

SEMI-PARAMETRIC DYNAMIC TIME SERIES MODELLING WITH APPLICATIONS TO DETECTING NEURAL DYNAMICS

BY FABIO RIGAT AND JIM Q. SMITH

University of Warwick

This paper illustrates novel methods for nonstationary time series modeling along with their applications to selected problems in neuroscience. These methods are semi-parametric in that inferences are derived by combining sequential Bayesian updating with a non-parametric change-point test. As a test statistic, we propose a Kullback–Leibler (KL) divergence between posterior distributions arising from different sets of data. A closed form expression of this statistic is derived for exponential family models, whereas standard Markov chain Monte Carlo output is used to approximate its value and its critical region for more general models. The behavior of one-step ahead predictive distributions under our semi-parametric framework is described analytically for a dynamic linear time series model. Conditions under which our approach reduces to fully parametric state-space modeling are also illustrated. We apply our methods to estimating the functional dynamics of a wide range of neural data, including multi-channel electroencephalogram recordings, longitudinal behavioral experiments and in-vivo multiple spike trains recordings. The estimated dynamics are related to the presentation of visual stimuli, to the evaluation of a learning performance and to changes in the functional connections between neurons over a sequence of experiments.

Introduction. Stochastic modeling of dynamic processes is often implemented via models having time-dependent parameters [Hamilton (1994), West and Harrison (1997), Frühwirth-Schnatter (2006)]. For instance, the coefficients of state-space (SS) and hidden Markov (HM) time series models [Kalman (1960), West, Harrison and Migon (1985), West and Harrison (1997), Cappe, Moulines and Ryden (2005)] follow smooth Markovian processes defined either on their own past or on past values of other latent variables, whereas those of change-point (CP) models [Muller (1992), Stephens (1994), Loader (1996), Mira and Petrone (1996), Bélisle et al. (1998), Fearnhead and Liu (2007)] describe pure jump processes. When these dynamics are specified appropriately, these time series models can effectively capture nonstationarities induced by switches among different dependence regimes [Hamilton (1990), Shumway and Stoffer (1991), Robert, Celeux and Diebolt (1993), Albert and Chib (1993), McCulloch and Tsay (1994), Kim (1994), Ghahramani and Hinton (2000), Frühwirth-Schnatter (2001)], by smooth changes of the model parameters through time [Harrison and Stevens (1976), West

Received October 2008; revised July 2009.

Key words and phrases. Dynamic time series modeling, change-point testing, Bayesian statistics, statistics for neural data.

and Harrison (1986)] or by the occurrence of abrupt changes in the data dependence structure [Page (1955), Smith (1975), Carlin, Gelfand and Smith (1992), Ferger (1995), Chib (1998)].

This paper illustrates theory and applications of a novel sequential method for estimating semi-parametrically the coefficients of time series models having time-dependent parameters. Our approach is motivated by applications where little is known about the factors driving the data dynamics. Here we focus on selected problems in neuroscience where the data exhibit periods of smooth change interlaced with occasional large jumps. We model this type of data by combining sequential Bayesian updating with a nonparametric change-point test. Sequential change-point testing is in fact a well established field which can be traced back at least to the seminal works of Page (1954), Kemp (1957), Barnard (1959) and Page (1961) in statistical process control. We propose testing for significant changes of a model's parameters using a novel Kullback–Leibler (KL) divergence [Kullback and Leibler (1951), Kullback (1997)] between their one-step ahead predictive distributions. The null distribution of this KL statistic reflects the concentration of the joint posterior density when new data are generated using the assumed model likelihood with parameter values drawn from their current posterior distribution. The semi-parametric nature of our method stands in the fact that the value of this test statistic does not depend on the model's parameters, which are integrated out in the calculation of the KL divergence.

With respect to the SS and HM families, our approach does not describe the parameters' dynamics using auxiliary regression equations depending on known predictors. With respect to CP models, we do not assume that parameter values between successive change points are constant. Instead, we induce a time-dependent parameter process by adopting different updating strategies depending on whether the KL statistic lies within its critical region or not. In the former case, the parameters' joint distribution is updated via Bayes' theorem. In the latter case, updating is carried out by matching the first two marginal moments of the current joint posterior probability density to the prior for the next time point. This second strategy, which does not carry the full information content of a posterior distribution to the future, substantiates the notion that a change point in the parameter values has been detected.

With respect to SS, HM and CP models, the advantages of our approach are twofold. First, in SS and HM models inferences and predictions are sensitive to the form of the state evolution equations [Frühwirth-Shnatter (1995), Bengtsson and Cavanaugh (2006)]. Therefore, an exploratory semi-parametric approach is a natural choice for a first analysis of time series data when a specific parametrization of the likelihood function is chosen but no reliable information about the evolution of its parameters is available [Robinson (1983), Härdle, Lütkepohl and Chen (1997)]. This is typically the case for many biological systems, where dynamic responses to novel experimental conditions are difficult to anticipate. Second, the joint distribution of the model's parameters is updated also between successive change-points,

allowing for a reduction of uncertainty and for smooth changes of the parameter estimates over time.

From a computational perspective, our approach is motivated by observing that fully Bayesian sequential inference for a model's time-dependent parameters and for a latent multiple change-point process is impractical unless marginal likelihoods can be calculated explicitly. Otherwise, the Bayes factors measuring the strength of evidence in the data about the occurrence of change-points can only be approximated numerically [Han and Carlin (2001)]. Current methods for calculating these approximations require knowledge of normalizing constants which may be hard to obtain and they also require estimating the exact value of a posterior probability density at one point, which is ideally chosen as one of the posterior modes [Newton and Raftery (1994), Gelfand and Dey (1994), Chib (1995), Chib (1998), Frühwirth-Shnatter (2006)]. Our approach represents a practical alternative to these methods in that point estimates of a latent change-point process are derived without using marginal likelihoods.

Section 1 of this paper includes its methodological developments. A general time series framework is introduced and the KL test is illustrated. A closed form expression of the KL statistic for exponential family models is derived and examples are presented. Markov chain Monte Carlo (MCMC) simulation [Gelfand and Smith (1990), Tierney (1994)] is used to approximate the exact critical region of the KL statistic under the null hypothesis. This approximation is chosen as it only requires the assumed data sampling distribution and the standard MCMC output. We present a simulation study showing that the power of the KL change-point test is unaffected in practice by adopting these MCMC approximations when using a conjugate Bernoulli model. A sequential algorithm summarizing the computational steps involved in the implementation of our method is presented. The behavior of location and spread of the one-step ahead predictive distributions arising from our method is described analytically for a conjugate Gaussian linear dynamic model. Conditions are given so that our semi-parametric approach reduces to fully parametric state-space dynamic time series modeling. In Sections 2, 3 and 4 our method is applied to estimating three different types of neural dynamics. First, we analyze a multivariate time series of electroencephalogram (EEG) recordings [Delorme et al. (2002), Makeig et al. (2002)] to reconstruct the time-varying functional relationships among different brain areas. Second, we estimate semi-parametrically a learning curve using a univariate binary time series arising from a longitudinal behavioral experiment [Smith et al. (2004)]. Finally, our method is applied to estimating the functional dynamics of networks of neurons using in-vivo experimental multiple spike trains recordings [Buzsáki (2004)].

1. Sequential time series modeling and Kullback–Leibler change-point testing. Let $\{Y_i\}_{i=1}^N$ represent a sequence of K -dimensional time series $Y_i \in \mathcal{Y}$ of random variables $Y_{i,k,t}$ measured at the time points $t = t_{i,1} < t_{i,2} < \dots < t_{i,n_i}$ with $t_{i,n_i} < t_{i+1,1}$ and $k = 1, \dots, K$. The distinction between the N time series

is relevant when we allow for the occurrence of time gaps between them. This situation arises, for instance, when N consecutive trials are run sequentially interposed by resting periods. When $n_i = 1$ for all values of i , we effectively have a single K -dimensional time series of length N measured at the time points $t_{i,1}$. In this paper the time series data $Y_i = y_i$ are assumed to be generated by a finite-dimensional model $P(y_i | \theta_{i-1}, y^{0:(i-1)})$, such as a vector auto-regressive (VAR) model with shared coefficients θ_{i-1} within each of the N periods. The probability density $f(\theta_{i-1} | y^{0:(i-1)})$ here represents a distribution of the model coefficients given the initial conditions y^0 and all past observations up to and including period $i - 1$. Note that, although we allow the parameter values to vary in time, neither the functional form of the likelihood nor the interpretation of its coefficients change over time.

Within this framework, dynamic modeling consists of specifying a transfer map, taking as arguments the posterior density $f(\theta_{i-1} | y^{0:(i-1)})$, the time series data y_i and possibly other fixed hyper-parameters α and returning the density $f(\theta_i | y^{0:i})$ for $i = 1, \dots, N$. Various characterizations of analogous maps are given in Smith (1990, 1992). For instance, in standard state-space models, this transfer map is defined by indexing the prior distribution for θ_i using the coefficients θ_{i-1} and a set of hyper-parameters. In Markov switching and finite mixture time series models, this transfer map is again derived by parametric modeling of the joint density of the coefficients θ_i and θ_{i-1} conditional on the location of a sequence of change-points [Frühwirth-Schnatter (2006)]. Here we provide an overview of a transfer map which integrates sequential Bayesian inference and change-point testing, leaving to Section 1.1 the detailed description of an appropriate test statistic.

Let $\hat{\theta}_{i-1}$ be a current point estimate of the model's parameters at time $t_{i,1}$. When $i = 1$ these are prior summaries, whereas for $i > 1$ these estimates incorporate evidence from past data as described below. If the data y_i are generated under significantly different parameter values with respect to period $i - 1$, we say a change-point has occurred. In this case we define a transfer prior

$$(1.1) \quad \theta_i \sim h(\theta_i | \hat{\theta}_{i-1}),$$

taking as arguments the current parameter estimates and returning a prior density $h(\cdot)$ for the model's coefficients θ_i . Among the many possible formulations of this prior, we let its hyperparameters be functions of the first two marginal moments of the current posterior density. Similar forms of prior moment matching have been used for dynamic point process modeling by Gamerman (1992) and for multi-process dynamic linear models by West and Harrison (1997). This partial information transfer from the posterior distribution ensures that the moment-matched priors allocate most of their mass around the current marginal posterior means, but upon detecting a change-point, the dependencies among different models' parameters, skewness, kurtosis and the other higher-order moments are all reset to their values prior to observing any data. Equation (1.1) represents a partially specified

state evolution density where neither the exact form of the prior nor the time of occurrence of the change-points are given a priori. Specific choices for the prior density depend on the structure of the time series model being entertained and on the interpretation of its parameters.

When no change-points are detected prior to observing the data $Y_i = y_i$, under (1.1) the joint posterior density of the model's parameters is

$$(1.2) \quad f(\theta_i | y^{0:i}, \alpha) \propto \begin{cases} f(\theta_i | y^{0:(i-1)}) P(y_i | \theta_i, y^{0:(i-1)}), & \text{if } y_i \in \Psi_i(\alpha), \\ h(\theta_i | \hat{\theta}_{i-1}) P(y_i | \theta_i, y^{0:(i-1)}), & \text{if } y_i \notin \Psi_i(\alpha). \end{cases}$$

Here $\Psi_1(\alpha) = \mathcal{Y}$ and for $i = 2, \dots, N$ the sets $\Psi_i(\alpha) \subseteq \mathcal{Y}$ include the time series Y_i which are inconsistent with their observed past $y^{0:(i-1)}$ under the current estimates of the model parameters and the hyper-parameters α .

Implementation of (1.2) presents two related challenges. First, it is essential to formulate the rejection sets $(\Psi_2(\alpha), \dots, \Psi_N(\alpha))$ in terms of a low-dimensional statistic of the data and of the hyper-parameters α . Second, it must be possible to derive the distribution of such a statistic over the sample space so as to provide at least a sequential approximation of the rejection sets for any value of α . A natural way to overcome these challenges is to view $(\Psi_2(\alpha), \dots, \Psi_N(\alpha))$ as the α -level critical regions of a sequential change-point test based on an appropriate statistic. The transfer map is thus completely specified by the prior (1.1) together with a choice of this test statistic.

1.1. *A Kullback–Leibler change-point statistic.* The Kullback–Leibler divergence [Kullback and Leibler (1951)] is a well-known information-theoretic criterion with many applications in statistics, such as density estimation [Hall (1987), Hastie (1987)], model selection [Akaike (1978), Akaike (1981), Carota, Parmigiani and Polson (1996), Goutis and Robert (1998)], experimental design [Lindley (1956), Stone (1959)] and the construction of uninformative priors [Bernardo (1979)]. Its geometric properties have been thoroughly explored by Critchley, Marriott and Salmon (1994). The change-point statistic proposed in this work has a complementary function to the KL divergence when used to support model selection. Instead of testing which of two competing model structures best predicts one given set of data, here we construct a statistic detecting whether the same parameter values could have likely generated two sets of data given a common model structure.

As change-point test statistic we adopt a Kullback–Leibler divergence

$$(1.3) \quad \begin{aligned} \text{KL}(y^{0:(i+1)}) &= \int_{\Theta} \log\left(\frac{f(\theta_i | y^{0:i})}{f(\theta_i | y^{0:(i+1)})}\right) f(\theta_i | y^{0:i}) d\theta_i \\ &= \log(E_{\theta_i | y^{0:i}}(P(y_{i+1} | \theta_i, y^{0:i}))) \\ &\quad - E_{\theta_i | y^{0:i}}(\log(P(y_{i+1} | \theta_i, y^{0:i}))), \end{aligned}$$

where the expectations in (1.3) are taken with respect to the posterior density $f(\theta_i|y^{0:i})$. The right-hand side of (1.3) is finite when the likelihood function is bounded away from zero and infinity for all values of the model's parameters and when their posterior density is proper. In this case (1.3) is a nonnegative convex function measuring the discrepancy between the posterior densities $f(\theta_i|y^{0:i})$ and $f(\theta_i|y^{0:(i+1)})$ over their common support Θ . Prior to observing the data $Y_{i+1} = y_{i+1}$, (1.3) is a random variable in which distribution under the null hypothesis depends on that of the future data Y_{i+1} via the likelihood $P(Y_{i+1}|\theta_i, y^{0:i})$. The following sections focus on the interpretation and on the computation of (1.3).

1.1.1. *Interpretation of the KL statistic and of the change-points.* The scalar hyper-parameter α of the joint posterior (1.2) has the interpretation of the type-1 error probability for the change-point test using the statistic (1.3). The rejection sets can be written explicitly as intervals $\Psi_i(\alpha) = (l_{i,\alpha}, u_{i,\alpha})$ representing the α -level highest probability interval for the random variable (1.3) under the hypothesis of no change over period i .

When α is low and (1.3) lies below the value $l_{i,\alpha}$, the likelihood of the observed data is almost a constant in the parameters θ_i over the range of their current posterior density. In other terms, the parameter values maximizing the likelihood of the observed y_{i+1} conditionally on the past data $y^{0:i}$ are given almost zero probability by the posterior distribution under the hypothesis of no change. If the change-point statistic lies above $u_{i,\alpha}$, the parameter values maximizing the likelihood of the data y_{i+1} are associated to substantial values of the current joint posterior density, but they are far from its global maximum. In this case the joint posterior density of all data $y^{0:(i+1)}$ under the hypothesis of no change is bimodal, indicating that the latest batch of data y_{i+1} are not adequately explained by the current parameter values. In both cases the value of the statistic (1.3) indicates that, in light of the data $y^{0:(i+1)}$, the posterior density of the model's parameters $f(\theta_i|y^{0:(i+1)})$ significantly departs from its assumed form under the hypothesis that sequential Bayesian updating is adequate.

When $\alpha = 0$, no change-point is ever detected, so that the model's parameters are updated sequentially only via Bayes' rule. On the other end, if $\alpha = 1$, a change in the parameter values is systematically detected at every time point. In this second limiting case the method proposed in this work is equivalent to a fully parametric first order Markov state-space model in which state evolution equations have the form (1.1).

1.1.2. *Computation of the change-point statistic.* The test statistic (1.3) is similar in spirit to the cumulative Bayes factors proposed in West (1986) and West and Harrison (1986), with the practical advantage that the computation of marginal likelihoods is not required. However, in general, neither the value of (1.3) nor the rejection sets $(\Psi_2(\alpha), \dots, \Psi_N(\alpha))$ may be available in closed form, so that numerical approximations may be required. In these cases, at each time period these approximations can be calculated without incurring in additional computational cost

using a sequence of parameter values $\{\theta_i^m\}_{m=1}^M$ generated using a Markov chain Monte Carlo algorithm [Gelfand and Smith (1990), Smith and Roberts (1993), Tierney (1994)] having as its target the current posterior probability density. Using this technique, the value of (1.3) is approximated by the average

$$(1.4) \quad \text{KL}(y^{0:(i+1)}) \approx \log\left(\frac{\sum_{m=1}^M p_{i+1}^m}{M}\right) - \frac{\sum_{m=1}^M \log(p_{i+1}^m)}{M},$$

where $p_{i+1}^m = P(y_{i+1}|\theta_i^m, y^{0:i})$ is the likelihood of the data y_{i+1} given the parameter values θ_i^m and the past data $y^{0:i}$. Using (1.4), the null distribution of (1.3) can be approximated as follows:

- (i) for each draw θ_i^m generate a pseudo-realization y_{i+1}^m using the joint sampling distribution $P(Y_{i+1}|\theta_i^m, y^{0:i})$;
- (ii) compute the statistic $\text{KL}(y_m^{0:(i+1)})$, where $y_m^{0:(i+1)} = (y_0, \dots, y_i, y_{i+1}^m)$, using its Monte Carlo approximation (1.4).

The empirical distribution of the sequence $\{\text{KL}(y_m^{0:(i+1)})\}_{m=1}^M$ approximates that of the KL statistic (1.3) under the hypothesis of no change. Therefore, the empirical $(\frac{\alpha}{2}, 1 - \frac{\alpha}{2})$ th percentiles of the sequence $\{\text{KL}(y_m^{0:(i+1)})\}_{m=1}^M$ approximate the rejection sets $\Psi_i(\alpha) = (l_{i,\alpha}, u_{i,\alpha})$ for any given value of α .

1.2. Change-point test power and sample size. When the time series $\{Y_i\}_{i=1}^N$ have substantially different lengths, the power of the change-point test based on the KL statistic is theoretically unchanged. For any value of α , this invariance is ensured by the behavior of the posterior distribution at the denominator of (1.3). When the data $Y_{i+1} = y_{i+1}$ carries a large amount of information about the coefficients of model $P(Y_{i+1}|\theta_i, y^{0:i})$, their joint posterior distribution under the hypothesis of no change concentrates by a corresponding large amount, so that the distribution of the KL divergence concentrates over large values. If Y_{i+1} is not expected to carry much additional information about the model parameters, for instance, due to its small sample size n_{i+1} , the null distribution of the KL discrepancy is concentrated over small nonnegative values. This mechanism represents an automatic adaptation of the critical region $\Psi_{i+1}(\alpha)$ of the KL test, ensuring that its power does not vary with the sample size of the data sequentially accrued over time.

Although this property is sufficiently clear in theory, it is an open question whether the power of the test is significantly affected when our method is implemented using the MCMC approximations outlined above. Here we briefly investigate this issue by simulation using a conjugate Bernoulli model. One hundred thousand simulations were run. For each simulation, two sample sizes n_1 and n_2 were independently generated as independent draws from a discrete uniform distribution on the integers $(1, \dots, M)$ with $M = 100$. A success probability π was also

independently generated for each simulation using a uniform distribution on the interval $(0, 1)$. Conditionally on (n_1, n_2, π) , two independent samples of Bernoulli random variables were generated, $Y_1 \sim \text{Ber}(\pi, n_1)$ and $Y_2 \sim \text{Ber}(\pi, n_2)$. For each simulation, a sample of size 5000 was generated from the conjugate posterior $\text{Beta}(1 + \sum_{j=1}^{n_1} Y_{1,j}, 1 + n_1 - \sum_{j=1}^{n_1} Y_{1,j})$ to compute the Monte Carlo approximation of the KL statistic and of the end-points of its 95% probability interval under the hypothesis of no change. For this simulation study the type-1 error probability of the test was fixed to $\alpha = 0.2$. Under this sampling scheme, with $n_i^* = n_i - 1$ for $i = 1, 2$, the random variables $\frac{n_1^*}{M-1}$ and $\frac{n_2^*}{M-1}$ are independent and approximately uniform on $(0, 1)$, so that the distribution of the statistic $Z = \log(\frac{n_1^*}{n_2^*})$ is approximately standard double-exponential. If the power of the KL change-point test is in practice not affected by the values of (n_1, n_2, π) , the distribution of Z for the group of simulations where the hypothesis of no change is accepted should be standard double exponential. Figure 1 represents with a solid line the empirical cumulative distribution function (CDF) of Z for the 79,743 simulations where a significant change was not detected. The two dashed lines in the same figure represent the point-wise 99% probability intervals for the CDF of a standard double exponential

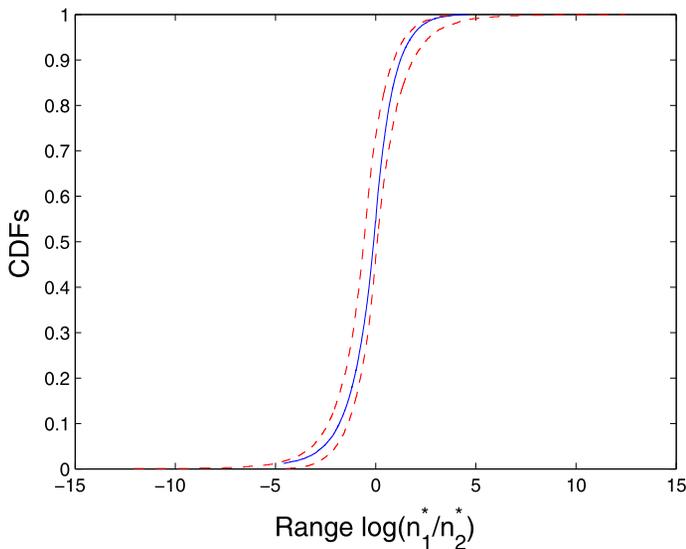


FIG. 1. The solid line represents the empirical cumulative distribution function (CDF) of the random variable Z for the 79,743 simulations where the hypothesis of no change was accepted. The dashed lines represent the approximate end-points of the point-wise 99% probability intervals for the CDF of a standard double exponential random variable. Acceptance of the hypothesis of no change did not cause a significant departure of the distribution of Z from that of a standard double exponential distribution, suggesting that the power of the change-point test is not significantly affected by different sample sizes (n_1, n_2) .

random variable. Since at each point the former CDF always lies within its 99% interval, this simulation study suggests that for the Bernoulli model the power of the change-point test is not significantly affected by different sample sizes (n_1, n_2) .

1.3. *Sequential fitting and change-point testing algorithm.* This section provides a summary of the computational steps involved by the dynamic modeling method illustrated so far. Despite not addressing any model-specific issues such as the explicit form of posterior distributions, we aim at providing here a general blueprint for implementing our method starting from the first sample y_1 :

- (i) Upon observing the data y_1 , derive the posterior density

$$f(\theta_1|y^{0:1}) \propto h(\theta_1|\hat{\theta}_0)P(y_1|\theta_1, y_0),$$

where $\hat{\theta}_0$ represents an estimate of the parameter values as reflected by the initial conditions y_0 .

- (ii) Having observed data y_2 , compute the statistic $KL(y^{0:2})$ and its rejection interval $\Psi_1(\alpha) = (l_{1,\alpha}, u_{1,\alpha})$ as described in Section 1.1.
- (iii) If $l_{1,\alpha} < KL(y^{0:2}) < u_{1,\alpha}$, no change-point is detected. In this case the prior density for θ_2 is the posterior at point (i) and the posterior density for θ_2 derived using Bayes' rule is

$$f(\theta_2|y^{0:2}, \alpha) \propto f(\theta_2|y^{0:1})P(y_2|\theta_2, y^{0:1}).$$

- (iv) Otherwise, match the first two posterior moments $\hat{\theta}_1$ to those of the prior for θ_2 and again apply Bayes' rule, deriving the conditional posterior density

$$f(\theta_2|\hat{\theta}_1, y^{0:2}, \alpha) \propto h(\theta_2|\hat{\theta}_1)P(y_2|\theta_2, y^{0:1}).$$

In case (iv) above, the sequentially estimated change-point process up to and including times $(1, 2)$ reports one change at time 2. Consistently with the interpretation of the KL statistic, the model parameters are updated using all data starting from the last detected change-point, if any. When a change is detected at level $1 - \alpha$, the new parameter values are updated using their conditional posterior distribution under the transfer prior (1.1) and the likelihood of the latest batch of data.

1.4. *Change-point KL statistic for exponential family models.* Several properties of the KL divergence for exponential family models have been explored by McCulloch (1988). Here we show that in this circumstance also the divergence (1.3) has a closed form. In this case the algorithm illustrated in Section 1.2 is simplified, as only the critical intervals $\Psi_i(\alpha) = (l_{i,\alpha}, u_{i,\alpha})$ need being approximated. Without loss of generality, in what follows we assume that no change-point is detected prior to period i . Also, we let Y_i be a 1 dimensional sample of conditionally independent observations with length n_i and joint density [Diaconis and Ylvisaker (1979)]

$$(1.5) \quad P(Y_i|\theta_i) = \prod_{j=1}^{n_i} a(Y_{i,j})e^{Y_{i,j}\theta_i - b(\theta_i)},$$

where θ_i is a scalar canonical parameter. Diaconis and Ylvisaker (1979) show that each element of Y_i has mean and variance

$$E(Y_{i,j}|\theta_i) = \frac{\partial b(\theta_i)}{\partial \theta_i}, \quad V(Y_{i,j}|\theta_i) = \frac{\partial^2 b(\theta_i)}{\partial \theta_i^2}.$$

Using the prior

$$f(\theta_i | n_0, y_0) = c(n_0, S_0) e^{S_0 \theta_i - n_0 b(\theta_i)},$$

where $S_0 = n_0 y_0$ for scalar n_0 and y_0 , the posterior for θ_i given the past data $y^{0:i}$ has conjugate density

$$(1.6) \quad f(\theta_i | n(i), y^{0:i}) = c(n(i), S(i)) e^{n(i)((S(i)/n(i))\theta_i - b(\theta_i))},$$

where $n(i) = \sum_{j=0}^i n_j$, $S(i) = \sum_{j=0}^i n_j \bar{y}_j$ and \bar{y}_j represents the arithmetic mean of sample y_j . Using the results of Gutiérrez-Peña (1997), the posterior mean and variance of θ_i are

$$E(\theta_i | n(i), S(i)) = \frac{\partial H(n(i), S(i))}{\partial S(i)}, \quad V(\theta_i | n(i), S(i)) = \frac{\partial H(n(i), S(i))}{\partial S(i)^2},$$

and the posterior mean and variance of the function $b(\theta_i)$ are

$$E(b(\theta_i) | n(i), S(i)) = \frac{\partial H(n(i), S(i))}{\partial n(i)},$$

$$V(b(\theta_i) | n(i), S(i)) = \frac{\partial H(n(i), S(i))}{\partial n(i)^2},$$

where $H(n(i), S(i)) = -\log(c(n(i), S(i)))$. Using these results, we derive the following explicit form for the KL divergence (1.3):

THEOREM. *When the posterior density for the coefficients θ_i has form (1.6), given the data up to and including y_{i+1} , the Kullback–Leibler statistic (1.3) is*

$$(1.7) \quad \text{KL}(y^{0:(i+1)}) = \log\left(\frac{c(n(i), S(i))}{c(n(i+1), S(i+1))}\right) - S_{i+1} \frac{\partial H(n(i), S(i))}{\partial S(i)} + n_{i+1} \frac{\partial H(n(i), S(i))}{\partial n(i)},$$

where the terms on the right-hand side of (1.7) are defined above.

PROOF. By letting the posterior densities $f(\theta_i | n(i), S(i))$ and $f(\theta_i | n(i+1), S(i+1))$ have form (1.6), the KL (1.3) becomes

$$\text{KL}(y^{0:(i+1)}) = \log\left(\frac{c(n(i), S(i))}{c(n(i+1), S(i+1))}\right) - S_{i+1} E(\theta_i) + n_{i+1} E(b(\theta_i)).$$

For exponential family models, the expectations $E(\theta_i)$ and $E(b(\theta_i))$ with respect to $f(\theta_i | y^{0:i})$ are reported above. By substituting these expressions, equation (1.7) obtains the following. \square

EXAMPLE 1.1. When Y_i is a Gaussian random variable with mean μ_i and precision λ_i , its distribution can be written in the form (1.5) using the two-dimensional statistic

$$Y_i^* = [Y_i, Y_i^2]$$

and the canonical parameter

$$\theta_i = [\theta_{1,i}, \theta_{2,i}] = \left[\lambda_i \mu_i, -\frac{\lambda_i}{2} \right]$$

with

$$a(Y_i^*) = (2\pi)^{-1/2},$$

$$b(\theta_i) = -\frac{1}{2} \log(\theta_{2,i}) - \frac{\theta_{1,i}^2}{\theta_{2,i}}.$$

The conjugate prior for (μ_i, λ_i) is Normal-Gamma $N(\mu_i | \gamma, \lambda_i(2\alpha - 1)) \text{Ga}(\lambda_i | \alpha, \beta)$ with coefficients $\alpha > 0.5, \beta > 0, \gamma \in \mathcal{R}$ and normalizing constant [Bernardo and Smith (2007)]

$$c(n_0, S_0) = \left(\frac{2\pi}{n_0} \right)^{1/2} \frac{S_{2,0}^{S_{1,0}/2} / 2}{\Gamma((n_0 + 1)/2)},$$

where $n_0 = 2\alpha - 1, y_0^* = [y_{1,0}^*, y_{2,0}^*] = [\gamma, \frac{2\beta}{2\alpha-1} + \gamma^2], S_{1,0} = n_0 y_{1,0}^*$ and $S_{2,0} = n_0 y_{2,0}^*$. Upon observing the realization (y_1, \dots, y_i) , the normalizing constant of the corresponding conjugate posterior is

$$c(n(i), S(i)) = \left(\frac{2\pi}{n(i)} \right)^{1/2} \frac{S(2, i)^{S(1, i)/2} / 2}{\Gamma((n(i) + 1)/2)},$$

where $n(i) = n_0 + i, S(1, i) = S_{1,0} + \sum_{j=1}^i y_j$ and $S(2, i) = S_{2,0} + \sum_{j=1}^i y_j^2$. When also y_{i+1} is observed, using (1.7), the KL statistic can be written as

$$\begin{aligned} \text{KL}(y^{0:(i+1)}) &= \log \left(\Gamma \left(\frac{n(i+1) + 1}{2} \right) / \Gamma \left(\frac{n(i) + 1}{2} \right) \right) + \frac{1}{2} \log \left(\frac{n(i+1)}{n(i)} \right) \\ &+ \log \left(\frac{S(2, i)^{S(1, i)/2}}{S(2, i+1)^{S(1, i+1)/2}} \right) - \frac{y_{i+1}}{2} \log \left(\frac{S(2, i)}{2} \right) \\ &- y_{i+1}^2 \frac{S(1, i)}{S(2, i)} + \frac{1}{2n(i)} + \Gamma \left(\frac{n(i) + 1}{2} \right) \frac{\partial \Gamma((n(i) + 1)/2)}{\partial n(i)}. \end{aligned}$$

EXAMPLE 1.2. Let Y_i be a sample of size n_i of conditionally independent Bernoulli random variables with success probabilities $\{\pi_i\}_{i=1}^N$. The canonical representation of the Bernoulli probability mass function obtains, by letting $\theta_i = \log(\frac{\pi_i}{1-\pi_i}), b(\theta_i) = \log(1 + e^{\theta_i})$ and $a(Y_i) = 1$. The conjugate prior for π_i

is $\text{Beta}(S_0, m_0)$, where $m_0 = n_0 - S_0$. Upon observing (y_1, \dots, y_i) , the conjugate posterior is $\text{Beta}(S(i), m(i))$, where $S(i) = \sum_{j=0}^i S_j$, $n(i) = \sum_{j=0}^i n_j$ and $m(i) = n(i) - S(i)$. When also y_{i+1} is observed, the KL statistic (1.7) has form

$$\begin{aligned} \text{KL}(y^{0:(i+1)}) &= \log\left(\frac{\prod_{k=1}^{n(i)} (n(i) + k) \prod_{w=1}^{n(i)-S(i)} (m(i) + w)}{\prod_{j=1}^{S(i)} (S(i) + j)}\right) \\ &\quad - S_{i+1} \frac{\Gamma(S(i))}{\Gamma(m(i))} \frac{\partial(\Gamma(m(i))/\Gamma(S(i)))}{\partial S(i)} \\ &\quad + n_{i+1} \left(\frac{\partial(\Gamma(n(i))\Gamma(m(i)))}{\partial n(i)}\right) / (\Gamma(n(i))\Gamma(m(i))), \end{aligned}$$

where $m(i) = n(i) - S(i)$.

EXAMPLE 1.3. Let Y_i represent the random number of events of a given kind observed within a time interval $(t_{i,1}, t_{i,n_i}]$ of fixed length. For this example we assume that the latter is identical for all samples $i = 1, \dots, N$. Let the random times at which the events take place be distributed according to a homogeneous Poisson process with intensity λ_i , so that the distribution of Y_i is Poisson with parameter $\lambda_i^* = \lambda_i(t_{i,n_i} - t_{i,1})$. The canonical form of the Poisson distribution has parameter $\theta_i = \log(\lambda_i^*)$ and functions $a(Y_i) = \frac{1}{Y_i!}$, $b(\theta_i) = e^{\theta_i}$. The conjugate prior for λ_i^* is Gamma with parameters $\text{Ga}(S_0, n_0)$ having mean y_0 and variance $\frac{y_0}{n_0}$. Upon observing (y_1, \dots, y_i) , the conjugate posterior for λ_i^* is $\text{Ga}(S(i), n(i))$ with $S(i) = S_0 + \sum_{j=1}^i y_j$, $n(i) = n_0 + i$. When also y_{i+1} is observed, using (1.7), the KL statistic has form

$$\begin{aligned} \text{KL}(y^{0:(i+1)}) &= \log\left(\frac{S(i)n(i)^{S(i)}}{n(i+1)^{S(i+1)}}\right) \\ &\quad + y_{i+1} \left(\log(n(i)) - \left(\frac{\partial\Gamma(S(i))}{\partial S(i)}\right) / \Gamma(S(i))\right) - \frac{S(i)}{n(i)}. \end{aligned}$$

1.5. *Effect of change-points on predictive densities.* In this section we illustrate analytically the effect of detecting a change-point on the one-step ahead predictive density using the transfer prior (1.1) and a conjugate Gaussian dynamic linear model. For each value of i , in what follows we let the scalar random variable Y_i be distributed as $N(\mu_i, \sigma_i^2)$. Analogously to Example 1.1, the prior distribution for $\theta_i = (\mu_i, \sigma_i^2)$ is taken as the conjugate Normal-inverse Gamma

$$\begin{aligned} \mu_i &\sim N(\hat{\mu}_{i^*-1}, \sigma_i^2), \\ \sigma_i^2 &\sim \text{IGa}\left(\frac{\nu}{2}, \frac{\nu}{2} \hat{\sigma}_{i^*-1}^2\right), \end{aligned}$$

where $1 \leq i^* < i$ is the time of the last detected change-point and $(\hat{\mu}_{i^*-1}, \hat{\sigma}_{i^*-1}^2)$ represent the estimated mean and variance of the joint posterior density at time i^* .

If $i^* = 1$, (μ_0, σ_0^2) represents a fixed initial condition. Here the prior density of the variance is

$$f(\sigma_i^2 | \nu, \hat{\sigma}_{i^*-1}^2) = \frac{(\nu \hat{\sigma}_{i^*-1}^2 / 2)^{\nu/2}}{\Gamma(\nu/2)} \sigma_i^{-2(\nu/2+1)} e^{-(\nu \hat{\sigma}_{i^*-1}^2) / (2\sigma_i^2)}.$$

Under this formulation, the prior expectation of the mean is $\hat{\mu}_{i^*-1}$ and that of the variance is $\frac{\nu/2}{\nu/2-1} \hat{\sigma}_{i^*-1}^2$. It follows that the one-step ahead marginal predictive distribution is a noncentral Student- t . In absence of change-points prior to time i , the predictive density is

$$(1.8) \quad Y_{i+1} \sim t_{\nu+i} \left(\tilde{\mu}_i, \frac{i+1}{i+2} \frac{\nu+i}{2} \frac{1}{\tilde{\sigma}_i^2} \right),$$

where

$$\begin{aligned} \tilde{\mu}_i &= \frac{1}{1+i} \mu_0 + \frac{i}{1+i} \bar{y}^{(1:i)}, \\ \tilde{\sigma}_i^2 &= \frac{\nu}{2} \sigma_0^2 + \frac{i}{2} s_{(1:i)}^2 + \frac{i}{i+1} (\mu_0 - \bar{y}^{(1:i)})^2, \end{aligned}$$

and $(\bar{y}^{(1:i)}, s_{(1:i)}^2)$ represent respectively the sample mean and variance of the data $y^{1:i}$. If a change-point is detected by the KL statistic (1.3) at time $1 < i^* < i$, under the transfer prior (1.1), the conditional predictive density is

$$(1.9) \quad Y_{i+1} \sim t_{\nu+i-(i^*-1)} \left(\tilde{\mu}_i^*, \frac{i-i^*+2}{i-i^*+3} \frac{\nu+i-i^*+1}{2} \frac{1}{(\tilde{\sigma}_i^2)^*} \right),$$

where

$$\begin{aligned} \tilde{\mu}_i^* &= \frac{1}{i-i^*+2} \hat{\mu}_{i^*-1} + \frac{i-(i^*-1)}{i-i^*+2} \bar{y}^{(i^*:i)}, \\ (\tilde{\sigma}_i^2)^* &= \frac{\nu}{2} (\hat{\sigma}_{i^*-1}^2)^2 + \frac{i-(i^*-1)}{2} s_{(i^*:i)}^2 + \frac{i-(i^*-1)}{i-i^*+2} (\hat{\mu}_{i^*-1} - \bar{y}^{(i^*:i)})^2. \end{aligned}$$

Since the mean and variance of the noncentral Student- t random variable with density $t_\nu(\mu, \sigma^2)$ are respectively equal to μ and to $\frac{\nu}{\nu+2} \sigma^2$, equations (1.8) and (1.9) provide a characterization of the one-step ahead posterior predictive moments as a function of the time of the last detected change-point and of the inverse-Gamma prior coefficient ν . For $i^* > 1$ the predictive mean is less influenced by the sample mean of the data preceding the change-point, $\bar{y}^{1:(i^*-1)}$, and it is more heavily influenced by $\bar{y}^{i^*:i}$, that is, the sample mean of the data from the change-point on. When a change-point is detected, the predictive variance is larger with respect to the case of no change. Its relative increase is a decreasing function of the difference $(i - i^*)$, which measures how far in time the change-point occurred, and it is an increasing function of the coefficient ν , which measures the strength of the prior at the initial time.

This behavior is consistent with the intuition that predictions ensuing from a dynamic time series model should discount the information content of remote data and focus on more recent data when significant dynamics occur. In absence of an autoregressive model structure, as in the present section, the distinction between remote and recent data is entirely left to the timing of the detected change-points.

2. Analysis of multivariate EEG recordings. This section presents an application of the methods discussed above to estimating neural functional dynamics using multivariate electroencephalogram (EEG) recordings. The data analyzed here arise from a sequence of 80 identical tests each having length of approximately four seconds with sampling rate of 128 points per second. During each test, the same subject was to press a button when a green square appeared in a specific screen location [Makeig et al. (2002)]. Previous analyses of these data have emphasized aspects of time-dependent interactions among different EEG channels, such as an increased overall synchronization of different brain areas after presentation of the visual stimulus [Delorme et al. (2002)]. The multidimensional EEG time series are modeled here as a discrete time Gaussian stochastic process, in which randomness is thought of as arising from the intrinsic variability of the brain activity and from the presence of experimental artifacts. We describe the dynamic functional relationships among different brain areas using the time-dependent means and covariance matrices indexing the data likelihood.

The 32 EEG channels record neural activity arising from seven functionally distinct brain areas, that are the frontal (F), central (C), central-parietal (CP), parietal (P), temporal (T), parietal-occipital (PO) and occipital (O) lobes. At each time point the recording channels targeting each of the seven brain areas were averaged within each trial and then across trials so as to obtain a seven-dimensional time series. The rationale for this preprocessing is that recordings within each brain area exhibit similar patterns within and across trials so that, for the purpose of our analysis, averaging yields a lower dimensional signal less affected by channel-specific recording noise. These trial-averaged EEG recordings are represented in Figure 2. The activity of the different areas prior to the presentation of the visual cue are tightly synchronized, exhibiting oscillations of high amplitude around frequency 10 Hz and fast low-amplitude oscillations at 60 Hz. Due to their low amplitude, the latter are hard to see in Figure 2. The lower frequency oscillations are consistent with the so-called α band reflecting eye movements. The higher frequency and lower amplitude oscillations are due to the alternating current being used in this experiment, suggesting an imperfect electrode grounding.

The trial-averaged signal at time i , Y_i , is modeled as $N_7(\mu_i, \Sigma_i)$. To derive Bayesian inferences for the mean vector and for the covariance matrix, we use the conjugate Normal-inverse Wishart prior:

$$(2.1) \quad \mu_i \sim N_7(\hat{\mu}_{i^*-1}, \Sigma_i),$$

$$(2.2) \quad \Sigma_i \sim IW_7(9, I_7 \hat{\Sigma}_{i^*-1}),$$

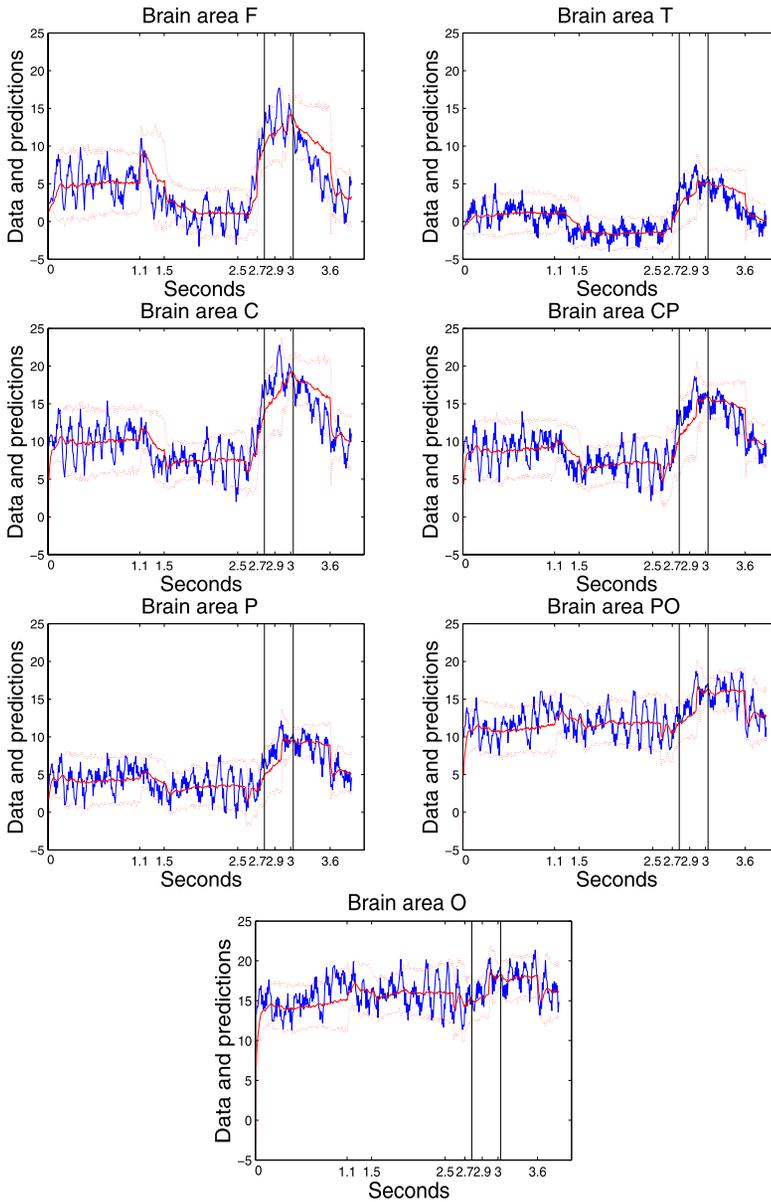


FIG. 2. EEG recordings (blue), one-step ahead marginal posterior point predictions and 95% posterior predictive intervals for each brain area (red). The estimated change-point times are marked on the horizontal axis of each plot. The two vertical lines represent respectively the average stimulus and response times. The brain activity is reduced roughly at half of the initial phase of the experiment and it increases when the cue is presented. The sharpest increases are detected in the frontal (F) and central (C) lobes, followed by the central-parietal (CP), parietal (P) and temporal (T) lobes. The estimated change in activity in the parietal-occipital (PO) and occipital (O) areas is far less pronounced.

where $1 \leq i^* < i$ is the time of the last detected change-point prior to time i . The marginal prior expectations are matched to the corresponding estimated marginal posterior moments at time $i^* - 1$ consistently with (1.1). At time i these prior distributions are updated using Bayes' theorem, taking into account all data points within the interval $[i^* + 1, i]$. Therefore, by combining the KL test with a static Bayesian update, this dynamic model retains a memory of past mean and covariance estimates from the last detected change-point onward. From this perspective, this model can be thought of as a form of time-varying vector autoregression which at time i has order $i - i^*$.

For this analysis, the initial conditions μ_0 and Σ_0 were set respectively equal to the null vector and to the identity matrix. The hyper-parameter of the posterior density was set at $\alpha = 0.01$, so as to detect only the most prominent changes. The number of degrees of freedom of the Inverse Wishart density is set so that predictive intervals of length consistent with the set value of α are not excessively inflated when a change is detected. The distribution of the KL statistic and its value were approximated at each time i using the last 500 Gibbs sampler draws of the mean and of the covariance matrix.

Along with the data, Figure 2 shows the one-step ahead marginal posterior point predictions and their 95% highest posterior predictive intervals for each of the seven brain areas. The estimated change-point times are marked along the horizontal axes. The predictions emphasize a downward shift in brain activity taking place roughly at half of the initial phase of the experiments, followed by a sharp increase corresponding to the cue presentation, a downward trend following the motor response and a stabilization of the EEG signals toward the end of the experiments. The first two change-points identify a transition during the first part of the experiment toward a state of more intense attention. The third to sixth change-points capture an abrupt increase in neural activity related to the presentation of the visual cue, whereas the last change-point indicates a return to a baseline activity. The sharp increases in the activity of the frontal and central areas during the generation of the response are consistent with their characterization as executive and motor centers of the brain. The intermediate increase in activity of the temporal and parietal lobes, mainly involved in speech, hearing, memory and in the integration of sensory inputs, reflects the mild involvement of their functions in the execution of the task entailed by this experimental protocol. The mild response to the visual stimulus of the parietal-occipital and occipital areas, including the visual cortex, is somewhat surprising. An analogous analysis of the trial-averaged EEG data from the eight distinct channels recording from these two areas reveals a consistently higher activity of the occipital channels with respect to the parietal-occipital ones but no significant change in response to the visual stimulus.

Figure 3 depicts the estimates of the time-dependent variance and covariance functions for each brain area. Whole segments represent periods during which their respective 95% highest posterior intervals do not intersect zero. The estimated change-point times are marked on the horizontal axis of each plot, as in

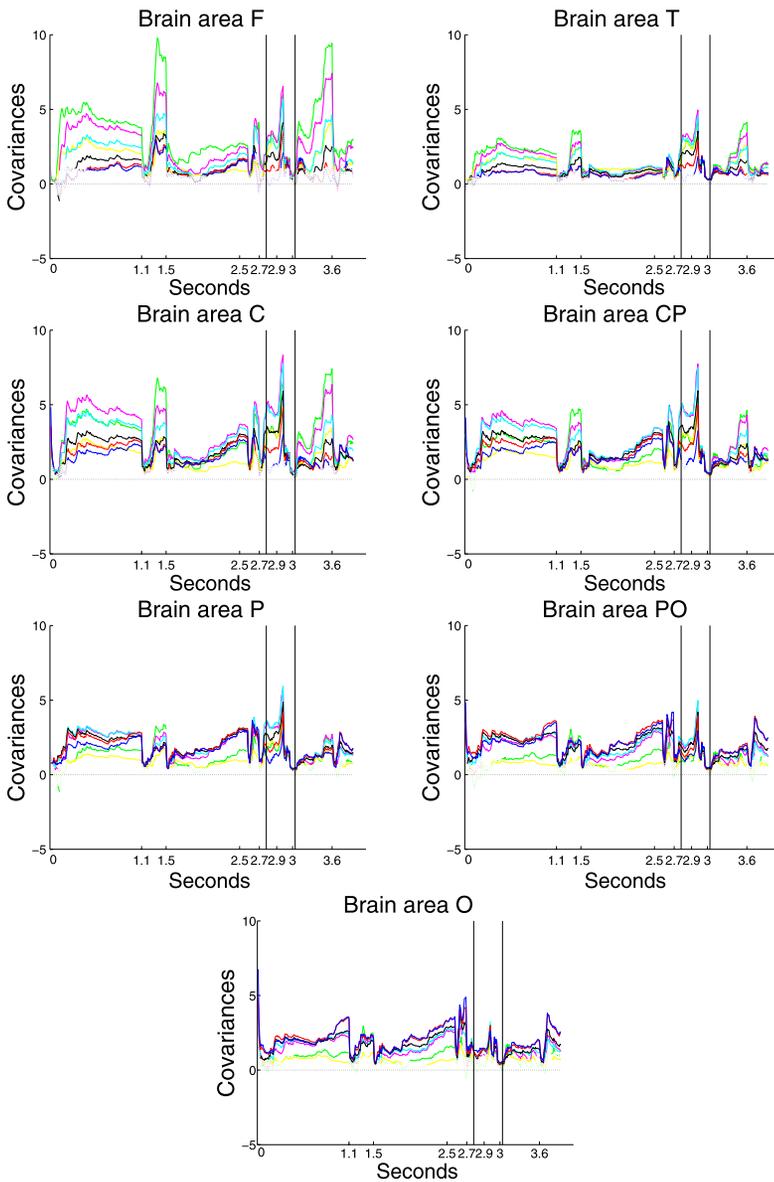


FIG. 3. Estimates of the time-dependent variance and covariance functions for the frontal (green), temporal (yellow), central (magenta), central-parietal (cyan), parietal (black), parietal-occipital (red) and occipital (blue) lobes. Whole segments represent periods during which their 95% posterior intervals do not intersect zero. The estimated change-point times are marked on the horizontal axis of each plot. The estimated covariances are almost always positive and time-varying, representing different levels of cooperative activity of the seven brain areas over time. The covariance functions are also spatially ordered, the strongest relationships being estimated between physically adjacent brain areas.

Figure 2. All estimated variances and covariances vary over time, indicating that a time-dependent covariance matrix is an appropriate modeling assumption for this data. The estimated covariances are almost always positive, suggesting that the activity of the seven brain areas is dynamically cooperative as found by DeLorme et al. (2002). An unexpected feature of the estimated covariance functions is their spatial ordering over time, the strongest relationships being estimated between adjacent brain areas. Since neither in the Gaussian likelihood nor the priors (2.1)–(2.2) include a spatial component, these estimates suggest a close correspondence between the detected functional relationships and the anatomical structure of the brain.

3. Estimation of a learning curve. The data analyzed in this section arises from a sequence of 55 trials during which a macaque monkey performed a location-scene association task [Wirth et al. (2003)]. The learning curve is represented by the time-dependent estimates of the trials' success probabilities. Smith et al. (2004) introduced a parametric state-space model for inferring the learning performance using longitudinal behavioral experiments. The learning curve is thereby modeled using univariate binary time series data along with a logit link for each trial's success probability and a Gaussian state evolution equation for the parameters' dynamics. In this section we use the same Bernoulli sampling distribution for the binary trial outcomes as in Smith et al. (2004) and we estimate the dynamics of its success probability over time using the semi-parametric method illustrated in Section 1. A first difference between our model and that of Smith et al. (2004) is that we do not use a nonlinear link function, thus imposing fewer constraints on the shape of the learning curve. A second difference is that the results of Smith et al. (2004) are based on a smoothing algorithm using both past and future data to obtain estimates at present times, whereas our method uses past observed values and simulated current data to update the distribution of the success probability. In the following analyses the success probability of the first trial was given a uniform prior, whereas the transfer prior (1.1) was implemented using a conjugate Beta prior. The data were analyzed under different values for the hyper-parameter α within the range (0.01, 0.9), respectively requiring from strong to weak evidence for detecting a change-point. The distribution of the KL statistic under the null hypothesis of no change was approximated using ten thousand Monte Carlo samples from the Beta posterior distribution of each trial's success probability. For this data, the smoothed state-space estimates of Smith et al. (2004) indicate that with 90% confidence the success probability significantly exceeds its chance value 0.25 from trial 23 onward, whereas their unsmoothed estimates indicate that the chance value is significantly exceeded from trial 27 onward. Figure 4 shows our estimates of the success probabilities under the four selected values of $\alpha = 0.9, 0.5, 0.1, 0.01$. From these estimates we conclude that learning has effectively taken place from trial 29 onward, that is, after observing a total of 7 successes yielding an empirical cumulative success rate of 0.24. Figure 4 also compares our dynamic estimates with the

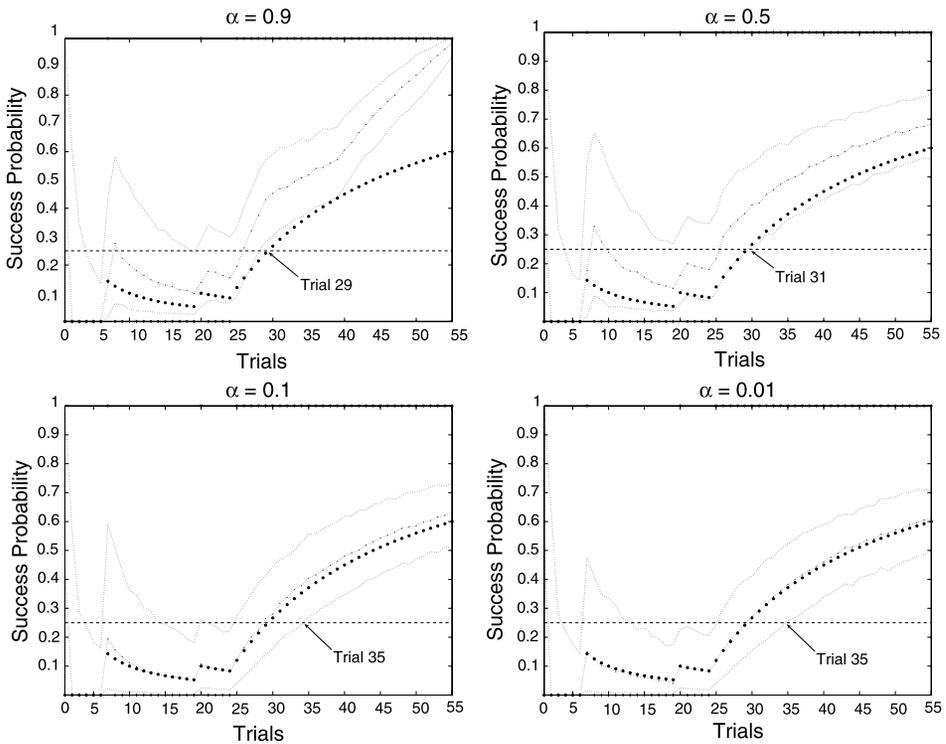


FIG. 4. *Macaque monkey binary data and semi-parametric estimates of their time-dependent success probabilities using $\alpha = 0.9, 0.5, 0.1, 0.01$. The binary data are represented as vertical ticks along the lower and upper horizontal axes. Asterisks represent the cumulative proportion of successful trials. The sequence of estimates of the success probabilities describe the macaque's learning curve over time. The first trial at which the learning curve lies above its chance level 0.25, indicating that learning has effectively taken place, is number 29. Lower values of α require more extreme values of the KL statistic for detecting change-points, making our estimates of the learning curve progressively closer to the empirical cumulative success rates.*

empirical cumulative proportion of successful trials, which is represented by asterisks. As the value of α decreases, so does the number of detected change-points. In particular, under the uniform prior for the initial success probability when $\alpha \leq 0.1$, our estimates of the learning curve are roughly equivalent to the empirical proportion of cumulative successes.

4. Dynamic modeling of functional neuronal networks. This example illustrates the application of the method presented in Section 1 in the context of a model for networks of spiking neurons. During the experiments analyzed here, the neural activity of a small section of a sheep's temporal cortex is recorded in vivo on a millisecond time frame using a multi-electrode array [Kendrick et al. (2001)]. The goal of these experiments was to investigate in detail the activity of brain areas

associated with memory. Along each of 77 disconnected experiments, a sheep is shown either a blank screen or two images. In the latter case, a reward is given when one of a set of “familiar faces” is correctly identified. It is important to note that, even within small brain areas, these experimental techniques only record the activity of a relatively small fraction of neurons. Therefore, these data do not allow reconstructing direct physical interactions among neurons but only functional relationships among relatively distant recording electrodes.

Introductions to the neuronal physiology and to neuronal modeling are presented in Fienberg (1974) and Brillinger (1988). Recent surveys of the state-of-the-art in multiple spike trains modeling can be found in Iyengar (2001), Brown, Kass and Mitra (2004), Kass, Ventura and Brown (2005), Okatan, Wilson and Brown (2005), Rao (2005) and Rigat, de Gunst and ven Pelt (2006). Dynamic point process neuronal models based on fully parametric state-space representations have been proposed by Eden et al. (2004), Truccolo et al. (2005), Brown and Barbieri (2006), Srinivansan et al. (2006) and Eden and Brown (2008).

4.1. *Binary network model.* In what follows each element of the experimental time series $\{Y_i\}_{i=1}^{77}$ is $Y_{i,k,t_{i,j(i)}} = 1$ if neuron k fires at time $t_{i,j(i)}$ during trial i and $Y_{i,k,t_{i,j(i)}} = 0$ otherwise with $j(i) = 1, \dots, n_i$. We model the joint sampling distribution of the multiple spike train data for trial i , Y_i , as a Bernoulli process with renewal [Rigat, de Gunst and ven Pelt (2006)]. The joint probability of a given realization y_i is

$$(4.1) \quad P(Y_i = y_i | \pi_i) = \prod_{t=t_{i,1}}^{t_{i,n_i}} \prod_{k=1}^K \pi_{i,k,t}^{y_{i,k,t}} (1 - \pi_{i,k,t})^{1-y_{i,k,t}}.$$

For model (4.1) to be biologically interpretable, the firing probability of neuron k at time $t_{i,j(i)}$ during trial i , $\pi_{i,k,t_{i,j(i)}}$, is defined as a one-to-one nondecreasing mapping of a real-valued voltage function $v_{i,k,t_{i,j(i)}}$ onto the interval $(0, 1)$. The function $v_{i,k,t_{i,j(i)}}$ represents the unnormalized difference of electrical potential across the membrane of neuron k at time $t_{i,j(i)}$. Let $\tau_{i,k,t_{i,j(i)}}$ be the last spiking time of neuron k prior to time $t_{i,j(i)}$ during trial i , that is,

$$\tau_{i,k,t_{i,j(i)}} = \begin{cases} 1, & \text{if } \sum_{\tau=1}^{t_{i,j(i)}} Y_{i,k,\tau} = 0 \text{ or } t_{i,j(i)} = 1, \\ \max\{1 \leq \tau < t_{i,j(i)} : Y_{i,k,\tau} = 1\}, & \text{otherwise} \end{cases}$$

and the voltage function is modeled as

$$(4.2) \quad v_{i,k,t_{i,j(i)}} = \sum_{l=1}^K \beta_{i,k,l} \sum_{w=\tau_{i,k,t_{i,j(i)}}}^{t_{i,j(i)}-1} y_{i,l,w}.$$

The spiking probabilities are linked to (4.2) via the logistic mapping

$$\pi_{i,k,t_i,j(i)} = \frac{e^{v_{i,k,t_i,j(i)}}}{1 + e^{v_{i,k,t_i,j(i)}}}.$$

The coefficients $\beta_{i,k,l}$ represent the strength of the functional relationship from neuron l to neuron k during trial i . When $\beta_{i,k,l}$ is positive during trial i , the firing activity of neuron l promotes that of neuron k , whereas when it is negative, firing of l inhibits that of k . When neurons l and k are physically connected to each other, the coefficients $\beta_{i,k,l}$ and $\beta_{i,l,k}$ represent direct functional connections. When the two neurons are not directly connected to each other, these network coefficients summarize a functional relationship possibly arising from a long chain of neurons in which activity cannot be currently recorded by the MEA technique. The coefficients $\beta_{i,k,k}$ represent the spontaneous spiking rate of neuron k during trial i . The last summation term in equation (4.2) indicates that the membrane potential of a neuron is assumed to be influenced only by the spiking activity of the other neurons during its last inter-spike interval. In this simple model we do not take into account the occurrence of leakage currents across the neuronal membrane [Plesser and Gerstner (2000)], so that the effect of the spikes produced by neuron l on the voltage function does not decrease over time.

For each trial $i = 1, \dots, N$ we use a Metropolis sampler to produce approximate posterior inferences for the K^2 model parameters. For each experiment, we run a neuron-wise random scan update with independent Gaussian random walk proposals for twenty-five thousand iterations. The initial prior for the parameters of all experiments is Gaussian with zero mean, standard deviation 1 and zero covariance for all pairs of neurons. Conditionally on the data $y^{0:i}$ and on the current posterior estimates, upon observing the outcome of the i th + 1 experiment, y_{i+1} , we use the KL statistic (1.3) to test whether a significant change occurred in any of the model's parameters. The occurrence of such changes and the corresponding parameter estimates indicate statistically significant variations of different aspects of the neural activity.

4.2. Analysis of sheep multiple spike trains. In this section we analyze the spiking activity of the 7 most active electrodes among the 64 recording channels. The plot on the left in Figure 5 shows the number of spikes recorded from these 7 electrodes along all 77 experiments. The panel on the right shows the mean spiking rates for each electrode and experiment, which reflect the overall low spiking rates typical of this type of measurement. The co-occurrence of relatively high firing rates for all electrodes suggests that the most prominent connections among the underlying neurons may be mutually excitatory functional relationships.

Table 1 displays summaries of the point estimates of the network coefficients across all experiments. Each cell reports the proportions of experiments during which each of the network coefficients were found either significantly excitatory or inhibitory. Significance here denotes experiments during which both 95% end points of the posterior interval of a pair-wise functional connection lie respec-

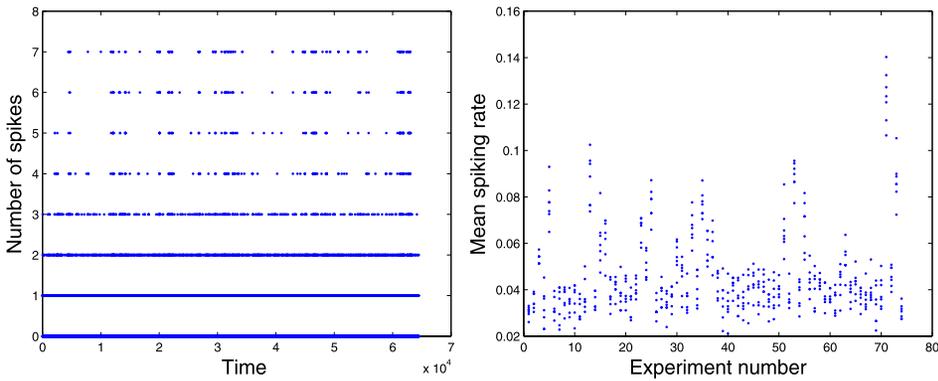


FIG. 5. Each dot in the left panel marks the number of recorded spikes of the 7 most active electrodes for each millisecond of the 77 experiments. Each dot in the right panel marks the proportion of milliseconds during which each electrode recorded a spike during each experiment. The range of these mean firing rates is 0.02–0.14, reflecting the low overall spiking rates typical for this type of recording. Clusters of points associated to relatively high mean spiking rates suggest that the underlying neurons may be mostly connected via mutually excitatory functional relationships.

tively above or below zero. The self-dependence coefficients on the main diagonal are always found significant and negative, representing the well-known property of neural refractoriness. The excitatory functional connection from electrode 3 toward 6 is most prominent, being significant over approximately 63% of the experi-

TABLE 1

Relative number of experiments during which both end points of the 95% posterior interval for any of the pair-wise functional connection coefficients lie respectively above or below zero, identifying significant excitatory (left proportion) or inhibitory (right proportion) relations

$i \setminus j$	1	2	3	4	5	6	7
1	0.00 1.00	0.10 0.28	0.28 0.35	0.30 0.12	0.30 0.30	0.08 0.29	0.08 0.38
2	0.01 0.46	0.00 1.00	0.32 0.35	0.05 0.60	0.25 0.25	0.30 0.14	0.41 0.12
3	0.25 0.13	0.28 0.14	0.00 1.00	0.08 0.30	0.25 0.12	0.01 0.68	0.08 0.28
4	0.09 0.40	0.05 0.36	0.34 0.14	0.00 1.00	0.33 0.12	0.08 0.13	0.10 0.36
5	0.28 0.10	0.08 0.42	0.51 0.13	0.08 0.62	0.00 1.00	0.08 0.28	0.25 0.25
6	0.10 0.39	0.10 0.13	0.63 0.00	0.08 0.51	0.05 0.30	0.00 1.00	0.08 0.40
7	0.09 0.41	0.10 0.14	0.30 0.28	0.08 0.14	0.08 0.28	0.05 0.39	0.00 1.00

The self-dependence coefficients on the main diagonal are always found significant and negative, representing the well-known property of neural refractoriness. Bold entries represent functional connections which are found significant over more than 50% of the experiments. The excitatory functional connection from electrode 3 toward 6 is most prominent, being significant over approximately 63% of the experiments. The most prominent inhibitory connections relate electrode 4 to 5, which is significant over 62% of the experiments, and electrode 6 to 3, which is found significant over 68% of the experiments.

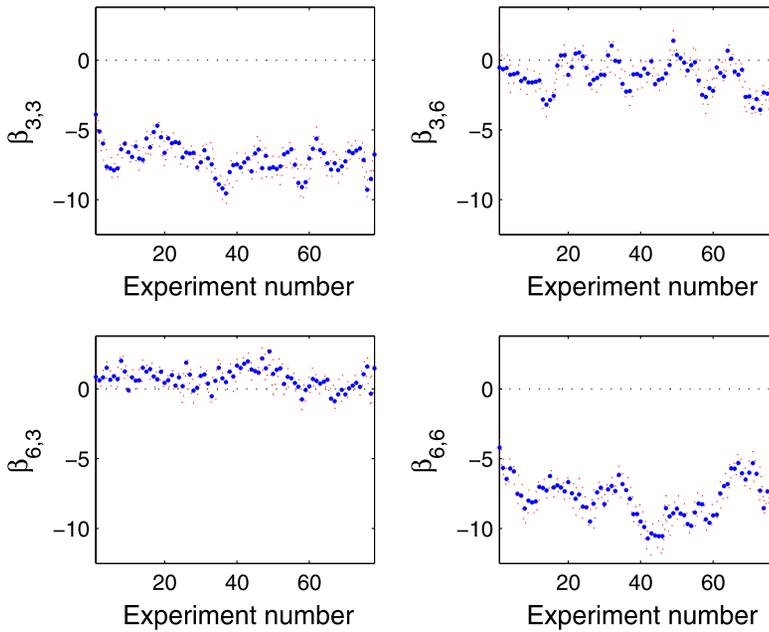


FIG. 6. Point estimates and 95% highest posterior intervals for the self-dependence parameters of electrodes 3 and 6 (main diagonal) and of their pair-wise functional connections $\beta_{3,6}$ and $\beta_{6,3}$. These two electrodes exhibit a comparable level of refractoriness over all experiments. The estimated correlations over experiments between the self-dependence coefficients $\beta_{3,3}$ and those of $\beta_{3,6}$ is -0.28 and that between $\beta_{6,6}$ and $\beta_{6,3}$ is -0.29 , suggesting that neural self-inhibition may tend to compensate for excitations and inhibitions supplied by functionally connected cells.

ments. The most prominent inhibitory connections are found significant over 62% and 68% of the experiments and they relate respectively electrodes 4 to 5 and 6 to 3. Note that the time series model (4.2) identifies a directed cyclic graph (DCG) of pair-wise functional relationships where the connections $i \rightarrow j$ and $j \rightarrow i$ are captured by distinct coefficients, so that the proportion of $3 \rightarrow 6$ significant excitatory connections and that of $6 \rightarrow 3$ significant inhibitory connections are not constrained to add up to one. Figure 6 illustrates in detail the point estimates and the 95% highest posterior intervals for the most prominent excitatory connection, $3 \rightarrow 6$, together with those of both electrodes' self-dependence and of the mostly inhibitory connection $6 \rightarrow 3$. The estimated correlation over experiments between the self-dependence coefficients $\beta_{3,3}$ and those of $\beta_{3,6}$ is -0.28 and that between $\beta_{6,6}$ and $\beta_{6,3}$ is -0.29 , suggesting that neural self-inhibition may tend to compensate for excitations and inhibitions supplied by the other recorded functionally connected cells.

5. Discussion. This work is motivated by the challenges encountered in constructing time series models when the factors driving the dynamics of their pa-

rameters are not well understood. The semi-parametric method illustrated here provides flexible time-dependent estimates without relying on explicit modeling of these dynamics. For exploratory data analyses, such as those presented in Sections 2, 3 and 4, these estimates may suffice to address specific scientific questions. Otherwise, appropriate measures of dependence between these time-dependent estimates and experimental factors of interest provide a principled basis for more precise formulations of the parameters' dynamics. Describing the exact form of such dependence measures is very much context-dependent and it lies outside of the scope of this work.

A distinctive feature of the modeling approach proposed here is that it combines elements of sequential Bayesian learning and conditional frequentist inference along the lines of Guttman (1967), Box (1980), Berger, Brown and Wolpert (1994), Meng (1994), Gelman, Meng and Stern (1996), Berger and Bayarri (1997), Spiegelhalter et al. (2002), Bayarri and Morales (2003), Kuhnert, Mergesen and Tesar (2003) and Bayarri and Berger (2004), among others. A general treatment of such pragmatic combination of frequentist and Bayesian ideas for model criticism can be found in Chapter 8 of O'Hagan and Forster (1999). From this perspective, our method is a "Bayesianly justifiable" procedure [Rubin (1984)] because only those future unobserved data that are consistent with the current conditional posterior distribution of the model's parameters are relevant for approximating the distribution of the KL change-point statistic (1.3).

The latter reflects a notion of change-point as an observation which, on the basis of the chosen model with its prior and the observations accrued so far, is "surprising" from a predictive point of view. Note that this characterization does not depend on the parametrization of the state space nor on the unobservable sample paths of latent states, but it depends only on the predictives on observables. Defining models and their properties via their one step ahead predictive statements has been recommended, among others, by Geisser and Eddy (1979) and San Martini and Spezzaferri (1984) for predictive model selection, by Dawid (1984) in his prequential inference, by West and Harrison (1986) for monitoring the adequacy of Bayesian forecasting models and by Smith (1992) for comparing the characteristics of different forecasting models. More recently, optimal predictive model selection criteria have been proposed by Barbieri and Berger (2004).

The results presented in Section 3 revealed a substantial dependence of the estimated learning curve with respect to the value of the hyper-parameter α . It is important to recall that this hyper-parameter measures how extreme a value of the KL statistic is needed for detecting a change-point. Therefore, a dependence of its corresponding estimated change-point process on the value of α is to be expected, with lower values of this hyper-parameter yielding less numerous change-points and vice versa. From this perspective, our method is not meant to be fully automatic and parameter estimates derived using different values of α should be inspected to gauge their sensitivity in the context of the specific time series model being entertained.

In this work, a single change-point process common to all model's parameters is used to define their conditional posterior distribution. Should the data provide evidence of changes of only some parameters, the posterior distributions for the unchanging coefficients would not make the most efficient use of the data. It is important to note that while in principle any subset of model parameters can be associated to a distinct change-point process, the limitations for implementing multivariate change-point process inference within our framework are eminently practical. This is because marginal likelihoods for each subset of model parameters having a different change-point process are required to approximate the distribution of their change-point test statistic. For classes of models where marginal likelihoods are available in closed form, this work can be extended by introducing a random variable identifying groups of coefficients sharing a common change-point process.

Posterior simulation via Markov chain Monte Carlo algorithms has been used in this work to fit multivariate time series models and to approximate critical values of the KL statistic. Although the current implementation of our method is operationally realistic, these computationally intensive methods are in fact rather impractical for an iterative process of model formulation and criticism. Currently two directions are being pursued to improve the computational efficiency of our method. On the one hand, faster resampling methods such as particle filters [Doucet, De Freitas and Gordon (2001)] and approximate Bayesian computation [Marjoram et al. (2003)] can be adopted. Alternatively, analytical posterior approximations can be adopted [Tierney and Kadane (1986)]. For instance, in the context of sequential time series modeling, Koyama, Perez-Bolde and Kass (2008) recently proposed a Laplace–Gauss posterior approximation that obviates the use of cumbersome resampling techniques.

Acknowledgments. The authors acknowledge the support of the Centre for Research in Statistical Methodology (CRiSM) at the University of Warwick and of the Warwick Centre for Analytical Science during the development of this work. We wish to thank Arnaud Delorme for sharing the EEG recordings analyzed in Section 2. We also wish to thank Professor Jiangfeng Feng and collaborators for providing the multiple spike trains analyzed in Section 4.

REFERENCES

- AKAIKE, H. (1978). On the likelihood of a time series model. *The Statistician* **27** 217–235.
- AKAIKE, H. (1981). Likelihood of a model and information criteria. *J. Econometrics* **16** 3–14.
- ALBERT, J. H. and CHIB, S. (1993). Bayes inference via Gibbs sampling for autoregressive time series subject to Markov mean and variance shifts. *J. Bus. Econom. Statist.* **11** 1–15.
- BARBIERI, M. M. and BERGER, J. O. (2004). Optimal predictive model selection. *Ann. Statist.* **32** 870–897. MR2065192
- BARNARD, G. A. (1959). Control charts and stochastic processes. *J. Roy. Statist. Soc. Ser. B* **21** 239–271.

- BAYARRI, M. J. and BERGER, J. O. (2004). The interplay between Bayesian and frequentist analysis. *Statist. Sci.* **19** 58–80. [MR2082147](#)
- BAYARRI, M. J. and MORALES, J. (2003). Bayesian measures of surprise for outlier detection. *J. Statist. Plann. Inference* **111** 3–22. [MR1955869](#)
- BÉLISLE, P., JOSEPH, L., MACGIBBON, B., WOLFSON, D. B. and DU BERGER, R. (1998). Change-point analysis of neuron spike train data. *Biometrics* **54** 113–123.
- BENGTSOON, T. and CAVANAUGH, J. E. (2006). An improved Akaike information criterion for state-space model selection. *Comput. Statist. Data Anal.* **50** 2635–2654. [MR2227324](#)
- BERGER, J. O. and BAYARRI, M. L. (1997). Measures of surprise in Bayesian analysis. ISDS Discussion Paper 46, Duke Univ.
- BERGER, J. O., BROWN, L. and WOLPERT, R. L. (1994). A unified conditional frequentist and Bayesian test for fixed and sequential hypothesis testing. *Ann. Statist.* **22** 1787–1807. [MR1329168](#)
- BERNARDO, J. (1979). Expected information as expected utility. *Ann. Statist.* **7** 686–690. [MR0527503](#)
- BERNARDO, J. M. and SMITH, A. F. M. (2007). *Bayesian Theory*. Wiley, Chichester, UK. [MR1274699](#)
- BOX, G. E. P. (1980). Sampling and Bayes' inference in scientific modelling and robustness. *J. Amer. Statist. Assoc.* **143** 383–430. [MR0603745](#)
- BRILLINGER, D. R. (1988). Some statistical methods for random processes data from seismology and neurophysiology. *Ann. Statist.* **16** 1–54. [MR0924855](#)
- BROWN, E. N. and BARBIERI, R. (2006). Dynamic analyses of neural representations using the state-space modeling paradigm. In: *The Cell Biology of Addiction* (B. Madras, M. Von Zastrow, C. Colvis, J. Rutter, D. Shurtleff and J. Pollock, eds.). Cold Spring Harbor Laboratory Press, New York.
- BROWN, E., KASS, R. E. and MITRA, P. P. (2004). Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nature Neuroscience* **7** 456–461.
- BUZSÁKI, G. (2004). Large scale recording of neuronal ensembles. *Nature Neuroscience* **7** 446–451.
- CAPPE, O., MOULINES, E. and RYDEN, T. (2005). *Inference in Hidden Markov Models*. Springer, New York. [MR2159833](#)
- CARLIN, B. P., GELFAND, A. E. and SMITH, A. F. M. (1992). Hierarchical Bayesian analysis of changepoint problems. *App. Statist.* **41** 389–405.
- CAROTA, C., PARMIGIANI, G. and POLSON, N. (1996). Diagnostic measures for model criticism. *J. Amer. Statist. Assoc.* **91** 753–762. [MR1395742](#)
- CHIB, S. (1995). Marginal likelihood from the Gibbs output. *J. Amer. Statist. Assoc.* **90** 1313–1321. [MR1379473](#)
- CHIB, S. (1998). Estimation and comparison of multiple change-point models. *J. Econometrics* **86** 221–241. [MR1649222](#)
- CRITCHLEY, F., MARRIOTT, P. and SALMON, M. (1994). Preferred point geometry and the local differential geometry of the Kullback–Leibler divergence. *Ann. Statist.* **22** 1587–1602. [MR1311991](#)
- DAWID, A. P. (1984). Present position and potential developments: Some personal views. Statistical theory: The prequential approach. *J. Roy. Statist. Soc. Ser. A* **147** 278–292. [MR0763811](#)
- DELORME, A., MAKEIG, S., FABRE-THORPE, M. and SEJNOWSKI, T. J. (2002). From single trial EEG to brain area dynamics. *Neurocomputing* **44** 1057–1064.
- DIACONIS, P. and YLVIKAKER, D. (1979). Conjugate priors for exponential families. *Ann. Statist.* **7** 269–281. [MR0520238](#)
- DOUCET, A., DE FREITAS, N. and GORDON, N. J. (2001). *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York. [MR1847783](#)
- EDEN, U. T. and BROWN, E. N. (2008). Continuous-time filters for state estimation from point process models of neural data. *Statist. Sinica*. [MR2468269](#)
- EDEN, U. T., FRANK, L. M., BARBIERI, R., SOLO, V. and BROWN, E. N. (2004). Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Comput.* **16** 971–998.

- FEARNHEAD, P. and LIU, Z. (2007). On-line inference for multiple change points. *J. Roy. Statist. Soc. Ser. B* **69** 589–605. [MR2370070](#)
- FERGER, D. (1995). Nonparametric tests for nonstandard change-point problems. *Ann. Statist.* **23** 1848–1861. [MR1370310](#)
- FIENBERG, S. E. (1974). Stochastic models for single neuron firing trains: A survey. *Biometrics* **30** 399–427. [MR0359082](#)
- FRÜHWIRTH-SHNATTER, S. (1995). Bayesian model discrimination and Bayes factors for linear Gaussian state-space models. *J. Roy. Statist. Soc. Ser. B* **1** 237–246. [MR1325388](#)
- FRÜHWIRTH-SHNATTER, S. (2001). Markov chain Monte Carlo estimation of classical and dynamic switching and mixture models. *J. Amer. Statist. Assoc.* **96** 194–209. [MR1952732](#)
- FRÜHWIRTH-SHNATTER, S. (2006). *Finite Mixture and Markov Switching Models*. Springer, New York. [MR2265601](#)
- GAMERMAN, D. (1992). A dynamic approach to the statistical analysis of point processes. *Biometrika* **79** 39–50. [MR1158516](#)
- GEISSER, S. and EDDY, W. F. (1979). A predictive approach to model selection. *J. Amer. Statist. Assoc.* **74** 153–160. [MR0529531](#)
- GELFAND, A. E. and DEY, E. K. (1994). Bayesian model choice: Asymptotics and exact calculations. *J. Amer. Statist. Assoc.* **56** 501–514. [MR1278223](#)
- GELFAND, A. E. and SMITH, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409. [MR1141740](#)
- GELMAN, A., MENG, X. L. and STERN, H. S. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statist. Sinica* **6** 733–807. [MR1422404](#)
- GHAHRAMANI, Z. and HINTON, G. E. (2000). Variational learning for switching state-space models. *Neural Comput.* **12** 831–864.
- GOUTIS, C. and ROBERT, C. (1998). Model choice in generalised linear models: A Bayesian approach via Kullback–Leibler projections. *Biometrika* **85** 29–37. [MR1627250](#)
- GUTIÉRREZ-PEÑA, E. (1997). Moments for the canonical parameter of an exponential family under a conjugate distribution. *Biometrika* **84** 727–732. [MR1603960](#)
- GUTTMAN, I. (1967). The use of the concept of a future observation in goodness-of-fit problems. *J. Roy. Statist. Soc. Ser. B* **29** 83–100. [MR0216699](#)
- HALL, P. (1987). On Kullback–Leibler loss and density estimation. *Ann. Statist.* **15** 1491–1519. [MR0913570](#)
- HAMILTON, J. D. (1990). Analysis of time series subject to changes in regime. *J. Econometrics* **45** 39–70. [MR1067230](#)
- HAMILTON, J. D. (1994). *Time Series Analysis*. Princeton Univ. Press, New Jersey. [MR1278033](#)
- HAN, C. and CARLIN, B. P. (2001). Markov chain Monte Carlo methods for computing Bayes factors: A comparative review. *J. Amer. Statist. Assoc.* **96** 1122–1132.
- HÄRDLE, W., LÜTKEPOHL, H. and CHEN, R. (1997). A review of nonparametric time series analysis. *International Statistical Review* **65** 49–72.
- HARRISON, P. J. and STEVENS, C. F. (1976). Bayesian forecasting. *J. Roy. Statist. Soc. Ser. B* **38** 205–247. [MR0655429](#)
- HASTIE, T. (1987). A closer look at the deviance. *Amer. Statist.* **41** 16–20. [MR0882765](#)
- IYENGAR, S. (2001). The analysis of multiple neural spike trains. In *Advances in Methodological and Applied Aspects of Probability and Statistics* (N. Bolakrishnan, ed.) 507–524. Taylor and Francis, New York. [MR1977526](#)
- KALMAN, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* **82** 35–45.
- KASS, R. E., VENTURA, V. and BROWN, E. (2005). Statistical issues in the analysis of neuronal data. *Journal of Neurophysiology* **94** 8–25.
- KEMP, K. W. (1957). Formulae for calculating the operating characteristic and the average sample number of some sequential tests. *J. Roy. Statist. Soc. Ser. B* **20** 379–386.

- KENDRICK, K. M., DA COSTA, A. P., LEIGH, A. E., HINTON, M. R. and PEIRCE, J. W. (2001). Sheep don't forget a face. *Nature* **414** 165–166.
- KIM, C. J. (1994). Dynamic linear models with Markov-switching. *J. Econometrics* **60** 1–22. [MR1247815](#)
- KOYAMA, S., PEREZ-BOLDE, L. C. and KASS, R. E. (2008). Approximate methods for state-space models: The Laplace–Gaussian filter. Submitted for publication.
- KUHNERT, P. M., MERGESEN, K. and TESAR, P. (2003). Bridging the gap between different statistical approaches: An integrated framework for modelling. *International Statistical Review* **71** 335–368.
- KULLBACK, S. (1997). *Information Theory and Statistics*. Dover, New York. [MR1461541](#)
- KULLBACK, S. and LEIBLER, R. A. (1951). On information and sufficiency. *Ann. Math. Statist.* **22** 79–86. [MR0039968](#)
- LINDLEY, D. (1956). On a measure of the information provided by an experiment. *Ann. Math. Statist.* **27** 986–1005. [MR0083936](#)
- LOADER, C. R. (1996). Change point estimation using nonparametric regression. *Ann. Statist.* **24** 1667–1678. [MR1416655](#)
- MAKEIG, S., WESTERFIELD, M., JUNG, T. P., ENGHOFF, S., TOWSEND, J., COURCHESNE, E. and SEJNOWSKI, T. J. (2002). Dynamic brain sources of visual evoked responses. *Science* **295** 690–694.
- MARJORAM, P., MOLITOR, J., PLAGNOL, V. and TAVARÉ, S. (2003). Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **100** 15324–15328.
- MCCULLOCH, R. E. (1988). Information and the likelihood function in exponential families. *Amer. Statist.* **42** 73–75. [MR0936686](#)
- MCCULLOCH, R. E. and TSAY, R. S. (1994). Statistical analysis of economic time series via Markov switching models. *J. Time Ser. Anal.* **15** 523–539. [MR1263893](#)
- MENG, X. L. (1994). Posterior predictive p -values. *Ann. Statist.* **22** 1142–1160. [MR1311969](#)
- MIRA, A. and PETRONE, S. (1996). Bayesian hierarchical nonparametric inference for change point problems. *Bayesian Statistics* **5** 693–703. [MR1425440](#)
- MULLER, H. G. (1992). Change-points in nonparametric regression analysis. *Ann. Statist.* **20** 737–761. [MR1165590](#)
- NEWTON, M. A. and RAFTERY, A. E. (1994). Approximate Bayesian inference by the weighted likelihood bootstrap. *J. Roy. Statist. Soc. Ser. B* **56** 3–56. [MR1257793](#)
- O'HAGAN, T. and FORSTER, J. (1999). *Kendall's Advanced Theory of Statistics* **2B**. Arnold, London, UK.
- OKATAN, M., WILSON, M. A. and BROWN, E. N. (2005). Analysing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput.* **17** 1927–1961.
- PAGE, E. S. (1954). An improvement to Wald's approximation for some properties of sequential tests. *J. Roy. Statist. Soc. Ser. B* **16** 136–139. [MR0065879](#)
- PAGE, E. S. (1955). A test for a change in a parameter occurring at an unknown point. *Biometrika* **42** 523–527. [MR0072412](#)
- PAGE, E. S. (1961). Cumulative sum charts. *Technometrics* **3** 1–9. [MR0119344](#)
- PLESSER, H. E. and GERSTNER, W. (2000). Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Comput.* **12** 367–384.
- RAO, R. P. N. (2005). Hierarchical Bayesian inference in networks of spiking neurons. In *Advances in NIPS* **17**. MIT Press, MA.
- RIGAT, F., DE GUNST, M. and VEN PELT, J. (2006). Bayesian modelling and analysis of spatio-temporal neuronal networks. *Bayesian Anal.* **1** 733–764. [MR2282205](#)
- ROBERT, C. P., CELEUX, G. and DIEBOLT, J. (1993). Bayesian estimation of hidden Markov chains: A stochastic implementation. *Statist. Probab. Lett.* **16** 77–83. [MR1208503](#)
- ROBINSON, P. M. (1983). Non-parametric estimation for time series models. *J. Time Ser. Anal.* **4** 185–208.

- RUBIN, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *Ann. Statist.* **12** 1151–1172. [MR0760681](#)
- SAN MARTINI, A. and SPEZZAFERRI, F. (1984). A predictive model selection criterion. *J. Roy. Statist. Soc. Ser. B* **46** 296–383. [MR0781890](#)
- SHUMWAY, R. H. and STOFFER, D. S. (1991). Dynamic linear models with switching. *J. Amer. Statist. Assoc.* **86** 763–769. [MR1147103](#)
- SMITH, A. C., LOREN, M. F., WYRTH, S., YANIKE, M., HU, D., KUBOTA, Y., GRAYBIEL, A. M., SUZUKI, W. A. and BROWN, E. M. (2004). Dynamic analysis of learning in behavioural experiments. *Journal of Neuroscience* **24** 447–461.
- SMITH, A. F. M. (1975). A Bayesian approach to inference about a change-point in a sequence of random variables. *Biometrika* **62** 407–416. [MR0381115](#)
- SMITH, A. F. M. and ROBERTS, G. O. (1993). Bayesian computations via the Gibbs sampler and related Markov chain Monte Carlo methods. *J. Roy. Statist. Soc. Ser. B* **55** 3–23. [MR1210421](#)
- SMITH, J. Q. (1990). Non-linear state space models with partially specified distributions on states. *J. Forecast.* **9** 137–149.
- SMITH, J. Q. (1992). A comparison of the characteristics of some Bayesian forecasting models. *International Statistical Reviews* **60** 75–85.
- SPIEGELHALTER, D. J., BEST, N. G., CARLIN, P. B. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit. *J. Roy. Statist. Soc. Ser. B* **64** 583–639. [MR1979380](#)
- SRINIVANSAN, L., EDEN, U. T., WILLSKY, A. S. and BROWN, E. N. (2006). A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural Comput.* **18** 2465–2494. [MR2256113](#)
- STEPHENS, D. A. (1994). Bayesian retrospective multiple-change-point identification. *App. Statist.* **43** 159–178.
- STONE, M. (1959). Application of a measure of information to the design and comparison of regression experiments. *Ann. Math. Statist.* **30** 55–70. [MR0106528](#)
- TIERNEY, L. (1994). Markov chains for exploring posterior distributions. *Ann. Statist.* **22** 1701–1762. [MR1329166](#)
- TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* **81** 84–86. [MR0830567](#)
- TRUCCOLO, W., EDEN, U. T., FELLOWS, M. R., DONOGHUE, J. P. and BROWN, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *Journal of Neurophysiology* **93** 1074–1089.
- WEST, M. (1986). Bayesian model monitoring. *J. Roy. Statist. Soc. Ser. B* **48** 70–78. [MR0848052](#)
- WEST, M. and HARRISON, P. J. (1986). Monitoring and adaptation in Bayesian forecasting models. *J. Amer. Statist. Assoc.* **81** 741–750.
- WEST, M. and HARRISON, P. J. (1997). *Bayesian Forecasting and Dynamic Models*, 2nd ed. Springer, New York. [MR1482232](#)
- WEST, M., HARRISON, P. J. and MIGON, H. S. (1985). Dynamic generalised linear models and Bayesian forecasting. *J. Amer. Statist. Assoc.* **80** 73–83. [MR0786598](#)
- WIRTH, S., YANIKE, M., LOREN, M. F., SMITH, A. C., BROWN, E. M. and SUZUKI, W. A. (2003). Single neurons in the monkey hippocampus and learning of new associations. *Science* **300** 1578–1584.

DEPARTMENT OF STATISTICS AND
CENTRE FOR ANALYTICAL SCIENCE
UNIVERSITY OF WARWICK
COVENTRY CV4 7AL
UK
E-MAIL: f.rigat@warwick.ac.uk
J.Q.Smith@warwick.ac.uk