

Neural Networks for Scalar Input and Functional Output

Sidi Wu¹, Cédric Beaulac² and Jiguo Cao¹

¹Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby, BC, Canada.

²Département de Mathématiques, Université du Québec à Montréal, Montréal, Québec, Canada.

Abstract

The regression of a functional response on a set of scalar predictors can be a challenging task, especially if there is a large number of predictors, or the relationship between those predictors and the response is nonlinear. In this work, we propose a solution to this problem: a feed-forward neural network (NN) designed to predict a functional response using scalar inputs. First, we transform the functional response to a finite-dimension representation and construct an NN that outputs this representation. Then, we propose to modify the output of an NN via the objective function and introduce different objective functions for network training. The proposed models are suited for both regularly and irregularly spaced data, and a roughness penalty can be further applied to control the smoothness of the predicted curve. The difficulty in implementing both those features lies in the definition of objective functions that can be back-propagated. In our experiments, we demonstrate that our model outperforms the conventional function-on-scalar regression model in multiple scenarios while computationally scaling better with the dimension of the predictors.

Keywords: Functional data analysis, Functional response, Functional principal component analysis, Machine learning

1 Introduction

Functional data analysis (FDA) is a rapidly developing branch of statistics which targets at the theory and analysis of functional variables. As the atom of FDA, functional variables or functional data refer to curves, surfaces and any random variables taking values in an infinite dimensional space, such as time and spatial space [15, 28]. A basic and commonly recognized framework in FDA is to regard the functional data as realizations of an underlying stochastic process [38], and indeed, a large fraction of data coming from different fields can be characterized as functional data. Figure 1, for example, illustrates the fertility rates over the age of females for 92 countries. Each curve is treated as a smooth function of age and can be viewed as functional data. The fertility curves were measured at several time points, each of which represents an age group in increasing order. This age-specific fertility rate (ASFR) data was collected from the United Nations Gender Information (UNGEN) database by Mehrotra and Maity [35]. They estimated the observations based on the surveys conducted between 2000 and 2005. The data set is publicly available at https://github.com/suchitm/fosr_clust.




Fig. 1 Age-specific fertility rates for 92 countries around the world.

Among different aspects of FDA, functional regression has received the most attention in applications and methodological development [25]. Functional regression, in general, can be divided into three types: (1) scalar-on-function regression (regression analysis of scalar responses on functional predictors), (2) function-on-function regression (predictors and responses of the regression are both functional), and (3) function-on-scalar (FoS) regression (regression of functional responses on scalar predictors, also named as the varying-coefficient model). While plenty of work has been done for the first two scenarios, studies discussing functional response regression remain scarce.

Functional linear models (FLM), a series of extensions of the classic linear models, are the most conventional and widely applied methods in dealing with regression problems. A general form of FLM for FoS regression can be written as

$$Y(t) = \mathbf{X}\boldsymbol{\beta}(t) + \epsilon(t), \quad (1)$$

where $Y(t)$ is the functional response, and \mathbf{X} is a vector of scalar covariates. The unknown vector $\boldsymbol{\beta}(t)$ consists of parameter functions varying over t , and $\epsilon(t)$ stands for the random error term. In FDA, the consensual approach is to represent functional data with a linear combination of a finite number of known basis functions. Ramsay & Silverman [28] introduced projecting functional response with basis functions like B-splines for fitting the FoS regression model. Chiou et al. [6, 7] proposed to use functional principal component analysis (FPCA) for dimensional reduction and represent the functional response with eigenfunctions obtained via spectral decomposition of the covariance function of $Y(t)$. Both approaches summarize the information carried by the functional variable $Y(t)$ to a finite-dimensional vector of scalar representations (B-spline basis coefficients or functional principal component scores (FPC scores)). In solving linear functional regression problems, taking FoS regression as an example, the linear relation between the scalar predictors and the functional response is indeed captured by fitting a linear relation between the predictors and the scalar representation.

Some nonlinear approaches have been developed to handle more complicated regression settings. Zhang and Wang [43] combined the FoS with the additive models to gain the varying-coefficient additive model for functional data which does not require the linearity assumption between the scalar predictors and the functional response. Afterwards, an extension to this varying-coefficient additive model, named as functional additive mixed (FAM) model, was established by Scheipl et al. [13] with a general form

$$Y(t) = \sum_{r=1}^R f_r(\mathbf{X}_r, t) + \epsilon(t), \quad (2)$$

for functional response $Y(t)$. Each term in the additive predictor $f_r(\mathbf{X}_r, t)$ is a function of t and a subset \mathbf{X}_r of the complete predictor set \mathbf{X} including scalar and functional covariates. This extensive framework can consist of both linear and nonlinear effects of functional and scalar covariates that may vary smoothly over the index of the functional response. Naturally, this model was further extended, by Scheipl et al [14], to the generalized functional additive mixed model in order to take care of non-Gaussian functional response. Although several nonlinear attempts have been made, the existing regression methods for FDA have been predominantly linear [38]. Considering the emerging trend of complicatedly structured functional data, there is an increasing demand to develop more nonlinear approaches to FDA.

In this work, we propose a solution to the FoS problem which is able to handle a large number of predictors and the nonlinear relation between the scalar predictors and the functional response. Our solution borrows from the machine learning literature, which follows a trend in adapting machine learning techniques to known statistical problems such as survival analysis [3, 21, 23] or causal inference [24, 33]. We adapt the neural network (NN) architecture for functional data.

In some existing works, efforts have been made to combine deep neural networks to the field of functional data analysis. For instance, Rossi et al. [30, 31] firstly explored the idea of applying NNs to functional data by constructing a functional neural network (FNN) with functional neurons in the first hidden layer for functional inputs. FNN was then extended by Thind et al. [36, 37] to feed both functional and scalar covariates as inputs and outputs a scalar response. Rao et al. [29] equipped FNN with geographically weighted regression and spatial autoregressive technique to handle regression problems with spatially correlated functional data. Meanwhile, Wang et al. [40] proposed a nonlinear function-on-function model using a fully connected NN. Previously, Yao et al. [42] developed an NN with a new basis layer whose hidden neurons are micro neural networks, to perform a parsimonious dimension reduction for functional inputs using information relevant to the scalar target.

Most of the mentioned works are focused on building NNs with functional inputs and scalar outputs. In this work, we consider the other side of the coin. We design an NN meant to predict a functional response and to the best of our knowledge, this study is the first attempt to solve the FoS regression problem using artificial neural networks. Because the standard machine learning techniques are designed for finite-dimensional feature vectors, we propose to encode the information contained in the intrinsically infinite-dimensional functional response to a finite-dimensional vector of scalar representations. Different from other nonlinear approaches targeting at FoS regression problems, our methods, in particular, focus on studying the nonlinear association between the predictors and the scalar representations of the functional response, to further reveal the relation between the scalar predictors and the functional response. In this way, we maintain the interpretability of the relation between the scalar predictors and the functional response, despite of the usage of NNs.

Challenged by the ASFR data introduced previously, in this work, we are interested in conducting a FoS analysis to accurately predict the functional curves using scalar predictors and also reveal any potential nonlinear relation between the predictors and the response. For each of the 92 observations (countries), the functional curve of fertility rates is associated with 15 demographic and socioeconomic variables averaging over the information available on Gapminder in a country-level manner from 2000 to 2005 [35]. These 15 scalar covariates consist of age at first marriage, under-5 mortality, maternity deaths per 1000 women, cervical cancer deaths per 100,000 women, female labor force participation, male to female ratio (women aged from 15-49), contraception prevalence (women aged from 15-49), life expectancy, mean years

of school (women % men, women aged from 15-34), female's body mass index (BMI), the number of births attended by trained birth staff (% of total), gross domestic product (GDP) per capita, the proportion of dollar billionaires per 1 million people, health expenditures (% of GDP) and the amount of alcohol its populace consumes. All the covariates downloaded from the listed data source have been previously standardized, and many predictor pairs among them are found highly correlated.

Due to the difficulty in directly determining whether there is a nonlinear relation between a scalar variable and a random function, we pay more attention to the association between the scalar covariates and the scalar representations summarizing the information of the functional response. Regarding the scalar representations, we choose to use the basis coefficients which are estimated by approximating the response function with a linear expansion of the most common B-spline basis functions. Besides the initial choice of the basis coefficients, we also attempt to implant the FPC scores as the scalar representations in the real application.




Fig. 2 Plots of the second B-spline basis coefficient (estimated using the functional response $Y(t)$) versus two scalar covariates, including female age and under-5 mortality, separately. The red lines are the smoothing curves estimated using R function `lowess()`.

Accordingly, with the ASFR data set, instead of picturing the relation between the infinite-dimensional response curve and the 15 scalar covariates, we obtain visualizations of one B-spline basis coefficient against one predictor for all possible combinations. The representation process is purely based on the fertility trajectories with no contribution from the 15 predictors. Six B-spline basis functions are employed, which is recommended by the cross-validation we performed. Figure 2 displays the in-pair associations for some selected basis coefficient-predictor pairs. We incorporate a locally-weighted smoothing curve

in red to help visualizing if the association is linear or not. The existence of nonlinear associations between some coefficient-predictor pairs is endorsed, as shown with the female age predictor in Figure 2. On the other hand, the association pattern of the second basis coefficient to under-5 mortality is basically linear as depicted in Figure 2. As both linear and nonlinear patterns are detected among all possible basis coefficient-predictor sets, we would like to set up a procedure with the ability to include every available covariate to precisely predict the response trajectory, while simultaneously capturing the true relationships between the predictors and the scalar representations of the functional response.

Three main contributions are made in this paper. First, we construct a general framework of NN for functional response which aims at predicting curves using the input through a feed-forward NN architecture by slightly modifying the output of the NN via the objective function. Such modifications are readily applicable to other types of networks such as convolutional neural networks (CNN) [45], and long short-term memory (LSTM) model [16, 17, 20]. Second, we pay attention to the behavior of the functional output during the network training process and propose to hold their natural smoothness property by adding a roughness penalty in the objective function for training NN. Lastly, we make numerous comparisons under various scenarios to analyze and conclude the conditions of use for different versions of our methods.

The biggest challenge to the implementation of the aforementioned contributions is to ensure the NN models can be trained using backpropagation as a standard NN [32]. Theoretically, we need to define the flow of information from the input to the objective function as a sequence of operations differentiable with respect to the weights of the NN. However, in practice, because we rely on an auto-differentiation package, such as **Keras** [8], to train the NN, we need to define the objective function as a sequence of operations provided by that package. The available differential operations are incapable of dealing with derivatives of the integral of the infinite-dimensional functional response with respect to NN weights. Besides, some common features when fitting functional data, such as roughness penalties, were obviously not designed with that consideration in mind. For example, it is common to penalize the integrated second-order derivative of a functional curve for smoothing in FDA. But a roughness penalty cannot be simply integrated into the objective function when training the NN. Overcoming that challenge is a central part of these contributions mentioned above and to do so, we needed to be creative and adapt the objective functions in various ways detailed later in this manuscript.

The remainder of this paper is organized as follows. In Section 2, we detail the proposed NN with a functional response, with an additional discussion on how to control the smoothness of the predicted curves. A brief discussion about the computational costs of NNs and the FLM with functional response is provided in Section 3. Section 4 contains the results of prediction in the real application using different methods. In section 5, we conduct simulation

studies, under linear and nonlinear scenarios, and compare the proposed models with the existing FoS model in both the predictive accuracy for discretely observed time points and the ability to reconstruct the true response trajectories over a continuum. Lastly, some concluding remarks are given in Section 6.

2 Methodology

Suppose we have N subjects and for the i -th subject, the input is a set of numerical variables $\mathbf{X}_i = \{X_{i1}, X_{i2}, \dots, X_{iP}\}$, and the output is a functional variable $Y_i(t), t \in \mathcal{T}$ in the $L^2(t)$ space. Functional data is often assumed to lie in the L^2 space because it ensures that the data is square-intergrable, allowing the optimization of the statistical models with respect to the squared-loss criterion in the FoS regression framework. In reality, the functional response $Y_i(t)$ is usually measured in a discrete manner, for instance, at m_i time points or locations. Therefore, instead of observing the full trajectory of $Y_i(t)$, for the i -th subject, we obtain m_i pairs of observations $\{t_{ij}, Z_i(t_{ij})\}, j = 1, 2, \dots, m_i$ and $Z_i(t_{ij}) = Y_i(t_{ij}) + \epsilon_i(t_{ij})$, where $\epsilon_i(t_{ij})$ is the i.i.d. observation error. To simplify the situation, in the following discussion, we assume that for all subjects, the functional term $Y_i(t)$'s are observed at the same m dense and equally-spaced time points. In other words, we assume that $m_1 = m_2 = \dots = m_N = m$ and $t_{1j} = t_{2j} = \dots = t_{Nj} = t_j$ for $j = 1, 2, \dots, m$. A summary of notations used throughout the manuscript is provided in Table S1 in the supplementary document.

2.1 Neural network with functional response: mapping to basis coefficients (NNBB)

To overcome the difficulty in modeling infinite-dimensional data, a common pipeline in FDA is to summarize the information of functions $\{Y_i(t)\}_{i=1}^N$ into a set of finite-dimensional vectors of coefficients using some basis representation method, and consequently,

$$Y_i(t) = \sum_{k=1}^{K_b} c_{ik} \theta_k(t) = \boldsymbol{\theta}' \mathbf{C}_i, \quad (3)$$

where K_b is the number of basis functions. It is common practice to select K_b by cross-validation [28]. The vector $\boldsymbol{\theta}$ contains the basis functions $\theta_1(t), \dots, \theta_{K_b}(t)$ from a selected basis system, such as the Fourier basis system or B-spline system (in our study, we mainly used the latter), and $\theta_k(t)$ also belongs to L^2 space. \mathbf{C}_i is a K_b -dimensional vector of basis coefficients $\{c_{ik}\}_{k=1}^{K_b}$ which lies in \mathbb{R}^{K_b} and needs to be determined. These basis coefficients serve as parameters of this model. Without the smooth underlying functions $Y_i(t)$'s, the discrete observations $Z_i(t_{ij})$'s are used to estimate $\{c_{ik}\}_{k=1}^{K_b}$ by fitting $Z_i(t_{ij}) = \sum_{k=1}^{K_b} c_{ik} \theta_k(t_{ij}) + \epsilon_i(t_{ij})$. The least square estimator of c_{ik} , denoted

by c_{ik}° , is obtained by minimizing the sum of squared error (SSE) criterion $\text{SSE}(\mathbf{Z}|\mathbf{C}) = \sum_{i=1}^N \sum_{j=1}^m \left(Z_i(t_{ij}) - \sum_{k=1}^{K_b} c_{ik} \theta_k(t_{ij}) \right)^2$.

Eq.(3) implies that $Y(t)$ can be approximated by a linear combination of the basis functions $\theta(t)$'s, which carry the fixed modes of variations over \mathcal{T} with real-valued basis coefficients. Learning how the predictors X 's regress on the variation of $Y(t)$, as a result, can be naturally replaced with learning how the predictors X 's regress on the set of unknown basis coefficients c 's. Therefore we propose to set the basis coefficients c 's as a function of X 's. Let $F(\cdot)$ be a mapping function from \mathbb{R}^P to \mathbb{R}^{K_b} so that \mathbf{X} can be mapped to the basis coefficients as

$$\mathbf{C}_i = F(\mathbf{X}_i), \quad (4)$$

which in turn can be used to map \mathbf{X} to the functional response $Y(t)$ with $Y_i(t) = \boldsymbol{\theta}' F(\mathbf{X}_i)$. Consequently, we can learn the nonlinear relation between \mathbf{X}_i and $Y_i(t)$ by setting $F(\cdot)$ to be a nonlinear function. Here we propose a dense feed-forward NN as the mapping function $F(\cdot)$, where the basis coefficients $\{c_{i1}, c_{i2}, \dots, c_{iK_b}\}$ are the outputs of the NN, then the model can be expressed as

$$\mathbf{C}_i = \text{NN}_\eta(\mathbf{X}_i) = g_{L+1} \left(\cdots g_1 \left(\sum_{p=1}^P w_{1p} X_{ip} + b_1 \right) \right), \quad (5)$$

where g_1, \dots, g_{L+1} are the activation functions at each layer with L being the number of hidden layers, and η denotes the NN parameter set consisting of weights $\{w_{\ell p}\}_{\ell=1}^{L+1}$ and bias $\{b_\ell\}_{\ell=1}^{L+1}$ of all layers. The NN is optimized by minimizing the standard objective function calculating the mean squared error (MSE) as $L_C(\eta) = 1/n_{\text{train}} \sum_{i=1}^{n_{\text{train}}} \sum_{k=1}^{K_b} (\hat{c}_{ik} - c_{ik})^2$, where n_{train} is the number of samples in the training set, and in application the *true* basis coefficients c_{ik} 's are replaced with c_{ik}° 's.

We named this model NNBB as it is an NN, with B-spline coefficient output (B) and trained by minimizing the MSE of those B-spline coefficients (B). With the exception of the output layer, the NNBB uses a conventional NN architecture and thus, typical approaches for hyper-parameter tuning and training can be applied. We summarize the training process of NNBB in Algorithm 1. The optimized NN is used for prediction where it takes the new scalar inputs in the test set and then outputs the predicted basis coefficients $\hat{\mathbf{C}}_{\text{new}}$. The predicted functional response $\hat{\mathbf{Y}}(t)$ is further constructed as $\hat{\mathbf{Y}}(t) = \boldsymbol{\theta}' \hat{\mathbf{C}}_{\text{new}}$.

2.2 Neural network with functional response: mapping to FPC scores (NNSS)

Apart from basis expansion, the other popular approach for dimension reduction is FPCA. To be specific, let $\mu(t)$ and $K(t, t') = \text{cov}(Y(t), Y(t'))$ be the

Algorithm 1: Training NNBB

Input: $\mathbf{X}_i = \{X_{i1}, X_{i2}, \dots, X_{iP}\}$, $\{Z_i(t_{ij})\}_{j=1}^m$ in the training set n_{train}

Output: NN with the optimized parameter set $\hat{\eta}$

- 1 set up hyper-parameters, including:
 - for smoothing $Z_i(t_{ij})$: basis functions vector $\boldsymbol{\theta} = [\theta_1(t), \dots, \theta_{K_b}(t)]$, number of basis functions K_b
 - for training NN: number of hidden layers (L), number of neurons per hidden layer, activation functions (g_1, \dots, g_{L+1}), number of epochs (E), batch size, NN optimizer (with a learning rate ρ), etc.
 - 2 **forall** $i \in n_{\text{train}}$ **do** obtain $\{c_{ik}^o\}_{k=1}^{K_b}$ by projecting $\{Z_i(t_{ij})\}_{j=1}^m$ to the selected set of basis functions $\boldsymbol{\theta}$
 - 3 randomly initialize NN parameter set and get $\eta = \eta_{\text{initial}}$
 - 4 train a fully-connected NN with inputs \mathbf{X}_i and outputs $\{c_{ik}^o\}_{k=1}^{K_b}$
 - for** $e = 1$ **to** E **do**
 - I. *forward propagation*
 - (i) pass \mathbf{X}_i through the NN and get $\{\hat{c}_{ik}\}_{k=1}^{K_b}$ for all $i \in n_{\text{train}}$
 - (ii) calculate $L(\eta) = 1/n_{\text{train}} \sum_{i=1}^{n_{\text{train}}} \sum_{k=1}^{K_b} (\hat{c}_{ik} - c_{ik}^o)^2$
 - II. *backward propagation*
 - (i) update NN parameter set as $\eta^* = \eta - \rho \frac{\partial L(\eta)}{\partial \eta}$
 - III. **if** $e < E$ **then** set $\eta = \eta^*$, and then repeat I. & II.
 else return η^*
 - end**
 - 5 **return** NN with the estimated parameter set $\hat{\eta} = \eta^*$
-

mean and covariance functions of the underlying function $Y(t)$, and accordingly, the spectral decomposition of the covariance function is $K(t, t') = \sum_{k=1}^{\infty} \gamma_k \phi_k(t) \phi_k(t')$, where $\{\gamma_k, k \geq 1\}$ are the eigenvalues in decreasing order with $\sum_k \gamma_k < \infty$ and ϕ_k 's are corresponding eigenfunctions, also named functional principle components (FPCs), with restriction of $\int \phi_k^2(t) dt = 1$ for all k . Following the Karhunen-Loève expansion, the i -th observed functional response $Y_i(t)$ can be represented as $Y_i(t) = \mu(t) + \sum_{k=1}^{\infty} \xi_{ik} \phi_k(t)$. Denote $\tilde{Y}_i(t) = Y_i(t) - \mu(t)$ as the centered functional response for the i -th subject, we can express the expansion as

$$\tilde{Y}_i(t) = \sum_{k=1}^{\infty} \xi_{ik} \phi_k(t), \quad (6)$$

where $\xi_{ik} = \int \{Y_i(t) - \mu(t)\} \phi_k(t) dt$ is the k -th FPC score for $Y_i(t)$ with zero mean and $\text{cov}(\xi_{ik}, \xi_{il}) = \gamma_k \mathbf{1}(k = l)$. Representing functions with eigenfunctions eigen-decomposed from $\text{cov}(Y(t), Y(t'))$ and the corresponding FPC

scores can be considered as a special case of basis expansion, where eigenfunctions play the role of basis functions which, however, are unknown and need to be numerically calculated using the observed functional data. Given a desired proportion of variance explained (τ) by the FPCs, $\tilde{Y}_i(t)$ can be approximated arbitrarily well by a finite number of the leading FPCs, where

$$\tilde{Y}_i(t) \approx \sum_{k=1}^{K_\tau} \xi_{ik} \phi_k(t), \quad (7)$$

and K_τ is a valued parameter truncating the FPCs and reducing the dimension of $Y(t)$. K_τ is determined by τ as the smallest integer satisfying $\sum_{k=1}^{K_\tau} \gamma_k \geq \tau$. Because of the fast decay rate of γ_k , K_τ is usually a small integer. In practice, FPCA can be easily performed in \mathbf{R} with the help of some FDA-related packages, e.g., **fd**a [27] and **fdapace** [4]. The **fd**a package requires the users to firstly smooth the discrete observations $Z_i(t_{ij})$'s and then apply FPCA on the smoothed functional data for $\{\xi_{ik}^\circ\}_{k=1}^{K_\tau}$ and $\{\hat{\phi}_k(t)\}_{k=1}^{K_\tau}$, an estimator of $\{\xi_{ik}\}_{k=1}^{K_\tau}$ and $\{\phi_k(t)\}_{k=1}^{K_\tau}$, respectively, while the **fdapace** package is able to estimate $\{\xi_{ik}\}_{k=1}^{K_\tau}$ and $\{\phi_k(t)\}_{k=1}^{K_\tau}$ with observed data $Z_i(t_{ij})$'s directly.

It is clearly shown in Eq.(7) that $\tilde{Y}_i(t)$ can be effectively approximated by a linear combination of the top eigenfunctions $\{\phi_k(t)\}_{k=1}^{K_\tau}$, with the major modes of variations among $\tilde{Y}_i(t)$ captured. Following the idea in Wang & al. [40], under the regression setting, in order to learn the impact of different values of \mathbf{X}_i on the variation of $\tilde{Y}_i(t)$, we propose to set the coefficients $\boldsymbol{\xi}_i = \{\xi_{i1}, \xi_{i2}, \dots, \xi_{iK_\tau}\}$ to be a function of \mathbf{X}_i , by formulation, $\boldsymbol{\xi}_i = F(\mathbf{X}_i)$ and consequently we have

$$\tilde{Y}_i(t) = \boldsymbol{\phi}' F(\mathbf{X}_i), \quad (8)$$

where $\boldsymbol{\phi}$ is a K_τ -dimensional vector cataloging the K_τ leading FPCs. To model the nonlinear relation, similarly, we proposed to set the mapping function $F(\cdot)$ as a dense feed-forward NN, and then we can write

$$\boldsymbol{\xi}_i = \text{NN}_\eta(\mathbf{X}_i) = g_{L+1} \left(\cdots g_1 \left(\sum_{p=1}^P w_{1p} X_{ip} + b_1 \right) \right), \quad (9)$$

where g_1, \dots, g_{L+1} are the activation functions at each layer with L being the number of hidden layers, and η denotes the NN parameter set consisting of weights $\{w_{\ell p}\}_{\ell=1}^{L+1}$ and bias $\{b_\ell\}_{\ell=1}^{L+1}$ of all layers. Likewise, $\text{NN}_\eta(\mathbf{X}_i)$ is trained by minimizing the MSE loss function $L_\boldsymbol{\xi}(\eta) = 1/n_{\text{train}} \sum_{i=1}^{n_{\text{train}}} \sum_{k=1}^{K_\tau} (\hat{\xi}_{ik} - \xi_{ik})^2$, where n_{train} stands for the number of training set subjects, and the *true* FPC scores ξ_{ik} 's are replaced by ξ_{ik}° 's.

Following the same naming convention, we refer to this model has NNSS because it is an NN model build with output being the FPC scores (S) trained on the scores (S) themselves. The complete training process of NNSS is summarized in Algorithm 2. Again, the hyper-parameter tuning and model training

Algorithm 2: Training NNSS

Input: $\mathbf{X}_i = \{X_{i1}, X_{i2}, \dots, X_{iP}\}$, $\{Z_i(t_{ij})\}_{j=1}^m$ in the training set n_{train}

Output: NN with the optimized parameter set $\hat{\eta}$

- 1 set up hyper-parameters, including:
 - for performing FPCA: basis functions vector $\boldsymbol{\theta} = [\theta_1(t), \dots, \theta_{K_b}(t)]$, number of basis functions K_b , the desired proportion of variance explained τ for determining K_τ
 - for training NN: number of hidden layers (L), number of neurons per hidden layer, activation functions (g_1, \dots, g_{L+1}), number of epochs (E), batch size, NN optimizer (with a learning rate τ), etc.
 - forall** $i \in n_{\text{train}}$ **do**
 - 2 Estimate the mean function $\mu(t)$ and the covariance function $K(t, t')$ using $\{Z_i(t_{ij})\}_{j=1}^m$
 - 3 Perform eigen-decomposition on $\hat{K}(t, t')$ and get $\{\hat{\phi}_k(t)\}_{k=1}^{K_\tau}$ and the corresponding FPC scores $\{\xi_{ik}^\circ\}_{k=1}^{K_\tau}$
 - end**
 - 4 randomly initialize NN parameter set and get $\eta = \eta_{\text{initial}}$
 - 5 train a fully-connected NN with inputs \mathbf{X}_i and outputs $\{\xi_{ik}^\circ\}_{k=1}^{K_\tau}$
 - for** $e = 1$ **to** E **do**
 - I. *forward propagation*
 - (i) pass \mathbf{X}_i through the NN and get $\{\hat{\xi}_{ik}\}_{k=1}^{K_\tau}$ for all $i \in n_{\text{train}}$
 - (ii) calculate $L(\eta) = 1/n_{\text{train}} \sum_{i=1}^{n_{\text{train}}} \sum_{k=1}^{K_\tau} (\hat{\xi}_{ik} - \xi_{ik}^\circ)^2$
 - II. *backward propagation*
 - (i) update NN parameter set as $\eta^* = \eta - \varrho \frac{\partial L(\eta)}{\partial \eta}$
 - III. **if** $e < E$ **then** set $\eta = \eta^*$, and then repeat I. & II.
 else return η^*
 - end**
 - 6 **return** NN with the estimated parameter set $\hat{\eta} = \eta^*$
-

of NNSS are the same as of a conventional NN. Meanwhile, the returned NN with optimized parameters set $\hat{\eta}$ is used for prediction where it takes the new scalar inputs in the test set and then outputs the predicted FPC scores $\hat{\boldsymbol{\xi}}_{\text{new}}$. The predicted functional response $\hat{\mathbf{Y}}(t)$ is further recovered as $\hat{\mathbf{Y}}(t) = \hat{\boldsymbol{\mu}}(t) + \hat{\boldsymbol{\phi}}' \hat{\boldsymbol{\xi}}_{\text{new}} = \hat{\boldsymbol{\mu}}(t) + \hat{\boldsymbol{\phi}}' \text{NN}(\mathbf{X}_{\text{new}} | \hat{\eta})$, where $\hat{\boldsymbol{\phi}}$ is the vector consisting of the estimated FPCs $\hat{\phi}_1(t), \dots, \hat{\phi}_{K_\tau}(t)$.

2.3 Modification to the objective function (NNBR and NNSR)

In the two previous sections, we define two NNs outputting basis coefficients or FPC scores, and those outputs are further used to construct the predicted

response variable. Let us discuss a way to build an objective function that directly uses the response variable in order to estimate the parameters. For simplicity, we only discuss the B-spline model of Section 2.1 but know that a similar concept can also be applied to the FPCA model described in Section 2.2. In brief, in Section 2.1, we first fit a B-spline model on the observed functional response to estimate a set of basis coefficients for the response in the training set and then we train an NN to predict these basis coefficients using the predictors.

In this section, we propose to modify the objective function in order to bypass the initial estimation of basis coefficients. The key idea of the new objective function is to directly minimize the prediction error of the response variable. In other words, instead of minimizing the MSE between \hat{c}_{ik} and c_{ik} , here we first transform the NN output, the B-splines basis coefficients, into the predicted response $\hat{Y}_i(t)$ and then minimize the MSE between $\hat{Y}_i(t)$ and $Y_i(t)$. What supports us to implement such objective function is the fact that we rely on differentiable operations to build this new objective function. Doing so guarantees that we can rely on the back-propagation algorithm to train the NN using readily available packages.

Suppose that we want to fit the functional response with a B-spline made of K_b basis functions $\boldsymbol{\theta}$ and K_b associated basis coefficients \mathbf{C} . Then the predicted response $\hat{Y}_i(t)$ at time t is the vector product between $\boldsymbol{\theta}$ and \mathbf{C}_i as illustrated in Eq. (3). This means that the relation between the predicted response $\hat{Y}_i(t)$ and the predicted basis coefficients $\hat{\mathbf{C}}_i$ is linear, thus we can easily compute the derivative of $\hat{Y}_i(t)$ with respect to the coefficients. This further indicates that if we observe the response at time t then we are able to compute the gradient of $(Y_i(t) - \hat{Y}_i(t))^2$ with respect to the basis coefficients and therefore we can also compute the gradient with respect to the parameters η of the NN that outputs those basis coefficients.

More generally, assuming the response is observed at m time points, t_1, t_2, \dots, t_m for every subject, we can generate a matrix of basis functions evaluated at those time points, denoted as Θ , which is a $K_b \times m$ matrix where each row represents a basis function with entries taking the values of that basis function at each of the m time points. Consequently, we can obtain the predicted response $\hat{Y}_{(n_{\text{train}} \times m)}$ for every m observed time points and for all n_{train} subjects by doing a simple matrix multiplication

$$\hat{Y}_{(n_{\text{train}} \times m)} = \hat{\mathbf{C}}_{(n_{\text{train}} \times K_b)} \Theta_{(K_b \times m)}, \quad (10)$$

where $\hat{\mathbf{C}}$ is a matrix where a row contains the K_b NN-estimated coefficients for a single observed subject. In our proposed model, $\hat{\mathbf{C}}$ is the output produced by the NN function as explained in Section 2.1. However, we now modify the objective function to train the NN to minimize the MSE between the observed

response and the predicted one

$$L_{\mathbf{Y}}(\eta) = \frac{1}{n_{\text{train}}} \sum_{i=1}^{n_{\text{train}}} \sum_{j=1}^m (Y_i(t_{ij}) - \hat{Y}_i(t_{ij}))^2. \quad (11)$$

Note that in implementation, the discrete observation $Z_i(t_{ij}) = Y_i(t_{ij}) + \epsilon_i(t_{ij})$ will replace $Y_i(t_{ij})$, the underlying *true* functional curve. In this way, we entirely bypass the need to estimate the basis coefficients first and fit the parameters η with respect to the functional response directly. We call this variant NNBR because its architecture is an NN with basis coefficients output (B) fitted by minimizing the MSE of the response variable (R). The entire training process of NNBR is detailed in Algorithm 3. Similarly, we can get a prediction for a new input \mathbf{X}_{new} as $\hat{Y}(t) = \boldsymbol{\theta}' \hat{\mathbf{C}}_{\text{new}} = \boldsymbol{\theta}' \text{NN}(\mathbf{X}_{\text{new}} | \hat{\eta})$.

Algorithm 3: Training NNBR

Input: $\mathbf{X}_i = \{X_{i1}, X_{i2}, \dots, X_{iP}\}$, $\{Z_i(t_{ij})\}_{j=1}^m$ in the training set n_{train}

Output: NN with the optimized parameter set $\hat{\eta}$

- 1 set up hyper-parameters, including:
 - for smoothing $Z_i(t_{ij})$: basis functions vector $\boldsymbol{\theta} = [\theta_1(t), \dots, \theta_{K_b}(t)]$, number of basis functions K_b
 - for training NN: number of hidden layers (L), number of neurons per hidden layer, activation functions (g_1, \dots, g_{L+1}), number of epochs (E), batch size, NN optimizer (with a learning rate ρ), etc.
 - 2 Evaluate $[\theta_1(t), \dots, \theta_{K_b}(t)]$ at all observed timestamps $\{t_j\}_{j=1}^m$ and form the matrix $\Theta_{(K_b \times m)}$
 - 3 randomly initialize NN parameter set and get $\eta = \eta_{\text{initial}}$
 - 4 train a fully-connected NN with inputs \mathbf{X}_i and outputs $\{Z_i(t_{ij})\}_{j=1}^m$
 - for** $e = 1$ **to** E **do**
 - I. *forward propagation*
 - (i) pass \mathbf{X}_i through the NN and get $\{\hat{c}_{ik}\}_{k=1}^{K_b}$ for all $i \in n_{\text{train}}$
 - (ii) multiply $\{\hat{c}_{ik}\}_{k=1}^{K_b}$ with $\Theta_{(K_b \times m)}$ to get $\{\hat{Y}(t_{ij})\}_{j=1}^m$
 - (iii) calculate $L(\eta) = 1/n_{\text{train}} \sum_{i=1}^{n_{\text{train}}} \sum_{j=1}^m (Z_i(ij) - \hat{Y}_i(t_{ij}))^2$
 - II. *backward propagation*
 - (i) update NN parameter set as $\eta^* = \eta - \rho \frac{\partial L(\eta)}{\partial \eta}$
 - III. **if** $e < E$ **then** set $\eta = \eta^*$, and then repeat I. & II.
 else return η^*
 - end**
 - 5 **return** NN with the estimated parameter set $\hat{\eta} = \eta^*$
-

When comparing NNBR to NNBB, we note that they share a similar architecture but use different objective functions. A significant difference between both models is that in order to train the NNBB we need to first estimate the basis coefficients of the functional response. Conversely, training the NNBR requires a modification of the objective function which slightly increases the computational cost but meanwhile bypasses the need of a beforehand estimation of basis coefficients.

A similar process could be conducted with an NN that outputs FPC scores as described in Section 2.2, where this time it predicts response \hat{Y} using those scores. We named that model NNSR and describe the training process in Algorithm 4.

It is worth mentioning that the network architectures of NNBR and NNSR remain the same as a classic NN even though we have modified their objective functions. Such modifications only influence the network optimization with no changes to the sensitivity with respect to the hyper-parameters and the architecture (number of layers and depth). Additionally, the proposed modification that targets at the minimization of the difference between $Y_i(t)$ and $\hat{Y}_i(t)$ is surprisingly simple, convenient and computationally efficient. Compared to the operation that minimizes the difference between c_{ik} and \hat{c}_{ik} (NNBB) or between ξ_{ik} and $\hat{\xi}_{ik}$ (NNS), the modified objective function only requires an additional operation, namely a simple matrix multiplication as Eq. (10) in forward propagation and an additional step to compute the gradient of $(Y_i(t) - \hat{Y}_i(t))^2$ with respect to the basis coefficients outputted by NN in backward propagation. This only increases negligibly the computational cost of NNBR and NNSR when compared to NNBB or NNS.

2.4 Irregularly-spaced functional data

Earlier in this section we made the common assumption that the functional response $Y(t)$ is observed at the same m equally-spaced time points. While this was a useful assumption to make in order to explain how to train the model, it is actually not a necessary condition to fit any of the four models described above and we can train these models even with irregularly spaced functional response.

For the first two models we introduced, NNBB and NNS, in order to train the NN, we simply need an estimate for the basis coefficients or an estimate for the FPC scores. We can rely on some existing FDA literature to get those estimates [28, 41] for irregularly-spaced functional data.

For both models that utilize the modified objective function, NNBR and NNSR, it is a bit more complicated but not so much. Once again, let us focus on NNBR to simplify the explanations. In the setting with irregularly-spaced functional observations, the assumptions $m_1 = m_2 = \dots = m_N = m$ and $t_{1j} = t_{2j} = \dots = t_{Nj} = t_j$ no longer hold. Suppose that m_{irr} is the total number of time points with at least one observation given all training subjects and $\{t_{ij}\}_{j=1}^{m_{\text{irr}}}$ represents the union set of $\{t_{ij}\}_{j=1}^{m_i}$ for all i in the training set. The goal is to train the model using only the observations at $\{t_{ij}\}_{j=1}^{m_i}$ when the i -th

Algorithm 4: Training NNSR

Input: $\mathbf{X}_i = \{X_{i1}, X_{i2}, \dots, X_{iP}\}$, $\{Z_i(t_{ij})\}_{j=1}^m$ in the training set n_{train}

Output: NN with the optimized parameter set $\hat{\eta}$

- 1 set up hyper-parameters, including:
 - for performing FPCA: basis functions vector $\boldsymbol{\theta} = [\theta_1(t), \dots, \theta_{K_b}(t)]$, number of basis functions K_b , the desired proportion of variance explained τ for determining K_τ
 - for training NN: number of hidden layers (L), number of neurons per hidden layer, activation functions (g_1, \dots, g_{L+1}), number of epochs (E), batch size, NN optimizer (with a learning rate τ), etc.
 - forall** $i \in n_{\text{train}}$ **do**
 - 2 Estimate the mean function $\mu(t)$ and the covariance function $K(t, t')$ using $\{Z_i(t_{ij})\}_{j=1}^m$
 - 3 Perform eigen-decomposition on $\hat{K}(t, t')$ and get $\{\hat{\phi}_k(t)\}_{k=1}^{K_\tau}$
 - 4 Form $\hat{\Phi}_{(K_\tau \times m)}$, a matrix of eigenfunctions $[\hat{\phi}_1(t), \dots, \hat{\phi}_{K_\tau}(t)]$ evaluated at all observed time points $\{t_j\}_{j=1}^m$
 - end**
 - 5 randomly initialize NN parameter set and get $\eta = \eta_{\text{initial}}$
 - 6 train a fully-connected NN with inputs \mathbf{X}_i and outputs $\{Z_i(t_{ij})\}_{j=1}^m$
 - for** $e = 1$ **to** E **do**
 - I. *forward propagation*
 - (i) pass \mathbf{X}_i through the NN and get $\{\hat{\xi}_{ik}\}_{k=1}^{K_\tau}$ for all $i \in n_{\text{train}}$
 - (ii) multiply $\{\hat{\xi}_{ik}\}_{k=1}^{K_\tau}$ with $\hat{\Phi}_{(K_\tau \times m)}$ to get $\{\hat{Y}(t_{ij})\}_{j=1}^m$
 - (iii) calculate $L(\eta) = 1/n_{\text{train}} \sum_{i=1}^{n_{\text{train}}} \sum_{j=1}^m (Z_i(ij) - \hat{Y}_i(t_{ij}))^2$
 - II. *backward propagation*
 - (i) update NN parameter set as $\eta^* = \eta - \rho \frac{\partial L(\eta)}{\partial \eta}$
 - III. **if** $e < E$ **then** set $\eta = \eta^*$, and then repeat I. & II.
 else return η^*
 - end**
 - 7 **return** NN with the estimated parameter set $\hat{\eta} = \eta^*$
-

subject is observed at this time. To achieve this, we need to generate a matrix Θ_{irr} of size $K_b \times m_{\text{irr}}$ with the k -th row containing entries of the k -th chosen basis function evaluated at $\{t_{ij}\}_{j=1}^{m_{\text{irr}}}$. This matrix Θ_{irr} will be employed for the calculation of $\hat{Y}_i(t_{ij})$ at all m_{irr} time points with observation(s) by taking the matrix multiplication $\text{NN}_\eta(\mathbf{X})\Theta_{\text{irr}}$. However, because for each subject i , $Y_i(t_{ij})$'s are originally observed at only $\{t_{ij}\}_{j=1}^{m_i}$ instead of all m_{irr} time points, when computing the objective function, we need to eliminate the contribution of those unobserved subject-time point pairs by multiplying them by 0, and

accordingly

$$\begin{aligned} L_{\mathbf{Y}_{\text{irr}}}(\eta) &= \frac{1}{n_{\text{train}}} \sum_{i=1}^{n_{\text{train}}} \sum_{j=1}^{m_i} \left(Y_i(t_{ij}) - \hat{Y}_i(t_{ij}) \right)^2 \\ &= \frac{1}{n_{\text{train}}} \sum_{i=1}^{n_{\text{train}}} \sum_{j=1}^{m_{\text{irr}}} \left(Y_i(t_{ij}) - \hat{Y}_i(t_{ij}) \right)^2 \cdot \mathbb{1}(Y_i(t_{ij}) \text{ is observed}), \end{aligned} \quad (12)$$

which assures that the unobserved subject-time point pairs do not contribute to the gradient of the objective function.

Because of the common strategy used in FDA, which represents the infinite-dimensional curve using a finite basis set, we are able to construct the curve over the entire interval \mathcal{T} . This allows us to compute the predictive accuracy even if the test set contains observations at time points previously unseen, as long as they are within the interval of time points observed in the training set.

2.5 Roughness penalty

When predicting the functional response using NNs, we need to pay attention not only to the predictive accuracy but also the smoothness of the predicted trajectories, as it is standard in FDA. In practice, we observe that the performance of the NN mapping with basis expansion is influenced by the number of basis functions selected. Usually, the more basis functions we use to re-express the functional response, the higher accuracy can be gained when predicting the response at a set of given time points.

In FDA, it is common to set the number of basis functions to be less than the number of the observed time/location points with enough valid observations. However, in the cases where the observed time points are limited, we can consider increasing the number of basis functions to benefit the predictive accuracy. This benefit brought by more basis functions comes at the sacrifice of smoothness of the fitted functional trajectories, as they introduce in more variations. This is to be expected; because the prediction response for a specific time point $\hat{Y}_i(t)$ is the inner product of the basis coefficients and basis functions that are non-zero at t , then the more non-zero basis functions we have the more flexible that point estimate is. To control the smoothness of fitted curves without limiting the number of basis functions, we follow the common idea in FDA and propose to add some classic roughness penalty to the objective function of the NN. This simple addition ensures the smoothness of the predicted functional curves without putting any burden to the differentiable operation (since the roughness penalty terms we consider are also linearly related to the NN outputs).

There are multiple types of roughness penalties that can be applied to smooth the functional curves. In our implementation, we focus on two conventional and popular approaches: penalizing with the second derivative of the function, and penalizing directly the basis coefficients, to relieve the roughness concern with the NN-fitted functional response. Notice that we do not need

to define roughness penalties for NNBB and NNSS as the smoothing is done when fitting the curve a priori.

2.5.1 Penalize the second-order derivative of $Y(t)$

When the smoothness of fit becomes a concern and the roughness penalty turns out to be a necessity, the most common and popular method would be penalizing the α -th order derivative of the function $Y(t)$. The squared second derivative of a function $Y(t)$ at t reveals its curvature at t , therefore it is natural to measure the roughness of the $\hat{Y}_i(t)$ by taking the integrated square of its second derivative [28]. Adding this penalty term to the objective function of the NN leads to

$$L_{\text{pen}}(\eta) = \frac{1}{n_{\text{train}}} \sum_{i=1}^{n_{\text{train}}} \left(\sum_{j=1}^m (Y_i(t_{ij}) - \hat{Y}_i(t_{ij}))^2 + \lambda \int_{\mathcal{T}} \left(\frac{d^2 \hat{Y}_i(t)}{dt^2} \right)^2 dt \right). \quad (13)$$

The parameter λ acts as a smoothness controller for balancing the trade-off between fitting to the data and the variability of the predicted function [28]. When λ becomes larger, the more emphasis will be put on the smoothness of the fitted curves. On the contrary, for a smaller λ , the fitted curve tends to be more wiggly as there is less penalty placed on its roughness. The selection of the optimal λ can be achieved by using cross-validation, where either the subjects or the time points are randomly divided into sub-samples (in practice, we partition the time points).

Unfortunately, we cannot back-propagate the gradient through the integral of Eq. (13), and thus we will approximate the integral with a summation over the domain \mathcal{T} . Because we are able to generate $\hat{Y}_i(t)$, we can approximate this integral with as many points as we deemed necessary

$$L_{\text{pen}}(\eta) = \frac{1}{n_{\text{train}}} \sum_{i=1}^{n_{\text{train}}} \left(\sum_{j=1}^m (Y_i(t_{ij}) - \hat{Y}_i(t_{ij}))^2 + \frac{\lambda T}{Q-1} \sum_{q=2}^Q \left(\frac{d^2 \hat{Y}_i(t_{iq})}{dt_{iq}^2} \right)^2 \right), \quad (14)$$

where $t_{iq} = t_q$ for all i and $\{t_q\}_{q=1}^Q$ are equally spaced time points covering the entire domain \mathcal{T} with length T . Notice that we do not need to actually compute those second order derivatives, because

$$\begin{aligned} \hat{Y}_i(t) &= \sum_{k=1}^{K_b} \hat{c}_{ik} \theta_k(t) \\ \Rightarrow \frac{d^2 \hat{Y}_i(t)}{dt^2} &= \sum_{k=1}^{K_b} \hat{c}_{ik} \frac{d^2 \theta_k(t)}{dt^2}, \end{aligned} \quad (15)$$

then what we really need are the second order derivatives of the basis functions, which are much easier to compute and readily available in the **fd**a package. Therefore, we can back-propagate the gradient of the objective function of Eq. (14) with respect to η without any problems.

2.5.2 Penalize the basis coefficients \mathcal{C}

Penalizing directly on the basis coefficients was firstly introduced by Eilers & Marx [12] when adjacent B-splines are used to re-express the functional variable in a regression problem. Compared to penalizing the derivative, this idea reduces the dimensionality of the smoothing problem to K_b , the number of basis functions, instead of N , the number of observations [12].

The basic structure of this penalty is the difference of a set of consecutive basis coefficients $\Delta^2 c_k = c_k - 2c_{k-1} + c_{k-2}$. This difference has a strong connection with the second derivative of the fitted function, which is revealed by the simple formula for derivatives of B-splines given by De Boor [10]. The summation of the squared differences for all three consecutive coefficients $\{c_k, c_{k-1}, c_{k-2}\}, k = 3, \dots, K_b$ would be the main component of the penalty term, with the strength of the penalty further controlled by a tuning parameter λ . Like previously, the penalty term is added to the objective function of NN as:

$$L_{pen}(\eta) = \frac{1}{n_{train}} \sum_{i=1}^{n_{train}} \left(\sum_{j=1}^m (Y_i(t_{ij}) - \hat{Y}_i(t_{ij}))^2 + \lambda \sum_{k=3}^{K_b} (\Delta^2 c_{ik})^2 \right). \quad (16)$$

Similarly, the optimal hyper-parameter λ is selected using k -fold cross-validation.

Empirically, when fitting one of the B-spline models, especially the NNBR model, including a large number of basis functions, more than m for example, leads to a better predictive accuracy, but also results in a set of very rough predicted curves. When $K_b > m$ the predicted curves quickly become *wiggly*, and thus we would increase the number of basis such that $K_b \gg m$ and impose a roughness penalty at the same time to keep the curves smooth. In some cases, this leads to the top performer in terms of predictive accuracy. Unfortunately, using $K_b \gg m$ and a roughness penalty could make the complex hyper-parameter tuning process even more difficult. Nonetheless, we believe it is important to provide roughness penalties and justifications for those when providing a technique to fit functional data.

We implemented the four models and the two roughness penalties in R [26]. The functional component of our implementation relies on the **fd**a package and the NN component uses the R implementation of **Keras**. Our implementations of those models along the real data example are available online on the second author's GitHub page.

3 Computational complexity

An advantage of using an NN as the link function instead of a linear combination is that it might have a lower computational cost. More specifically, when using a standard training procedure for both models, the NN approach scales better with the number of predictors. Indeed, we usually find an exact solution for the FoS model which involves the inversion of a matrix, resulting in worse scaling with respect to the number of predictors. On the flip side, the algorithm used to train NN scales linearly with the number of predictors.

Start by looking at the FoS model, we rely on the formulation established in Ramsay & Silverman [28], chapter 14. This book contains various information about the computational details of the solutions of the FoS optimization problem under various parameterizations. However, its simplest form is quite similar to the least square solution of a traditional regression where:

$$\hat{\mathbf{B}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}, \quad (17)$$

which involves the inversion of a $P \times P$ matrix that requires a number of operations proportional to P^3 , say $O(P^3)$ in big O notation, using Gaussian elimination. The Strassen algorithm [34] manages to get that order down to $O(P^{2.8})$ and the more recent Coppersmith-Winograd algorithm [9] gets the order of matrix multiplication down to $O(P^{2.37})$ for a matrix of dimension $P \times P$.

The `fda` package uses the `solve()` function included in the default R language which in turn employs the linear algebra package (LAPACK) [1] that relies on Gaussian elimination to inverse matrices. Other than the matrix inversion, the matrix multiplication component of the solution involves a $P \times P$ matrix, a $P \times N$ matrix and a $N \times m$ matrix resulting in a number of operations of order $O(PN)$. Including the matrix multiplication part of the solution the whole FoS estimation requires a number of operations of order $O(P^3N)$. The more complete solutions that consider the smoothness of the predicted curves involve the inversion of a matrix of dimension $K_p P \times K_p P$ where K_p is the number of basis of the functional parameters $\beta(t)$. It means scaling polynomial with power 3 with the number of predictors P when using Gaussian elimination. Hence, we consider that the number of operations needed to obtain the exact solution to the FoS problem requires a number of operations of order $O(P^3N)$.

As far as NNs are concerned, according to Hastie & al. [18], the computational cost of training an NN is of order $O(PNn_wE)$ when back-propagation is adopted [32] to estimate the gradient and when a gradient-based approach is applied to fit the NN. In the formulation above, N is the number of observations, P is the number of predictors, n_w is the number of weights (hidden neurons) in the NN and E the number of training epochs.

Therefore, if we strictly focus on how the run time of an NN scales with N and P , we are looking at a linear scaling in both cases, $O(PN)$. Consequently, when comparing computational costs, the NN model has a pretty significant

advantage when it comes to its scaling with respect to the number of predictors. Because of this better scaling with respect to the number of predictors, we can claim that our proposed models are better equipped than the FoS regression to deal with data sets containing lots of predictors.

However, let us be a little more nuanced. It is clear that under some circumstances we can fit extremely large NNs which could lead to a slow fitting process. Additionally, the main difference between the traditional ways to fit these models is that we pursue an exact solution for the linear problem and on the contrary we utilize a gradient-based approach to find a solution in the case of NNs. As a result, for a linear model, such as FoS, we do have worse scaling with respect to P but we have guarantees that the optimal solution is reached. In theory, it would be possible to fit a linear model, such as FoS, with a gradient-based approach to improve how its computational cost scales with respect to P , though this is certainly not the standard optimization approach.

4 Real data application

We evaluate the predictive performance of the proposed methods, along with the conventional FoS model and the FAM model on the ASFR data set introduced in Section 1. We consider to compare the proposed models with the FAM model, $Y_i(t) = \sum_{p=1}^P f_p(X_{ip}, t) + \epsilon_i(t)$, because our methods can easily and naturally account for the nonlinear effects of all possible interactions simultaneously, while FAM is predominantly focused on learning the nonlinear relation between the functional response and each of the scalar predictors individually (it is computationally difficult and expensive for FAM to learn all possible interactions from a multi-dimensional aspect). For each of the models, we proceed with 20 repetitions of random subsampling validation: randomly dividing the data set into a training set and a test set, with 80% and 20% of the total samples assigned to them, respectively.

Following the same tuning strategy usually applied in classic NNs, the hyper-parameters of all models (except FAM) in comparison are firstly tuned using 5-fold cross-validation in order to fairly improve their performances during actual training. We start the hyper-parameter tuning with FoS, considering it involves only one hyper-parameter K_b , the number of basis functions. To reduce the computational time consumed by the tuning processes of the NN-based models, the optimal basis system selected for the FoS model is then used for the rest of the models. Afterwards, for each NN-based model, we perform a grid search on all network hyper-parameters simultaneously by taking a list of possible values for each of the hyper-parameters and running a 5-fold cross-validation for all combinations. An NN consisting of 2 hidden layers with 50 and 30 neurons, respectively, is selected. 99% proportion of variance explained is recommended by cross-validation as the threshold to truncate the FPCs scores in NNSS and NNSR. For NNBR with penalized objective function (NNBR(P)), we apply 10 basis functions and the second-order derivative roughness penalty, together with smoothing parameter $\lambda = 10^{-7}$, which is

similarly suggested by the 5-fold cross-validation from a set of possible values $\{10^{-i}\}_{i=1}^8$. Considering the number of basis functions is not dramatically large, the fitted curves are already quite smooth without any roughness penalty. Thus, a relatively small λ is acceptable as more emphasis ought to be placed on the fitting to data for an overall smallest $L_{pen}(\eta)$. A summary of the optimal hyper-parameters determined for each model is provided in Table S2 in the supplementary document. This hyper-parameter tuning approach is also applied in the simulations of Section 5. The FoS model used in this experiment is coming from the R-package `fd`, and FAM model is trained with the help of function `pffr()` in the `refund` package.

Table 1 Table of MSEPs of 20 random test sets for various models with ASFR data set.

Methods	Mean	Std. Dev.	<i>p</i> -value of <i>t</i> -test
FoS	1177.98	513.58	-
NNBB	1031.87	294.87	0.18
NNSS	992.50	296.68	0.01
NNBR	1060.50	210.63	0.20
NNBR(P)	1059.07	208.27	0.19
NNSR	1047.46	277.37	0.11
FAM	1317.33	919.62	0.28

Table 1 shows the predictive performance of each model using the mean squared error of prediction (MSEP) over the 20 replications. The predictive accuracy is measured by the MSEP averaged across the number of samples and the number of observed time points in the test set. A two-sided paired *t*-test is later conducted to compare the MSEPs of the 20 replicates of our models to the 20 MSEPs of the FoS model. It is clearly shown that our proposed models consistently outperform FoS and FAM models in predicting the fertility curves, with both smaller means and standard deviations (SDs) of the prediction errors of 20 random test sets, demonstrating their advantages in capturing both linear and nonlinear effects of all predictors on the functional response. NNSS model, as indicated in Table 1, has the best predictive performance, which interestingly implies that the FPC scores are more informative than the basis coefficients in representing $Y(t)$ with this data set, and consequently the effects of covariates on functional response are better described by the relations between the predictors and FPC scores instead of basis coefficients. In addition, the prediction error of NNBR optimized by the penalized objective function is lower than that of the non-penalized NNBR, indicating that including a relatively large number of basis functions, along with a roughness penalty term added to the objective function, can help improve the prediction without loss on the smoothness of the predicted curves.

As previously mentioned in Section 1, several covariates, such as female age and GDP per capita, are seemingly not linearly related to many of the basis coefficients representing the fertility trajectory. Therefore we focus on the basis

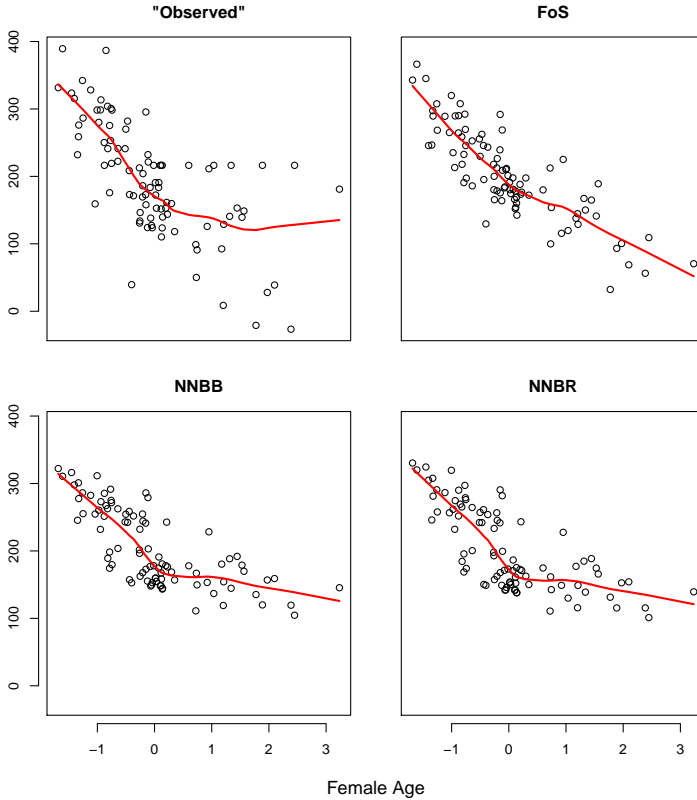


Fig. 3 A comparison of the relations between the covariate female age and the $Y(t)$ -estimated second basis coefficient ($\hat{c}_{2,Y(t)}$), FoS-predicted second basis coefficient ($\hat{c}_{2,\text{FoS}}$), NNBB-predicted second basis coefficient ($\hat{c}_{2,\text{NNBB}}$) and NNBR-predicted second basis coefficient ($\hat{c}_{2,\text{NNBR}}$), respectively.

coefficients \hat{C}_{FoS} , \hat{C}_{NNBB} and \hat{C}_{NNBR} predicted by FoS, NNBB and NNBR (they are the models using basis coefficients for regression), respectively. The relations between the predictor female age and each of the obtained second basis coefficients, including $Y(t)$ -estimated second basis coefficient $\hat{c}_{2,Y(t)}$ and the model-predicted ones, are displayed in Figure 3. We can see that the non-linear pattern for female age and the second basis coefficient is better recovered by NNBB and NNBR, while the FoS-predicted basis coefficient $\hat{c}_{2,\text{FoS}}$ is more likely linearly related to female age. Likewise, we select under-5 mortality out of those covariates that have highly likely linear associations with some of the basis coefficients, and Figure 4 reveals how under-5 mortality relates to different model-predict second basis coefficient \hat{c}_2 . It is noticed there are plainly linear trends between under-5 mortality and the second basis coefficients predicted by the three models mentioned, while under-5 mortality has been shown linearly related to the second basis coefficient estimated by $Y(t)$. It is not surprising that FoS reconstructs the most similar pattern to the one built on the

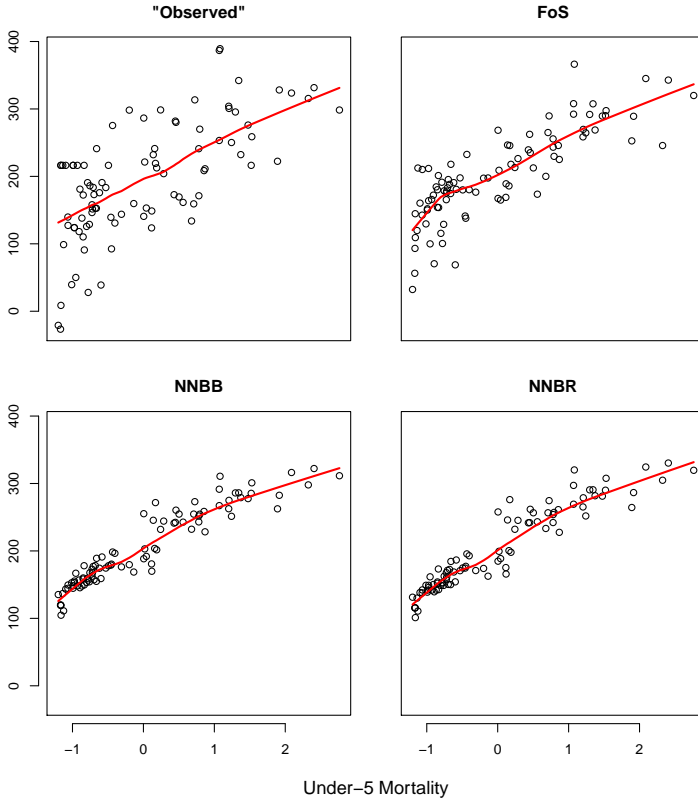


Fig. 4 A comparison of the relations between the covariate Under-5 Mortality and the $Y(t)$ -estimated second basis coefficient ($\hat{c}_{2,Y(t)}$), FoS-predicted second basis coefficient ($\hat{c}_{2,\text{FoS}}$), NNBB-predicted second basis coefficient ($\hat{c}_{2,\text{NNBB}}$) and NNBR-predicted second basis coefficient ($\hat{c}_{2,\text{NNBR}}$), respectively.

$Y(t)$ -estimated basis coefficient, but the proposed models also master the linear information when recovering the associations, which proves that the NN-based models have the ability to take good care of both the nonlinear and linear relations between the predictors and scalar representations simultaneously. Due to the relatively small number of observations, it is worth mentioning that the relations between covariates and basis coefficients may not be fully revealed and visual misunderstanding could occur.

5 Simulation studies

5.1 Generating data

We are interested in data sets that have multiple predictors and a potentially nonlinear relation between the scalar predictors and the functional response. We were inspired by the variable selection literature to design our data generators. More precisely, our data generators are similar to those proposed by

Wang & al. [39] and Barber & al. [2]. The basic concept is to build the functional response using a linear combination of a set of K random curves $\psi_k(t)$ and K associated coefficients $\zeta_k(\mathbf{X})$:

$$Y(t) = \sum_{k=1}^K \zeta_k(\mathbf{X})\psi_k(t). \quad (18)$$

and we are going to explore multiple approaches to generate the random curves $\psi_k(t)$'s and multiple ways to use the predictors \mathbf{X} to build coefficients $\zeta_k(\mathbf{X})$'s.

Firstly, we concentrate on how to generate the curves $\psi_k(t)$'s. Wang & al. [39] employ spline functions constructed with B-splines. Due to the requirement of their simulation, they use rather simple order 4 B-splines with 1 interior knot corresponding to 5 basis functions. To create a simple scenario for visualization purpose, we consider a configuration where a random curve $\psi_k(t)$ is set to be a single B-spline basis function $B_k(t)$:

$$\psi_k(t) = B_k(t). \quad (19)$$

Here, the applied B-spline basis system is of order of 4, while the number of interior knots are calculated using the number of predictors (number of predictors -4). The total number of basis functions, in this scenario, is equivalent to the number of scalar covariates. The main purpose of this configuration is to visualize if models are able to recover the coefficients $\zeta_k(\mathbf{X})$'s. Because in this example, a single covariate only affect a single B-spline basis function, we can fix $\zeta_k(\mathbf{X})$ to be a nonlinear function of the coefficients \mathbf{X} and to see if the trained model captured that nonlinear effect. We notably use that configuration for our **Design 1**.

The second configuration is a somewhat more realistic case where the covariates affect the response over the entire time interval. In this configuration, each random curve $\psi_k(t)$ will be built using B-splines of order 4 with 9 interiors knots corresponding to 13 basis function $\{B_l(t)\}_{l=1}^{13}$. The 13 associated coefficients β_l 's are randomly generated from a normal distribution. Thus, for the B-splines configuration we have:

$$\psi_k(t) = \sum_{l=1}^{13} \beta_{k,l} B_l(t). \quad (20)$$

This configuration is used in our **Design 2, 3 & 4**.

In terms of the coefficients $\zeta_k(\mathbf{X})$'s we will look at three configurations. Because NNs are known to capture nonlinear relationships, it is important to design nonlinear functions $\zeta_k(\cdot)$'s. First, for visualization, we apply a polynomial function to the continuous predictors. Specifically, a subset of continuous covariates, denoted by \mathbf{X}_{poly} , is randomly selected and each covariate in \mathbf{X}_{poly} is further transformed by a polynomial function with either second or third degree (half of the selected covariates are processed by quadratic functions and the rest by cubic functions), and accordingly $\zeta_k(\mathbf{X}) = \text{polynomial}(X_k)$

for $X_k \in \mathbf{X}_{\text{poly}}$, otherwise $\zeta_k(\mathbf{X}) = X_k$. This configuration is used in **Design 1 & 2** with a visualization example displayed in **Design 1**. Our second configuration is also nonlinear but this time it is a bit more complex. We define a 3-hidden-layer NN with sigmoid (which is nonlinear) activation functions and random weights. For this configuration, $\zeta_k(\mathbf{X})$ is the k -th output of an NN taking \mathbf{X} as its input. These coefficients are used for **Design 3**. Finally, we consider the case where $Y(t)$ is a linear combination of curves with \mathbf{X} being the coefficients directly. In other words, the k -th coefficient is simply the k -th predictor: $\zeta_k(\mathbf{X}) = X_k$. This is a scenario where we expect the FoS to outperform the models we propose, but we want to visit this example regardless. We call this the linear configuration and use it for **Design 4**.

Finally, we add a random noise, which is normally distributed with zero mean and a variance of 2 to every data point $Y_i(t_{ij})$.

5.2 Results

Four different simulation designs are considered to illustrate the advantages of the proposed methods. For each design, we generated data sets of 2000 observations and randomly sampled 1800 training observations and 200 testing observations. This random subsampling procedure was repeated 20 times. We opted to experiment with 20 scalar predictors, different random curve configurations and different relations between the coefficients and predictors, either linear or nonlinear, all with some random noise. To mimic the real-world scenario, the actual observations for $Y(t)$ were simulated discretely at 40 equally spaced time points $\{t_j\}_{j=1}^{40}$ that entirely cover $\mathcal{T} = [0, 1]$. Similar to the real application, we compare the proposed NN-based methods, including NNBB, NNSS, NNBR and NNSR, to the novel FoS model (we exclude FAM model because its training cost is tremendously high and based on the results of the initial replicates, it performs poorly compared to the other models). The predictive accuracy was evaluated using the MSEP on the test set. We also report the p -value of the two-sided paired t -test of the MSEPs of each of our methods to that of the FoS.

Design 1 (nonlinear scenario): We generate 20 random predictors ($K = 20$), where X_k 's are i.i.d. uniform random variables from $[a, b]$ with $a \in \{-4, -3, -2, -1, 0\}$, $b \in \{3, 4, 5, 6, 7\}$ for all k . We apply the polynomial transformation to 50% of the continuous variables by setting $\zeta_k(\mathbf{X}) = \text{polynomial}(X_k)$ with degree of 2 and 3 for $k = 8, 10, 12, 13, 14$ and $k = 1, 3, 4, 7, 9$, respectively, and $\zeta_k(\mathbf{X}) = X_k$ otherwise. The random curves $\psi_k(t)$'s, in this case, are set to be the B-spline basis functions as $\psi_k(t) = B_k(t)$ for all k , and accordingly the functional response is simulated as $Y(t) = \sum_{k=1}^{20} \zeta_k(\mathbf{X})B_k(t)$.

Table 2 summarizes the means and SDs of the MSEPs on the same testing observations achieved by different models in 20 replicates for **Design 1**. In the nonlinear scenario, we observe that the predictive performances of all the NN-based models surpass that of the FoS Model under 1% significance level. Two proposed models using basis coefficients as the output of NN exhibit the

Table 2 Table of MSEPs of 20 random test sets for various models with data generated by **Design 1**.

Methods	FoS	NNBB	NNSS	NNBR	NNSR
Mean	49.20	5.14	5.98	4.95	5.91
Std. Dev.	1.64	4.11	0.23	0.11	0.19
<i>p</i> -value	-	<2.2e-16	<2.2e-16	<2.2e-16	<2.2e-16

best-level performances in predicting the response curve, with smaller means of MSEPs compared to those of the NN-based models outputting FPC scores. This is expected because the functional response in this design is generated as a directly linear combination of the polynomial-transformed predictors and B-spline basis functions, and given this special setting, we are able to visualize the relations as some describable patterns for some selected coefficient-predictor pairs. This design highlights a situation where using FoS would be extremely problematic and where any of our proposed NN models would provide a significant improvement.

Next, we compare NNBB, NNBR with FoS for their abilities in recovering the underlying relationships for different coefficient-predictor pairs. The relations between a predictor and basis coefficients achieved by the aforementioned models are visualized through the scatter plots of each of the basis coefficients against the predictor. Remember that we designed this simulation specifically to visualize how NN models learn $\mathbf{C} = \text{NN}_\eta(\mathbf{X})$ and to compare it with the $\zeta_k(\mathbf{X})$ we designed.

We randomly select X_{12} , one of the polynomial-transformed covariates, and c_{12} , the basis coefficient corresponding to the 12-th B-spline basis function, and show their true relation, as $\zeta_{12}(\cdot)$, together with the relations reconstructed by the three mentioned approaches using scatter plots of true c_{12} and \hat{c}_{12} predicted by each of the models against X_{12} in Figure 5. Undoubtedly, NNBB and NNBR both precisely capture the \cap -shape between the true c_{12} and X_{12} , while FoS replaces the nonlinear trend with a downwards linear pattern. Being interested in a more complex truth, we also select X_4 , one of the covariates transformed with a cubic polynomial, coupled with c_4 , the basis coefficient corresponding to the 4-th B-spline, and provide the scatter plots illustrating the true and model-reconstructed relations between X_4 and c_4 in Figure 6. As expected, NNBB and NNBR perform well in overall recovering the nonlinear shape while NNBR behaves more accurately in capturing the local curvature occurring in the left component of the true pattern. It is not surprising that FoS again learns a linear relation because it is designed to only learn such relations. Likewise, the associations between X_5 and c_5 's by multiple models are revealed in Figure 7, but differently, the true relation for this X_5 - c_5 pair is linear. We can observe, NNBB and NNBR perform similarly as FoS in retrieving the linear pattern but with comparably thicker bandwidths, indicating that our models can also successfully detect the linear relation for some coefficient-predictor pairs but with higher variance. Figure 5, 6 and 7 highlight the utility of our methods




Fig. 5 Scatter plots of the true c_{12} (ζ_{12} as per the generator), FoS-predicted $\hat{c}_{12,\text{FoS}}$, NNBB-predicted $\hat{c}_{12,\text{NNBB}}$, and NNBR-predicted $\hat{c}_{12,\text{NNBR}}$ against X_{12} in **Design 1**, from left to right respectively.

in recovering the true underlying relations, especially the nonlinear relations between the basis coefficients and the predictors, and the success in learning the true relationships contributes to their fabulous performances on predicting the response trajectories. The ability to recover the true underlying nonlinear relations is the main benefit of our proposed approaches.

Design 2 (nonlinear scenario): 20 random predictors ($K = 20$) are generated, with X_k 's being binary variables for $k = 1, 3, 5, 7$, eight-level categorical variables for $k = 2, 4, 6$, and i.i.d. uniform random variables from $[a, b]$ with $a \in \{-4, -3, -2, -1, 0\}$, $b \in \{3, 4, 5, 6, 7\}$ for $k \geq 8$. Likewise, approximate 50% of the continuous variables are later transformed by polynomial functions with different degrees, where $\zeta_k(X) = \text{polynomial}(X_k)$ with the second and the third degree for $k = 14, 16, 17, 18$ and $k = 8, 10, 19, 20$, separately, and $\zeta_k(X) = X_k$ for the remainder. Then we construct the response curve as $Y(t) = \sum_{k=1}^{20} \zeta_k(\mathbf{X})\psi_k(t)$, coupled with $\psi_k(t) = \sum_{l=1}^{13} \beta_{k,l}B_l(t)$, where $\{B_l(t)\}_{l=1}^{13}$ are the B-spline basis functions with order 4 and $\beta_{k,l}$ are i.i.d. random variables following the normal distribution $\mathcal{N}(0, 4)$.




Fig. 6 Scatter plots of the true c_4 (ζ_4 as per the generator), FoS-predicted $\hat{c}_{4,\text{FoS}}$, NNBB-predicted $\hat{c}_{4,\text{NNBB}}$, and NNBR-predicted $\hat{c}_{4,\text{NNBR}}$ against X_4 in **Design 1**, from left to right respectively.

Table 3 Table of MSEPs of 20 random test sets for various models with data generated by **Design 2**.

Methods	FoS	NNBB	NNSS	NNBR	NNSR
Mean	4559.20	38.33	286.23	36.38	265.50
Std. Dev.	159.01	9.18	79.68	18.76	16.05
p -value	-	<2.2e-16	<2.2e-16	<2.2e-16	<2.2e-16

The means and SDs of the MSEPs on the testing observations for all methods in comparison can be found in Table 3. We observe that all our proposed methods remain superior to the FoS method in this setting. The NN-based models with basis-coefficient output continue to be the top-tier performers, especially the NNBR model trained using the response variable directly. The performance of the FoS is still significantly lower than the NN models and still has the highest variance. For both of these designs, some coefficient functions ζ_k 's have significant nonlinear curvature being polynomial functions of degree




Fig. 7 Scatter plots of the true c_5 (ζ_5 as per the generator), FoS-predicted $\hat{c}_{5,\text{FoS}}$, NNBB-predicted $\hat{c}_{5,\text{NNBB}}$, and NNBR-predicted $\hat{c}_{5,\text{NNBR}}$ against X_5 in **Design 1**, from left to right respectively.

2 or 3. Consequently, the dominance of the NN models is expected, even now with a more realistic scenario where a predictor affects the response on its entire domain \mathcal{T} .

Design 3 (nonlinear scenario): We continue to set $K = 20$, and X_k 's are generated as binary variables, eight-level categorical variables and i.i.d. uniform random variables from $[a, b]$ with $a \in \{-4, -3, -2, -1, 0\}$, $b \in \{3, 4, 5, 6, 7\}$ for different sets of k (same as **Design 2**). We trigger the nonlinear configuration by passing the predictors \mathbf{X} through the 3-hidden-layer NN introduced in Section 5.1, resulting in $\zeta(\mathbf{X}) = \text{NN}(\mathbf{X})$. Same as **Design 2**, we choose to use the random curves $\psi_k(t) = \sum_{l=1}^{13} \beta_{k,l} B_l(t)$, where $\{B_l(t)\}_{l=1}^{13}$ are the B-spline basis functions with order 4 and $\beta_{k,l} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 4)$. In consequence, the functional response is generated following $Y(t) = \sum_{k=1}^{20} \text{NN}(\mathbf{X}) \psi_k(t)$.

Table 4 reports the means and SDs of MSEPs on the test sets for all models with data generated by **Design 3**. It is shown that most models perform similarly on predicting the functional curve. The $\zeta_k(\mathbf{X})$'s produced by the NN generator are barely nonlinear, having only small curvature with respect to

Table 4 Table of MSEPs of 20 random test sets for various models with data generated by **Design 3**.

Methods	FoS	NNBB	NNSS	NNBR	NNSR
Mean	4.17	4.18	4.16	4.13	4.17
Std. Dev.	0.02	0.03	0.02	0.03	0.02
p -value	-	1.1e-02	6.2e-02	2.7e-08	1.8e-01

the predictors. This leads to a reduction in the performance gap between our NN models and the FoS model. However, the NNBR model comes ahead once again, being significantly superior to the FoS model.

Design 4 (linear scenario): In the last design, we consider setting up a scenario where the functional response $Y(t)$ is linearly associated to the scalar predictors. 20 predictors are generated in the manner that $\{X_k\}_{k=1}^{20}$ are binary variables, eight-level categorical variables and i.i.d. uniform random variables from $[a, b]$ with $a \in \{-4, -3, -2, -1, 0\}$, $b \in \{3, 4, 5, 6, 7\}$ for different subsets of $k = 1, 2, \dots, 20$ (same as **Design 2 & 3**). The linear setting is simply achieved with $\zeta_k(\mathbf{X}) = X_k$. We continue with the random curves $\psi_k(t) = \sum_{l=1}^{13} \beta_{k,l} B_l(t)$, where $\{B_l(t)\}_{l=1}^{13}$ are the order 4 B-spline basis functions, together with $\beta_{k,l} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 4)$. The functional response is generated as a linear combination of the random curves and predictors, following $Y(t) = \sum_{k=1}^{20} X_k \psi_k(t)$.

Table 5 Table of MSEPs of 20 random test sets for various models with data generated by **Design 4**.

Methods	FoS	NNBB	NNSS	NNBR	NNSR
Mean	4.01	4.19	6.66	4.07	5.94
Std. Dev.	0.05	0.06	0.15	0.05	0.05
p -value	-	3.2e-08	4.3e-13	7.8e-07	2.2e-09

Table 5 displays the results of all models under the linear setting described. In the linear scenario, we expect FoS to be the top performer and indeed, it has the lowest prediction error compared to all the proposed methods. This is reasonable as the $Y(t)$ is constructed linearly with respect to the predictors in this design. However, despite the proposed approaches being significantly worse, NNBB and NNBR remain competitive with strong performances closely following the one of FoS. From an absolute perspective, the gap between NNBR and FoS is much smaller in the scenario favouring FoS (**Design 4**), with FoS being marginally ahead, than it is in the scenario favouring NNBR (**Design 1 & 2**), with NNBR being dramatically ahead.

6 Conclusions and Discussion

In the article, we introduced a new solution for the regression of a functional response on scalar predictors, which consistently outperforms the current models when the relation between the functional response and the scale predictors is nonlinear. We designed and manipulated the standard feed-forward NN to produce two types of output, either basis coefficients or FPC scores, both of them being the traditional techniques to analyse functional data. The proposed modification to the objective function allows us to train the model directly with the functional response variable thus bypassing the necessity to firstly estimate those coefficients using conventional FDA approaches. The modified objective function enables the transformation of a scalar layer to a functional output by constructing functional curves using a linear operation, which ensures the applicability of back-propagation. In this way, an NN with the proposed functional output layer behaves similarly as a regular NN and allows the usage of typical cross-validation techniques for hyper-parameter tuning. Such modifications can be directly implemented to the output layer of any existing NN to produce a functional response, promoting the utilization of various deep learning techniques in a functional regression framework. We implemented all of our proposed models in a way they can be trained on both regularly and irregularly spaced domain. Additionally, we provided the necessary tools to control the smoothness of the predicted response curves by implementing two different roughness penalties. Furthermore, our family of models scales better with the number of predictors.

Based on a real data application, we demonstrated a case where the models we propose are superior to already established techniques such as FoS and FAM. Moreover, through several simulation studies, we not only showed the superior predictive power of our NN-based approaches, but also their strong ability to recover nonlinear relation between the predictors and the coefficients representing the response curves.

On the other hand, the developed methods rely on a large number of hyper-parameters, including the number of hidden layers, the number of neurons in each of the hidden layers, the number of basis functions (NN output size), the number of training epochs, etc. This is the main weakness of our approach since conducting a grid search on that space is particularly annoying and time-consuming. For both B-spline expansion models, the performance varied dramatically from one hyper-parameter configuration to another and this is certainly what we need to bring up. Comparatively, the FPCA-based models varied only slightly while changing the number of principal components. Given once we have the first few leading principal components, we can capture the vast majority of the variability between curves. The performances were also much more stable across various parameterizations, though those models were rarely top performers. In contrast, the FoS model is the easiest to fit among all the models applied, albeit it performed poorly in nonlinear situations.

Multiple directions of further research are considered at the moment. Our proposed models can be extended to predict multidimensional (mainly two-dimensional) functional response where $t \in \mathbb{R}^q$, $q > 1$. In this scenario, we can borrow the basis expansion or the FPCA technique to compress the multidimensional functional data to a vector of finite scalar basis coefficients or FPC scores [5, 11, 19, 22, 44]. Additionally, our current models rely on the existing NN architecture with a scalar output layer, which requires us to firstly project the functional response to some finite-dimensional coefficients and then feed the NN with the scalar representations obtained. Extending the NN to a more dynamic architecture allowing a functional output layer could be a more appropriate tool for such type of regression problems. This could be achieved by defining functional neurons and functional layers, some concepts we are currently exploring. Furthermore, a combination of our proposed NN models with existing literature that tackles the SoF problem can be explored to create a general and complete framework for using NN to analyze and solve regression problems for various forms of functional data.

Supplementary Materials

The supplementary document includes a list of all notations used in the manuscript and model configurations used in the application. We also provide the R codes and data required to reproduce the numerical results in the real application.

Acknowledgements

The authors would like to acknowledge the financial support of the Canadian Statistical Sciences Institute (CANSSI) and the Natural Sciences and Engineering Research Council of Canada (NSERC).

Declarations: The authors have no conflicts of interest to declare.

References

- [1] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.
- [2] Rina Foygel Barber, Matthew Reimherr, and Thomas Schill. The function-on-scalar lasso with applications to longitudinal gwas. *Electronic Journal of Statistics*, 11(1):1351–1389, 2017.
- [3] Cédric Beaulac, Jeffrey S Rosenthal, and David Hodgson. A deep latent-variable model application to select treatment intensity in survival analysis. *Proceedings of the Machine Learning for Health (ML4H)*

Workshop at NeurIPS 2018, 2018.

- [4] Cody Carroll, Alvaro Gajardo, Yaqing Chen, Xionghao Dai, Jianing Fan, Pantelis Z. Hadjipantelis, Kyunghye Han, Hao Ji, Hans-Georg Mueller, and Jane-Ling Wang. *fdapace: Functional Data Analysis and Empirical Dynamics*, 2020. R package version 0.5.5.
- [5] Lu-Hung Chen and Ci-Ren Jiang. Multi-dimensional functional principal component analysis. *Statistics and Computing*, 27(5):1181–1192, 2016.
- [6] Jeng-Min Chiou, Hans georg Müller, and Jane ling Wang. Functional response models. *Statistica Sinica*, pages 675–693, 2004.
- [7] Jeng-Min Chiou, Hans-Georg Müller, and Jane-Ling Wang. Functional quasi-likelihood regression models with smooth random effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):405–423, 2003.
- [8] François Chollet et al. Keras. <https://keras.io>, 2015.
- [9] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. In *Proceedings of the nineteenth annual ACM symposium on Theory of Computing*, pages 1–6, 1987.
- [10] Carl d. Boor. *A Practical Guide to Splines*. Springer Verlag, New York, 1978.
- [11] Paul Dierckx. Computation of least-squares spline approximations to data over incomplete grids. *Computers & Mathematics with Applications*, 10(3):283–289, 1984.
- [12] Paul H. C. Eilers and Brian D. Marx. Flexible smoothing with B-splines and penalties. *Statistical Science*, 11(2):89–121, 1996.
- [13] Ana-Maria Staicu Fabian Scheipl and Sonja Greven. Functional additive mixed models. *Journal of Computational and Graphical Statistics*, 24(2):477–501, 2015.
- [14] Jan Gertheiss Fabian Scheipl and Sonja Greven. Generalized functional additive mixed models. *Electronic Journal of Statistics*, 10(1):1455–1492, 2016.
- [15] Frédéric Ferraty and Philippe Vieu. *Nonparametric Functional Data Analysis: Theory and Practice*. Springer-Verlag, New York, 2006.
- [16] Felix Alexander Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural Computation*, 12:2451–2471, 2000.

- [17] Klaus Greff, Rupesh K. Srivastava, Jan Koutník, Bas R. Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10):2222–2232, 2017.
- [18] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 2. Springer, 2009.
- [19] J. G. Hayes and J. Halliday. The least-squares fitting of cubic spline surfaces to general data sets. *IMA Journal of Applied Mathematics*, 14(1):89–103, 1974.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9:1735–1780, 1997.
- [21] Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, and Michael S Lauer. Random survival forests. *The Annals of Applied Statistics*, 2(3):841–860, 2008.
- [22] Andrada Ivanescu. A note on bivariate smoothing for two-dimensional functional data. *International Journal of Statistics and Probability*, 2(2), 2013.
- [23] Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval Kluger. Deepsurv: personalized treatment recommender system using a cox proportional hazards deep neural network. *BMC Medical Research Methodology*, 18(1):1–12, 2018.
- [24] Paola Lecca. Machine learning for causal inference in biological networks: perspectives of this challenge. *Frontiers in Bioinformatics*, page 45, 2021.
- [25] Jeffrey S. Morris. Functional regression. *Annual Review of Statistics and Its Application*, 2(1):321–359, 2015.
- [26] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2020.
- [27] J. O. Ramsay, Spencer Graves, and Giles Hooker. *fda: Functional Data Analysis*, 2021. R package version 5.5.1.
- [28] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis (Second Edition)*. Springer, 2005.
- [29] Aniruddha Rajendra Rao, Qiyao Wang, Haiyan Wang, Hamed Khorasgani, and Chetan Gupta. Spatio-temporal functional neural networks, 2020.

- [30] F. Rossi, B. Conan-Guez, and F. Fleuret. Functional data analysis with multi layer perceptrons. In *Proceedings of the 2002 International Joint Conference on Neural Networks. IJCNN'02 (Cat. No.02CH37290)*, volume 3, pages 2843–2848, 2002.
- [31] Fabrice Rossi and Brieuc Conan-Guez. Functional multi-layer perceptron: a non-linear tool for functional data analysis. *Neural Networks*, 18(1):45–60, 2005.
- [32] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [33] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- [34] Volker Strassen et al. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4):354–356, 1969.
- [35] Arnab Maity Suchit Mehrotra. Simultaneous variable selection, clustering, and smoothing in function-on-scalar regression. *Canadian Journal of Statistics*, 50(1):180–199, 2022.
- [36] Barinder Thind, Kevin Multani, and Jiguo Cao. Neural networks as functional classifiers. *arXiv preprint arXiv:2010.04305*, 2020.
- [37] Barinder Thind, Kevin Multani, and Jiguo Cao. Deep learning with functional inputs. *Journal of Computational and Graphical Statistics*, 0:1–27, 2022.
- [38] Jane-Ling Wang, Jeng-Min Chiou, and Hans-Georg Müller. Functional data analysis. *Annual Review of Statistics and Its Application*, 3(1):257–295, 2016.
- [39] Lifeng Wang, Guang Chen, and Hongzhe Li. Group scad regression analysis for microarray time course gene expression data. *Bioinformatics*, 23(12):1486–1494, 2007.
- [40] Q. Wang, H. Wang, C. Gupta, A. Rao, and H. Khorasgani. A non-linear function-on-function model for regression with time series data. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 232–239. IEEE Computer Society, 2020.
- [41] Fang Yao, Hans-Georg Müller, and Jane-Ling Wang. Functional data analysis for sparse longitudinal data. *Journal of the American statistical association*, 100(470):577–590, 2005.

- [42] Junwen Yao, Jonas Mueller, and Jane-Ling Wang. Deep learning for functional data analysis with adaptive basis layers. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, volume 139, pages 11898–11908. PMLR, 2021.
- [43] Xiaoke Zhang and Jane-Ling Wang. Varying-coefficient additive models for functional data. *Biometrika*, 102(1):15–32, 12 2014.
- [44] Lan Zhou and Huijun Pan. Principal component analysis of two-dimensional functional data. *Journal of Computational and Graphical Statistics*, 23(3):779–801, 2014.
- [45] Xinliang Zhu, Jiawen Yao, and Junzhou Huang. Deep convolutional neural network for survival analysis with pathological images. In *2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 544–547. IEEE, 2016.