



**University of  
Zurich**<sup>UZH</sup>

**Zurich Open Repository and  
Archive**

University of Zurich  
Main Library  
Strickhofstrasse 39  
CH-8057 Zurich  
[www.zora.uzh.ch](http://www.zora.uzh.ch)

---

Year: 2016

---

## A unified framework of constrained regression

Hofner, Benjamin ; Kneib, Thomas ; Hothorn, Torsten

**Abstract:** Generalized additive models (GAMs) play an important role in modeling and understanding complex relationships in modern applied statistics. They allow for flexible, data-driven estimation of covariate effects. Yet researchers often have a priori knowledge of certain effects, which might be monotonic or periodic (cyclic) or should fulfill boundary conditions. We propose a unified framework to incorporate these constraints for both univariate and bivariate effect estimates and for varying coefficients. As the framework is based on component-wise boosting methods, variables can be selected intrinsically, and effects can be estimated for a wide range of different distributional assumptions. Bootstrap confidence intervals for the effect estimates are derived to assess the models. We present three case studies from environmental sciences to illustrate the proposed seamless modeling framework. All discussed constrained effect estimates are implemented in the comprehensive R package `mboost` for model-based boosting.

DOI: <https://doi.org/10.1007/s11222-014-9520-y>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-103599>

Journal Article

Accepted Version

Originally published at:

Hofner, Benjamin; Kneib, Thomas; Hothorn, Torsten (2016). A unified framework of constrained regression. *Statistics and Computing*, 26(1-2):1-14.

DOI: <https://doi.org/10.1007/s11222-014-9520-y>

# A Unified Framework of Constrained Regression

Benjamin Hofner · Thomas Kneib · Torsten Hothorn

Received: date / Accepted: date

**Abstract** Generalized additive models (GAMs) play an important role in modeling and understanding complex relationships in modern applied statistics. They allow for flexible, data-driven estimation of covariate effects. Yet researchers often have a priori knowledge of certain effects, which might be monotonic or periodic (cyclic) or should fulfill boundary conditions. We propose a unified framework to incorporate these constraints for both univariate and bivariate effect estimates and for varying coefficients. As the framework is based on (functional gradient descent) boosting methods, variables can be selected intrinsically, and effects can be estimated for a wide range of different distributional assumptions. We present three case studies from environmental sciences. The first on air pollution illustrates the use of monotonic and periodic effects in the context of an additive Poisson model. The second case study highlights the use of bivariate cyclic splines to model activity profiles of roe deer. The third case study demonstrates how to estimate the complete conditional distribution function of deer–vehicle collisions with the help of monotonicity constraints, and a cyclic constraint is considered for the seasonal variation of collision numbers. All discussed constrained effect estimates are im-

plemented in the comprehensive R package **mboost** for model-based boosting.

**Keywords** boosting · bivariate constraints · cyclic constraints · generalized additive models · monotonic constraints · periodic effects

## 1 Introduction

When statistical models are used, certain assumptions are made, either for convenience or to incorporate the researchers' assumptions on the shape of effects, e.g., because of prior knowledge. A common, yet very strong assumption in regression models is the linearity assumption. The effect estimate is constrained to follow a straight line. Despite the widespread use of linear models, it often may be more appropriate to relax the linearity assumption.

Let us consider a set of observations  $(y_i, \mathbf{x}_i^\top)$ ,  $i = 1, \dots, n$ , where  $y_i$  is the response variable and  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(L)})$  consists of  $L$  possible predictors of different nature, such as categorical or continuous covariates. To model the dependency of the response on the predictor variables, we consider a structured additive regression (STAR) model, where  $\mathbb{E}(y|\mathbf{x}) = h(\eta(\mathbf{x}))$  with (known) response function  $h$ , and a *structured additive predictor*  $\eta(\mathbf{x})$  of the form

$$\eta(\mathbf{x}) = \beta_0 + \sum_{l=1}^L f_l(\mathbf{x}) \quad (1)$$

is used. The functions  $f_l(\cdot)$  depend on one or more predictors contained in  $\mathbf{x}$ . Examples include linear effects, categorical effects, and smooth effects. More complex models with functions that depend on multiple variables such as random effects, varying coefficients, and

---

B. Hofner  
Institut für Medizininformatik, Biometrie und Epidemiologie,  
Friedrich-Alexander-Universität Erlangen-Nürnberg,  
Waldstraße 6, D-91054 Erlangen, Germany.  
E-mail: benjamin.hofner@fau.de

T. Kneib  
Lehrstuhl für Statistik, Georg-August-Universität Göttingen,  
Platz der Göttinger Sieben 5, D-37073 Göttingen, Germany.

T. Hothorn  
Institut für Sozial- und Präventivmedizin,  
Abteilung Biostatistik, Universität Zürich,  
Hirschengraben 84, CH-8001 Zürich, Switzerland.

bivariate effects, can be expressed in this framework as well (for details, see [Fahrmeir et al., 2004](#)). If all generic functions represent smooth effects, then a special case of STAR models—the generalized additive model (GAM) as introduced by [Hastie and Tibshirani \(1986, 1990\)](#)—can be considered.

A convenient way to fit STAR models is given by component-wise functional gradient descent boosting ([Bühlmann and Yu, 2003](#)), which minimizes the empirical risk function (e.g., the negative log-likelihood). The boosting algorithm is especially attractive due to its intrinsic variable selection properties ([Kneib et al., 2009](#); [Hofner et al., 2011a](#)) and the ease of combining a wide range of modeling alternatives in a single model specification. Furthermore, a single estimation framework can be used for a very wide range of distributional assumptions (or even for distribution-free approaches, such as quantile regression; see [Fenske et al., 2011](#)). In the context of conditional transformation models (?), complete distribution functions can be modeled by boosting.

STAR models offer great flexibility but typically result in smooth yet otherwise unconstrained effect estimates  $\hat{f}_i$ . In the first case study presented, we model the effect of air pollution on daily mortality in São Paulo (Section 5.1). Additionally, we control for environmental conditions (temperature, humidity) and model both the seasonal pattern and the long-term trend. Furthermore, we consider the effect of the pollutant of interest, SO<sub>2</sub>. In modeling the seasonal pattern of mortality related to air pollution, the effect should be continuous over time, and huge jumps for effects only one day apart would be unrealistic. Thus, the first and last days of the year should be continuously joined. However, smooth effect estimates are known to be unstable in the boundaries because less information is available there than in the middle of the domain. Hence, we use smooth functions with a cyclic constraint (Section 3.2). This has two effects. First, it allows us to fit a plausible model as we avoid jumps at the boundaries. Second, the estimation at the boundaries is stabilized as we exploit the cyclic nature of the data.

From a biological point of view, it seems reasonable to expect an increase in mortality with increasing concentration of the pollutant SO<sub>2</sub>. Linear effects are monotonic but do not offer enough flexibility in this case because they do not allow for changes in the effect of the pollutant on the number of deaths that might occur, for example, owing to a safe threshold or saturation effects at higher concentrations. Smooth effects, on the other hand, offer more flexibility, but monotonicity might be violated. To bridge this gap, smooth monotonic effects can be used. Recently, [Hofner et al. \(2011b\)](#) proposed a framework based on boosting to incorporate

both unconstrained smooth effects and monotonicity-constrained smooth effects (Sec. 3.3).

In a second case study, we aim at modeling the activity of roe deer in Bavaria, Germany, given environmental conditions, such as temperature, precipitation, and depth of snow; animal-specific variables, such as age and sex; and a temporal component. The latter reflects the animals’ day/night rhythm as well as seasonal patterns. We model the temporal effect as a smooth bivariate effect as the days change throughout the year, i.e., the solar altitude changes in the course of a day and with the seasons. Cyclic constraints for both variables (time of the day and calendar day) should be used. Hence, we have a bivariate periodic effect  $f(t_{\text{hours}}, t_{\text{days}})$  (Section 4.2). As male and female animals differ strongly in their temporal activity profiles, we additionally use sex as a binary effect modifier  $f(t_{\text{hours}}, t_{\text{days}})I_{(\text{sex} = \text{male})}$ , i.e., we have a varying coefficient surface with a cyclicity constraint for the smooth bivariate effect. Additionally, the effects of environmental variables are allowed to smoothly vary over time but are otherwise unconstrained.

In a third case study, we go beyond a model for the mean activity by modeling the conditional distribution of a surrogate of roe deer activity: the number of deer–vehicle collisions per day. In the framework of conditional transformation models (?), we fit daily distributions of the number of such collisions and penalize differences in these distributions between subsequent days. A monotonic constraint is needed to fit the conditional distribution, while a cyclic constraint should be used for the seasonal effect of deer–vehicle collisions. These two conditions yield a tensor product of two univariate constrained effects.

The remainder of the paper is organized as follows. Model estimation based on boosting is briefly introduced in Section 2. Monotonic effects, cyclic P-splines, and P-splines with boundary constraints are introduced in Section 3, where we also briefly introduce varying coefficients. An extension of monotonicity and cyclicity constraints to bivariate P-splines is given in Section 4. We present the three case studies described above at the end of the paper. An overview of past and present developments of constrained regression models is given in Online Resource 1 (App. ??). Some mathematical details are found in Online Resource 1 (App. ??); details on the data sets are given in other appendices.

## 2 Model Estimation Based on Boosting

To fit a STAR model (1) by boosting, one starts with a constant model, e.g.,  $\hat{\eta}(\mathbf{x}) \equiv 0$ , and computes the

negative gradient  $\mathbf{u} = (u_1, \dots, u_n)^\top$  of the loss function evaluated at each observation. An appropriate loss function is guided by the fitting problem. For Gaussian regression models, one may use the quadratic loss function, and for generalized linear models, the negative log-likelihood. In the Gaussian regression case, the negative gradient  $\mathbf{u}$  equals the standard residuals; in other cases,  $\mathbf{u}$  can be regarded as “working residuals”. In the next step, each model component  $f_l, l = 1, \dots, L$ , of the structured additive model (1) is fitted separately to the negative gradient  $\mathbf{u}$  by penalized least squares. Only the model component that best describes the negative gradient is updated by adding a small proportion of the fit (e.g., 10%) to the *current* model fit. New residuals are computed, and the whole procedure is iterated until a fixed number of iterations is reached. The final model  $\hat{\eta}(\mathbf{x})$  is defined as the sum of all models fitted in this process. As only *one* modeling component in each boosting iteration is updated, variables are selected by stopping the boosting procedure after an appropriate number of iterations. This is usually done using cross-validation techniques.

For each of the model components, a corresponding regression model that is applied to fit the residuals has to be specified, the so-called base-learner. Hence, the base-learners resemble the generic representations  $f_l$  and determine which functional form each of the components can take. In the following sections, we briefly introduce smooth effect estimates and derive special base-learners for fitting constrained effect estimates. For details on functional gradient descent boosting and specification of base-learners, see [Bühlmann and Hothorn \(2007\)](#) and [Hofner et al. \(2012\)](#).

### 3 Constrained Regression

#### 3.1 Estimating Smooth Effects

For the sake of simplicity in the remainder of this paper, we will consider an arbitrary continuous predictor  $x$  and a single base-learner  $f_l$  only when we drop the function index  $l$ . To model smooth effects of continuous variables, we utilized penalized B-splines (i.e., P-splines). These were introduced by [Eilers and Marx \(1996\)](#) for nonparametric regression and were later transferred to the boosting framework by [Schmid and Hothorn \(2008\)](#). Considering observations  $\mathbf{x} = (x_1, \dots, x_n)^\top$  of a single variable  $x$ , a non-linear function  $f(x)$  can be approximated as

$$f(x) = \sum_{j=1}^J \beta_j B_j(x; \delta) = \mathbf{B}(x)^\top \boldsymbol{\beta},$$

where  $B_j(\cdot; \delta)$  is the  $j$ th B-spline basis function of degree  $\delta$ . The basis functions are defined on a grid of  $J - (\delta - 1)$  inner knots  $\xi_1, \dots, \xi_{J-(\delta-1)}$  with additional boundary knots (and usually a knot expansion in the boundary knots) and are combined in the vector  $\mathbf{B}(x) = (B_1(x), \dots, B_J(x))^\top$ , where for simplicity  $\delta$  was dropped. For more details on the construction of B-splines, we refer the reader to [Eilers and Marx \(1996\)](#). The function estimates can be written in matrix notation as  $\hat{f}(\mathbf{x}) = \mathbf{B}\hat{\boldsymbol{\beta}}$ , where the design matrix  $\mathbf{B} = (\mathbf{B}(x_1), \dots, \mathbf{B}(x_n))^\top$  comprises the B-spline basis vectors  $\mathbf{B}(x)$  evaluated for each observation  $x_i, i = 1, \dots, n$ . The function estimate  $\hat{f}(\mathbf{x})$  might adapt the data too closely and might become too erratic. To enforce smoothness of the function estimate, an additional penalty is used that penalizes large differences of the coefficients of adjacent knots. Hence, for a continuous response  $\mathbf{u}$  (here the negative gradient vector), we can estimate the function by minimizing a penalized least-squares criterion

$$\mathcal{Q}(\boldsymbol{\beta}) = (\mathbf{u} - \mathbf{B}\boldsymbol{\beta})^\top (\mathbf{u} - \mathbf{B}\boldsymbol{\beta}) + \lambda \mathcal{J}(\boldsymbol{\beta}; d) \quad (2)$$

with a quadratic difference penalty on the coefficients

$$\mathcal{J}(\boldsymbol{\beta}; d) = \boldsymbol{\beta}^\top \mathbf{D}_{(d)}^\top \mathbf{D}_{(d)} \boldsymbol{\beta}, \quad (3)$$

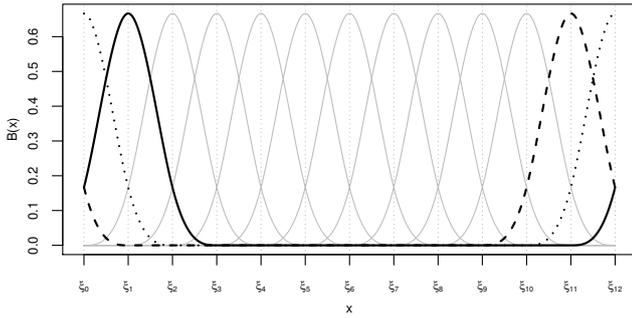
where  $d$  is the order of the difference penalty for the P-spline and  $\lambda$  is the smoothing parameter that governs the trade-off between smoothness and closeness to the data. The difference matrices  $\mathbf{D}_{(d)}$  are constructed such that they lead to the appropriate differences: first- and second-order differences result from a matrices of the form

$$\mathbf{D}_{(1)} = \begin{pmatrix} -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & & \\ & & & & -1 & 1 \end{pmatrix} \quad (4)$$

and

$$\mathbf{D}_{(2)} = \begin{pmatrix} 1 & -2 & 1 & & & \\ & \ddots & \ddots & \ddots & & \\ & & & & & \\ & & & & 1 & -2 & 1 \end{pmatrix}, \quad (5)$$

where empty cells are equal to zero. Higher order difference penalties can be easily derived. Difference penalties of order one penalize the deviation from a constant. Second-order differences penalize the deviation from a straight line. In general, differences of order  $d$  penalize deviations from polynomials of order  $d - 1$ . The unpenalized effects, i.e., the constant ( $d = 1$ ) or the straight line ( $d = 2$ ) are called the null space of the penalty. The null space remains unpenalized, even in the limit of  $\lambda \rightarrow \infty$ .



**Fig. 1** Illustration of cyclic P-splines of degree three, with 11 inner knots and boundary knots  $\xi_0$  and  $\xi_{12}$ . The gray curves correspond to B-splines. The black curves correspond to B-splines that are “wrapped” at the boundary knots, leading to a cyclic representation of the function. The dashed B-spline basis depends on observations in  $[\xi_9, \xi_{12}] \cup [\xi_0, \xi_1]$  and thus on observations from both ends of the range of the covariate. The same holds for the other two black B-spline curves (solid and dotted).

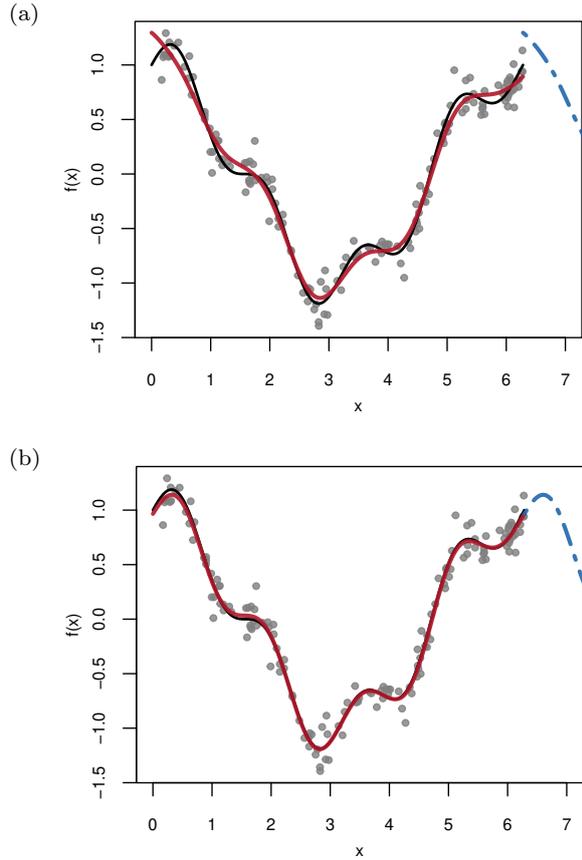
### 3.2 Estimating Cyclic Smooth Effects

P-splines with a cyclic constraint (Eilers and Marx, 2010) can be used to model periodic, seasonal data. The cyclic B-spline basis functions are constructed without knot expansion (Figure 1). The B-splines are “wrapped” at the boundary knots. The boundary knots  $\xi_0$  and  $\xi_J$  (equal to  $\xi_{12}$  in Fig. 1) play a central role in this setting as they specify the points where the function estimate should be smoothly joined. If  $x$  is, for example, the time during the day, then  $\xi_0$  is 0:00, whereas  $\xi_J$  is 24:00. Defining the B-spline basis in this fashion leads to a cyclic B-spline basis with the  $(n \times (J+1))$  design matrix  $\mathbf{B}_{\text{cyclic}}$ . The corresponding coefficients are collected in the  $((J+1) \times 1)$  vector  $\boldsymbol{\beta} = (\beta_0, \dots, \beta_J)$ .

Specifying a cyclic basis guarantees that the resulting function estimate is continuous in the boundary knots. However, no smoothness constraint is imposed so far. This can be achieved by a cyclic difference penalty, for example,  $\mathcal{J}_{\text{cyclic}}(\boldsymbol{\beta}) = \sum_{j=0}^J (\beta_j - \beta_{j-1})^2$  (with  $d = 1$ ) or  $\mathcal{J}_{\text{cyclic}}(\boldsymbol{\beta}) = \sum_{j=0}^J (\beta_j - 2\beta_{j-1} + \beta_{j-2})^2$  (with  $d = 2$ ), where the index  $j$  is “wrapped”, i.e.,  $j := J+1+j$  if  $j < 0$ . Thus, the differences between  $\beta_0$  and  $\beta_J$  or even  $\beta_{J-1}$  are taken into account for the penalty. Hence, the boundaries of the support are stabilized, and smoothness in and around the boundary knots is enforced. This can also be seen in Figure 2. The non-cyclic estimate (Fig. 2(a)) is less stable at the boundaries. As a consequence, the ends do not meet. The cyclic estimate (Fig. 2(b)), in contrast, is stabilized at the boundaries, and the ends are smoothly joined.

In matrix notation the penalty can be written as

$$\mathcal{J}_{\text{cyclic}}(\boldsymbol{\beta}; d) = \boldsymbol{\beta}^\top \tilde{\mathbf{D}}_{(d)}^\top \tilde{\mathbf{D}}_{(d)} \boldsymbol{\beta}, \quad (6)$$



**Fig. 2** (a) Non-cyclic and (b) cyclic P-splines. The red curve is the function estimated P-spline function. The blue, dashed curve is the same function shifted one period to the right. The data were simulated from a cyclic function with period  $2\pi$ :  $f(x) = \cos(x) + 0.25 \sin(4x)$  (black line) and realizations with additional normally distributed errors ( $\sigma = 0.1$ ; gray dots). The cyclic estimate is closer to the true function and more stable in the boundary regions, and the ends meet (see (b) at  $x = 2\pi$ ).

with difference matrices

$$\tilde{\mathbf{D}}_{(1)} = \begin{pmatrix} 1 & & -1 \\ -1 & 1 & \\ & \ddots & \ddots \\ & & -1 & 1 \end{pmatrix} \quad (7)$$

and

$$\tilde{\mathbf{D}}_{(2)} = \begin{pmatrix} 1 & & 1 & -2 \\ -2 & 1 & & 1 \\ 1 & -2 & 1 & \\ & \ddots & \ddots & \ddots \\ & & 1 & -2 & 1 \end{pmatrix}, \quad (8)$$

where empty cells are equal to zero. Coefficients can then be estimated using the penalized least-squares criterion (2), where the design matrix and the penalty

matrix are replaced with the corresponding cyclic counterparts, i.e.,  $\mathcal{Q}(\boldsymbol{\beta}) = (\mathbf{u} - \mathbf{B}_{\text{cyclic}}\boldsymbol{\beta})^\top (\mathbf{u} - \mathbf{B}_{\text{cyclic}}\boldsymbol{\beta}) + \lambda \mathcal{J}_{\text{cyclic}}(\boldsymbol{\beta}; d)$ .

As mentioned in Section 3.1, P-splines have a null space, i.e., an unpenalized effect, which depends on the order of the differences in the penalty. However, cyclic P-splines have a null space that includes only a constant, irrespective of the order of the difference penalty. Globally seen, i.e., for the complete function estimate, the order of the penalty plays no role (even in the limit  $\lambda \rightarrow \infty$ ). Locally, however, the order of the difference penalty has an effect. For example, with  $d = 2$ , the estimated function is penalized for deviations from linearity and hence, locally approaches a straight line (with increasing  $\lambda$ ).

### 3.3 Estimating Monotonic Effects

To achieve a smooth, yet monotonic function estimate, Eilers (2005) introduced P-splines with an *additional asymmetric difference penalty*. The penalized least-squares criterion (2) becomes

$$\mathcal{Q}(\boldsymbol{\beta}) = (\mathbf{u} - \mathbf{B}\boldsymbol{\beta})^\top (\mathbf{u} - \mathbf{B}\boldsymbol{\beta}) + \lambda_1 \mathcal{J}(\boldsymbol{\beta}; d) + \lambda_2 \mathcal{J}_{\text{asym}}(\boldsymbol{\beta}; c), \quad (9)$$

with the quadratic difference penalty of order  $d$  as in usual P-splines (Eq. 3) and the additional asymmetric difference penalty of order  $c$

$$\mathcal{J}_{\text{asym}}(\boldsymbol{\beta}; c) = \sum_{j=c+1}^J v_j (\Delta^c \beta_j)^2 = \boldsymbol{\beta}^\top \mathbf{D}_{(c)}^\top \mathbf{V} \mathbf{D}_{(c)} \boldsymbol{\beta}, \quad (10)$$

where the difference matrix  $\mathbf{D}_{(c)}$  is constructed as in Equations (4) and (5). The choice of  $c$  implies the type of the additional constraint: monotonicity for  $c = 1$  or convexity/concavity for  $c = 2$ . In the remainder of this article, we restrict our attention to monotonicity constraints; however, one can also consider concave constraints. The asymmetric penalty looks very much like the P-spline penalty (3) with the important distinction of weights  $v_j$ , which are specified as

$$v_j = \begin{cases} 0 & \text{if } \Delta^c \beta_j > 0 \\ 1 & \text{if } \Delta^c \beta_j \leq 0. \end{cases} \quad (11)$$

The weights are collected in the diagonal matrix  $\mathbf{V} = \text{diag}(\mathbf{v})$ . With  $c = 1$ , this enforces monotonically *increasing* functions. Changing the direction of the inequalities in the distinction of cases leads to monotonically *decreasing* functions. As the weights (11) depend on the coefficients  $\boldsymbol{\beta}$ , a solution to (9) can only be found by iteratively minimizing  $\mathcal{Q}(\boldsymbol{\beta})$  with respect to  $\boldsymbol{\beta}$ , where

the weights  $\mathbf{v}$  are updated in each iteration. The estimation converges if no further changes in the weight matrix  $\mathbf{V}$  occur. The penalty parameter  $\lambda_2$ , which is associated with the additional constraint (10), should be chosen quite large (e.g.,  $10^6$ ; Eilers, 2005) and resembles the researcher's *a priori* assumption of monotonicity. Larger values are associated with a stronger impact of the monotonic constraint on the estimation. A detailed discussion of monotonic P-splines in the context of boosting models, together with an empirical evaluation, is given in Hofner *et al.* (2011b). In their presented framework, the authors also derive an asymmetric difference penalty for monotonicity-constrained, ordered categorical effects.

One can also use asymmetric difference penalties on differences of order 0, i.e., on the coefficients themselves, to achieve smooth positive or negative effect estimates. This idea can also be used to fit smooth effect estimates with an arbitrarily fixed co-domain by specifying either upper or lower bounds or both bounds at the same time.

*Improved Fitting Method* An alternative to iteratively minimizing the penalized least-squares criterion (9) to obtain smooth monotonic estimates is given by quadratic programming methods (Goldfarb and Idnani, 1982, 1983). The criterion we want to minimize is given by the penalized least-squares criterion (2) with the additional constraint

$$\mathbf{D}_{(c)}\boldsymbol{\beta} \geq \mathbf{0}$$

with difference matrix  $\mathbf{D}_{(c)}$  as defined above and null vector  $\mathbf{0}$  (of appropriate dimension). To change the direction of the constraint, e.g., to obtain monotonically decreasing functions, one can use the negative difference matrix  $-\mathbf{D}_{(c)}$ . The results obtained by quadratic programming are (virtually) identical to the results obtained by iteratively solving (9), but the computation time can be greatly reduced.

### 3.4 Estimating Effects with Boundary Constraints

In some cases, e.g., for extrapolation, it might be of interest to impose boundary constraints, such as constant or linear boundaries, to higher order splines. These constraints can be enforced by using a strong penalty on, e.g., the three outer spline coefficients on each side of the range of the data or on one side only. Constant boundaries are obtained by a strong penalty on the first-order differences, while a strong second-order difference penalty results in linear boundaries. Technically,

this can be obtained by an additional penalty

$$\begin{aligned} \mathcal{J}_{\text{boundary}}(\boldsymbol{\beta}; e) &= \sum_{j=c+1}^J v_j (\Delta^e \beta_j)^2 \\ &= \boldsymbol{\beta}^\top \mathbf{D}_{(e)}^\top \mathbf{V}^{(3)} \mathbf{D}_{(e)} \boldsymbol{\beta}, \end{aligned} \quad (12)$$

where  $\mathbf{D}_{(e)}$  is a difference matrix of order  $e$  (cf. Eq. (4) and (5)). The weight  $v_j^{(3)}$  is one if the corresponding coefficient is subject to a boundary constraint. Thus, here the first and the last three elements of  $\mathbf{v}^{(3)}$  are equal to one, and the remaining weights are equal to zero. The weight matrix  $\mathbf{V}^{(3)} = \text{diag}(\mathbf{v}^{(3)})$ . Boundary constraints can be successfully imposed on P-splines as well as on monotonic P-splines by adding the penalty (12) to the respective penalized least-squares criterion. A quite large penalty parameter  $\lambda_3$  associated with the boundary constraint is chosen (e.g.,  $10^6$ ). For an application of modeling the gas flow in gas transmission networks using monotonic effects with boundary constraints, see [Sobotka et al. \(2013\)](#).

### 3.5 Varying Coefficients

Varying-coefficient models (VCMs) allow one to model flexible interactions in which the regression coefficients of a predictor vary smoothly with one or more other variables, the so-called effect modifiers ([Hastie and Tibshirani, 1993](#)). The varying coefficient term can be written as  $f(x, z) = x \cdot \beta(z)$ , where  $z$  is the effect modifier and  $\beta(\cdot)$  a smooth function of  $z$ . Technically, varying coefficients can be modeled by fitting the interaction of  $x$  and a basis expansion of  $z$ . Thus, we can use all discussed spline types, such as simple P-splines, monotonic splines, or cyclic splines, to model  $\beta(z)$ . Furthermore, bivariate P-splines as discussed in the following section can be facilitated as well.

## 4 Constrained Effects for Bivariate P-Splines

### 4.1 Bivariate P-spline Base-learners

Bivariate, or tensor product, P-splines are an extension of univariate P-splines that allow one to model smooth effects of two variables. These can be used to model smooth interaction surfaces, most prominently spatial effects. A bivariate B-spline of degree  $\delta$  for two variables  $x_1$  and  $x_2$  can be constructed as the product of two univariate B-spline bases  $B_{jk}(x_1, x_2; \delta) = B_j^{(1)}(x_1; \delta) \cdot B_k^{(2)}(x_2; \delta)$ . The bivariate B-spline basis is formed by all possible products  $B_{jk}$ ,  $j = 1, \dots, J$ ,  $k = 1, \dots, K$ . Theoretically, different numbers of knots for  $x_1$  ( $J$ ) and

$x_2$  ( $K$ ) are possible, as well as B-spline basis functions with different degrees  $\delta_1$  and  $\delta_2$  for  $x_1$  and  $x_2$ , respectively. A bivariate function  $f(x_1, x_2)$  can be approximated as  $f(x_1, x_2) = \sum_{j=1}^J \sum_{k=1}^K \beta_{jk} B_{jk}(x_1, x_2) = \mathbf{B}(x_1, x_2)^\top \boldsymbol{\beta}$ , where the vector of B-spline bases for variables  $(x_1, x_2)$  equals  $\mathbf{B}(x_1, x_2) = \left( B_{11}(x_1, x_2), \dots, B_{1K}(x_1, x_2), B_{21}(x_1, x_2), \dots, B_{JK}(x_1, x_2) \right)^\top$ , and the coefficient vector

$$\boldsymbol{\beta} = (\beta_{11}, \dots, \beta_{1K}, \beta_{21}, \dots, \beta_{JK})^\top. \quad (13)$$

The  $(n \times JK)$  design matrix then combines the vectors  $\mathbf{B}(\mathbf{x}_i)$  for observations  $\mathbf{x}_i = (x_{i1}, x_{i2})$ ,  $i = 1, \dots, n$ , such that the  $i$ th row contains  $\mathbf{B}(\mathbf{x}_i)$ , i.e.,  $\mathbf{B} = \left( \mathbf{B}(\mathbf{x}_1), \dots, \mathbf{B}(\mathbf{x}_i), \dots, \mathbf{B}(\mathbf{x}_n) \right)^\top$ . The design matrix  $\mathbf{B}$  can be conveniently obtained by first evaluating the univariate B-spline bases  $\mathbf{B}^{(1)} = \left( B_j^{(1)}(x_{i1}) \right)_{\substack{i=1, \dots, n \\ j=1, \dots, J}}$  and  $\mathbf{B}^{(2)} = \left( B_k^{(2)}(x_{i2}) \right)_{\substack{i=1, \dots, n \\ k=1, \dots, K}}$  of the variables  $x_1$  and  $x_2$  and subsequently constructing the design matrix as

$$\mathbf{B} = \left( \mathbf{B}^{(1)} \otimes \mathbf{e}_K^\top \right) \odot \left( \mathbf{e}_J^\top \otimes \mathbf{B}^{(2)} \right), \quad (14)$$

where  $\mathbf{e}_K = (1, \dots, 1)^\top$  is a vector of length  $K$  and  $\mathbf{e}_J = (1, \dots, 1)^\top$  a vector of length  $J$ . The symbol  $\otimes$  denotes the Kronecker product and  $\odot$  denotes the element-wise product. Definitions of both products are given in Online Resource 1 (App. ??).

As for univariate P-splines, a suitable penalty matrix is required to enforce smoothness. The bivariate penalty matrix can be constructed from separate, univariate difference penalties for  $x_1$  and  $x_2$ , respectively. Consider the  $(J \times J)$  penalty matrix  $\mathbf{K}^{(1)} = (\mathbf{D}^{(1)})^\top \mathbf{D}^{(1)}$  for  $x_1$ , and the  $(K \times K)$  penalty matrix  $\mathbf{K}^{(2)} = (\mathbf{D}^{(2)})^\top \mathbf{D}^{(2)}$  for  $x_2$ . The penalties are constructed using difference matrices  $\mathbf{D}^{(1)}$  and  $\mathbf{D}^{(2)}$  of (the same) order  $d$ . However, different orders of differences  $d_1$  and  $d_2$  could be used if this is required by the data at hand. The combined difference penalty can then be written as the sum of Kronecker products

$$\mathcal{J}_{\text{tensor}}(\boldsymbol{\beta}; d) = \boldsymbol{\beta}^\top (\mathbf{K}^{(1)} \otimes \mathbf{I}_K + \mathbf{I}_J \otimes \mathbf{K}^{(2)}) \boldsymbol{\beta}, \quad (15)$$

with identity matrices  $\mathbf{I}_J$  and  $\mathbf{I}_K$  of dimension  $J$  and  $K$ , respectively.

With the response vector  $\mathbf{u}$ , models can then be estimated by optimizing the penalized least-squares criterion in analogy to univariate P-splines:

$$\mathcal{Q}(\boldsymbol{\beta}) = (\mathbf{u} - \mathbf{B}\boldsymbol{\beta})^\top (\mathbf{u} - \mathbf{B}\boldsymbol{\beta}) + \lambda \mathcal{J}_{\text{tensor}}(\boldsymbol{\beta}; d), \quad (16)$$

with design matrix (14) and penalty (15). For more details on tensor product splines, we refer the reader to [Wood \(2006b\)](#). [Kneib et al. \(2009\)](#) give an introduction to tensor product P-splines in the context of boosting.

## 4.2 Estimating Bivariate Cyclic Smooth Effects

Based on bivariate P-splines, cyclic constraints in both directions of  $x_1$  and  $x_2$  can be straightforwardly implemented. One builds the univariate, *cyclic* design matrices  $\mathbf{B}_{\text{cyclic}}^{(1)}$  and  $\mathbf{B}_{\text{cyclic}}^{(2)}$  for  $x_1$  and  $x_2$ , respectively. The bivariate design matrix then is  $\mathbf{B}_{\text{cyclic}} = (\mathbf{B}_{\text{cyclic}}^{(1)} \otimes \mathbf{e}_K^\top) \odot (\mathbf{e}_J^\top \otimes \mathbf{B}_{\text{cyclic}}^{(2)})$ , as in Equation (14).

With the univariate, *cyclic* difference matrices  $\tilde{\mathbf{D}}^{(1)}$  for  $x_1$  and  $\tilde{\mathbf{D}}^{(2)}$  for  $x_2$  (cf. Eq. (7) and (8)), we obtain cyclic penalty matrices  $\mathbf{K}^{(1)} = (\tilde{\mathbf{D}}^{(1)})^\top \tilde{\mathbf{D}}^{(1)}$  and  $\mathbf{K}^{(2)} = (\tilde{\mathbf{D}}^{(2)})^\top \tilde{\mathbf{D}}^{(2)}$ . Thus, in analogy to the usual bivariate P-spline penalty (Eq. 15), the bivariate cyclic penalty can be written as  $\mathcal{J}_{\text{cyclic, tensor}}(\boldsymbol{\beta}) = \boldsymbol{\beta}^\top (\mathbf{K}^{(1)} \otimes \mathbf{I}_K + \mathbf{I}_J \otimes \mathbf{K}^{(2)}) \boldsymbol{\beta}$ , in analogy to Equation (13). Estimation is then a straightforward application of the penalized least-squares criterion as in (Eq. 16) with cyclic design and penalty matrices. An example of bivariate, cyclic splines is given in Section 5.2, where a cyclic surface is used to estimate the combined effect of time (during the day) and calendar day on roe deer activity.

## 4.3 Estimating Bivariate Monotonic Effects

As shown in Hofner (2011), it is not sufficient to add monotonicity constraints to the marginal effects because the resulting interaction surface might still be non-monotonic. Thus, instead of the marginal functions, the complete surface needs to be constrained in order to achieve a monotonic surface. Therefore, we utilized bivariate P-splines and added monotonic constraints for the row- and column-wise differences of the matrix of coefficients  $\mathcal{B} = (\beta_{jk})_{j=1, \dots, J; k=1, \dots, K}$ . As proposed by Bollaerts *et al.* (2006), one can use two independent asymmetric penalties to allow different directions of monotonicity, i.e., increasing in one variable, e.g.,  $x_1$ , and decreasing in the other variable, e.g.,  $x_2$ , or one can use different prior assumptions of monotonicity reflected in different penalty parameters  $\lambda$ . Let  $\mathbf{B}$  denote the  $(n \times JK)$  design matrix (14) comprising the bivariate B-spline bases of  $\mathbf{x}_i$ , and let  $\boldsymbol{\beta}$  denote the corresponding  $(JK \times 1)$  coefficient vector (13). Monotonicity is enforced by the asymmetric difference penalties

$$\begin{aligned} \mathcal{J}_{\text{asym},1}(\boldsymbol{\beta}; c) &= \sum_{j=c+1}^J \sum_{k=1}^K v_{jk}^{(1)} (\Delta_1^c \beta_{jk})^2 \\ &= \boldsymbol{\beta}^\top (\mathbf{D}^{(1)} \otimes \mathbf{I}_K)^\top \mathbf{V}^{(1)} (\mathbf{D}^{(1)} \otimes \mathbf{I}_K) \boldsymbol{\beta}, \\ \mathcal{J}_{\text{asym},2}(\boldsymbol{\beta}; c) &= \sum_{j=1}^J \sum_{k=c+1}^K v_{jk}^{(2)} (\Delta_2^c \beta_{jk})^2 \\ &= \boldsymbol{\beta}^\top (\mathbf{I}_J \otimes \mathbf{D}^{(2)})^\top \mathbf{V}^{(2)} (\mathbf{I}_J \otimes \mathbf{D}^{(2)}) \boldsymbol{\beta}, \end{aligned}$$

where  $\Delta_1^c$  are the column-wise and  $\Delta_2^c$  the row-wise differences of order  $c$ , i.e.,  $\Delta_1^1 \beta_{jk} = \beta_{jk} - \beta_{(j-1)k}$  and  $\Delta_2^1 \beta_{jk} = \beta_{jk} - \beta_{j(k-1)}$ , etc. Thus,  $\mathcal{J}_{\text{asym},1}$  is associated with constraints in the direction of  $x_1$ , while  $\mathcal{J}_{\text{asym},2}$  acts in the direction of  $x_2$ . The corresponding difference matrices are denoted by  $\mathbf{D}^{(1)}$  and  $\mathbf{D}^{(2)}$ . The weights  $v_{jk}^{(l)}$ ,  $l = 1, 2$  are specified in analogy to (11), i.e., with  $c = 1$ , we obtain monotonically increasing estimates with weights  $v_{jk}^{(l)} = 1$  if  $\Delta_l^c \beta_{jk} \leq 0$ , and  $v_{jk}^{(l)} = 0$  otherwise. Changing the inequality sign leads to monotonically decreasing function estimates. Differences of order  $c = 2$  lead to convex or concave constraints. For the matrix notation, the weights are collected in the diagonal matrices  $\mathbf{V}^{(l)} = \text{diag}(\mathbf{v}^{(l)})$ . The constraint estimation problem for monotonic surface estimates in matrix notation becomes

$$\begin{aligned} \mathcal{Q}(\boldsymbol{\beta}) &= (\mathbf{u} - \mathbf{B}\boldsymbol{\beta})^\top (\mathbf{u} - \mathbf{B}\boldsymbol{\beta}) + \lambda_1 \mathcal{J}_{\text{tensor}}(\boldsymbol{\beta}; d) \\ &\quad + \lambda_{21} \mathcal{J}_{\text{asym},1}(\boldsymbol{\beta}; c) + \lambda_{22} \mathcal{J}_{\text{asym},2}(\boldsymbol{\beta}; c), \end{aligned} \quad (17)$$

where  $\mathcal{J}_{\text{tensor}}(\boldsymbol{\beta}; d)$  is the standard bivariate P-spline penalty of order  $d$  with corresponding penalty parameter  $\lambda_1$ , and  $\lambda_{21}$  and  $\lambda_{22}$  are the penalty parameters associated with constraints in the direction of  $x_1$  and  $x_2$ , respectively. Hence, setting either of the two penalty parameters to zero results in an unconstrained estimate in this direction, with a constraint in the other direction. For example, by setting  $\lambda_{21} = 0$ , one could get a surface that is monotone in  $x_2$  for each value of  $x_1$  but is not necessarily monotone in  $x_1$ .

*Model Fitting* Model estimation can be achieved by using either the iterative algorithm to solve Equation (17) or quadratic programming methods as in the univariate case described in Section 3.3. In the latter case, we minimize the penalized least-squares criterion (16) subject to the constraints  $(\mathbf{D}^{(1)} \otimes \mathbf{I}_K) \boldsymbol{\beta} \geq \mathbf{0}$  and  $(\mathbf{I}_J \otimes \mathbf{D}^{(2)}) \boldsymbol{\beta} \geq \mathbf{0}$ , i.e., we constrain the row-wise or column-wise differences to be non-negative. As above, multiplying the difference matrices by  $-1$  leads to monotonically decreasing estimates. Using the two constraints is equivalent to requiring

$$\begin{pmatrix} \mathbf{D}^{(1)} \otimes \mathbf{I}_K \\ \mathbf{I}_J \otimes \mathbf{D}^{(2)} \end{pmatrix} \boldsymbol{\beta} \geq \mathbf{0}.$$

## 5 Case Studies

### 5.1 São Paulo Air Pollution

In this study, we examined the effect of air pollution in São Paulo on mortality. Saldiva *et al.* (1995) investigated the impact of air pollution on mortality caused by respiratory problems of elderly people (over 65 years

of age). We concentrated on the effect of  $\text{SO}_2$  on mortality of elderly people. We considered a Poisson model for the number of respiratory deaths of the form

$$\log(\boldsymbol{\mu}) = \mathbf{x}^\top \boldsymbol{\beta} + f_1(\text{day of the year}) + f_2(\text{time}) + f_3(\text{SO}_2),$$

where the expected number of death due to respiratory causes  $\boldsymbol{\mu}$  is related to a linear model with respect to covariates  $\mathbf{x}$ , such as temperature, humidity, days of week, and non-respiratory deaths. Additionally, we wanted to adjust for temporal changes; the study was conducted over four successive years from January 1994 to December 1997. This allows us to decompose the seasonal effect into a smooth cyclicity-constrained effect for the day of the year ( $f_1$ ) and a smooth long-term trend for the variation over the years ( $f_2$ ). Finally, we added a smooth effect  $f_3$  of the pollutant's concentration. Smooth estimates of the effect of  $\text{SO}_2$  on respiratory deaths behaved erratically. This seems unreasonable as an increase in the air pollutant should not result in a decreased risk of death. Hence, a monotonically increasing effect should lead to a more stable model that can be interpreted. Recently, [Leitenstorfer and Tutz \(2007\)](#) estimated a similar model for this data set. They used an approach of likelihood-based boosting ([Tutz and Binder, 2006](#)) to estimate a model with a monotonic effect but without seasonal effect estimates for the temporal model component. Thus, the long-term trend captured both seasonal variation and long-term variation. Here, we combined both monotonicity constraints and cyclic effects in one model. Details of the data set and the used model can be found in Online Resource 1 (App. ??).

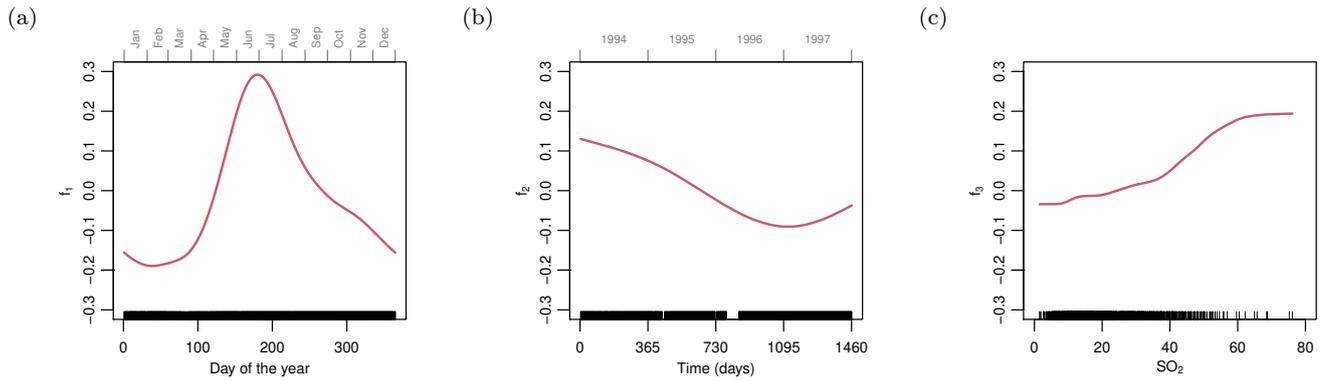
**Results** In addition to our boosting approach ('mboost'), which is implemented in the package **mboost** ([Hofner et al., 2012](#); [Hothorn et al., 2013a](#)), we fitted a model using the monotonic approach of [Leitenstorfer and Tutz \('LT2007'\)](#), and an unconstrained additive model with the function **gam** in the R package **mgcv** ([Wood, 2006a, 2010](#)) ('mgcv').

As we used a cyclic constraint for the seasonal effect, the ends of the function estimate meet, i.e., day 365 and day 1 are smoothly joined (see [Figure 3\(a\)](#)). The effect showed a clear peak in the cool and dry winter months (May to August in the southern hemisphere) and a decreased risk of mortality in the warm summer months. This is in line with the results of other studies (e.g., [Saldiva et al., 1995](#)). In the trend over the years ([Figure 3\(b\)](#)), mortality decreased from 1994 to early 1997 and increased thereafter. However, one should keep in mind that this trend needs to be combined with the periodical effect to form the complete temporal pattern.

The estimated smooth effect for the pollutant  $\text{SO}_2$  resulting from the model 'mboost' ([Figure 3\(c\)](#)) indicated that an increase of the pollutant's concentration does not result in a (substantially) higher mortality up to a concentration of  $40 \mu\text{g}/\text{m}^3$ . From this point onward, a steep increase in the expected mortality was observed, which flattened again for concentrations above  $60 \mu\text{g}/\text{m}^3$ . Hence, a dose-response relationship was observed, where higher pollutant concentrations result in a higher expected mortality. At the same time, the model indicated that increasing pollutant concentrations are almost harmless until a threshold is exceeded, and that the harm of  $\text{SO}_2$  is not further increased after reaching an upper threshold. In an investigation of the effect of  $\text{PM}_{10}$  [Saldiva et al. \(1995\)](#) found no "safe" threshold in their study of elderly people in São Paulo. They also investigated the effect of  $\text{SO}_2$  but did not report on details, such as possible threshold values, in this case. The more recent study on the effect of air pollution in São Paulo on children ([Conceição et al., 2001](#)) used only linear effects for pollutant concentrations. Hence, no threshold values can be estimated.

The linear effects of the model 'mboost' (results not presented here) showed a negative effect of humidity, which indicated that higher humidity reduces the expected number of deaths. The minimum temperature had no impact on mortality (i.e., it was not selected in the boosting algorithm) or at least no additional effect to the seasonal component, which might capture a temperature effect as well. Regarding the days of the week, mortality was higher on Monday than on Sunday and was even lower on all other days. This result might be due to different behavior and thus personal exposure to the pollutant on weekends or, more likely, due to a lag in recording on weekends.

Despite the possible difference in the shape from year to year of the estimated seasonal effects resulting from 'LT2007' and 'mgcv', the estimates were very similar compared to those of the model 'mboost' ([Fig. ??](#)), in which the effects have the same shape every year. Yet, the estimation of the complete time pattern without decomposition into the trend effect and the periodical, seasonal effect was less stable. Especially at the boundaries, models 'LT2007' and 'mgcv' were highly variable or perhaps even behaved erratically. Modeling the trend and the periodic effect separately may have the disadvantage that some of the small-scale changes (e.g., around day 730) are missed. However, without this decomposition, models do not allow a direct inspection of the seasonal effect throughout the year. Hence, decomposing influence of time into seasonal effects and smooth long-term effects seems highly preferable as it



**Fig. 3** Estimated effects of the model ‘mboost’ all centered around zero: (a) cyclic seasonal effect, (b) unconstrained long-term trend, and (c) monotonic effect of  $SO_2$ .

offers a stable, yet flexible method to model the data and allows an easier and more profound interpretation.

Both monotonic approaches showed a similar pattern in the effect of  $SO_2$  (Fig. ??). Estimates from non-monotonic model (‘mgcv’) were very wiggly for small values up to a concentration of  $40 \mu g/m^3$ , which seems unreasonable. For a detailed discussion, see [Leitenstorfer and Tutz \(2007\)](#).

Finally, all considered models had almost the same linear effects for the covariates (results not shown here) or showed at least the same direction and magnitude of the effect. Hence, we can conclude that the linear effects in this model are very stable and are hardly influenced by the fitting method, i.e., boosting vs. penalized iteratively weighted least-squares (P-IWLS; see e.g., [Wood, 2008](#)), or by constraints, i.e., monotonic or cyclic constraints, that are used to model the data.

## 5.2 Activity of Roe Deer (*Capreolus capreolus*)

In the Bavarian Forest National Park (Germany), the applied wildlife management strategy is regularly examined. Part of the strategy involves trying to understand the activity profiles of the various species, including lynx, wild boar, and roe deer. The case study here focuses on the activity of European roe deer (*Capreolus capreolus*). According to [Stache et al. \(2013\)](#), animal activity is influenced by exogenous factors, such as the azimuth of the sun (i.e., day/night rhythm and seasons), temperature, precipitation, and depth of snow. Another important role is played by endogenous factors, such as the species (e.g., reflected in their diet; roe deer are browsers), age, and sex. Additionally, as roe deer tend to be solitary animals, a high level of individual specific variation in activity is to be expected. The activity data was recorded using telemetry collars with an acceleration sensor unit. The activity is repre-

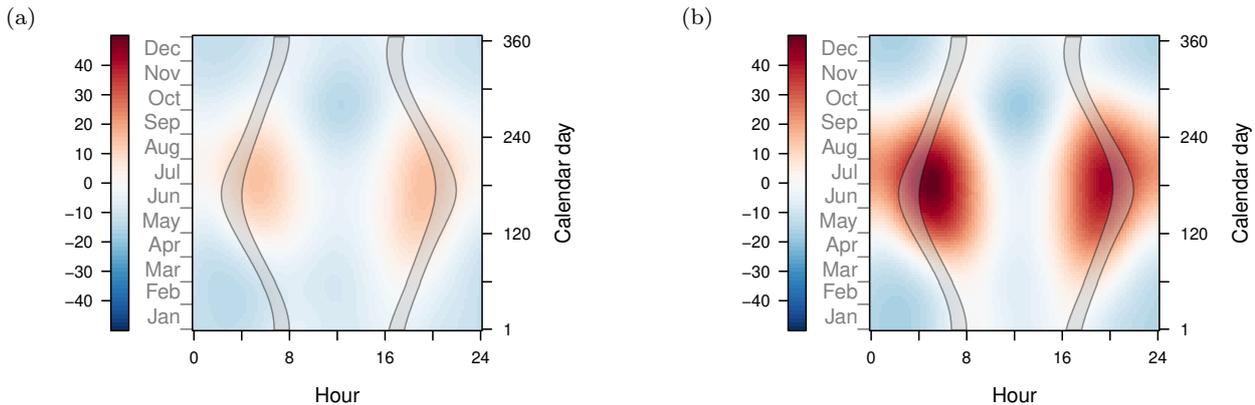
sented by a number ranging from 0 to 510, where higher values represent higher activity.

Activity profiles for the day and for the year were provided. As earlier analysis showed, the activity of males and females differs greatly. Hence, sex should be considered as an effect modifier in the analysis by defining sex-specific activity profiles. We considered a Gaussian model with the additive predictor

$$\begin{aligned} \mathbb{E}(\text{activity}|\cdot) = & \mathbf{x}^\top \boldsymbol{\beta} + \text{temp} \cdot f_1(t_{\text{days}}) \\ & + \text{depth of snow} \cdot f_2(t_{\text{days}}) \\ & + \text{precipitation} \cdot f_3(t_{\text{days}}) \\ & + f_4(t_{\text{hours}}, t_{\text{days}}) \\ & + I_{(\text{sex} = \text{male})} f_5(t_{\text{hours}}, t_{\text{days}}) + b_{\text{roe}}, \end{aligned}$$

where  $\mathbf{x}$  contains the categorical covariates sex, type of collar, year of observation, and age. Temperature, depth of snow, and precipitation entered the model rescaled to  $|x| \leq 1$  by dividing the variables by the respective absolute maximum values. The effects of temperature ( $f_1$ ), depth of snow ( $f_2$ ), and precipitation ( $f_3$ ) depend on the calendar day ( $t_{\text{days}}$ ). An interaction surface ( $f_4$ ) for time of the day ( $t_{\text{hours}}$ ) and calendar day ( $t$ ) was specified to flexibly model the daily activity profiles throughout the year. An additional effect for male roe deer was specified with  $f_5$ . Finally, a random intercept  $b_{\text{roe}}$  for each roe deer was included. Details on the data set and on the model specification can be found in [Online Resource 1 \(App. ??\)](#).

**Results** In the resulting model, six of ten base-learners were selected. The largest contribution to the model fit was given by the smooth interaction surfaces  $f_4$  and  $f_5$ , which represent the time-dependent activity profiles for male and female roe deer (Figure 4). The individual activity of the roe deer  $b_{\text{roe}}$  substantially contributed to the total predicted variation, with a range of approximately 20 units (not depicted here). The time-varying



**Fig. 4** Influence of time on roe deer activity. Combined effect of calendar day and time of day (a) for female roe deer ( $= f_4$ ) and (b) for male roe deer ( $= f_4 + f_5$ ), together with twilight phases (gray). White areas depict the mean activity level throughout the year; blue shading represents decreased activity, and red shading represents increased activity.

effects of temperature and depth of snow and the effect of the type of collar had a lower impact on the recorded activity of roe deer.

The activity profiles (Fig. 4) showed that roe deer were most active in and around the twilight phases in the mornings and evenings. This holds for the whole year and for both males and females. In general, the activity profiles of female and male roe deer were very similar, but male activity was much higher and had more variability. The activity of roe deer was strongly influenced by the season: During summer, the activity was much higher throughout the entire day. The phase of least activity was around noon. This behavior was enhanced in autumn. In spring, activity was more evenly distributed throughout the daytime, and lowest activity occurred during the hours after midnight.

The effects of climatic variables are depicted in Figure 5. A higher temperature led to lower activity (negative effect of temperature), except from January to March, when higher temperatures led to higher activity (positive effect of temperature). The depth of snow had a negative effect on roe deer activity throughout the year, i.e., deeper snow led to lower activity. The effect of snow depth was stronger in the summer months (when there is hardly any snow), and less strong in January and February even though the snow depth was the greatest. Precipitation had no effect on roe deer activity according to our model.

### 5.3 Deer–vehicle Collisions in Bavaria

Important areas of application for both monotone and cyclic base-learners are conditional transformation models (?). Here, we describe the distribution

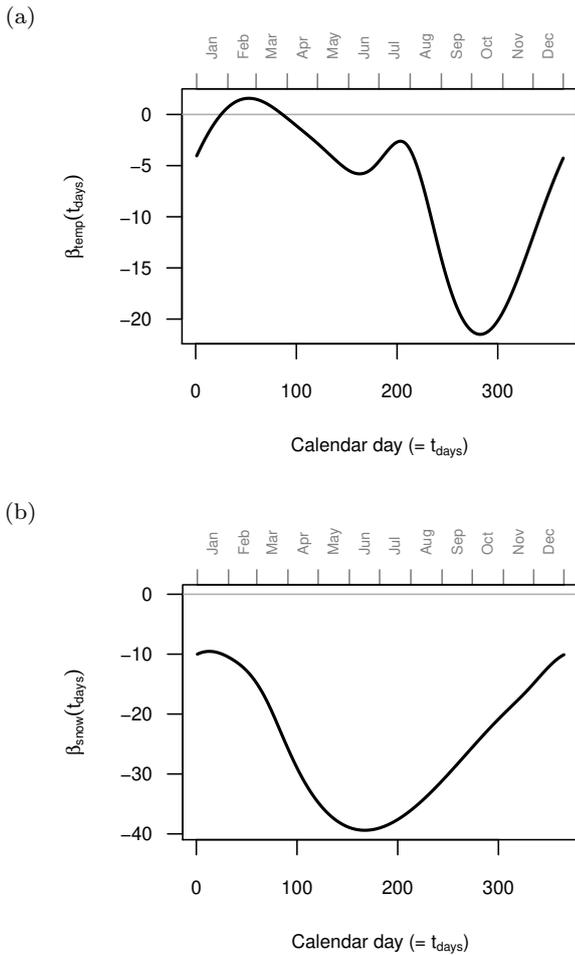
of the number of deer–vehicle collisions (DVC) that took place throughout Bavaria, Germany, for each day  $k = 1, \dots, 365$  of the year 2006 (see Hothorn *et al.*, 2012, for a more detailed description of the data), i.e.,

$$\mathbb{P}(\text{number of DVCs} \leq y | \text{day} = k) = \Phi(h(y|k)),$$

where  $\Phi$  is the distribution function of the standard normal distribution. The conditional transformation function  $h$  is parametrized as

$$h(y|k) = (\mathbf{B}_{\text{day}}(k) \otimes \mathbf{B}_{\text{DVC}}(y))\beta,$$

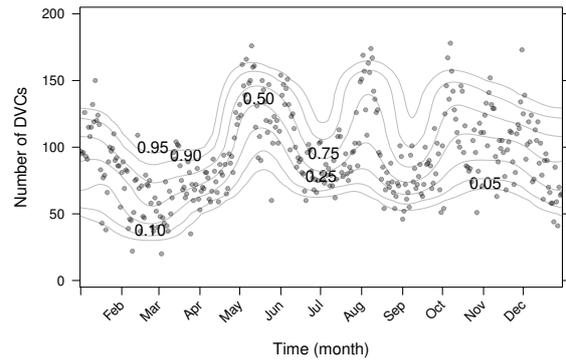
where  $\mathbf{B}_{\text{day}}$  is a cyclic B-spline transformation for the day of the year (where Dec 31 and Jan 1 should match) and  $\mathbf{B}_{\text{DVC}}$  is a B-spline transformation for the number of deer–vehicle collisions. The Kronecker product  $\mathbf{B}_{\text{day}}(k) \otimes \mathbf{B}_{\text{DVC}}(y)$  defines a bivariate tensor product spline, which is fitted under smoothness constraints in both dimensions. Since the transformation function  $h(y|k)$  must be monotone in  $y$  for all days  $k$  (otherwise  $\Phi(h(y|k))$  is not a distribution function), we imposed a monotonicity constraint on the second term, i.e., we required monotonicity with respect to the number of deer–vehicle collisions but not with respect to time. The model was fitted by minimizing a scoring rule for probabilistic forecasts (e.g., Brier score or log score). Here, we applied the boosting approach described in (?) and display the corresponding quantile functions over the course of the year 2006 in Figure 6. Three peaks (territorial movement at beginning of May, rut at end of July/ beginning of August, and early in October) were identified. While the first two peaks are expected, the significance of the third peak in October remains to be discussed with ecologists. Over the year, not only the mean but also higher moments of the distribution of the number of deer–vehicle collisions varied over time.



**Fig. 5** Time-varying effects (i.e.,  $\beta(t_{\text{days}})$ ) of (a) temperature and (b) depth of snow. Note that both variables are rescaled, i.e.,  $\beta(t_{\text{days}})$  is the maximal effect. On a given day, the effects of temperature and depth of snow are linear. The higher the amplitude for a given day, the stronger the effect will be.

## 6 Concluding Remarks

In this article, we extended the flexible modeling framework based on boosting to allow inclusion of monotonic or cyclic constraints for certain variables. The monotonicity constraint on continuous variables leads to monotonic, yet smooth effects. Many other approaches to monotonic modeling result in non-smooth function estimates (e.g., Dette *et al.*, 2006; de Leeuw *et al.*, 2009; ?). In many application contexts such as life sciences in the broadest sense, however, we feel that smooth effect estimates are more plausible and hence preferable. In the limit, i.e., if the ‘true effect’ is monotonic and if there is little noise in the data, our approach leads to P-spline estimates. Hence, in this case, it does not matter whether monotonic constraints are used or not. Monotonic effects can be furthermore applied to bivariate P-splines. In this case, one can specify different mono-



**Fig. 6** Number of deer-vehicle collisions (DVCs) in Bavaria, Germany, for each day of the year 2006. Superimposed lines depict the conditional quantiles (5%, 10%, 25%, 50%, 75%, 90%, 95%) of this distribution for each day of the year.

tonicity constraints for each variable separately. However, it is ensured that the resulting interaction surface is monotonic (as specified). Monotonicity constraints might be especially useful in, but are not necessarily restricted to, data sets with relatively few observations or noisy data. Furthermore, as seen in the São Paulo air pollution data, the introduction of monotonicity constraints can help to estimate more appropriate models that can be interpreted. In the context of conditional transformation models monotonicity constraints are an essential ingredient as we try to estimate distribution functions, which are per definition monotonic.

Cyclic estimates can be easily used to model, for example, seasonal effects. The resulting estimate is a smooth effect estimate, where the boundaries are smoothly matched. Cyclic effects can be applied straightforwardly to model surfaces where the boundaries in each direction should match if cyclic tensor product P-splines are used. One possible application was given in the analysis of roe deer activity, where a varying coefficient model with smooth, time-varying effect and additional cyclic constraint was used. The idea of cyclic effects could also be extended to ordinal covariates with a temporal, periodic effect — such as days of the week. As for smooth, cyclic effects, the estimate should not depend on the ‘gap’. It should not matter if Sunday is chosen as the first day of the week (as in the USA) or if Monday is chosen (as common in Europe). If no penalty is applied, this naturally holds. However, we would often like to avoid large jumps from day to day as these are not reasonable. Using the ordinal penalty as introduced Hofner *et al.* (2011a) resolves this problem. At the same time, however, the first and last category play a special role as they do not have two neighbors in this coding. Introducing a cyclic penalty as

for the smooth effects would help in solving this problem and would stabilize the estimation of cyclic ordinal variables.

Finally, both restrictions — monotonic and cyclic constraints — can be mixed in one model: Some of the covariates are monotonicity restricted, others have cyclic constraints and the rest is modeled, for example, as smooth effects without further restrictions or as linear effects. Such a model was used to model the São Paulo air pollution data. We can even combine cyclic constraints for one variable and monotonicity constraints for the other variable to fit a bivariate effect as shown in the deer–vehicle collision example. The result is an estimate that is monotone in  $y$  given the day and has a periodic behavior over the days given the quantile (specified by  $y$ ).

Both monotonic P-splines and cyclic P-splines integrate seamlessly in the functional gradient descent boosting approach as implemented in **mboost** (Hofner et al., 2012; Hothorn et al., 2013a). This allows a single framework for fitting possible complex models. Additionally, the idea of asymmetric penalties for adjacent coefficients can be transferred from P-splines to ordinal factors (Hofner et al., 2011b), which can be integrated in the boosting framework as well.

**Acknowledgements** We thank the “Laboratório de Poluição Atmosférica Experimental, Faculdade de Medicina, Universidade de São Paulo, Brasil”, and Julio M. Singer for letting us use the data on air pollution in São Paulo. We thank Marco Heurich from the Bavarian Forest National Park, Grafenau, Germany, for the roe deer activity data, and Karen A. Brune for linguistic revision of the manuscript.

## References

- Ayer M, Brunk HD, Ewing GM, Reid WT, Silverman E (1955). “An empirical distribution function for sampling with incomplete information.” *Annals of Mathematical Statistics*, **26**, 641–647.
- Bacchetti P (1989). “Additive isotonic models.” *Journal of the American Statistical Association*, **84**, 289–294.
- Barlow RE, Bartholomew DJ, Bremner JM, Brunk HD (1972). *Statistical inference under order restrictions: The theory and application of isotonic regression*. John Wiley & Sons, London.
- Beliakov G (2000). “Shape preserving approximation using least squares splines.” *Approximation Theory and its Applications*, **16**(4), 80–98.
- Bollaerts K, Eilers PHC, van Mechelen I (2006). “Simple and multiple P-splines regression with shape constraints.” *British Journal of Mathematical and Statistical Psychology*, **59**, 451–469.
- Brunk HD (1955). “Maximum likelihood estimates of monotone parameters.” *Annals of Mathematical Statistics*, **26**, 607–616.
- Bühlmann P, Hothorn T (2007). “Boosting algorithms: Regularization, prediction and model fitting.” *Statistical Science*, **22**, 477–505.
- Bühlmann P, Yu B (2003). “Boosting with the  $L_2$  loss: Regression and classification.” *Journal of the American Statistical Association*, **98**, 324–339.
- Conceição GMS, Miraglia SGEK, Kishi HS, Saldiva PHN, Singer JM (2001). “Air pollution and child mortality: A time-series study in São Paulo, Brazil.” *Environmental Health Perspectives*, **109**, 347–350.
- de Leeuw J, Hornik K, Mair P (2009). “Isotone optimization in R: Pool-adjacent-violators algorithm (PAVA) and active set methods.” *Journal of Statistical Software*, **32**, 5.
- Dette H, Neumeyer N, Pilz KF (2006). “A simple non-parametric estimator of a strictly monotone regression function.” *Bernoulli*, **12**, 469–490.
- Dette H, Scheder R (2006). “Strictly monotone and smooth nonparametric regression for two or more variables.” *Canadian Journal of Statistics*, **34**, 535–562.
- Eilers PHC (2005). “Unimodal smoothing.” *Journal of Chemometrics*, **19**, 317–328.
- Eilers PHC, Gampe J, Marx BD, Rau R (2008). “Modulation models for seasonal time series and incidence tables.” *Statistics in Medicine*, **27**, 3430–3441.
- Eilers PHC, Marx BD (1996). “Flexible smoothing with B-splines and penalties (with discussion).” *Statistical Science*, **11**, 89–121.
- Eilers PHC, Marx BD (2010). “Splines, knots, and penalties.” *Wiley Interdisciplinary Reviews: Computational Statistics*, **2**, 637–653.
- Fahrmeir L, Kneib T, Lang S (2004). “Penalized structured additive regression: A Bayesian perspective.” *Statistica Sinica*, **14**, 731–761.
- Fang Z, Meinshausen N (2011). “LASSO ISOTone for high dimensional additive isotonic regression.” *Journal of Computational and Graphical Statistics*. To appear, URL <http://arxiv.org/abs/1006.2940v1>.
- Fenske N, Kneib T, Hothorn T (2011). “Identifying risk factors for severe childhood malnutrition by boosting additive quantile regression.” *Journal of the American Statistical Association*, **106**, 494–510.
- Fok CCT, Ramsay JO (2006). “Fitting curves with periodic and nonperiodic trends and their interactions with intensive longitudinal data.” In “Models for Intensive Longitudinal Data,” pp. 109–123. Oxford University Press.
- Friedman JH (1991). “Multivariate adaptive regression splines.” *The Annals of Statistics*, **19**, 1–50.

- Gentle JE (2007). *Matrix algebra: Theory, computations, and applications in statistics*. Springer, New York.
- Goldfarb D, Idnani A (1982). “Dual and primal-dual methods for solving strictly convex quadratic programs.” In “Numerical Analysis,” pp. 226–239. Springer-Verlag, Berlin.
- Goldfarb D, Idnani A (1983). “A numerically stable dual method for solving strictly convex quadratic programs.” *Mathematical Programming*, **27**, 1–33.
- Gu C (2002). *Smoothing spline ANOVA models*. Springer, New York.
- Harvey A, Koopman SJ (1993). “Forecasting hourly electricity demand using time-varying splines.” *Journal of the American Statistical Association*, **88**, 1228–1236.
- Harvey A, Koopman SJ, Riani M (1997). “The modeling and seasonal adjustment of weekly observations.” *Journal of Business & Economic Statistics*, **15**, 354–368.
- Harville DA (2008). *Matrix algebra from a statisticians perspective*. Springer, New York.
- Hastie T, Tibshirani R (1986). “Generalized additive models.” *Statistical Science*, **1**, 297–310.
- Hastie T, Tibshirani R (1990). *Generalized additive models*. Chapman & Hall / CRC, London.
- Hastie T, Tibshirani R (1993). “Varying-coefficient models.” *Journal of the Royal Statistical Society. Series B*, **55**, 757–796.
- He X, Ng P (1999). “COBS: qualitatively constrained smoothing via linear programming.” *Computational Statistics*, **14**, 315–338.
- He X, Shi P (1998). “Monotone B-spline smoothing.” *Journal of the American Statistical Association*, **93**.
- Heckman NE, Ramsay JO (2000). “Penalized regression with model-based penalties.” *The Canadian Journal of Statistics*, **28**, 241–258.
- Hofner B (2011). *Boosting in structured additive models*. Ph.D. thesis, LMU Mnchen. Verlag Dr. Hut, Mnchen, URL <http://nbn-resolving.de/urn:nbn:de:bvb:19-138053>.
- Hofner B, Hothorn T, Kneib T, Schmid M (2011a). “A framework for unbiased model selection based on boosting.” *Journal of Computational and Graphical Statistics*, **20**, 956–971.
- Hofner B, Mayr A, Robinzonov N, Schmid M (2012). “Model-based boosting in R – A hands-on tutorial using the R package mboost.” *Computational Statistics*. doi:10.1007/s00180-012-0382-5. Online First.
- Hofner B, Müller J, Hothorn T (2011b). “Monotonicity-constrained species distribution models.” *Ecology*, **92**, 1895–1901.
- Hothorn T, Brandl R, Müller J (2012). “Large-scale model-based assessment of deer-vehicle collision risk.” *PLoS ONE*, **7**(2), e29510. doi:10.1371/journal.pone.0029510.
- Hothorn T, Bühlmann P, Kneib T, Schmid M, Hofner B (2013a). *mboost: Model-based boosting*. R package version 2.2-2, URL <http://CRAN.R-project.org/package=mboost>.
- Hothorn T, Kneib T, Bühlmann P (2013b). “Conditional transformation models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. doi:10.1111/rssb.12017.
- Jones RH, Brelford WM (1967). “Time series with periodic structure.” *Biometrika*, **54**, 403–408.
- Kneib T, Hothorn T, Tutz G (2009). “Variable selection and model choice in geosadditive regression models.” *Biometrics*, **65**, 626–634.
- Koenker R, Ng P (2005a). “Frisch-Newton algorithm for sparse quantile regression.” *Acta Mathematicae Applicatae Sinica, English Series*, **21**, 225–236.
- Koenker R, Ng P (2005b). “Inequality constrained quantile regression.” *Sankhya: The Indian Journal of Statistics (2003-)*, **67**, 418–440.
- Leitenstorfer F, Tutz G (2007). “Generalized monotonic regression based on B-splines with an application to air pollution data.” *Biostatistics*, **8**, 654–673.
- Lewis PAW, Stevens JG (1991). “Nonlinear modeling of time series using multivariate adaptive regression splines (MARS).” *Journal of the American Statistical Association*, **86**, 864–877.
- Ng P, Maechler M (2007). “A fast and efficient implementation of qualitatively constrained quantile smoothing splines.” *Statistical Modeling*, **7**, 315–328.
- Ramsay JO (1988). “Monotone regression splines in action.” *Statistical Science*, **3**, 425–441.
- Ramsay JO, Dalzell CJ (1991). “Some tools for functional data analysis.” *Journal of the Royal Statistical Society. Series B*, **53**, 539–572.
- Saldiva P, Pope CI, Schwartz J, Dockery D, Lichtenfels A, Salge J, Barone I, Bohm G (1995). “Air pollution and mortality in elderly people: A time-series study in São Paulo, Brazil.” *Archives of Environmental Health*, **50**, 159–164.
- Schmid M, Hothorn T (2008). “Boosting additive models using component-wise P-splines.” *Computational Statistics & Data Analysis*, **53**, 298–311.
- Sobotka F, Mirkov R, Hofner B, Eilers P, Kneib T (2013). “Modelling flow in gas transmission networks using shape-constrained expectile regression.” To appear in . . .
- Spix C, Heinrich J, Dockery D, Schwartz J, Vlksch G, Schwinkowski K, Clen C, Wichmann HE (1993). “Air pollution and daily mortality in Erfurt, East

- Germany, 1980-1989." *Environmental Health Perspectives*, **101**, 518–526.
- Stache A, Heller E, Hothorn T, Heurich M (2013). "Activity patterns of European roe deer (*Capreolus capreolus*) are strongly influenced by individual behaviour." *Folia zoologica*, **62**, 67 – 75.
- Tutz G, Binder H (2006). "Generalized additive modelling with implicit variable selection by likelihood-based boosting." *Biometrics*, **62**, 961–971.
- Tutz G, Leitenstorfer F (2007). "Generalized smooth monotonic regression in additive modeling." *Journal of Computational and Graphical Statistics*, **16**, 165–188.
- van Eeden C (1956). "Maximum likelihood estimation of ordered probabilities." *Koninkl. Nederl. Akademie van Wetenschappen, Proc. Ser. A.*, **59**, 444–455.
- Welham SJ, Cullis BR, Kenward MG, Thompson R (2006). "The analysis of longitudinal data using mixed model l-splines." *Biometrics*, **62**, 392–401.
- Wood SN (2006a). *Generalized additive models: An introduction with R*. Chapman & Hall / CRC, London.
- Wood SN (2006b). "Low-rank scale-invariant tensor product smooths for generalized additive mixed models." *Biometrics*, **62**, 1025–1036.
- Wood SN (2008). "Fast stable direct fitting and smoothness selection for generalized additive models." *Journal of the Royal Statistical Society. Series B*, **70**, 495–518.
- Wood SN (2010). *mgcv: GAMs with GCV/AIC/REML smoothness estimation and GAMMs by PQL*. R package version 1.7-2, URL <http://CRAN.R-project.org/package=mgcv>.
- Zhang D, Lin X, Sowers M (2000). "Semiparametric regression for periodic longitudinal hormone data from multiple menstrual cycles." *Biometrics*, **56**, 31–39.
- Zhang D, Lin X, Sowers M (2007). "Two-stage functional mixed models for evaluating the effect of longitudinal covariate profiles on a scalar outcome." *Biometrics*, **63**, 351–362.