

Robust deep neural network estimation for multi-dimensional functional data*

Shuoyang Wang and Guanqun Cao and for the Alzheimer’s Disease
Neuroimaging Initiative

Department of Biostatistics, Yale University, U.S.A.
e-mail: shuoyang.wang@yale.edu

Department of Mathematics and Statistics, Auburn University, U.S.A.
e-mail: gzc0009@auburn.edu

Abstract: In this paper, we propose a robust estimator for the location function from multi-dimensional functional data. The proposed estimators are based on the deep neural networks with ReLU activation function. At the meanwhile, the estimators are less susceptible to outlying observations and model-misspecification. For any multi-dimensional functional data, we provide the uniform convergence rates for the proposed robust deep neural networks estimators. Simulation studies illustrate the competitive performance of the robust deep neural network estimators on regular data and their superior performance on data that contain anomalies. The proposed method is also applied to analyze 2D and 3D images of patients with Alzheimer’s disease obtained from the Alzheimer Disease Neuroimaging Initiative database.

MSC2020 subject classifications: Primary 62G05, 62G08; secondary 62G35.

Keywords and phrases: Functional data analysis, deep neural networks, M-estimators, rate of convergence, ReLU activation function, ADNI databas.

Received March 2022.

Contents

1	Introduction	6462
2	The model and the robust deep neural network estimator	6464
	2.1 FDA model	6464
	2.2 Robust deep neural network estimator	6464
3	Theoretical properties of the RDNN estimator	6466
	3.1 Definitions and notations	6466
	3.2 Assumptions	6466
	3.3 Unified rate of convergence	6467
4	Implementation	6468
	4.1 Neural network’s architecture selection	6468
	4.2 Training neural networks	6468
5	Simulation	6469

*Wang’s and Cao’s research was partially supported by NSF award DMS 1736470. Cao’s research was also partially supported by the Simons Foundation under Grant #849413.

5.1	2D simulation	6469
5.2	3D simulation	6471
6	Real data analysis	6474
7	Discussion	6475
	Appendix	6475
	A.1 Technical lemmas and proofs	6475
	A.2 Proof of Theorem 3.1	6481
	Acknowledgments	6486
	References	6486

1. Introduction

We consider the problem of robust estimation of the location function from a collection of functional observations defined over \mathbb{R}^d ($d \geq 1$) a multi-dimensional domain. To be precise, let $\xi = \{\xi(\mathbf{X}) : \mathbf{X} \in \mathcal{I}\}$ be a compactly supported random field, i.e., a real-valued second-order stochastic process on a compact set $\mathcal{I} \subset \mathbb{R}^d$. Such data are nowadays commonly referred to as functional data. In many applications, data are collected over one-dimensional domains (i.e., $d = 1$) such as time-varying trajectories and relevant research has been enjoying considerable popularity. The readers are referred to some monographs [19, 25, 12, 9] for a comprehensive overview of functional data analysis (FDA). Thanks to the improved capabilities of data recording and storage, as well as advances in scientific computing and data science, many new forms of functional data have emerged. Instead of traditional unidimensional functional data, multi-dimensional functional data becomes increasingly common in various fields, such as geographical science and neuroscience. For example, for the early detection and tracking of Alzheimer’s disease, the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu) contains each individual’s 3D brain-scans. Despite the promising of multi-dimensional functional data, statistical methods for such data are limited, except for very few existing works, for example, [6, 29, 7].

A fundamental problem in FDA is the estimation of central tendency, yet most current estimation procedures either lack robustness with respect to the many kinds of anomalies one can encounter in the functional setting or only focus on the robustness for unidimensional scenario. The fact that robust estimation has not been widely investigated for multi-dimensional scenario is certainly not owing to lack of interesting applications, but to the greater technical difficulty to handle such loss function for multi-dimensional functional data and establish their theoretical properties.

To give some background on our proposed method for multi-dimensional functional data, we first review several relevant robust FDA methods that have been developed for analyzing unidimensional functional data. [2, 13] proposed robust estimators for the functional principal components by adapting the projection pursuit approach and based on MM estimation, respectively. [18] established a robust version of spline-based estimators for a linear functional regression model. [21] proposed a robust procedure based on convex and non-convex loss functions

in functional linear regression models. Recently, [14, 15] proposed robust estimators and associated simultaneous confidence bands for the mean function of functional data using least absolute deviation and M-estimation, respectively.

We notice that there are few existing works on robust methods for analyzing so-called two-way functional data which consist of a data matrix whose row and column domains are both structured, as when the data are time series collected at different locations in space. For example, [28] develop a robust regularized singular value decomposition method for analyzing such special type functional data. It is formulated as a penalized loss minimization problem and a pre-decided two-way roughness penalty function is used to ensure smoothness along each of the two functional domains. As this method is only designed for the special two-way functional data, it can not be adopted to the general multi-dimensional FDA directly. Furthermore, a lack of theoretical analysis provides inadequate assurance to robust methods practitioners.

To remedy these deficiencies, we introduce the first class of optimal robust location estimators based on the deep neural network (DNN) method. DNN is one of the most promising and vibrant areas in deep learning. DNN has been recently applied in various nonparametric regression problems, they have been shown to successfully overcome the curse of dimensionality in nonparametric regression; see [20, 3, 16, 17]. There are also some works proposed for deep learning algorithms for FDA from the statistical point of view [22]. Based on the sparsely connected DNN, [26] proposed a DNN estimator for the mean function from functional data based on the least squares neural network regression. However, none of them works on the robust statistics, not to mention the proven theoretical results for robust FDA.

The contributions of the present paper are three-fold. First, to the best of our knowledge, this is the first work on proposing DNN based robust estimator for FDA. We propose a broad class of M-type RDNN (robust DNN) estimators to estimate location functions for multi-dimensional functional data. Compared with the prior work [26], which can only handle the least-square estimation of the mean function of Gaussian functional data, this paper is able to deal with non-Gaussian and contaminated multi-dimensional functional data. Furthermore, the loss function considered here is an absolutely continuous convex function, which leads to a flexible choice of loss function and provide also quantile and expectile estimators. Second, RDNN estimators come with theoretical guarantees. In particular, we study the rate of convergence of the estimator under weak assumptions and show that the estimator is rate-optimal even for any d -dimensional functional data. By borrowing the strength from the DNN, the convergence rate of the proposed RDNN estimator does not depend on the dimension d if the true function is an intrinsically low-dimensional function. Finally, our analyses are fully nonparametric. At the meanwhile, RDNN estimator does not suffer the curse-of-dimensionality which is a classical drawback in the traditional nonparametric regression framework.

The paper is structured as follows. Section 2 provides the model setting in FDA and introduces multilayer feed-forward artificial neural networks and discusses mathematical modeling. The implementation on hyperparameter selec-

tions is also included in Section 2. The theoretical properties of the proposed RDNN estimator can be found in Section 3. Section 4 provides the detailed implementation on neural network's architecture selecting and training. In Section 5, the finite sample performance of proposed neural network estimator is demonstrated through simulation studies.. The proposed method is applied to the spatially normalized positron emission tomography (PET) data from ADNI in Section 6 and we make some concluding remarks in Section 7. Technical proofs are collected in the Appendix.

2. The model and the robust deep neural network estimator

2.1. FDA model

Let the process $\{\xi(\mathbf{X}), \mathbf{X} \in [0, 1]^d\}$ be L^2 , i.e. $E \int_{[0,1]^d} \xi^2(\mathbf{X}) d\mathbf{X} < \infty$. In the classical FDA setting, $d = 1$ refers to the index variable as time. When $d = 2, 3$, it could also be a spatial variable, such as in image or geoscience applications. We model the multi-dimensional functional data as noisy sampled points from a collection of trajectories that are assumed to be independent realizations of a smooth random function $\xi(\mathbf{X})$, with unknown mean function $f_0(\mathbf{X}) = E\{\xi(\mathbf{X})\}$. Without loss of generality, we assume $\|f_0\|_\infty$ is bounded by for some universal constant C . We consider a version of the model that incorporates uncorrelated measurement errors. Let ξ_1, \dots, ξ_n denote n independent and identically distributed (i.i.d.) copies of ξ at points $\mathbf{X} = (X_1, \dots, X_d)$. Our goal is to recover the mean function $f_0(\mathbf{X}_j)$ from the noisy observations of the discretized functional data:

$$Y_{ij} = \xi_i(\mathbf{X}_j) + e_i(\mathbf{X}_j), \quad i = 1, 2, \dots, n, j = 1, 2, \dots, N_i, \quad (2.1)$$

where $e_i(\mathbf{X}_j)$ are random noise variables. In [27, 5, 4], it is assumed that the noise variables $e_i(\mathbf{X}_j)$ are independent of the ξ_i and i.i.d. with zero mean and finite variance. However, we allow for correlated errors that are not necessarily independent of the functional curves.

In terms of mean-deviations, model (2.1) can be equivalently written as

$$Y_{ij} = f_0(\mathbf{X}_j) + \epsilon_i(\mathbf{X}_j), \quad i = 1, 2, \dots, n, j = 1, 2, \dots, N_i, \quad (2.2)$$

where $\epsilon_i(\mathbf{X}_j) = \xi_i(\mathbf{X}_j) - E\{\xi_i(\mathbf{X}_j)\} + e_i(\mathbf{X}_j)$ denotes the error process associated with the i -th response evaluated at \mathbf{X}_j . The problem is thus reformulated as a regression problem with repeated measurements and possibly correlated errors. In the following, for simple notations, we consider the equally number of observations for each subject design ($N_i \equiv N$). The main results can be easily extended to irregular number design.

2.2. Robust deep neural network estimator

We first briefly introduce the necessary notations and terminologies used in the neural networks. Popular choice of activation functions includes rectified linear

unit (ReLU), sigmoid, and tanh. In this article, we will mainly focus on neural networks with the ReLU activation function, i.e., $\sigma(x) = (x)_+$ for $x \in \mathbb{R}$. For any real vector $\mathbf{y} = (y_1, \dots, y_r)^\top$, define the shift activation function $\sigma(\mathbf{y}) = (\sigma(y_1), \dots, \sigma(y_r))^\top$. For an integer $L \geq 1$ and $\mathbf{p} = (p_0, p_1, \dots, p_L, p_{L+1}) \in \mathbb{N}^{L+2}$, let $\mathcal{F}(L, \mathbf{p})$ denote the class of DNN, with L hidden layers and p_l nodes on hidden layer l , for $l = 1, \dots, L$. We consider the feed-forward neural network class, and any $f \in \mathcal{F}(L, \mathbf{p})$ has a composition structure, i.e.,

$$f(\mathbf{x}) = \mathbf{W}_L \sigma(\mathbf{W}_{L-1} \dots \sigma(\mathbf{W}_1 \sigma(\mathbf{W}_0 \mathbf{x} + \mathbf{u}_0) + \mathbf{u}_1) + \dots + \mathbf{u}_{L-1}) + \mathbf{u}_L, \quad \mathbf{x} \in \mathbb{R}^d, \tag{2.3}$$

where $\mathbf{W}_l \in \mathbb{R}^{p_{l+1} \times p_l}$ are weight matrices and $\mathbf{u}_l \in \mathbb{R}^{p_l}$ are shift vectors, for $l = 1, \dots, L$. Owing to the large capacity of neural network class, it tends to overfit the training dataset easily. To avoid the overfitting and reduce the computational burden, we train the robust estimator using the following s-sparse ReLU DNN class:

$$\begin{aligned} & \mathcal{F}(L, \mathbf{p}, s) \\ = & \left\{ f \in \mathcal{F}(L, \mathbf{p}) : \sum_{l=0}^L \|\mathbf{W}_l\|_0 + \|\mathbf{u}_l\|_0 \leq s, \max_{l=0, \dots, L} \|\mathbf{W}_l\|_\infty \vee \|\mathbf{u}_l\|_\infty \leq 1, \right. \\ & \left. \|f\|_\infty \leq C \right\}, \end{aligned} \tag{2.4}$$

where $\|\cdot\|_\infty$ denotes the maximum-entry norm of a matrix/vector or supnorm of a function, $\|\cdot\|_0$ denotes the number of non-zero entries of a matrix or vector, $s > 0$ controls the number of nonzero weights and shift vectors, C controls the range of network f . The selecting procedures of unknown tuning parameters (L, \mathbf{p}, s) shall be given in Section 4. To simplify the notations, we write \mathcal{F} instead of $\mathcal{F}(L, \mathbf{p}, s)$ in the following.

In the regression model, the common objective is to find an optimal estimator by minimizing a loss function. In the DNN setting, this coincides with training neural networks by minimizing the empirical risk over all the training data. In particular, given the networks in (2.4), the proposed RDNN estimator is defined as

$$\hat{f} = \arg \min_{f \in \mathcal{F}} \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \rho(Y_{ij} - f(\mathbf{X}_j)), \tag{2.5}$$

where ρ is some convex nonnegative loss function satisfying $\rho(0) = 0$ and \mathcal{F} is some function class. This formulation is very general, allowing the flexibility in the choice of the loss function, so that better resistance towards outlying observations is achieved. One of the well-known examples of such loss functions is Huber's loss function given by $\rho_k(x) = x^2/2\mathbb{I}(|x| \leq k) + k(|x| - k/2)\mathbb{I}(|x| > k)$, where $\mathbb{I}(\cdot)$ is the indicator function, and $k > 0$ controls the blending of square and absolute losses. Furthermore, the symmetry of the loss function in (2.5) is not required, such versatile estimators may be readily incorporated into the present framework. Indeed, to estimate conditional quantiles, one would only need to select the loss function as $\rho(x) = x(\tau - \mathbb{I}(x < 0))$ for some $\tau \in (0, 1)$.

Finally, the asymptotic properties of quantile estimators are covered by the theory developed in Section 3.

3. Theoretical properties of the RDNN estimator

3.1. Definitions and notations

For any vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^N$, define the scaled N -inner product $\langle \mathbf{a}, \mathbf{b} \rangle_N = N^{-1} \sum_{j=1}^N a_j b_j$ and the associated N -norm $\|\mathbf{a}\|_N = \sqrt{N^{-1} \sum_{j=1}^N a_j^2}$. Define the ball of β -Hölder functions with radius K as

$$\mathcal{C}_d^\beta(D, K) = \left\{ f : D \subset \mathbb{R}^d \rightarrow \mathbb{R} : \sum_{\alpha: |\alpha| < \beta} \|\partial^\alpha f\|_\infty + \sum_{\alpha: |\alpha| = \lfloor \beta \rfloor} \sup_{\mathbf{x}, \mathbf{y} \in D, \mathbf{x} \neq \mathbf{y}} \frac{|\partial^\alpha f(\mathbf{x}) - \partial^\alpha f(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|_\infty^{\beta - \lfloor \beta \rfloor}} \leq K \right\},$$

where $\partial^\alpha = \partial^{\alpha_1} \dots \partial^{\alpha_d}$ with $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}^d$ and $|\alpha| := |\alpha|_1$.

We assume the true location function f_0 has the natural composition structure, i.e.,

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

where $g_\ell : [a_\ell, b_\ell]^{d_\ell} \rightarrow [a_{\ell+1}, b_{\ell+1}]^{d_{\ell+1}}$, $g_\ell = (g_{\ell j})_{j=1, \dots, d_{\ell+1}}^\top$, $\ell = 1, \dots, q$, with unknown parameters d_ℓ and q . We assume each $g_{\ell j}$ is β_ℓ -Hölder function with radius K_ℓ . Let t_ℓ be the maximal number of variables on which each of the $g_{\ell j}$ depends on t_ℓ , and $t_\ell \leq d_\ell$. Since $g_{\ell j}$ is also t_ℓ -variate, the true underlying function space becomes

$$\begin{aligned} &\mathcal{G}(q, \mathbf{d}, \mathbf{t}, \beta, \mathbf{K}) \\ &:= \left\{ f = g_q \circ \dots \circ g_0 : g_\ell = (g_{\ell j})_j : [a_\ell, b_\ell]^{d_\ell} \rightarrow [a_{\ell+1}, b_{\ell+1}]^{d_{\ell+1}}, \right. \\ &\quad \left. g_{\ell j} \in \mathcal{C}_{t_\ell}^{\beta_\ell}([a_\ell, b_\ell]^{t_\ell}, K_\ell), |a_\ell|, |b_\ell| \leq K_\ell \right\}, \end{aligned} \tag{3.1}$$

with $\mathbf{d} := (d_0, \dots, d_{q+1})$, $\mathbf{t} := (t_0, \dots, t_q)$, $\beta := (\beta_0, \dots, \beta_q)$, $\mathbf{K} := (K_0, \dots, K_q)$ and $\beta_\ell^* := \beta_\ell \prod_{k=\ell+1}^q (\beta_k \wedge 1)$.

3.2. Assumptions

In this section, we develop the convergence rate of the proposed RDNN estimator in (2.5). For simple notations, \log denotes 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 . For simplicity, let $\epsilon_{ij} = \epsilon_i(\mathbf{X}_j)$.

We now introduce the main assumptions:

- (A1) The true regression function $f_0 \in \mathcal{G}(q, \mathbf{d}, \mathbf{t}, \beta, \mathbf{K})$.
- (A2) The RDNN estimator $\hat{f} \in \mathcal{F}(L, \mathbf{p}, s)$, where $L \asymp \log(nN^\nu)$, $s \asymp (nN^\nu)^{\frac{1}{\theta+1}}$, $\min_{l=1, \dots, L} p_l \asymp (nN^\nu)^{\frac{1}{\theta+1}}$, for $\theta = \min_{\ell=0, \dots, q} \frac{2\beta_\ell^*}{t_\ell}$ and $\nu \geq 0$.

- (A3) The loss function $\rho(\cdot)$ is an absolutely continuous convex function on \mathbb{R} with derivative $\psi(\cdot) := \rho'(\cdot)$ existing almost everywhere.
- (A4) There exist finite constants κ and c_1 such that for all $x \in \mathbb{R}$ and $|x'| < \kappa$, $|\psi(x + x') - \psi(x)| < c_1$.
- (A5) There exist a finite constant c_2 such that $\sup_{j \leq N} \mathbb{E}\{|\psi(\epsilon_{1j} + u) - \psi(\epsilon_{1j})|^2\} < c_2|u|$, as $|u| \rightarrow 0$.
- (A6) $\sup_{j \leq N} \mathbb{E}\{(\psi(\epsilon_{1j}))^2\} = O(N^{-\nu})$, for some constant $\nu \geq 0$, and $\mathbb{E}\{\psi(\epsilon_{1j})\} = 0$. There exist finite constants δ_j , $j = 1, \dots, N$ such that $0 < \inf_{j \leq N} \delta_j \leq \sup_{j \leq N} \delta_j < \infty$ and $\sup_{j \leq N} |\mathbb{E}\{\psi(\epsilon_{1j} + u)\} - \delta_j u| = o(u)$, as $|u| \rightarrow 0$.

Assumption (A1) is a natural definition for the neural network, which is fairly flexible and many well known function classes are contained in it. For example, 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 $g_0(x_1, \dots, x_d) = (f_1(x_1), \dots, f_d(x_d))$, $g_1(x_1, \dots, x_d) = \sum_{i=1}^d x_i$, and $g_2(\cdot) = h(\cdot)$. Assumption (A2) depicts the architecture and parameters' setting in the network space. To use discontinuous score functions, Assumptions (A3)-(A6) impose some regularity on the error process and its finite-dimensional distributions. In particular, Assumption (A3) guarantees the existence of the solution of the optimization problem in (2.5). Most of the loss functions chosen in practice satisfy this condition, such as the Huber loss function. Assumptions (A4) and (A5) require boundedness and some regularity of the score function, which are standard conditions for M-estimation procedures for FDA, see the similar conditions required in [14]. For the first part of Assumption (A6), when considering the classical L_2 loss, it essentially makes sure the maximal value of the covariance function is finite and decreases when the number of measurements increases. They are standard regularity conditions for the covariance functions in FDA literature, see [4, 14, 26] for example. The second part of Assumption (A6) essentially requires that for any $j = 1, \dots, N$, function $h_j(u) = \mathbb{E}\{\psi(\epsilon_{1j} + u)\}$, is differentiable with strictly positive derivative at the origin. This is a necessary condition for the minimum to be well-separated in the limit. Assumption (A6) on the score function ψ is also standard conditions in M-estimation for functional data literature, see [14, 10]. It is also not stringent assumptions for errors, for example, ϵ_{ij} 's following a zero mean Gaussian process or mixture Normal-Cauchy distribution. We provide more detailed examples for ϵ_{ij} 's in Section 5.

3.3. Unified rate of convergence

The following theorem establishes the unified convergence rate of the RDNN estimator \hat{f} for any multi-dimensional functional data under the empirical norm. Its proof and some technical lemmas are provided in the Appendix.

Theorem 3.1. *Under Assumptions (A1)-(A6), we have*

$$\|\hat{f} - f_0\|_N^2 = O_p(nN^\nu)^{-\frac{\theta}{\theta+1}} \log^6(nN^\nu), \quad (3.2)$$

where $\nu \geq 0$, $\theta = \min_{\ell=0, \dots, q} \frac{2\beta_\ell^*}{t_\ell}$.

It is interesting to observe that our proposed RDNN estimator enjoys the same asymptotic as the least squares DNN estimator [26] does. Specifically, the convergence rate for M-type DNN estimator in the functional regression model depends on the smoothness, i.e., β_ℓ^* , and the intrinsic dimension, i.e., t_ℓ , of the true mean function f_0 , and the decay rate of the maximal value of the variance function $E\{(\psi(\epsilon_{1j}))^2\}$.

4. Implementation

Different from classical nonparametric estimators, \hat{f} has no analytical expression or basis expansion expression. The proposed robust estimator is constructed using the neural network class which is fully characterized by the architectures (L, \mathbf{p}, s) . We now provide the detailed implementation procedure for the proposed estimator in (2.5).

4.1. Neural network's architecture selection

In the DNNs' computations, tuning parameters are crucial as they control the overall behavior of the proposed estimator and the learning process. The tuning parameters are so-called network architecture parameters, which include the number of layers L , the number of hidden neurons within these layers \mathbf{p} , and sparse parameter s . There are fairly rich literature discussing the optimization selection, such as grid search, random search, and Bayesian optimization. Nevertheless, the selection of network architecture parameters has been rarely discussed. In practice, some model selection methods such as cross-validation may have good performances, but with huge computational burdens. For this reason, considering both the computational efficiency and the theoretical guarantee, we select architecture parameters based on the assumptions in Theorem 3.1. Particularly, we choose $L = \lceil 0.5 \log(nN^{1/2}) \rceil$, $p_l = \lceil 10n^{1/2}N^{1/4} \rceil$, $s = \lceil 5n^{1/2}N^{1/4} \rceil L$. We also select $\nu = \theta = 0.5$ in Assumption (A2). This choice of ν and θ includes a large scope of true function classes. Note that in our considered sparse neural network space \mathcal{F} , the sparse parameter s should be carefully selected. When designing the network architecture practically, the dropout rate is suggested as $\lceil 5n^{1/2}N^{1/4} \rceil (\lceil 10n^{1/2}N^{1/4} \rceil)^{-1}$ in each layer during the optimization procedure.

4.2. Training neural networks

The minimization in (2.5) is generally a computational cumbersome optimization problem owing to non-linearities and non-convexities. The most commonly used solution utilises stochastic gradient descent (SGD) to train a neural network. SGD uses a batch of a specific size, that is, a small subset of the data (typical size $B = 2^2$ to 2^{10}) is randomly drawn at each iteration of optimization to evaluate the gradient, to alleviate the computation hurdle. Our input size of network is nN , thus we choose relatively large batches B from 256 to 512 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 step of the derivative at each epoch is controlled by the learning rate which is 0.001. The readers are referred to recent monographs ([8]) 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 ([1]). To overcome these challenges, we use a variant gradient based optimization algorithm Adam. Different from the classical SGD procedures, 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 ([11]), and is widely used in network training for FDA, such as [26]. In our numerical studies, Adam provides the best results and is the most computationally efficient among other gradient based algorithms. In the real applications, we recommend Adam algorithm for finding RDNN estimators in (2.5).

5. Simulation

To illustrate the finite sample performance of the introduced RDNN estimators based on our proposed neural networks method, we conduct substantial simulations for both 2D and 3D functional data. All experiments are conducted in R. We summarize R codes and examples for the proposed RDNN algorithms on GitHub (<https://github.com/FDASTATAUBURN/RDNN>).

5.1. 2D simulation

The 2D functional data are generated from the model:

$$Y_{ij} = f_0(\mathbf{X}_j) + \epsilon_{ij}, \quad (5.1)$$

where the true mean function $f_0(\mathbf{x}_j) = -8 [1 + \exp \{ \cot(x_{1j}^2) \cos(2\pi x_{2j}) \}]^{-1}$, and $\mathbf{x}_j = (j_1/N_2, j_2/N_2)$, $1 \leq j_1, j_2 \leq N_2$, are the equally spaced grid points on $[0, 1]^2$, and $N_2^2 = N$. The error term is $\epsilon_{ij} = \eta(\mathbf{X}_j) + e_{ij}$, where $\eta(\cdot)$ is generated from a Gaussian process, with zero mean and covariance function $G_0(\mathbf{x}_j, \mathbf{x}_{j'}) = \sum_{k=1}^2 \cos(2\pi(x_{kj} - x_{kj'}))$, $j, j' = 1, \dots, N$. The measurement errors e_{ij} 's are i.i.d. standard normal random variables.

Under the proposed functional model (5.1), we introduce outlier hyper-surfaces to the generated functional sample by randomly contaminating a subset, R^o , of the original sample. The contamination proportion r is chosen to be 0, 0.1 and 0.2. The similar simulation setting has been considered in [14, 15]. We consider the following four types of outliers, i.e., two surface outliers and two heavy-tailed distributed outliers. They mimic the types of noised data encountered in the real dataset in Section 6.

Case 1: Stripe outliers To simulate outliers on a stripe in 2D regions, the contamination occurs on a line segment $a_0 \times \mathcal{I}$, that is,

$$Y_{ij^*}^o = Y_{ij^*} + \epsilon_{ij^*}^o, \quad i \in R^o, \quad j_1^*/N_2 = a_0, \quad j_2^*/N_2 \in \mathcal{I},$$

where $\epsilon_{ij^*}^o \sim U(10, 20)$. In this simulation, $a_0 = 0.2$, and we choose (i) $\mathcal{I} = \cup_{k=1}^5 [\frac{2k-2}{10}, \frac{2k-1}{10})$, and (ii) $\mathcal{I} = [0, 1]$.

Case 2: Square outliers To simulate outliers on a consecutive 2D region, the contamination occurs on a square $[a_0, a_1]^2$, that is,

$$Y_{ij^*}^o = Y_{ij^*} + \epsilon_{ij^*}^o, \quad i \in R^o, \quad (j_1^*/N_2, j_2^*/N_2) \in [a_0, a_1]^2$$

where $\epsilon_{ij^*}^o \sim U(10, 20)$. In the simulation, we choose (i) $[a_0, a_1]^2 = [0.1, 0.3]^2$, and (ii) $[a_0, a_1]^2 = [0.1, 0.5]^2$.

Case 3: Mixture Normal-Cauchy To simulate outliers with heavy-tailed distribution, the distribution of $\epsilon_{ij^*}^o$'s follow a mixture of a normal distribution $N(0, 1)$ and a Cauchy distribution with location 0 and scale 0.5. The mixing weights for Cauchy distribution are (i) 30%, and (ii) 50%.

Case 4: Mixture Normal-Slash Similar to previous case, but using a mixture of a normal distribution $N(0, 1)$ and a Slash distribution with location 0 and scale 0.5. The mixing weights for Slash distribution are (i) 30%, and (ii) 50%.

We consider sample size $n = 50, 100, 200$. For each image, let $N_2 = 10$, implicating the number of observational points (pixels) is set to be $N = N_2^2 = 100$. The network architecture is determined in a data driven way as suggested in Section 4.1, and we use Huber's loss function with tuning parameter 1 for RDNN estimator in (2.5). The results of each setting are based on 100 Monte Carlo simulations. Figures 1 presents heat maps of a typical set of the true mean function and abnormal observations, along with the estimations of RDNN and DNN estimators. From Table 1, we can see that when training the clean data, DNN method has comparable L_2 risks with RDNN estimators. These risks decrease as the sample size n increases. However, when contamination is involved, Table 2 shows that the risks of DNN estimators elevated drastically, while RDNN ones keep consistent results. In addition, although increasing either contamination rate r or contamination areas on a stripe raises the risks, we can see that RDNN estimators perform steady and remains relatively small L_2 risks even given 20% data contain anomalies. From Table 2, we can also see that when contamination occurs in a square region, the same trend is revealed, as previous discussion. It is worth mentioning that when $r = 0.2$, for the contaminated region $[0.1, 0.5]^2$, DNN estimators has extremely large risks, which are more than 10 times of ones of RDNN. Similar findings can be concluded from Table 3, where the random errors following non-Gaussian heavy-tail distributions. The RDNN estimator best mitigates the effect of this contamination relative to its competitors. Overall, the present simulation experiments suggest that RDNN perform well in clean data and safeguard against outlying observations either in the form of outlying surfaces or heavy-tailed measurement errors.

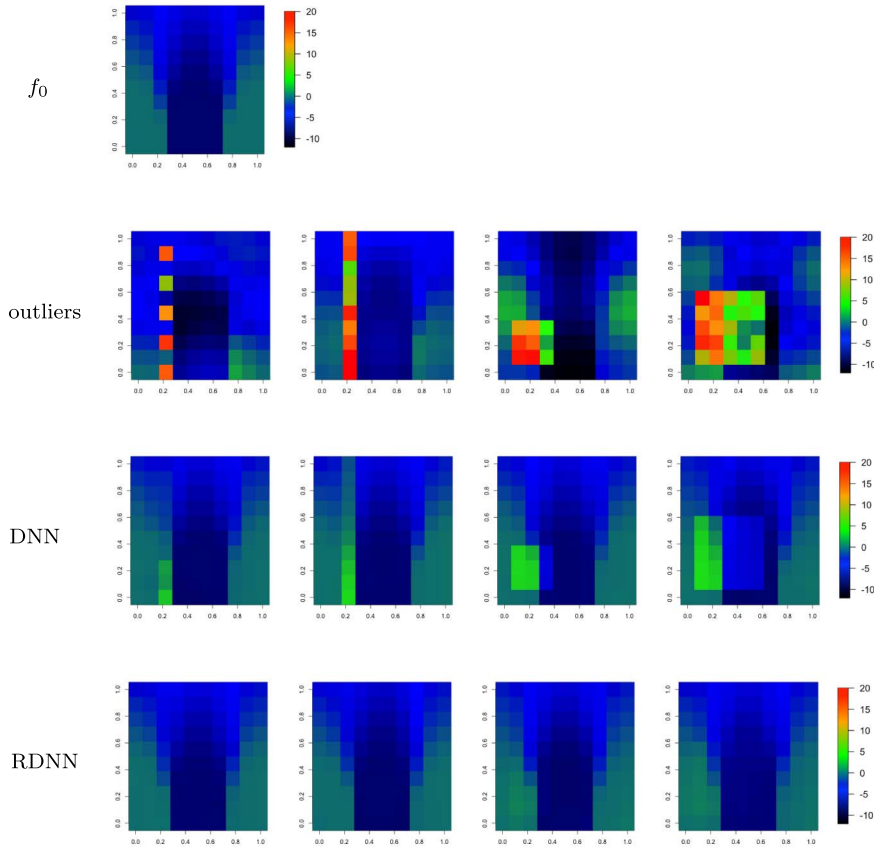


FIG 1. 2D simulation for Cases 1 and 2 on contaminated stripes and Cases 3 and 4 on square regions. The first row: true function f_0 ; The second row to forth row present the contaminated data Y^o , DNN estimations, RDNN estimations. From left to right, the observed data are generated from Case 1 (i) and (ii), Case 2 (i) and (ii).

TABLE 1
Empirical L_2 risk of 2D uncontaminated data with standard errors in brackets.

n	RDNN	DNN
50	0.114 (0.040)	0.125 (0.049)
100	0.059 (0.029)	0.055 (0.028)
200	0.034 (0.016)	0.031 (0.017)

5.2. 3D simulation

For 3D simulation, the functional data are generated from the model (5.1). The true mean function is $f_0(\mathbf{x}_j) = f_0(x_{1j}, x_{2j}, x_{3j}) = \exp\left(\frac{1}{3}x_{1j} + \frac{1}{3}x_{2j} + \sqrt{x_{3j} + 0.1}\right)$, where $\mathbf{x}_j = (x_{1j}, x_{2j}, x_{3j}) = (j_1/N_3, j_2/N_3, j_3/N_3)$, $1 \leq j_1, j_2, j_3 \leq N_3$, are

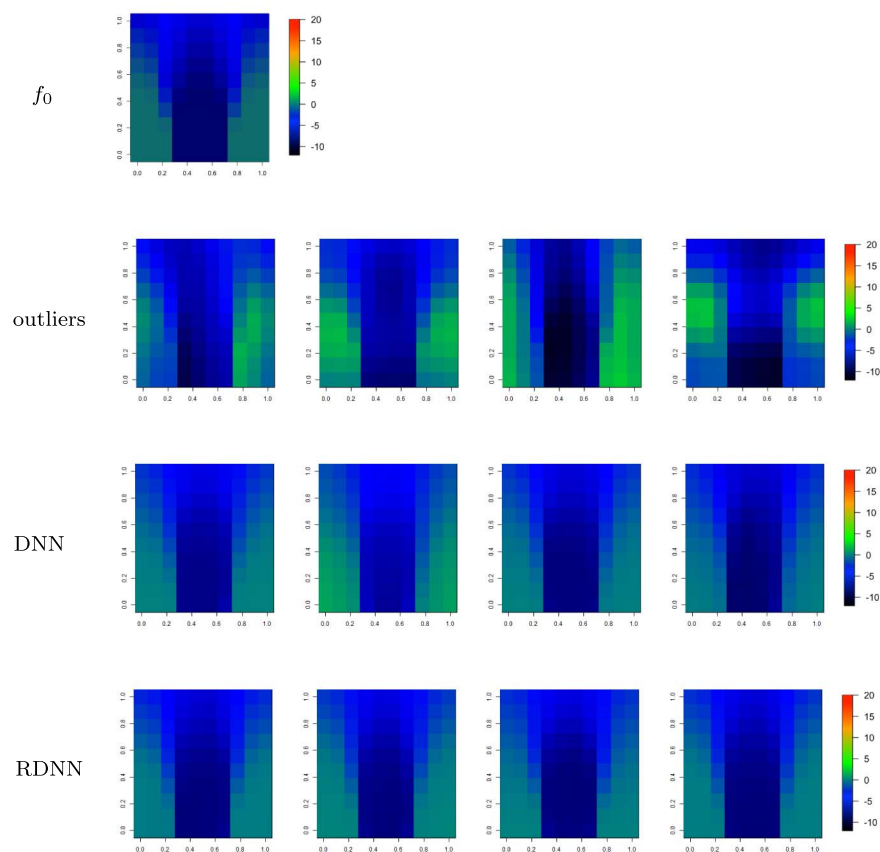


FIG 2. 2D simulation for mixed Cauchy and mixed Slash distribution. The first row: true function f_0 ; The second row to forth row present the contaminated data Y^o , DNN estimations, RDNN estimations. From left to right, the observed data are generated from Case 3 (i) and (ii), Case 4 (i) and (ii).

equally spaced grid points in $[0, 1]^3$ and $N = N_3^3 = 5^3$. Generate $\eta(\cdot)$ from a Gaussian process, with zero mean and covariance function $G_0(\mathbf{x}_j, \mathbf{x}_{j'}) = \sum_{k=1}^3 \cos(2\pi(x_{kj} - x_{kj'}))$, $j, j' = 1, \dots, N$, and the measurement errors e_{ij} 's are i.i.d. random variables generated from standard normal distribution. To contaminate the clean data, we apply the similar settings in Section 5.1.

Case 5: To simulate outliers on a consecutive 3D region, the contamination occurs on a square $[a_0, a_1]^3$, that is,

$$Y_{ij^*}^o = Y_{ij^*} + \epsilon_{ij^*}^o, \quad i \in R^o, \quad (j_1^*/N_3, j_2^*/N_3, j_3^*/N_3) \in [a_0, a_1]^3$$

where $\epsilon_{ij^*}^o \sim U(10, 20)$. In the simulation, we choose $[a_0, a_1]^3 = [0.10, 0.20]^3$

TABLE 2
Empirical L_2 risk of 2D contaminated data in Cases 1 and 2 with standard errors in brackets.

contaminated regions	n	$r = 0.1$		$r = 0.2$		
		RDNN	DNN	RDNN	DNN	
stripe	$\cup_{k=1}^5 \left[\frac{2k-2}{10}, \frac{2k-1}{10} \right)$	50	0.115 (0.048)	0.179 (0.078)	0.128 (0.055)	0.329 (0.095)
		100	0.055 (0.023)	0.102 (0.033)	0.065 (0.033)	0.252 (0.055)
		200	0.032 (0.015)	0.081 (0.024)	0.041 (0.018)	0.257 (0.043)
	[0, 1]	50	0.137 (0.066)	0.311 (0.081)	0.151 (0.051)	0.864 (0.164)
		100	0.064 (0.027)	0.240 (0.055)	0.088 (0.029)	0.842 (0.112)
		200	0.036 (0.015)	0.226 (0.032)	0.064 (0.023)	0.848 (0.084)
square	$[0.1, 0.3]^2$	50	0.118 (0.048)	0.260 (0.093)	0.154 (0.071)	0.664 (0.170)
		100	0.065 (0.033)	0.195 (0.059)	0.069 (0.025)	0.754 (0.107)
		200	0.038 (0.019)	0.195 (0.042)	0.054 (0.022)	0.752 (0.091)
	$[0.1, 0.5]^2$	50	0.151 (0.060)	0.657 (0.159)	0.234 (0.080)	2.014 (0.297)
		100	0.078 (0.042)	0.533 (0.111)	0.134 (0.063)	2.070 (0.248)
		200	0.046 (0.023)	0.550 (0.091)	0.108 (0.042)	2.172 (0.191)

TABLE 3
Empirical L_2 risk of 2D contaminated data in Cases 3 and 4 with standard errors in brackets.

error types	n	mixing weight = 30%		mixing weight = 50%	
		RDNN	DNN	RDNN	DNN
Cauchy	50	0.186 (0.069)	0.665 (0.959)	0.191 (0.069)	1.343 (3.193)
	100	0.097 (0.044)	0.289 (0.265)	0.104 (0.586)	0.586 (0.799)
	200	0.051 (0.029)	0.140 (0.175)	0.053 (0.024)	0.104 (0.066)
Slash	50	0.142 (0.065)	0.456 (0.686)	0.136 (0.071)	0.949 (2.022)
	100	0.074 (0.033)	0.419 (0.948)	0.071 (0.033)	0.822 (1.533)
	200	0.054 (0.027)	0.304 (0.617)	0.055 (0.029)	0.544 (1.004)

and $[0.10, 0.30]^3$ for different contamination proportions.

Case 6: *Mixture Normal–Cauchy* Similar to case 3, the distribution of ϵ_{ij}^o 's follow a mixture of a normal distribution $N(0, 1)$ and a Cauchy distribution with location 0 and scale 0.5. The mixing weights for Cauchy distribution are (i) 30%, and (ii) 50%.

Case 7: *Mixture Normal–Slash* Similar to case 4, the distribution of ϵ_{ij}^o 's follow a mixture of a normal distribution $N(0, 1)$ and a Slash distribution with location 0 and scale 0.5. The mixing weights for Slash distribution are (i) 30%, and (ii) 50%.

The results of each setting are based on 100 Monte Carlo simulations for sample sizes are 50, 100, and 200. For reference, Table 4 shows the average of empirical L_2 risks for clean data. We find that when data are clean, both of RDNN and DNN provide comparable estimations results, and the empirical risk decreases as the sample size increases. Tables 5 and 6 report the average of empirical L_2 risks for cases 6 and 7. As expected, non-robust DNN estimator has explosive risks, which are around three times of those for uncontaminated data. Similar to the 2D cases, either enlarging the contaminated region or the contamination proportion increases risk with DNN estimators. The precision of the RDNN estimator is kept at the same level as all outlier types and the clean

TABLE 4
Empirical L_2 risk of 3D uncontaminated data with standard errors in brackets.

n	RDNN	DNN
50	0.103 (0.050)	0.090 (0.045)
100	0.055 (0.033)	0.047 (0.023)
200	0.027 (0.013)	0.026 (0.018)

TABLE 5
Empirical L_2 risk of 3D contaminated data for cases 5 with standard errors in brackets.

Contaminated regions	n	$r = 0.1$		$r = 0.2$	
		RDNN	DNN	RDNN	DNN
[0.10, 0.20] ³	50	0.111 (0.049)	0.204 (0.066)	0.119 (0.052)	0.515 (0.107)
	100	0.056 (0.028)	0.155 (0.041)	0.078 (0.033)	0.539 (0.067)
	200	0.033 (0.018)	0.148 (0.029)	0.049 (0.017)	0.571 (0.058)
[0.10, 0.30] ³	50	0.118 (0.060)	0.463 (0.104)	0.173 (0.055)	1.598 (0.212)
	100	0.066 (0.032)	0.472 (0.092)	0.135 (0.052)	1.925 (0.160)
	200	0.042 (0.017)	0.478 (0.077)	0.103 (0.033)	1.942 (0.156)

TABLE 6
Empirical L_2 risk of 3D contaminated data for cases 6 and 7 with standard errors in brackets.

error types	n	mixing weight = 30%		mixing weight = 50%	
		RDNN	DNN	RDNN	DNN
Cauchy	50	0.130 (0.072)	0.526 (1.421)	0.134 (0.073)	0.804 (2.805)
	100	0.066 (0.035)	0.459 (0.953)	0.062 (0.036)	0.535 (1.295)
	200	0.043 (0.023)	0.163 (0.267)	0.045 (0.026)	0.418 (0.907)
Slash	50	0.128 (0.062)	0.753 (2.220)	0.125 (0.057)	0.787 (1.938)
	100	0.066 (0.042)	0.403 (0.887)	0.068 (0.049)	0.760 (1.458)
	200	0.049 (0.036)	0.321 (0.771)	0.047 (0.030)	0.587 (1.312)

dataset. This provides strong evidence that the proposed RDNN estimator is less sensitive to the presence of outliers, maintaining precision. In the worst case, the risks of RDNN estimator have increased no more than four times compared with the clean data scenarios, however, the non-robust ones have increased around 20 times.

6. Real data analysis

The dataset used in the preparation of this article were obtained from the ADNI database (adni.loni.usc.edu). The ADNI is a longitudinal multicenter 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 85 patients in AD group. This PET dataset has been spatially normalized and post-processed. These AD patients have three to six times doctor visits and we only select the PET scans obtained in the third visits. Patients' age ranges from 59 to 88 and average age is 76.49. All scans were reoriented into $79 \times 95 \times 69$ voxels, which means each patient has 69 sliced 2D images with 79×95 pixels. For 2D case, it indicates that each subject has $N = 7,505 = 79 \times 95$ observed

pixels for each selected image slice.

In this imaging dataset, we observe that there exists a few abnormal observations, which have different pattern from the majority of data. In Figure 3, the first row demonstrates the averaged images of the 20-th, 30-th, 40-th, and 50-th slices across all patients. In the second row, images are taken from different individuals, where extreme small values showing in certain regions, which lead to blur boundaries. For the 2D case, we select the 20-th, 30-th, 40-th and 50-th slices from 69 slices for each patient, and apply the proposed RDNN for each slice, respectively, with loss function $\rho_\tau(x) = x(\tau - \mathbb{I}(x < 0))$ with $\tau = 0.1, 0.3, 0.5, 0.7, 0.9$. The neural network (2.5) is trained through optimizer Adam with architecture parameters (L, p, s) selected as discussed in Section 4.1. We used 100 epochs and 128 as batch size given different data. Based on the images, we obtain the proposed RDNN estimators for each slice, and also recover the image with the original resolution 79×95 pixels and a higher resolution 128×128 . To visualize the estimates, Figures 4 provides the heat maps of the RDNN estimator of different quantiles for all four slices in 2D scenario, Figure 5 depicts the same estimates but with a finer resolution (128×128). For 3D scenario, we combine all the four slices together, hence, the 3D data totally contains $79 \times 95 \times 4$ voxels. We first obtain the RDNN estimators with the original resolution and recover them also in a higher resolution $128 \times 128 \times 4$. Figures 6 and 7 depict the RDNN estimators in the original resolution and higher resolution for each slice and quantile, respectively. The estimated quantiles serve to confirm the suspected multi-modality in this imaging data. According to the heat maps, in 20-th, 30-th, and 40-th slices, higher quantiles significantly differ from lower ones in that there are much larger value presenting in the bottom regions. In particular, for 50-th slice, higher quantiles can be easily distinguished from lower ones in terms of overall larger values and wider boundaries.

7. Discussion

In this work, we resolve the robust estimation for functional data on multi-dimensional domains via the promising technique from the deep learning. By properly choosing network architecture, our estimator achieves the optimal non-parametric convergence rate in empirical norm. To the best of our knowledge, the present work is the first work on multi-dimensional functional data robust estimation with theoretical justification for robust deep learning estimators. Numerical analysis demonstrates that our approach is useful in recovering the signal for imaging data given the existing of anomalies.

Appendix

A.1. Technical lemmas and proofs

We first introduce some lemmas which are useful in the proof of Theorem 3.1.




FIG 3. The first row are the averaged images for 20-th, 30-th, 40-th and 50-th slices across all patients. The rest are some abnormal data for each slices from some patients.

Lemma A.1 (Lemma A.2 of [26]). *For any composite regression function in the class $\mathcal{G}(q, \mathbf{d}, \mathbf{t}, \beta, \mathbf{K})$ and large M , such that $M \geq \max_{\ell=0, \dots, q} \{(\beta_\ell + 1)^{t_\ell} \vee (K_\ell + 1)e^{t_\ell}\}$, there exists an estimator $f^* \in \mathcal{F}$ such that*

$$\inf_{f^* \in \mathcal{F}} \|f^* - f_0\|_\infty^2 \leq C_I M^{-\theta}, \quad (\text{A.1})$$

where $\theta = \min_{\ell=0, \dots, q} \frac{2\beta_\ell^*}{t_\ell}$ and C_I is a constant. The parameters of the neural network f^* satisfy:

- (i) $L \leq D_0 D_1 \log M - 1$;
- (ii) $V^2 := \{\prod_{l=0}^{L+1} (p_l + 1)\}^2 \leq 4(d+1)^2 (12rM)^{2D_0 D_1 \log M}$;
- (iii) $s \leq D_0 D_2 M \log M - 1$,

where $D_0 = 2 \max_\ell \frac{t_\ell + \beta_\ell}{\beta_\ell}$, $D_1 = 4 \max_\ell (1 + \lceil \log(t_\ell \vee \beta_\ell) \rceil)(q+1)$, $D_2 = 290 \times \sum_{\ell=0}^q d_{\ell+1} (t_\ell + \beta_\ell + 1)^{3+t_\ell}$ and $r = \max_\ell d_{\ell+1} (t_\ell + \lceil \beta_\ell \rceil)$.

Remark 1. [26] focused on deriving Lemma A.1 which provides the approximation error, and the estimation error can be easily derived under the Gaussian distribution assumption. In contrast, given the general convex semicontinuous loss function and non-Gaussian distribution, we have to first show the minimization of convex semi-continuous functions in the unit balls in Lemma A.2 below. Then, we work on the estimation error rate $\|\hat{f} - f^*\|_N^2 = O_p(C_{nN})$ given the convex semi-continuous loss function, which is most challenging and breakthrough step. Lemma A.1 is the intermediate step but not the major technical contribution in the current one.

Lemma A.2. *Consider Euclidean space \mathbb{R}^N endowed with inner product $\langle \cdot, \cdot \rangle_N$ and associated norm $\|\cdot\|_N$. Let $L(\cdot) : \mathbb{R}^N \rightarrow \mathbb{R}^+$ be a convex lower semicontinuous function. If $L(\mathbf{0}) < \inf_{\|\mathbf{a}\|_N=1} L(\mathbf{a})$ for $\mathbf{a} \in \mathbb{R}^N$, then there exists an $\tilde{\mathbf{a}}$ (may*




FIG 4. 2D quantile estimators with 79×95 pixels. From the left to the right: the 20-th, 30-th, 40-th, and 50-th slices. From the top to the bottom: (10% , 30% , 50% , 70%, 90%)-quantiles.

not be unique), such that $\tilde{\mathbf{a}} = \inf_{\mathbf{a} \in \mathbb{R}^N} L(\mathbf{a})$, and the minimizer is attained in the unit ball $\{\mathbf{a} : \|\mathbf{a}\|_N \leq 1\}$.

Proof. Note that the given space is a Hilbert space. Let $s = \inf_{\|\mathbf{a}\|_N=1} L(\mathbf{a})$. Since L is bounded below by 0, we have $0 \leq s < \infty$. Denote $A = \{\mathbf{a} : L(\mathbf{a}) \leq s\}$ and $B = \{\mathbf{a} : \|\mathbf{a}\|_N \leq 1\}$.




FIG 5. 2D quantile estimators with 128×128 pixels. From the left to the right: the 20-th, 30-th, 40-th, and 50-th slices. From the top to the bottom: (10%, 30%, 50%, 70%, 90%)-quantiles.

We first prove $A \subset B$. Suppose $\exists \bar{\mathbf{a}} \in A$, and $\bar{\mathbf{a}} \notin B$, which implies $L(\bar{\mathbf{a}}) \leq s$ and $\|\bar{\mathbf{a}}\|_N > 1$. By convexity assumption,

$$L\left(\frac{\bar{\mathbf{a}}}{\|\bar{\mathbf{a}}\|_N}\right) = L\left(\frac{1}{\|\bar{\mathbf{a}}\|_N}\bar{\mathbf{a}} + \frac{\|\bar{\mathbf{a}}\|_N - 1}{\|\bar{\mathbf{a}}\|_N}\mathbf{0}\right) \leq \frac{1}{\|\bar{\mathbf{a}}\|_N}L(\bar{\mathbf{a}}) + \frac{\|\bar{\mathbf{a}}\|_N - 1}{\|\bar{\mathbf{a}}\|_N}L(\mathbf{0}) < s,$$




FIG 6. 3D quantile estimators with 79×95 pixels. From the left to the right: the 20-th, 30-th, 40-th, and 50-th slices. From the top to the bottom: (10% , 30% , 50% , 70%, 90%)-quantiles.

where the last inequality is obtained by assumption $L(\mathbf{0}) < s$. Since $\frac{\bar{\mathbf{a}}}{\|\bar{\mathbf{a}}\|_N}$ is on the unit ball, this leads to the contradiction, which implies $A \subset B$.

We then prove A is a compact set. In particular, we only need to prove the boundedness and closedness in the Euclidean space. Indeed, boundedness is implied by the fact that $A \subset B$. For any sequence $\{\mathbf{a}_m\}_{m \geq 1}$, define $\mathbf{a}^* =$




FIG 7. 3D quantile estimators with 128×128 pixels. From the left to the right: the 20-th, 30-th, 40-th, and 50-th slices. From the top to the bottom: (10% , 30% , 50% , 70% , 90%)-quantiles.

$\lim_{m \rightarrow \infty} \mathbf{a}_m$, by the lower semicontinuity of L , $L(\mathbf{a}^*) \leq \liminf_{m \rightarrow \infty} L(\mathbf{a}_m) \leq s$, which implies $\mathbf{a}^* \in A$. Thus, the closedness is proven.

Finally, we prove that there exists a vector $\tilde{\mathbf{a}} \in A$, such that $\tilde{\mathbf{a}} = \inf_{\mathbf{a} \in \mathbb{R}^N} L(\mathbf{a})$, which implies that the minimizer is attained in B . Since $0 \leq \inf_{\mathbf{a} \in \mathbb{R}^N} L(\mathbf{a}) \leq s < \infty$, there exists $\{\tilde{\mathbf{a}}_m\}_{m \geq 1}$, such that $\lim_{m \rightarrow \infty} L(\tilde{\mathbf{a}}_m) = \inf_{\mathbf{a} \in \mathbb{R}^N} L(\mathbf{a})$. Without

loss of generality, assume $L(\tilde{\mathbf{a}}_m) \leq s$ for all m , i.e. $\tilde{\mathbf{a}}_m \in A$. Since A is compact, there exists a vector $\tilde{\mathbf{a}} \in A$, such that $\|\mathbf{a}_m - \tilde{\mathbf{a}}\|_N \rightarrow 0$. Thus, by the lower semicontinuity of L ,

$$L(\tilde{\mathbf{a}}) \leq \liminf_{m \rightarrow \infty} L(\tilde{\mathbf{a}}_m) = \inf_{\mathbf{a} \in \mathbb{R}^N} L(\mathbf{a}),$$

i.e. $\tilde{\mathbf{a}} \in A \subset B$ is the minimzer of L . \square

A.2. Proof of Theorem 3.1

Proof. Let $L_n(f)$ denote the objective function, that is,

$$L_n(f) = \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \rho(Y_{ij} - f(\mathbf{X}_j)).$$

Let $g(\mathbf{X}_j) = f^*(\mathbf{X}_j) - f(\mathbf{X}_j)$ and it is easy to see that minimizing $L_n(f)$ is equivalent to minimizing

$$L_n(g) = \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \rho(\epsilon_{ij} + R_j + g(\mathbf{X}_j)).$$

where $R_j = f_0(\mathbf{X}_j) - f^*(\mathbf{X}_j)$.

Denoting $C_{nN} = (nN^\nu)^{-\frac{\theta}{\theta+1}} \log^6(nN^\nu)$, we aim to show that for every $\epsilon > 0$, there exists a $\gamma_\epsilon \geq 1$ (in the following we simply write γ instead of γ_ϵ), such that

$$\lim_{n \rightarrow \infty} \Pr \left(\inf_{\|g\|_N = \gamma} L_n(C_{nN}^{1/2} g) \geq L_n(\mathbf{0}) \right) \geq 1 - \epsilon. \quad (\text{A.2})$$

Define $\mathcal{F}^* = \{f - f^* : f \in \mathcal{F}\}$, such that for any $g \in \mathcal{F}^*$, $(g_1, \dots, g_N)^\top \in [-2C, 2C]^N \subset \mathbb{R}^N$. We first show the convexity and lower semicontinuity of L_n on \mathcal{F}^* . The convexity follows from Assumption (A3) and the convexity of the map $(g_1, \dots, g_N)^\top \rightarrow \|g\|_N^2$. The lower semi-continuity can be derived from the Assumption (A3) that ρ is convex and continuous, thus also lower semicontinuous. Note that by definition, since $f^*, f \in \mathcal{F}$, we have $\|f^* - f\|_N \leq \|f^* - f\|_\infty \leq 2C$. Therefore, when equality $L_n(\mathbf{0}) = \inf_{\|g\|_N = \gamma} L_n(C_{nN}^{1/2} g)$ holds, the convex lower semicontinuous L_n automatically guarantees the minimizer is attained for $\hat{f} \in \mathcal{F}$ such that $\|\hat{f} - f^*\|_N \leq \gamma C_{nN}^{1/2}$.

We now establish (A.2). To this end, we decompose

$$\begin{aligned} & L_n(C_{nN}^{1/2} g) - L_n(\mathbf{0}) \\ &= \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \rho(\epsilon_{ij} + R_j + C_{nN}^{1/2} g(\mathbf{X}_j)) - \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \rho(\epsilon_{ij} + R_j) \\ &= I_1(g) + I_2(g) + I_3(g), \end{aligned}$$

where

$$\begin{aligned}
 I_1(g) &= \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \int_{R_j}^{R_j + C_{nN}^{1/2} g(\mathbf{X}_j)} \mathbb{E}\{\psi(\epsilon_{ij} + u)\} du \\
 I_2(g) &= \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \int_{R_j}^{R_j + C_{nN}^{1/2} g(\mathbf{X}_j)} [\{\psi(\epsilon_{ij} + u) - \psi(\epsilon_{ij})\} \\
 &\quad - \mathbb{E}\{\psi(\epsilon_{ij} + u) - \psi(\epsilon_{ij})\}] du \\
 I_3(g) &= C_{nN}^{1/2} \frac{1}{nN} \sum_{i=1}^n \sum_{j=1}^N \psi(\epsilon_{ij}) g(\mathbf{X}_j).
 \end{aligned}$$

By the superadditivity of the infimum we have the lower bound

$$\inf_{\|g\|_N = \gamma} \{L_n(C_{nN}^{1/2} g) - L_n(\mathbf{0})\} \geq \inf_{\|g\|_N = \gamma} I_1(g) + \inf_{\|g\|_N = \gamma} I_2(g) + \inf_{\|g\|_N = \gamma} I_3(g).$$

We need to determine the order of each one of the three terms.

By Schwarz Inequality, we obtain

$$|I_3(g)| \leq \frac{C_{nN}^{1/2}}{n} \left\{ \frac{1}{N} \sum_{j=1}^N \left| \sum_{i=1}^n \psi(\epsilon_{ij}) \right|^2 \right\}^{1/2} \left\{ \frac{1}{N} \sum_{j=1}^N |g(\mathbf{X}_j)|^2 \right\}^{1/2}.$$

By Assumption (A6), we have

$$\begin{aligned}
 \mathbb{E} \left\{ \frac{1}{N} \sum_{j=1}^N \left| \sum_{i=1}^n \psi(\epsilon_{ij}) \right|^2 \right\} &= \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^n \sum_{k=1}^n \mathbb{E}\{\psi(\epsilon_{ij})\psi(\epsilon_{kj})\} \\
 &= \frac{n}{N} \sum_{j=1}^N \mathbb{E}\{|\psi(\epsilon_{ij})|^2\} = O(nN^{-\nu}).
 \end{aligned}$$

Using Markov's inequality, we thus find $\left\{ \frac{1}{N} \sum_{j=1}^N \left| \sum_{i=1}^n \psi(\epsilon_{ij}) \right|^2 \right\}^{1/2} = O_p(n^{1/2} N^{-\nu/2})$. The second factor is naturally bounded by γ . Combining these two bounds and note that $(nN^\nu)^{-1} \ll C_{nN}$, we obtain

$$\sup_{\|g\|_N = \gamma} |I_3(g)| = C_{nN}^{1/2} O_p(1) (nN^\nu)^{-1/2} \gamma \leq O_p(\gamma C_{nN}). \tag{A.3}$$

When n is large enough, for all $g \in \mathcal{F}^*$,

$$C_{nN}^{1/2} \max_{j \leq N} |g(\mathbf{X}_j)| \leq 2C_{nN}^{1/2} C = o(1),$$

and by Lemma A.1,

$$\max_{j \leq N} |R_j| \leq \|f_0 - f^*\|_\infty \leq C_I M^{-\theta}.$$

When $M^{-\theta} = O(1)C_{nN}$, we have $\max_{j \leq N} |R_j| = o(1)$ when n is large.

Applying Assumption (A6) yields, for large n

$$\begin{aligned} & \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \mathbf{E}\{\psi(\epsilon_{ij} + u)\}du \\ &= \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \{\delta_j u + o(u)\}du \\ &= \frac{C_{nN}}{2}\delta_j|g(\mathbf{X}_j)|^2\{1 + o(1)\} + \delta_j C_{nN}^{1/2}R_j g(\mathbf{X}_j)\{1 + o(1)\}. \end{aligned}$$

Hence,

$$\begin{aligned} & \frac{1}{N} \sum_{j=1}^N \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \mathbf{E}\{\psi(\epsilon_{ij} + u)\}du \\ &= \frac{C_{nN}}{2N} \sum_{j=1}^N \delta_j |g(\mathbf{X}_j)|^2 \{1 + o(1)\} + \frac{C_{nN}^{1/2}}{N} \sum_{j=1}^N \delta_j R_j g(\mathbf{X}_j) \{1 + o(1)\} \\ &\geq I_{11}(g) + I_{12}(g), \end{aligned}$$

where $I_{11}(g) := 1/2 \inf_j \delta_j C_{nN} \|g\|_N^2$ and $I_{12}(g) := N^{-1} C_{nN}^{1/2} \sum_{j=1}^N \delta_j R_j |g(\mathbf{X}_j)|^2 \{1 + o(1)\}$. We find an upper bound for $|I_{12}(g)|$. By the Schwarz inequality,

$$\begin{aligned} \sup_{\|g\|_N=\gamma} |I_{12}(g)| &\leq \sup_{\|g\|_N=\gamma} \left(\sup_j \delta_j \right) \frac{C_{nN}^{1/2}}{N} \sum_{j=1}^N |g(\mathbf{X}_j)| |R_j| \{1 + o(1)\} \\ &\lesssim C_{nN}^{1/2} \|R\|_N \sup_{\|g\|_N=\gamma} \|g\|_N \{1 + o(1)\} \\ &\lesssim \gamma C_{nN} \{1 + o(1)\}, \end{aligned}$$

where the last inequality is derived by $\|R\|_N^2 \leq \|f_0 - f^*\|_\infty \leq M^{-\theta} \asymp C_{nN}$. Therefore,

$$\inf_{\|g\|_N=\gamma} I_{11}(g) + I_{12}(g) = O(1)C_{nN}\gamma^2\{1 + O(\gamma^{-1}) + o(1)\}$$

for all $g \in \mathcal{F}$.

In the following, we show that for every $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} \Pr \left(\sup_{\|g\|_N=\gamma} |I_2(g)| \geq \epsilon C_{nN} \right) = 0. \tag{A.4}$$

Consider the class of real-valued functions $h_g : \mathbb{R}^N \rightarrow \mathbb{R}$ given by

$$h_g(\mathbf{w}) = \frac{1}{N} \sum_{j=1}^N \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \{\psi(w_j + u) - \psi(w_j)\}du.$$

Let \mathbb{P}_n be the empirical measure associated with $(\epsilon_{i1}, \dots, \epsilon_{iN}), i = 1, \dots, n$. We rewrite $I_2(g)$ as $I_2(g) = (\mathbb{P}_n - \mathbb{P})h_g$, where $\mathbb{P}h_g$ is the expectation for h_g . Let $\mathcal{F}^* = \{f - f^* : f \in \mathcal{F}\}$, where \mathcal{F} is the sparse neural network class defined in (2.4). By Markov’s inequality, we have

$$\Pr \left(\sup_{g \in \mathcal{F}^*} |(\mathbb{P}_n - \mathbb{P})h_g| > \epsilon C_{nN} \right) \leq \frac{\mathbb{E}\{|\sup_{g \in \mathcal{F}^*} n^{1/2}|(\mathbb{P}_n - \mathbb{P})h_g|\}}{n^{1/2}C_{nN}\epsilon}. \tag{A.5}$$

Thus it suffices to show that the right-hand side of the above inequality tends to zero as $n \rightarrow \infty$.

Let $\mathcal{N}_{[]}(\epsilon, \mathcal{G}, L_2(\mathbb{P}))$ be the ϵ -bracketing number for a set of functions \mathcal{G} in the $L_2(\mathbb{P})$ -norm and define the ϵ -bracketing integral,

$$J_{[]}(\epsilon, \mathcal{G}, L_2(\mathbb{P})) = \int_0^\epsilon \{\log \mathcal{N}_{[]}(\epsilon, \mathcal{G}, L_2(\mathbb{P}))\}^{1/2} du.$$

By Lemma 19.36 of [23], for any class of real-valued functions \mathcal{H} such that $\mathbb{E}(|h|^2) < \delta^2$ and $\|h\|_\infty < C_3$ for all $h \in \mathcal{H}$, we have

$$\mathbb{E} \left(\sup_{h \in \mathcal{H}} n^{1/2}|(\mathbb{P}_n - \mathbb{P})h| \right) \leq c_0 J_{[]}(\delta, \mathcal{H}, L_2(\mathbb{P})) \left(1 + \frac{J_{[]}(\delta, \mathcal{H}, L_2(\mathbb{P}))C_3}{\delta^2 n^{1/2}} \right).$$

Hence, in the following, we first determine a bound on the bracketing number and then we estimate δ and C_3 . By our previous discussion, $\max_{j \leq N} |R_j| \rightarrow 0$ and $C_{nN}^{1/2} \times \sup_{g \in \mathcal{F}^*} \|g\|_\infty \rightarrow 0$. Hence, for any $(g_1, g_2) \in \mathcal{F}^* \times \mathcal{F}^*$, by Assumption (A4),

$$\begin{aligned} & \sup_{\mathbf{w} \in \mathbb{R}^N} |h_{g_1}(\mathbf{w}) - h_{g_2}(\mathbf{w})| \\ &= \sup_{\mathbf{w} \in \mathbb{R}^N} \left| \frac{1}{N} \sum_{j=1}^N \int_{R_j + C_{nN}^{1/2}g_2(\mathbf{x}_j)}^{R_j + C_{nN}^{1/2}g_1(\mathbf{x}_j)} \{\psi(w_j + u) - \psi(w_j)\} du \right| \\ &\leq 2c_1 C_{nN}^{1/2} \|g_1 - g_2\|_\infty. \end{aligned} \tag{A.6}$$

Let $\mathcal{N}(\epsilon, \mathcal{F}^*, \|\cdot\|_\infty)$ denote the ϵ -covering number in the sup-norm. By (A.6) and Theorem 2.7.11 of [24], we now have

$$\mathcal{N}_{[]}(\epsilon, \{h_g, g \in \mathcal{F}^*\}, L_2(\mathbb{P})) \leq \mathcal{N}(\epsilon/(4c_1 C_{nN}^{1/2}), \mathcal{F}^*, \|\cdot\|_\infty). \tag{A.7}$$

By Lemma 5 of [20] and Lemma A.1, we have

$$\begin{aligned} & \log \mathcal{N}(\epsilon/(4c_1 C_{nN}^{1/2}), \mathcal{F}^*, \|\cdot\|_\infty) \\ &\leq (s + 1) \log(2\epsilon^{-1}(4c_1 C_{nN}^{1/2})(L + 1)V^2) \\ &\asymp O(1)M \log M \left\{ \log^2 M + \log(C_{nN}^{1/2}\epsilon^{-1}) \right\}, \end{aligned} \tag{A.8}$$

where the last inequality follows the conditions for L, s and V given in Lemma A.1.

Next, we estimate δ and C_3 . By the Schwarz inequality and Assumption (A5), we have

$$\begin{aligned} \mathbb{E}\{|h_g|^2\} &\leq N^{-2} \mathbb{E} \left\{ \left| \sum_{j=1}^N \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \{\psi(\epsilon_{ij} + u) - \psi(\epsilon_{ij})\} du \right|^2 \right\} \\ &\leq N^{-1} \sum_{j=1}^N \mathbb{E} \left\{ \left| \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \{\psi(\epsilon_{ij} + u) - \psi(\epsilon_{ij})\} du \right|^2 \right\} \\ &\leq N^{-1} \sum_{j=1}^N C_{nN}^{1/2} |g(\mathbf{X}_j)| \mathbb{E} \left\{ \int_{R_j}^{R_j+C_{nN}^{1/2}g(\mathbf{X}_j)} \{\psi(\epsilon_{ij} + u) - \psi(\epsilon_{ij})\}^2 du \right\} \\ &\leq C_{nN}^{1/2} \max_{j \leq N} |g(\mathbf{X}_j)| N^{-1} \sum_{j=1}^N c_2 \{|R_j|^2 + C_{nN} |g(\mathbf{X}_j)|^2\} \\ &\leq 2c_2 C_{nN}^{3/2}, \end{aligned}$$

where the last inequality follows by the definition of neural network space \mathcal{F} in (2.4). Since this estimate is uniform on \mathcal{F} , we take $\delta^2 = c_0 C_{nN}^{3/2}$, where the constant c_0 will be specified later. In addition, by Assumption (A5) and Lemma A.1, we have

$$\sup_{g \in \mathcal{F}^*} \sup_{\mathbf{w} \in \mathbb{R}^N} |h_g(\mathbf{w})| \leq c_1 C_{nN}^{1/2} \sup_{g \in \mathcal{F}^*} \left\{ \frac{1}{N} \sum_{j=1}^N |g(\mathbf{X}_j)|^2 \right\}^{1/2} \leq c_1 C_{nN}^{1/2}.$$

Let $c_0 = (2c_2) \vee c_1$, then $C_3 = c_0 C_{nN}^{1/2}$ and $C_3/\delta^2 = C_{nN}^{-1}$. With our estimate of δ and the bound on the bracketing number implied by (A.7) and (A.8), when we take $M^{-\theta} \asymp C_{nN}$, the bracketing integral over $n^{1/2} C_{nN}$ is bounded as follows

$$\begin{aligned} &(n^{1/2} C_{nN})^{-1} J(\delta, \mathcal{F}^*, L_2(\mathbb{P})) \\ &\lesssim C_{nN}^{-1/4} n^{-1/2} M^{1/2} (\log^{3/2} M + (\log M \log C_{nN}^{-1/4})^{1/2}) \\ &= O(1) \left(n^{\theta/(4\theta+4)} \log^{-3/2} n \right) n^{-1/2} \left(n^{1/(2\theta+2)} \log^{-3/\theta} n \right) \log^2 n \\ &= O(1) n^{-\theta/(4\theta+4)} \log^{(\theta-6)/(2\theta)} n = o(1), \end{aligned}$$

which also implies $C_3(\delta^2 n^{1/2})^{-1} J(\delta, \mathcal{G}, L_2(\mathbb{P})) = o(1)$. Consequently, inequality (A.5) entails (A.4). By the previous discussion, (A.2) holds. Given the convexity and lower semicontinuity of L_n on \mathcal{F} , by Lemma A.2, inequality (A.2) entails the existence of a minimizer $\hat{g} = f^* - \hat{f}$, such that $\|\hat{g}\|_N^2 = O_p(C_{nN})$. Note that by definition, since $f^*, f \in \mathcal{F}$, we have $\|g\|_N \leq \|f^* - f\|_\infty \leq 2C$. Therefore, when $L_n(\mathbf{0}) = \inf_{\|g\|_N = \gamma} L_n(C_{nN}^{1/2} g)$, the convex lower semicontinuous L_n automatically guarantees the minimizer is attained in $\{g : \|g\|_N \leq 2C C_{nN}\}$. Using the one-to-one relation between g and f , if (A.2) holds, we obtain $\|f - f^*\|_N^2 =$

$O_p(C_{nN})$. By Lemma A.1, we have

$$\begin{aligned}\|\widehat{f} - f\|_N^2 &\leq 2\|\widehat{f} - f^*\|_N^2 + 2\|f^* - f\|_N^2 &= O_p(C_{nN} + M^{-\theta}) \\ &= O_p(C_{nN} + M^{-\theta}) = O_p(C_{nN}).\end{aligned}$$

which implies Theorem 3.1. \square

Acknowledgments

We thank an Associate Editor and two referees for their constructive comments which have helped to improve the presentation of the paper.

Data used in preparation of this article were obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

References

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. *Proceedings of the 36th International Conference on Machine Learning*, 97:242–252, 2019. [MR4001274](#)
- [2] Juan Lucas Bali, Graciela Boente, David E Tyler, and Jane-Ling Wang. Robust functional principal components: A projection-pursuit approach. *Annals of Statistics*, 39(6):2852–2882, 2011. [MR3012394](#)
- [3] B. Bauer and M. Kohler. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47:2261–2285, 2019. [MR3953451](#)
- [4] G. Cao, L. Wang, Y. Li, and L. Yang. Oracle-efficient confidence envelopes for covariance functions in dense functional data. *Statistica Sinica*, 26:359–383, 2016. [MR3468356](#)
- [5] G. Cao, L. Yang, and D. Todem. Simultaneous inference for the mean function of dense functional data. *Journal of Nonparametric Statistics*, 24:359–377, 2012. [MR2921141](#)
- [6] Kehui Chen and Hans-Georg Müller. Modeling repeated functional observations. *J. Amer. Statist. Assoc.*, 107(500):1599–1609, 2012. [MR3036419](#)
- [7] Lu-Hung Chen and Ci-Ren Jiang. Multi-dimensional functional principal component analysis. *Stat. Comput.*, 27(5):1181–1192, 2017. [MR3647091](#)
- [8] Jianqing Fan, Cong Ma, and Yiqiao Zhong. A selective overview of deep learning. *arXiv preprint arXiv:1904.05526*, 2019. [MR4255194](#)
- [9] T. Hsing and R. Eubank. *Theoretical foundations of functional data analysis, with an introduction to linear operators*. Wiley Series in Probability and Statistics. John Wiley & Sons, Ltd., Chichester, 2015. [MR3379106](#)

- [10] Ioannis Kalogridis. Asymptotics for m-type smoothing splines with non-smooth objective functions. *arXiv*, page <https://arxiv.org/abs/2002.04898>, 2020. [MR4447316](#)
- [11] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *In the 3rd International Conference on Learning Representations (ICLR)*, 2015.
- [12] P. Kokoszka and M. Reimherr. *Introduction to functional data analysis*. Texts in Statistical Science Series. CRC Press, Boca Raton, FL, 2017. [MR3793167](#)
- [13] Seokho Lee, Hyejin Shin, and Nedret Billor. M-type smoothing spline estimators for principal functions. *Computational Statistics & Data Analysis*, 66:89–100, 2013. [MR3064026](#)
- [14] Italo R. Lima, Guanqun Cao, and Nedret Billor. M-based simultaneous inference for the mean function of functional data. *Annals of the Institute of Statistical Mathematics*, 71:577–598, 2019. [MR3968850](#)
- [15] Italo R. Lima, Guanqun Cao, and Nedret Billor. Robust simultaneous inference for the mean function of functional data. *TEST*, 28:785–803, 2019. [MR3992138](#)
- [16] R. Liu, B. Boukai, and Z. Shang. Optimal nonparametric inference via deep neural network. *Journal of Mathematical Analysis and Applications*, 505:125561, 2021. [MR4300990](#)
- [17] R. Liu, Z. Shang, and G. Cheng. On deep instrumental variables estimate. *arXiv:2004.14954*, 2021.
- [18] Ricardo A Maronna and Victor J Yohai. Robust functional linear regression based on splines. *Computational Statistics & Data Analysis*, 65:46–55, 2013. [MR3064942](#)
- [19] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis, Second Edition*. Springer Series in Statistics, New York, 2005. [MR2168993](#)
- [20] J. Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020. [MR4134774](#)
- [21] Hyejin Shin and Seokho Lee. An RKHS approach to robust functional linear regression. *Statistica Sinica*, 26:255–272, 2016. [MR3468352](#)
- [22] Barinder Thind, Kevin Multani, and Jiguo Cao. Deep learning with functional inputs. *arXiv preprint arXiv:2006.09590*, 2020.
- [23] A. W. van der Vaart. *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998. [MR1652247](#)
- [24] Aad W. van der Vaart and Jon A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics. [MR1385671](#)
- [25] J.L. Wang, J. M. Chiou, and H. G. Müller. Functional data analysis. *Annual Review of Statistics and Its Application*, 3:257–295, 2016.
- [26] Shuoyang Wang, Guanqun Cao, and Zuofeng Shang. Estimation of the mean function of functional data via deep neural networks. *Stat*, e393, 2021. [MR4319004](#)
- [27] F. Yao, H. G. Müller, and J. L. Wang. Functional data analysis for

- sparse longitudinal data. *Journal of the American Statistical Association*, 100:577–590, 2005. [MR2160561](#)
- [28] Lingsong Zhang, Haipeng Shen, and Jianhua Z. Huang. Robust regularized singular value decomposition with application to mortality data. *The Annals of Applied Statistics*, 7(3):1540–1561, 2013. [MR3127958](#)
- [29] Lan Zhou and Huijun Pan. Principal component analysis of two-dimensional functional data. *Journal of Computational and Graphical Statistics*, 23(3):779–801, 2014. [MR3224656](#)