

Generalized functional additive mixed models

Fabian Scheipl

*Institut für Statistik
Ludwig-Maximilians-Universität München
Ludwigstrasse 33
Munich, Germany*
e-mail: fabian.scheipl@stat.uni-muenchen.de

Jan Gertheiss

*Institut für Angewandte Stochastik und Operations Research
TU Clausthal
Erzstrasse 1
Clausthal-Zellerfeld, Germany*
e-mail: jan.gertheiss@tu-clausthal.de

and

Sonja Greven

*Institut für Statistik
Ludwig-Maximilians-Universität München
Ludwigstrasse 33
Munich, Germany*
e-mail: sonja.greven@stat.uni-muenchen.de

Abstract: We propose a comprehensive framework for additive regression models for non-Gaussian functional responses, allowing for multiple (partially) nested or crossed functional random effects with flexible correlation structures for, e.g., spatial, temporal, or longitudinal functional data as well as linear and nonlinear effects of functional and scalar covariates that may vary smoothly over the index of the functional response. Our implementation handles functional responses from any exponential family distribution as well as many others like Beta- or scaled and shifted t -distributions. Development is motivated by and evaluated on an application to large-scale longitudinal feeding records of pigs. Results in extensive simulation studies as well as replications of two previously published simulation studies for generalized functional mixed models demonstrate the good performance of our proposal. The approach is implemented in well-documented open source software in the `pffr` function in R-package `refund`.

Received October 2015.

1. Introduction

Data sets in which measurements consist of curves or images instead of scalars – i.e., functional data – are becoming ever more common in many areas of

application. This is due to the increasing affordability and deployment of sensors like accelerometers or spectrometers, high-throughput imaging technologies and automated logging equipment that continuously records conditions over time. Recent methodological development in this area has been rapid and intense, see Morris [13] for a review of the state of the art for regression models for functional data.

In this work, we extend the general framework for functional additive mixed models for potentially correlated functional Gaussian responses described in Scheipl et al. [19] to non-Gaussian functional responses. The development is motivated by and evaluated on an animal husbandry dataset in which the feeding behavior of growing-finishing pigs was monitored continuously over 3 months [11, 3]. These are non-Gaussian functional data in the sense that the underlying probability of feeding is assumed to be a continuous function over time, while the available data are sequences of binary indicators (“feeding: yes/no”) evaluated at the temporal resolution of the sensors. Aggregating these binary indicators over given time intervals, we get time series of counts or proportions for which (truncated) Poisson, Negative Binomial, Binomial or Beta distributional assumptions could be appropriate. Another example of non-Gaussian functional data with large practical relevance would be continuously valued functional responses with heavy-tailed measurement errors for which a scaled t -distribution could be appropriate.

We briefly summarize the most relevant prior work on regression models for non-Gaussian functional responses. The fundamental work of Hall et al. [6] relates observed functional binary or count data to a latent Gaussian process (GP) through a link function. This is the underlying idea of almost all the works that follow. They differ primarily in 1) how the latent functional processes are represented (i.e., either spline, wavelet or functional principal component (FPC) representations or full GP models), 2) to what extent additional covariate information can be included, 3) whether they allow the modelling of dependencies between and along the functional responses, 4) which distributions are available for the responses, 5) whether the functional data has to be available on a joint regular grid, and 6) in the availability of documented and performant software implementations. In Hall et al. [6], the latent GP is represented in terms of its FPCs and neither correlated responses nor covariates are accommodated. No implementation is publicly available. A Bayesian variant of Hall et al. [6] is provided by van der Linde [25]. Zhu et al. [32] describe a robustified wavelet-based functional mixed model with scalar covariates for error-contaminated continuous functional responses on regular grids. No implementation was publicly available at the time of writing. Serban et al. [20] extend the approach of Hall et al. [6] to multilevel binary data without covariates and provide a rudimentary problem-specific implementation. Wang and Shi [26] describe an empirical Bayesian approach for latent Gaussian process regression models for data from exponential families with scalar and concurrent functional covariates. The available implementation does not accommodate covariate effects and is limited to binary data with curve-specific i.i.d. functional random effects, see Section 4.3 for a systematic comparison with our proposal based on a replication of their simulation

study. Li et al. [9] present a model for concurrent binary and continuous functional observations. Association between the two is modeled via cross-correlated (latent) FPC scores and cannot take into account any other covariate effects or dependency structures. A similar approach is described in Tidemann-Miller [24, ch. 3]. Goldsmith et al. [4] develop a fully Bayesian approach which uses latent FPCs in a spline basis representation to represent multilevel functional random effects and linear functional effects of scalar covariates. They provide an implementation for binary outcomes with a logit link function, see Section 4.4 for a systematic comparison with our proposal based on a replication of their simulation study. Gertheiss et al. [3] describe a marginal GEE-type approach for correlated binary functional responses with concurrent functional covariates. Brockhaus et al. [2] develop a framework for estimating flexible functional regression models via boosting with similar flexibility as ours, implemented in R [15] package `FDboost` [1]. The package additionally implements functional quantile regression models, which are not included in the framework we present here. Since their implementation is based on a component-wise gradient boosting algorithm, they cannot provide hypothesis tests or resampling-free construction of confidence intervals, however, and also have to rely on computationally intensive resampling methods for hyperparameter tuning.

Compared to previous work, the novel contribution of this work is the development, implementation and evaluation of a comprehensive maximum likelihood-based inferential framework for generalized functional additive mixed models (GFAMMs) for potentially correlated functional responses. Our proposal accommodates diverse latent-scale correlation structures as well as flexible modeling of the conditional mean structure with multiple linear and non-linear effects of both scalar and functional covariates and their interactions. Our proposal is implemented in the full generality described here for both regular grid data and sparse or irregularly observed functional responses in the `pffr` function in R package `refund` [7]. Available response distributions include all exponential family distributions as well as Beta, scaled and shifted t -, Negative Binomial, Tweedie and zero-inflated Poisson distributions, each with various link functions, as well as cumulative threshold models for ordered categorical responses. With the exception of Brockhaus et al. [2], none of the previous proposals in this area achieve anything close to this level of generality, not to mention offer publicly available, widely applicable open-source implementations. Since our framework is a natural extension of previous work done on generalized additive mixed models for scalar data and builds on the high performing, flexible implementations available for them, we can directly make use of many results from this literature, such as improved confidence intervals and tests for smooth effects [10, 29].

The remainder of this work is structured as follows: Section 2 introduces notation and the theoretical and inferential framework for our model class. Section 3 presents results for our application. Section 4 summarizes results of our extensive validation on synthetic data and of our partial replication of the simulation studies of Wang and Shi [26] and Goldsmith et al. [4]. Section 5 concludes.

2. Model

2.1. Model structure

In what follows, we discuss structured additive regression models of the general form

$$y_i(t) \sim \mathcal{F}(\mu_i(t), \boldsymbol{\nu})$$

$$g(\mu_i(t)) = \eta_i(t) = \sum_{r=1}^R f_r(\mathcal{X}_{ri}, t), \tag{1}$$

where, for each t , $y_i(t)$, $i = 1, \dots, n$, is a random variable from some distribution \mathcal{F} with conditional expectation $E(y_i(t)|\mathcal{X}_i, t, \boldsymbol{\nu}) = \mu_i(t)$ observed over a domain \mathcal{T} and an optional vector of nuisance parameters $\boldsymbol{\nu}$. We use $g(\cdot)$ to denote the known link function. Our implementation allows analysts to choose \mathcal{F} from the exponential family distributions as well as Tweedie, Negative Binomial, Beta, ordered categorical, zero-inflated Poisson and scaled and shifted t -distributions. Note that, in the special case of ordered categorical responses, $\mu_i(t)$ is not the conditional mean of the response itself, but that of a latent variable whose value determines the response category.

Each term in the additive predictor is a function of a) the index t of the response and b) a subset \mathcal{X}_r of the complete covariate set \mathcal{X} potentially including scalar and functional covariates and (partially) nested or crossed grouping factors. Note that the definition also includes functional random effects $b_g(t)$ for a grouping variable g with M levels. These are modeled as realizations of a mean-zero Gaussian random process on $\{1, \dots, M\} \times \mathcal{T}$ with a general covariance function $K^b(m, m', t, t') = \text{Cov}(b_{g,m}(t), b_{g,m'}(t'))$ that is smooth in t , where m, m' denote different levels of g . Table 1 shows how a selection of the most frequently used effect types fit into this framework.

TABLE 1
A selection of possible model terms $f_r(\mathcal{X}_r, t)$ for model (1). All effects can be constant in t as well.

\mathcal{X}_r	type of effect	$f_r(\mathcal{X}_r, t)$
\emptyset (none)	smooth intercept	$\beta_0(t)$
scalar covariate z	linear or smooth effect	$z\beta(t), f(z, t)$
two scalars z_1, z_2	linear or smooth interaction	$z_1z_2\beta(t),$ $z_1f(z_2, t),$ $f(z_1, z_2, t)$
functional covariate $x(s)$	linear or smooth (historical) functional effect	$\int x(s)\beta(s, t)ds,$ $\int_{l(t)}^{u(t)} x(s)\beta(s, t)ds,$ $\int F(x(s), s, t)ds$
functional covariate $v(t)$	concurrent effects	$v(t)\beta(t), f(v(t), t)$
functional covariates $v(t), w(t)$	concurrent interactions	$v(t)w(t)\beta(t),$ $f(v(t), w(t), t)$
grouping variable g	functional random intercept	$b_g(t)$
grouping variable g , scalar z	functional random slope	$zb_g(t)$
curve indicator i	smooth functional residual	$e_i(t)$

We approximate each term $f_r(\mathcal{X}_r, t)$ by a linear combination of basis functions given by the tensor product of marginal bases for \mathcal{X}_r and t . Since the basis has to be rich enough to ensure sufficient flexibility, Section 2.4 describes a penalized likelihood approach that stabilizes estimates by suppressing variability of the effects that is not strongly supported by the data and finds a data-driven compromise between goodness of fit and simplicity of the fitted effects.

2.2. Data and notation

In practice, functional responses $y_i(t)$ are observed on a grid of T_i points $\mathbf{t}_i = (t_{i1}, \dots, t_{iT_i})^\top$ which can be irregular and/or sparse. Let $y_{il} = y_i(t_l)$ and $\mathbf{y}_i = (y_{i1}, \dots, y_{iT_i})$. To fit the model, we form $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top)^\top$ and $\mathbf{t} = (\mathbf{t}_1^\top, \dots, \mathbf{t}_n^\top)^\top$, two $N = \sum_{i=1}^n T_i$ -vectors that contain the concatenated observed responses and their argument values, respectively. Let \mathcal{X}_{ri} contain the observed values of \mathcal{X}_r associated with a given \mathbf{y}_i . Model (1) can then be expressed as

$$y_{il} \sim \mathcal{F}(\mu_{il}, \boldsymbol{\nu})$$

$$g(\mu_{il}) = \sum_{r=1}^R f_r(\mathcal{X}_{ri}, t_{il}) \tag{2}$$

for $i = 1, \dots, n$ and $l = 1, \dots, T_i$. Let $f(\mathbf{t})$ denote the vector of function evaluations of f for each entry in the vector \mathbf{t} and let $f(\mathbf{x}, \mathbf{t})$ denote the vector of evaluations of f for each combination of rows in the vectors or matrices \mathbf{x}, \mathbf{t} . In the following, we let $T_i \equiv T$ to simplify notation, but our approach is equally suited to data on irregular grids. For regular grids, each observed value in any \mathcal{X}_r that is constant over t is simply repeated T times to match up with the corresponding entry in the $N = nT$ -vector \mathbf{y} .

2.3. Tensor product representation of effects

We approximate each term $f_r(\mathcal{X}_r, t)$ by a linear combination of basis functions defined on the product space of the two spaces: one for the covariates in \mathcal{X}_r and one over \mathcal{T} , where each marginal basis is associated with a corresponding marginal penalty. A very versatile method to construct basis function evaluations on such a joint space is given by the row tensor product of marginal bases evaluated on \mathcal{X}_r and \mathbf{t} [e.g. 27, ch. 4.1.8]. Let $\mathbf{1}_d = (1, \dots, 1)^\top$ denote a d -vector of ones. The row tensor product of an $m \times a$ matrix \mathbf{A} and an $m \times b$ matrix \mathbf{B} is defined as the $m \times ab$ matrix $\mathbf{A} \odot \mathbf{B} = (\mathbf{A} \otimes \mathbf{1}_b^\top) \cdot (\mathbf{1}_a^\top \otimes \mathbf{B})$, where \otimes denotes the Kronecker product and \cdot denotes element-wise multiplication. Specifically, for each of the terms,

$$f_r(\mathcal{X}_r, \mathbf{t}) \approx \underbrace{(\boldsymbol{\Phi}_{\mathbf{x}r})}_{N \times K_{\mathbf{x}r}} \odot \underbrace{(\boldsymbol{\Phi}_{\mathbf{t}r})}_{N \times K_{\mathbf{t}r}} \underbrace{\boldsymbol{\theta}_r}_{K_{\mathbf{x}r} K_{\mathbf{t}r} \times 1} = \boldsymbol{\Phi}_r \boldsymbol{\theta}_r, \tag{3}$$

where $\Phi_{\mathbf{x}_r}$ contains the evaluations of a suitable marginal basis for the covariate(s) in \mathcal{X}_r and $\Phi_{\mathbf{t}_r}$ contains the evaluations of a marginal basis in \mathbf{t} with K_{x_r} and K_{t_r} basis functions, respectively. The shape of the function is determined by the vector of coefficients θ_r . A corresponding penalty term can be defined using the Kronecker sum of the marginal penalty matrices $\mathbf{P}_{\mathbf{x}_r}$ and $\mathbf{P}_{\mathbf{t}_r}$ associated with each basis [27, ch. 4.1], i.e.

$$\begin{aligned} \text{pen}(\theta_r | \lambda_{t_r}, \lambda_{x_r}) &= \theta_r^T \mathbf{P}_r(\lambda_{t_r}, \lambda_{x_r}) \theta_r, \\ \text{where } \mathbf{P}_r(\lambda_{t_r}, \lambda_{x_r}) &= \lambda_{x_r} \mathbf{P}_{\mathbf{x}_r} \otimes \mathbf{I}_{K_{t_r}} + \lambda_{t_r} \mathbf{I}_{K_{x_r}} \otimes \mathbf{P}_{\mathbf{t}_r}. \end{aligned} \quad (4)$$

$\mathbf{P}_{\mathbf{x}_r}$ and $\mathbf{P}_{\mathbf{t}_r}$ are known and fixed positive semi-definite penalty matrices and λ_{t_r} and λ_{x_r} are positive smoothing parameters controlling the trade-off between goodness of fit and the smoothness of $f_r(\mathcal{X}_r, \mathbf{t})$ in \mathcal{X}_r and \mathbf{t} , respectively.

This approach is extremely versatile and powerful as it allows analysts to pick and choose bases and penalties best suited to the problem at hand. Any basis and penalty over \mathcal{T} (for example, incorporating monotonicity or periodicity constraints) can be combined with any basis and penalty over \mathcal{X}_r . We provide some concrete examples for frequently encountered effects: For a functional intercept $\beta_0(t)$, $\Phi_{\mathbf{x}_r} = \mathbf{1}_N$ and $\mathbf{P}_{\mathbf{x}_r} = \mathbf{0}$. For a linear functional effect of a scalar covariate $z\beta(t)$, $\Phi_{\mathbf{x}_r} = \mathbf{z} \otimes \mathbf{1}_T$ and $\mathbf{P}_{\mathbf{x}_r} = \mathbf{0}$, where $\mathbf{z} = (z_1, \dots, z_n)^T$ contains the observed values of a scalar covariate z . For a smooth functional effect of a scalar covariate $f(z, t)$, $\Phi_{\mathbf{x}_r} = [\phi_h(z_i)]_{\substack{1 \leq i \leq n \\ 1 \leq h \leq K_{x_r}}} \otimes \mathbf{1}_T$ and $\mathbf{P}_{\mathbf{x}_r}$ is the penalty associated with the basis functions $\phi_h(z)$.

For a linear effect of a functional covariate $\int_{l(t)}^{u(t)} x(s)\beta(s, t)ds$,

$$\Phi_{\mathbf{x}_r} = \begin{pmatrix} \mathbf{W}_s \cdot \mathbf{X} \\ N \times S \quad N \times S \quad S \times K_{x_r} \end{pmatrix} \Phi_s,$$

where \mathbf{W}_s contains suitable quadrature weights for the numerical integration and zeroes for combinations of s and t not inside the integration range $[l(t), u(t)]$. $\mathbf{X} = [x_i(s_k)]_{\substack{1 \leq i \leq n \\ 1 \leq k \leq S}} \otimes \mathbf{1}_T$ contains the repeated evaluations of $x(s)$ on the observed grid (s_1, \dots, s_S) and $\Phi_s = [\phi_h(s_k)]_{\substack{1 \leq k \leq S \\ 1 \leq h \leq K_{x_r}}}$ the evaluations of basis functions $\phi_h(s)$ for the basis expansion of $\beta(s, t)$ over s . $\mathbf{P}_{\mathbf{x}_r}$ is again simply the penalty matrix associated with the basis functions $\phi_h(s)$.

For a functional random slope $z b_g(t)$ for a grouping factor g with $K_{x_r} = M$ levels, $\Phi_{\mathbf{x}_r} = (\text{diag}(\mathbf{z})\mathbf{G}) \otimes \mathbf{1}_T$, where $\mathbf{G} = [\delta(g_i = m)]_{\substack{1 \leq i \leq n \\ 1 \leq m \leq M}}$ is an incidence matrix mapping each observation to its group level. $\mathbf{P}_{\mathbf{x}_r}$ is typically \mathbf{I}_M for i. i. d. random effects, but can also be any inverse correlation or covariance matrix defining the expected similarity of levels of g with each other, defined for example via spatial or temporal correlation functions or other similarity measures defined between the levels of g .

In general, $\Phi_{\mathbf{t}_r}$ and $\mathbf{P}_{\mathbf{t}_r}$ can be any matrix of basis function evaluations $\Phi_{\mathbf{t}_r} = \mathbf{1}_n \otimes [\phi_k(t_l)]_{\substack{1 \leq l \leq T \\ 1 \leq k \leq K_{t_r}}}$ over t with a corresponding penalty matrix. For effects that are constant over t , we simply use $\Phi_{\mathbf{t}_r} = \mathbf{1}_N$ and $\mathbf{P}_{\mathbf{t}_r} = \mathbf{0}$.

Note that for data on irregular t -grids, the stacking via " $\otimes \mathbf{1}_T$ " in the expressions above simply has to be replaced by using a different suitable repetition

pattern. Scheipl et al. [19] discuss tensor product representations of many additional effects available in our implementation.

As for other penalized spline models, fits are usually robust against increases in K_{tr} and K_{xr} once a sufficiently large number of basis functions has been chosen since it is the penalty parameters $\lambda_{tr}, \lambda_{xr}$ that control the effective degrees of freedom of the fit, c.f. Ruppert [18], Wood [27, Ch. 4.1.7]. Pya and Wood [14] describe and implement a test procedure to determine sufficiently large basis dimensions that is also applicable to our model class.

2.4. Inference and implementation

Let $\Phi = [\Phi_1 | \dots | \Phi_R]$ contain the concatenated design matrices Φ_r associated with the different model terms and $\theta = (\theta_1^\top, \dots, \theta_R^\top)^\top$ the respective stacked coefficient vectors θ_r . Let $\lambda = (\lambda_{t1}, \lambda_{x1}, \dots, \lambda_{tR}, \lambda_{xR})$. The penalized log-likelihood to be maximized is given by

$$\ell_p(\theta, \lambda, \nu | \mathbf{y}) = \ell(\mu, \nu | \mathbf{y}) - \frac{1}{2} \sum_{r=1}^R \text{pen}(\theta_r | \lambda_{tr}, \lambda_{xr}) \tag{5}$$

where $\mu = g^{-1}(\Phi\theta)$ and $\ell(\mu, \nu | \mathbf{y}) = \sum_{i=1}^n \sum_{l=1}^{T_i} \ell(\mu_{il}, \nu | y_{il})$ is the log-likelihood function implied by the respective response distribution $\mathcal{F}(\mu_{il}, \nu)$ given in (2).

Recent results indicate that REML-based inference for this class of penalized regression models can be preferable to GCV-based optimization [16, 28, Section 4]. Wood [30] describes a numerically stable implementation for directly optimizing a Laplace-approximate marginal likelihood of the smoothing parameters λ which we summarize below. Note that this is equivalent to REML optimization in the Gaussian response case. The idea is to iteratively perform optimization over λ and then estimate basis coefficients θ given λ via standard penalized likelihood methods.

We first define a Laplace approximation $\ell_{LA}(\lambda, \nu | \mathbf{y}) \approx \int \ell_p(\theta, \lambda, \nu | \mathbf{y}) d\theta$ to the marginal (restricted) ML score suitable for optimization. Let $\mathbf{H} = -\partial^2 \ell_p(\theta, \lambda, \nu | \mathbf{y}) / (\partial\theta\partial\theta^T)$ denote the negative Hessian of the penalized log-likelihood (5), and let $\tilde{\theta} = \arg \max_{\theta} \ell_p(\theta, \lambda, \nu | \mathbf{y})$ for fixed λ, ν with $\tilde{\mu} = g^{-1}(\Phi\tilde{\theta})$. The approximate marginal (restricted) ML criterion can then be defined as

$$\ell_{LA}(\lambda, \nu | \mathbf{y}) = \ell(\tilde{\mu}, \nu | \mathbf{y}) - \frac{1}{2} \sum_{r=1}^R \left(\text{pen}(\tilde{\theta}_r | \lambda_{tr}, \lambda_{xr}) - \log |\mathbf{P}_r(\lambda_{tr}, \lambda_{xr})|^+ \right) - \frac{1}{2} \log |\mathbf{H}| + \frac{1}{2} d_{\emptyset} \log(2\pi), \tag{6}$$

where $|\mathbf{A}|^+$ denotes the generalized determinant, i.e., the product of the positive eigenvalues of \mathbf{A} , and d_{\emptyset} is the sum of the rank deficiencies of the $\mathbf{P}_r(\lambda_{tr}, \lambda_{xr})$, $r = 1, \dots, R$.

The R package `mgcv` [27] provides an efficient and numerically stable implementation for the iterative optimization of (6) that 1) finds $\tilde{\theta}$ for each candidate

λ via penalized iteratively re-weighted least squares, 2) computes the gradient and Hessian of $\ell_{LA}(\lambda, \nu | \mathbf{y})$ w.r.t. $\log(\lambda)$ via implicit differentiation, 3) updates $\log(\lambda)$ via Newton’s method and 4) estimates ν either as part of the maximization for λ or based on the Pearson statistic of the fitted model. See Wood [28, 30], Wood et al. [31] for technical details and numerical considerations.

Confidence intervals [10] and tests [29] for the estimated effects are based on the asymptotic normality of ML estimates. Confidence intervals use the variance of the distance between the estimators and true effects. Similar to the idea behind the MSE, this accounts for both the variance and the penalization induced bias (shrinkage effect) of the spline estimators.

We provide an implementation of the full generality of the proposed models in the `pffr()` function in the R package `refund` [7] and use the `mgcv` [28] implementation of (scalar) generalized additive mixed models as computational backend. The `pffr()` function provides a formula-based wrapper for the definition of a broad variety of functional effects as in Section 2.3 as well as convenience functions for summarizing and visualizing model fits and generating predictions.

3. Application to pigs’ feeding behavior

3.1. Data

We use data on pigs’ feeding behavior collected in the ICT-AGRI era-net project “PIGWISE”, funded by the European Union. In this project, high frequency radio frequency identification (HF RFID) antennas were installed above the troughs of pig barns to register the feeding times of pigs equipped with passive RFID tags [11]. We are interested in modelling the binary functional data that arises from measuring the proximity of the pig to the trough (yes-no) every 10 secs over 102 days available for each pig. The raw data for one such pig, called ‘pig 57’ whose behavior we analyse in depth is shown in Figure B.1 in Appendix B. Such models of (a proxy of) individual feeding behavior can be useful for ethology research as well as monitoring individual pigs’ health status and/or quality of the available feed stock, c.f. Gertheiss et al. [3]. Available covariates include the time of day, the age of the pig (i.e, the day of observation), as well as the barn’s temperature and humidity.

3.2. Model

Due to pronounced differences in feeding behavior between individual pigs [c.f. 3, Figure 4], we focus on pig-wise models and model the observed feeding rate for a single pig for each day $i = 1, \dots, 102$ as a smooth function over the daytime t . For our model, we aggregate the originally observed binary feeding indicators $\tilde{y}_i(t)$, which are measured every 10 seconds, over 10 min intervals into $y_i(t)$. This temporal resolution is sufficient for the purposes of our modeling effort. We then model

$$y_i(t) | \mathcal{X}_i \sim \text{Bin}(60, \pi = \mu_i(t));$$

$$\mu_i(t) = \text{logit}^{-1} \left(\sum_{r=1}^R f_r(\mathcal{X}_{ri}, t) \right).$$

As is typical for applied problems such as this, many different distributional assumptions for $y_i(t)|\mathcal{X}_i$ and specifications of $\sum_{r=1}^R f_r(\mathcal{X}_{ri}, t)$ are conceivable. In this case, both Beta and Negative Binomial distributions yielded less plausible and more unstable fits than the Binomial distribution. With respect to the additive predictor, the effect of temperature or humidity (i.e., $\mathcal{X}_{ri} = \text{hum}_i(t)$), for example, could be modeled as a linear functional effect $\int_{l(t)}^{u(t)} \text{hum}(s)\beta(s, t)ds$, which represents the cumulative effect of humidity exposure over the time window $l(t) \leq t \leq u(t)$, or the concurrent effect of humidity, either linearly via $\text{hum}(t)\beta_h(t)$ or non-linearly via $f(\text{hum}(t), t)$. Auto-regressive and lagged effects of previous feeding behavior (i.e., $\mathcal{X}_{ri} = y_i(t)$) could be modeled similarly either as the cumulative effect of previous feeding over a certain time window ($\int_{t-\delta}^{t-10\text{min}} y(s)\beta(s, t)ds$) or nonlinear ($f(y_i(t-\delta), t)$) or linear ($y_i(t-\delta)\beta(t)$) effects of the lagged responses with a pre-specified time lag δ . Finally, our implementation also offers diverse possibilities for accounting for aging effects and day-to-day variations in behavior (i.e., $\mathcal{X}_{ri} = i$). Aging effects that result in a gradual change of daily feeding rates over time could be represented as $i\beta(t)$ for a gradual linear change or as $f(i, t)$ for a nonlinear smooth change over days i . Less systematic day-to-day variations can be modeled as daily functional random intercepts $b_i(t)$, potentially auto-correlated over i to encourage similar shapes for neighboring days. In this application, the performances of most of these models on the validation set were fairly similar and the subsequent section presents results for a methodologically interesting model that was among the most accurate models on the validation set.

3.3. Results

In what follows, we present detailed results for an auto-regressive binomial logit model with smoothly varying day effects for pig 57 (see Figure B.1, Appendix B):

$$\text{logit}(\mu_i(t)) = \beta_0(t) + f(i, t) + \int_{t-3\text{h}}^{t-10\text{min}} y_i(s)\beta(t, s)ds.$$

This model assumes that feeding behavior during the previous 3 hours affects the current feeding rate. We use periodic P-spline bases over t to enforce similar $f_r(\mathcal{X}_{ri}, t)$ for $t = 00:00\text{h}$ and $t = 23:59\text{h}$ for all r . The term $f(i, t)$ can be interpreted as a non-linear aging effect on the feeding rates or as a functional random day effect whose shapes vary smoothly across days i . This model contains $\dim(\boldsymbol{\lambda}) = 5$ smoothing parameters (1 for the functional intercept, 2 each for the daily and auto-regressive effects), $\dim(\boldsymbol{\theta}) = 106$ spline coefficients (24 for the intercept, 25 for the autoregressive effect, 56 for the day effect) and $N = nT = 9648$ observations in the training data. We leave out every third day to serve as external validation data for evaluating the generalizability of the

model. The fit takes about 1 minute on a modern desktop PC. Appendix B contains detailed results for an alternative model specification, Section 4.1 describes results for synthetic data with similar structure and size.

Beside giving insight into pigs' feeding patterns and how these patterns change as the pigs age, the model reported here enables short-term predictions of feeding probabilities for a given pig based on its previous feeding behavior. This could be very helpful when using the RFID system for surveillance and early identification of pigs showing unusual feeding behavior. Unusual feeding behavior is then indicated by model predictions that are consistently wrong; i.e., the pig is not behaving as expected. Such discrepancies can then indicate problems such as disease or low-quality feed stock. For the auto-regressive model discussed here, only very short-term predictions 10 minutes into the future can be generated easily as $y_i(t - 10\text{min})$ is required as a covariate value. Other model formulations without such auto-regressive terms or larger lead times of auto-regressive terms will allow more long-term forecasting. More long-term forecasts for auto-regressive models could also be achieved by using predictions $\hat{y}_i(t)$ instead of observed values of $y_i(t)$ as inputs for the forecast. These can be generated sequentially for timepoints farther and farther in the future, but errors and uncertainties will accumulate correspondingly. A detailed analysis of this issue is beyond the scope of this paper, Tashman [22] gives an overview. In addition to the lagged response value, also the aging effect $f(i, t)$ is needed. Here, we could either use the value from the preceding day $i - 1$, if it can be assumed that the pig's (expected) behavior does not change substantially from one day to the next; or do some extrapolation.

Figure 1 shows fitted and observed values for 12 selected days from the training data (top) as well as estimated and observed values for 12 selected days from the validation data (bottom). The model is able to reproduce many of the feeding episodes, in the sense that peaks in the estimated probability curves mostly line up well with observed spikes of $y_i(t)$. This model explains about 24% of the deviance and, taking the mean over all days and time-points, achieves a Brier score of about 0.024 on both the training data and the validation data, i.e., we see no evidence for overfitting with this model specification. This model performs somewhat better on the validation data than an alternative model specification with i. i. d. random functional day effects $b_i(t)$, c.f. Appendix B, Figure B.2.

Figure 2 shows the estimated components of the additive predictor for the model with smoothly varying day effects. The functional intercept in the left top panel of Figure 2 shows two clear peaks of the feeding rate around 10h and 18h, and very low feeding activity from 19h to 23h and from 2h to 6h. A two-peak profile has been identified previously as the typical feeding behavior of pigs [12, 5]. Our analysis exploiting the rich information provided by the new RFID antennas used [11], however, gives pig-specific results, whereas previous studies typically reported on group characteristics; see, e.g., Hyun et al. [8] and Guillemet et al. [5]. Furthermore, existing results usually aggregated over time periods.

The estimated smoothly varying day effect (top right panel) shows increased feeding rates in the early morning in the first few days and a corresponding

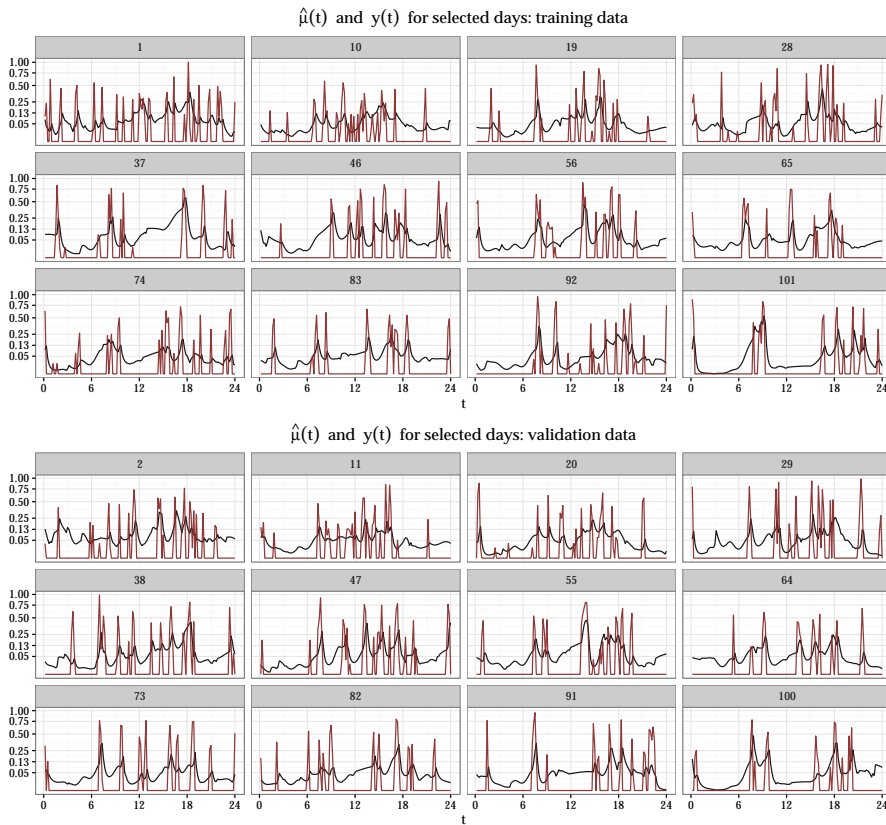


FIG 1. Training (top) and validation data (bottom) results for pig 57 for selected days with a smoothly varying day effect. Black lines for fitted (top) and predicted (bottom) values $\hat{\mu}_i(t)$, red for observed feeding rate $y_i(t)/60$. Numbers above each panel denote the day. Vertical axis on $\sqrt{\cdot}$ -scale.

strong reduction in feeding rates in the early morning towards the end of the fattening period, as well as a tendency for increased feeding activity to concentrate more strongly in two periods around 9h and 21h towards the end of the observation period, a pattern also visible in the raw data shown in Figure B.1 (Appendix B). The cumulative auto-regressive effect of feeding during the previous 3 hours is displayed in the bottom row of Figure 2. Unsurprisingly, feeding behavior in the immediate past is associated positively with current feeding (c.f. the blue region for $s - t \approx 0$), especially during off-peak times in the early morning where a higher propensity for feeding is not already modeled by the global functional intercept. The model also finds a negative association between prior feeding during the night and feeding in the very early hours of the morning that takes place 1 to 3 hours later (c.f. the red region around $t = 0$). Confidence intervals for all three effects show that they can be estimated fairly precisely.

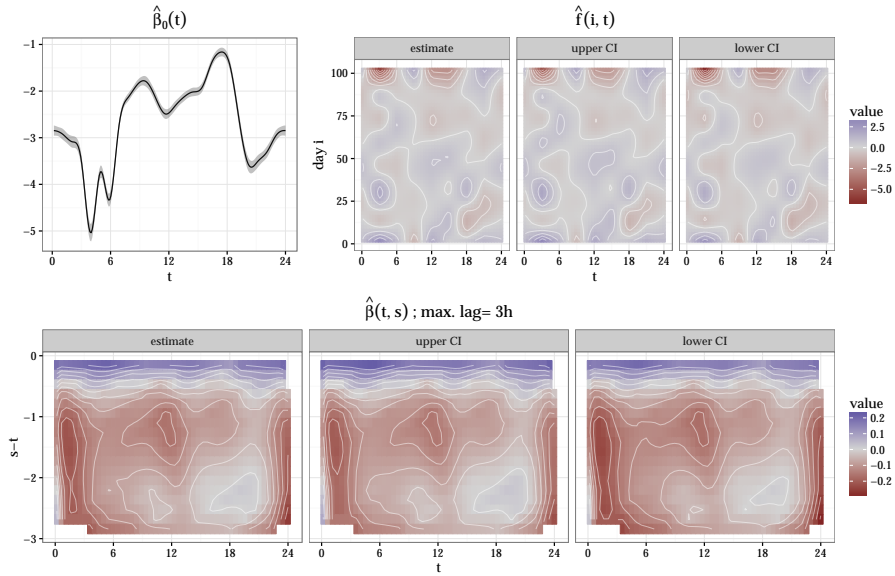


FIG 2. Top row: Estimated functional intercept and smoothly varying day effects for the model with smoothly varying day effect. Bottom: Estimated coefficient surface for cumulative auto-regressive effect. Point-wise intervals are ± 2 standard errors.

4. Simulation study

This section describes results on binomial (Section 4.1) and Beta-, t - and negative binomial-distributed data (4.2). Sections 4.3 and 4.4 present replications of (parts of) the simulation studies on functional random intercept models in Wang and Shi [26] and Goldsmith et al. [4], respectively. Note that neither of these competing approaches are able to fit nonlinear effects of scalar covariates, effects of functional covariates, functional random slopes or correlated functional random effects in their current implementation, while `pffr` does. Both approaches are implemented only for binary data, not the broad class of distributions available for `pffr`-models.

To evaluate results on artificial data, we use the relative root integrated mean square error `rRIMSE` of the estimates evaluated on equidistant grid points over \mathcal{T} :

$$\text{rRIMSE}(\hat{f}_r(\mathcal{X}_r, t)) = \sqrt{\frac{|\mathcal{T}|}{nT} \sum_{i,l} \frac{(\hat{f}_r(\mathbf{x}_{ri}, t_l) - f(\mathbf{x}_{ri}, t_l))^2}{\text{sd}(f(\mathbf{x}_{ri}, t))^2}},$$

where $\text{sd}(f(\mathbf{x}_{ri}, t))$ is the empirical standard deviation of $[f(\mathbf{x}_{ri}, t_l)]_{l=1, \dots, T}$, i.e., the variability of the estimand values over t . We use the relative RIMSE rather than $\text{RIMSE}(\hat{f}_r(\mathcal{X}_r, t)) = \sqrt{\frac{|\mathcal{T}|}{nT} \sum_{i,l} (\hat{f}_r(\mathbf{x}_{ri}, t_l) - f_r(\mathbf{x}_{ri}, t_l))^2}$ in order

to make results more comparable between estimands with different numerical ranges and to give a more intuitively accessible measure of the estimation error, i.e., the relative size of the estimation error compared to the variability of the estimand across \mathcal{T} .

All results presented in this section can be reproduced with the code supplement available for this work.

4.1. Synthetic data: Binomial additive mixed model

We describe an extensive simulation study to evaluate the quality of estimation for various potential model specifications for the PIGWISE data presented in Section 3.1. We simulate Binomial data $y_i(t) \sim B(\text{trials}, \pi = \text{logit}^{-1}(\eta_i(t)))$ with trials = 10, 60, 200 for $n = 100$ observations on $T = 150$ equidistant grid-points over $[0, 1]$ for a variety of true additive predictors $\eta_i(t) = \sum_{r=1}^R f_r(\mathcal{X}_{ri}, t)$. The subsequent paragraph describes the various $f_r(\mathcal{X}_{ri}, t)$ we investigate.

The additive predictor always includes a functional intercept $\beta_0(t)$. For modelling day-to-day differences in feeding rates, the additive predictor includes an i. i. d. functional random intercept $b_i(t)$ (setting: **ri**) or a smoothly varying day effect $f(i, t)$ (setting: **day**). Note that the “true” functional random intercepts were generated with Laplace-distributed spline coefficients in order to mimic the spiky nature of the application data. For modeling the auto-regressive effect of past feeding behavior, the additive predictor can include one of

- a time-varying nonlinear term $f(\tilde{y}_i(t - .005), t)$ (**lag**),
- a time-constant nonlinear term $f(\tilde{y}_i(t - .005))$ (**lag.c**),
- a cumulative linear term over the previous 0.3 time units $\int_W \tilde{y}_i(s)\beta(t, s)ds$; $W = \{s : t - 0.3 \leq s < t\}$ (**ff.3**),
- a cumulative linear term over the previous 0.6 time units $\int_W \tilde{y}_i(s)\beta(t, s)ds$; $W = \{s : t - 0.6 \leq s < t\}$ (**ff.6**),

where $\tilde{y}_i(s) = (y_i(s) - \bar{y}(s))$ denotes the centered responses and $\bar{y}(s)$ the mean response function. The simulated data also contains humidity and temperature profiles generated by drawing random FPC score vectors from the respective empirical FPC score distributions estimated from the (centered) humidity and temperature data provided in the PIGWISE data. For modeling the effect of these functional covariates, the additive predictor includes either a nonlinear concurrent interaction effect $f(\text{hum}_i(t), \text{temp}_i(t), t)$ (**ht**) or a nonlinear time-constant concurrent interaction effect $f(\text{hum}_i(t), \text{temp}_i(t))$ (**ht.c**).

We simulate data based on additive predictors containing each of the terms described above on its own, as well as additive predictors containing either functional random intercepts (**ri**) or a smoothly varying day effect (**day**) and one of the remaining terms, for a total of 21 different combinations. For all of these settings, we vary the amplitude of the global functional intercept (small: $\text{logit}^{-1}(\beta_0(t)) \in [0.06, 0.13] \forall t$; intermediate: $\in [0.04, 0.19]$; large: $\in [0.02, 0.34]$). For the settings with just a single additional term we also varied the absolute effect size of these terms. For small effect sizes the respective $\exp(f_r(\mathcal{X}_{ri}, t)) \in [0.5, 2] \forall t$, is $\in [0.2, 5]$ for intermediate effect sizes and $\in [0.1, 10]$ for large




FIG 3. $rRIMSE(\hat{\eta}(t))$ and coverages for the different models for synthetic binomial data (60 trials, $\text{logit}^{-1}(\beta_0(t)) \in [0.04, 0.19]$).

effect sizes. We ran 30 replicates for each of the resulting 333 combinations of simulation parameters, for a total of 9990 fitted models. Since the settings’ effects are large and systematic and we use a fully crossed experimental design, 30 replicates are more than enough to draw reliable qualitative conclusions from the results here.

We focus on the results for trials = 60 with intermediate amplitude of the global intercept as these settings are most similar to the application. Results for the other settings are qualitatively similar, with the expected trends: performance improves for more informative data (i.e, more “trials”) and larger effect sizes and worsens for less informative data with fewer “trials” and smaller effect sizes.

Figure 3 shows $rRIMSE(\hat{\eta}(t))$ (left panel) and average point-wise coverages (right panel) for the various settings – it is easy to see that the concurrent functional effects (**humtemp**, **humtemp.c**) and especially the daily random intercepts (**ri**) are the hardest to estimate. Further study shows that concurrent interaction effects are very sensitive to the rank of the functional covariates involved: both humidity and temperature are practically constant in this data and show little variability, so effects are hard to estimate reliably in this special case due to a lack of variability in the covariate. The daily random intercepts are generated from a Laplace distribution of the coefficients to yield “peaky” functions as seen in the application, not the Gaussian distribution our model assumes. Figure A.3 in Appendix A.1 displays a typical fit for the setting with functional random intercepts and a cumulative autoregressive effect (**ri+ff.6**) and shows that $rRIMSE$ s around .2 correspond to useful and sensible fits in this setting. Also note that performance improves dramatically if the Gaussianity assump-

tion for the random intercept functions is met, c.f. the results in Sections 4.3 and 4.4 for Gaussian random effects in some much simpler settings.

Note that we are trying to estimate a latent smooth function for each day based only on the feeding episodes for that day for `ri` and, through smoothing, based on each day and its neighboring days in the case of `day`. The results show that this strategy can work rather well, albeit with a high computational cost in the case of `ri`: On our hardware, models with such functional random intercepts usually take between 5 minutes and 2 hours to fit, while models with a smooth day effect usually take between 15 seconds and 15 minutes, depending on which covariate effects are present in the model. See Figure A.1 in Appendix A.1 for details on computation times.

As seen in the replication studies in Sections 4.3 and 4.4, estimating random intercepts with our approach also works as well or better than competing approaches in more conventional settings with multiple replicates per group level rather than the one considered here with a single curve per group level. The results here still represent quite successful fits – Figure 4 shows results for a typical `day+ff.6`-model that yields about the median RIMSE for that setting. Most estimates are very close to the “truth”.

In terms of CI coverage, we find that CIs for $\hat{\eta}_i(t)$ (Figure 3, right panel) are usually close to their nominal 95% level.

Appendix A.1 shows typical fits for some of the remaining settings. In summary, the proposed approach yields useful point estimates for most terms under consideration and confidence intervals with coverages consistently close to the nominal level on data with a similar structure as the application.

4.2. Synthetic data: Beta-, negative binomial-, $t(3)$ -distribution

4.2.1. Data Generating Process

We generate data with responses with scaled $t(3)$ -distributed errors, $\text{Beta}(\alpha, \beta)$ -distributed responses and Negative Binomial $\text{NB}(\mu, \theta)$ -distributed responses. We generate data with $n = 100, 300$ observations, each with signal-to-noise ratios $\text{SNR} = 1, 5$. SNR is not easily adjustable for Negative Binomial data since variance and mean cannot be specified separately. In our simulation, we use $g(\cdot) = \log(\cdot)$, with $\theta = 0.5$, so that $\text{Var}(y(t)) = \text{E}(y(t)) + 2\text{E}(y(t))^2$. See Appendix A.2 for details.

Functional responses are evaluated on $T = 60$ equidistant grid-points over $[0, 1]$. We investigate performance for 4 different additive predictors:

- `int`: $\text{E}(y(t)) = g^{-1}(\beta_0(t))$
- `smoo`: index-varying non-linear effect: $\text{E}(y(t)) = g^{-1}(\beta_0(t) + f(x, t))$
- `te`: non-linear interaction: $\text{E}(y(t)) = g^{-1}(\beta_0(t) + f(x_1, x_2))$ via tensor product spline
- `ff`: functional covariate: $\text{E}(y(t)) = g^{-1}(\beta_0(t) + \int x(s)\beta(t, s)ds)$

50 replicates were run for each of the 32 settings for Beta and $t(3)$ and each of the 16 settings for NB.

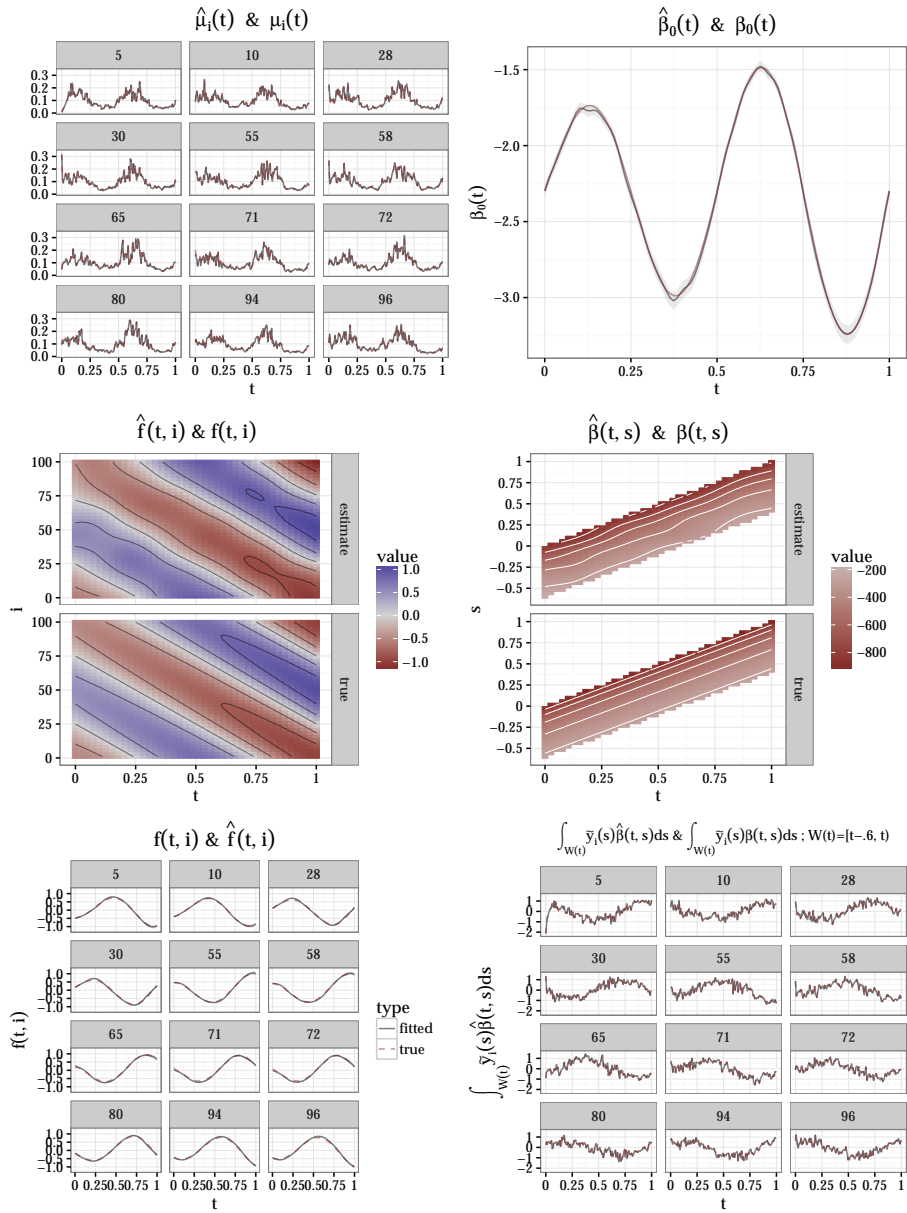


FIG 4. Typical model fit for setting *day+ff.6*: smoothly varying day effect with a cumulative auto-regressive term. Left column shows true (red, dashed) and estimated (black, solid) conditional expectation (top) for 12 randomly selected days in the top row, estimated (top panel) and true (bottom panel) smoothly varying day-daytime effect (middle and bottom rows). Right column shows true and estimated global intercept (top), estimated (middle, top panel) and true (middle, bottom panel) coefficient surface for the cumulative auto-regressive effect and true and estimated contributions of the cumulative auto-regressive effect to the additive predictor (bottom). For this model, $rRIMSE(\hat{\eta}(t)) = 0.054$, while the median $rRIMSE$ for this setting is 0.069

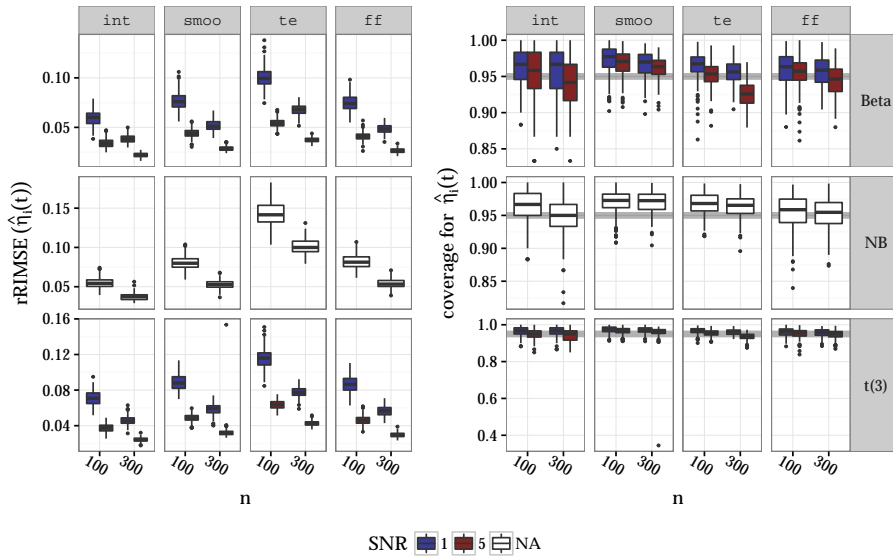


FIG 5. Relative RIMSE (left) and achieved coverage (right, nominal 95%) for estimated additive predictor. Rows for the three distributions and columns for the 4 model settings. Blue for high noise settings with SNR = 1, red for SNR = 5, white for NB without available SNR. Horizontal axis represents the varying number of observations. Fat grey horizontal line denotes nominal 95% coverage.

4.2.2. Results

Figure 5 shows relative RIMSEs (left figure) and achieved point-wise coverages averaged across t and i (right figure, nominal 95%) for the estimated additive predictor $\hat{\eta}_i(t)$ for each setting and replicate. Both point estimation and uncertainty quantification work well in these settings. For the point estimates, we observe the expected patterns of increasing precision with bigger data sets and decreasing noise levels. The systematic under-coverage for $n = 300$, SNR= 5 we observe for the Beta-settings for the model with a time-constant nonlinear interaction effect (third column, first row) seems of little practical relevance – in this setting, the point estimates are so close to the true effects and estimation uncertainty becomes so small that the CIs almost vanish, and in any case observed coverages mostly remain above 90% for nominal 95%.

Figure 6 shows relative RIMSEs (left) and coverages (right) for the estimated effects in the three settings that include more than an intercept. Again, we observe the expected patterns of increasing precision with bigger data sets and decreasing noise levels. Estimates are generally fairly precise and most coverages are close to or greater than the nominal 95%-level. Systematic under-coverage occurs for the scalar interaction effect (middle column) for the Beta and $t(3)$ -distributions, but note again that given the high precision of the point estimates the fact that CIs tend to be too narrow is of little practical relevance. This could

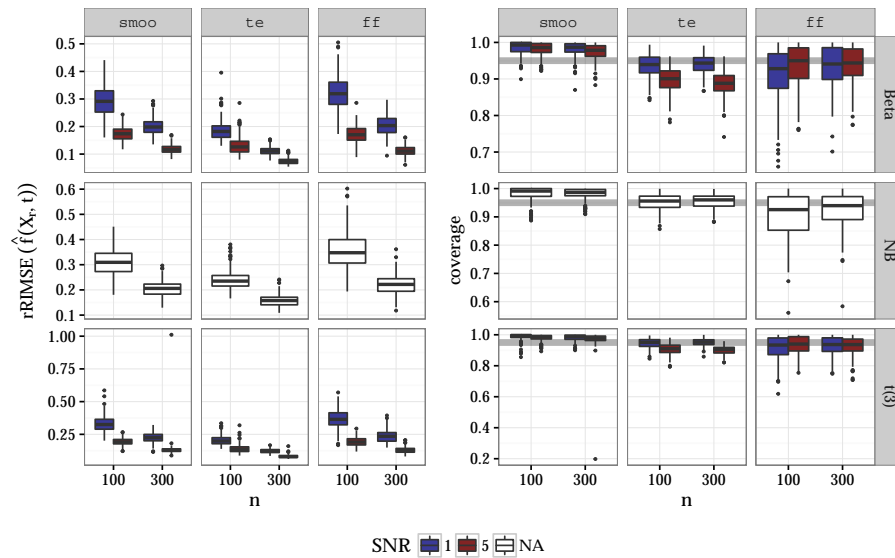


FIG 6. Relative RIMSE (left) and achieved coverage (right, nominal 95%) for estimated effects. Columns for the three distributions and rows for settings with index-varying nonlinear effect, nonlinear interaction effect and linear function-on-function effect, respectively. Blue for high noise settings with SNR = 1, red for SNR = 5, white for NB without available SNR. Horizontal axis represents the varying number of observations.

also indicate a basis specification that is too small to fit the true effect without approximation bias. We also observe a single replicate of setting `smoo` with $t(3)$ -distributed errors with much larger rRIMSE and coverage around 20% (bottom row, leftmost panels, rightmost boxplots), indicating that the robustification by specifying a t -distribution may not always be successful.

Median computation times for these fairly easy settings were between 3 and 49 seconds. Estimating the bivariate non-linear effect for setting `smoo` took the longest, with some fits taking up to 8 minutes. We observed no relevant or systematic differences in computation times between the three response distributions we used.

Appendix A.2 shows some typical results for these settings and gives additional details on the data generating process as well as tabular representations of the performance measures in Figures 5 and 6.

4.3. Replication of simulation study in Wang and Shi [26]

This section describes our results of a replication of the simulation study described in Section 4.1 of Wang and Shi [26]. We ran 20 replicates per setting. The data comes from a logit model for binary data with just a functional intercept and observation-specific smooth functional random effects. The observation-specific smooth functional random effects $b_i(t)$ are realizations of a Gaussian

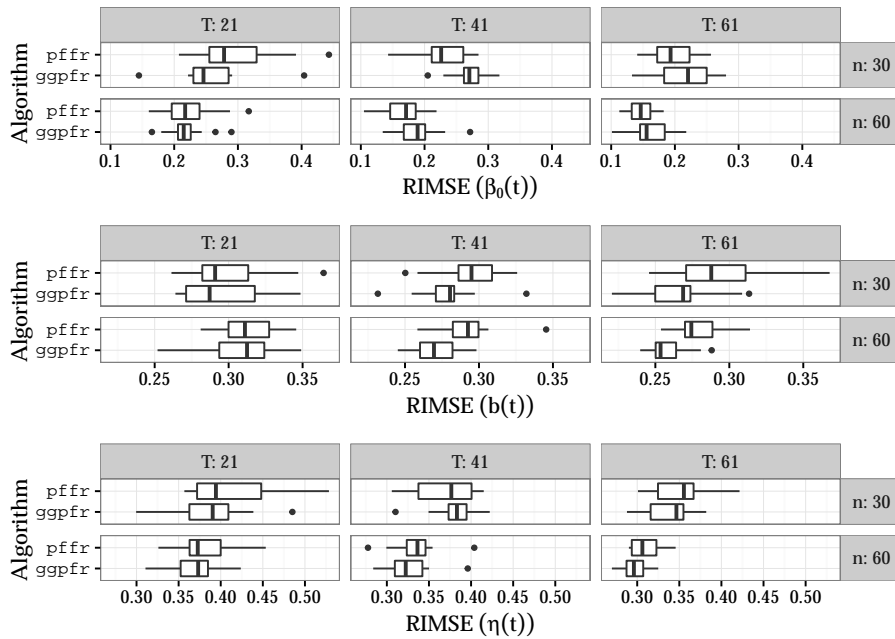


FIG 7. Replication of Wang and Shi [26]: RIMSE for the estimated effects and the additive predictor. Top: functional intercept; middle: functional random effects; bottom: sum of the two. Each graph shows results for $n = 30, 60$ (rows) and $T = 21, 41, 61$ (columns). Upper boxplots for *pffr*, lower for *ggprf*.

process with “squared exponential” covariance structure, the functional intercept $\beta_0(t)$ is a cubed sine-function, see Wang and Shi [26, their Section 4.1.] for details. The model and data generating process are thus $P(y_i(t_l) = 1) = \text{logit}^{-1}(\eta_i(t))$ with $\eta_i(t) = \beta_0(t_l) + b_i(t_l)$; $i = 1, \dots, n$; $l = 1, \dots, T$.

Note that Wang & Shi’s MATLAB [23] implementation of their proposal (denoted by *ggprf*) uses the same “squared exponential” covariance structure for fitting as that used for generating the data, while our *pffr*-implementation uses a B-spline basis (8 basis functions per subject) to represent the $b_i(t)$. We used the same basis dimension for $\beta_0(t)$ for *ggprf* as Wang & Shi.

Boxplots in Figure 7 show RIMSEs for the estimated effects and the additive predictor for each of the 20 replicates per combination of settings. While *ggprf* tends to yield more precise estimates for $b_i(t)$, it often does not do as well as *pffr* in estimating the functional intercept. As a consequence, the two approaches yield very similar errors for $\eta_i(t)$, even though the data generating process is tailored towards *ggprf* and more adversarial for *pffr*. Note that we were not able to reproduce the mean RIMSE values for $\text{logit}(\eta(t))$ reported by Wang and Shi [26] despite the authors’ generous support. Our mean RIMSE results for *ggprf* were 0.37, 0.33, 0.30 while Wang and Shi [26] reported 0.32, 0.26, 0.24 for $n = 60$; $T = 21, 41, 61$, respectively.




FIG 8. Replication of Wang and Shi [26]: Observed CI coverage for $\eta(t) = \text{logit}(\beta_0(t) + b_i(t))$ for $n = 30, 60$ (rows) and $T = 21, 41, 61$. Upper boxplots for **pffr**, lower for **ggppfr**. Vertical fat gray line denotes nominal 90% coverage.




FIG 9. Replication of Wang and Shi [26]: Computation times for $n = 30, 60$ (rows) and $T = 21, 41, 61$ (columns). Upper boxplots for **pffr**, lower for **ggppfr**. Horizontal axis on \log_{10} -scale.

Figure 8 shows observed CI coverage for nominal 90% CIs for the two approaches evaluated over $\hat{\eta}_i(t)$. Both approaches mostly do not achieve nominal coverage for these very small data sets, but coverages for **pffr** converge towards the nominal level much faster as n and T increase than those for **ggppfr** and show fewer and smaller incidences of severe under-coverage. Figure 9 shows computation times for the two approaches. The median computation time for **ggppfr** is about three times that of **pffr** for $T = 21$, which increases to a factor of 13-19 for $n = 60; T = 41, 61$ and a factor of 40-48 for $n = 30; T = 41, 61$ ¹. Note that computation times for **ggppfr** are given for a single run using a pre-specified number of basis functions for estimating $\beta_0(t)$ here. But since the BIC-based selection of the basis dimension proposed by Wang and Shi [26] requires k runs for selecting between k candidate numbers of basis functions, actual computation times for **ggppfr** in practical applications will increase roughly k -fold.

Our replication shows that **pffr** achieves mostly similar estimation accuracies with better CI coverage for a data generating process tailored specifically to **ggppfr**'s strengths. Depending on the MATLAB version that is used, **pffr** does so in similar or much, much shorter time.

¹Note that these computation times seem to be extremely sensitive towards different MATLAB versions – in personal communication, Bo Wang reported average computation times of 3 minutes for $T = 21, N = 60$ (our result for **ggppfr**: ≈ 19), 4 minutes for $T = 41, N = 30$ (our result: ≈ 200) and 10 minutes for $T = 61, N = 60$ (our result: ≈ 540) for a single run using MATLAB 7.4 on an Intel Core Duo CPU (2.53GHz, 3GB RAM) instead of MATLAB 7.12 which we used.

4.4. Replication of Simulation Study in Goldsmith et al. [4]

This section describes our results of a replication of the simulation study described in Web Appendix 2 of Goldsmith et al. [4]. We ran 20 replicates per setting. The data comes from a logit model for binary data with a functional intercept, a functional linear effect of a scalar covariate and observation-specific smooth functional random effects. The observation-specific smooth functional random effects $b_i(t)$ are drawn using 2 trigonometric functional principal component functions, the functional intercept $\beta_0(t)$ is a trigonometric function as well and the functional coefficient function $\beta_1(t)$ associated with a $N(0, 25)$ -distributed scalar covariate x is a scaled normal density function, see Goldsmith et al. [4, Web Appendix, Section 2] for details. Note that `genfpca`, the `rstan` [21] implementation provided by Goldsmith et al., uses the same number of FPCs (i.e. two) for the fit as used for generating the data, while our `pffr`-implementation uses 10 cubic B-spline basis functions to represent each $b_i(t)$. The model and data generating process are thus $P(y_i(t_l) = 1) = \text{logit}^{-1}(\eta_i(t))$ with $\eta_i(t) = \beta_0(t) + \beta_1(t)x_i + b_i(t)$; $i = 1, \dots, n$; $l = 1, \dots, T$.

Figure 10 shows RIMSE for (top to bottom) the estimated functional intercept, the estimated functional coefficient, the functional random effect and the combined additive predictor. It seems that `pffr` yields somewhat more precise estimates for $\beta_0(t)$ in this setting, while the estimation performance of the two approaches for the functional coefficient and the random effects is broadly similar although the results for `pffr` for $\beta_1(t)$ show more variability. `pffr` yields slightly better estimates of the total additive predictor $\eta_i(t)$ in this setting. Note that Goldsmith et al. [4] report mean integrated square error (MISE, in their notation) while we report the RIMSE. Taking this into account, the `genfpca` results reported here correspond very closely to theirs. Figure 11 shows observed confidence [credibility] interval (CI) coverage for nominal 95% CIs for the two approaches evaluated over $\hat{\eta}_i(t)$. Both approaches mostly achieve close to nominal coverage in this setting, but while coverages for `pffr` are mostly greater than the nominal 95% for $T = 50$ with some rare problematic fits with coverages below 75%, `genfpca` shows small but systematic under-coverage for $T = 100$. Figure 12 shows computation times for the two approaches. The median computation time for `genfpca` is about 4 times that of `pffr` for $n = 50$, and 0.80 to 0.94 that of `pffr` for $n = 100$.

In summary, our replication shows that `pffr` achieves mostly similar estimation accuracies and coverages in much shorter time or roughly the same time for a data generating process tailored specifically to `genfpca`'s strengths.

4.5. Summary of simulation results

We achieve useful and mostly reliable estimates for many complex settings in data structured like the PIGWISE data in acceptable time. Simulation results for less challenging settings and other response distributions than Binomial are good to excellent as well. Our partial replications of the simulation studies in




FIG 10. Replication of Goldsmith et al. [4]: RIMSE for the estimated effects and the additive predictor. Top: functional intercept; 2nd row: functional coefficient; 3rd row: functional random effects; bottom: sum of the three. Each graph shows results for $n = 50, 100$ (rows) and $T = 50, 100$ (columns). Upper boxplots for *pffr*, lower for *genfpca*.




FIG 11. Replication of Goldsmith et al. [4]: Observed CI coverage for $\eta_i(t) = \text{logit}(\beta_0(t) + \beta_1(t)x_i + b_i(t))$ for $n = 50, 100$ (rows) and $T = 50, 100$. Upper boxplots for *pffr*, lower for *genfpca*. Vertical fat gray line denotes nominal 95% coverage.




FIG 12. Replication of Goldsmith et al. [4]: Computation times for $n = 50, 100$ (rows) and $T = 50, 100$ (columns). Upper boxplots for *pffr*, lower for *genfpca*. Horizontal axis on \log_{10} -scale.

Wang and Shi [26] and Goldsmith et al. [4] in Sections 4.3 and 4.4, respectively, show that our implementation achieves highly competitive results in similar or much shorter computation times even though the settings are tailored towards the competing earlier approaches which are also less flexible than ours by quite a large margin.

5. Discussion

This work introduces a comprehensive framework for generalized additive mixed models (GAMM) for non-Gaussian functional responses. Our implementation extends all the flexibility of GAMMs for dependent scalar responses to dependent functional responses and functional covariates, even for response distributions outside the exponential family such as robust models based on the t -distribution. Dependency structures can be spatial, temporal or hierarchical. Simulation and application results show that our approach provides reliable and precise inference about the underlying latent processes.

The work presented here opens up promising new avenues of inquiry – one challenge that we have already begun to work on is to improve the speed and memory efficiency of the underlying computational engine to be able to fit the huge data sets increasingly common in functional data analysis. Along these lines, more efficient computation of simultaneous instead of pointwise confidence intervals for functional effects as well as a more detailed investigation of the performance of the available approximate pointwise confidence intervals for noisy and small non-Gaussian data is another important field of inquiry. An extension of the approach presented here to models with multiple additive predictors controlling different aspects of the responses' distribution – like the generalized additive models for location, scale and shape introduced for scalar response by Rigby and Stasinopoulos [17] or zero-inflation and hurdle models – is yet another promising generalization of the ideas we have presented here.

Appendix A: Additional simulation details

Section A.2 provides some more details on the data generating process used in Section 4.2 and shows some typical effect estimates in these settings.




FIG A.1. Computation times for the different models. Time axis on \log_2 -scale.

A.1. Simulation 1: Computation times, details and examples

A.1.1. Computation times

Figure A.1 shows computation times for the various settings. Note that the times are plotted on binary log-scale. One fit for a random intercept model (**ri**) did not converge within 200 iterations and ran for more than 16 hours.

A.1.2. Computational details

We use the following (spline) bases to construct $\Phi_{\mathbf{x}}$ for the different effects:

- **lag.c**, **lag**: 20 cubic thin plate splines over $\tilde{y}(t)$
- **humtemp**, **humtemp.c**: 50 bivariate cubic thin plate splines over the joint space of **hum** and **temp**.
- **day**: 8 cubic thin plate splines over $i = 1, \dots, 100$.
- **ri**: dummy variables for each curve $i = 1, \dots, 100$.

The coefficient surfaces for **ff.3**, **ff.6** are estimated with 30 bivariate cubic thin plate splines over $\mathcal{T} \times \mathcal{T}$. The functional intercept (**int**) uses 40 cubic cyclic P-splines with first order difference penalty over \mathcal{T} , while the functional random intercepts (**ri**) use 9 of those per curve. For all other terms, we use 8 cubic P-splines with first order difference penalty for $\Phi_{\mathbf{t}}$.

A.1.3. Typical model fits

Figures A.2 to A.4 show graphical summaries of typical model fits for three difficult settings of the simulated binomial data discussed in Section 4.1. It is clear to see that estimating unstructured daily functional random intercepts is a harder task than estimating a smoothly varying aging effect (compare the accuracy of $\hat{f}(t, i)$ in Figures A.2 and A.4 to that of $\hat{b}_i(t)$ in Figure A.3), and that concurrent functional nonlinear interaction effects can not be recovered very reliably in this setting. Also note that we can model very rough latent response processes $\mu_i(t)$ for these data by including, e.g., a (nonlinear) auto-regressive term as in the top left panel of Figure A.2.

A.2. Simulation study 2: Details & typical fits

A.2.1. Data generating process for simulation study 2

We use different signal-to-noise ratios (SNR) to control the difficulty of the simulation settings. For Beta-distributed responses with $g(\cdot) = \text{logit}^{-1}(\cdot)$, distribution parameters α, β are determined by approximate SNR via $\phi = \text{SNR} \frac{m_\mu}{v_\mu} - 1$, with $m_\mu = N^{-1} \sum_{i,t} \mu_{it}(1 - \mu_{it})$, where v_μ is the sample variance of the simulated μ_{it} and $\boldsymbol{\mu} = \text{logit}^{-1}(\boldsymbol{\eta})$. Beta parameters for generating \mathbf{y} are then given by $\boldsymbol{\alpha} = \phi \boldsymbol{\mu}$ and $\boldsymbol{\beta} = \phi - \boldsymbol{\alpha}$. $\boldsymbol{\eta}$ is first scaled linearly to the interval $[-1.5, 1.5]$.

A.2.2. Computational details

We use the following spline bases to construct $\boldsymbol{\Phi}_x$ for the different effects:

- **smoo**: 8 cubic thin plate splines over the range of x
- **te**: 45 bivariate cubic thin plate splines over the joint space of x_1 and x_2
- **ff**: 5 cubic P-splines with first order difference penalty over \mathcal{S}

The functional intercept (**int**) uses 40 cubic cyclic P-splines with first order difference penalty over \mathcal{T} . For all other terms, we use 5 cubic P-splines with first order difference penalty for $\boldsymbol{\Phi}_t$.

A.2.3. Typical fits for simulation study 2

Figures A.5 to A.7 show some typical fits for these data generating processes. Figure A.5 shows estimated effects for a Beta-response model with a nonlinear effect of a scalar covariate (**smoo**, $n = 300$, SNR = 1), Figure A.6 shows estimated effects for a Negative Binomial-response model with a linear function-on-function effect (**ff**, $n = 100$) and Figure A.7 shows estimated effects for a $t(3)$ -response model with a nonlinear interaction effect of two scalar covariates (**te**, $n = 100$, SNR = 5)

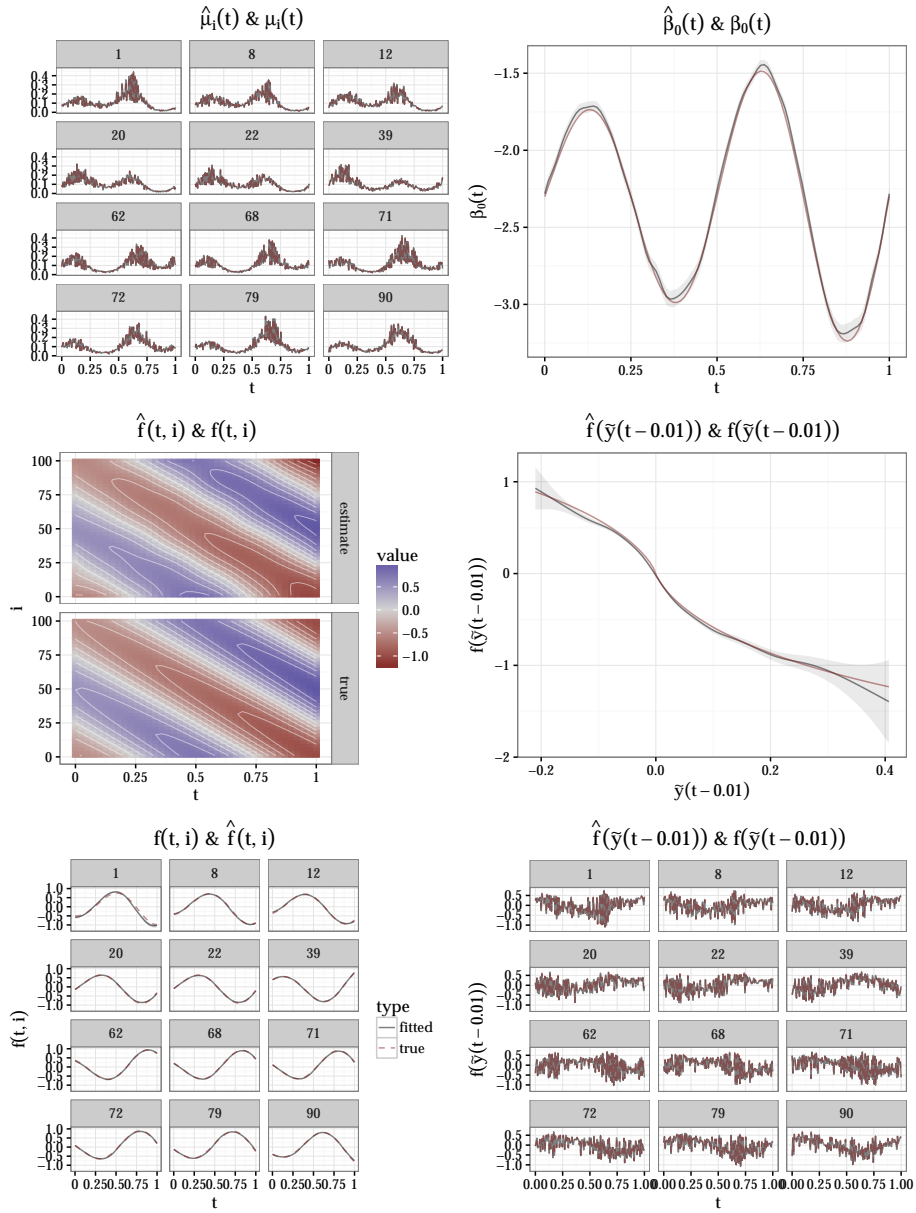


FIG A.2. Typical model fit for setting *day+lag.c*: smoothly varying day effect with a nonlinear auto-regressive term. Left column shows true (dashed, red) and estimated (solid, black) conditional expectation for 12 randomly selected days in the top row and estimated and true smoothly varying day-daytime effect in the middle and bottom panels. Right column shows true and estimated global intercept (top) and true and estimated nonlinear auto-regressive effect (middle) and its true and estimated contributions to the additive predictor. For this model, $rRIMSE(\hat{\eta}_i(t)) = 0.05$, about the same as the median $rRIMSE$ for this setting.

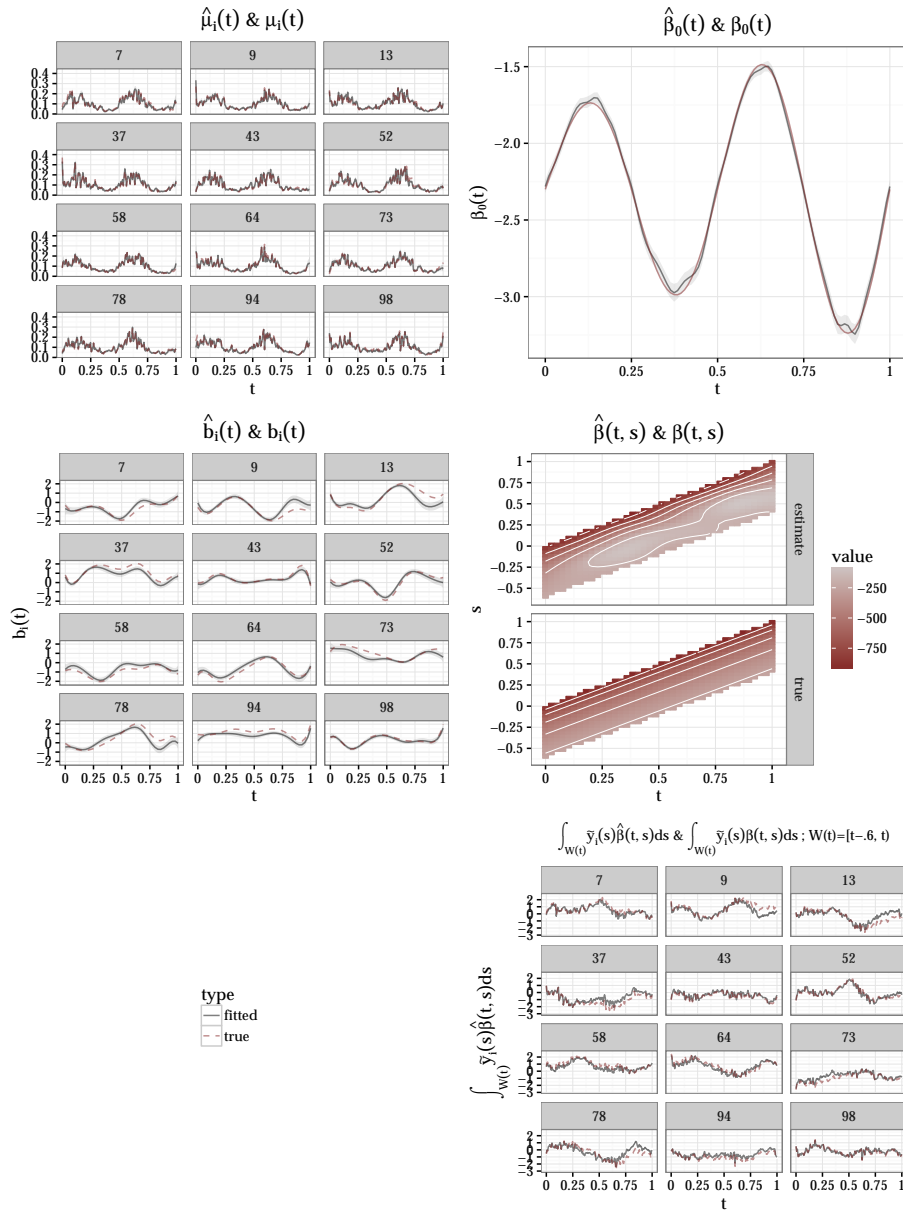


FIG A.3. Typical model fit for setting $ri+ff.6$: random functional intercepts for each day with a cumulative auto-regressive term. Left column shows true (dashed, red) and estimated (solid, black) conditional expectation (top row) and estimated and true random intercepts (middle). Right column shows true and estimated global intercept (top) and estimated (middle, top panel) and true (middle, bottom panel) coefficient surface for the cumulative auto-regressive effect and true and estimated contributions of the cumulative auto-regressive effect to the additive predictor (bottom) for 12 randomly selected days. For this model, $rRIMSE(\hat{\eta}(t)) = 0.2$, while the median $rRIMSE$ for this setting is 0.19.

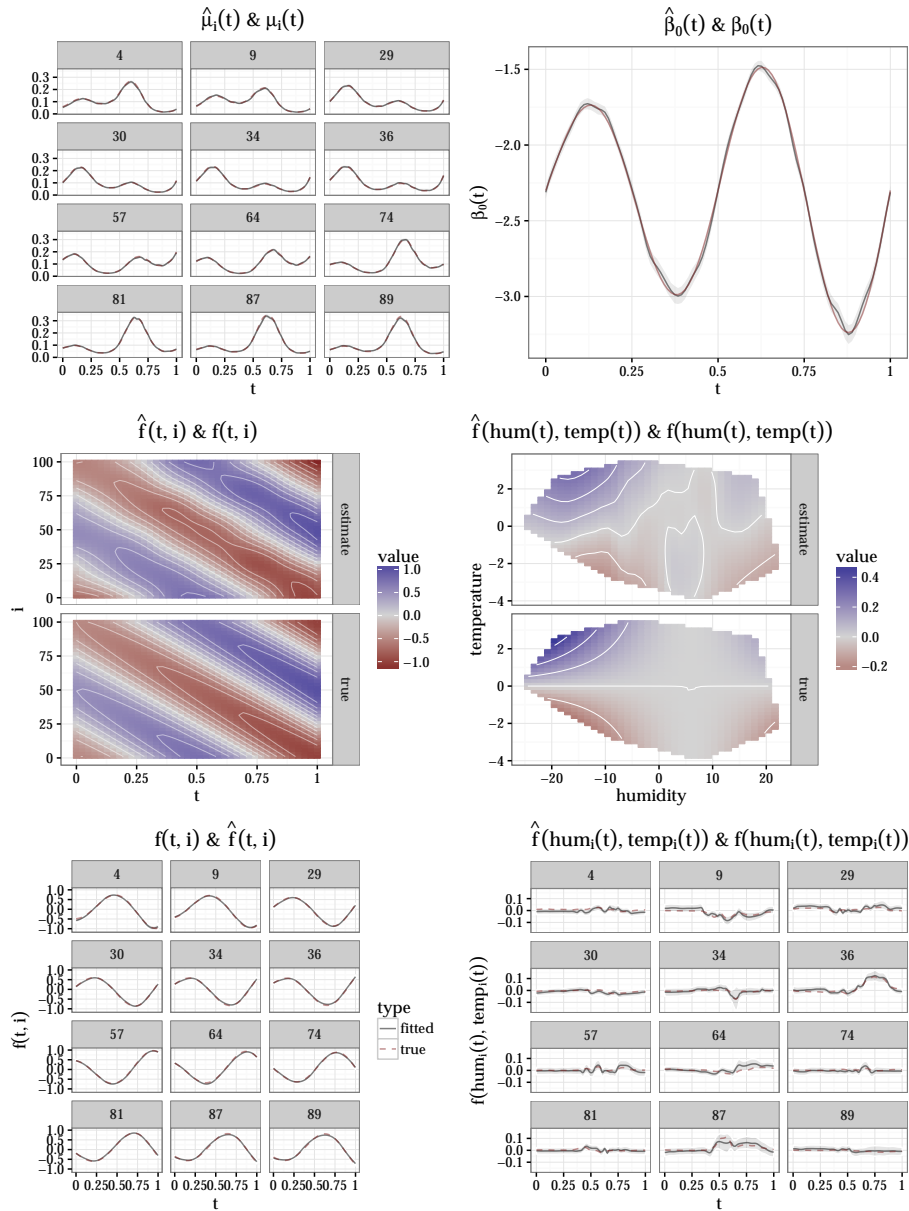


FIG A.4. Typical model fit for setting `day+humtemp.c`: smoothly varying day effect with a time-constant concurrent nonlinear interaction effect of humidity and temperature. Left column shows true (dotted) and estimated (solid) conditional expectation (top) for 12 randomly selected days in the top row, estimated (top panel) and true (bottom panel) smoothly varying day-daytime effect (middle), and true and estimated contributions of the time-constant humidity-temperature interaction effect to the additive predictor (bottom). Right column shows true and estimated global intercept (top) and estimated (middle, top panel) and true (middle, bottom panel) effect surface for the humidity-temperature interaction. For this model, $r\text{RIMSE}(\hat{\eta}(t)) = 0.05$, while the median $r\text{RIMSE}$ for this setting is 0.05.




FIG A.5. Typical model fit for a Beta-distributed model with a nonlinear effect of scalar covariate (*smoo*, $n = 300$, $SNR = 1$) with $rRIMSE(\hat{\eta}(t)) \approx 0.047$ and coverage ≈ 0.98 for $\hat{\eta}(t)$. Red line in the left panel shows true intercept function, light grey ribbon gives approximate pointwise 95% interval. On the right, top panel shows estimated nonlinear effect of scalar covariate $\hat{f}(x, t)$ and bottom panel the true $f(x, t)$.




FIG A.6. Typical model fit for a NB-distributed model with a linear effect of a functional covariate (*ff*, $n = 100$) with $rRIMSE(\hat{\eta}(t)) \approx 0.083$ and coverage ≈ 0.97 for $\hat{\eta}(t)$. Red dotted line in the left panel shows true intercept function, light grey ribbon gives approximate pointwise 95% interval. On the right, top panel shows estimated $\hat{\beta}(s, t)$ and bottom panel the true $\beta(s, t)$ for the function-on-function effect.




FIG A.7. Typical model fit for a t_3 -distributed model with a nonlinear interaction effect of two scalar covariates (te , $n = 300$, $SNR = 5$) with $rRIMSE(\hat{\eta}(t)) \approx 0.043$ and coverage ≈ 0.94 for $\hat{\eta}(t)$. Red dotted line in the left panel shows true intercept function, light grey ribbon gives approximate pointwise 95% interval. On the right, top panel shows estimated smooth interaction $\hat{f}(x_1, x_2)$ and bottom panel the true $f(x_1, x_2)$.

A.2.4. Tabular results for simulation study 2

Tables 2 and 3 show median $rRIMSE$ s and coverages for simulation study 4.2 for the entire additive predictor and the estimated covariate effects, respectively.

A.3. Computational details

Code to reproduce results for Section 4 is included in a code supplement.

Results described here were computed on various Linux PCs and servers under R-3.2.3 to 3.2.5 with `mgcv` 1.8-5 to 1.8-7, `refundDevel` 0.1-11 to 0.1-15.

Appendix B: PIGWISE model: Alternatives and criticism

This section shows the raw data used for the PIGWISE application (Figure B.1) and compares the result of a different model specification to that of the main article.

Beyond the smoothly varying functional day effects used in the analysis in Section 3.2, we considered two alternative random effect specifications: auto-correlated functional random day effects with a marginal $AR(1)$ -structure with

TABLE 2
 Tabular display of results in Figure 5. Median and 25% and 75% quantiles for relative RIMSE and pointwise coverages for the additive predictor $\hat{\eta}(t)$.

Family	Set	SNR	n	Median rRIMSE($\hat{\eta}(t)$)	q25	q75	Median coverage	q25	q75		
Beta	int	1	100	0.06	0.05	0.06	0.97	0.95	0.98		
			300	0.04	0.04	0.04	0.97	0.93	0.98		
		5	100	0.03	0.03	0.04	0.96	0.93	0.98		
			300	0.02	0.02	0.02	0.94	0.92	0.97		
		smoo	1	100	0.08	0.07	0.08	0.98	0.96	0.99	
				300	0.05	0.05	0.06	0.97	0.95	0.98	
	5		100	0.04	0.04	0.05	0.97	0.96	0.98		
			300	0.03	0.03	0.03	0.96	0.95	0.97		
	te		1	100	0.10	0.09	0.11	0.97	0.96	0.98	
				300	0.07	0.06	0.07	0.96	0.95	0.97	
		5	100	0.05	0.05	0.06	0.95	0.94	0.96		
			300	0.04	0.04	0.04	0.93	0.91	0.94		
		ff	1	100	0.07	0.07	0.08	0.96	0.94	0.98	
				300	0.05	0.05	0.05	0.96	0.94	0.97	
	5		100	0.04	0.04	0.04	0.96	0.94	0.97		
			300	0.03	0.02	0.03	0.95	0.93	0.96		
	NB		int	NA	100	0.05	0.05	0.06	0.97	0.95	0.98
					300	0.04	0.03	0.04	0.95	0.93	0.97
		smoo		100	0.08	0.07	0.09	0.97	0.96	0.98	
				300	0.05	0.05	0.06	0.97	0.96	0.98	
		te		100	0.14	0.13	0.15	0.97	0.96	0.98	
				300	0.10	0.09	0.11	0.97	0.95	0.98	
		ff	100	0.08	0.08	0.09	0.96	0.94	0.97		
				0.05	0.05	0.06	0.95	0.94	0.97		
300			0.05	0.05	0.06	0.95	0.94	0.97			
			0.03	0.02	0.03	0.95	0.93	0.96			
t(3)			int	1	100	0.07	0.06	0.08	0.97	0.95	0.98
					300	0.05	0.04	0.05	0.97	0.95	0.98
	5	100		0.04	0.03	0.04	0.95	0.93	0.97		
		300		0.02	0.02	0.03	0.95	0.92	0.97		
	smoo	1		100	0.09	0.08	0.10	0.98	0.96	0.99	
				300	0.06	0.05	0.06	0.97	0.96	0.98	
		5	100	0.05	0.05	0.05	0.97	0.96	0.98		
			300	0.03	0.03	0.03	0.97	0.95	0.97		
		te	1	100	0.12	0.11	0.12	0.97	0.96	0.98	
				300	0.08	0.07	0.08	0.96	0.95	0.97	
	5		100	0.06	0.06	0.07	0.95	0.94	0.97		
			300	0.04	0.04	0.04	0.94	0.93	0.95		
ff	1		100	0.09	0.08	0.09	0.96	0.94	0.98		
			300	0.06	0.05	0.06	0.96	0.94	0.97		
	5	100	0.05	0.04	0.05	0.96	0.94	0.97			
		300	0.03	0.03	0.03	0.95	0.93	0.96			

auto-correlation 0.8 over the days as well as i. i. d. random day effects $b_i(t)$. We also considered models with no day effects at all.

Beyond the 3h-cumulative auto-regressive effect of previous feeding used in the analysis in Section 3.2, we also fit models with a 6h-cumulative auto-regressive effect ($\int_{t-6h}^{t-10min} y_i(s)\beta(t, s)ds$) as well as non-linear time-constant auto-regressive effects ($f(y_i(t - 10min))$) and time-varying ($f(y_i(t - 10min), t)$). We also considered models with no auto-regressive effects at all.

To model possible effects of humidity and temperature, we considered additive non-linear time-varying ($f(hum(t), t) + f(temp(t), t)$) and time-constant

TABLE 3
 Tabular display of results in Figure 6. Median and 25% and 75% quantiles for relative RIMSE and pointwise coverages for the estimated effects $\hat{f}(\mathcal{X}_r, t)$.

Family	Set	SNR	n	Median			Median			
				rRIMSE($\hat{f}(\mathcal{X}_r, t)$)			coverage			
Beta	smoo	1	100	0.29	0.25	0.33	0.99	0.97	1.00	
			300	0.20	0.18	0.22	0.99	0.97	1.00	
		5	100	0.17	0.16	0.19	0.99	0.97	1.00	
			300	0.12	0.11	0.13	0.98	0.96	0.99	
		te	1	100	0.18	0.16	0.20	0.94	0.92	0.96
			300	0.11	0.10	0.12	0.94	0.92	0.96	
	5	100	0.13	0.11	0.15	0.90	0.88	0.92		
		300	0.07	0.07	0.08	0.89	0.87	0.91		
	ff	1	100	0.32	0.28	0.36	0.93	0.87	0.97	
			300	0.20	0.18	0.23	0.94	0.90	0.99	
		5	100	0.17	0.15	0.19	0.95	0.90	0.99	
			300	0.11	0.10	0.12	0.94	0.91	0.98	
NB		smoo	NA	100	0.31	0.27	0.35	0.99	0.97	1.00
				300	0.21	0.18	0.22	0.99	0.97	1.00
	te	100	0.23	0.22	0.26	0.96	0.93	0.97		
		300	0.16	0.14	0.17	0.96	0.94	0.97		
	ff	100	0.35	0.31	0.40	0.93	0.85	0.97		
		300	0.22	0.19	0.24	0.94	0.89	0.97		
t(3)	smoo	1	100	0.32	0.29	0.36	1.00	0.98	1.00	
			300	0.22	0.20	0.25	0.99	0.97	1.00	
		5	100	0.19	0.18	0.21	0.99	0.97	0.99	
			300	0.13	0.12	0.14	0.98	0.96	0.99	
		te	1	100	0.20	0.18	0.22	0.95	0.93	0.97
			300	0.12	0.11	0.13	0.95	0.94	0.97	
	5	100	0.14	0.12	0.15	0.91	0.89	0.93		
		300	0.08	0.07	0.09	0.90	0.88	0.92		
	ff	1	100	0.36	0.32	0.41	0.93	0.87	0.98	
			300	0.24	0.20	0.26	0.94	0.89	0.98	
		5	100	0.19	0.17	0.22	0.94	0.90	0.99	
			300	0.13	0.11	0.14	0.94	0.90	0.97	

$(f(\text{hum}(t)) + f(\text{temp}(t)))$ concurrent effects, as well as corresponding concurrent interaction effects $f(\text{hum}(t), \text{temp}(t))$ and $f(\text{hum}(t), \text{temp}(t), t)$, respectively. We also considered models with no effects of humidity and/or temperature at all.

We estimated models for all 100 combinations (4 day effects, 5 auto-regressive effects, 5 humidity/temperature effects) of the different effect specifications given above. As in the main analysis, we excluded every third day from the training sample to serve as a validation set. Analysis of mean Brier scores achieved on the training data showed that all 100 models fit the training data similarly well, with slightly better fits for models with i. i. d. or $AR(1)$ functional random effects compared to models with no day effects or a smooth day effect. Analysis of mean Brier scores on the validation data, however, showed that only very small differences in predictive accuracy between models with a smooth day effect and no day effects at all exist and also revealed a strong decrease in predictive performance for models with humidity-temperature effects, especially in combination with cumulative auto-regressive effects. All in all, the majority of models performed worse in terms of predictive Brier score than a simple functional intercept model $y_i(t) \sim B(60, \text{logit}^{-1}(\beta_0(t)))$, which had a mean predictive Brier




FIG B.1. Observed feeding episodes for pig 57. Horizontal axis gives time of day, vertical axis represents days. Black dots show observed feeding episodes as measured by proximity to the trough (yes-no).

score of 0.026, compared to a mean predictive Brier score of 0.029 for the i. i. d. functional random intercept model discussed below and a mean predictive Brier score of 0.024 of a model like the one discussed in Section 3.3 with a smoothly varying day effect better suited for generating interpolating predictions for missing days instead of functional random day effects, which are constant 0 for days in the test set. We conclude that overfitting of complex effects could be a serious concern for the model class we propose, as shown by the large differences in Brier scores between training and validation data for many of the models with complicated effect structures. Note, however, that even though the i. i. d. random effect model yields inferior predictions, it may be better suited for estimating explanatory and interpretable models than the one discussed in Section 3.2, as the larger flexibility of the i. i. d. functional random day effects is presumably more effective at modeling the peaky and irregular temporal dependencies in the observations than a smoothly varying aging effect.

We present detailed results for an alternative model specification with i. i. d. functional random day effects $b_i(t)$ instead of a smooth aging effect $f(i, t)$ below. Figure B.2 shows fitted and observed values for 12 selected days from the training data (top) as well as predicted and observed values for 12 selected days from the validation data (bottom). It is easy to see that the model is able to reproduce many of the feeding episodes in the training data, in the sense that peaks in the estimated probability curves mostly line up well with observed spikes of $y_i(t)$. This model explains about 41% of the deviance and achieves a Brier score of about 0.019 taking the mean over all days and time-points. Prediction (lower panel) with this model is more challenging (Brier score: 0.029), and succeeds only partially. For example, while the predictions for day 73 are mostly very good, predictions for the evening hours of day 11 or around noon on day 82 are far off the mark. Note, however, that $\hat{b}_i(t) \equiv 0$ for the validation data, as $E(b_i(t)) \equiv 0 \forall t, i$. Models that include a smooth day effect $f(i, t)$ instead, which can be interpolated for the days in the calibration data, are more successful at prediction, see Section 3.3.

Figure B.3 shows the estimated components of the additive predictor for this alternative model specification. The functional intercept in the top left panel of Figure B.3 is fairly similar in shape to Figure 2, but overall the base rate is

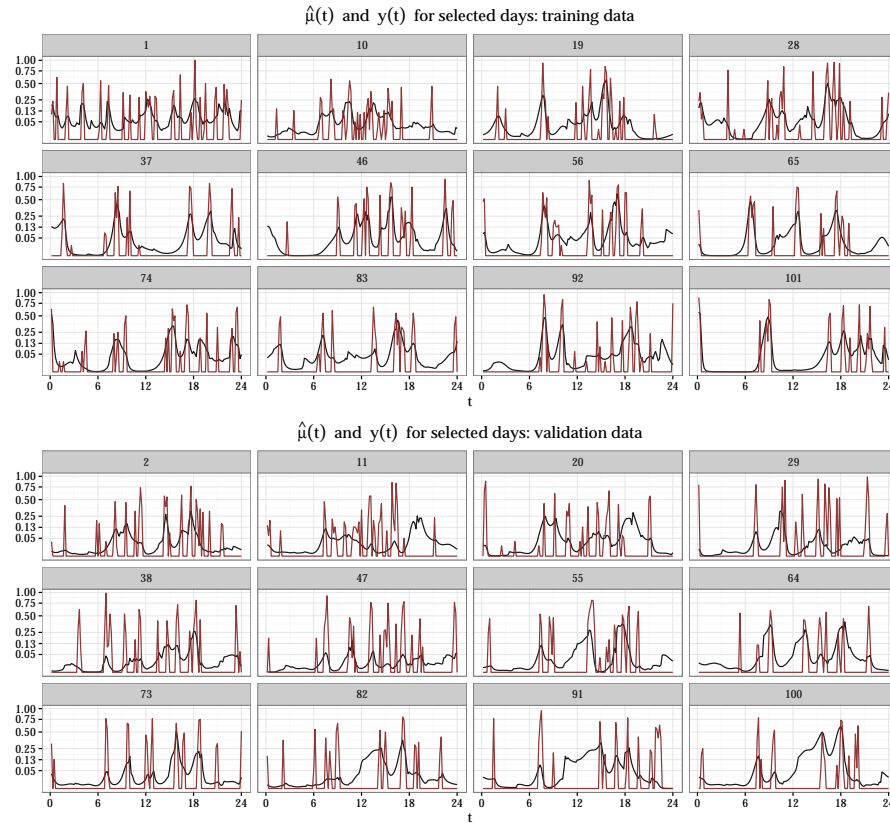


FIG B.2. Fitted (top) and predicted (bottom) values for pig 57 for selected days for an alternative model specification with i. i. d. functional random day effects instead of a smooth aging effect. Black lines for fitted (top) and predicted (bottom) values $\hat{\mu}_i(t)$, red for observed feeding rate $y_i(t)/60$. Numbers above each panel give the day. Vertical axis on $\sqrt{\cdot}$ -scale.

estimated to be much higher and the peak before 6h is much less pronounced. Estimated random day effects are shown in the top right panel of Figure B.3. In terms of absolute size, these are much larger than the values of the smooth aging effect depicted in Figure 2. Effect sizes for $b_i(t)$ are often unrealistically large, with some estimates going as low as -20 (-20 on the logit scale corresponds to a reduction of the probability by a factor of about $2 \cdot 10^{-9}$), presumably an artifact of the very low feeding activity between midnight and early morning leading to (quasi-)complete separability of the data. The vertical axis in Figure B.3 is cut off at -10 in order to showcase the typical “peaky” structure of the $\hat{b}_i(t)$ which would otherwise be hidden by the larger vertical axis required to show the 12 random intercept functions reaching minimal values below -10 in their entirety. The cumulative auto-regressive effect of feeding during the previous 3 hours displayed in the bottom row of Figure B.3 is quite similar to the estimate for the original model shown in Figure 2 but with a weaker positive

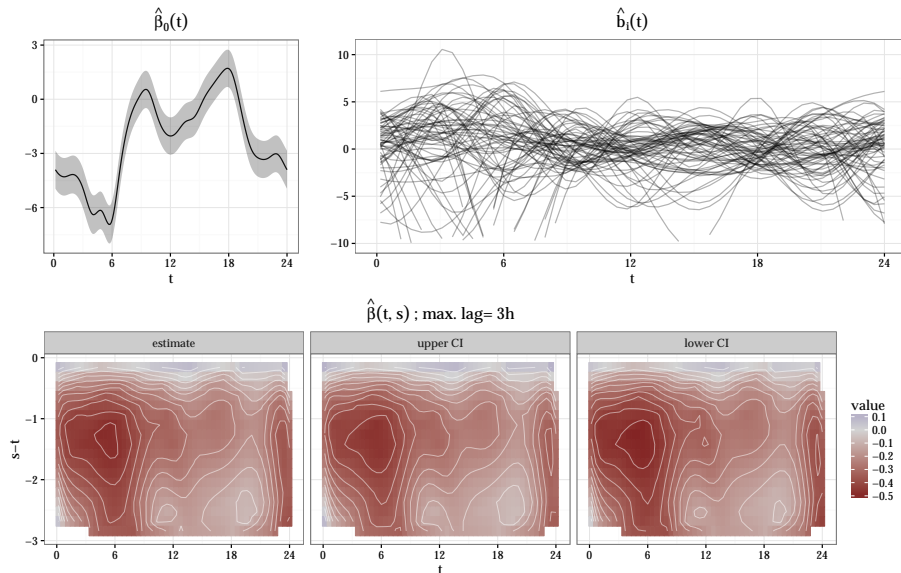


FIG B.3. *Estimated effects for an alternative model specification with i. i. d. functional random day effects instead of a smooth aging effect. Top row: Estimated functional intercept and functional random day effects. Bottom row: Estimated coefficient surface for cumulative autoregressive effect. Intervals are ± 2 standard errors. Vertical axis for functional random day effects is truncated at -10 to show structure of typical results.*

association between feeding in the immediate past and current feeding and a stronger negative association for feeding episodes in the early morning hours.

Acknowledgements

Sonja Greven and Fabian Scheipl were funded by Emmy Noether grant GR 3793/1-1 from the German Research Foundation. We appreciate Engel Hessel's permission to use the PIGWISE data collected in the framework of the ICT-AGRI era-net project PIGWISE "Optimizing performance and welfare of fattening pigs using High Frequent Radio Frequency Identification (HF RFID) and synergistic control on individual level" (Call for transnational research projects 2010). We are grateful for Bo Wang's patient and generous support for our attempt to reproduce the `ggpfr` results. We are indebted to the associate editor and three anonymous reviewers whose constructive remarks helped to improve the manuscript.

Supplementary Material

Code to reproduce results for Section 4

(doi: [10.1214/16-EJS1145SUPP](https://doi.org/10.1214/16-EJS1145SUPP); .zip). Contains code for the simulations and replication studies in Section 4 as well an R Markdown document with graphs of typical results for the Binomial data in Section 4.1

References

- [1] BROCKHAUS, S. (2015). *FDboost: Boosting functional regression models*. R package version 0.0-8.
- [2] BROCKHAUS, S., F. SCHEIPL, T. HOTHORN, and S. GREVEN (2015). The functional linear array model. *Statistical Modelling* 15(3), 279–300. [MR3349797](#)
- [3] GERTHEISS, J., V. MAIER, E. F. HESSEL, and A.-M. STAIKU (2015). Marginal functional regression models for analyzing the feeding behavior of pigs. *Journal of Agricultural, Biological, and Environmental Statistics* 20(3), 353–370. [MR3396556](#)
- [4] GOLDSMITH, J., V. ZIPUNNIKOV, and J. SCHRACK (2015). Generalized multilevel function-on-scalar regression and principal component analysis. *Biometrics* 71(2), 344–353. [MR3366239](#)
- [5] GUILLEMET, R., J. Y. DOURMAD, and M. C. MEUNIER-SALAÜN (2006). Feeding behavior in primiparous lactating sows: Impact of a high-fiber diet during pregnancy. *Journal of Animal Science* 84, 2474–2481.
- [6] HALL, P., H.-G. MÜLLER, and F. YAO (2008). Modelling sparse generalized longitudinal observations with latent Gaussian processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 70(4), 703–723. [MR2523900](#)
- [7] HUANG, L., F. SCHEIPL, J. GOLDSMITH, J. GELLAR, J. HAREZLAK, M. W. MCLEAN, B. SWIHART, L. XIAO, C. CRAINICEANU, and P. REISS (2015). *refund: Regression with Functional Data*. R package version 0.1-12.
- [8] HYUN, Y., M. ELLIS, F. K. MCKEITH, and E. R. WILSON (1997). Feed intake pattern of group-housed growing-finishing pigs monitored using a computerized feed intake recording system. *Journal of Animal Science* 75, 1443–1451.
- [9] LI, H., J. STAUDENMAYER, and R. J. CARROLL (2014). Hierarchical functional data with mixed continuous and binary measurements. *Biometrics* 70(4), 802–811. [MR3295741](#)
- [10] MARRA, G. and S. N. WOOD (2012). Coverage properties of confidence intervals for generalized additive model components. *Scandinavian Journal of Statistics* 39(1), 53–74. [MR2896791](#)
- [11] MASELYNE, J., W. SAEYS, B. D. KETELAERE, K. MERTENS, J. VANGYTE, E. F. HESSEL, S. MILLET, and A. VAN NUFFEL (2014). Validation of a high frequency radio frequency identification (HF RFID) system for registering feeding patterns of growing-finishing pigs. *Computers and Electronics in Agriculture* 102, 10–18.
- [12] MONTGOMERY, G. W., D. S. FLUX, and J. R. CARR (1978). Feeding patterns in pigs: The effects of amino acid deficiency. *Physiology & Behavior* 20, 693–698.
- [13] MORRIS, J. S. (2015). Functional regression. *Annual Review of Statistics and its Applications* 2, 321–359.
- [14] PYA, N. and S. N. WOOD (2016). A note on basis dimension selection

- in generalized additive modelling. Technical report. <http://arxiv.org/abs/1602.06696>.
- [15] R DEVELOPMENT CORE TEAM (2015). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
 - [16] REISS, P. T. and R. T. OGDEN (2009). Smoothing parameter selection for a class of semiparametric linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71(2), 505–523. [MR2649608](#)
 - [17] RIGBY, R. A. and D. M. STASINOPOULOS (2005). Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 54(3), 507–554. [MR2137253](#)
 - [18] RUPPERT, D. (2002). Selecting the number of knots for penalized splines. *Journal of computational and graphical statistics* 11(4), 735–757. [MR1944261](#)
 - [19] SCHEIPL, F., A.-M. STAICU, and S. GREVEN (2015). Functional additive mixed models. *Journal of Computational and Graphical Statistics* 24(3), 477–501. [MR3357391](#)
 - [20] SERBAN, N., A.-M. STAICU, and R. J. CARROLL (2013). Multilevel cross-dependent binary longitudinal data. *Biometrics* 69(4), 903–913. [MR3146786](#)
 - [21] STAN DEVELOPMENT TEAM (2014). *rstan: R interface to Stan, Version 2.6.0*.
 - [22] TASHMAN, L. J. (2000). Out-of-sample tests of forecasting accuracy: an analysis and review. *International Journal of Forecasting* 16(4), 437–450.
 - [23] THE MATHWORKS, INC. (2012). *MATLAB 7.12.0.635*. The MathWorks, Inc., Natick, Massachusetts.
 - [24] TIDEMANN-MILLER, B. A. (2014). *Statistical Modeling of Multivariate Functional Data that Exhibit Complex Correlation Structures*. Ph. D. thesis, North Carolina State University. [MR3321761](#)
 - [25] VAN DER LINDE, A. (2009). A Bayesian latent variable approach to functional principal components analysis with binary and count data. *AStA Advances in Statistical Analysis* 93(3), 307–333. [MR2545698](#)
 - [26] WANG, B. and J. Q. SHI (2014). Generalized Gaussian process regression model for non-gaussian functional data. *Journal of the American Statistical Association* 109(507), 1123–1133. [MR3265685](#)
 - [27] WOOD, S. N. (2006). *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC. [MR2206355](#)
 - [28] WOOD, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B* 73(1), 3–36. [MR2797734](#)
 - [29] WOOD, S. N. (2013). On p-values for smooth components of an extended generalized additive model. *Biometrika* 100(1), 221–228. [MR3034335](#)
 - [30] WOOD, S. N. (2014). General smooth additive modeling. In T. Kneib, F. Sobotka, J. Fahrenholz, and H. Irmer (Eds.), *Proceedings of the 29th International Workshop on Statistical Modelling (IWSM)*, Volume 1, pp. 55–60. Georg-August-Universität Göttingen.

- [31] WOOD, S. N., N. PYA, and B. SÄFKEN (2015). Smoothing parameter and model selection for general smooth models. Technical report. <http://arxiv.org/abs/1511.03864>. MR3334416
- [32] ZHU, H., P. J. BROWN, and J. S. MORRIS (2011). Robust, adaptive functional regression in functional mixed model framework. *Journal of the American Statistical Association* 106(495), 1167–1179. MR2894772