Spline-based nonparametric regression for periodic functions and its application to directional tuning of neurons *

Cari G. Kaufman, Valérie Ventura and Robert E. Kass

October 6, 2004

Abstract

The activity of neurons in the brain often varies systematically with some quantitative feature of a stimulus or action. A well-known example is the tendency of the firing rates of neurons in the primary motor cortex to vary with the direction of a subject's arm or wrist movement. When this movement is constrained to vary in only two dimensions, the direction of movement may be characterized by an angle, and the neuronal firing rate can be written as a function of this angle. The firing rate function has traditionally been fit with a cosine, but recent evidence suggests that departures from cosine tuning occur frequently. We report here a new nonparametric regression method for fitting periodic functions and demonstrate its application to the fitting of neuronal data. The method is an extension of Bayesian Adaptive Regression Splines (BARS) and applies both to normal and non-normal data, including Poisson data, which commonly arise in neuronal applications. We compare the new method to a periodic version of smoothing splines and some parametric alternatives and find the new method to be especially valuable when the smoothness of the periodic function varies unevenly across its domain.

KEYWORDS: nonparametric regression, periodic functions, Bayesian Adaptive Regression Splines, smoothing splines, neuronal data

1 INTRODUCTION

A striking finding from the past forty years of neurophysiological research is that the firing rates of neurons in many parts of the brain are "tuned" to features of a stimulus or action. In a variety of experiments in which the voltage changes of indivdual neurons have been recorded using microelectrodes, neuronal activity has been shown to vary systematically with some externally-defined quantitative variable, e.g., References [1-7]. We consider here the relationship of neuronal activity in the primary motor cortex of a monkey to the monkey's two-dimensional hand movement direction, which we characterize by an angle θ between zero and 2π radians. Taking the firing rate of a neuron to be λ, the scientific problem is to describe how λ varies with θ. The statistical problem, therefore, is to define a periodic functional form for $\lambda \equiv \lambda(\theta)$, the so-called "directional tuning curve," and to fit it to data.

In the neurophysiology literature, a standard approach to fitting directional tuning curves to motor cortical data for two dimensional movements is as follows [8]. If θ is the direction of movement in a given time interval of length T, let $Y(\theta)$ be the number of times the neuron fires during that interval. The count $Y(\theta)$ is assumed to be normally distributed with mean $T \cdot \lambda(\theta)$ with

$$\lambda(\theta) = \mu + \alpha \cos(\theta - \theta_0) \tag{1}$$
$$= \mu + \beta_1 \cos(\theta) + \beta_2 \sin(\theta). \tag{2}$$

Linear regression is used to fit the resulting *cosine tuning function*. The parameter θ_0 in (1) is commonly referred to as the neuron's "preferred direction" and is the key component in the pioneering method for predicting movement based on neuronal activity, the Population Vector Algorithm [9, 10]. However, the linear parameterization in (2) is easier to fit and can easily be used to derive an estimate of θ_0 if needed.

There are two difficulties with this approach. First, the counts $Y(\theta)$ may not be even approximately normally distributed and may not have homogeneous variances with respect to θ. Second, there is the more substantial problem that, in some cases, the cosine tuning function does not fit adequately. One approach to solving the latter problem is to use more general parametric forms, and these may be used successfully for many data sets [11]. A still more flexible approach is to use nonparametric methods. In this paper we describe the fitting of periodic functions by spline-based generalized nonparametric regression.

Figure 1 illustrates several of the methods discussed here using data from two neurons recorded in the primary motor cortex of a rhesus monkey performing a movement task (described in Section 4). The monkey made wrist movements to eight targets located radially in a plane; hence, θ takes eight values. It is clear that the firing rate λ varies systematically with the angle of movement θ. The data from the first neuron are approximately Poisson-distributed, with sample means and variances that fluctuate together from nearly zero to over 100 spikes per second, whereas the data from the second neuron are approximately normally distributed with approximately constant variance. For both neurons the cosine model fits relatively well, although it clearly misses detailed features of the tuning curve. Attempts to improve the fit parametrically as suggested by Amirikian and Georgopoulos [11] are described in Section 2.1, but they are only marginally effective here.

Spline-based generalized nonparametric regression has recently been reviewed by Hansen and Kooperberg [12]. Spline fitting requires regularization in some form. Two general strategies are to penalize the likelihood, thereby shrinking the spline basis coefficients toward zero, or to reduce the number of basis elements by selecting locations for a relatively small number of knots. The former, used by smoothing splines, is effective in a wide range of settings [13].

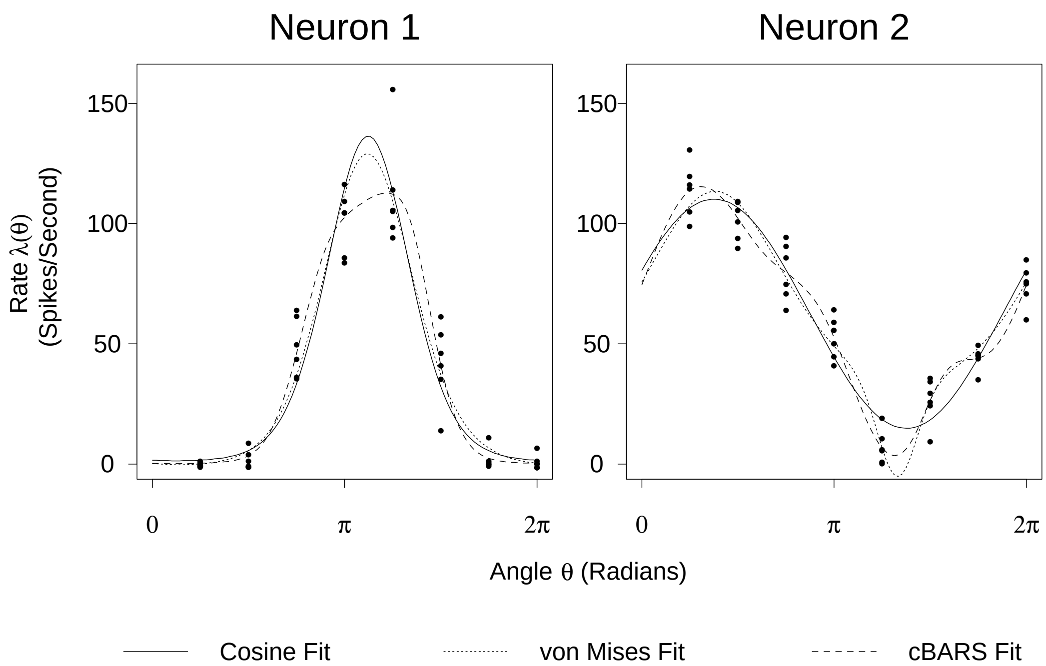

Figure 1: Fits to activity of two motor cortex neurons. Each data point represents the observed firing rate of a neuron in the motor cortex of a monkey during one repetition of a wrist movement to a particular target at angle θ. Firing rates are shown in spikes per second and the recording interval was 200 milliseconds. The cosine fits use the cosine function in (2) and the von Mises fits use the parametric forms in (3) for Neuron 1 and (4) for Neuron 2. The cosine and von Mises parametric regressions were fit by numerically maximizing the likelihood function, based on a Poisson density for Neuron 1 and a normal density for Neuron 2. Poisson and normal versions of the cBARS algorithm provide better fits. (See Section 4).

However, global penalization can be too crude to accommodate sharp local fluctuations in the function being fitted, which has led many researchers to devise algorithms for selecting a small set of knot locations [12]. To search through large spaces of alternative models (corresponding to alternative sets of knot locations), Bayesian methods using Markov chain Monte Carlo (MCMC) techniques have been shown to be effective [12, 14, 15, 17, 18]. Among these methods is Bayesian Adaptive Regression Splines (BARS), proposed by DiMatteo, Genovese and Kass [14]. Because BARS has worked well in many applications [19, 21, 22], and it has been shown to be clearly superior to competitors in some examples [14], we developed a periodic version of BARS, which we call cBARS ("c" for circular), to deal with periodic explanatory variables such as the angle of movement θ. The first purpose of this paper is to report our implementation and application of cBARS.

As mentioned, BARS and related methods perform better than global penalization methods when the function being fitted has sharp variations. When the smoothness of the function is more homogeneous, this advantage disappears. We compared cBARS to a periodic version of smoothing splines to determine whether this result also holds in the periodic case. An extension of smoothing splines to periodic functions, which we will refer to as circular smoothing splines, has been described by Wahba [24] and implemented by Wang, Ke and Brown [25]. In addition, Wang and Ke [26] have made publicly available their ASSIST software, which may be used for a wide range of smoothing spline methods, including circular smoothing splines. The second purpose of this paper, then, is to provide some comparison of cBARS with circular smoothing splines in fitting directional tuning curves to firing rate data from motor cortical neurons.

Section 2 describes parametric and nonparametric methods for fitting periodic regression models. Section 3 compares the efficiency of these methods using simulated data. Before we conclude in Section 5, Section 4 returns to the experimental data shown in Figure 1 and gives further details about the experimental design and the application of each method to these data.

2 METHODS

In the case of a designed experiment in which a monkey is trained to make straight-line movements to specific targets (as in the standard two-dimensional "center-out" task, e.g., [8, 6]), then although the movement time T may be on the order of 200–500 milliseconds, we may take the direction of movement θ to be constant throughout this interval. When the monkey makes unrestrained movements in two-dimensions, the angle of movement θ will be measured along with the firing count $Y(\theta)$, and T usually will need to be smaller, on the order of 50 ms, say, to capture temporal variations in movement direction. Here, we assume the data come from a designed experiment, although the methods we describe could also be applied to data involving free arm movements.

Let $Y_i(\theta_j)$ be the number of times a neuron fires during the ith trial (repetition of the experiment) at angle θ_j, where $j = 1, \ldots, J$, and J is the number of distinct angles. Although neuron spike trains typically follow non-Poisson point processes, neuronal spike count data in large intervals are often reasonably well approximated by the Poisson distribution [27, 28, 29]. Sometimes, as for the second neuron in Figure 1, the normal distribution appears to be a better approximation. The methods we discuss in this paper are relatively easy to modify to handle various distributional assumptions.

2.1 Parametric Models

The cosine tuning function in (1) may be generalized in various ways. Amirikian and Georgopoulos [11] describe a class of functions based on the form of the von Mises probability density function. For example, the von Mises-like functions used to fit the data in Figure 1 are

$$\lambda(\theta) = \mu + \beta \exp(\kappa \cos(\theta - \tau + \eta \cos(\theta - \tau))) \quad (3)$$

for the first neuron, and

$$\lambda(\theta) = \mu + \beta_1 \exp(\kappa_1 \cos(\theta - \tau_1)) + \beta_2 \exp(\kappa_2 \cos(\theta - \tau_2)) \quad (4)$$

for the second neuron. The function in (3) includes parameters controlling the baseline firing rate (μ), the amplitude (β), width (κ) and location (τ) of the mode, and the skewness about the mode (η). The function in (4) includes parameters corresponding to two modes and does not allow for skewness. We chose the parametric form in (3) for the first neuron in Figure 1 because by eye the firing rates corresponding to the middle two angles (π and $5\pi/4$ radians) seemed to suggest skewness. We chose the parametric form in (4) for the second neuron because by eye it appeared that the width and amplitude of the mode were not equal to the width and amplitude of the anti-mode. Some other von Mises-like functions are suggested in [11]. We explored several of these, but none improved the fit.

To fit tuning curves like (3) or (4), the likelihood function may be maximized numerically using nonlinear optimization software. When the data are normally distributed, this is equivalent to minimizing the sum of squared errors.

2.2 Circular Smoothing Splines

Nonparametric methods are sometimes preferable to the parametric methods of the previous section. They are more flexible and may be more convenient, for example, in the case that a parametric approach requires delicate nonlinear optimization. One nonparametric option is smoothing splines, which we describe briefly in this section. Another option is a regression spline, and in the next section we discuss a Bayesian method for choosing the knots.

In the case of the standard nonparametric regression model $Y_i(\theta_j) = \lambda(\theta_j) + \epsilon_{ij}$, with independent and normally distributed errors, the smoothing spline estimate of λ is the twice continuously differentiable function $\hat{\lambda}$ which minimizes

$$\frac{1}{IJ}\sum_{i=1}^{I}\sum_{j=1}^{J}(y_i(\theta_j) - \hat{\lambda}(\theta_j))^2 + \nu \int_0^1 (\hat{\lambda}''(\theta))^2 d\theta. \quad (5)$$

The global smoothing parameter ν determines the trade-off between goodness of fit (in terms of mean squared error) and smoothness (in terms of the integrated squared second derivative of $\hat{\lambda}$) [23]. It may be chosen by any of several existing data-based methods, such as generalized cross validation, generalized maximum likelihood and unbiased risk estimation; given a particular value of ν, the unique minimizer of (5) is a natural cubic spline with knots at the distinct values of θ_j and spline coefficients that are shrunk according to ν [24].

When the regression function λ is restricted to be in the collection of all periodic functions on $[0, 1]$ of the form

$$\lambda(\theta) = \sqrt{2}\sum_{k=1}^{\infty} a_k \cos 2\pi k\theta + \sqrt{2}\sum_{k=1}^{\infty} b_k \sin 2\pi k\theta \quad (6)$$

with $\sum_{k=1}^{\infty}(a_k^2 + b_k^2)(2\pi k)^{2m} < \infty$, the solution to (5) is a periodic spline basis expansion

$$\lambda(x) = \sum_{i=1}^{N} \alpha_i K(x, x_i), \quad (7)$$

where $K(x, x_i)$ is the reproducing kernel

$$K(x, x_i) = \sum_{k=1}^{\infty} \frac{2}{(2\pi k)^4} \cos(2\pi k(x - x_i)), \tag{8}$$

[24, 25].

For non-normal data, the negative log likelihood is used in place of the mean squared error in (5), but the solution is the same as in (7). Wang and Ke [26] have written software for R and S-PLUS called ASSIST that allows one to fit a variety of smoothing spline models, including circular smoothing splines, to data with various distributions (Bernoulli, binomial, Poisson, gamma and normal). Standard options are also available to determine the smoothing parameter ν in (5).

2.3 Circular BARS

Bayesian Adaptive Regression Splines (BARS, DiMatteo, Genovese, and Kass [14]) is a Bayesian method for fitting regression splines with an unknown number of knots and knot locations. Di Matteo et al. showed that BARS is sometimes more efficient than other leading spline-based alternatives, such as the Bayesian version of Multivariate Adaptive Regression Splines (Bayesian MARS) of Denison, Mallick and Smith [15] and the Spatially Adaptive Regression Splines (SARS) of Zhou and Shen [16]. Some additional interpretation, contrasting BARS with Bayesian MARS, is given by Kass and Wallstrom [17]. Because BARS is specifically designed to allocate additional knots to regions of rapid functional variation, it is especially effective in fitting functions that vary slowly throughout most of their domain but have one or more sudden jumps or peaks. This behavior is observed in many neuronal firing rate intensity functions [19]. In this section we describe the model on which BARS is based and the algorithm used to fit it. We give details of the implementation when the underlying function is assumed to be periodic.

To make inference about a function $\lambda(\theta)$ (here, the firing rate), BARS fits the spline-based generalized nonparametric regression model for the spike counts $Y_i(\theta_j)$

$$Y_i(\theta_j) \sim p(y_i \mid \lambda(\theta_j)), \tag{9}$$

where λ is a spline having knots at unknown locations $\xi = (\xi_1, \ldots, \xi_k)$, and p is a specified distribution. For a given knot set ξ, we may write $\lambda(\theta)$ in terms of basis functions $b_{\xi,h}(\theta)$ as $\lambda(\theta) = \sum_h b_{\xi,h}(\theta)\beta_{\xi,h}$, where the index of summation h depends on the spline basis being used; cubic splines and the natural spline basis have been used in most applications of BARS. However, to allow BARS to fit periodic functions, we use the periodic spline basis defined by the reproducing kernel (8). Suppose the knots are $\xi_1, \ldots, \xi_k$ in $[0, 2\pi]$. Then a periodic spline basis of order 1 has k basis functions, excluding an intercept, and the i^{th} basis element is a cosine with maximum value occurring at the knot ξ_i.

For a given knot set ξ, model (9) poses a relatively easy estimation problem; for exponential-family responses (such as Poisson) it becomes a generalized linear model. The difficult problem is to determine ξ. BARS begins by placing prior distributions on the number and locations of the knots, generally Poisson and uniform, respectively. An MCMC algorithm is then used to generate draws from the posterior distribution on the knot sets.

Key features of the MCMC implementation of BARS include (i) a reversible-jump chain [20] on ξ after integrating the marginal density

$$p(y|\xi) = \int p(y \mid \beta_\xi, \xi)\pi(\beta_\xi \mid \xi)d\beta_\xi \tag{10}$$

for observations $y = (y_1, \ldots, y_n)$, (ii) continuous proposal distributions for ξ, in contrast to the proposal distributions used by Bayesian MARS, which allow knots only at the data points, and (iii) a locality heuristic for the proposals that attempts to place potential new knots near existing knots. Items (ii) and (iii) together allow proposal knots to be placed close to one another, which is advantageous when there is a sudden jump or peak in the function. For the circular version of BARS, we must also ensure that the proposal distribution for new knot locations in the MCMC algorithm is also periodic. We proceed as follows: a current knot ξ_i is drawn randomly from the set of current knots. A proposal knot ξ^c is given by

$$\xi^c = [\xi_i + 2\pi(X - 0.5)],\ X \sim Beta(\alpha, \alpha), \tag{11}$$

where the operator [ ] signifies modulo on $[0, 2\pi]$, and α is an adjustable tuning parameter in the MCMC algorithm controlling how close the proposal knot ξ^c tends to be to the current knot ξ_i. Changing α changes the proportion of proposal knots which are accepted and may be adjusted to improve the mixing of the chain. Note that $Beta(\alpha, \alpha)$ has support on $[0, 1]$ and is symmetric about 0.5, so that the distribution of ξ^c is periodic on $[0, 2\pi]$ and symmetric about the current knot ξ. Clearly (11) can easily be scaled for periodic domains other than $[0, 2\pi]$.

For each sample $\xi^{(g)}$ drawn from the posterior distribution of ξ, a sample of regression spline coefficients $\beta_\xi^{(g)}$ is obtained via generalized linear regression. The posterior distribution of (ξ, β_ξ) is not of primary interest, however; what we want is a sample from the posterior distribution of $\lambda(\tilde{\theta})$ for any specified $\tilde{\theta}$ in the range of the function, which can be obtained by calculating $\lambda^{(g)}(\tilde{\theta}) = \sum b_{\xi,h}^{(g)}(\tilde{\theta})\beta_{\xi,h}^{(g)}$. Either the posterior mean or median across knot sets may be used as a point estimate of $\lambda(\tilde{\theta})$, and (pointwise) credible intervals may likewise be obtained using sample quantiles.

3 SIMULATION STUDY

We conducted a small simulation study to compare cBARS with circular smoothing splines and cosine regression for the four periodic functions shown in Figure 2. We did not include a comparison with the von Mises-like functions of Section 2.1 because the nonlinear optimization needed to fit such functions is difficult to automate. The first test function is the cosine function

$$\lambda_1(\theta) = 50 + 25\cos(\theta - \pi/2). \tag{12}$$

The next two test functions are the actual cBARS fits to neurons from the experiment described in Section 4. The fourth function has a sharp peak, with

$$\lambda_4(\theta) = 50 + 25\sin(\theta - 1.2\pi) + 75\exp(-10(15(\theta - \pi)/2\pi)^2). \tag{13}$$

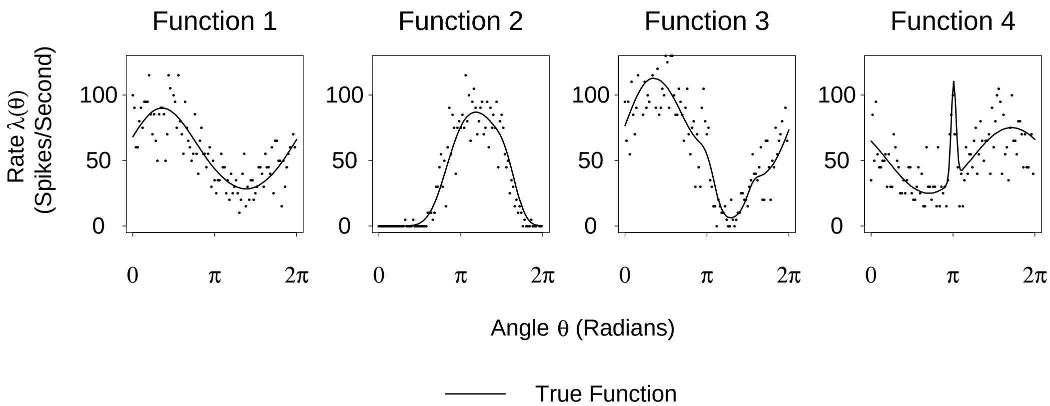

Figure 2: Four different test functions were used in the simulation study (solid curves). For each function, 1000 datasets were generated, each of size 100 data points. These data points were drawn from a Poisson distribution with mean function corresponding to one of the four test functions. Points in the figure are the data drawn in the first iteration of the simulation.

For each test function $\lambda_k(\theta)$, our efficiency results are based on 1000 datasets of size 100, simulated from a Poisson model with mean $\lambda_k(\theta)$ and with θ taking 100 equally spaced values in $[0, 2\pi]$. A simulated data set for each function is also shown in Figure 2.

The cosine model was fit to the simulated data via the `glm` function in S-PLUS, specifying Poisson data and either the identity or log link function. We tried both link functions for each test function except $\lambda_2(\theta)$, for which the identity link is problematic, because the many zero data values often yield predicted rates that are negative. The smoothing spline was fit via the `ssr` function in ASSIST, specifying Poisson data and with smoothing parameter determined by generalized maximum likelihood. Finally, cBARS was fit by modifying recent C code implementing BARS, as described in Section 2.3. [1] The prior distributions on the knots took their number to be Poisson distributed with mean five and their locations to be uniformly distributed on $[0, 2\pi]$. The adjustable parameter α in the distribution for proposal knots (11) was 25, which means that the proposal knot ξ^c will be within approximately $\pi/4$ radians of the current knot ξ with probability 95%. We found this choice of α allowed the MCMC algorithm to mix well in the types of data we examined. For each simulated dataset, a burn-in period of 100 iterations was used; the next 1000 iterations were treated as a sample from the posterior distribution, and the fitted curves according to each of the knot sets in this sample were averaged pointwise to create the fit to that dataset.

The integrated squared error between the true curve and a fit is the squared area between the two curves; it was approximated on a grid of size 100 and averaged across the 1000 simulated datasets to produce the mean integrated squared errors (MISE) reported in Table 1.

For the cosine test function $\lambda_1(\theta)$, the cosine parametric model fits well, but for the other

[1] The cBARS code may be obtained via www.stat.cmu.edu/˜ kass.

Table 1: Mean integrated squared errors (MISE) for each curve fitting method and four periodic test functions. Simulation errors were all less than .0013.

Simulation Test Function	1	2	3	4
Cosine (Log Link)	0.53	3.83	3.67	5.74
Cosine (Identity Link)	0.30	NA	1.88	5.33
Circular Smoothing Spline	0.50	0.90	1.40	3.84
Circular BARS	0.47	0.59	1.12	1.80

three functions the nonparametric methods improve the fits substantially. In every case cBARS performs better than circular smoothing splines. As expected, for $\lambda_2(\theta)$ and $\lambda_4(\theta)$, for which the functional variation is non-uniform, circular smoothing splines perform considerably worse than cBARS: the MISE for circular smoothing splines are, respectively, 50% and 110% greater than those for cBARS. Figure 3 displays typical fitted curves for $\lambda_2(\theta)$ and $\lambda_4(\theta)$.

4 APPLICATION TO EXPERIMENTAL DATA

We applied the methods of Section 2 to data from an experiment involving wrist movement reported by Kakei, Hoffman and Strick [6]. In this experiment, a monkey was trained to move its wrist in one of eight directions located radially about a circle, while its forearm was restrained in either a pronated (palm down), a supinated (palm up), or a middle (palm inward) position. Meanwhile, the firing times of single neurons in the primary motor cortex were recorded. This experimental paradigm allowed the researchers to dissociate the effects of the direction of movement and the specific muscle needed to make the movement. First, tuning curves were fit to the firing rates for each forearm position (in [6], modeled with a cosine), then the direction at which the fitted curve achieved its maximum (the "preferred direction") was compared across forearm positions. Neurons with only small changes in preferred direction were dubbed "extrinsic like," because their tuning properties depended only on movement direction, while neurons with large shifts were dubbed "muscle like," because their tuning properties changed with the muscle needed to make the movement. This study gave evidence that variables defined on both extrinsic and intrinsic (physiological) coordinate systems are encoded by neurons in the primary motor cortex, addressing a longstanding scientific controversy.

In the Introduction we noted that cBARS provides better fits than parametric models to the data from the two neurons shown in Figure 1. The data for Neuron 1 appeared to be approximately Poisson distributed and we fit Poisson regression models by numerically maximizing the likelihood using the `nlminb` function in S-PLUS. (We used the log link to fit the cosine model and the identity link to fit the von Mises model.) The data for Neuron 2 were more closely approximated by a normal distribution with constant variance, and so to find the maximum likelihood estimates numerically, we used the least squares criterion. Note that to fit (4), we

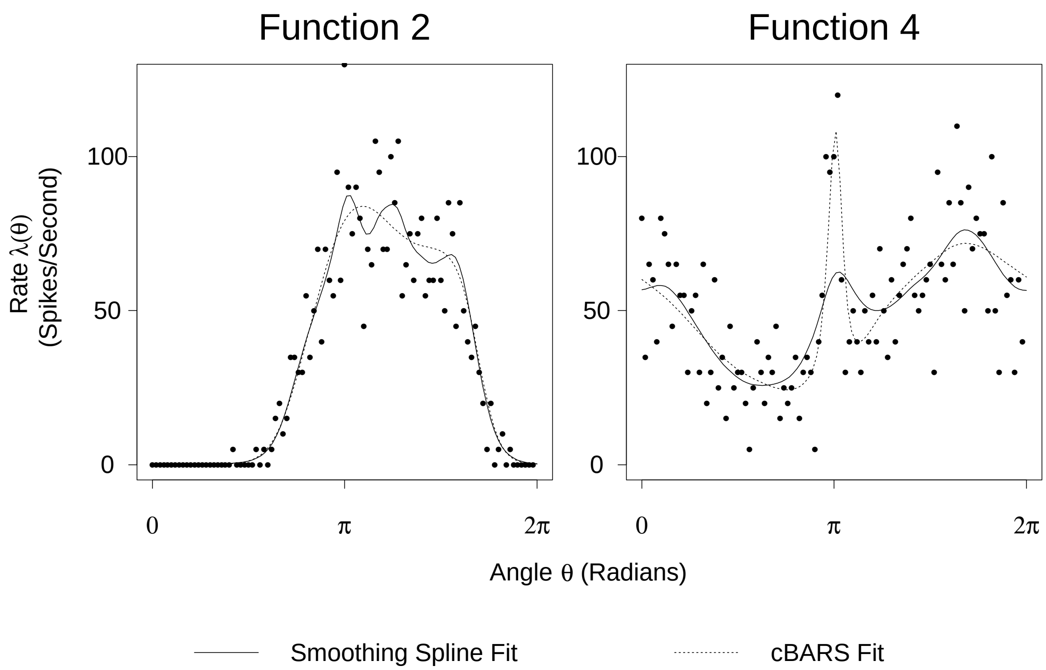

Figure 3: Fits from circular smoothing splines and cBARS for a single data set in the simulation study with periodic test functions 2 and 4. Both functions have non-uniform variation across their domains, which leads to better performance by cBARS.

constrained τ_1 and τ_2 to be of opposite signs to stabilize the optimization. For Neuron 1 the Poisson deviances for the Cosine, von Mises, circular smoothing spline and cBARS fits were, respectively, 57.5, 57.1, 37.5 and 36.0. For Neuron 2 the mean squared errors were 226, 161, 132 and 130. Circular smoothing splines were nearly identical to those of cBARS in terms of fitted values and deviances and are therefore omitted from Figure 1.

5 DISCUSSION

We have extended BARS methodology to the fitting of periodic functions and applied the new algorithm, cBARS, to the fitting of directional tuning curves from neuronal data. The new method uses a periodic spline basis, as in the approach of [25], which applies smoothing splines rather than regression splines with empirically-defined knots. BARS smooths data adaptively and thereby is able to produce fitted curves that vary slowly in one part of the domain but rapidly in another. So, too, can cBARS adapt to sudden jumps or peaks in a periodic function. Our simulation study showed that cBARS can provide better fits than circular smoothing splines. On the other hand, when the variation in the periodic function is nearly uniform throughout its domain, circular smoothing splines are likely to produce fits that are nearly the same as those from cBARS, as for the neuronal data in Figure 1.

Cosine tuning functions are widely used in the analysis of neurophysiological data, and in algorithms for brain-controlled robotic devices [30]. Recent evidence using 16 directions of movement has indicated that directional tuning curves often deviate from cosine form [11]. The experiment described in Section 4 involved only eight directions, and therefore, departures from cosine tuning are less apparent. Nonetheless, as Figure 1 shows, they are sometimes still discernible. The extent to which they are important neurophysiologically has yet to be determined. There is already some evidence that allowing for non-cosine tuning curves may improve brain-controlled robotic devices [31]. The work reported here indicates that incorporation of periodic nonparametric regression methods into algorithms for such devices may be desirable. These methods may also be useful in other applications where the data is inherently periodic, such as biological monitoring.

References

[1] McNaughton BL, Barnes CA, O'Keefe J. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely moving rats. *Experimental Brain Research* 1983; **52**(1):41–49.

[2] Andersen RA, Essick GK, Siegel RM. Encoding of spatial location by posterior parietal neurons. *Science* 1985; **230**(4724):456–458.

[3] Schwartz AB, Kettner RE, Georgopoulos AP. Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *The Journal of Neuroscience* 1988; **8**(8):2913–2927.

[4] Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 1993; **261**(5124):1055–1058.

[5] Lewis JE, Kristan WB. Representation of touch location by a population of leech sensory neurons. *Journal of Neurophysiology* 1998; **80**(5):2584–2592.

[6] Kakei S, Hoffman DS, Strick PL. Muscle and movement representations in the primary motor cortex. *Science* 1999; **285**(5436):2136–2139.

[7] Moran DW, Schwartz AB. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology* 1999; **82**(5):2676–2692.

[8] Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience* 1982; **2**(11):1527–1537.

[9] Georgopoulos AP, Schwartz AB, Kettner RE. Neuronal population coding of movement direction. *Science* 1986; **233**(4771):1416–1419.

[10] Georgopoulos AP, Kettner RE, Schwartz AB. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by neuronal population. *The Journal of Neuroscience* 1988; **8**(8):2928–2937.

[11] Amirikian B, Georgopoulos AP. Directional tuning profiles of motor cortical cells. *Neuroscience Research* 2000; **36**(1):73–79.

[12] Hansen MH, Kooperberg, C. Spline adaptation in extended linear models. *Statistical Science* 2002; **17**(1):2–20.

[13] Ruppert D, Wand M, Carrol R. *Nonparametric and semiparametric regression modeling.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2003.

[14] DiMatteo I, Genovese CR, Kass RE. Bayesian curve-fitting with free-knot splines. *Biometrika* 2001; **88**(4):1055–1071.

[15] Denison DGT, Mallick BK, Smith AFM. Automatic Bayesian curve fitting. *Journal of the Royal Statistical Society, Series B, Methodological* 1998; **60**(2):333–350.

[16] Zhou S, Shen X. Spatially Adaptive Regression Splines and accurate knot selection schemes. *Journal of the American Statistical Association* 2001; **96**(453):247–259.

[17] Kass RE, Wallstrom G. Comment on "Spline adaptation in extended linear models." *Statistical Science* 2002; **17**(1):24–29.

[18] Smith M, Kohn R. Nonparametric regression using Bayesian variable selection. *Journal of Econometrics* 1996; **75**(2):317–344.

[19] Kass RE, Ventura V, Cai C. Statistical smoothing of neuronal data. *NETWORK: Computation in Neural Systems* 2003; **14**(1):5–15.

[20] Green, PJ. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 1995; **82**(4):711–732.

[21] Wallstrom G, Kass RE, Miller A, Cohn JF, Fox NA. Correction of ocular artifacts in EEG using Bayesian Adaptive Regression Splines. In *Case studies in Bayesian statistics, Vol VI,* Gastonis CG, Kass RE, Carriquiry A, Higdon D, Pauler D, Verdinelli I (eds). Springer-Verlag: New York, 2002.

[22] Zhang X, Roeder K, Wallstrom G, Devlin B. Integration of association statistics over genomic regions using Bayesian Adaptive Regression Splines. *Human Genomics* 2003; **1**(1):20–29.

[23] Hastie T, Tibshirani R, Friedman J. *The elements of statistical learning: Data mining, inference, and prediction.* Springer-Verlag: New York, 2001.

[24] Wahba G. *Spline models for observational data.* CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM: Philadelphia, 1990.

[25] Wang Y, Ke C, Brown MB. Shape invariant modeling of circadian rhythms with random effects and smoothing spline ANOVA decompositions. *Biometrics* 2003; **59**(4):804-812.

[26] Wang Y, Ke C. ASSIST: A suite of S functions implementing spline smoothing techniques. http://www.pstat.ucsb.edu/faculty/yuedong/software.html[3 July 2004].

[27] Ventura V, Carta R, Kass RE, Gettner SN, Olson CR. Statistical analysis of temporal evolution in single-neuron firing rates. *Biostatistics* 2001; **3**(1):1–20.

[28] Kass RE, Ventura V. A spike-train probability model. *Neural Computation* 2001; **13**(8):1713–1720.

[29] Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation* 2002; **14**(2):325–46.

[30] Brockwell AE, Rojas A, Kass, RE. Recursive Bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology* 2004; **91**(4):1899–1907.

[31] Gao Y, Black MJ, Bienenstock E, Wu W, Donoghue JP. A quantitative comparison of linear and non-linear models of motor cortical activity for the encoding and decoding of arm motions. In *1st International IEEE/EMBS conference on neural engineering*, Capri, Italy, March 20-22, 2003; 189–192.