

FUNCTIONAL ADDITIVE REGRESSION

BY YINGYING FAN¹, GARETH M. JAMES AND PETER RADCHENKO²

University of Southern California

We suggest a new method, called *Functional Additive Regression*, or FAR, for efficiently performing high-dimensional functional regression. FAR extends the usual linear regression model involving a functional predictor, $X(t)$, and a scalar response, Y , in two key respects. First, FAR uses a penalized least squares optimization approach to efficiently deal with high-dimensional problems involving a large number of functional predictors. Second, FAR extends beyond the standard linear regression setting to fit general nonlinear additive models. We demonstrate that FAR can be implemented with a wide range of penalty functions using a highly efficient coordinate descent algorithm. Theoretical results are developed which provide motivation for the FAR optimization criterion. Finally, we show through simulations and two real data sets that FAR can significantly outperform competing methods.

1. Introduction. The univariate functional regression situation, where one models the relationship between a scalar response, Y , and a functional predictor, $X(t)$, has recently received a great deal of attention. A few examples include [2, 6, 7, 15, 16, 18–20, 22, 23, 30]. See Chapter 15 of [32] for a thorough discussion of the issues involved with fitting such data.

Most work in this area involves different approaches for fitting the functional linear regression model,

$$(1) \quad Y_i = \int \beta(t) X_i(t) dt + \varepsilon_i, \quad i = 1, \dots, n.$$

For notational convenience, we assume throughout this paper that the response and predictors have been centered so the intercept can be ignored. Model (1) provides a natural extension of linear regression to the functional domain but it has two significant limitations. First, it assumes a single predictor, while functional regression situations involving a large number of predictors, $X_{i1}(t)$, $X_{i2}(t)$, \dots , $X_{ip}(t)$, are becoming increasingly common. For example, [36] analyzes two gene expression data sets measured over time, which involve only a small number of patients but tens of thousands of functional predictors. Second, (1) is relatively inflexible

Received November 2014; revised May 2015.

¹Supported by NSF CAREER Award DMS-11-50318 and USC Marshall summer research funding.

²Supported in part by NSF Grant DMS-12-09057.

MSC2010 subject classifications. Primary 62G08; secondary 62G20.

Key words and phrases. Functional regression, shrinkage, single index model, variable selection.

because it assumes a linear relationship between the predictor and response. Just as in the standard regression setting more accurate fits can often be produced by modeling a nonlinear relationship.

In this paper, we address both of these limitations using a functional additive regression framework of the form

$$(2) \quad Y_i = \sum_{j=1}^p f_j(X_{ij}) + \varepsilon_i, \quad i = 1, \dots, n,$$

where the f_j 's are general nonlinear functions of $X_{ij}(t)$. There has been some previous work extending the classical functional regression model. James and Silverman [23] proposed an index model to implement a nonlinear functional regression, and, more recently, both [14] and [7] extended this work to a fully nonparametric setting and provided further theoretical motivation. However, all of these approaches are primarily intended for the univariate setting, where $p = 1$. Lian [25] did consider a multivariate setting involving both functional and scalar predictors, but with only a single functional predictor, so the corresponding model does not extend to (2). James and Silverman [23] proposed a kernel based method for fitting (2), which works well in low-dimensional situations. However, they do not attempt to perform any kind of variable selection. As a result, the method suffers from computational and statistical issues when p is large, such as for the gene expression data in [36]. Zhu et al. [40] proposed a Bayesian variable selection approach for selecting and estimating important functional predictors in a classification setting. However, while their method can potentially be implemented on a large number of functions, it still assumes a linear relationship between the response and predictors. Finally, a recent paper [11] considers a more general form of (2) where the response is also functional. Their approach appears to work well but the paper does not provide any theoretical results. See also [13, 17, 28] for additional recent developments on functional regression models with multiple functional covariates under various model settings.

Fitting (2) in the high-dimensional setting poses a couple of significant complications. First, in order to make the problem feasible, we must assume sparsity in the predictor space, that is, that most of the predictors are unrelated to the response. Thus, we need an approach that can automatically perform high-dimensional variable selection on nonlinear functions. Second, (2) involves estimating functions, $f_j(x)$, of functional predictors, $X_{ij}(t)$. Even in the univariate situation, involving a single predictor, there has been little research on this problem and the best approach is unclear. Most current methods involve using the first few functional principal component scores of $X_{ij}(t)$ as a finite-dimensional predictor space [31]. However, the principal component scores are computed independently from the response, in an unsupervised fashion, so there is no a priori reason to believe that these scores will correspond to the best dimensions for the regression problem.

In this paper, we suggest a new penalized least squares method called *Functional Additive Regression*, or FAR, for fitting a nonlinear functional additive model. FAR makes three important contributions. First, it efficiently fits high-dimensional functional models while simultaneously performing variable selection to identify the relevant predictors. This is an area that has historically received very little attention in the functional domain, but the importance of the connections between functional and high-dimensional statistics are just starting to become clear. See, for example, the recent conference on this topic [4].

Second, FAR extends beyond the standard linear regression setting to fit general nonlinear additive models. FAR models $f_j(x)$ as a nonlinear function of a one-dimensional linear projection of $X_{ij}(t)$; a functional version of the *single index model* approach. Our method uses a supervised fit to automatically project the functional predictors into the best one-dimensional space. We believe this is an important distinction because projecting into the unsupervised PCA space is currently the dominant approach in functional regressions, even though it is well known that this space need not be optimal for predicting the response.

Third, FAR can be implemented using a wide range of penalty functions and a highly efficient coordinate descent algorithm. In the linear case, we establish a number of theoretical results, which show that, under suitable conditions and for an appropriately chosen penalty function, FAR is guaranteed to asymptotically choose the correct model as n and p go to infinity. Theoretical investigation for the nonlinear FAR approach presents some serious additional challenges, because the regression functions, f_j , are estimated rather than known. We allow the number of functional predictors, p , to grow faster than the number of observations, n , and establish asymptotic bounds on the ℓ_2 estimation error for each of the estimated regression functions. The difficulties associated with the high-dimensional nature of the functional data are exacerbated by the large number of estimated components in the additive regression model for the response. Moreover, the functional aspect of the data (infinite dimensional predictors) adds further complexity to the already very challenging problem. Our method of proof uses ideas from the estimation theory for high-dimensional additive models [5, 21, 29]. However, the proof itself is new, rather than a compilation of existing results.

Our paper is set out as follows. In Section 2, we develop the FAR method for performing high-dimensional functional regression. Section 2.1 uses functional index models to motivate the FAR model. Then Section 2.2 presents the optimization criterion and an efficient coordinate descent algorithm for fitting FAR in the linear regression setting. Finally, Section 2.3 extends the algorithm to the nonlinear regression framework. In Section 3, we provide a number of theoretical results. We first prove that, under appropriate conditions, the linear version of FAR will asymptotically include all the true signal variables and remove all the noise predictors from the model. In addition, we provide an asymptotic bound on the estimation error of the signal functions, $f_j(x)$, under the vector infinity norm, and

show that the FAR estimator is asymptotically normal. In the nonlinear setting, we establish the rate of convergence, with respect to the ℓ_2 distance, for the estimates of the regression functions, $f_j(x)$, corresponding to each of the predictors. We also investigate the variable selection properties of our estimator and show that, under some conditions, it can recover the index set of the signal predictors. Extensive simulation results are presented in Section 4. We compare FAR to other functional regression methods and demonstrate its superior performance in many settings. Finally, we apply FAR to both medium and high-dimensional real data sets in Section 5, and end with a discussion in Section 6.

2. Functional additive regression. Let $\mathbf{f}_j = (f_j(X_{1j}), \dots, f_j(X_{nj}))^T$. Then our general approach for fitting (2) is to minimize the following penalized regression criterion over $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_p$:

$$(3) \quad \frac{1}{2n} \left\| \mathbf{Y} - \sum_{j=1}^p \mathbf{f}_j \right\|_2^2 + \sum_{j=1}^p \rho_{\lambda_n} \left(\frac{1}{\sqrt{n}} \|\mathbf{f}_j\|_2 \right),$$

where $\mathbf{Y} = (Y_1, \dots, Y_n)^T$, $\rho_{\lambda_n}(t)$ is a penalty function, λ_n is the regularization parameter and $\|\mathbf{f}_j\|_2 = \sqrt{\mathbf{f}_j^T \mathbf{f}_j}$. To aid the presentation, we drop the subscript and use $\|\cdot\|$ to denote the ℓ_2 norm of a vector in the future. Although it may not be immediately obvious from this formulation, we show that minimizing (3) will in general automatically implement variable selection by shrinking a subset of the \mathbf{f}_j 's to exactly zero. In this article, we explore general concave functions for ρ , with the ℓ_1 penalty $\rho_\lambda(t) = \lambda t$ considered as a special case. There is by now a substantial literature demonstrating the advantages of concave penalty functions for high-dimensional problems [8–10, 26, 27].

We assume that the trajectories of functional predictors, $X_{ij}(t)$, are fully observed. Our methodology and theoretical results can be extended to the case of densely observed predictors under additional smoothness and regularity assumptions. However, for the clarity of the exposition we do not investigate this case in the paper.

2.1. Functional index models. Minimizing (3) requires specifying the form of $f_j(x)$. A limitation of linear functional regression models is that they can perform poorly when there is a nonlinear relationship between $X(t)$ and Y . However, the infinite-dimensional nature of $X(t)$ makes it challenging to model a nonlinear relationship between the predictor and response. As a result, relatively few papers have investigated this extension. Most methods focus on approximating $X(t)$ using its first few functional principal components and then implementing nonlinear fits using the principal component scores as predictors [31]. However, this unsupervised approach has the usual limitation; the directions which explain $X(t)$ best may not be the most appropriate for predicting the response.

In the multivariate setting, index models are commonly used for providing non-linear fits to high-dimensional data. For a centered response, the standard single index model can be expressed in the form $Y = g(\boldsymbol{\beta}^T \mathbf{X}) + \varepsilon$, where $g(x)$ is a general nonlinear function and $\boldsymbol{\beta}$ is a norm one vector representing the best single direction to project the predictors into. A key advantage of the index model formulation is that $\boldsymbol{\beta}$ is chosen in a supervised fashion, incorporating both the response and predictors, potentially providing more accurate fits. Index models can be naturally extended to functional predictors using the formulation $f_j(X_{ij}) = g_j(\int \beta_j(t) X_{ij}(t) dt)$, where $g_j(x)$ and $\beta_j(t)$ are both nonparametric smooth functions, and the integral is well-defined. Functional single index models have been considered previously. For example, [1, 3, 7, 14, 23], all fit index models to functional data, but these previous approaches all concentrate on the $p = 1$ problem.

Using this nonlinear representation, the FAR model (2) can be expressed as

$$(4) \quad Y_i = \sum_{j=1}^p g_j \left(\int \beta_j(t) X_{ij}(t) dt \right) + \varepsilon_i.$$

For identifiability, in addition to centering the response, we also center the regression functions: $\sum_{i=1}^n g_j(\int \beta_j(t) X_{ij}(t) dt) = 0$ for all j . Note that index functions β_j are only identifiable up to multiplications by nonzero constants, however, our focus is on estimating f_j rather than β_j . The general FAR optimization criterion (3) becomes

$$(5) \quad \frac{1}{2n} \left\| \mathbf{Y} - \sum_{j=1}^p g_j \left(\int \beta_j(t) \mathbf{X}_j(t) dt \right) \right\|^2 + \sum_{j=1}^p \rho_{\lambda_n} \left(\frac{1}{\sqrt{n}} \|\mathbf{f}_j\| \right),$$

where $\mathbf{X}_j(t) = (X_{1j}(t), \dots, X_{nj}(t))^T$ and $g_j(\int \beta_j(t) \mathbf{X}_j(t) dt) = (f_j(X_{1j}), \dots, f_j(X_{nj}))^T$.

2.2. Linear FAR. Our approach for minimizing (5) is easiest to understand by first considering the situation where $f_j(x)$ is taken to be linear. Hence, in this section we develop FAR in the setting where $g_j(x)$ is set to the identity function, in which case FAR reduces to a multivariate functional linear regression model.

2.2.1. FAR criterion. We assume without loss of generality that each predictor is observed over the range $0 \leq t \leq 1$. Hence, in the linear setting,

$$(6) \quad f_j(X_{ij}) = \int_0^1 \beta_j(t) X_{ij}(t) dt,$$

where $\beta_j(t)$ is an unknown smooth coefficient function, and the FAR optimization criterion becomes

$$(7) \quad \frac{1}{2n} \left\| \mathbf{Y} - \sum_{j=1}^p \int_0^1 \beta_j(t) \mathbf{X}_j(t) dt \right\|^2 + \sum_{j=1}^p \rho_{\lambda_n} \left(\frac{1}{\sqrt{n}} \|\mathbf{f}_j\| \right),$$

where $\mathbf{X}_j(t) = (X_{1j}(t), \dots, X_{nj}(t))^T$.

Given an orthonormal basis $\{b_l(t)\}$, the functional predictors and the corresponding regression coefficients can be decomposed as

$$(8) \quad X_{ij}(t) = \sum_{l=1}^{\infty} \theta_{ijl} b_l(t), \quad \beta_j(t) = \sum_{l=1}^{\infty} \eta_{0,jl} b_l(t),$$

where θ_{ijl} and $\eta_{0,jl}$ are the coefficients of $X_{ij}(t)$ and $\beta_j(t)$ corresponding to the l th basis function $b_l(t)$, respectively. Using (8), the j th additive component has the following representation

$$(9) \quad f_j(X_{ij}) = \int_0^1 X_{ij}(t) \beta_j(t) dt = \sum_{l=1}^{\infty} \theta_{ijl} \eta_{0,jl}.$$

In order for the functions optimizing (7) to have nontrivial solutions, some form of smoothness constraint must be imposed on the $\beta_j(t)$'s. Two standard approaches are to include a smoothness penalty in the optimization criterion or alternatively to restrict the functions to some low-dimensional class. In this setting, either approach could be adopted but we use the latter method. Specifically, for a given sequence of integers $q_n = o(n)$ depending only on the sample size n , write $\boldsymbol{\eta}_{0j} = (\eta_{0,j1}, \dots, \eta_{0,jq_n})^T$ and $\boldsymbol{\theta}_{ij} = (\theta_{ij1}, \dots, \theta_{ijq_n})^T$. Thus, the j th additive component $f_j(X_{ij})$ can be approximately as $\boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_{0j}$. Denote by e_{ij} the approximation error, that is,

$$(10) \quad e_{ij} = f_j(X_{ij}) - \boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_{0j} = \sum_{l=q_n+1}^{\infty} \theta_{ijl} \eta_{0,jl}.$$

Then by the Cauchy–Schwarz inequality and Condition 1 in Appendix B, uniformly across all $i = 1, \dots, n$ and $j \in \mathfrak{M}_0$,

$$(11) \quad |e_{ij}|^2 \leq \sum_{l=q_n+1}^{\infty} \eta_{0,jl}^2 l^{-4} \sum_{l=q_n+1}^{\infty} \theta_{ijl}^2 l^4 \leq C^2 q_n^{-4} \sum_{l=q_n+1}^{\infty} \eta_{0,jl}^2 \leq \tilde{C} C^2 q_n^{-4},$$

where C and \tilde{C} are two positive constants defined in Condition 1. Thus, for large enough q_n , the approximation error is uniformly small.

Let Θ_j be an $n \times q_n$ matrix whose rows are formed by $\{\boldsymbol{\theta}_{ij}, i = 1, \dots, n\}$. Then, if q_n is large enough, $f_j(X_{ij}) \approx \boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_{0j}$ and (7) can be approximated by

$$(12) \quad \frac{1}{2n} \left\| \mathbf{Y} - \sum_{j=1}^p \Theta_j \boldsymbol{\eta}_j \right\|^2 + \sum_{j=1}^p \rho_{\lambda_n} \left(\frac{1}{\sqrt{n}} \|\Theta_j \boldsymbol{\eta}_j\| \right).$$

Note that the $\boldsymbol{\eta}_j$'s must be estimated, but the Θ_j 's are calculated from the fully observed trajectories of the functional predictors, $X_{ij}(t)$. Hence, we fit FAR by minimizing (12) over $\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_p$.

2.2.2. *FAR algorithm.* The criterion given by (12) is still $p \times q_n$ dimensional, so is potentially challenging to optimize over, even if p is only of moderate size. However, in this form our FAR criterion is closely related to the standardized group lasso [35] which allows us to develop an efficient algorithm to fit FAR. In particular, a distinct advantage of (12) is that, when using the Lasso penalty $\rho_{\lambda_n}(t) = \lambda_n t$, there is a simple closed form expression for computing its minimum over $\boldsymbol{\eta}_j$.

PROPOSITION 1. *If $\rho_{\lambda_n}(t) = \lambda_n t$, then the solution to (12) satisfies $\hat{\mathbf{f}}_j = \Theta_j \hat{\boldsymbol{\eta}}_j$ where*

$$\hat{\boldsymbol{\eta}}_j = \left(1 - \frac{\sqrt{n}\lambda_n}{\|S_j \mathbf{R}_j\|}\right)_+ (\Theta_j^T \Theta_j)^{-1} \Theta_j^T \mathbf{R}_j,$$

$S_j = \Theta_j (\Theta_j^T \Theta_j)^{-1} \Theta_j^T$, $\mathbf{R}_j = \mathbf{Y} - \sum_{k \neq j} \Theta_k \hat{\boldsymbol{\eta}}_k$, and $z_+ = \max(0, z)$ represents the positive part of z .

The derivation of Proposition 1 involves simple algebra and similar results are proved in [33] and [35] so we do not provide the proof here. Proposition 1 suggests Algorithm 1, a simple but very efficient coordinate descent algorithm for minimizing (12) when $\rho_{\lambda_n}(t) = \lambda_n t$.

We repeat this algorithm over a grid of values for λ , using the previous values for the $\hat{\boldsymbol{\eta}}_j$'s to initialize the parameters for the new λ . Since the parameters change very little for a small change in λ , the algorithm generally converges very rapidly. Note that the S_j 's only need to be computed once for all values of λ so the computation at each step of the algorithm is extremely fast. In addition, it is clear from Proposition 1 that (12) will decrease at each step. This approach has the advantage of decomposing the estimation of $\hat{\mathbf{f}}_j$ into two simple, and separate, steps. First, compute the unshrunk estimate $\hat{\mathbf{P}}_j$ and second, apply the shrinkage factor α_j . When $\alpha_j = 0$ then the j th predictor is absent from the model. Our FAR algorithm has similarities to the SpAM algorithm [33] but SpAM cannot model functional data.

For a general penalty function, $\rho_{\lambda_n}(t)$, we use the local linear approximation method proposed in [42] to solve (12). The penalty function can be approximated

Algorithm 1 Linear FAR algorithm

0. Initialize $\hat{\boldsymbol{\eta}}_j = \mathbf{0}$ and $S_j = \Theta_j (\Theta_j^T \Theta_j)^{-1} \Theta_j^T$, for $j \in \{1, \dots, p\}$.
 1. Fix all $\hat{\mathbf{f}}_k$ for $k \neq j$. Compute the residual vector $\mathbf{R}_j = \mathbf{Y} - \sum_{k \neq j} \hat{\mathbf{f}}_k$.
 2. Let $\hat{\mathbf{P}}_j = S_j \mathbf{R}_j$ represent the unshrunk estimate for \mathbf{f}_j .
 3. Let $\hat{\mathbf{f}}_j = \alpha_j \hat{\mathbf{P}}_j$ where $\alpha_j = (1 - \lambda_n \sqrt{n} / \|\hat{\mathbf{P}}_j\|)_+$ is a shrinkage parameter.
 4. Center $\hat{\mathbf{f}}_j \leftarrow \hat{\mathbf{f}}_j - \text{mean}(\hat{\mathbf{f}}_j)$.
 5. Repeat steps 1 through 4 for $j = 1, 2, \dots, p$ and iterate until convergence.
-

as $\rho_{\lambda_n}(\|\mathbf{f}\|/\sqrt{n}) \approx \rho'_{\lambda_n}(\|\mathbf{f}^*\|/\sqrt{n})\|\mathbf{f}\|/\sqrt{n} + C$, where \mathbf{f}^* is some vector that is close to \mathbf{f} and $C = \rho_{\lambda_n}(\|\mathbf{f}^*\|/\sqrt{n}) - \rho'_{\lambda_n}(\|\mathbf{f}^*\|/\sqrt{n})\|\mathbf{f}^*\|/\sqrt{n}$ is a constant. Hence, the only required change to the FAR algorithm for optimizing over general penalty functions is to replace the calculation of α_j in step 3 by

$$\alpha_j = \left(1 - \rho'_{\lambda_n} \left(\frac{1}{\sqrt{n}} \|\hat{\mathbf{f}}_j\| \right) \sqrt{n} / \|\hat{\mathbf{P}}_j\| \right)_+,$$

where $\hat{\mathbf{f}}_j$ represents the most recent estimate for \mathbf{f}_j . The initial estimate of $\hat{\mathbf{f}}_j$ can be obtained by using the Lasso penalty. This simple approximation allows the FAR algorithm to be easily applied to a wide range of penalty functions.

2.3. *Nonlinear FAR.* We now consider the more general nonlinear setting (4) where $g_j(x)$ is estimated as part of the fitting process. Since $\beta_j(t)$ corresponds to a direction that we project $X_{ij}(t)$ into we impose the constraint $\|\beta_j\|_2 = 1$. Note that β_j are still not uniquely identifiable, however, our focus is on estimating the regression functions, f_j , rather than the index functions. We assume that $g_j(x)$ can be well approximated by a d_n -dimensional basis $\mathbf{h}(x)$ such that $g_j(x) \approx \mathbf{h}(x)^T \boldsymbol{\xi}_j$. Using this basis, representation (5) can be expressed as

$$(13) \quad \frac{1}{2n} \left\| \mathbf{Y} - \sum_{j=1}^p H_j \boldsymbol{\xi}_j \right\|^2 + \sum_{j=1}^p \rho_{\lambda_n} \left(\frac{1}{\sqrt{n}} \|H_j \boldsymbol{\xi}_j\| \right),$$

where H_j is an n by d_n matrix who's i th row is given by $\mathbf{h}(\boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j)^T$.

We use an iterative algorithm to approximately minimize (13) over $\boldsymbol{\xi}_j$ and $\boldsymbol{\eta}_j$. First, given current estimates for the $\boldsymbol{\eta}_j$'s we minimize (13) over $\boldsymbol{\xi}_j$. Second, given current estimates for the $\boldsymbol{\xi}_j$'s we minimize the sum of squares term

$$(14) \quad \sum_{i=1}^n \left(Y_i - \sum_{j=1}^p \mathbf{h}(\boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j)^T \hat{\boldsymbol{\xi}}_j \right)^2$$

over $\boldsymbol{\eta}_j$. Note that we do not include the penalty ρ_{λ_n} when estimating $\boldsymbol{\eta}_j$ because the $\boldsymbol{\eta}_j$'s are providing a direction in which to project $X_{ij}(t)$ so are constrained to be norm one. Hence, applying a shrinkage term would be inappropriate.

Formally, the nonlinear FAR algorithm can be summarized as follows (Algorithm 2).

One of the appealing aspects of this approach is that, for fixed \hat{H}_j , (12) and (13) are equivalent so estimation of the $\boldsymbol{\xi}_j$'s in step 2 can be achieved using the linear FAR algorithm from Section 2.2.2. Minimization of (14) in step 3 can be approximately achieved using a first-order Taylor series approximation of $g_j(x)$. We provide details on this minimization and on computing initial values for the $\boldsymbol{\eta}_j$'s in Appendix A.

Algorithm 2 Nonlinear FAR algorithm

0. Initialize $\hat{\eta}_j$ for $j \in \{1, \dots, p\}$ using the linear FAR algorithm.
 1. Compute \hat{H}_j using the current estimates for η_j .
 2. Estimate $\hat{\xi}_j$ for $j \in \{1, \dots, p\}$ by minimizing (13) given the current values of \hat{H}_j .
 3. Conditional on the $\hat{\xi}_j$'s from step 2, estimate the η_j 's by minimizing (14).
 4. Repeat steps 1 through 3 and iterate until convergence.
-

Potentially one could compute the nonlinear FAR algorithm for each possible λ . However, we have found that a more efficient approach is to compute initial estimates for η_j , minimize (13) over ξ_j for each possible value of λ , choose the ξ_j 's corresponding to the "best" value of λ , estimate the η_j 's for only this one set of parameters, and iterate. This approach means that, for each iteration, the minimization of (14) only needs to be performed for a single value of λ . The choice of λ can be made using a variety of methods, as discussed in the next section.

2.4. Selecting tuning parameters. Both the linear and nonlinear versions of FAR require choosing the tuning parameter, λ . As with all penalized regression methods, there are several possible methods one could adopt. Popular approaches include, BIC, AIC or cross-validation. The BIC and AIC methods require the calculation of the effective degrees of freedom. For the Lasso, it has been shown that an unbiased estimate for this quantity is the number of nonzero coefficients [41]. One could potentially use the same value for FAR. However, given FAR's more complicated structure it is not clear that this is still an appropriate estimate. Computing the effective degrees of freedom for FAR is a topic for future research. For our simulations and one real data example, we selected λ using a separate validation data set. For the other real data example, we selected λ using the 20-fold cross-validation method, since there were not enough data points to be used as validation data.

3. Theory.

3.1. Linear theory. Denote by $\mathfrak{M}_0 = \{j : \beta_j(t) \neq 0, 1 \leq j \leq p\}$ the set of true functional predictors and let s_n represent the cardinality of \mathfrak{M}_0 . By minimizing the FAR criterion (12), we aim to identify the set \mathfrak{M}_0 and accurately estimate functions $\beta_j(t)$ for $j \in \mathfrak{M}_0$. In this section, we discuss the theoretical properties of FAR in the setting where the f_j 's are linear functions, that is, $f_j(X_{ij}) = \int_0^1 X_{ij}(t)\beta_j(t) dt$. In particular, we present two theorems, both of which are conditional on the observed predictors, $X_{ij}(t)$, $i = 1, \dots, n$, $j = 1, \dots, p$. Theorem 1 concerns FAR's model selection properties. We show that, with probability tending to one, FAR can remove all noise predictors from the fitted model. Theorem 1 also places an

error bound on the estimated f_j 's under the vector infinity norm, where $j \in \mathfrak{M}_0$. Our second result, Theorem 2 shows the asymptotic normality of the estimator.

In order to prove these results, we make two sets of assumptions. The first set of conditions relates to the level of accuracy in our basis approximations of $X_{ij}(t)$ and $\beta_j(t)$. The second set of conditions concerns the shape of the penalty function, the strength of the signal and the correlation structure of the predictors. Explicit conditions can be found in Appendix B.

Let $\boldsymbol{\eta}_0 = (\boldsymbol{\eta}_{0,1}, \dots, \boldsymbol{\eta}_{0,p}) \in R^{pq_n}$ with $\boldsymbol{\eta}_{0j}$ representing the true coefficient vector in the basis representation $f_j(X_{ij}) = \boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_{0j} + e_{ij}$. For any index set $S \subset \{1, \dots, p\}$, we use $\boldsymbol{\eta}_S$ to denote the vector formed by stacking vectors $\boldsymbol{\eta}_j$, $j \in S$ one underneath each other, and Θ_S to denote the matrix formed by stacking the matrices Θ_j , $j \in S$ one after another. Moreover, we standardize each column of Θ such that they all have ℓ_2 -norm \sqrt{n} . Theorem 1 below shows that FAR possesses the oracle property for model selection.

THEOREM 1. *Assume that $q_n + \log p = O(n\lambda_n^2)$, $\lambda_n n^\alpha q_n \sqrt{s_n} \rightarrow 0$, and $\log(pq_n) = o(n^{1-2\alpha} s_n^{-1} q_n^{-2})$ with α defined in Condition 2(B). Further assume that $s_n q_n^{-2} = o(\lambda_n)$, then under Conditions 1 and 2, with probability tending to 1 as $n \rightarrow \infty$, there exists a local minimizer $\hat{\boldsymbol{\eta}}$ of (12) such that:*

- (1) $\hat{\boldsymbol{\eta}}_{\mathfrak{M}_0^c} = 0$,
- (2) $\|\hat{\boldsymbol{\eta}}_{\mathfrak{M}_0} - \boldsymbol{\eta}_{0\mathfrak{M}_0}\|_\infty \leq c_0^{1/2} n^{-\alpha} q_n^{-1/2}$,

where $\|\cdot\|_\infty$ stands for the infinity norm of a vector.

Although Theorem 1 is on a local minimizer of the linear FAR criterion (12), it has been proved by [26] that any local minimizer will fall within statistical precision of the true parameter vector under appropriate conditions on the penalty function. Part 2 of Theorem 1 concerns the approximation accuracy of the basis coefficients rather than the functions themselves. However, the result extends naturally. Denote by $\hat{\mathbf{f}}_j = \Theta_j \hat{\boldsymbol{\eta}}_j$ and $\mathbf{f}_{0j} = (f_j(X_{j1}), \dots, f_j(X_{jn}))^T$, respectively, the estimated and true values of the j th functional component, both evaluated at the n training data points. Then the corollary below follows immediately from Theorem 1 and Condition 1.

COROLLARY 1. *Suppose the conditions in Theorem 1 are satisfied. Then with probability tending to 1 as $n \rightarrow \infty$, there exists a FAR estimate such that $\hat{\mathbf{f}}_j = 0$ for $j \notin \mathfrak{M}_0$, and*

$$\max_{j \in \mathfrak{M}_0} \frac{1}{\sqrt{n}} \|\hat{\mathbf{f}}_j - \mathbf{f}_{0j}\|_2 \leq C_2 n^{-\alpha},$$

where C_2 is some positive constant.

Theorem 2 shows the asymptotic normality of the FAR estimators that correspond to signal variables. As with Theorem 1, we first provide the result for the $\hat{\eta}_j$'s and then extend to the functions.

THEOREM 2. *Assume that the conditions in Theorem 1 hold and in addition, $\rho'_{\lambda_n}(a_n/2) = o(a_n n^{\alpha-1/2} s_n^{-1/2})$, $\sup_{t \geq a_n/2} \rho''_{\lambda_n}(t) = O(n^{-1/2})$, $s_n = o(n^{2\alpha})$ and $s_n q_n^{-2} = o(n^{-1/2})$. Then with probability tending to 1 as $n \rightarrow \infty$, there exists a strict local minimizer $\hat{\eta}$ of (12) such that $\hat{\eta}_{\mathfrak{M}_0^c} = 0$ and*

$$\mathbf{c}^T [(\Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0})^{1/2} (\hat{\eta}_{\mathfrak{M}_0} - \eta_{0,\mathfrak{M}_0}) + n(\Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0})^{-1/2} \mathbf{v}_{0,\mathfrak{M}_0}] \xrightarrow{\mathcal{D}} N(0, \sigma^2),$$

where $\mathbf{c} \in R^{q_n s_n}$ satisfies $\mathbf{c}^T \mathbf{c} = 1$ and $\mathbf{v}_{0,\mathfrak{M}_0}$ is a vector formed by stacking the vectors $\mathbf{v}_{0,k} = \rho'_{\lambda_n}(\frac{1}{\sqrt{n}} \|\Theta_k \eta_{0,k}\|) \frac{1}{\sqrt{n}} \frac{\Theta_k^T \Theta_k \eta_{0,k}}{\|\Theta_k \eta_{0,k}\|}$, $k \in \mathfrak{M}_0$ underneath each other.

Let $f_{0j}^* = \theta_j^{*T} \eta_{0j}$ and $\hat{f}_j^* = \theta_j^{*T} \hat{\eta}_j$, with $\theta_j^* \in R^{q_n}$ the coefficient vector when projecting a given new observation, $X_j^*(t)$, onto the basis function, $\mathbf{b}(t)$. Then as q_n increases, f_{0j}^* better approximates $f_j(X_j^*)$ for each fixed $j = 1, \dots, p$. Define $\mathbf{f}_0^* = (f_{01}^*, \dots, f_{0p}^*)^T$ and $\hat{\mathbf{f}}^* = (\hat{f}_1^*, \dots, \hat{f}_p^*)^T$. Taking $\mathbf{c} = (\Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0})^{-1/2} \Theta^* \tilde{\mathbf{c}}_0$ with $\Theta^* = \text{diag}(\theta_1^*, \dots, \theta_{s_n}^*) \in R^{(q_n s_n) \times s_n}$ in Theorem 2 and $\tilde{\mathbf{c}}_0$ a vector in R^{s_n} , we have the following asymptotic normality of \mathbf{f}_0^* .

COROLLARY 2. *Assume that the conditions in Theorem 2 hold. Then with probability tending to 1 as $n \rightarrow \infty$, there exists a FAR estimate such that $\hat{\mathbf{f}}_{\mathfrak{M}_0^c}^* = 0$. Moreover,*

$$\tilde{\mathbf{c}}_0^T [\hat{\mathbf{f}}_{\mathfrak{M}_0}^* - \mathbf{f}_{0,\mathfrak{M}_0}^* + n(\Theta^*)^T (\Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0})^{-1} \mathbf{v}_{0,\mathfrak{M}_0}] \xrightarrow{\mathcal{D}} N(0, \sigma^2),$$

where $\tilde{\mathbf{c}}_0$ is a vector in R^{s_n} satisfying $\tilde{\mathbf{c}}_0^T (\Theta^*)^T (\Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0})^{-1} \Theta^* \tilde{\mathbf{c}}_0 = 1$, and $\mathbf{v}_{0,\mathfrak{M}_0}$ is defined in Theorem 2.

3.2. Nonlinear theory. Throughout this section, we focus on the minimizer of the nonlinear FAR criterion with the ℓ_1 penalty function. We treat all the predictors as deterministic. For identifiability purposes, we assume that the true regression functions, f_{0j} , as well as the response vector, are centered, that is, $\sum_{i=1}^n f_{0j}(X_{ij}) = 0$ and $\sum_{i=1}^n Y_i = 0$. As a result, the corresponding estimates, \hat{f}_j , are automatically centered as well.

We use cubic B-splines to approximate the true ‘‘link’’ functions, g_{0j} . Given a candidate index vector η_j , the B-spline basis for representing a candidate link function for the j th predictor is constructed using uniformly placed knots on the interval $[\min_i \eta_j^T \theta_{ij}, \max_i \eta_j^T \theta_{ij}]$. The corresponding row vector valued basis function is denoted by $\mathbf{h}_{\eta_j,j}$. We denote by \mathcal{F}_j^0 the class of candidate regression functions for the j th predictor. More specifically, $\mathcal{F}_j^0 = \{f(\cdot) =$

$\mathbf{h}_{\eta_j, j}(\eta_j^T \cdot) \boldsymbol{\xi}$, $\sum_{i=1}^n f(\boldsymbol{\theta}_{ij}) = 0$, $\eta_j \in \mathbb{R}^{q_n}$, $\boldsymbol{\xi} \in \mathbb{R}^{d_n}$, $\|\eta_j\| = 1$. If \tilde{f} and \check{f} belong to \mathcal{F}_j^0 , we denote by $\|\tilde{f} - \check{f}\|_n$ the ℓ_2 distance between these two functions with respect to the empirical probability measure corresponding to $\boldsymbol{\theta}_{1j}, \dots, \boldsymbol{\theta}_{nj}$. More specifically, $\|\tilde{f} - \check{f}\|_n^2 = n^{-1} \sum_{i=1}^n (\tilde{f}(\boldsymbol{\theta}_{ij}) - \check{f}(\boldsymbol{\theta}_{ij}))^2$. We refer to the estimated regression function for the j th predictor as $\hat{f}_j(\cdot) = \hat{g}_j(\hat{\eta}_j^T \cdot)$. The corresponding true regression functions are referred to as f_{0j} . We slightly abuse the notation and write $\|\hat{f}_j - f_{0j}\|_n$ for the ℓ_2 distance between \hat{f}_j and f_{0j} with respect to the empirical probability measure corresponding to the j th predictor:

$$\begin{aligned} \|\hat{f}_j - f_{0j}\|_n^2 &= \frac{1}{n} \sum_{i=1}^n [\hat{f}_j(\boldsymbol{\theta}_{ij}) - f_{0j}(X_{ij})]^2 \\ &= \frac{1}{n} \sum_{i=1}^n \left[\hat{g}_j(\hat{\eta}_j^T \boldsymbol{\theta}_{ij}) - g_{0j} \left(\int_0^1 \beta_j(t) X_{ij}(t) dt \right) \right]^2. \end{aligned}$$

As before, we write \mathfrak{M}_0 for the index set of the signal predictors, that is, $\mathfrak{M}_0 = \{j : 1 \leq j \leq p_n, f_{0j} \neq 0\}$. Note that this set depends on n , but we will refrain from using an additional subscript for simplicity of the notation. We use $\hat{\mathfrak{M}}_n$ to denote the corresponding estimated set, $\{j : 1 \leq j \leq p_n, \hat{f}_j \neq 0\}$. Let $s_n = |\mathfrak{M}_0|$. A *universal constant* is interpreted as a constant that does not depend on n or any of the other parameters that appear in the corresponding expression. Given expressions E_1 and E_2 , we use $E_1 \gtrsim E_2$ to mean that there exists a positive universal constant c , such that $E_1 \geq cE_2$. We write $E_1 \asymp E_2$ when both $E_1 \gtrsim E_2$ and $E_2 \gtrsim E_1$ are satisfied.

The results provided below establish the rate of convergence for the estimated regression functions. To derive these results, we impose a number of regularity conditions on the components of the FAR model. We also impose a version of the *compatibility condition*, which is commonly used in high-dimensional additive models [5, 29]. The proofs, as well as a more detailed discussion of the conditions, are provided in Appendix C.

THEOREM 3. *Suppose that Conditions 3 and 4 are satisfied. Let $q_n \gtrsim d_n \gtrsim \log \log n$. Then there exists a universal constant c , such that for $\lambda_n \geq c(n^{-1/2} q_n^{1/2} + n^{-1/2} \sqrt{\log p_n})$, the following bound holds with probability tending to one, as n tends to infinity:*

$$(15) \quad \sum_{j=1}^{p_n} \|\hat{f}_j - f_{0j}\|_n = O(s_n \lambda_n + s_n d_n^{-2} + s_n^2 n^{1/2} d_n^{-4} q_n^{-1/2}).$$

The following corollary focuses on the choice of q_n and d_n that yields the fastest rate of convergence. Note that the case $q_n/d_n = o(1)$ is not covered in the statement of Theorem 3. However, it follows from the proof of the theorem that such settings correspond to an error bound that is inferior to the one presented below.

COROLLARY 3. *Suppose that Conditions 3 and 4 are satisfied. Let $q_n \asymp d_n \asymp (s_n n)^{1/5}$. Then there exists a universal constant c , such that for $\lambda_n \geq c(s_n^{1/10} n^{-2/5} + n^{-1/2} \sqrt{\log p_n})$, the following bound holds with probability tending to one, as n tends to infinity:*

$$\sum_{j=1}^{p_n} \|\hat{f}_j - f_{0j}\|_n = O(s_n \lambda_n).$$

We now turn to the variable selection properties of the nonlinear FAR estimator. Methods that use ℓ_2 regularization are known to typically produce models containing a large number of noise predictors ([5], Chapter 7, e.g.). To alleviate this problem, we follow the popular approach of thresholding the initial estimator. We define the thresholded FAR estimator as follows: $\tilde{f}_j = \hat{f}_j I\{\|\hat{f}_j\|_n > \lambda_n\}$, $j = 1, \dots, p_n$. Note that the threshold parameter is taken equal to the tuning parameter λ_n , which is used to compute the initial estimators, \hat{f}_j . Thus, we do not introduce any new tuning parameters at the thresholding stage. Let $\tilde{\mathfrak{M}}_n$ denote the index set of the corresponding nonzero regression function estimates, $\{j : 1 \leq j \leq p_n, \tilde{f}_j \neq 0\}$. Recall that $s_n = |\mathfrak{M}_0|$. The next result provides bounds for the estimation error of the thresholded FAR approach and for the corresponding number of selected predictors.

THEOREM 4. *Under the assumptions of Corollary 3, there exists a universal constant c , such that for $\lambda_n \geq c(s_n^{1/10} n^{-2/5} + n^{-1/2} \sqrt{\log p_n})$, the following bounds hold with probability tending to one, as n tends to infinity:*

$$|\tilde{\mathfrak{M}}_n| = O(s_n) \quad \text{and}$$

$$\sum_{j=1}^{p_n} \|\tilde{f}_j - f_{0j}\|_n = O(s_n \lambda_n).$$

Now consider the case where the components of the FAR model do not depend on n . More specifically, suppose that the number of signal predictors, $|\mathfrak{M}_0|$, and the signal regression functions, $\{f_{0k}\}_{k \in \mathfrak{M}_0}$, are fixed and do not change with n . The estimation error bound in Theorem 4 implies that, with probability tending to one, our estimator has zero false negatives, while the number of false positives stays bounded. This variable selection result can be strengthened by increasing the threshold from λ_n to $\tau \lambda_n$, for a sufficiently large τ . The next corollary demonstrates that the corresponding thresholded estimator can correctly recover the index set of the relevant predictors.

COROLLARY 4. *Suppose that the components of the FAR model do not depend on n . Suppose also that the assumptions of Corollary 3 are satisfied. Then there*

exist universal constants τ_0 and c , such that, provided $\tau \geq \tau_0$, $\lambda_n \geq c(s_n^{1/10} n^{-2/5} + n^{-1/2} \sqrt{\log p_n})$ and $\lambda_n = o(1)$, we have

$$\widetilde{\mathfrak{M}}_n = \mathfrak{M}_0,$$

with probability tending to one, as n goes to infinity.

4. Simulations. In this section, we compare the performance of FAR to several alternative linear and nonlinear functional approaches in a series of simulation studies. We consider the linear setting in Section 4.1, while Section 4.2 contains our nonlinear results.

4.1. *Linear additive models.* We first generated the functional predictors, $X_{ij}(t)$, from a 4-dimensional Fourier basis $\mathbf{b}(t) = (1, \sqrt{2} \sin(\pi t), \sqrt{2} \sin(2\pi t), \sqrt{2} \sin(3\pi t))^T$, plus an error term:

$$X_{ij}(t_k) = \mathbf{b}(t_k)^T \boldsymbol{\theta}_{ij} + w_{ijk}, \quad w_{ijk} \sim N(0, \sigma_x^2), \quad \boldsymbol{\theta}_{ij} \sim N(0, I),$$

where $\sigma_x = 0.5$, and each predictor was observed at 200 equally spaced time points, $0 = t_1, t_2, \dots, t_{200} = 1$. The basis coefficients, $\boldsymbol{\theta}_{ij}$, and the error terms, w_{ijk} , were all sampled independently from each other. The first s_n coefficient functions, $\beta_1(t), \dots, \beta_{s_n}(t)$, were also generated, from the same basis function, $\beta_j(t) = \mathbf{b}(t)^T \boldsymbol{\eta}_j$, while the remaining $p - s_n$ predictors were noise variables with $\beta_j(t) = 0$. For each $j = 1, \dots, s_n$, the coefficient vector $\boldsymbol{\eta}_j$ were first independently generated from a multivariate standard normal distribution and then rescaled to have ℓ_2 norm equal to 1. The responses were then generated from (2) with $f_j(x)$ computed using (6). We tested a total of six linear settings corresponding to different numbers of observations, predictors and noise levels.

To ensure a fair real world comparison, where the true functional form of $\beta_j(t)$ would be unknown, we implemented the linear version of FAR using an orthogonal cubic spline basis, rather than the true Fourier basis. We tested FAR using both the SCAD [8] and the Lasso penalty functions but found that the former penalty generally gave superior predictive ability so only report the SCAD results here. We compared FAR to three competing methods. The first was a functional principal components analysis (FPCA) based approach produced by decomposing the predictors into functional principal components, selecting the first K components and finally using the resulting PCA scores to fit linear regression models to the response. Since only s_n of the predictor functions were associated with the response, we fit the linear regressions to the FPCA scores using the group SCAD penalty function to produce sparse fits, where the K principal components for each predictor were grouped together.

Our second approach involved implementing the additive modeling method (ADD) of [16]. ADD fits an additive model with the same general form as (2). A key difference relative to FAR is that ADD uses a kernel based fitting method

and a forward selection procedure to iteratively add functional predictors to the model. The final method, SIR, is described in [3]. This method first computes the wavelet coefficients on a single predictor function, then applies the SIR [24] dimension reduction method to the resulting coefficients, and finally a linear regression is fit using the reduced dimensions as the predictors. This approach is not designed for multiple predictor functions so we adapted it by computing the reduced dimensions marginally for each predictor and then performing a multiple linear regression on all the resulting dimensions.

The tuning parameters for the various methods were chosen by minimizing prediction error on a separately generated validation data set with identical characteristics to the training data. FAR had two tuning parameters; λ and the dimension of the orthogonal cubic spline basis for fitting $\beta_j(t)$. We fitted FAR separately for each possible basis dimension, and then selected the value (between 5 and 10) which gave the smallest prediction error on the validation set. The FPCA method had two tuning parameters; λ , the penalty level for the group SCAD fit, and K , the number of principal components used for each predictor. We used the same value of K for all predictors. To select K , we first identified a number K_{\max} such that the first K_{\max} scores of each predictor express at least 99% of the total variation of this predictor, and then selected K as the value (between 1 and K_{\max}) which minimized prediction error on the validation data. The SIR method had one tuning parameter; the number of directions into which each predictor was projected. We considered up to 4 directions for each predictor, and selected the number of directions as the one with the lowest prediction error on the validation set.

For each simulation setting, we fitted each method to 100 different training sets and recorded the false positive rate (FPR), false negative rate (FNR), average prediction error on a separate test data set (Mean PE) and the standard error of the mean PE (SE PE). The FPR records the fraction of noise predictors incorrectly included in the model while the FNR corresponds to the fraction of signal variables incorrectly excluded. The simulation results are summarized in Table 1. Prediction errors that were either the best or were not statistically worse than the best result are shown in bold font. Note that because of the extremely computationally intensive nature of the ADD and SIR methods it was not feasible to compute fits for p larger than about 10. In fact, in the $p = 600$ and 2000 settings the FPCA, ADD and SIR comparison methods were all too slow to implement, and thus we only report the results for FAR. In terms of prediction error, FAR was superior to all of the competing methods in most simulation settings. The FPCA method was the best competitor followed by SIR and finally ADD. The only setting where FPCA was superior was the situation where $\sigma_y = 2$ and $p = 100$, which had high noise and high dimensionality. For the ultra-high dimensional setting of $p = 2000$, FAR still does a reasonably good job in variable selection. Note that when fitting FAR, since each functional predictor is approximated using a spline basis, the dimensionality in the linear FAR criterion is in fact much higher than p . For example, if a 5-dimensional spline basis is used, the dimensionality is in fact $5p = 10,000$.

TABLE 1
Comparison of FAR to three alternative methods in five linear simulation settings

		FAR	FPCA	ADD	SIR
$n = 60$	FN	0.0000	0.0000	NA	NA
$p = 10$	FP	0.0250	0.1067	NA	NA
$s_n = 4$	Mean PE	1.4834	1.6558	2.6474	2.3318
$\sigma_y = 1$	SE PE	0.0285	0.0274	0.0298	0.0275
$n = 60$	FN	0.0225	0.005	NA	NA
$p = 10$	FP	0.05	0.1633	NA	NA
$s_n = 4$	Mean PE	2.6805	2.7979	3.3462	6.2968
$\sigma_y = 2$	SE PE	0.0267	0.0264	0.0296	0.0857
$n = 80$	FN	0.0067	0.1917		
$p = 100$	FP	0.0743	0.0454		
$s_n = 6$	Mean PE	2.0176	3.5502		
$\sigma_y = 1$	SE PE	0.0548	0.0353		
$n = 80$	FN	0.0483	0.0067		
$p = 100$	FP	0.1896	0.1569		
$s_n = 6$	Mean PE	3.7051	3.3250		
$\sigma_y = 2$	SE PE	0.0548	0.0353		
$n = 100$	FN	0.0700			
$p = 600$	FP	0.0432			
$s_n = 8$	Mean PE	3.6423			
$\sigma_y = 1$	SE PE	0.0910			
$n = 100$	FN	0.1925			
$p = 2000$	FP	0.0171			
$s_n = 8$	Mean PE	4.6422			
$\sigma_y = 1$	SE PE	0.0871			

4.2. *Nonlinear models.* We examined three different simulation settings with the responses generated from the nonlinear model (4). The standard deviation, σ_x , the predictors, $X_{ij}(t)$, and coefficient curves, $\beta_j(t)$, were all produced in an identical fashion to the linear setting. To produce a sparse relationship between the predictors and the response, we set $g_j(x) = 0$ for $j = 3, 4, \dots, p$. The remaining two curves were chosen as $g_1(x_1) = x_1$ and $g_2(x_2) = -x_2 + \sin(x_2)$. Note that these functions were not generated from a B-spline basis so the FAR fit contains bias in the estimates for both $\beta_j(t)$ and $g_j(x)$; a real world situation where the data is unlikely to exactly correspond to the FAR model. The sample size was fixed at $n = 100$, and the model errors were independently generated from a Gaussian distribution with mean zero and standard deviation $\sigma = 0.5$.

We compared the nonlinear version of FAR to the same three competing methods as in the linear setting. However, to account for the nonlinear relationships between the response and predictors, we implemented FPCA by applying the SpAM

method [33] to the principal component scores. SpAM essentially fits a penalized version of Generalized Additive Models (GAM), allowing for automatic variable selection in a nonlinear but additive regression situation. We adapted SpAM slightly to implement a group penalization where all K PCs for a given predictor were penalized together. The SIR method was still implemented using the linear regression approach from the previous section while the kernel approach of ADD already produced a nonlinear fit so these last two methods did not require any adaptations to the new setting. In each simulation, we again fit the methods to 100 separate data sets and used a separate validation data set, with identical characteristics to the training data, to select the tuning parameters. The nonlinear setting increased by one the number of tuning parameters for the FAR and FPCA methods; d , the basis dimension for $g_j(x)$. For both methods, we chose d by computing the validation error rates for values between 5 and 10, selecting the optimal value and then using this dimension to compute $g_j(x)$. To reduce the computational cost for FAR, we selected q , the dimension of the spline basis for $\beta_j(t)$, as the value (between 5 and 10) which gave the best hold out accuracy on the predictors in the validation set. In particular, we held out 20% of each predictor’s time points, computed the least squares fit to the remaining time points for each possible basis dimension, and then selected the value of q which gave the lowest error rate on the held-out points.

The simulation results are summarized in Table 2, with bold font indicating the statistically best prediction errors. As with the linear setting it was not computationally feasible to implement ADD or SIR for dimensionality p larger than the sample size n . In the low-dimensional setting of $p = 5$, SIR produced the lowest mean prediction error with FAR the second best. For the higher-dimensional setting of $p = 50$, the mean prediction error of SIR increased dramatically and was

TABLE 2
Comparison of FAR to three alternative methods in three nonlinear simulation settings

		FAR	FPCA	ADD	SIR
$n = 100$	FN	0.0000	0.0000		
$p = 5$	FP	0.1833	0.1300		
$\sigma_y = 0.5$	Mean PE	0.9792	1.3108	1.7408	0.8688
	SE PE	0.0132	0.0174	0.0074	0.0049
$n = 100$	FN	0.0000	0.0000		
$p = 50$	FP	0.0171	0.1138		
$\sigma_y = 0.5$	PE	1.1068	1.3907	1.8965	3.5062
	SE PE	0.0164	0.0156	0.0110	0.0309
$n = 100$	FN	0.0000	0.0000		
$p = 120$	FP	0.0064	0.0697		
$\sigma_y = 0.5$	Mean PE	1.2108	1.5164		
	SE PE	0.0157	0.0159		

the largest among all competitors. In the last two settings, FAR was significantly superior to all three competing methods, with FPCA generally providing the next best results. However, we remark that the FPCA method is significantly slower than FAR in these nonlinear settings due to the extra tuning parameter.

5. Real data.

5.1. *Hollywood stock exchange data.* The goal for this analysis was to compare the accuracy of FAR and FPCA in predicting the total box office revenue (over the first ten weeks after release) for 262 movies. We use pre-release trading histories from the *Hollywood Stock Exchange (HSX)*, one of the best known online virtual stock markets, as our functional predictors. The Hollywood stock exchange has nearly 2 million active participants worldwide. Each trader is initially endowed with \$2 million virtual currency and can increase his or her net worth by strategically selecting and trading movie stocks (i.e., buying low and selling high). Figure 1 shows the HSX trading histories, between 52 and 10 weeks prior to a movie's release, for a sample of 15 out of the 262 movies in our data set. Each curve represents the traders' collective daily average predictions of the box

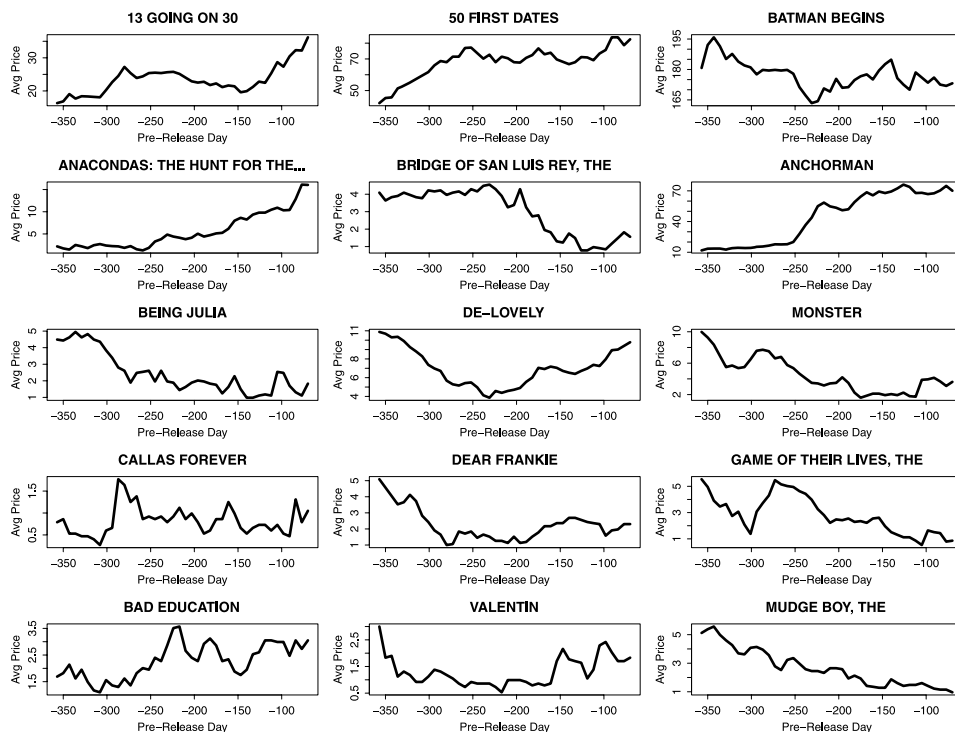


FIG. 1. Trading histories for a sample of movies from the HSX data set.

office revenue that the movie will generate after it is released. In addition to the *Daily Average* curves, we also observed four additional predictors for each movie: *Accounts Trading*; *Accounts Trading Short*; *Shares Held Short*; *Shares Traded Sell*.

We only consider HSX curves from 10 weeks prior to release date because the goal is to form accurate revenue predictions early enough to affect strategic decisions, such as, advertising budget, locations of theater release, etc. We implemented the nonlinear versions of both FAR and FPCA on the log revenues as this appeared to give superior results for both methods. For FAR, we needed to select 3 tuning parameters, λ , q and d , and for FPCA we also had 3 tuning parameter, λ , K and d . Hence, we randomly divided the 262 movies into three approximately equal partitions. The methods were trained on the first group over grids of the tuning parameters, the second group was used to select the final tuning parameters and out of sample error rates were computed on the final group.

The mean hold out (log) prediction error, averaged over 50 random partitions, was 2.45 for FAR, while the FPCA error rate was higher at 2.66. The standard error in the difference between the FAR and FPCA methods over the 50 random partitions was 0.10. Both FAR and FPCA chose *Daily Average* in all 50 partitions, with the average model size of FAR being 1.92 and the average model size of FPCA being 1.86. The mean hold out (log) prediction error on the test movies using the null model is 4.75, indicating that using these functional predictors from the trading histories indeed improves the prediction results.

Figure 2 plots the 50 estimated $\beta(t)$ and $g(x)$ functions corresponding to the *Daily Average* variable with the solid red lines representing the average effect. Most of the curves show remarkably consistent patterns; $g(x)$ is estimated as a strictly increasing, but nonlinear function, and $\beta(t)$ places approximately zero weight on the earlier trading history and a larger positive weight on roughly the final month under consideration. These curves conform to our intuition that the trading history closest to release date provides the strongest prediction accuracy and that there is a positive correlation between HSX curves and movie revenues.




FIG. 2. The $\beta(t)$ and $g(x)$ curves corresponding to the *Daily Average* variable in the HSX data.

The nonlinear shape of $g(x)$ also suggests that a linear model would not provide accurate results for this data.

5.2. MEG data. Our second data set consisted of Magnetoencephalography (MEG) recordings for 20 subjects conducted at the Center for Clinical Neurosciences, University of Texas Health Science Center at Houston. The MEG readings for each subject were recorded over 248 “channels” at 356 equally spaced time points. Each channel measured the intensity level of the magnetic field at a particular point on the brain. Multiple trials, consisting of reading a patient a word and measuring the MEG over time, were recorded for each patient. We averaged the trials for each patient to produce 248 functional predictors, one for each channel. The response of interest was whether the patient was left (14 subjects) or right (6 subjects) brain dominated. We coded $Y = 1$ and $Y = -1$, respectively, for left- and right-brained subjects. Some channels were missing for some patients and were removed from the study, leaving a total of $p = 199$ predictors.

This was a very challenging data set because the ratio of predictors to observations was 10:1. We first fit the linear version of FAR to the full data set using a five-dimensional basis for $\beta_j(t)$. The tuning parameter, λ , was chosen as the point which minimized the classification error using 20-fold cross-validation. In this setting, FAR selected only a five variable model (Channels 3, 138, 139, 167 and 220), which corresponded to a 20% cross-validated error rate. Figure 3 displays the $\beta(t)$ curves for each selected channel. All five channels put the bulk of their weight on the early time points. Channel 3 appears to provide the majority of the predictive power with smaller contributions from Channels 138 and 167. In particular $\beta_3(t)$ represents a contrast between early and late time points. Hence, people who start low in Channel 3 and end high are predicted to be left-brained while the opposite is true for right-brained patients.

We also fit the nonlinear version of FAR. Given the small number of observations and the extra demands of fitting a nonlinear regression method we felt it




FIG. 3. Plots of $\beta(t)$ for linear FAR on the MEG data.

was prudent to first perform a marginal pre-screening to select a smaller subset of predictors for the final analysis. The marginal screening was performed by running nonlinear FAR, using a 7-dimensional basis function, separately on each of the 194 predictors that linear FAR did not choose and selecting the 45 best predictors in terms of marginal prediction accuracy. Nonlinear FAR was then run on the 50 predictors, including the 5 selected by linear FAR. 20-fold cross validation was again used to select the tuning parameter, resulting in five channels being selected. The channels were not the same as those selected by linear FAR. The cross-validated error rate was 25%, suggesting that linear FAR may have a slight advantage on this data.

6. Discussion. FAR extends the recent linear penalized regression literature by incorporating functional predictors and modeling general nonlinear relationships. It has several advantages over current functional regression methods. First, the penalized approach automatically deals with high-dimensional data using an efficient coordinate descent algorithm. Second, the single index formulation provides a nonlinear supervised method for projecting the predictors into a lower-dimensional space, providing more accurate results than the traditional linear unsupervised PCA approach. Third, our theoretical results suggest that FAR should provide accurate variable selection and prediction results and the simulation results show that FAR outperforms traditional approaches.

There are three obvious possible extensions for FAR. The first is to incorporate FAR into the generalized linear models setting. Conceptually, such an extension could be achieved by replacing the sum of squares term in (5) with the log likelihood and then using a modified version of the coordinate descent algorithm to maximize the criterion. The second possible extension would be to replace the single index model with a multiple index model of the form, $f_j(X_{ij}) = \sum_{k=1}^K g_{jk}(\int \beta_{jk}(t) X_{ij}(t) dt)$. This would increase the flexibility of FAR to model more general nonlinear relationships. Finally, FAR could be extended to model functional responses in addition to functional predictors.

APPENDIX A: DETAILS OF THE NONLINEAR FAR ALGORITHM

In the initialization step (step 0) of this algorithm, some of the η_j 's will likely be set to zero. This suggests that the corresponding predictors do not appear related to the response. However, the initialization assumes a linear model. It is conceivable that a response that appears unimportant using a linear model will become statistically significant using a nonlinear model. Hence, if η_j is estimated to be zero in step 0 we instead set η_j equal to the loading vector of the first principal component of Θ_j . This estimate is the direction that explains the most variability in $X_{ij}(t)$ so is the most natural unsupervised projection and allows for potential nonlinear relationships to be detected in step 2.

To implement step 3 of the FAR algorithm, we minimize (14) with respect to the η_j 's. Directly minimizing (14) is difficult due to the nonlinearity of the functions $g_j(t) \approx \mathbf{h}(t)^T \hat{\xi}_j$. To overcome this difficulty, we observe that, with the estimate $\hat{\xi}_j$ from step 2 and the current value $\eta_{j,\text{old}}$ of η_j , the first-order approximation of $g(\theta_{ij}^T \eta_j) \approx \mathbf{h}(\theta_{ij}^T \eta_j)^T \hat{\xi}_j$ is

$$(16) \quad \mathbf{h}(\theta_{ij}^T \eta_j)^T \hat{\xi}_j \approx \mathbf{h}(\theta_{ij}^T \eta_{j,\text{old}})^T \hat{\xi}_j + \mathbf{h}'(\theta_{ij}^T \eta_{j,\text{old}})^T \hat{\xi}_j \cdot \theta_{ij}^T (\eta_j - \eta_{j,\text{old}}).$$

Thus, we can approximate (14) as

$$(17) \quad \sum_{i=1}^n \left(R_i - \sum_{j=1}^p \mathbf{h}'(\theta_{ij}^T \eta_{j,\text{old}})^T \hat{\xi}_j \cdot \theta_{ij}^T (\eta_j - \eta_{j,\text{old}}) \right)^2,$$

where $R_i = Y_i - \sum_{j=1}^p \mathbf{h}(\theta_{ij}^T \eta_{j,\text{old}})^T \hat{\xi}_j$, that is, the residual for the i th observation from step 2 of the algorithm in the current iteration. The above approximation (17) is a quadratic function of η_j and can be minimized easily. Hence, the new value of η_j is updated as the minimizer of (17). We also note that if the estimate $\hat{\xi}_j$ from step 2 is $\mathbf{0}$, then the corresponding value of η_j will not be updated.

APPENDIX B: TECHNICAL CONDITIONS OF THEOREMS 1–2

We make the following assumption on the functional predictors $X_{ij}(t)$ and the corresponding regression coefficients $\beta_j(t)$.

CONDITION 1. (A) Functional predictors, $\{X_{ij} : [0, 1] \rightarrow \mathbb{R}, i = 1, \dots, n, j = 1, \dots, p_n\}$, belong to a Sobolev ellipsoid of order two: there exists a universal constant C , such that $\sum_{k=1}^\infty \theta_{ijk}^2 k^4 \leq C^2$ for all $i = 1, \dots, n, j = 1, \dots, p_n$.

(B) The true coefficient functions satisfy $\max_{j \in \mathfrak{M}_0} \int_0^1 \beta_j^2(t) dt \leq \tilde{C}$ with \tilde{C} some positive constant.

Note that the linear FAR model can be written as

$$(18) \quad Y_i = \sum_{j=1}^p \Theta_j \eta_j + \varepsilon_i^*,$$

where $\varepsilon_i^* = \varepsilon_i + \sum_{j=1}^p e_{ij}$ with e_{ij} defined in (10). When $j \in \mathfrak{M}_0^c$, $\beta_j(t) = 0$ and thus the approximation error e_{ij} in (9) disappears. Thus, in view of (11), the approximation error satisfies that

$$\left| \sum_{j=1}^p e_{ij} \right| \leq \sum_{j \in \mathfrak{M}_0} |e_{ij}| \leq C s_n q_n^{-2},$$

uniformly over all $i = 1, \dots, n$.

Our second set of conditions concern the shape of the penalty function, the strength of the signal and the correlation structure of the predictors.

CONDITION 2. (A) For any fixed $\lambda > 0$, $\rho_\lambda(t)$ is concave and nondecreasing in $[0, \infty)$, and has nonincreasing first derivative $\rho'_\lambda(t)$. Further, $\rho'_\lambda(0+) > 0$.

(B) Let $a_n = \min_{j \in \mathfrak{M}_0} \|\Theta_j \boldsymbol{\eta}_{0,j}\|/\sqrt{n}$. It holds that $n^\alpha a_n \rightarrow \infty$ with $\alpha \in (0, \frac{1}{2})$.

(C) It holds that $\rho'_{\lambda_n}(a_n/2) = o(n^{-\alpha} q_n^{-1} s_n^{-1/2})$ and $\sup_{t \geq a_n/2} \rho''_{\lambda_n}(t) = o(1)$.

(D) There exists a positive constant c_0 such that

$$(19) \quad c_0 \leq \min_{j \in \mathfrak{M}_0} \Lambda_{\min} \left(\frac{1}{n} \Theta_j^T \Theta_j \right) < \Lambda_{\max} \left(\frac{1}{n} \Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0} \right) \leq c_0^{-1},$$

where Λ_{\min} and Λ_{\max} are the smallest and largest eigenvalues of a matrix, respectively. Further, we have

$$(20) \quad \max_{j \in \mathfrak{M}_0} \|\Theta_j (\Theta_j^T \Theta_j)^{-1} \Theta_j^T \Theta_{\mathfrak{M}_0} (\Theta_{\mathfrak{M}_0}^T \Theta_{\mathfrak{M}_0})^{-1}\|_{\infty,2} < \frac{\sqrt{c_0} \rho'_{\lambda_n}(0+)}{2\sqrt{n} \rho'_{\lambda_n}(a_n/2)},$$

where for a matrix B , $\|B\|_{\infty,2} = \sup_{\|\mathbf{x}\|_\infty=1} \|B\mathbf{x}\|_2$ with \mathbf{x} a vector.

(E) The model errors $\varepsilon_i, i = 1, \dots, n$ are independent and identically distributed as $N(0, \sigma^2)$.

Condition 2(A) requires that the penalty functional, $\rho_\lambda(t)$, is concave and singular at 0. Many penalty functions proposed in the literature such as the hard thresholding penalty, SCAD [8] and SICA [27] all satisfy this condition. From (9), we see that Condition 2(B) places a lower bound on the signal strength of the true predictors $j \in \mathfrak{M}_0$. In particular, it assumes that the weakest signal, a_n , can decay with sample size but the decay rate cannot be faster than $n^{-\alpha}$. Condition 2(C) is a mild condition which can be easily satisfied by penalty functions with flat tails. For instance, if $\lambda_n = o(a_n/2)$, then for SCAD penalty, it can be verified from the definition that $\rho'_{\lambda_n}(a_n/2) = 0$ and $\rho''_{\lambda_n}(t) = 0$ for all $t \geq a_n/2$, and thus Condition 2(C) is satisfied. Although Condition 2(C) assumes the existence of the second-order derivative for $\rho_{\lambda_n}(t)$, it can be relaxed to the existence of the first-order derivative by using the local concavity definition in [27]. Condition 2(D) relates to the design matrix for the signal predictors, $\Theta_{\mathfrak{M}_0}$. We assume that the eigenvalues for the design matrix corresponding to true predictors are bounded from below and above. If $\Theta_{\mathfrak{M}_0}$ is orthogonal, then (19) is satisfied with $c_0 = 1$. The upper bound in condition (20) depends on the penalty function through the ratio $\rho'_{\lambda_n}(0+)/\rho'_{\lambda_n}(a_n/2)$, which is larger than 1 for concave penalties and equal to 1 for the group Lasso penalty, $\rho_{\lambda_n}(t) = \lambda_n t$. For instance, if $\lambda_n = o(a_n)$, then $\rho'_{\lambda_n}(0+)/\rho'_{\lambda_n}(a_n/2) = \infty$ for SCAD penalty and thus (20) is satisfied automatically. The detailed proofs of Theorems 1 and 2 are in the supplementary materials [12].

APPENDIX C: TECHNICAL CONDITIONS AND PROOF OF THEOREMS 3–4

C.1. Conditions. Given an orthonormal basis expansion for $\beta_j(t)$, that is, $\beta_j(t) = \sum_{l=1}^\infty \eta_{jl}^* b_l(t)$, we will define $\boldsymbol{\eta}_j^* = (\eta_{j1}^*, \dots, \eta_{jq_n}^*)^T$. We will also define

$f_j^*(\theta_{ij}) = \mathbf{h}_{\eta_j^*, j}(\theta_{ij}^T \eta_j^*) \xi_j^*$, where ξ_j^* is chosen to minimize $\sum_{i=1}^n [\mathbf{h}_{\eta_j^*, j}(\theta_{ij}^T \eta_j^*) \xi_j^* - g_{0j}(\theta_{ij}^T \eta_j^*)]^2$ over $\xi \in \mathbb{R}^{d_n}$ with the constraint $\sum_i f_j^*(\theta_{ij}) = 0$. Note that f_j^* , ξ_j^* , η_j^* and θ_{ij} depend on n , but we omit the corresponding subscripts for the simplicity of the notation. The following are the technical conditions for the theory in Section 3.2. A discussion of the conditions is given below.

CONDITION 3. (A) Functional predictors, $\{X_{ij} : [0, 1] \rightarrow \mathbb{R}, i = 1, \dots, n, j = 1, \dots, p_n\}$, belong to a Sobolev ellipsoid of order two: there exists a universal constant C , such that $\sum_{k=1}^\infty \theta_{ijk}^2 k^4 \leq C^2$ for all $i = 1, \dots, n, j = 1, \dots, p_n$.

(B) The true index functions, $\{\beta_j(t), j \in \mathfrak{M}_0\}$, satisfy $\int_0^1 \beta_j^2(t) dt = 1$.

(C) Errors ε_i are independent and uniformly sub-Gaussian.

(D) The true link functions, g_{0j} , are twice continuously differentiable and are bounded, together with their first and second derivatives, uniformly over $j \in \mathfrak{M}_0$ and n .

(E) For each η with $\|\eta\| = 1$ and each $j \leq p_n$ let $Q_{\eta, j, n}$ denote the empirical distribution associated with the index values $\eta^T \theta_{1j}, \dots, \eta^T \theta_{nj}$. Assume that there exist corresponding probability distributions $P_{\eta, j, n}$, each with bounded support and a positive continuous density, such that the densities are bounded both above and away from zero uniformly over j and n , and

$$(21) \quad \sup_{u \in \mathbb{R}, \|\eta\|=1, 1 \leq j \leq p_n} |Q_{\eta, j, n}(-\infty, u] - P_{\eta, j, n}(-\infty, u]| = o(d_n^{-1}).$$

Condition 3(A) is identical to Condition 1(A), imposed for the linear FAR theory. It is a common smoothness requirement in nonparametric regression, when the orthogonal basis approach is used, as discussed, for example, in Chapter 8 in [37]. Condition 3(B) is imposed for identifiability. Conditions 3(C) and (D) are typical in high-dimensional regression and nonparametric regression problems, respectively. The reason we require uniformity is to handle the situation where the number of signal predictors grows with n . Again, uniformity is needed to handle the growing number of signal predictors. Condition 3(E) ensures that the candidate index values, $\eta^T \theta_{1j}, \dots, \eta^T \theta_{nj}$, have sufficiently regular distributions. Assumptions of this form are typical in spline estimation [39], for example.

We impose two more assumptions below. Condition 4(A) is a natural generalization of the *compatibility condition* used in high-dimensional additive models, for example, in [29] and Section 8.4. in [5]. Note that because we do not use a smoothness penalty in our estimation approach, the smoothness penalty does not appear in the compatibility condition. Condition 4(B) is a version of the standard regularity condition on the behavior of the sum of squares function near its minimum. Assumptions of this form have been imposed in the single index model literature, for example, [38]. We again require uniformity over $j \in \mathfrak{M}_0$ to handle the growing number of signal predictors.

CONDITION 4. (A) There exists a positive universal constant ϕ^2 for which the following holds. If functions $\{h_j, j = 1, \dots, p_n\}$ are such that each h_j is a difference of two functions in \mathcal{F}_j^0 , and inequality $\sum_{j \in \mathfrak{M}_0^c} \|h_j\|_n \leq 3 \sum_{j \in \mathfrak{M}_0} \|h_j\|_n$ is satisfied, then the following inequality holds: $\sum_{j \in \mathfrak{M}_0} \|h_j\|_n^2 \leq \|\sum_{j=1}^{p_n} h_j\|_n^2 / \phi^2$.

(B) There exist positive universal constants τ, c_1 and c_2 , such that for all sufficiently large n and each $f_{\eta_j, j}(\cdot) = \mathbf{h}_{\eta_j, j}(\boldsymbol{\eta}_j^T \cdot) \boldsymbol{\xi}$ with $\|\boldsymbol{\eta}_j\| = \|\boldsymbol{\eta}_j^*\|$ and $j \in \mathfrak{M}_0$, inequalities $\|f_{\eta_j, j} - f_j^*\|_n \leq \tau$ and $\|\boldsymbol{\eta}_j - \boldsymbol{\eta}_j^*\| < \|\boldsymbol{\eta}_j + \boldsymbol{\eta}_j^*\|$ imply $\|\boldsymbol{\eta}_j - \boldsymbol{\eta}_j^*\| \leq c_1 \|f_{\eta_j, j} - f_j^*\|_n$ and $\|f_{\boldsymbol{\eta}_j^*, j} - f_j^*\|_n \leq c_2 \|f_{\eta_j, j} - f_j^*\|_n$.

C.2. Preliminaries. We start by deriving a bound on the error due to our approximation of index functions β_j and link functions g_{0j} . Observe that

$$\begin{aligned} \left| \int_0^1 \beta_j(t) X_{ij}(t) dt - \boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j^* \right|^2 &= \left| \sum_{k=q_n+1}^{\infty} \eta_{jk}^* k^{-2} k^2 \theta_{ijk} \right|^2 \\ &\leq \sum_{k=q_n+1}^{\infty} (\eta_{jk}^*)^2 k^{-4} \sum_{k=q_n+1}^{\infty} \theta_{ijk}^2 k^{-4}. \end{aligned}$$

Condition 3(A) implies that the right-most sum is bounded by a universal constant C^2 . Also note that $\sum_{l=q_n+1}^{\infty} (\eta_{jl}^*)^2 k^{-4} \leq q_n^{-4} \sum_{k=q_n+1}^{\infty} (\eta_{jk}^*)^2 \leq q_n^{-4}$, by Condition 3(B). Thus, if we set $I_{ij} = \int_0^1 \beta_j(t) X_{ij}(t) dt$, then the bound $|I_{ij} - \boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j^*| \leq C q_n^{-2}$ holds for all n, i , and $j \in \mathfrak{M}_0$. Hence, if we let \tilde{C} be the uniform bound over the first derivatives in Condition 3(D), then

$$(22) \quad |g_{0j}(I_{ij}) - g_{0j}(\boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j^*)| \leq \tilde{C} |I_{ij} - \boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j^*| = O(q_n^{-2}),$$

uniformly over i and $j \in \mathfrak{M}_0$. Set $M_j = \sup_t |g_{0j}''(t)|$ for $j \in \mathfrak{M}_0$, and note that constants M_j are uniformly bounded by Condition 3(D). Taking advantage of the approximation bounds for the cubic B-splines (e.g., Corollary 6.21 in [34]), we then have

$$(23) \quad n^{-1} \sum_{i=1}^n (g_{0j}(\boldsymbol{\theta}_{ij}^T \boldsymbol{\eta}_j^*) - f_j^*(\boldsymbol{\theta}_{ij}))^2 = O(d_n^{-4} M_j^2) = O(d_n^{-4}),$$

uniformly over $j \in \mathfrak{M}_0$. Combining inequalities (22) and (23), we deduce $\|f_{0j} - f_j^*\|_n = O(q_n^{-2} + d_n^{-2}) = O(d_n^{-2})$, uniformly over j . Note that for $j \in \mathfrak{M}_0^c$, both f_{0j} and f_j^* are zero. Consequently,

$$(24) \quad \sum_{j=1}^{p_n} \|f_j^* - f_{0j}\|_n = O(s_n d_n^{-2}).$$

This gives us a useful bound on the approximation error.

We will write $f_0(\mathbf{X}_i)$ for $\sum_{j=1}^{p_n} f_{0j}(X_{ij})$; we also write $\hat{f}(\boldsymbol{\theta}_i)$ for $\sum_{j=1}^{p_n} \hat{f}_j(\boldsymbol{\theta}_{ij})$ and define f^* by analogy. To be consistent with the standard least-squares estimation notation, we will write $(\varepsilon, f)_n$ for $n^{-1} \sum_{i=1}^n \varepsilon_i f(\boldsymbol{\theta}_i)$. We will need the following result, which is proved in the supplementary material [12].

LEMMA 1. Define $r_n = n^{-1/2} q_n^{1/2} + n^{-1/2} \sqrt{\log p_n}$. There exists a positive universal constant C_1 , such that

$$(25) \quad (\varepsilon, \hat{f} - f^*)_n \leq C_1 s_n r_n^2 + C_1 r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n,$$

with probability tending to one.

C.3. Main body of the proof. Let $\|y - f\|_n^2$ denote $n^{-1} \sum_{i=1}^n (Y_i - f(\boldsymbol{\theta}_i))^2$ and let $\|f\|_n^2$ denote $n^{-1} \sum_{i=1}^n f(\boldsymbol{\theta}_i)^2$. Consider the following simple identity:

$$(26) \quad \|y - \hat{f}\|_n^2 - \|y - f^*\|_n^2 = \|\hat{f} - f_0\|_n^2 - \|f^* - f_0\|_n^2 - 2(\varepsilon, \hat{f} - f^*)_n.$$

Note that $\|y - \hat{f}\|_n^2 + \lambda_n \sum_{j=1}^{p_n} \|\hat{f}_j\|_n - \|y - f^*\|_n^2 - \lambda_n \sum_{j=1}^{p_n} \|f_j^*\|_n \leq 0$ by the definition of \hat{f} . Let e_n denote the approximation error, $\|f^* - f_0\|_n$. Inequality (26) then implies

$$\|\hat{f} - f_0\|_n^2 + \lambda_n \sum_{j=1}^{p_n} \|\hat{f}_j\|_n \leq e_n^2 + 2(\varepsilon, \hat{f} - f^*)_n + \lambda_n \sum_{j=1}^{p_n} \|f_j^*\|_n.$$

By Lemma 1, the above inequality yields

$$(27) \quad \begin{aligned} & \|\hat{f} - f_0\|_n^2 + \lambda_n \sum_{j=1}^{p_n} \|\hat{f}_j\|_n \\ & \leq e_n^2 + 2C_1 s_n r_n^2 + 2C_1 r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n + \lambda_n \sum_{j=1}^{p_n} \|f_j^*\|_n, \end{aligned}$$

with probability tending to one.

Case (i). Consider the event $e_n^2 + C_1 s_n r_n^2 \geq r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n$.

Note that $e_n^2 = O(s_n^2 d_n^{-4})$ by (24). Thus, $\sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n = O(s_n^2 n^{1/2} d_n^{-4} \times q_n^{-1/2} + s_n r_n)$. Consequently, $\sum_{j=1}^{p_n} \|\hat{f}_j - f_{0j}\|_n = O(d_n^{-4} q_n^{-1/2} + s_n r_n + s_n d_n^{-2})$, which implies the stochastic bound in display (15).

Case (ii). Consider the event $e_n^2 + C_1 s_n r_n^2 < r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n$.

Using inequality $\|\hat{f} - f^*\|_n^2 \leq 2\|\hat{f} - f_0\|_n^2 + 2e_n^2$ together with (27), we get

$$\begin{aligned} & \|\hat{f} - f^*\|_n^2 + 2\lambda_n \sum_{j=1}^{p_n} \|\hat{f}_j\|_n \\ & \leq 4e_n^2 + 4C_1s_nr_n^2 + 4C_1r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n + 2\lambda_n \sum_{j=1}^{p_n} \|f_j^*\|_n. \end{aligned}$$

On the event $e_n^2 + C_1s_nr_n^2 < r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n$ the above inequality simplifies to

$$\begin{aligned} & \|\hat{f} - f^*\|_n^2 + 2\lambda_n \sum_{j=1}^{p_n} \|\hat{f}_j\|_n \\ (28) \quad & \leq 4(C_1 + 1)r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n + 2\lambda_n \sum_{j=1}^{p_n} \|f_j^*\|_n. \end{aligned}$$

Because we assume $r_n = O(\lambda_n)$, we can rewrite inequality (28) as

$$(29) \quad \|\hat{f} - f^*\|_n^2 = \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n O(\lambda_n).$$

Inequality (28) also gives

$$\sum_{j \in \mathfrak{M}_0^c} \|\hat{f}_j\|_n \leq 2\lambda_n^{-1}(C_1 + 1)r_n \sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n + \sum_{j \in \mathfrak{M}_0} \|\hat{f}_j - f_j^*\|_n.$$

Consequently,

$$\begin{aligned} \sum_{j \in \mathfrak{M}_0^c} \|\hat{f}_j - f_j^*\|_n & \leq 2\lambda_n^{-1}(C_1 + 1)r_n \sum_{j \in \mathfrak{M}_0^c} \|\hat{f}_j - f_j^*\|_n \\ & \quad + [2\lambda_n^{-1}(C_1 + 1)r_n + 1] \sum_{j \in \mathfrak{M}_0} \|\hat{f}_j - f_j^*\|_n, \end{aligned}$$

which, provided $\lambda_n \geq 4(C_1 + 1)r_n$, implies

$$(30) \quad \sum_{j \in \mathfrak{M}_0^c} \|\hat{f}_j - f_j^*\|_n \leq 3 \sum_{j \in \mathfrak{M}_0} \|\hat{f}_j - f_j^*\|_n.$$

This allows us to apply the compatibility condition, 4(A), to $\hat{f} - f^*$. It follows that $s_n^{-1}(\sum_{j \in \mathfrak{M}_0} \|\hat{f}_j - f_j^*\|_n)^2 \leq \|\hat{f} - f^*\|_n^2/\phi^2$, which, by (30), yields $s_n^{-1}(\sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n)^2 \leq 16\|\hat{f} - f^*\|_n^2/\phi^2$. Stochastic bound (29) then gives $\sum_{j=1}^{p_n} \|\hat{f}_j - f_j^*\|_n = O(s_n\lambda_n)$, and hence $\sum_{j=1}^{p_n} \|\hat{f}_j - f_0\|_n = O(s_n\lambda_n + s_nd_n^{-2})$, which again implies the bound in display (15). This completes the proof of Theorem 3.

Under the assumptions of Theorem 4, the error bound in the statement of Theorem 3 simplifies to $\sum_{j=1}^{p_n} \|\hat{f}_j - f_{0j}\|_n = O(s_n \lambda_n)$. Consequently, on the sets of probability tending to one,

$$(31) \quad \sum_{j \in \mathfrak{M}_0^c} \|\tilde{f}_j\|_n \leq \sum_{j \in \mathfrak{M}_0^c} \|\hat{f}_j - f_{0j}\|_n = O(s_n \lambda_n).$$

Using bound (31) and the fact that $\|\tilde{f}_j\|_n > \lambda_n$ for $j \in \tilde{\mathfrak{M}}_n$, we can deduce $|\mathfrak{M}_0^c \cap \tilde{\mathfrak{M}}_n| = O(s_n)$. This implies $|\tilde{\mathfrak{M}}_n| \leq |\mathfrak{M}_0| + |\mathfrak{M}_0^c \cap \tilde{\mathfrak{M}}_n| = O(s_n)$. Also note that

$$\sum_{j \in \mathfrak{M}_0} \|\tilde{f}_j - f_{0j}\|_n \leq \sum_{j \in \mathfrak{M}_0} (\lambda_n + \|\hat{f}_j - f_{0j}\|_n) = O(s_n \lambda_n).$$

The above bound, together with (31), yields the error bound in Theorem 4.

Acknowledgments. We would like to thank the Center for Clinical Neurosciences, University of Texas Health Science Center at Houston for the use of their MEG data.

SUPPLEMENTARY MATERIAL

Supplementary material for: Functional additive regression (DOI: [10.1214/15-AOS1346SUPP](https://doi.org/10.1214/15-AOS1346SUPP); .pdf). Due to space constraints, the proofs of Theorems 1 and 2 and Lemma 1 are relegated to the supplement [12].

REFERENCES

- [1] AIT-SAÏDI, A., FERRATY, F., KASSA, R. and VIEU, P. (2008). Cross-validated estimations in the single-functional index model. *Statistics* **42** 475–494. [MR2465129](#)
- [2] ALTER, O., BROWN, P. O. and BOTSTEIN, D. (2000). Singular value decomposition for genome-wide expression data processing and modeling. *Proc. Natl. Acad. Sci. USA* **97** 10101–10106.
- [3] AMATO, U., ANTONIADIS, A. and DE FEIS, I. (2006). Dimension reduction in functional regression with applications. *Comput. Statist. Data Anal.* **50** 2422–2446. [MR2225577](#)
- [4] BONGIORNO, E., GOIA, A. and SALINELLI, E. (2014). Contributions in infinite-dimensional statistics and related topics. Societa Editrice Esculapio.
- [5] BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data. Methods, Theory and Applications*. Springer, Heidelberg. [MR2807761](#)
- [6] CARDOT, H., FERRATY, F. and SARDA, P. (2003). Spline estimators for the functional linear model. *Statist. Sinica* **13** 571–591. [MR1997162](#)
- [7] CHEN, D., HALL, P. and MÜLLER, H.-G. (2011). Single and multiple index functional regression models with nonparametric link. *Ann. Statist.* **39** 1720–1747. [MR2850218](#)
- [8] FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- [9] FAN, J. and LV, J. (2011). Nonconcave penalized likelihood with NP-dimensionality. *IEEE Trans. Inform. Theory* **57** 5467–5484. [MR2849368](#)
- [10] FAN, J. and PENG, H. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *Ann. Statist.* **32** 928–961. [MR2065194](#)

- [11] FAN, Y., FOUTZ, N., JAMES, G. M. and JANK, W. (2014). Functional response additive model estimation with online virtual stock markets. *Ann. Appl. Stat.* **8** 2435–2460. [MR3292504](#)
- [12] FAN, Y., JAMES, G. M. and RADCHENKO, P. (2015). Supplement to “Functional additive regression.” DOI:[10.1214/15-AOS1346SUPP](#).
- [13] FEBRERO-BANDE, M. and GONZÁLEZ-MANTEIGA, W. (2013). Generalized additive models for functional data. *TEST* **22** 278–292. [MR3062258](#)
- [14] FERRATY, F., GOIA, A., SALINELLI, E. and VIEU, P. (2013). Functional projection pursuit regression. *TEST* **22** 293–320. [MR3062259](#)
- [15] FERRATY, F. and VIEU, P. (2003). Curves discrimination: A nonparametric functional approach. *Comput. Statist. Data Anal.* **44** 161–173. [MR2020144](#)
- [16] FERRATY, F. and VIEU, P. (2009). Additive prediction and boosting for functional data. *Comput. Statist. Data Anal.* **53** 1400–1413. [MR2657100](#)
- [17] GOIA, A. (2012). A functional linear model for time series prediction with exogenous variables. *Statist. Probab. Lett.* **82** 1005–1011. [MR2910049](#)
- [18] HALL, P., POSKITT, D. S. and PRESNELL, B. (2001). A functional data-analytic approach to signal discrimination. *Technometrics* **43** 1–9. [MR1847775](#)
- [19] HALL, P., REIMANN, J. and RICE, J. (2000). Nonparametric estimation of a periodic function. *Biometrika* **87** 545–557. [MR1789808](#)
- [20] HASTIE, T. and MALLOWS, C. (1993). Comment on “A statistical view of some chemometrics regression tools.” *Technometrics* **35** 140–143.
- [21] HUANG, J., HOROWITZ, J. L. and WEI, F. (2010). Variable selection in nonparametric additive models. *Ann. Statist.* **38** 2282–2313. [MR2676890](#)
- [22] JAMES, G. M. (2002). Generalized linear models with functional predictors. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 411–432. [MR1924298](#)
- [23] JAMES, G. M. and SILVERMAN, B. W. (2005). Functional adaptive model estimation. *J. Amer. Statist. Assoc.* **100** 565–576. [MR2160560](#)
- [24] LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *J. Amer. Statist. Assoc.* **86** 316–342. [MR1137117](#)
- [25] LIAN, H. (2011). Functional partial linear model. *J. Nonparametr. Stat.* **23** 115–128. [MR2780819](#)
- [26] LOH, P.-L. and WAINWRIGHT, M. J. (2015). Regularized M -estimators with nonconvexity: Statistical and algorithmic theory for local optima. *J. Mach. Learn. Res.* **16** 559–616. [MR3335800](#)
- [27] LV, J. and FAN, Y. (2009). A unified approach to model selection and sparse recovery using regularized least squares. *Ann. Statist.* **37** 3498–3528. [MR2549567](#)
- [28] MAS, A. and PUMO, B. (2007). The ARHD model. *J. Statist. Plann. Inference* **137** 538–553. [MR2298956](#)
- [29] MEIER, L., VAN DE GEER, S. and BÜHLMANN, P. (2009). High-dimensional additive modeling. *Ann. Statist.* **37** 3779–3821. [MR2572443](#)
- [30] MÜLLER, H.-G. and STADTMÜLLER, U. (2005). Generalized functional linear models. *Ann. Statist.* **33** 774–805. [MR2163159](#)
- [31] MÜLLER, H.-G. and YAO, F. (2008). Functional additive models. *J. Amer. Statist. Assoc.* **103** 1534–1544. [MR2504202](#)
- [32] RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer, New York. [MR2168993](#)
- [33] RAVIKUMAR, P., LAFFERTY, J., LIU, H. and WASSERMAN, L. (2009). Sparse additive models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 1009–1030. [MR2750255](#)
- [34] SCHUMAKER, L. L. (2007). *Spline Functions: Basic Theory*, 3rd ed. Cambridge Univ. Press, Cambridge. [MR2348176](#)
- [35] SIMON, N. and TIBSHIRANI, R. (2012). Standardization and the group Lasso penalty. *Statist. Sinica* **22** 983–1001. [MR2987480](#)

- [36] STOREY, J. D., XIAO, W., LEEK, J. T., TOMPKINS, R. G. and DAVIS, R. W. (2005). Significance analysis of time course microarray experiments. *Proc. Natl. Acad. Sci. USA* **102** 12837–12842.
- [37] WASSERMAN, L. (2006). *All of Nonparametric Statistics*. Springer, New York. [MR2172729](#)
- [38] YU, Y. and RUPPERT, D. (2002). Penalized spline estimation for partially linear single-index models. *J. Amer. Statist. Assoc.* **97** 1042–1054. [MR1951258](#)
- [39] ZHOU, S., SHEN, X. and WOLFE, D. A. (1998). Local asymptotics for regression splines and confidence regions. *Ann. Statist.* **26** 1760–1782. [MR1673277](#)
- [40] ZHU, H., VANNUCCI, M. and COX, D. D. (2010). A Bayesian hierarchical model for classification with selection of functional predictors. *Biometrics* **66** 463–473. [MR2758826](#)
- [41] ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2007). On the “degrees of freedom” of the lasso. *Ann. Statist.* **35** 2173–2192. [MR2363967](#)
- [42] ZOU, H. and LI, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *Ann. Statist.* **36** 1509–1533. [MR2435443](#)

DATA SCIENCES AND OPERATIONS DEPARTMENT
MARSHALL SCHOOL OF BUSINESS
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA 90089
USA
E-MAIL: fanyingy@marshall.usc.edu
gareth@marshall.usc.edu
radchenk@marshall.usc.edu