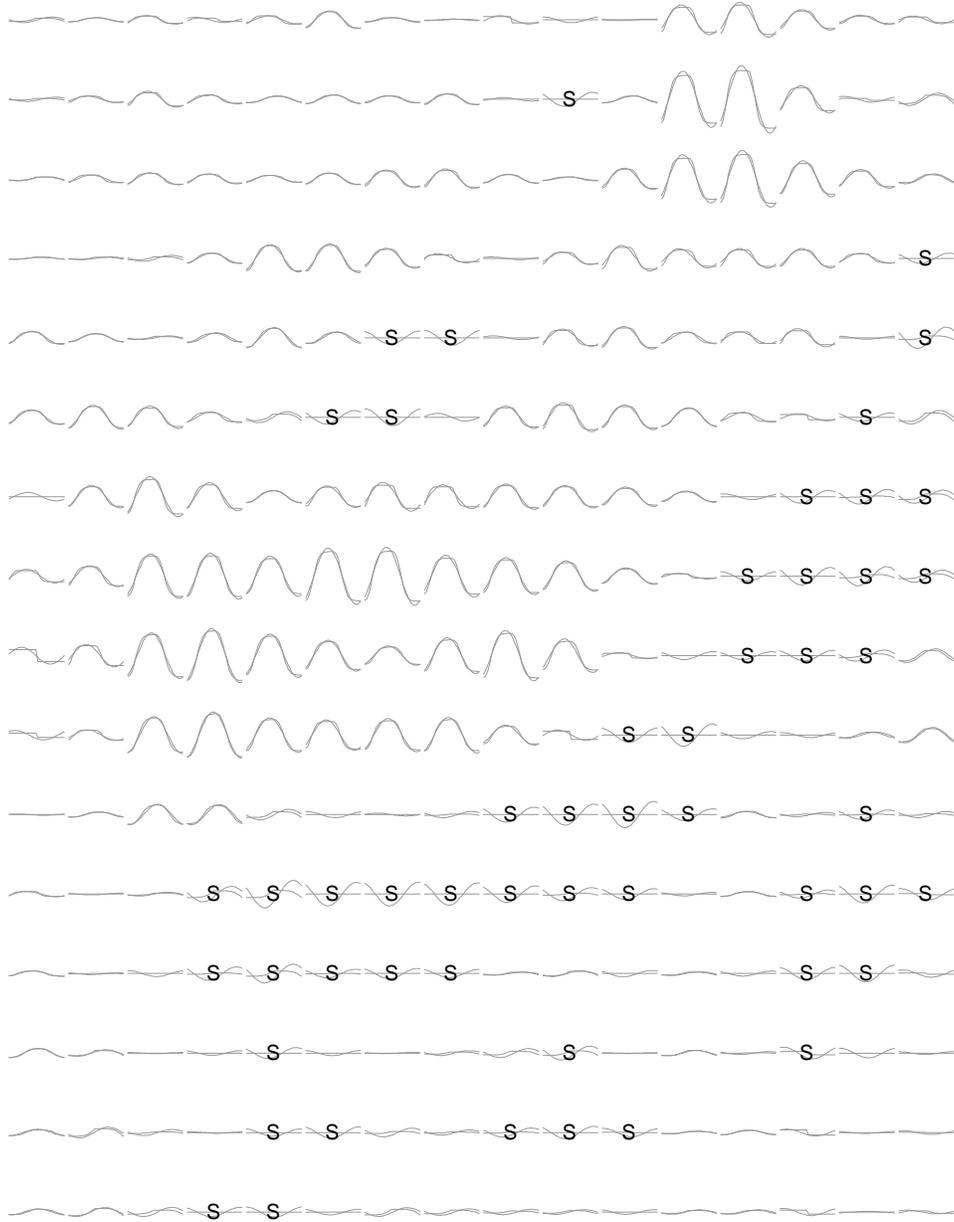


Figure 13: *Bootstrap simulation comparing the poisson convolution model to sinusoid fits. In the areas marked “S”, the sinusoid model dominates the poisson model. Elsewhere they are not-significantly different.*

the bootstrap to obtain a 90% confidence interval for ΔRSS at each pixel.

We obtain a bootstrap sampling distribution by sampling blocks of residuals from the periodic spline fits, which is a richer fit than either of the two under consideration. These are pooled across all pixels, sampled with replacement, and pasted back onto the spline fits as before. For each of 300 such realizations, ΔRSS is computed at each pixel.

Figure 13 shows at which of the pixels the sinusoid or poisson dominates, and those for which the comparisons were inconclusive. These choices were based on the pixelwise 90% confidence intervals. For example, if the lower 5% value of ΔRSS were positive, that would favor the poisson (this never occurred), while if the upper 5% value were negative, the sinusoid would be favored.

We see that either the sinusoids dominate, or there is no significant difference. The sinusoids dominate in low-amplitude areas which are of lesser interest. In many of these cases the convolution models are effectively flat. The monotonicity constraint implicit in the convolution models means that sinusoids with a wide range of phase delays may fit the data well but cannot be modeled by the convolution model. This appears to explain many of these cases.

6.3 Hemodynamic spline versus sinusoid

We carried out an identical comparison between the sinusoid model and the hemodynamic spline convolution model. The results are shown in figure 14. Here it appears that the hemodynamic spline dominates in the active areas, while the sinusoid prevails as before in non-active areas where different phases appear to be called for. However, the differences between the hemodynamic spline and sinusoids are much more modest than those realized by the periodic spline model for the time courses.

6.4 Further confirmations

To confirm that these results are not specific to the example data used, similar analyses were performed on a reference data set. Peter Jezzard kindly provided us with access to a copy of the data analyzed in Friston et al. [1994b], which is also the example data set used and described in Lange and Zeger [1997]. Our findings were essentially the same as here; periodic spline fits were a significant improvement over sinusoids in the areas of activation, and exhibited a biphasic feature much like we have seen here. Details are given in Crellin [1996].

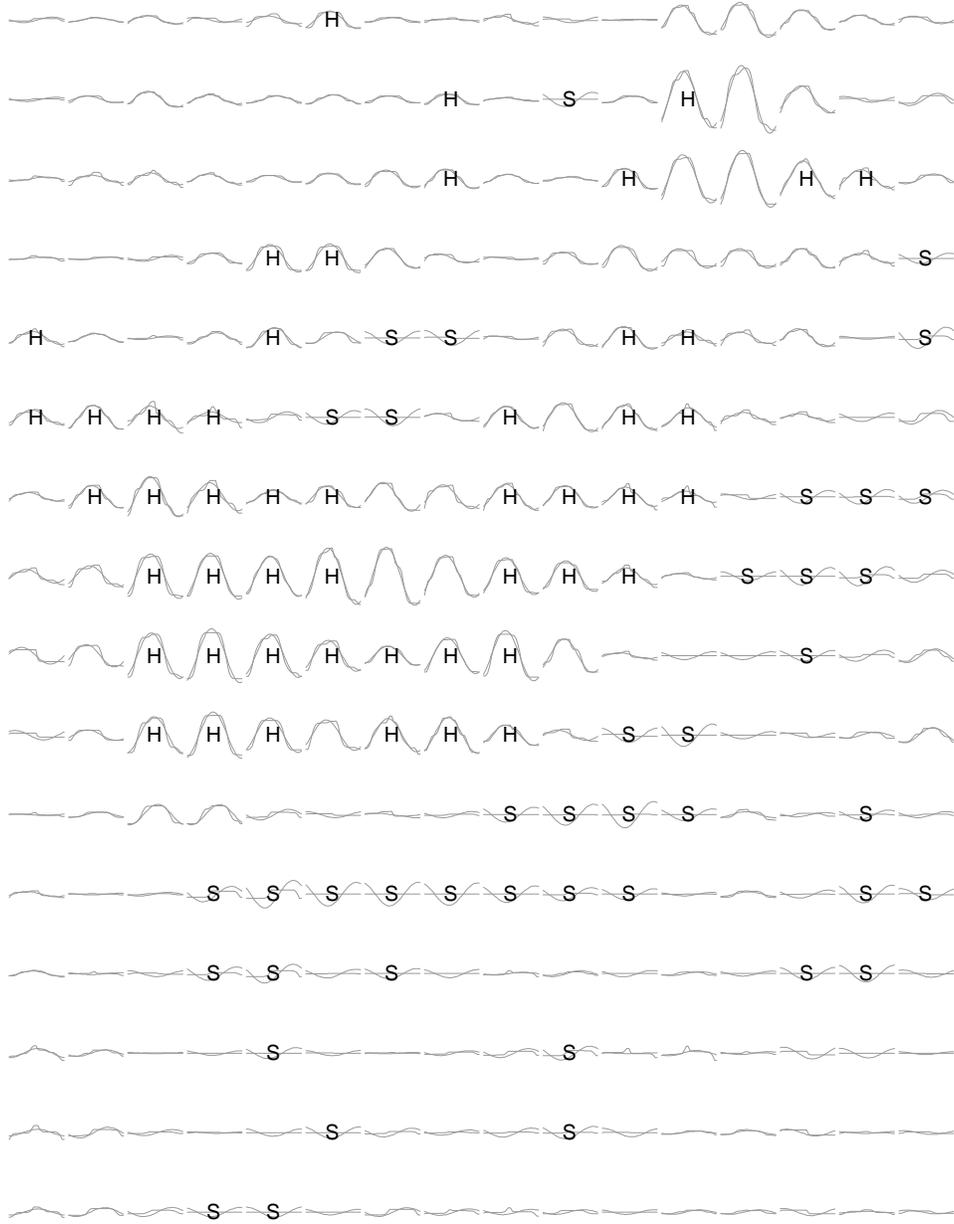


Figure 14: *Bootstrap simulation comparing the hemodynamic spline convolution model to sinusoid fits. In the areas marked “S”, the sinusoid model dominates, while in the areas marked “H”, the hemodynamic spline model dominates. Elsewhere they are not-significantly different.*

7 Discussion

This paper shows several important results about fMRI time course data. The first is that the convolution model imposes intrinsic structural constraints which limits its ability to capture the patterns seen in the periodic spline fits to the data. One possible modification which could improve matters would be to allow the hemodynamic response function to take negative values, at the cost of more difficult physiological interpretation¹. Models for the hemodynamic response function in current use are almost uniformly positive (Poisson, Gamma or Gaussian density functions), and as such will not capture the patterns seen in the periodic spline fits.

Another important result is that fitting the Poisson-based convolution model is shown to perform no better than fitting a sinusoid with arbitrary phase and amplitude at the activation frequency. However, fitting the spline-based convolution model does outperform the sinusoidal fits, despite the inherent constraints of the convolution model. The periodic spline fits to the time course data can differ substantially from the sinusoid fits, and we have given some evidence that these differences are real. They show biphasic behavior, which persists when the number of knots are increased, and also occur in spatially connected patterns.

References

- P.A. Bandettini, A. Jesmanowicz, E.C. Wong, and J.S. Hyde. Processing strategies for time-course data sets in functional MRI of the human brain. *Magnetic Resonance in Medicine*, 30(2):390–397, 1993.
- G.M. Boynton, S.A. Engel, G.H. Glover, and D.J. Heeger. Linear systems analysis of fMRI in human v1. *Journal of Neuroscience*, 16:4207–4221, 1996.
- N. Crellin. *Modeling Image Sequences, with particular application to fMRI data*. PhD thesis, Statistics Department, Stanford University, 1996.
- C. de Boor. *A Practical Guide to Splines*. Applied Mathematical Sciences. Springer-Verlag, New York, 1978.

¹David Heeger, in personal communication, proposes the following possible mechanism: immediately following neural activation but before vascular replenishment, there is a higher level of deoxyhemoglobin at activation sites. This could be modeled by negative hemodynamic response values. This usually happens faster than images are rescanned however.

- K.J. Friston, C.D. Frith, R.S.J. Frackowiak, and R. Turner. Characterizing dynamic brain responses with fMRI: A multivariate approach. *NeuroImage*, 2(2):166–172, 1995a.
- K.J. Friston, C.D. Frith, R. Turner, and R.S.J. Frackowiak. Characterizing evoked hemodynamics with fMRI. *NeuroImage*, 2(2):157–165, 1995b.
- 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(1):45–53, 1995c.
- K.J. Friston, A.P. Holmes, K.J. Worsley, J.-P. Poline, C.D. Frith, and R.S.J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2(4):189–210, 1994a.
- K.J. Friston, P. Jezzard, and R. Turner. Analysis of functional MRI time-series. *Human Brain Mapping*, 1(2):153–171, 1994b.
- G. Golub and C. Van Loan. *Matrix computations*. Johns Hopkins University Press, 1983.
- K.K. Kwong, J.W. Belliveau, D.A. Chesler, I.E. Goldberg, R.M. Weisskoff, B.P. Poncelet, D.N. Kennedy, B.E. Hoppel, M.S. Cohen, R. Turner, H. Cheng, T.J. Brady, and B.R. Rosen. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Science*, 89:5675–5679, 1992.
- N. Lange and S.L. Zeger. Non-linear fourier time series analysis for human brain mapping by functional magnetic resonance imaging. *Applied Statistics*, 46(1):1–29, 1997.
- C. Lawrence, J.L. Zhou, and A.L. Tits. User’s guide for CFSQP version 2.3: A C code for solving (large scale) constrained nonlinear (minimax) optimization problems, generating iterates satisfying all inequality constraints. Technical Report Institute for Systems Research TR-94-16r1, University of Maryland, College Park, MD 20742, USA, 1994.
- S. Ogawa, D.W. Tank., R. Menon, J.M. Ellerman, S.G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Science*, 89:5951–5955, 1992.
- K.J. Worsley and K.J. Friston. Analysis of fMRI time-series revisited – again. *NeuroImage*, 2(3):173–181, 1995.