ORIGINAL ARTICLE

Estimation of the mean function of functional data via deep neural networks

Shuoyang Wang[1] | Guanqun Cao[1] | Zuofeng Shang[2] | for the Alzheimer's Disease Neuroimaging Initiative

[1]Department of Mathematics and Statistics, Auburn University, Auburn, Alabama, 36849, USA

[2]Department of Mathematics and Statistics, New Jersey Institute of Technology, Newark, New Jersey, 07102, USA

Correspondence
Guanqun Cao, Department of Mathematics and Statistics, Auburn University, 221 Parker Hall, Auburn, AL 36849, USA.
Email: gzc0009@auburn.edu

Funding information
NSF, Grant/Award Numbers: DMS 1736470, DMS 1764280, DMS 1821157

In this work, we propose a deep neural networks-based method to perform non-parametric regression for functional data. The proposed estimators are based on sparsely connected deep neural networks with rectifier linear unit (ReLU) activation function. We provide the convergence rate of the proposed deep neural networks estimator in terms of the empirical norm. Through Monte Carlo simulation studies, we examine the finite sample performance of the proposed method. Finally, the proposed method is applied to analyse positron emission tomography images of patients with Alzheimer's disease obtained from the Alzheimer Disease Neuroimaging Initiative database.

KEYWORDS

ADNI database, functional data analysis, neural networks, non-parametric regression, rate of convergence, ReLU activation function

1 | INTRODUCTION

Functional data refer to curves or functions; that is, the data for each variable are viewed as smooth curves, surfaces or hypersurfaces evaluated at a finite subset of some interval in 1D and 2D (e.g., some period of time, some range of pixels or voxels and so on). Functional data mean intrinsically infinite dimensional but are usually measured discretely. The high intrinsic dimensionality of these data poses challenges for both theory and computation. Functional data analysis (FDA) has been a topic of increasing interest in the statistics community for recent decades. Ramsay and Silverman (2005) and J. Wang et al. (2016) gave a comprehensive overview of FDA.

In FDA problems, estimation of mean functions is the fundamental first step; see Cardot (2000), Ferraty and Vieu (2006) and Rice and Wu (2001), for example. Various methods exist that allow to estimate the regression function non-parametrically. Rice and Wu (2001) adopted the mixed effect models where the mean function and the eigenfunctions were represented with B-splines and the spline coefficients were estimated by the expectation–maximization (EM) algorithm; Yao et al. (2005a) applied the local linear smoothers to estimate the mean and the covariance functions. Morris and Carroll (2006) generalized the linear mixed model to the functional mixed model framework, with model fitting done by using a Bayesian wavelet-based approach. In Cao et al. (2012), a polynomial spline estimator is proposed for the mean function of functional data together with a simultaneous confidence band. These non-parametric methods apply the pre-specified basis expansion, for example, polynomial splines, local linear smoothers, wavelet and so on, to fit the unknown mean function. The convergence rates achieve either the optimal non-parametric rate or parametric rates depending on how dense the observed points are for each subject.

Even though FDA has received considerable attention over the last decade, most of the existing approaches still focus on 1D functional data. The high intrinsic dimensionality of these data poses challenges for both theory and computation; these challenges vary with how the functional data were sampled. Hence, few are developed for general multidimensional functional data. Recently, several attempts have been made to extend these non-parametric methods for spatial and image data. Y. Wang et al. (2020) used bivariate splines over triangulations to handle an irregular domain of the images that is common in brain imaging studies. The proposed spline estimators of the mean functions are shown to be consistent and asymptotically

---

Abbreviations: FDA, functional data analysis; FPCA, functional principal component analysis; DNN, deep neural networks.

normal. However, the triangularized bivariate splines are designed for 2D functions only. Extending spline basis functions for general d-dimensional data observed on an irregular domain is very sophisticated and becomes extremely complex as d increases. B. Wang et al. (2014) proposed a regularized Haar wavelet-based approach for the analysis of 3D brain image data in the framework of functional linear regression models.

Another popular method is functional principal component analysis (FPCA) which is an extension of multivariate principal component analysis; see Hall et al. (2006) and Yao et al. (2005b), for example. Recently, there are a few studies on 2D FDA. Zhou and Pan (2014) proposed a smooth FPCA for 2D functions on irregular planar domains; their approach is based on a mixed effects model that specifies the principal component functions as bivariate splines on triangulations and the principal component scores as random effects. Lila et al. (2016) proposed an FPCA model that can handle real functions observable on a 2D manifold. Chen and Jiang (2017) extended it to analyse functional/longitudinal data observed on a general d-dimensional domain. When applying FPCA, how to choose the number of eigenfunctions is an important practical issue without a satisfactory theoretical solution. Presumably, the larger the number of eigenfunctions, the more flexible the approximation would be, and hence, the closer to the true curve. However, a large number of eigenfunctions always result in a complex model which introduces difficulties to follow-up analysis.

For many years, the use of neural networks has been one of the most promising approaches in connection with applications related to approximation and estimation of multivariate functions (see, e.g., Anthony & Bartlett, 2009; Ripley, 2014). Recently, the focus is on multilayer neural networks, which use many hidden layers, and the corresponding techniques are called *deep learning*. Under the non-parametric regression model, via sparsely connected deep neural networks, Schmidt-Hieber (2020) and Liu et al. (2021) showed that the L_2 risk of the least squares neural network regression estimator achieves the same minimax rate of convergence (up to a logarithmic factor) as proposed in Stone (1982). Furthermore, this neural network estimator does not suffer the curse of dimensionality which is a classical drawback in the traditional non-parametric regression framework. Bauer and Kohler (2019) have also obtained similar results under the deep learning framework via a different activation function. Liu et al. (2019) further removed the logarithmic factors to achieve the exact optimal non-parametric rate.

Although considerable advances have been achieved in deep learning research, from the statistical perspective, its application and theoretical research is still in its infancy stage (Fan et al., 2019). There are many technical challenges left for statisticians. For example, the availability of scalable computing and stochastic optimization techniques are challenging for developing statistical asymptotic properties. Rossi et al. (2005) extended the radial basis function networks and multilayer perceptron models to functional data inputs. However, neither the consistency of the proposed estimator nor deep neural network architectures have been considered in their work. Recently, there are some works proposed for deep learning algorithms from the statistical point of view (Thind et al., 2020a, 2020b). Motivated by these desiderata, the main goal of this article is to provide a novel method of FDA in the deep neural network framework.

The contributions of the present paper are fourfold. First, to our best knowledge, this is the first work on proposing a deep neural networks (DNN)-based estimator for FDA. An R package 'FDADNN' has been developed and is available from the GitHub website. Second, we develop the convergence rate (in empirical norm) of the proposed neural networks estimator. It is well known that when the observed points come from a hypercube, that is, $[0, 1]^d$, $d = 3$ for 3D imaging study, the non-parametric convergence rates are slower than the optimal non-parametric rate. This means that no statistical procedure can perfectly recover the signal pointwisely. However, by borrowing the advantage from the *deep learning* domain, the convergence rate of the proposed DNN estimator does not depend on the dimension d. Third, our proposed DNN estimator is unified for any dimensional functional data, which has broader and more flexible applications. Finally, we do not assume additional or complex structure for the true mean function, for example, additive models or single-index models. As in the deep learning domain, the true regression functions are assumed to be constructed in a modular form and the modularity of the system can be fairly complex, which resolve the misspecification issue.

Different from the existing neural network literature on non-parametric regression (Bauer & Kohler, 2019; Liu et al., 2021; Schmidt-Hieber, 2020) which only handle i.i.d. data, we focus on FDA where each subject is a random curve in a hypercube. Because of this special data structure, the major challenge becomes to deal with the correlation among the N evaluation points in the framework of neural networks, which has not been achieved in the existing works. It is not surprising that the convergence rate decreases with n (the number of subjects) as well as N.

The paper is structured as follows. Section 2 provides the model setting in FDA, introduces multilayer feed-forward artificial neural networks and discusses mathematical modelling. The implementation on hyperparameter selections also is included in Section 3. The theoretical properties of the proposed DNN estimator can be found in Section 4. In Section 5, the finite sample performance of proposed neural network estimator is presented. The proposed method is applied to the spatially normalized positron emission tomography (PET) data from Alzheimer Disease Neuroimaging Initiative (ADNI) in Section 6 and some concluding remarks are provided in Section 7. The proof of the main result together with additional discussion can be found in Appendix S1.

2 | MODELS AND DNN ESTIMATOR

2.1 | FDA model

Denote by Y_{ij} the jth observation of the random curve $\xi_i(\cdot)$ at grid points $\mathbf{X}_{ij}$, $1 \le i \le n$, $1 \le j \le N_i$. For simple notations, we examine the equally spaced design; in other words, $\mathbf{X}_{ij} = \mathbf{X}_j = j/N$. The main results can be extended to irregularly spaced design. Without loss of generality, let

$\mathbf{X}_j = (X_{j1},...,X_{jd}) \in [0,1]^d$. For the ith subject, its sample path $\{\mathbf{X}_j, Y_{ij}\}$ consists of the noisy realization of the Gaussian process $\xi_i(\mathbf{X})$ in the sense that $Y_{ij} = \xi_i(\mathbf{X}_j) + \epsilon_i(\mathbf{X}_j)$, and $\{\xi_i(\mathbf{X}), \mathbf{X} \in [0,1]^d\}$ are i.i.d. copies of the process $\{\xi(\mathbf{X}), \mathbf{X} \in [0,1]^d\}$ which is L^2; that is, $E\int_{[0,1]^d} \xi^2(\mathbf{X})d\mathbf{X} < +\infty$. The error term $\epsilon_i(\mathbf{X}_j)$ has mean zero and finite variance.

In this work, we consider the following classical FDA model:

$$\begin{aligned} Y_{ij} &= \xi_i(\mathbf{X}_j) + \epsilon_i(\mathbf{X}_j) \\ &= f_0(\mathbf{X}_j) + \eta(\mathbf{X}_j) + \epsilon_i(\mathbf{X}_j), i=1,2,...,n, j=1,2,...,N, \end{aligned} \tag{1}$$

where $f_0 : \mathbb{R}^d \to \mathbb{R}$, $E(Y_{ij}) = f_0(\mathbf{X}_j)$, $\eta(\cdot)$ is a Gaussian process characterizing individual curve variations from $f_0(\cdot)$ with mean zero and $\mathrm{Cov}(\eta(\mathbf{X}_j), \eta(\mathbf{X}_{j'})) := G(\mathbf{X}_j, \mathbf{X}_{j'})$. Let $\epsilon_i(\mathbf{X}_j) = \tau(\mathbf{X}_j)\varepsilon_{ij}$, where ε_{ij}s are independent normal random variables and $\tau(\mathbf{X})$ is the standard deviation function bounded above zero for any $\mathbf{X} \in [0,1]^d$. By Mercer's theorem, the covariance function $G(\mathbf{X}, \mathbf{X}')$ has the following spectrum decomposition:

$$G(\mathbf{X}, \mathbf{X}') = \sum_{k=1}^{\infty} \lambda_k \psi_k(\mathbf{X}) \psi_k(\mathbf{X}'),$$

where $\{\lambda_k\}_{k=1}^{\infty}$ and $\{\psi_k(\mathbf{X})\}_{k=1}^{\infty}$ are the eigenvalues and eigenfunctions of $G(\mathbf{X}, \mathbf{X}')$, and $\{\psi_k(\mathbf{X})\}_{k=1}^{\infty}$ are orthonormal bases in $L_2([0,1]^d)$.

2.2 | Deep neural networks

Before conducting the estimation of the mean function f_0 in (1) via the DNN, let us briefly introduce the necessary notations and terminologies used in the neural networks. From the high level, typical DNN use a composition of a series of simple non-linear functions to model non-linearity; that is,

$$\mathbf{h}_L = \mathbf{g}_L \circ \mathbf{g}_{L-1} \circ ... \circ \mathbf{g}_1(\mathbf{x}), \mathbf{x} \in \mathbb{R}^d,$$

where $\circ$ denotes a composition of two functions and L is called the number of hidden layers or depth of a DNN model. One can define $\mathbf{h}_l = \mathbf{g}_l(\mathbf{h}_{l-1})$ for each $1 \leq l \leq L$ recursively and $\mathbf{h}_0 = \mathbf{x}$. 'Deep' in deep neural networks refers to the use of multiple layers in the network. In the feed-forward neural network, the information moves in only one direction forward from the input layers, through the hidden layers and to the output node layers. In this kind of neural nets, there is a specific choice of $\mathbf{g}_l$: $\mathbf{g}_l(\mathbf{h}_{l-1}) = \sigma(\mathbf{W}_l \mathbf{h}_{l-1})$, $l = 1,...,L$, where $\mathbf{W}_l$ is a $p_l \times p_{l+1}$ weight matrix in the lth layer and $\mathbf{p} = (p_0,...,p_{L+1})$ is the width vector. The non-linear function σ is called the activation function. Here, we study the popular rectifier linear unit (ReLU) activation function applied element-wise $[\sigma(\mathbf{x})]_j = \max(x_j, 0)$, $j = 1,...,d$. For any vector $\mathbf{v} = (v_1,...,v_d) \in \mathbb{R}^d$, define the shifted activation function $\sigma_\mathbf{v} : \mathbb{R}^d \to \mathbb{R}^d$ as $[\sigma_\mathbf{v}(\mathbf{x})]_j = \sigma(x_j - v_j)$, $j = 1,...,d$. We call $\mathbf{v}$ the activation vector. The DNN model is then defined as $f : \mathbb{R}^{p_0} \to \mathbb{R}^{p_{L+1}}$,

$$f(\mathbf{x}) = \mathbf{W}_L \sigma_{V_L} \mathbf{W}_{L-1} \sigma_{V_{L-1}} ... \mathbf{W}_1 \sigma_{V_1} \mathbf{W}_0 \mathbf{x}. \tag{2}$$

To fit networks with data generated from the d-dimensional hypercube functional data model, we must have $p_0 = d$ and $p_{L+1} = 1$. Without loss of generality, we assume for any $f \in \mathcal{F}$, its empirical L_2 norm is bounded; that is, $\|f\|_N = \left(N^{-1}\sum_{j=1}^N f^2(\mathbf{x}_j)\right)^{1/2} \leq C_f < \infty$.

Although the depth and width of the neural nets can be extremely large, overfitting and computational burden are serious problems in such networks. To overcome these issues, the networks are modelled by assuming that each unit will be active only for a small fraction of the data to avoid overfitting. Smaller weights in a neural network can result in a model that is more stable and less likely to overfit the training dataset, in turn having better performance when making a prediction on new data (Srivastava et al., 2014). Therefore, we assume that there are only a few non-zero network parameters. Equivalently, we define the sparse neural networks and add constrains on the maximum-entry norm and non-zero entries of the weight matrix $\mathbf{W}_l$ and the activation vector $\mathbf{v}_l$. The sparse neural networks for our functional data model are given by

$$\begin{aligned} &\mathcal{F}(L, \mathbf{p}, s) \\ &= \left\{ f(\cdot) \text{ of the form } (2) : \max_{l=0,...,L} \|\mathbf{W}_l\|_\infty + \|\mathbf{v}_l\|_\infty \leq 1, \sum_{l=0}^L \|\mathbf{W}_l\|_0 + \|\mathbf{v}_l\|_0 \leq s \right\}, \end{aligned} \tag{3}$$

where $s > 0$, $\|\cdot\|_\infty$ denotes the maximum-entry norm and $\|\cdot\|_0$ denotes the number of non-zero entries. Let $\mathbf{v}_0$ be a zero vector for simplicity. The selecting procedures of unknown tuning parameters $(L, \mathbf{p}, s)$ shall be given in Section 3.

2.3 | Deep neural network estimator

In the functional data regression model, the common objective is to find an optimal estimator by the least squares loss function. In the neural network setting, this coincides with training neural networks by minimizing the empirical risk over all the training data. In particularly, given the networks in (3) and denote $\mathcal{F} = \mathcal{F}(L, \mathbf{p}, s)$, the proposed DNN estimator is defined as

$$\hat{f} = \underset{f \in \mathcal{F}}{\arg\min} \frac{1}{N} \sum_{j=1}^{N} \left\{ \bar{Y}_{\cdot j} - f(\mathbf{X}_j) \right\}^2, \tag{4}$$

where $\bar{Y}_{\cdot j} = \frac{1}{n} \sum_{j=1}^{n} Y_{ij}$. Different from classical non-parametric estimators, $\hat{f}$ has no analytical expression or basis expansion expression. Hence, to better understand the reasons that this DNN estimator has excellent performance, we first project f_0 onto the network space $\mathcal{F}$, namely, $f^* := \arg\min_{f \in \mathcal{F}} \|f_0 - f\|_{\infty}$. In other words, f^* is the best possible approximation of f_0 in $\mathcal{F}$. Note that

$$\frac{1}{N} \sum_{j=1}^{N} \left(\bar{Y}_{\cdot j} - \hat{f}(\mathbf{X}_j) \right)^2 \leq \frac{1}{N} \sum_{j=1}^{N} \left(\bar{Y}_{\cdot j} - f^*(\mathbf{X}_j) \right)^2,$$

which is equivalent to

$$\frac{1}{N} \sum_{j=1}^{N} \left(f_0(\mathbf{X}_j) - \hat{f}(\mathbf{X}_j) + \bar{\rho}_{\cdot j} \right)^2 \leq \frac{1}{N} \sum_{j=1}^{N} \left(f_0(\mathbf{X}_j) - f^*(\mathbf{X}_j) + \bar{\rho}_{\cdot j} \right)^2,$$

where $\bar{\rho}_{\cdot j} = \frac{1}{n} \sum_{i=1}^{n} \rho_{ij} = \eta(\mathbf{X}_j) + \frac{1}{n} \sum_{i=1}^{n} \epsilon_i(\mathbf{X}_j)$. Hence, we follow the conventional approximation–estimation decomposition (or bias–variance trade-off) to decompose the empirical norm

$$\|\hat{f} - f_0\|_N = \frac{1}{N} \sum_{j=1}^{N} \left(\hat{f}(\mathbf{X}_j) - f_0(\mathbf{X}_j) \right)^2$$

as

$$\|\hat{f} - f_0\|_N \leq \underbrace{\frac{1}{N} \sum_{j=1}^{N} \left(f^*(\mathbf{X}_j) - f_0(\mathbf{X}_j) \right)^2}_{\text{approximation error}} + \underbrace{\frac{2}{N} \sum_{j=1}^{N} \left(\hat{f}(\mathbf{X}_j) - f^*(\mathbf{X}_j) \right) \bar{\rho}_{\cdot j}}_{\text{estimation error}}. \tag{5}$$

The above equation indicates that the empirical norm of the estimator is bounded by two items. The first item is the approximation error and essentially determined by the distance between the network class $\mathcal{F}$ and the true function class f_0, which can be arbitrarily small according to Yarotsky (2017). From the statistical point of view, the second item is the estimation error and is a weighted average of a random process. It is affected by the parameters in $\mathcal{F}$, the true mean function class and the characteristic of the error terms.

3 | IMPLEMENTATION

In this section, we discuss the detailed computational procedure for the proposed DNN estimator in (4). The following proposed computational procedure can be easily realized via R package 'FDADNN' which is available at https://github.com/FDASTATAUBURN/FDADNN.

3.1 | Neural network's architecture selection

Tuning parameters are crucial as they control the overall behaviour of the proposed estimator and the learning process. In machine learning, those parameters are called network architecture parameters. A neural network's architecture can simply be defined as the number of layers L and the number of hidden neurons within these layers $\mathbf{p}$. In our considered sparse neural network space $\mathcal{F}$, the sparse parameter s should also be carefully selected. Note that in the practice, it's unrealistic to control the exact number of inactive nodes, so instead of using sparse parameter s, we add an L_1 penalty to control the number of active nodes in each layer during the optimization procedure. Denote ζ as the L_1 regularization factor. In the following, we utilize ζ to replace the sparse parameter s in the numerical analysis. The ultimate goal is to find an optimal combination of $(L, \mathbf{p}, \zeta)$

that minimizes a pre-defined loss function to give better results. There is a fairly large amount of literature discussing the optimization selection, such as grid search, random search and Bayesian optimization. Considering the computational efficiency and statistical properties, we recommend the following data-adaptive selection procedure in the practical application. The further justification of the optimization algorithm is beyond the scope of this work and shall be anther interesting and challenging topic for the future work.

We set the same neuron numbers for each layer for simplicity, that is, $\mathbf{p} = (p,...,p)$, and we follow the rule that p is increasing as n and N are increasing. We use K-fold cross-validation to choose (L,p,ζ); that is,

$$\left(L_{opt}, p_{opt}, \zeta_{opt}\right) = \arg\min_{(L,p,s)\,\in\,\Theta} \sum_{k=1}^{K} \sum_{j=1}^{N} \left(\widehat{Y}_{.j}^{(-k)}(L,p,\zeta) - \bar{Y}_{.j}^{(k)}\right)^2,$$

where Θ is an architecture parameter space which contains pre-selected choices of (L,p,ζ). Typically, $K = 5$ or 10. For the kth cross-validation, at the jth grid point, $\widehat{Y}_{.j}^{(-k)}(L,p,\zeta)$ denotes the estimated output given (L,p,ζ) and $\bar{Y}_{.j}^{(k)}$ is the average of observations.

3.2 | Training neural networks

The minimization in (4) is usually done via stochastic gradient descent (SGD). In a way similar to gradient descent, in each update, a small subsample called a *batch*, which is typically of size $B = 32$ to 512, is randomly drawn and the gradient calculation is only on the subsample instead of the whole training dataset. This saves considerably the computational cost in calculation of gradients. By the law of large numbers, this stochastic gradient should be close to the full sample one, albeit with some random fluctuations. We choose $B = 32$ or 64 batches depending on the performance of convergence. A pass of the whole training set is called an *epoch*. Typical choices of epochs are 200, 300 and 500. The number of epochs defines the number of times that the learning algorithm works through the entire training dataset. The learning rate, which controls how much parameters are changed in a single update in computing gradients in neural networks, is chosen as 0.001. In practice, given different data structures and minimization algorithms, the neural networks training could be challenging. The readers are referred to recent monographs (Emmert-Streib et al., 2020; Fan et al., 2019) for a general discussion of these numerical challenges.

There are certainly some challenges for SGD to train DNN. For example, albeit good theoretical guarantees for well-behaved problems, SGD might converge very slowly; the learning rates are difficult to tune (Allen-Zhu et al., 2019). To address these challenges, several variants gradient-based optimization algorithms are introduced, such as `Adam`, `RMSprop` and `Adadelta`. Instead of the classical SGD procedure, `Adam` is a method for efficient stochastic optimization that only requires first-order gradients with little memory requirement. Hence, it is well suited for problems when there are large sample sizes and parameters (Kingma & Ba, 2015). In our numerical studies, `Adam` provides the best results and is the most computationally efficient among these candidates. We recommend `Adam` in the real-life applications.

4 | THEORETICAL PROPERTIES OF THE DNN ESTIMATOR

In this section, we develop the convergence rate of the proposed DNN estimator in (4). For simple notations, log means the logarithmic function with base 2. For sequences $(a_n)_n$ and $(b_n)_n$, $a_n \asymp b_n$ means $a_n \leq c_1 b_n$ and $a_n \geq c_2 b_n$ where c_1 and c_2 are absolute constants for any n. Let $\mathbf{C}_N = [G(\mathbf{X}_j, \mathbf{X}_{j'})/N]_{j,j'=1}^N$ be the $N \times N$ kernel matrix corresponding to the covariance function G. We now introduce the main assumptions:(A1) The true regression function $f_0 \in \mathcal{G}(q, \mathbf{d}, \mathbf{t}, \boldsymbol{\beta}, \mathbf{K})$. (The definition of $\mathcal{G}(q, \mathbf{d}, \mathbf{t}, \boldsymbol{\beta}, \mathbf{K})$ is given in Appendix S1.)

(A2) The standard deviation function $\tau(\cdot)$ is bounded for any $\mathbf{x} \in [0,1]^d$.

(A3) The eigenvalues of $G(\cdot\,,\,\cdot)$ satisfy $\lambda_1 \geq \lambda_2 \geq ... \geq 0$ and $\sum_{k=1}^{\infty} \lambda_k < \infty$. Moreover, the maximal eigenvalue of the kernel matrix $\mathbf{C}_N$ satisfies $\lambda_{1,N} = O(N^{-\varrho})$ for some constant $\varrho \geq 0$.

(A4) The DNN estimator $\hat{f} \in \mathcal{F}(L, \mathbf{p}, s)$, where $L \asymp \log(nN^\varrho)$, $s \asymp (nN^\varrho)^{\frac{1}{\theta+1}}$, $\min_{l=1,...,L} p_l \asymp (nN^\varrho)^{\frac{1}{\theta+1}}$, for $\theta = \min_{i=0,...,q} \frac{2\beta_i^*}{t_i}$.

Assumption (A1) is a natural definition for neural networks, which is fairly flexible and many well-known function classes are contained in it. For example, the additive model $f_0(\mathbf{x}) = \sum_{i=1}^{d} f_i(x_i)$ can be written as a composition of two functions $f_0 = g_1 \circ g_0$, with $g_0(\mathbf{x}) = (f_1(x_1),...,f_d(x_d))^\top$ and $g_1(\mathbf{x}) = \sum_{j=1}^{d} x_j$, such that $g_0 : [0,1]^d \to \mathbb{R}^d$ and $g_1 : \mathbb{R}^d \to \mathbb{R}$. Here, $\mathbf{d} = (d,d,1)$ and $\mathbf{t} = (1,d)$. The generalized additive model $f_0(\mathbf{x}) = h\left(\sum_{i=1}^{d} f_i(x_i)\right)$ can be written as a composition of three functions $f_0 = g_2 \circ g_1 \circ g_0$, with g_0, g_1 described above and $g_2 = h$.

Assumption (A2) is a standard assumption for the variance of measurement errors, which requires the bounded variance of measurement errors over the whole space. This assumption has been widely used in functional data non-parametric regression literature; see Cao et al. (2012) and Yao et al. (2005a), for example. Assumption (A3) is a standard eigenvalue assumption for the Mercer kernel, and it is a widely used assumption for covariance functions in FDA literature; see Cao et al. (2016) and Li and Hsing (2010), for example. We also provide two examples to demonstrate that (A3) is a reasonable assumption in the supporting information. By Braun (2006), Assumption (A3) trivially holds for $\varrho = 0$ (see

Propositions A.1 and A.2 in the supporting information) and may even hold for some positive ϱ as revealed by Example 1 in Section A in the supporting information. Assumption (A4) depicts the architecture and parameters' setting in the network space.

We assume a natural compositional function class for the true mean function f_0:

$$f_0 = g_q \circ g_{q-1} \circ \cdots \circ g_1 \circ g_0,$$

where $g_i : [a_i, b_i]^{d_i} \to [a_{i+1}, b_{i+1}]^{d_{i+1}}$, $g_i = (g_{ij})_{j=1,\ldots,d_{i+1}}^\top$, $i = 1, \ldots, q$, with unknown parameters d_i and q.

The following theorem establishes the convergence rate of the DNN estimator $\hat{f}$ under the empirical norm. Its proof and some technical lemmas will be provided in the supporting information.

Theorem 1. *Under Assumptions (A1)–(A4), with probability greater than $\left(1 - \frac{2}{nN^\varrho}\right)^{\lceil \log(nN^\varrho)\rceil + 1} \to 1$, we have*

$$\|\hat{f} - f_0\|_N^2 \le c(nN^\varrho)^{-\frac{\theta}{\theta+1}} \log^6(nN^\varrho), \tag{6}$$

where $\varrho \ge 0$, $\theta = \min_{i=0,\ldots,q} \frac{2\beta_i^*}{t_i}$, c is a constant only depending on t, d, β which are defined in (A.1) in Appendix S1.

5 | SIMULATION

To illustrate how the introduced non-parametric regression estimators based on our proposed neural networks method behave in case of finite sample sizes, we conduct substantial simulations for both 2D and 3D functional data. We use the empirical L_2 risk $N^{-1} \sum_{j=1}^N \left(f_0(\mathbf{X}_j) - \hat{f}(\mathbf{X}_j)\right)^2$ as the criterion to evaluate the performance of estimators.

5.1 | 2D simulation

In this simulation, the 2D images are generated from the model:

$$Y_{ij} = f_0(\mathbf{X}_j) + \eta(\mathbf{X}_j) + \epsilon_i(\mathbf{X}_j), \tag{7}$$

where $\mathbf{X}_j = (X_{1j}, X_{2j}) = (j_1/N_2, j_2/N_2)$, $1 \le j_1, j_2 \le N_2$ are equally spaced grid points on the $[0,1]^2$ and $N_2^2 = N$. To demonstrate the practical performance of our theoretical results, we consider the following two mean functions:

- Case 1: $f_0(x_{1j}, x_{2j}) = \frac{-8}{1 + \exp\left(\cot(x_{1j}^2)\cos(2\pi x_{2j})\right)}$,

- Case 2: $f_0(x_{1j}, x_{2j}) = \log\left(\sin(2\pi x_{1j}) + 2|\tan(2\pi x_{2j})| + 2\right)$,

and the corresponding images are shown in the first row of Figure 1. To simulate the within-subject dependence for each subject i, we generate $\eta(\cdot)$ from a Gaussian process with mean 0 and the covariance function $G_0(\mathbf{x}_j, \mathbf{x}_{j'}) = \sum_{k=1}^2 \cos(2\pi(x_{kj} - x_{kj'}))$, $j, j' = 1, \ldots, N$. We generate $\epsilon_i(\mathbf{x}_j) = \varepsilon_{ij} \sim^{\text{i.i.d.}} \mathcal{N}(0, \sigma^2)$ for $i = 1, \ldots, n, j = 1, \ldots, N$. The noise level is set to be $\sigma = 1, 2$. We consider the sample size $n = 50, 100, 200$ and for each image, let $N_2 = 15$ or 25, which means for each 2D image, the number of observational points (pixels) is set to be $N = N_2^2 = 225$ or 625. The neural network (4) is trained through the optimizer Adam with architecture parameters (L, p, ζ) selected from $L \in \{3, 4\}$, $p \in \{100, 300, 500, 1000, 2000\}$, $\zeta \in \{10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}\}$ and the learning rate 0.001. The tenfold cross-validation method discussed in Section 3 is applied to select the optimal architecture parameters in each Monte Carlo simulation. Epochs are selected from 300 to 500, and the batch size is chosen as 32. We find that the convergence of algorithms is promising.

The alternative approach for 2D case we considered is a 2D regression spline method (bivariate spline). With regard to the variety of modifications of this approach known in the literature, we focus on the version for 2D FDA in Lai and Wang (2003). Let $\mathbf{B}^\top(\mathbf{x}) = \{B_m(\mathbf{x})\}_{m \in \mathcal{M}}$ be the set of bivariate Bernstein basis polynomials, where $\mathcal{M}$ stands for an index set of Bernstein basis polynomials. Then we can represent any bivariate function $f(\mathbf{x})$ by $f(\mathbf{x}) \approx \mathbf{B}^\top(\mathbf{x})\boldsymbol{\gamma}$ where $\boldsymbol{\gamma}^\top = (\gamma_m, m \in \mathcal{M})$ is the bivariate spline coefficient vector. The estimator $\hat{f}_{BS}$ is implemented by the R package BPST, which was developed by Lai and Wang (2003).

The second and third rows in Figure 1 depict the proposed neural network estimator $\hat{f}_{DNN}$ and the bivariate spline estimator $\hat{f}_{BS}$ when $n = 200$, $N = 625$ and $\sigma = 1$. Table 1 summarizes the empirical L_2 risk and standard deviation of estimators $\hat{f}_{DNN}$ and $\hat{f}_{BS}$ under 100 simulations for two different noise levels. From the above figure and table, one can see that our method and the bivariate spline method have fairly similar
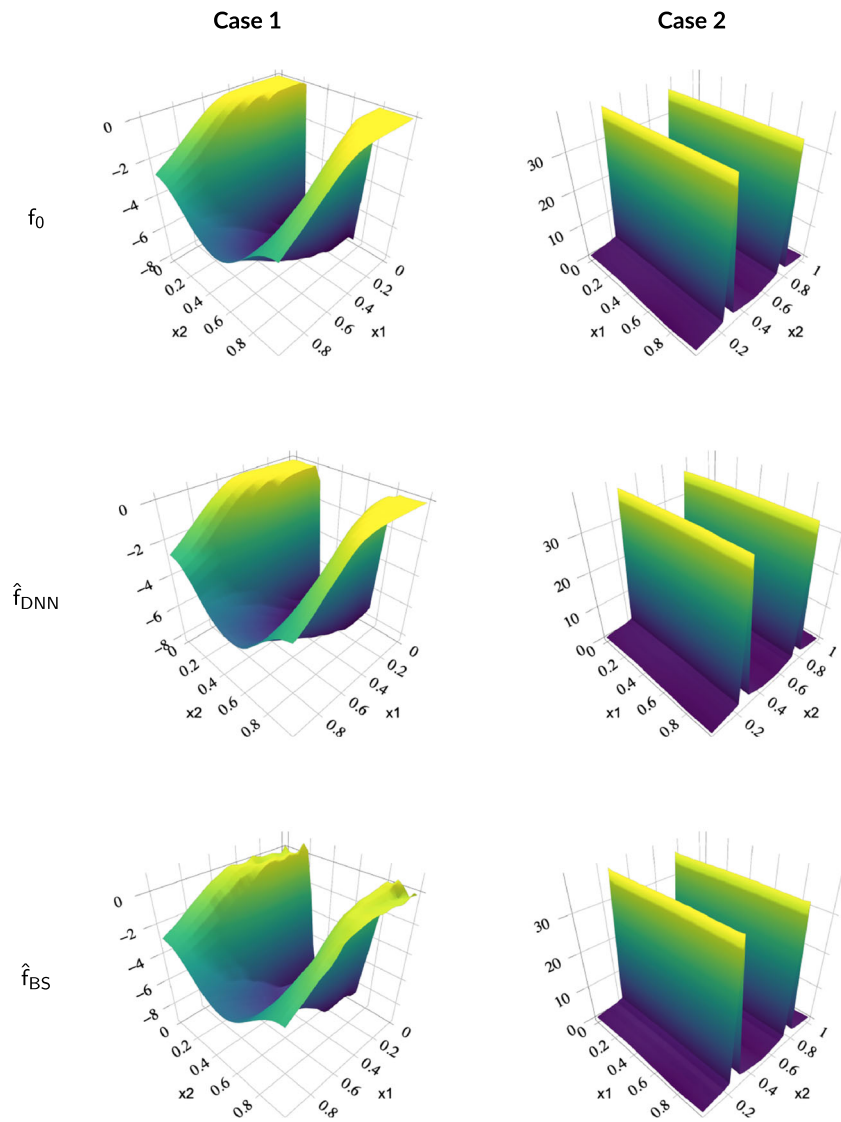
Case 1   Case 2

FIGURE 1 2D simulation. Left: from the top to bottom, they are the true function f_0 (Case 1) and its estimators $\hat{f}_{DNN}$ and $\hat{f}_{BS}$. Right: from the top to bottom, they are the true function f_0 (Case 2) and its estimators $\hat{f}_{DNN}$ and $\hat{f}_{BS}$ ($n = 200$, $N = 625$ and $\sigma = 1$)

estimation performances. As the bivariate spline estimator is able to achieve the optimal non-parametric convergence rate (Y. Wang et al., 2020), the comparable estimation results in Tables 1 and 2 also support the asymptotic convergence rate of our proposed estimator $\hat{f}_{DNN}$ in Theorem 1.

5.2 | 3D simulation

For 3D simulation, the images are generated from model (7) in 2D case. The true mean function is $f_0(x_{1j}, x_{2j}, x_{3j}) = \exp\left(\frac{1}{3}x_{1j} + \frac{1}{3}x_{2j} + \sqrt{x_{3j} + 0.1}\right)$, where $(x_{1j}, x_{2j}, x_{3j}) = \left(\frac{j_1}{N_3}, \frac{j_2}{N_3'}, \frac{j_3}{N_3''}\right)$, $1 \le j_1 \le N_3$, $1 \le j_2 \le N3'$, $1 \le j_3 \le N3''$ are equally spaced grid points in each dimension on $[0,1]^3$ and $N_3 N_3' N_3'' = N$. Here, we mimic the number of voxels of the real data, which usually have different values for N_3, $N3'$ and $N3''$. For each subject, the within-imaging dependence $\eta(\cdot)$ is generated from a Gaussian process with mean 0 and the covariance function $G_0(\mathbf{x}_j, \mathbf{x}_{j'}) = \sum_{k=1}^{3} \cos(2\pi(x_{kj} - x_{kj'}))$, $j, j' = 1, ..., N$. Measurement errors $\epsilon_i(\cdot)$ are generated the same as 2D case. We consider the sample size $n = 50, 100, 200$ and $N = 3000$ ($20 \times 15 \times 10$) and 4500 ($30 \times 15 \times 10$). Results of each setting are based on 100 simulations. The training of the neural networks architecture (L, p, s) follows the same procedures as in the 2D case. Architecture parameters (L, p, ζ) are selected from $L \in \{3, 4\}$, $p \in \{100, 300, 500, 1000, 2000, 5000\}$, $\zeta \in \{10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}\}$. The triangularized bivariate splines method proposed in Y. Wang et al. (2020) is designed for 2D functions only. Extending spline basis functions for 3D functional data is very sophisticated, and to our best knowledge, it is not available for 3D FDA yet. Hence, we only conduct 3D numerical analysis with our proposed DNN method. To exam the performance of the estimator $\hat{f}$, we also summarize the

TABLE 1 The average empirical L_2 risk and their standard deviations of the estimators of f_0 (Case 1) across 100 simulation runs (2D case)

$f_0(x_{1j}, x_{2j}) = \dfrac{-8}{1+\exp\left(\cot(x_{1j}^2)\cos(2\pi x_{2j})\right)}$

			DNN			Bivariate spline	
σ	N	n	L_2 risk	SD		L_2 risk	SD
1		50	0.1327	0.1905		0.6030	0.0418
	225	100	0.0797	0.1244		0.5757	0.0249
		200	0.0432	0.0574		0.5584	0.0120
		50	0.0770	0.0497		0.1497	0.0462
	625	100	0.0535	0.0368		0.1136	0.0214
		200	0.0352	0.0295		0.0987	0.0098
2		50	0.1880	0.1521		0.6564	0.1009
	225	100	0.0918	0.0793		0.6035	0.0619
		200	0.0593	0.0529		0.5765	0.0316
		50	0.1594	0.1555		0.2241	0.1218
	625	100	0.0862	0.0755		0.1430	0.0557
		200	0.0420	0.0412		0.1098	0.0232

TABLE 2 The average empirical L_2 risk and their standard deviations of the estimators of f_0 (Case 2) across 100 simulation runs (2D case)

$f_0(x_{1j}, x_{2j}) = \log\left(\sin(2\pi x_{1j}) + 2|\tan(2\pi x_{2j})| + 2\right)$

|			DNN			Bivariate spline	
σ	N	n	L_2 risk	SD		L_2 risk	SD
1		50	0.0731	0.0446		0.0804	0.0382
	225	100	0.0437	0.0249		0.0517	0.0186
		200	0.0254	0.0217		0.0351	0.0100
		50	0.0560	0.0206		0.0751	0.0351
	625	100	0.0351	0.0128		0.0541	0.0254
		200	0.0245	0.0085		0.0383	0.0110
2		50	0.1190	0.0975		0.1290	0.0950
	225	100	0.0829	0.0681		0.0931	0.0597
		200	0.0348	0.0276		0.0464	0.0251
		50	0.0573	0.0264		0.1213	0.0859
	625	100	0.0331	0.0132		0.0827	0.0630
		200	0.0139	0.0059		0.0502	0.0251

empirical L_2 risk and standard deviation of estimators $\hat{f}_{DNN}$ in Table 3. It is clear to find that the empirical risk decreases when sample sizes or observed voxel numbers increase for both noise levels, which supports our theoretical findings. The mean function f_0 and its DNN estimator are presented in Figure 2. It is easy to conclude that the DNN estimator follows the same pattern as the true mean function.

6 | ADNI PET DATA ANALYSIS

The dataset used in the preparation of this article was obtained from the ADNI database (adni.loni.usc.edu). The ADNI is a longitudinal multicentre study designed to develop clinical, imaging, genetic and biochemical biomarkers for the early detection and tracking of AD. From this database, we collect PET data from 79 patients in an AD group. This PET dataset has been spatially normalized and post-processed. These AD patients have three to six times of doctor visits, and we only select the PET scans obtained at the third visits. Patients' age ranges from 59 to 88, and the average age is 76.49. There are 33 females and 46 males among these 79 subjects. All the scans were reoriented into $79 \times 95 \times 69$ voxels, which

TABLE 3 The average empirical L_2 risk and their standard deviations of the DNN estimator of f_0 across 100 simulation runs (3D case)

σ	N	n	L_2 risk	SD
1		50	0.0028	0.0020
	3000	100	0.0011	0.0006
		200	0.0006	0.0004
		50	0.0007	0.0007
	4500	100	0.0005	0.0007
		200	0.0003	0.0004
2		50	0.0030	0.0024
	3000	100	0.0012	0.0007
		200	0.0007	0.0005
		50	0.0009	0.0007
	4500	100	0.0005	0.0008
		200	0.0003	0.0005

FIGURE 2 Two different angles (left and right panels) to view the true mean function and the DNN estimator in the 3D simulation case $(n = 200, N = 4500, \sigma = 1)$

means each patient has 69 sliced 2D images with 79×95 pixels. For the 2D case, it means each subject has $N = 7,505 = 79 \times 95$ observed pixels for each selected image slice. For the 3D case, the observed number of voxels for each patient's brain sample is $N = 79 \times 95 \times 69$, which is more than 0.5 million.

For the 2D case, we select the 20th, 40th and 60th slices from 69 slices for each patient. We first take average across 79 patients for each slice (the first row in Figure 3). Then, based on the averaged images, we obtain the proposed DNN estimators for each slice (the second row in Figure 3). We also recover the image in higher resolutions with 512×512 pixels, instead of the original 95×69 pixels for each slice (the third row in Figure 3). The neural network (4) is trained through the optimizer `Adam` with architecture parameters (L, p, ζ) selected from $L \in \{3, 4\}$, $p \in \{500, 1000\}$, $\zeta \in \{10^{-5}, 10^{-6}, 10^{-7}\}$ and the learning rate 0.001. The tenfold cross-validation method selects the optimal architecture parameters $L_{opt} = 3$, $p_{opt} = 1000$ and $\zeta_{opt} = 10^{-7}$. We used 300 to 500 epochs and 2 to 8 as batch sizes given different slices.

FIGURE 3 From top to bottom are averaged images $\{\bar{Y}_j\}_{j=1}^{7505}$, recovered images $\hat{f}(x_{1j}, x_{2j'})$, $j = 1, ..., 79$, $j' = 1, ..., 95$, recovered high-resolution (128×128) images $\hat{f}(x_{1j}, x_{2j})$, $j = 1, ..., 128$ and recovered images from 3D image. Left: the 20th slices; middle: the 40th slices; right: the 60th slices

In the 3D case, there are 79 patients, and totally $79 \times 95 \times 69$ voxels. Same as the 2D case, we first average the total 79 3D scans into one 3D scan and then perform neural networks to train the model based on the averaged 3D image. In the bottom row of Figure 3, we break down the recovered 3D image and show the recovered 20th, 40th and 60th slices. The neural network (4) is trained through the optimizer `Adam` with architecture parameters (L, p, ζ) selected from $L \in \{3, 4\}$, $p \in \{1000, 1500\}$, $\zeta \in \{10^{-5}, 10^{-6}, 10^{-7}\}$. The tenfold cross-validation method selects the optimal architecture parameters $L_{opt} = 4$, $p_{opt} = 1500$ and $\zeta_{opt} = 10^{-7}$. According to our numerical experience, we find 300 epochs and 64 batch sizes providing the reasonably well results.

In Figure 4, we also recover the image in higher resolutions with $128 \times 128 \times 128$ voxels, which means instead of the original $79 \times 95 \times 69$ voxels, we can provide the estimated image slices with higher resolutions (128×128 pixels, instead of the original 79×95 pixels) at finer grid points (128 points, instead of the original 69 points).

For model assessment, we split the data into to training and testing groups, with 60 patients and 19 patients separately. Table 4 summarizes the performance of the proposed estimator based on empirical L_2 risks. We conclude that the proposed functional regression model and the DNN estimator have satisfactory numerical results in this real application.

FIGURE 4 Recovered higher resolutions of selected nine slices in the 3D case

TABLE 4 The empirical L_2 risk ($\times 10^{-3}$) for ADNI data

	2D case			3D case
	20th	40th	60th	
Training	1.6	0.8	0.6	0.3
Testing	3.1	1.5	0.9	1.2

7 | DISCUSSION

In this work, we resolve the model misspecification issue in multidimensional FDA via the promising technique from the deep learning domain. By properly choosing network architecture, our estimator achieves the optimal non-parametric convergence rate in the empirical norm. To our best knowledge, this is the first piece of work in FDA, which yields attractive empirical convergence rate for multidimensional FDA and in the meanwhile is free from model misspecification. Numerical analysis demonstrates that our approach is useful in recovering the signal for imaging data. Some interesting future works may include the functional linear regression model and classification problems in the framework of DNN.

DATA AVAILABILITY STATEMENT

To show that Assumption (A3) is reasonable and can be satisfied, we provide two examples and their implementations in the supporting information. Some notations for the neural networks, technical lemmas and the detailed proof of Theorem 1 are also included.

Data used in the preparation of this article are openly from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 as a public–private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). For up-to-date information, see adni-info.org.

ORCID

Shuoyang Wang https://orcid.org/0000-0001-6963-6267
Guanqun Cao https://orcid.org/0000-0002-6454-3210
Zuofeng Shang https://orcid.org/0000-0003-1125-2302

REFERENCES

Allen-Zhu, Z., Li, Y., & Song, Z. (2019). A convergence theory for deep learning via over-parameterization. *Proceedings of the 36th International Conference on Machine Learning*, 97, 242–252.

Anthony, M., & Bartlett, P. (2009). *Neural network learning*. Cambridge: Cambridge University Press.

Bauer, B., & Kohler, M. (2019). On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47, 2261–2285.

Braun, M. L. (2006). Accurate error bounds for the eigenvalues of the kernel matrix. *Journal of Machine Learning Research*, 7, 2303–2328.

Cao, G., Wang, L., Li, Y., & Yang, L. (2016). Oracle-efficient confidence envelopes for covariance functions in dense functional data. *Statistica Sinica*, 26, 359–383.

Cao, G., Yang, L., & Todem, D. (2012). Simultaneous inference for the mean function of dense functional data. *Journal of Nonparametric Statistics*, 24, 359–377.

Cardot, H. (2000). Nonparametric estimation of smoothed principal components analysis of sampled noisy functions. *Journal of Nonparametric Statistics*, 12, 503–538.

Chen, L.-H., & Jiang, C.-R. (2017). Multi-dimensional functional principal component analysis. *Statistics and Computing*, 27(5), 1181–1192.

Emmert-Streib, F., Yang, Z., Feng, H., Tripathi, S., & Dehmer, M. (2020). An introductory review of deep learning for prediction models with big data. *Frontiers in Artificial Intelligence*, 3, 4.

Fan, J., Ma, C., & Zhong, Y. (2019). A selective overview of deep learning. arXiv preprint arXiv:1904.05526.

Ferraty, F., & Vieu, P. (2006). *Nonparametric functional data analysis: Theory and practice*, Springer Series in Statistics. New York: Springer.

Hall, P., Müller, H. G., & Wang, J. L. (2006). Properties of principal component methods for functional and longitudinal data analysis. *The Annals of Statistics*, 34, 1493–1517.

Kingma, D., & Ba, J. (2015). Adam: A method for stochastic optimization. In *In the 3rd International Conference on Learning Representations (ICLR)*.

Lai, M. J., & Wang, L. (2003). Bivariate penalized splines for regression. *Statistica Sinica*, 23, 1399–1417.

Li, Y., & Hsing, T. (2010). Uniform convergence rates for nonparametric regression and principal component analysis in functional/longitudinal data. *The Annals of Statistics*, 38, 3321–3351.

Lila, E., Aston, J. A. D., & Sangalli, L. (2016). Smooth principal component analysis over two-dimensional manifolds with an application to neuroimaging. *The Annals of Applied Statistics*, 10(4), 1854–1879.

Liu, R., Boukai, B., & Shang, Z. (2019). Optimal nonparametric inference via deep neural network. arXiv Preprint arXiv:1902.01687.

Liu, R., Shang, Z., & Cheng, G. (2021). On deep instrumental variables estimate. arXiv preprint arXiv:2004.14954.

Morris, J. S., & Carroll, R. J. (2006). Wavelet-based functional mixed models. *Journal of the Royal Statistical Society, Series B*, 68, 179–199.

Ramsay, J. O., & Silverman, B. W. (2005). *Functional data analysis, second edition*. New York: Springer Series in Statistics.

Rice, J. A., & Wu, C. O. (2001). Nonparametric mixed effects models for unequally sampled noisy curves. *Biometrics*, 57, 253–259.

Ripley, B. D. (2014). *Pattern recognition and neural networks*. Cambridge: Cambridge University Press.

Rossi, F., Delannay, N., Conan-Guez, B., & Verleysen, M. (2005). Representation of funcional data in neural networks. *Neurocomputing*, 64, 183–210.

Schmidt-Hieber, J. (2020). Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4), 1875–1897.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15, 1929–1958.

Stone, C. J. (1982). Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics, 10*(4), 1040–1053.

Thind, B., Multani, K., & Cao, J. (2020a). Deep learning with functional inputs. arXiv preprint arXiv:2006.09590.

Thind, B., Multani, K., & Cao, J. (2020b). Neural networks as functional classifiers. arXiv preprint arXiv:2010.04305.

Wang, B., Nan, B., Zhu, J., & Koeppe, R. (2014). Regularized 3D functional regression for brain image data via Haar wavelets. *The Annals of Applied Statistics, 8*, 1045–1064.

Wang, J. L., Chiou, J. M., & Müller, H. G. (2016). Functional data analysis. *Annual Review of Statistics and Its Application, 3*, 257–295.

Wang, Y., Wang, G., Wang, L., & Ogden, T. (2020). Simultaneous confidence corridors for mean functions in functional data analysis of imaging data. *Biometrics, 76*, 427–437.

Yao, F., Müller, H. G., & Wang, J. L. (2005a). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association, 100*, 577–590.

Yao, F., Müller, H. G., & Wang, J. L. (2005b). Functional linear regression analysis for longitudinal data. *The Annals of Statistics, 33*, 2873–2903.

Yarotsky, D. (2017). Error bounds for approximations with deep ReLU networks. *Neural Networks, 27*(94), 103–114.

Zhou, L., & Pan, H. (2014). Principal component analysis of two-dimensional functional data. *Journal of Computational and Graphical Statistics, 23*(3), 779–801.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.