

Bilinear Dynamical Systems

W. Penny¹, Z. Ghahramani² and K. Friston¹

Wellcome Department of Imaging Neuroscience¹,
Gatsby Computational Neuroscience Unit²,
University College, London WC1N 3BG.

February 11, 2005

Abstract

In this paper we propose the use of Bilinear Dynamical Systems (BDS) for model-based deconvolution of fMRI time series. The importance of this work lies in being able to deconvolve hemodynamic time-series, in an informed way, to disclose the underlying neuronal activity. Being able to estimate neuronal responses in a particular brain region is fundamental for many models of functional integration and connectivity in the brain. BDSs comprise a stochastic bilinear neurodynamical model specified in discrete-time and a set of linear convolution kernels for the hemodynamics. We derive an Expectation-Maximisation (EM) algorithm for parameter estimation, in which fMRI time series are deconvolved in an E-step and model parameters are updated in an M-Step. We report preliminary results that focus on the assumed stochastic nature of the neurodynamic model and compare the method to Wiener deconvolution.

Keywords: fMRI, Deconvolution, Connectivity, EM, Dynamical, Embedding

1 Introduction

Imaging neuroscientists have at their disposal a variety of imaging techniques for investigating human brain function [8]. The Electroencephalogram (EEG)

records electrical activity from electrodes placed on the scalp, the Magnetoencephalogram (MEG) records the magnetic field from sensors placed just above the head and functional Magnetic Resonance Imaging (fMRI) records changes in magnetic resonances due to variations in blood oxygenation. However, as the goal of brain imaging is to obtain information about the neuronal networks that support perceptual inference and cognition, one must first transform measurements from imaging devices into estimates of intracerebral electrical activity. Brain imaging methodologists are therefore faced with an inverse problem, ‘How can one make inferences about intracerebral neuronal processes given extracerebral/vascular measurements?’

Though not often expressed in this terminology, we argue that this problem is best formulated as a model-based spatio-temporal deconvolution problem. For EEG and MEG the required deconvolution is primarily spatial, and for fMRI it is primarily temporal. Although one can make minimal assumptions about the source signals by applying ‘blind’ deconvolution methods [23, 24], knowledge of the underlying physical processes can be used to great effect. This information can be implemented in a forward model that is inverted during deconvolution. In EEG/MEG, forward models make use of Maxwell’s equations governing propagation of electromagnetic fields [1], and in fMRI, forward models comprise hemodynamic processes as described by ‘Balloon’ models [4, 9, 28, 30].

In this paper, we propose a new state-space method [18] for model-based deconvolution of fMRI. The importance of this work lies in being able to deconvolve hemodynamic time-series, in an informed way, to disclose the underlying neuronal activity. Such estimates are required by many models of functional integration and connectivity in the brain. Typically, the influence of an experimental manipulation, on the coupling between two regions, is tested using a statistical model of the interaction between the experimental factor and neuronal activity in the source region. These interaction terms rest on being able to deconvolve the fMRI time-series. The procedures described below enable the construction of precise and informed interaction terms that can then be used to detect context sensitive coupling among brain regions. The interaction terms are variously known as Psychophysiological Interactions [10], moderator variables in Structural Equation Models [3], and bilinear inputs in Dynamic Causal Models [11].

The model we propose is called a Bilinear Dynamical System (BDS) and comprises the following elements. Experimental manipulation (input) causes changes in neuronal activation (state) which in turn cause changes in fMRI

signal (output). Experimental inputs fall into two classes (i) driving inputs which directly excite neurodynamics and (ii) modulatory inputs which change the neurodynamics. These modulatory inputs typically correspond to instructional or attention set, and therefore allow neurodynamic changes to be directly attributed to changes in experimental context. Readers of our earlier papers on Dynamic Causal Modelling (DCM) [11, 25] may well be experiencing a sense of ‘deja-vu’. This is no accident because BDS employs the same concepts of driving and modulatory inputs, neuronal states and outputs.

A key difference to DCM, however, is that BDS uses a stochastic neurodynamic model. This is biologically more realistic as regional dynamics are unlikely to be under full experimental control. Moreover, it has been established, across a variety of spatial scales, that neurodynamics can be meaningfully treated as stochastic processes, from Poisson processes describing spike generation [6] to statistical mechanical treatments of activity in cortical macrocolumns [21]. Stochastic components could also be considered as the consequence of local or global deterministic nonlinear dynamics of a possibly chaotic nature (see eg. [2]). From another perspective, stochastic components could be considered as reflecting input from remote regions [5].

In BDS, neurodynamics cause changes in fMRI signals via hemodynamic models specified in terms of a set of basis functions. Although more detailed differential equation models exist [4, 9, 28], the relationship between neuronal activation and fMRI signals can be formulated as a first-order convolution with a kernel expansion using basis functions (typically two or three). This kernel is the hemodynamic response function. This approach has enjoyed a decade of empirical success in the guise of the General Linear Model (GLM) [12, 19]. Moreover, this means that the state-output relation can be described by a linear convolution.

The BDS model may be viewed as a marriage of two formulations, the input-state relation being a stochastic DCM and the state-output relation being a GLM.

BDS has a similar form to models used in signal processing and machine learning. In particular, BDS is almost identical to a Linear Dynamical System (LDS) [29] with inputs. The only difference is that the inputs can change the linear dynamics in a bilinear fashion - hence the name ‘Bilinear’ Dynamical Systems. In such models, the hidden variable, which in our case corresponds to the neuronal time series, can be estimated using Kalman smoothing. Further, the parameters of the model can be estimated using

an Expectation-Maximisation (EM) algorithm described in [15] which needs only minor modification to accommodate the bilinear term.

The structure of the paper is as follows. In section 2 we define BDS and describe how it can be applied to deconvolve fMRI time series. We describe a novel EM algorithm for estimation of model parameters. In section 3 we present preliminary results on applying the models to synthetic data and data from an fMRI experiment.

1.1 Notation

We denote matrices and vectors with bold upper case and bold lower case letters respectively. All vectors are column vectors. \mathbf{X}^T denotes the transpose of \mathbf{X} and \mathbf{I}_k is the k -dimensional identity matrix. Brackets are used to denote entries in matrices/vectors. For example, $\mathbf{A}(i)$ denotes the i th row of \mathbf{A} , $\mathbf{A}(i, j)$ denotes the scalar entry in the i th row and j th column, and $\mathbf{a}(i)$ denotes the i th entry in \mathbf{a} . We define \mathbf{z}_N as a vector with a 1 as the first entry followed by $N - 1$ zeros and $\mathbf{0}_{I,J}$ as the zero matrix of dimension $I \times J$.

2 Bilinear Dynamical Systems

Although, in principle, we could define models for activity in a network of regions, in this paper we define a BDS model for activity in only a single region. BDS is an input-state-output model where the states correspond to ‘neuronal’ activations. Neuronal activity is defined mathematically below and can be thought of as that component of the Local Field Potential (LFP) to which fMRI is most sensitive [22]. BDS is defined, for a single region, with the following state-space equations where n indexes time

$$s_n = (a + \mathbf{b}^T \mathbf{u}_n) s_{n-1} + \mathbf{d}^T \mathbf{v}_n + w_n \quad (1)$$

$$\mathbf{x}_n = [s_n, s_{n-1}, s_{n-2}, \dots, s_{n-L+1}]^T \quad (2)$$

$$y_n = \boldsymbol{\beta}^T \boldsymbol{\Phi} \mathbf{x}_n + e_n \quad (3)$$

The first equation describes the stochastic neurodynamic model. Driving inputs \mathbf{v}_n cause an increase in neuronal activity s_n (a scalar) that decays with a time constant determined by $(a + \mathbf{b}^T \mathbf{u}_n)$. This time constant is determined by the value of the ‘intrinsic connection’, a , and ‘modulatory inputs’, \mathbf{u}_n . Driving inputs typically correspond to the presentation of experimental stimuli

and modulatory inputs typically correspond to instructional or attentional set. The strengths of the driving and modulatory effects are determined by the vectors of driving connections, \mathbf{d} , and modulatory connections, \mathbf{b} . The driving connections are also known as ‘input efficacies’. Neuronal activity is also governed by zero-mean additive Gaussian noise, w_n , having variance σ_w^2 .

Neuronal activation then gives rise to fMRI time series according to the second and third equations. The second equation defines an embedding process [32] in which neuronal activity over the last L time steps is placed in the ‘buffer’, \mathbf{x}_n . The vector \mathbf{x}_n is referred to as an ‘embedded’ time series where L is the embedding dimension. Neuronal activation, now described by \mathbf{x}_n , then gives rise to fMRI time series, y_n , via a linear convolution process described in the third equation. Equation 3 comprises a signal term $\beta^T \Phi \mathbf{x}_n$ and a zero-mean Gaussian error term e_n with variance σ_e^2 . The h th row of matrix Φ defines the h th convolution kernel (ie. basis of the hemodynamics response function) and $\beta(h)$ is the corresponding regression coefficient. Figure 1 shows a set of convolution kernels that have been used widely for the analysis of fMRI data in the context of the GLM. Linear combinations of these functions have been found empirically to account for most subject-to-subject and voxel-to-voxel variations in hemodynamic response [19, 20].

The model can perhaps be better understood by looking at the neuronal and hemodynamic time series that it generates. An example is shown in Figure 2. The model we have proposed, BDS, can be viewed as a combination of two models (i) a stochastic, discrete-time version of DCM to describe neurodynamics and (ii) the GLM to describe hemodynamics. GLMs are well established in the neuroimaging literature [8] and DCMs are becoming so. Consequently, the model choices we have made are informed by precedents established for GLMs and DCMs. For example, from previous work using GLMs [20], we know that the embedding dimension should be chosen to span a 20-30s period. With regard to neurodynamics, the discrete bilinear form affords analytic tractability while allowing nonlinear interactions between exogenous input and states. The splitting of exogenous inputs into driving and modulatory terms is also prompted by DCM. The motivation for this is that, for data collected from controlled experiments, one wishes to relate changes in connectivity to experimental manipulation. This distinction is echoed by the separate neurobiological mechanisms underlying driving versus modulatory activity [25].

The modulatory coefficients in equation 1, \mathbf{b} , are also known as ‘bilinear’ terms. This is because the dependent variable, s_n , is linearly dependent on

the product of s_{n-1} and \mathbf{u}_n . That \mathbf{u}_n and s_{n-1} combine in multiplicative fashion endows the model with ‘nonlinear’ dynamics that can be understood as a nonstationary linear system that changes according to experimental context. Importantly, because \mathbf{u}_n is known, it is straightforward to estimate how the dynamics are changing. We emphasise that the term ‘bilinear’ relates to interactions between a state and an input rather than among the states themselves.

2.1 Relation to biophysical models

In the previous section, we interpreted $(a + \mathbf{b}^T \mathbf{u}_n)$ as the time constant of decaying neuronal activity. However, our BDS model has a much more general relationship to underlying biophysical models of neuronal dynamics. Consider some biologically plausible model of neuronal activity that is formulated in continuous time $\dot{x} = f(x, \mathbf{u}, \boldsymbol{\theta})$ and allows for arbitrarily complicated and nonlinear effects of the state and exogenous inputs. The parameters of this biophysical model $\boldsymbol{\theta}$ (c.f. neural mass models [31]) can be related directly to biophysical processes. In our discrete-time formulation, under local linearity assumptions, $a = \exp(\Delta t \partial f / \partial x)$ and $\mathbf{b}(i) = \partial a / \partial \mathbf{u}(i) = \Delta t a \partial^2 f / \partial x \partial \mathbf{u}(i)$ where Δt is the sampling interval and i indexes the input. This means that our input-dependent autoregression coefficients can be regarded as a lumped representation of the underlying model parameters. This re-parameterization, in terms of a state-space model, precludes an explicit estimation of the original parameters. However, this does not matter because we are not interested in the parameters per se; we only require the conditional estimates of the states. The transformation of a dynamic formulation into a discrete BDS is a useful perspective because priors on the biophysical parameters can be used to specify priors on the autoregression coefficients, should they be needed.

2.2 Relation to GLMs

For models with a single input BDS reduces to a GLM if $s_n = v_n$, that is, if the neuronal activity is synonymous with the experimental manipulation (driving inputs). This highlights a key assumption of GLMs, that the dynamics of neuronal transients can be ignored.

For models with multiple inputs the relation between BDS and GLMs is more complex. Ignoring bilinear and noise terms, BDS can be more parsimonious

monious than the GLM in the sense of having fewer parameters. Say, for example, we are modelling hemodynamics with H basis functions. Then for M (driving) input variables, the GLM has HM parameters, whereas in BDS, there are $M + H + 1$ parameters. For $H = 3$, for example, BDS has fewer coefficients if $M \geq 3$.

The reason for this reduction is that, in BDS, the inputs affect only neurodynamics. And the strength of this effect, ignoring modulatory terms, is captured in a single parameter, the input efficacy. The relation between neurodynamics and hemodynamics is independent of which input excited neuronal activity. This is obviously more consistent with physiology where experimental manipulation only affects fMRI by changing neuronal activity.

2.3 Deconvolution

Deconvolution consists of estimating a neuronal time series, s_n , given an fMRI time series, y_n . It is possible to perform model-based deconvolution using BDS and a modified Kalman-filtering algorithm. The modification is necessary because, due to the modulatory terms, the state transition matrix is time-dependent. Also, the BDS model must be reformulated so that the output is an instantaneous function of the state. This section describes standard Kalman filtering and the modifications required for BDS.

In the Kalman-filtering approach, deconvolution progresses via iterative computation of the probability density $p(s_n|\mathbf{y}_1^n)$ where $\mathbf{y}_1^n = \{y_1, y_2, \dots, y_n\}$. That is, our estimate of neuronal activity at time step n is based on all the fMRI data up to time n (but not on future values).

The Kalman-filtering algorithm can be split into two steps that are applied recursively. The first step is a *time update* where the density $p(s_{n-1}|\mathbf{y}_1^{n-1})$ is updated to $p(s_n|\mathbf{y}_1^{n-1})$. In our model this will take into account the effects of inputs \mathbf{u}_n and \mathbf{v}_n and the natural decay of neuronal activity. The second step is a *measurement update* where the density $p(s_n|\mathbf{y}_1^{n-1})$ is updated to $p(s_n|\mathbf{y}_1^n)$ thereby taking into account the ‘new’ fMRI measurement y_n . The two steps together take us from time point $n - 1$ to time point n and are applied recursively to deconvolve the entire time-series.

A complication arises, however, because the Kalman filtering updates require that the output be an instantaneous function of the hidden state. That is, that y_n depend only on s_n and not on s_{n-1} , s_{n-2} etc, which is clearly not the case. But Kalman filtering can proceed if we use the embedded state variables \mathbf{x}_n . This will lead to an estimate of the multivariate den-

sity $p(\mathbf{x}_n|\mathbf{y}_1^n)$. Our desired density, $p(s_n|\mathbf{y}_1^n)$, is then just the first entry in $p(\mathbf{x}_n|\mathbf{y}_1^n)$. Whilst this is conceptually simple, a good deal of book-keeping is required to translate the BDS neurodynamical model into an ‘embedded space’. This is described in detail in Appendices A and B.

A second complication arises in that, in the standard Kalman update formulae [15], the inputs do not change the intrinsic dynamics. To account for this we introduce a time-dependent state-transition matrix that embodies changes in intrinsic dynamics due to modulatory inputs. This is also detailed in Appendices A and B.

It is also possible to perform deconvolution based on Kalman-smoothing rather than Kalman-filtering. Kalman-smoothing estimates neuronal activity by computing the density $p(s_n|\mathbf{y}_1^N)$. It is therefore based on the whole fMRI record, ie. past *and* future values and, as we shall see, provides more accurate deconvolutions in high noise environments.

2.4 Estimation

To run the deconvolution algorithms described in the previous section we must know the parameters of the model eg. β , a , \mathbf{b} and \mathbf{d} . If these parameters are unknown, as will nearly always be the case, they can be estimated using Maximum-Likelihood (ML) methods.

For BDS this can be implemented using an Expectation-Maximisation (EM) algorithm which we derive in Appendices A and B. The E-step uses a Kalman smoothing algorithm which contains only minor modifications to the algorithm presented in [15]. The M-step contains updates derived by considering estimation of the neurodynamic parameters (a , \mathbf{b} , \mathbf{d}) as a constrained linear regression problem [27]. The neurodynamic parameters are updated as shown in equation 25 and the hemodynamic parameters, β , are updated using equation 28. It is also possible to update the initial state estimates as described in [15], but this was not implemented for the empirical work in this paper.

It is also possible to implement ML parameter estimation by computing the likelihood using Kalman filtering (as shown in Appendix C), and then maximising this using standard Matlab optimisation algorithms such as Pseudo-Newton or Simplex methods [26]. However, preliminary simulations showed EM to be faster than these other ML methods.

2.5 Initialisation

The EM estimation algorithm is initialised as follows. In the limit of Zero Neuronal Noise (ZNN) the neuronal activity is given by

$$s_n = (a + \mathbf{b}^T \mathbf{u}_n) s_{n-1} + \mathbf{d}^T \mathbf{v}_n \quad (4)$$

and the predicted hemodynamic activity is

$$\hat{y}_n = \boldsymbol{\beta}^T \boldsymbol{\Phi} \mathbf{x}_n \quad (5)$$

We initially set a to a random number between 0 and 1, and \mathbf{b} such that $\mathbf{b}^T \mathbf{u}_n$ is between 0 and $1 - a$ for all n . The values of \mathbf{d} are also set to random numbers between 0 and 1. The first $\boldsymbol{\beta}$ coefficient is fixed at unity, and others, if there are any, are initialised to random numbers between 0 and 1. We then run a pseudo-Newton gradient descent optimisation (the function `fminunc` in Matlab) to minimise the discrepancy between the observed fMRI time series y_n and the estimated values \hat{y}_n under the ZNN assumption. Occasionally, parameter estimates result in an unstable model. We therefore repeat these initialisation steps until a stable model is returned. These parameters are then used as a starting point for EM.

3 Results

3.1 Simulations

We generate simulated data to demonstrate various properties of the deconvolution algorithm. These simulations are similar to the first set of simulations in [16]. We also compare our results to those obtained with Wiener deconvolution [17], which makes minimal assumptions about the source (that it has a flat spectral density [26, 16]).

A 250s time series of input events with sampling period $\Delta t = 0.5$ s was generated from a Poisson distribution with a mean interval between events of 12s. Events less than 2s apart were then removed. The half-life of neuronal transients was fixed to 1s by setting $a = 0.71$, the input efficacy was set to $d = 0.9$ and the neuronal noise variance was set to $\sigma_w^2 = 0.0001$. There were no modulatory coefficients and a single hemodynamic basis function (the ‘canonical’, shown in Figure 1) was used.

We then added observation noise to achieve a Signal to Noise Ratio (SNR) of unity (this is equivalent to the 100% noise case in [16]). This was achieved using $\sigma_e^2 = 0.015$. The generated time series is shown in Figure 3.

The parameters were estimated using EM to be $a = 0.72$ and $d = 0.88$. Wiener and BDS Kalman smoothing deconvolutions (the latter obtained using estimated parameters) are shown in Figure 4. Wiener deconvolution used the known value of σ_e^2 and BDS Kalman smoothing used the known values of σ_e^2 and σ_w^2 . The quality of the BDS deconvolution is considerably better than that using the method described in [16] (see eg. Figure 3D in [16]). This is because information about the paradigm (eg. experimental inputs) has been used. For Wiener deconvolution the correlation with the generated neuronal time series is $r = 0.520$ and for BDS Kalman smoothing it is $r = 0.998$.

We then ran a second simulation with the neuronal noise variance set to a high value, $\sigma_w^2 = 0.03$. This data is shown in the top panel of Figure 5. Neurodynamic parameters were estimated, using EM, to be $a = 0.68$ and $d = 0.83$. For comparison, ZNN initialisation (see section 2.5) produced parameter estimates $a = 0.45$ and $d = 1.13$. That is, the assumption of zero neuronal noise leads to an underestimate of the neuronal time constant and an overestimate of the input efficacy.

The neuronal time series, as estimated using BDS Kalman smoothing (bottom panel of Figure 5) had a correlation with the generated time series of $r = 0.775$. The ZNN estimate neuronal time series had a correlation of $r = 0.724$. The Wiener deconvolved time series (second panel in Figure 5) gave $r = 0.52$ (again). We also show, in the third panel of Figure 5, deconvolutions from BDS Kalman filtering.

Figure 5 shows that Wiener estimation recovers the intrinsic dynamics but misses the evoked responses, whereas BDS Kalman filtering recovers the evoked responses but misses the intrinsic dynamics. Deconvolution using BDS Kalman smoothing recovers both.

The smoother can capture the intrinsic dynamics because it uses more information than the filter. It updates the filter estimates of neuronal activity, at a given time point, using information in advance of that time point (ie. from the future). These updates are implemented using the ‘backward recursions’ described in Appendix B.1. Heuristically, the reason for the improvement is that the best estimates of neuronal activity are obtained using observed fMRI activity about 5 seconds or so in the future ie. at the peak of the hemodynamic response.

3.2 Single word processing fMRI

We now turn to the analysis of an fMRI data set recorded during a passive listening task, using epochs of single words presented at different rates. The experimental inputs in Figure 6 describe the paradigm in more detail. The driving inputs in the top panel indicate presentation of words in the subject’s headphones and modulatory inputs in the lower panels indicate epochs with different presentation rates.

We focus on a single time series in primary auditory cortex shown in Figure 7. This comprises 120 time points with a sampling period $\Delta t = 1.7\text{s}$. Further details of data collection are given in [11]. As the experimental inputs are specified at a higher temporal resolution than the fMRI acquisition, we upsampled the fMRI data by a factor of 4 prior to analysis. The input variables were convolved with a ‘canonical’ hemodynamic response function (see top panel of Figure 1) to form regressors in a GLM. Figure 7 shows the resulting GLM model fit.

The same input variables and hemodynamic basis function were then used to define a BDS model. The observation and state noise variances were set to $\sigma_e^2 = \sigma_w^2 = 0.1$, and the hemodynamic regression coefficient was set to $\beta = 1$. Figure 8 shows the resulting model fit which is clearly superior to the GLM model fit in Figure 7 (but see Discussion). Figure 9 shows neuronal activity as estimated using BDS Kalman smoothing. This includes both event-related responses (the ‘spikes’) and intrinsic activity (slow fluctuations).

Neurodynamic parameters were estimated, using EM, as $d = 0.80$, $a = 0.92$, $\mathbf{b}(1) = -0.44$, $\mathbf{b}(2) = -0.08$ and $\mathbf{b}(3) = -0.02$. Thus, epochs with faster stimulus presentations were estimated to have an increasingly inhibitory effect on event-related neurodynamics. This pattern was robust across a wide range of settings of the noise variance parameters. This finding is consistent with neuronal repetition suppression effects and is in agreement with Dynamic Causal Modelling of this data [11].

4 Discussion

We have proposed a new algorithm, based on a Bilinear Dynamical System (BDS), for model-based deconvolution of fMRI time series. The importance of the work is that hemodynamic time-series can be deconvolved, in an informed way, to disclose the underlying neuronal activity. Being able

to estimate neuronal responses in a particular brain region is fundamental for many models of functional integration and connectivity in the brain. Of course, these estimates can only describe those components of neuronal activity to which fMRI is sensitive. They should nevertheless be useful for making fMRI-based assessments of connectivity [10, 16].

BDS is fitted to data using a novel Expectation-Maximisation (EM) algorithm where deconvolution is instantiated in an E-step and model parameters are updated in an M-step. Simulations showed EM to be faster than maximum likelihood optimisation based on simplex or Pseudo-Newton methods. Deconvolution can be based either on a full E-step, using Kalman smoothing, or a partial E-step based on Kalman filtering. Kalman smoothing uses the full data record whereas Kalman filtering only uses information from the past.

Simulations showed that our model-based deconvolution is more accurate than blind deconvolution methods (Wiener filtering). This is because BDS uses information about the paradigm. We also observed the following trends. Wiener estimation recovers the intrinsic dynamics but misses the evoked responses, whereas BDS Kalman filtering recovers the evoked responses but misses the intrinsic dynamics. Deconvolution using BDS Kalman smoothing recovers both.

Simulations also suggest that if dynamics are indeed of a stochastic nature, as is assumed in BDS, then if we mistakenly assume deterministic dynamics, estimation of neuronal efficiencies and time constants will become inaccurate. This has implications for models that assume deterministic dynamics, such as DCM [11].

Our applications of BDS provide good examples of Kalman smoothing providing better deconvolutions than Kalman filtering. The reason is that smoothing also uses future observations and the best estimates of neuronal activity are obtained using observed fMRI activity about 5 seconds or so into the future, that is, at the peak of the hemodynamic response. This property should hold for any state-space model of fMRI (see eg. [28]).

A comparison of GLM and BDS model fits in figures 7 and 8 clearly shows that BDS is superior. Whilst this is somewhat encouraging, this observation should be tempered with a note of caution. This is because the BDS model is more complex and no penalty was paid for this during model fitting. The BDS model may therefore be overfitted. In particular, as fMRI time series are known to contain aliased cardiac and respiratory artifacts, the finer details that BDS picks up may be of artifactual rather than neuronal origin. To eliminate this possibility one would need to estimate model parameters

(especially the state noise variance) using Bayesian or cross-validation [33] methods. Fortunately, Bayesian methods have already been developed for Linear Dynamical Systems [14] that should not require too many changes to accommodate the bilinear term.

Recently, Riera et al. [28] have proposed a state-space model for fMRI time series which allows for nonlinear state-output relations. They showed how model parameters could be estimated using a maximum likelihood procedure based on Extended Kalman Filtering. Moreover, they showed how the parameters could be related to a stochastic differential equation implementation of the Balloon model. Thus, unlike BDS which assumes linear hemodynamics, their model can describe nonlinear properties such as hemodynamic refractoriness. This is important, as the neuronal refractoriness inferred using BDS in section 3.2, for example, may actually be of hemodynamic rather than neuronal origin. In fact, a Dynamic Causal Modelling analysis of this data [11], which allows for both neuronal and hemodynamic refractoriness, concluded that both effects were present.

The model proposed by Riera et al. [28] also differs from BDS in the input-state relations. In BDS, this is governed by driving inputs that excite linear neurodynamics which can be changed according to modulatory inputs. This means that neurodynamics and changes in them can be directly attributed to changes in experimental context. This modelling approach is therefore appropriate for designed experiments where the inputs are known. In the Riera model, the inputs are not assumed to be known. Instead, a more flexible radial basis function approach is used to make inferences about the onsets of linear neurodynamical processes.

In previous work we have proposed an fMRI deconvolution model based on the GLM [16]. This uses a forward model in which neuronal activity, represented using a temporal basis set and corresponding coefficients, is convolved with a known hemodynamic kernel to produce an fMRI time series. Observation noise is then added. Deconvolution is achieved by estimating the coefficients of the temporal basis functions. In [16] we used full-rank Fourier basis sets and overcomplete basis sets that used information about experimental design. Coefficients were estimated using priors and Parametric Empirical Bayes. A problem with the method, however, is that for J coefficients (where J is typically the length of the time series or longer) one must store and invert $J \times J$ covariance matrices. A computational benefit of the BDS approach is that, by making use of factorisations that derive from Markov properties of the state-space model, one need only manipulate

$L \times L$ covariance matrices, where L is the temporal embedding dimension ($L \ll J$). We have also argued, in section 3.2, that BDS is more parsimonious and biologically informed than the GLM.

The method we have proposed has been applied to deconvolve data at a single voxel. This is useful in providing a ‘source’ neuronal time series, for example, for the analysis of PPIs [10, 16]. More generally, however, one is interested in making inferences about changes in connectivity in neural networks extended over multiple regions. This requires simultaneous deconvolution at multiple voxels and, ideally, a full model-based spatio-temporal deconvolution. The application of state-space models to this more difficult problem is an exciting area of current research (see eg. [13]).

Acknowledgements

Will Penny and Karl Friston are funded by the Wellcome Trust.

References

- [1] S. Baillet, J.C. Mosher, and R.M. Leahy. Electromagnetic Brain Mapping. *IEEE Signal Processing Magazine*, pages 14–30, November 2001.
- [2] M. Breakspear and J.R. Terry. Nonlinear interdependence in neural systems: motivation, theory and relevance. *International Journal of Neuroscience*, 112:1263–1284, 2002.
- [3] C. Buchel and K.J. Friston. Modulation of connectivity in visual pathways by attention: Cortical interactions evaluated with structural equation modelling and fMRI. *Cerebral Cortex*, 7:768–778, 1997.
- [4] R.B. Buxton, E.C. Wong, and L.R. Frank. Dynamics of blood flow and oxygenation changes during brain activation: The Balloon Model. *Magnetic Resonance in Medicine*, 39:855–864, 1998.
- [5] O. David and K.J. Friston. A neural mass model for MEG/EEG: coupling and neuronal dynamics. *NeuroImage*, 20(3):1743–1755, 2003.
- [6] P. Dayan and L.F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, 2001.

- [7] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B*, 39:1–38, 1977.
- [8] R.S.J. Frackowiak, K.J. Friston, C. Frith, R. Dolan, C.J. Price, S. Zeki, J. Ashburner, and W.D. Penny, editors. *Human Brain Function*. Academic Press, 2003. 2nd Edition.
- [9] K.J. Friston. Bayesian estimation of dynamical systems: An application to fMRI. *NeuroImage*, 16:513–530, 2002.
- [10] K.J. Friston, C. Buchel, G.R. Fink, J. Morris, E. Rolls, and R. Dolan. Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage*, 6:218–229, 1997.
- [11] K.J. Friston, L. Harrison, and W.D. Penny. Dynamic Causal Modelling. *NeuroImage*, 19(4):1273–1302, 2003.
- [12] K.J. Friston, A.P. Holmes, K.J. Worsley, J.B. Poline, C. Frith, and R.S.J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2:189–210, 1995.
- [13] A. Galka, O. Yamashita, T. Ozaki, R. Biscay, and P. Valdes-Sosa. A solution to the dynamical inverse problem of EEG generation using spatiotemporal Kalman filtering. *NeuroImage*, 2004. In Press.
- [14] Z. Ghahramani and M.J. Beal. Propagation algorithms for Variational Bayesian learning. In T. Leen et al, editor, *NIPS 13*, Cambridge, MA, 2001. MIT Press.
- [15] Z. Ghahramani and G.E. Hinton. Parameter Estimation for Linear Dynamical Systems. Technical Report CRG-TR-96-2, Department of Computer Science, University of Toronto, 1996. Also available from <http://www.gatsby.ucl.ac.uk/~zoubin/papers.html>.
- [16] D.R. Gitelman, W.D. Penny, J. Ashburner, and K.J. Friston. Modeling regional and psychophysiological interactions in fMRI: the importance of hemodynamic deconvolution. *NeuroImage*, 19:200–207, 2003.
- [17] G.H. Glover. Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage*, 9:416–429, 1999.

- [18] S. Haykin. *Adaptive Filter Theory, 3rd Edition*. Prentice-Hall, 1996.
- [19] R.N.A. Henson. *Human Brain Function*, chapter 40. Academic Press, 2003. 2nd Edition.
- [20] R.N.A. Henson, M.D. Rugg, and K.J. Friston. The choice of basis functions in event-related fMRI. *NeuroImage*, 13(6):127, June 2001. Supplement 1.
- [21] L. Ingber. Statistical mechanics of multiple scales of neocortical interactions. In P. Nunez, editor, *Neocortical dynamics and human EEG rhythms*. Oxford University Press, 1995.
- [22] N.K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann. Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412:150–157, 2001.
- [23] S. Makeig, M. Westerfield, T.-P. Jung, S. Enghoff, J. Townsend, E. Courchesne, and T.J. Sejnowski. Dynamic brain sources of visual evoked responses. *Science*, 295:690–694, 2002.
- [24] M.J. McKeown, S. Makeig, G.G. Brown, T.P. Jung, S.S. Kindermann, A.J. Bell, and T.J. Sejnowski. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping*, 6:160–188, 1998.
- [25] W.D. Penny, K.E. Stephan, A. Mechelli, and K.J. Friston. Comparing Dynamic Causal Models. *NeuroImage*, 22(3):1157–1172, 2004.
- [26] W. H. Press, S.A. Teukolsky, W.T. Vetterling, and B.V.P. Flannery. *Numerical Recipes in C*. Cambridge, 1992.
- [27] C.R Rao and H. Toutenberg. *Linear Models: Least squares and alternatives*. Springer, 1995.
- [28] J. Riera, J. Watanabe, K. Iwata, N. Miura, E. Aubert, T. Ozaki, and R. Kawashima. A state-space model of the hemodynamic approach: nonlinear filtering of BOLD signals. *NeuroImage*, 21:547–567, 2004.
- [29] S. Roweis and Z. Ghahramani. A unifying review of linear gaussian models. *Neural Computation*, 11(2):305–346, 1999.

- [30] K.E. Stephan, L. Harrison, W.D. Penny, and K.J. Friston. Biophysical models of fMRI responses. *Current Opinion in Neurobiology*, 14(5):629–635, 2004.
- [31] P.A. Valdes, J.C. Jimenez, J. Riera, R. Biscay, and T. Ozaki. Nonlinear EEG analysis based on a neural mass model. *Biological Cybernetics*, 81:415–424, 1999.
- [32] A.S. Weigend and N.A. Gershenfeld. *Time series prediction: forecasting the future and understanding the past*. Addison-Wesley, 1994.
- [33] O. Yamashita, A. Galka, T. Ozaki, R. Biscay, and P. Valdes-Sosa. Recursive penalised least squares solution for dynamical inverse problems of EEG generation. *Human Brain Mapping*, 21:221–235, 2004.

A Embedding

It is convenient to re-write the neurodynamic model as

$$s_n = \tilde{\mathbf{b}}^T \tilde{\mathbf{u}}_n s_{n-1} + \mathbf{d}^T \mathbf{v}_n + w_n \quad (6)$$

and $\tilde{\mathbf{u}}_n$ is a vector of modulatory inputs augmented with a 1 as the first entry (note some inputs may be driving *and* modulatory in which case they appear in both $\tilde{\mathbf{u}}_n$ and \mathbf{v}_n) and $\tilde{\mathbf{b}}^T = [a, \mathbf{b}]^T$ absorbs a .

The state-space equations can then be written in terms of the embedded neuronal activity, \mathbf{x}_n , as

$$\begin{aligned} \mathbf{x}_n &= \mathbf{A}_n \mathbf{x}_{n-1} + \mathbf{D} \mathbf{v}_n + \tilde{\mathbf{w}}_n \\ y_n &= \mathbf{c}^T \mathbf{x}_n + e_n \end{aligned} \quad (7)$$

where $\mathbf{c}^T = \beta^T \Phi$. The embedded state transition matrix is

$$\mathbf{A}_n = \begin{bmatrix} \tilde{\mathbf{b}}^T \mathbf{F}_n \\ \Psi_L \end{bmatrix} \quad (8)$$

where

$$\mathbf{F}_n = \tilde{\mathbf{u}}_n \mathbf{z}_L^T \quad (9)$$

is a $(B+1) \times L$ matrix (where B is the number of modulatory inputs) and Ψ_L is the $(L-1) \times L$ delay matrix that fills the lower $L-1$ rows of \mathbf{A}_n . This ensures that the embedded time series are shifted one time step each time \mathbf{A}_n is applied. The input matrix is given by

$$\mathbf{D} = \begin{bmatrix} \mathbf{d}^T \\ \mathbf{0}_{L-1, M} \end{bmatrix} \quad (10)$$

where M is the number of driving inputs. The state noise in equation 7 is given by $\tilde{\mathbf{w}}_n = w_n \mathbf{z}_L^T$. The covariance of $\tilde{\mathbf{w}}_n$ is \mathbf{Q} and the only non-zero entry is $\mathbf{Q}(1, 1) = \sigma_w^2$.

If we have a BDS with a single driving input, no modulatory inputs and $L = 4$, $a = 0.92$ and $d = 0.80$ then the embedded neurodynamic model is

$$\begin{aligned} \mathbf{x}_n &= \mathbf{A}_n \mathbf{x}_{n-1} + \mathbf{D} \mathbf{v}_n + \mathbf{w}_n \\ \begin{bmatrix} s_n \\ s_{n-1} \\ s_{n-2} \\ s_{n-3} \end{bmatrix} &= \begin{bmatrix} 0.92 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} s_{n-1} \\ s_{n-2} \\ s_{n-3} \\ s_{n-4} \end{bmatrix} + \begin{bmatrix} 0.80 \\ 0 \\ 0 \\ 0 \end{bmatrix} v_n + \begin{bmatrix} w_n \\ 0 \\ 0 \\ 0 \end{bmatrix} \end{aligned} \quad (11)$$

Modulatory inputs would change the first entry in \mathbf{A}_n .

B EM algorithm

In probabilistic models with hidden variables maximum-likelihood learning can be implemented using an Expectation-Maximisation (EM) algorithm [7]. This requires us to maximise the auxiliary function

$$F = \int p(\mathbf{X}|\mathbf{y}) \log p(\mathbf{X}, \mathbf{y}) d\mathbf{X} \quad (12)$$

where $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n, \dots, \mathbf{x}_N\}$ are the hidden variables and $\mathbf{y} = \{y_1, \dots, y_n, \dots, y_N\}$ are the observed variables [15]. In BDS, the observed variables are the fMRI time series and the hidden variables are the neuronal activities. Using the Markov property we can write

$$p(\mathbf{X}, \mathbf{y}) = p(\mathbf{x}_1) \prod_{n=2}^N p(\mathbf{x}_n|\mathbf{x}_{n-1}) \prod_{n=1}^N p(y_n|\mathbf{x}_n) \quad (13)$$

The initial, transition and output probabilities are given by

$$p(\mathbf{x}_1) = \mathbf{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \quad (14)$$

$$p(\mathbf{x}_n|\mathbf{x}_{n-1}) = \mathbf{N}(\mathbf{A}_n \mathbf{x}_{n-1} + \mathbf{D} \mathbf{v}_n, \mathbf{Q}) \quad (15)$$

$$p(y_n|\mathbf{x}_n) = \mathbf{N}(\mathbf{c}^T \mathbf{x}_n, \sigma_e^2) \quad (16)$$

which define the observation model, state transition model and initial state distribution. In the above expression the quantity \mathbf{c} is as defined in equation 7. Therefore the joint log-likelihood is a sum of quadratic terms

$$\begin{aligned} L_J = \log p(\mathbf{X}, \mathbf{y}) &= - \sum_{n=1}^N \frac{1}{2\sigma_e^2} \left[(y_n - \mathbf{c}^T \mathbf{x}_n)^T (y_n - \mathbf{c}^T \mathbf{x}_n) \right] - \frac{N}{2} \log |\sigma_e^2| \quad (17) \\ &- \sum_{n=2}^N \frac{1}{2} \left[(\mathbf{x}_n - \mathbf{A}_n \mathbf{x}_{n-1} - \mathbf{D} \mathbf{v}_n)^T \mathbf{Q}^{-1} (\mathbf{x}_n - \mathbf{A}_n \mathbf{x}_{n-1} - \mathbf{D} \mathbf{v}_n) \right] \\ &- \frac{1}{2} \left[(\mathbf{x}_1 - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_1^{-1} (\mathbf{x}_1 - \boldsymbol{\mu}_1) \right] - \frac{1}{2} \log |\boldsymbol{\Sigma}_1| \\ &- \frac{N-1}{2} \log |\mathbf{Q}| - \frac{NL}{2} \log 2\pi \end{aligned}$$

In the above equation, the quantity \mathbf{Q} refers to an arbitrary covariance matrix. In Appendix B.2 we show how L_J , as a function of the state variables \mathbf{x}_n , simplifies for the \mathbf{Q} defined for the BDS model in Appendix A. The expectation over the terms in the above equation can be maximised as shown in the following sections.

B.1 E-Step

The objective of the E-step is to compute the probability of the hidden variables given the data. Because the initial, transition and output probability distributions are Gaussian this can be achieved by updating the conditional mean and conditional covariance. Following [15], we write the expected value of \mathbf{x}_n conditioned on all data up to time n as $\mathbf{x}_n^n \equiv E[\mathbf{x}_n|\mathbf{y}_1^n]$. Similarly, the corresponding covariance is given by $\Sigma_n^n \equiv \text{Var}[\mathbf{x}_n|\mathbf{y}_1^n]$.

Kalman filtering

The objective of Kalman filtering is compute the probability of the hidden variables given all observed variables up to that time point, that is, to compute, $p(\mathbf{x}_n|\mathbf{y}_1^n)$. Kalman filtering implements the recursive computation of \mathbf{x}_n^n and Σ_n^n from \mathbf{x}_{n-1}^{n-1} and Σ_{n-1}^{n-1} in two steps. Firstly, in the time update step

$$p(\mathbf{x}_n|\mathbf{y}_1^{n-1}) = \int p(\mathbf{x}_n|\mathbf{x}_{n-1})p(\mathbf{x}_{n-1}|\mathbf{y}_1^{n-1})d\mathbf{x}_{n-1} \quad (18)$$

This is implemented using

$$\begin{aligned} \mathbf{x}_n^{n-1} &= \mathbf{A}_n\mathbf{x}_{n-1}^{n-1} + \mathbf{D}\mathbf{v}_n \\ \Sigma_n^{n-1} &= \mathbf{A}_n\Sigma_{n-1}^{n-1}\mathbf{A}_n^T + \mathbf{Q} \end{aligned} \quad (19)$$

Secondly, in the measurement update step

$$p(\mathbf{x}_n|\mathbf{y}_1^n) = \frac{p(y_n|\mathbf{x}_n)p(\mathbf{x}_n|\mathbf{y}_1^{n-1})}{\int p(y_n|\mathbf{x}_n)p(\mathbf{x}_n|\mathbf{y}_1^{n-1})d\mathbf{x}_n} \quad (20)$$

This is simply Bayes rule where $p(\mathbf{x}_n|\mathbf{y}_1^{n-1})$ (from equation 18) describes our belief in \mathbf{x}_n before observing y_n . The measurement update is implemented using

$$\begin{aligned} \mathbf{K}_n &= \Sigma_n^{n-1}\mathbf{c}(\mathbf{c}^T\Sigma_n^{n-1}\mathbf{c} + \sigma_e^2)^{-1} \\ \mathbf{x}_n^n &= \mathbf{x}_n^{n-1} + \mathbf{K}_n(y_n - \mathbf{c}^T\mathbf{x}_n^{n-1}) \\ \Sigma_n^n &= \Sigma_n^{n-1} - \mathbf{K}_n\mathbf{c}^T\Sigma_n^{n-1} \end{aligned} \quad (21)$$

where \mathbf{K}_n , known as the Kalman gain matrix, operates as an adaptive step size parameter for each hidden variable. In the above expressions the quantity

\mathbf{c} is as defined in equation 7. The procedure is initialised using $\mathbf{x}_1^0 = \boldsymbol{\mu}_1$ and $\boldsymbol{\Sigma}_1^0 = \boldsymbol{\Sigma}_1$. These updates are exactly as described in [15] except that (i) \mathbf{A}_n is used instead of \mathbf{A} because our dynamics are input-dependent, and (ii) we use \mathbf{c} instead of \mathbf{C} because our observations are univariate.

Kalman smoothing

The objective of Kalman smoothing is compute the probability of the hidden variables given all observed variables, that is, to compute, $p(\mathbf{x}_n | \mathbf{y}_1^N)$. They are implemented using a set of ‘backward recursions’ which compute \mathbf{x}_{n-1}^N and $\boldsymbol{\Sigma}_{n-1}^N$ from the forward estimates \mathbf{x}_{n-1}^{n-1} , $\boldsymbol{\Sigma}_{n-1}^{n-1}$. Because these formulae are also almost identical to those described in [15], we do not reproduce them here. The only difference is that \mathbf{A}_n is used instead of \mathbf{A} .

Expectations

The M-step requires a number of expectations that can be derived from the E-step

$$\begin{aligned} \mathbf{m}_n &= \int p(\mathbf{x}_n | \mathbf{y}) \mathbf{x}_n d\mathbf{x}_n \\ \mathbf{P}_n &= \int p(\mathbf{x}_n | \mathbf{y}) \mathbf{x}_n \mathbf{x}_n^T d\mathbf{x}_n \\ \mathbf{P}_{n,n-1} &= \int p(\mathbf{x}_n, \mathbf{x}_{n-1} | \mathbf{y}) \mathbf{x}_n \mathbf{x}_{n-1}^T d\mathbf{x}_n \end{aligned} \quad (22)$$

These can be computed as shown in [15], with the minor modification that the updates for \mathbf{P}_n and $\mathbf{P}_{n,n-1}$ depend on \mathbf{A}_n instead of \mathbf{A} .

B.2 M-step for Neurodynamics

Because the only non-zero element in \mathbf{Q} (see appendix A) is the first entry, the joint log-likelihood can be written as a function of the neurodynamic parameters as follows

$$\begin{aligned} L_J &= -\frac{1}{2} \sum_{n=2}^N \left(\mathbf{x}_n(1) - \mathbf{x}_{n-1} \mathbf{F}_n^T \tilde{\mathbf{b}} - \mathbf{v}_n^T \mathbf{d} \right)^2 \\ &= -\frac{1}{2} \sum_{n=2}^N \left(\mathbf{x}_n(1) - [\mathbf{x}_{n-1}^T \mathbf{F}_n^T, \mathbf{v}_n^T] \boldsymbol{\theta} \right)^2 \end{aligned} \quad (23)$$

where $\mathbf{x}_n(1)$ is the first entry in \mathbf{x}_n and

$$\boldsymbol{\theta} = \begin{bmatrix} \tilde{\mathbf{b}} \\ \mathbf{d} \end{bmatrix} \quad (24)$$

Taking expectations and derivatives leads to the update

$$\hat{\boldsymbol{\theta}} = \left(\sum_n \tilde{\mathbf{S}}_n \right)^{-1} \left(\sum_n \tilde{\mathbf{v}}_n \right) \quad (25)$$

where

$$\tilde{\mathbf{S}}_n = \begin{bmatrix} \mathbf{F}_n \mathbf{P}_{n-1} \mathbf{F}_n^T & \mathbf{F}_n \mathbf{m}_{n-1} \mathbf{v}_n^T \\ \mathbf{v}_n \mathbf{m}_{n-1}^T \mathbf{F}_n^T & \mathbf{v}_n \mathbf{v}_n^T \end{bmatrix} \quad (26)$$

and

$$\tilde{\mathbf{v}}_n = [\mathbf{P}_{n,n-1}(1) \mathbf{F}_n^T, \mathbf{m}_n(1) \mathbf{v}_n^T] \quad (27)$$

The quantities \mathbf{m}_n , \mathbf{m}_{n-1} , \mathbf{P}_{n-1} and $\mathbf{P}_{n,n-1}$ are computed after Kalman smoothing, \mathbf{v}_n are the driving inputs and \mathbf{F}_n is a matrix derived from the modulatory inputs defined in equation 9.

B.3 M-step for Hemodynamics

The output kernel coefficients can be updated using

$$\hat{\boldsymbol{\beta}} = \left(\sum_n \tilde{\mathbf{P}}_n(1) \right)^{-1} \left(\sum_n \boldsymbol{\Phi} \mathbf{m}_n y_n \right) \quad (28)$$

where

$$\tilde{\mathbf{P}}_n = \boldsymbol{\Phi} \mathbf{P}_n \boldsymbol{\Phi}^T \quad (29)$$

The quantities \mathbf{m}_n and \mathbf{P}_n are computed after Kalman smoothing, $\boldsymbol{\Phi}$ is the matrix of hemodynamic basis functions and y_n is the observed fMRI time series.

C Likelihood

The likelihood is given by

$$\begin{aligned} L &= \log p(\mathbf{y}_1^N) \\ &= \log p(y_1) + \sum_{n=2}^N \log p(y_n | \mathbf{y}_1^{n-1}) \end{aligned} \tag{30}$$

where

$$p(y_n | \mathbf{y}_1^{n-1}) = N(\mathbf{c}^T \mathbf{x}_n^{n-1}, \mathbf{c}^T \boldsymbol{\Sigma}_n^{n-1} \mathbf{c} + \sigma_e^2) \tag{31}$$

The quantities \mathbf{x}_n^{n-1} and $\boldsymbol{\Sigma}_n^{n-1}$ are obtained from Kalman filtering (see Appendix B.1).

Figure captions

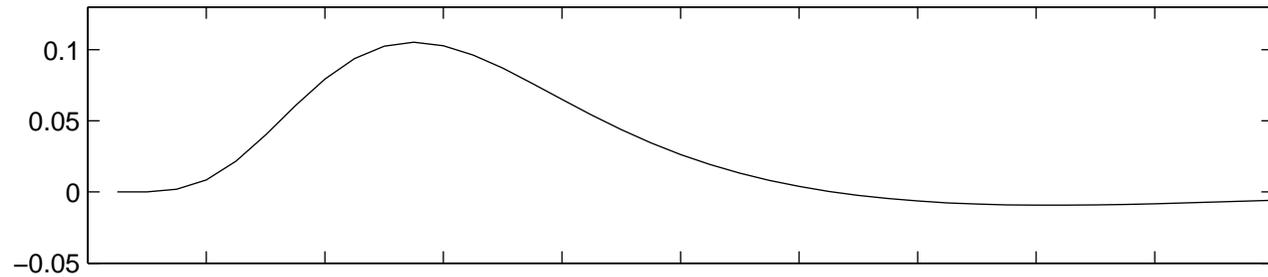
- 1. Hemodynamic Basis Functions** Top Panel: Canonical function, Middle Panel: Derivative of Canonical with respect to time, Bottom Panel: Derivative of Canonical with respect to dispersion. Linear combinations of these functions have been found empirically to account for most subject-to-subject and voxel-to-voxel variations in hemodynamic responses [19, 20].
- 2. Time series from exemplar BDS model** Driving input v_n (top panel) causes an increase in neuronal activity s_n (third panel). This activity decays with a time constant that changes according to a modulatory input u_n (second panel). In this instance the modulatory coefficient, b , is negative. Driving inputs typically correspond to the presentation of experimental stimuli and modulatory inputs typically correspond to instructional or attentional set. Neuronal activity then gives rise to the fMRI time series y_n (bottom panel) by convolving neuronal activity with a set of hemodynamic kernels (shown in Figure 1).
- 3. Low Neuronal Noise Data** Top Panel: Driving input v_n , Middle Panel: Neuronal activity s_n , Bottom Panel: fMRI data y_n .
- 4. Deconvolving Low Neuronal Noise Data** (a) Wiener Deconvolution (solid line), (b) BDS Kalman smoothing deconvolution (solid line) using estimated parameters $a = 0.72$ and $d = 0.88$. The simulated (true) neuronal activity is shown as a dotted line in both plots.
- 5. Deconvolving High Neuronal Noise Data** Top (First) Panel: fMRI data, Second Panel: Wiener Deconvolution (thick line), Third Panel: BDS Kalman Filtering (thick line), Bottom (fourth) Panel: BDS Kalman Smoothing (thick line). In the bottom 3 panels the original simulated neuronal activity is shown as a thin line. Wiener estimation recovers the intrinsic dynamics but misses the evoked responses (the ‘spikes’), whereas BDS Kalman filtering recovers the evoked responses but misses the intrinsic dynamics (slow fluctuations). Deconvolution using BDS Kalman smoothing recovers both.
- 6. Word fMRI inputs** Driving input v_n (top panel) and modulatory inputs u_n (lower panels). The delta functions in the top panel indicate presentation of words in the subject’s headphones and the modulatory

inputs in the lower panels indicate epochs with different presentation rates.

7. **GLM fit to Word fMRI data** fMRI time series from primary auditory cortex (thin line) and GLM model fit (thick line).
8. **BDS fit to Word fMRI data** fMRI time series from primary auditory cortex (thin line) and BDS model fit (thick line).
9. **BDS deconvolution of Word fMRI data** Estimation of neuronal activity using BDS Kalman smoothing. Event-related responses can be seen as spikes superimposed on intrinsic dynamics.

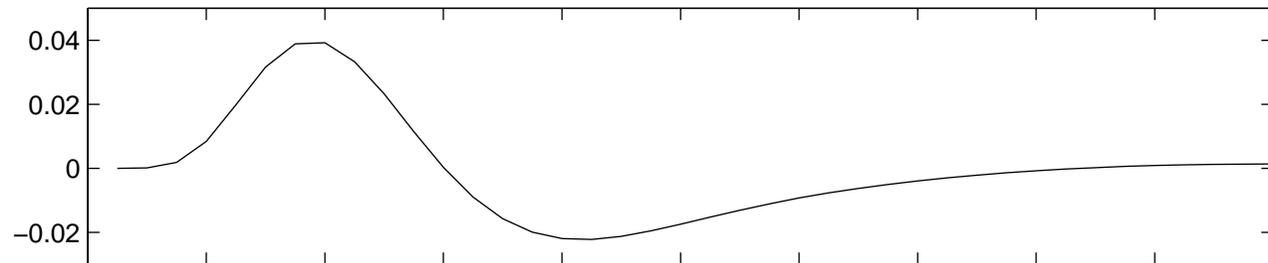
Canonical,

ϕ_1



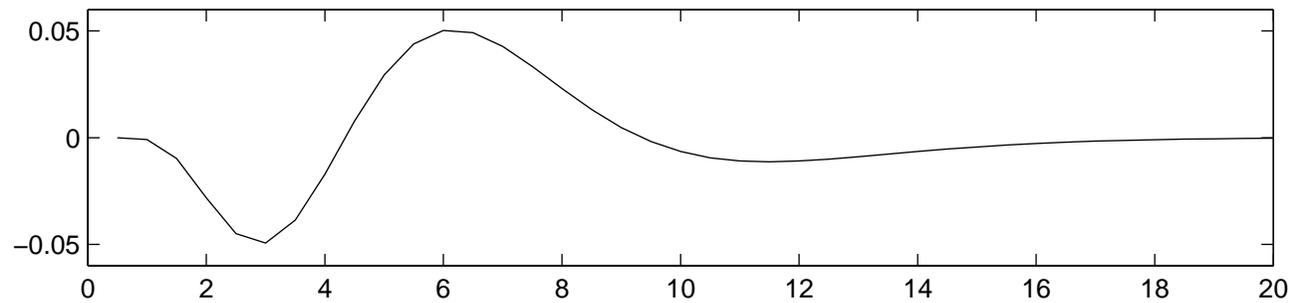
Temporal
Derivative,

ϕ_2

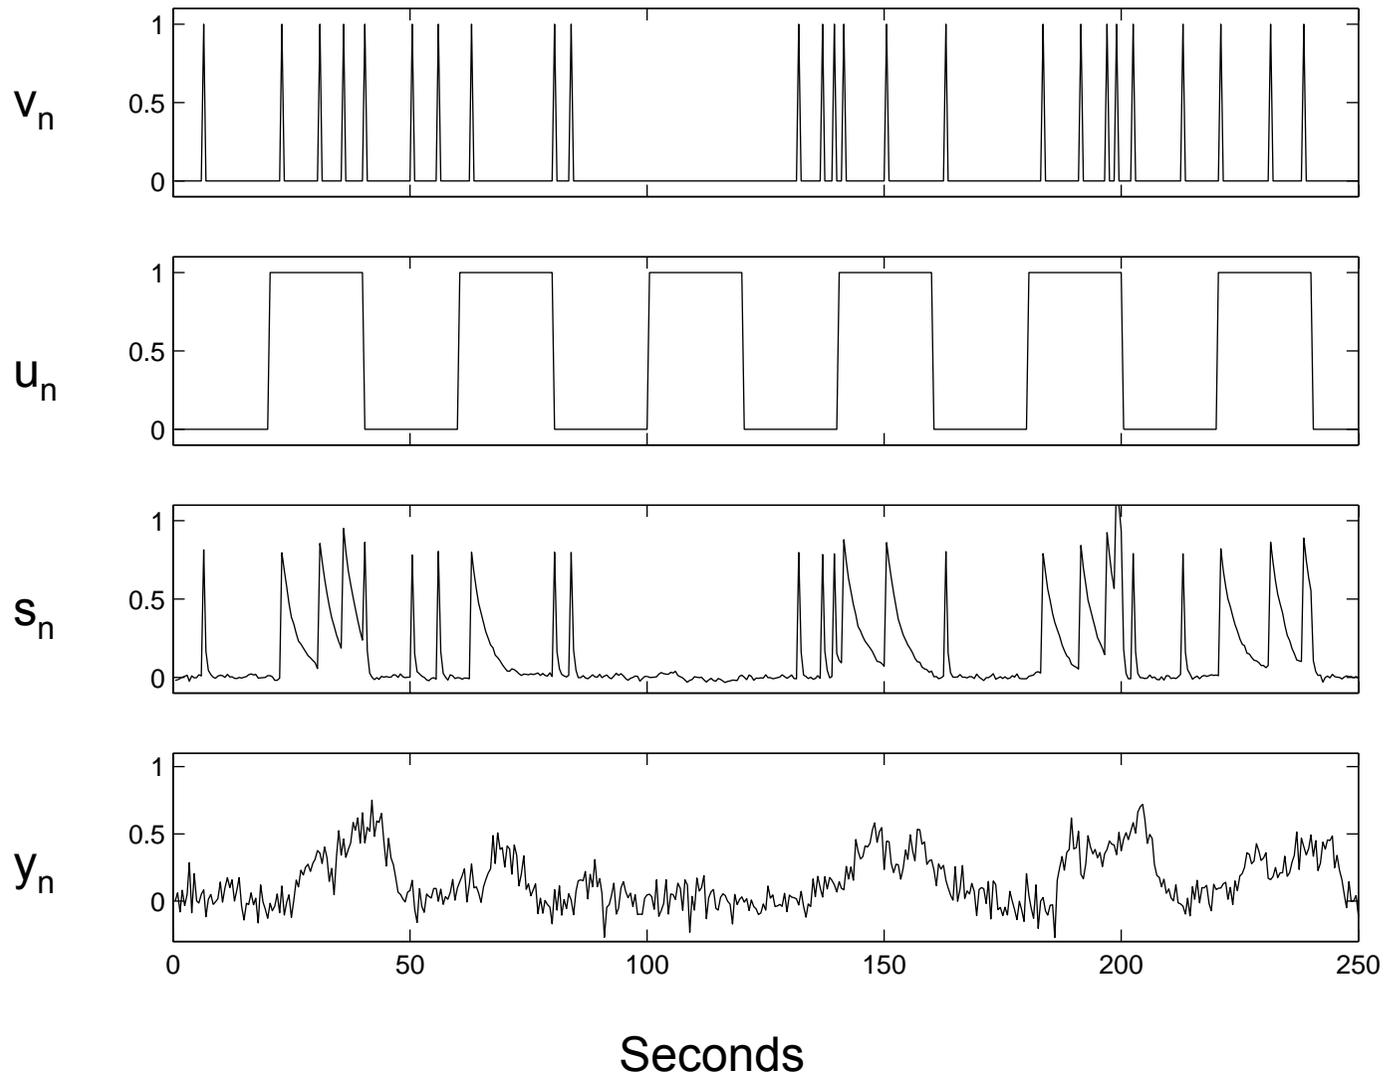


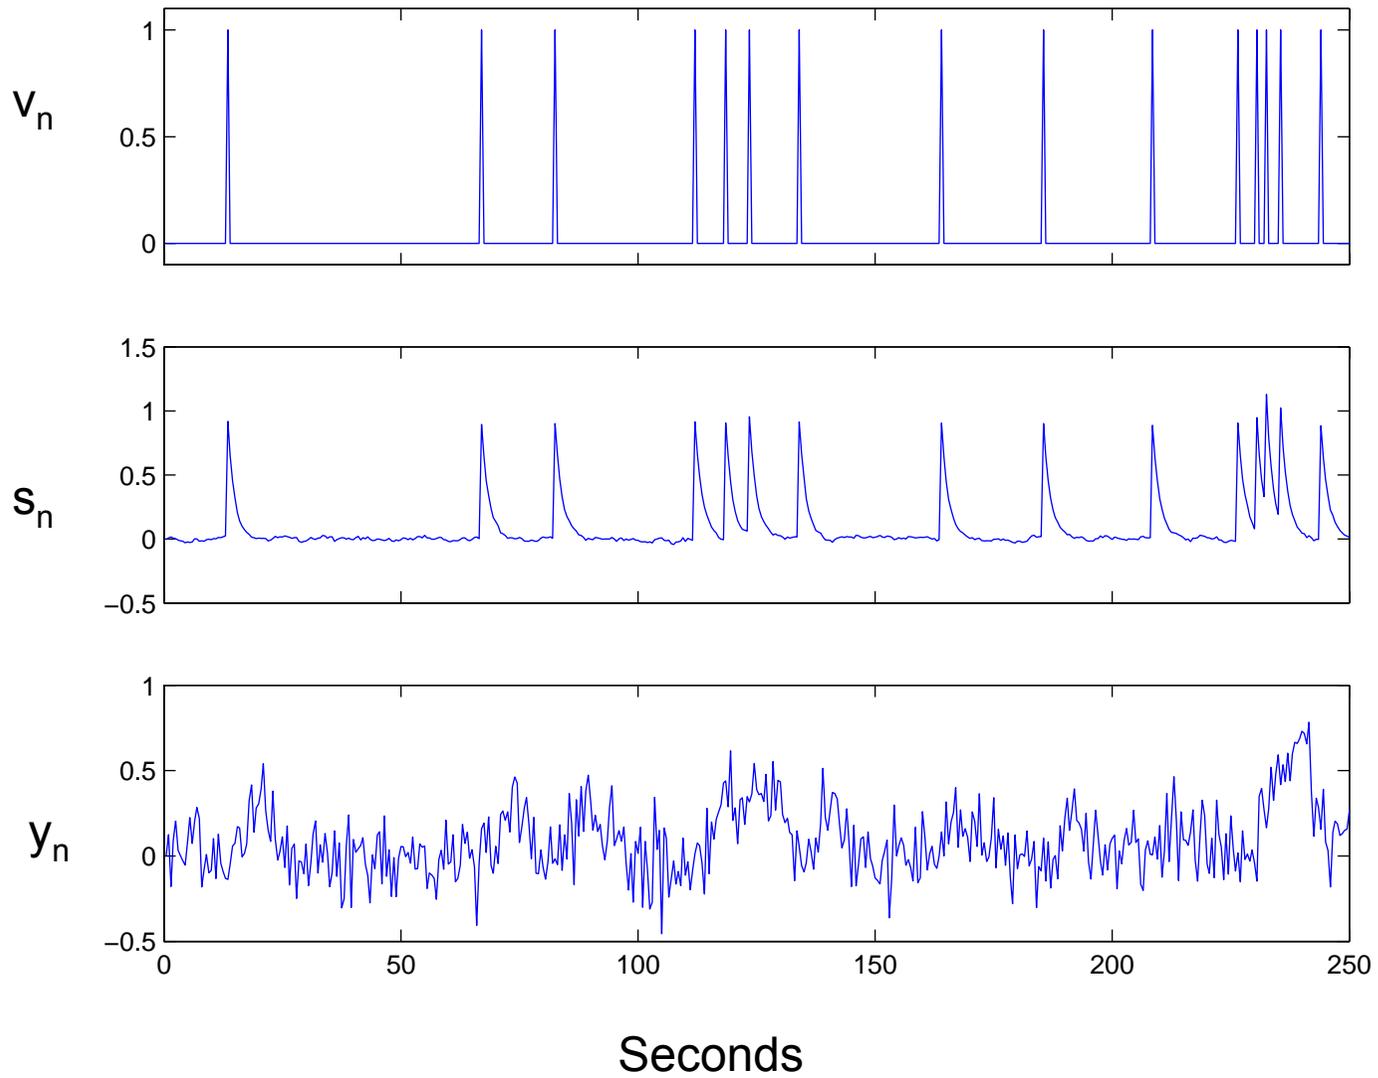
Dispersion
Derivative,

ϕ_3

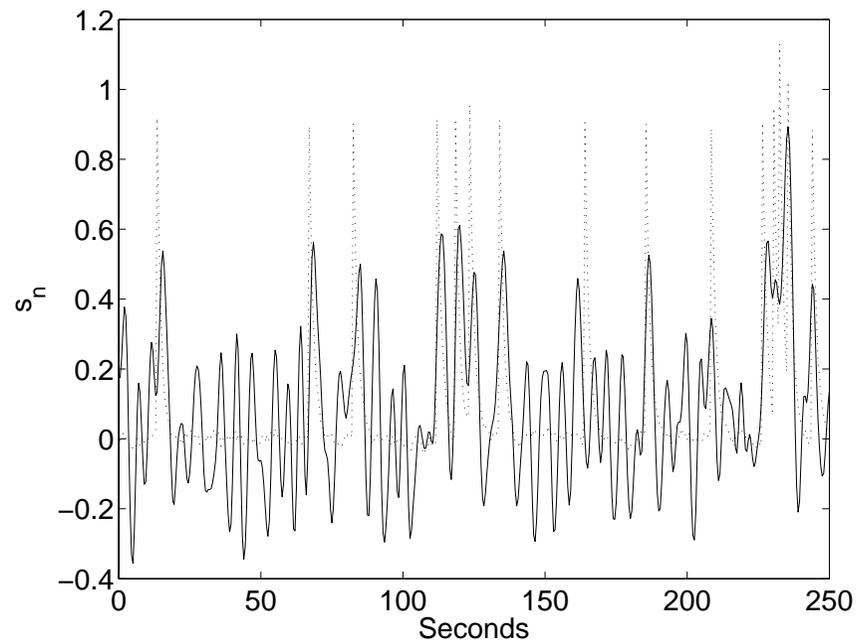


Seconds

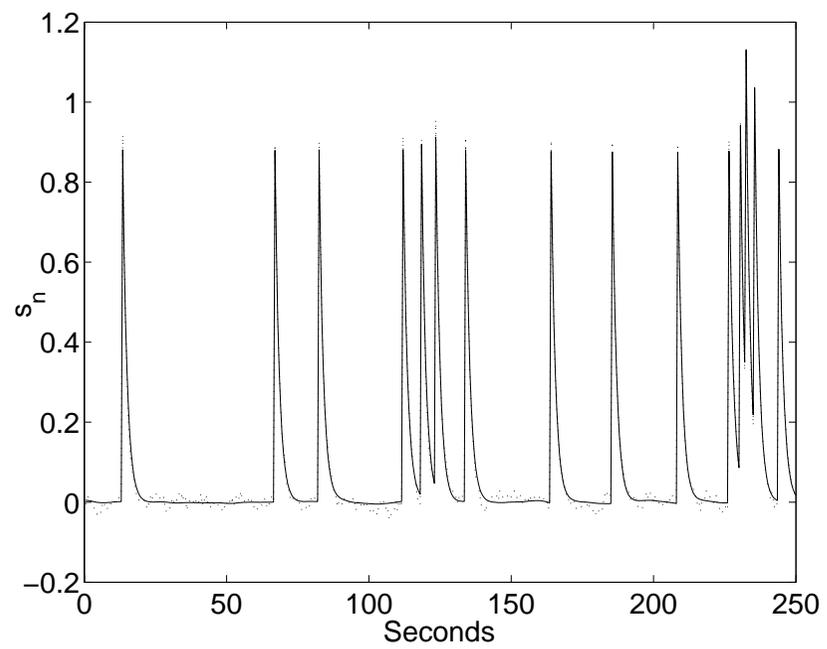




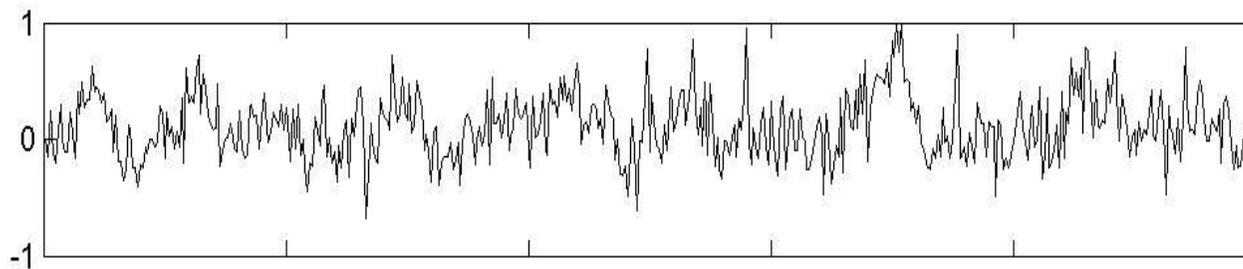
(a)



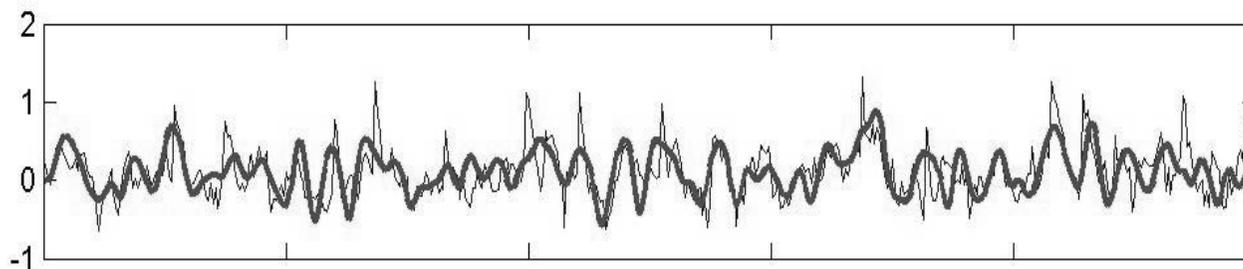
(b)



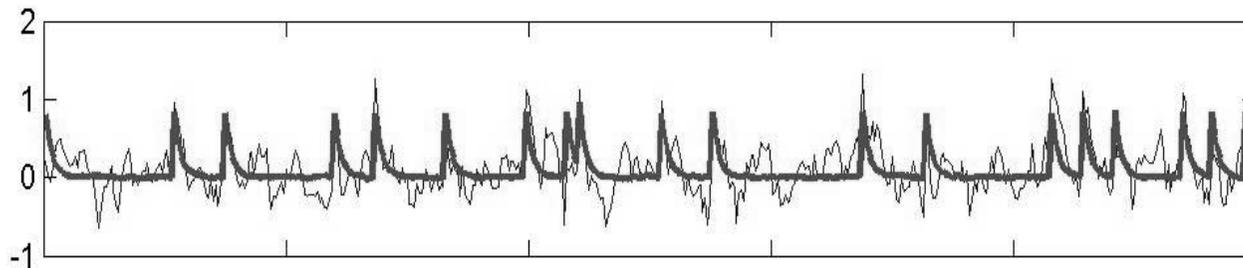
fMRI



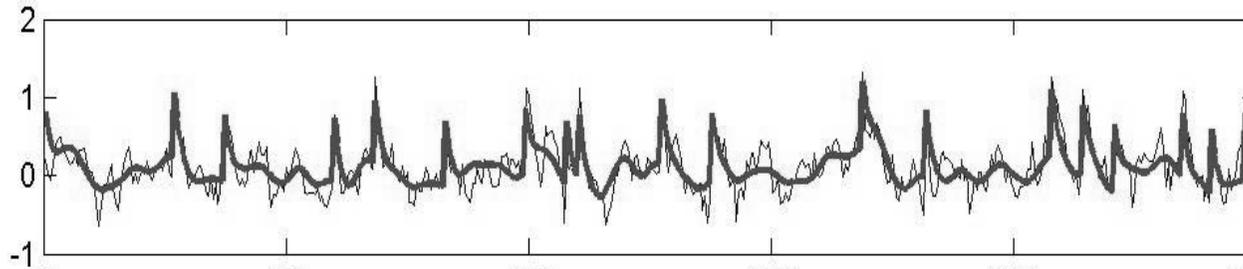
Wiener



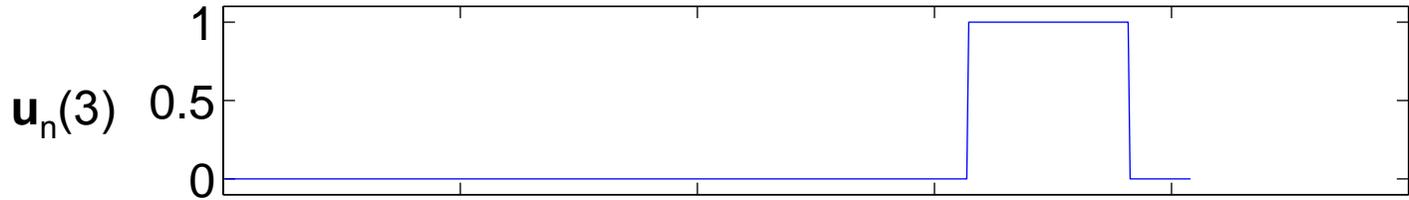
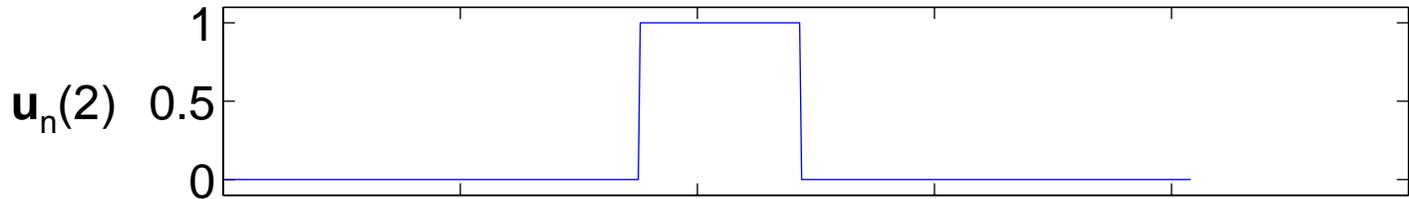
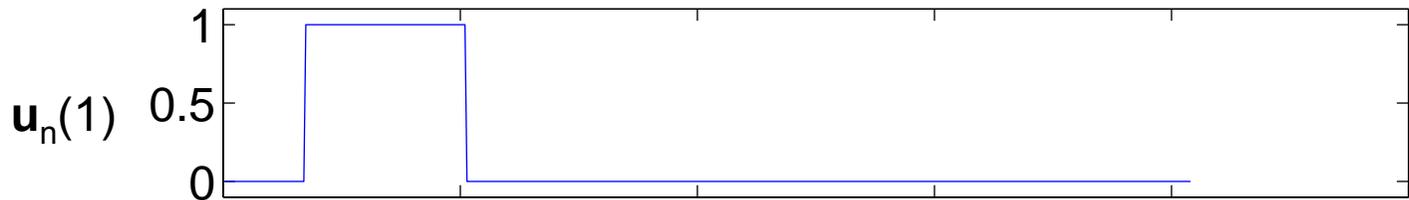
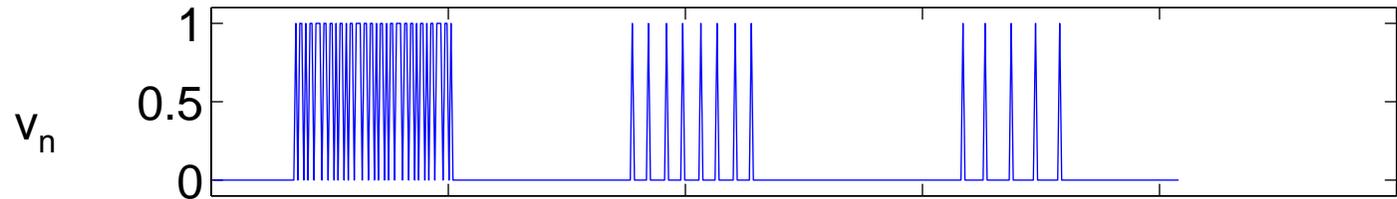
BDS KF



BDS KS



Seconds



0 50 100 150 200 250

Seconds

