

Neural Networks Perform Sufficient Dimension Reduction

Shuntuo Xu, Zhou Yu*

Key Laboratory of Advanced Theory and Application in Statistics and Data Science – MOE
School of Statistics, East China Normal University, Shanghai, China
oaksword@163.com, zyu@stat.ecnu.edu.cn

Abstract

This paper investigates the connection between neural networks and sufficient dimension reduction (SDR), demonstrating that neural networks inherently perform SDR in regression tasks under appropriate rank regularizations. Specifically, the weights in the first layer span the central mean subspace. We establish the statistical consistency of the neural network-based estimator for the central mean subspace, underscoring the suitability of neural networks in addressing SDR-related challenges. Numerical experiments further validate our theoretical findings, and highlight the underlying capability of neural networks to facilitate SDR compared to the existing methods. Additionally, we discuss an extension to unravel the central subspace, broadening the scope of our investigation.

Introduction

Neural networks have achieved significant success in a tremendous variety of applications (Lee and Abu-El-Haija 2017; Silver et al. 2018; Jumper et al. 2021; Brandes et al. 2022; Thirunavukarasu et al. 2023). As the most essential foundation, a feedforward neural network is typically constructed by a series of linear transformations and nonlinear activations. To be specific, a function f implemented by a feedforward neural network with L layers can be represented as

$$f(x) = \phi_L \circ \sigma_{L-1} \circ \phi_{L-1} \circ \cdots \circ \sigma_1 \circ \phi_1(x). \quad (1)$$

Here, \circ is the functional composition operator, $\phi_i(z) = W_i^\top z + b_i$ denotes the linear transformation and σ_i represents the elementwise nonlinear activation function. Despite a clear architecture, the formula (1) provides limited insight into how the information within the input data is processed by the neural networks. To comprehensively understand how neural networks retrieve task-relevant information, there have been extensive efforts towards unraveling their interpretability (Doshi-Velez and Kim 2017; Guidotti et al. 2018; Zhang et al. 2021). For instance, post-hoc interpretability (Lipton 2018) focused on the predictions generated by neural networks, while disregarding the detailed mechanism and feature importance. Ghorbani, Abid, and

Zou (2019) highlighted the fragility of this type of interpretation, as indistinguishable perturbations could result in completely different interpretations.

For the sake of striking a balance between the complexity of representation and the power of prediction, Tishby and Zaslavsky (2015) and Saxe et al. (2019) pioneered the interpretability of deep neural networks via the information bottleneck theory. Ghosh (2022) further linked the information bottleneck to sufficient dimension reduction (SDR) (Li 1991; Cook 1996), which is a rapidly developing research field in regression diagnostics, data visualization, pattern recognition and etc.

In this paper, we provide a theoretical understanding of neural networks for representation learning from the SDR perspective. Let $x \in \mathbb{R}^p$ represent the covariates and $y \in \mathbb{R}$ represent the response. Consider the following regression model

$$y = f_0(B_0^\top x) + \epsilon, \quad (2)$$

where $B_0 \in \mathbb{R}^{p \times d}$ is a nonrandom matrix with $d \leq p$, $f_0 : \mathbb{R}^d \rightarrow \mathbb{R}$ is an unknown function, and ϵ is the noise such that $\mathbb{E}(\epsilon|x) = 0$ and $\text{Var}(\epsilon|x) = \nu^2$ for some positive constant ν . Intuitively, the semiparametric and potentially multi-index model (2) asserts that the core information for regression is encoded in the low-dimensional linear representation $B_0^\top x$. In the literature of SDR, model (2) has been extensively studied. Based on model (2), Cook and Li (2002) proposed the objective of sufficient mean dimension reduction as

$$y \perp\!\!\!\perp \mathbb{E}(y|x)|B^\top x \quad (3)$$

for some matrix B , where $\perp\!\!\!\perp$ means statistical independence. Denote the column space spanned by B as Π_B , which is commonly referred to as the mean dimension-reduction subspace. It is evident that the matrix B satisfying (3) is far away from uniqueness. Hence, we focus on the intersection of all possible mean dimension-reduction subspace, which is itself a mean dimension-reduction subspace under mild conditions (e.g., when the domain of x is open and convex; see Cook and Li (2002)), and term it the central mean subspace. Agreeing to condition (3), B_0 defined in model (2) yields the central mean subspace, denoted as Π_{B_0} , under certain assumption. Statistical estimation and inference about the central mean subspace is the primary goal of SDR. Popular statistical methods for recovering central mean subspace

*Corresponding author.

include ordinary least squares (Li and Duan 1989), principal Hessian directions (Li 1992), minimum average variance estimation (Xia et al. 2002), semiparametric approach (Ma and Zhu 2013), generalized kernel-based dimension reduction (Fukumizu and Leng 2014), and many others. Although numerous studies have demonstrated the ability of neural networks to approximate complex functions (Hornik, Stinchcombe, and White 1989; Barron 1993; Yarotsky 2017; Shen, Yang, and Zhang 2021) and to adapt low-dimensional structures (Bach 2017; Bauer and Kohler 2019; Schmidt-Hieber 2020; Abbe, Adsera, and Misiakiewicz 2022; Jiao et al. 2023; Troiani et al. 2024), it is of subsequent interest to investigate whether neural networks are capable of correctly identifying the intrinsic structure encapsulated in B_0 , thereby deepening the interpretation of neural networks.

Our study was inspired by an observation that the weight matrix in the first layer, i.e., W_1 , could accurately detect the presence of B_0 in a toy data set, with the rank regularization. Specifically, consider the toy model $y = (B_0^\top x)^3 + \epsilon$ where $x \sim \text{Uniform}([0, 1]^{10})$, $B_0 = (1, -2, 0, \dots, 0)^\top$ and $\epsilon \sim \text{Normal}(0, 0.1^2)$. We trained neural networks using the least squares loss with $W_1 = W_{11}W_{12}$ where $W_{11} \in \mathbb{R}^{10 \times q}$ and $W_{12} \in \mathbb{R}^{q \times 64}$ for $q = 1, \dots, 10$. Evidently, the rank of W_1 did not exceed q . It was then observed that for each q , (i) B_0 was closely contained within the column space of W_{11} , as the absolute cosine similarity between B_0 and its projection on $\Pi_{W_{11}}$ was close to 1, and (ii) the leading eigenvector of $W_1W_1^\top$ closely aligned with B_0 (see Figure 1). This observation indicates that the first layer of the neural network may potentially discover the underlying low-dimensional intrinsic structure.

It is important to note that the application of neural networks for estimating the central mean subspace based on model (2) was previously explored by Kapla, Fertl, and Bura (2022) through a two-stage method. The first stage focused on obtaining a preliminary estimation of Π_{B_0} , which subsequently served as an initial point for the joint estimation of Π_{B_0} and f_0 in the second stage. Given our toy example, however, it is prudent to critically evaluate the necessity of the first stage. Furthermore, their work lacks a comprehensive theoretical guarantee. Additionally, another re-




Figure 1: Absolute cosine similarity between (i) B_0 and its projection on $\Pi_{W_{11}}$ (the dot line with triangle marks), (ii) B_0 and the leading eigenvector of $W_1W_1^\top$ (the solid line with square marks).

lated work conducted by Liang, Sun, and Liang (2022) concentrated on seeking nonlinear sufficient representations. In contrast, we focus more on revealing the fundamental nature of neural networks.

We in this paper show that, with suitable rank regularization, the first layer of a feedforward neural network conducts SDR in a regression task, wherein $d(W_{11}, B_0) \rightarrow 0$ in probability for certain distance metric $d(\cdot, \cdot)$. Furthermore, numerical experiments provide empirical evidences supporting this result, while demonstrating the efficiency of neural networks in addressing the issue of SDR.

Throughout this paper, we use $\|v\|_2$ to represent the Euclidean norm of a vector v . For a matrix A , $\|A\|_F$ is the Frobenius norm of A , $\pi_A = A(A^\top A)^{-1}A^\top$ denotes the projection matrix of A where A^\top is the generalized inverse of A , and Π_A stands for the linear space spanned by the columns of A . For a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$, $\|f\|_{L^2(\mu)}$ represents the L^2 norm of f with respect of a given probability measure μ , and $\|f\|_{L^\infty(\mathcal{X})} = \sup_{x \in \mathcal{X}} |f(x)|$ represents the supreme norm of f over a set \mathcal{X} . $\mathcal{B}(\mathcal{X})$ is the unit ball induced by \mathcal{X} , such that $\mathcal{B}(\mathcal{X}) \subset \mathcal{X}$ and $\|v\|_2 \leq 1$ for any $v \in \mathcal{X}$.

Theoretical Justifications

Population-Level Unbiasedness

Suppose the true intrinsic dimension d defined in model (2) is known and the covarites $x \in \mathcal{B}([0, 1]^p)$. As B_0 is not identifiable in (2) and (3), it is assumed without loss of generality that $B_0^\top B_0 = I_d$ where I_d is the identity matrix with d rows. By defining $\Psi_d = \{B \in \mathbb{R}^{p \times d} : B^\top B = I_d\}$, we have $B_0 \in \Psi_d$. In this paper, we consider the following neural network function class

$$\mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}} = \left\{ \begin{aligned} f(x) &= \phi_L \circ \sigma_{L-1} \circ \phi_{L-1} \circ \dots \circ \sigma_1 \circ \phi_1(B^\top x) : \\ B &\in \Psi_d, \phi_i(z) = W_i^\top z + b_i, W_i \in \mathbb{R}^{d_i \times d_{i+1}}, i = 1, \dots, L, \end{aligned} \right.$$

with $L = \mathcal{L}$, $\max_{i=1, \dots, L+1} d_i = \mathcal{M}$, $\sum_{i=1}^L (d_i + 1)d_{i+1} = \mathcal{S}$,

$$\|f\|_{L^\infty(\mathcal{B}([0, 1]^p))} \leq \mathcal{R} \}.$$

The activation functions $\sigma_i (i = 1, \dots, L-1)$ utilized are the rectified linear units, i.e., $\sigma_i(x) = \max(x, 0)$. We emphasize that $\mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}$ incorporates rank regularization of d in the first layer. For arbitrary $f \in \mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}$, we use $\mathcal{T}(f)$ to represent the component $B \in \Psi_d$ in the first layer of f .

For a regression task, the theoretical study is highly related to the smoothness of underlying conditional mean function (Yang and Barron 1999). Here, we introduce the following assumptions of model (2).

Assumption 1 (Smoothness). f_0 is a Hölder continuous function of order $\alpha \in (0, 1]$ with constant λ , i.e., $|f_0(x) - f_0(z)| \leq \lambda \|x - z\|_2^\alpha$ for any $x, z \in [0, 1]^d$. Additionally, $\|f_0\|_{L^\infty([0, 1]^d)} \leq \mathcal{R}_0$ for some constant $\mathcal{R}_0 \geq 1$.

Assumption 2 (Sharpness). For any scalar $\delta > 0$ and $B \in \Psi_d$, $\|\pi_B - \pi_{B_0}\|_F > \delta$ implies $\mathbb{E}\{\text{Var}[f_0(B_0^\top x)|B^\top x]\} > M(\delta)$ for some $M(\delta) > 0$.

Assumption 3. y is sub-exponentially distributed such that there exists $\tau > 0$ satisfying $\mathbb{E}[\exp(\tau|y|)] < \infty$.

Assumption 1 is a technical condition to study the approximation capability of neural networks (Shen, Yang, and Zhang 2020). Alternatively, other functional spaces, such as the Sobolev space, can also be employed for this purpose (Abdeljawad and Grohs 2022; Shen et al. 2022). Furthermore, Assumption 2 establishes a restriction on the sharpness of f_0 . Consider the case that f_0 is solely a constant function as zero. Then, a trivial neural network by setting all the parameters except B to zero perfectly fits f_0 , regardless of the value of B . With Assumption 2, it becomes difficult to accurately capture the overall behavior of $f_0(B_0^\top x)$ using a biased B . In other words, $B_0^\top x$ is sufficient and necessary for recovering $\mathbb{E}(y|x)$, i.e., Π_{B_0} is the central mean subspace. Similar condition was also adopted in Theorem 4.2 of Li and Dong (2009) to distinguish sufficient directions B_0 from other B . Assumption 3 is a commonly used condition for applying empirical process tools and concentration inequalities (Van der Vaart 2000; Zhu, Jiao, and Jordan 2022).

Theorem 1. Suppose that Assumptions 1 and 2 hold. Let $f^* = \operatorname{argmin}_{f \in \mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}} \mathbb{E}[y - f(x)]^2$, then

$$\Pi_{\mathcal{T}(f^*)} = \Pi_{B_0},$$

provided that \mathcal{R} is sufficiently large, and \mathcal{L} and \mathcal{M} tend to infinity.

Theorem 1 builds a bridge connecting neural networks and SDR. It demonstrates that neural networks indeed achieve representation learning as the first layer of the optimal neural network perfectly reaches the target of SDR at population level. Theorem 1 also inspires us to perform SDR based on neural networks with a minor adjustment for the first layer. The detailed proof of Theorem 1 can be found in Section Proofs.

Sample Estimation Consistency

We now investigate the theoretical property of the neural network-based sample-level estimator for SDR. Given sample observations $\mathcal{D}_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$, where (X_i, Y_i) is an independent copy of (x, y) for $i = 1, \dots, n$, the commonly used least squares loss is adopted, i.e.,

$$L_n(f) = \frac{1}{n} \sum_{i=1}^n [Y_i - f(X_i)]^2.$$

Denote the optimal neural network estimator at the sample level as

$$\hat{f}_n = \operatorname{argmin}_{f \in \mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}} L_n(f).$$

$\mathcal{T}(\hat{f}_n)$ is then the sample estimator approximately spanning the central mean subspace.

To examine the closeness between $\Pi_{\mathcal{T}(\hat{f}_n)}$ and Π_{B_0} , we define the distance metric $d(\cdot, \cdot)$ as

$$d(B, B_0) = \min_{Q \in \mathcal{Q}} \|B_0 - BQ\|_2,$$

where $B \in \Psi_d$ and \mathcal{Q} is the collection of all orthogonal matrices in $\mathbb{R}^{d \times d}$. We see that $d(B, B_0) = 0$ if and only if $\Pi_B = \Pi_{B_0}$. And we make the following assumption essentially another view of Assumption 2.

Assumption 4. For any positive scalar δ , $d(B, B_0) > \delta$ implies $\mathbb{E}\{\operatorname{Var}[f_0(B_0^\top x)|B^\top x]\} > M_1(\delta)$ for some $M_1(\delta) > 0$.

Theorem 2. Suppose the Assumptions 1, 3 and 4 hold. Then we have

$$d(\mathcal{T}(\hat{f}_n), B_0) \xrightarrow{P} 0,$$

when the depth $\mathcal{L} = O(n^{d/(2d+8\alpha)})$, the width $\mathcal{M} = O(1)$ and $\mathcal{R} \geq \mathcal{R}_0$.

Theorem 2 confirms that $\Pi_{\mathcal{T}(\hat{f}_n)}$ converges to Π_{B_0} in probability. As a result, under the least squares loss, the neural networks, with appropriate rank regularization, provide a consistent estimator of the central mean subspace. Therefore, it is promising to adopt the neural network function class $\mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}$ to address sufficient mean dimension reduction problems. More importantly, this approach offers several advantages over existing methods. Unlike SIR (Li 1991), SAVE (Cook and Weisberg 1991) and PHD (Li 1992), our method does not impose stringent probabilistic distributional conditions on x , such as linearity, constant variance and normality assumptions. Compared to MAVE (Xia et al. 2002), our method adopts the powerful neural networks to estimate the link function, thereby working similar or better than classical nonparametric tools based on first-order Taylor expansion. We present the proof of Theorem 2 in Section Proofs.

The results illustrated in Theorem 1 and 2 are contingent on the availability of true intrinsic dimension d . In the case where d is unknown, a natural modification for $\mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}$ is to set $B \in \mathbb{R}^{p \times p}$ without rank regularization. Under Assumptions 1 and 2, in this scenario, the optimal f^* at the population level still ensures that $\Pi_{\mathcal{T}(f^*)}$ encompasses Π_{B_0} , where the sharpness of f_0 plays a crucial role; see the Supplementary Material for more details.

Simulation Study

From the perspective of numerical experiments, we utilized simulated data sets to demonstrate that (i) the column space of $\mathcal{T}(\hat{f}_n)$ approached the central mean subspace Π_{B_0} as sample size n increased, (ii) the performance of neural network-based method in conducting SDR was comparable to classical methods. In some cases, the neural network-based method outperformed classical methods, particularly when the latent intrinsic dimension d was large. Five commonly used methods, sliced inverse regression (SIR), sliced average variance estimation (SAVE), principal Hessian directions (PHD), minimum average variance estimation (MAVE) and generalized kernel dimension reduction (GKDR), were included for comparisons.

The neural network-based method was implemented in Python using PyTorch. Specifically, we utilized a linear layer without bias term to represent the matrix B (recall that we suppose d is known), which was further appended by a fully-connected feedforward neural network with number of neurons being $h - h/2 - 1$. Here, we set $h = 64$ when the sample size was less than 1000, and set $h = 128$ when sample size was between 1000 and 2000. Overall, the neural network architecture was $p - d - h - h/2 - 1$. Code is available at <https://github.com/oaksword/DRNN>.

			NN	MAVE	GKDR	SIR	SAVE	PHD
Setting 1	$(n, p) = (100, 10)$	mean	0.135	0.160	0.388	0.897	0.315	0.623
		std	0.064	0.056	0.126	0.186	0.082	0.150
	$(n, p) = (200, 30)$	mean	0.276	0.214	1.012	0.946	0.298	0.807
		std	0.274	0.051	0.299	0.113	0.049	0.107
	$(n, p) = (300, 30)$	mean	0.120	0.130	0.542	0.900	0.337	0.709
		std	0.038	0.029	0.166	0.136	0.063	0.108
Setting 2	$\sigma = 0.1$	mean	0.296	0.730	1.130	0.665	0.319	1.550
		std	0.126	0.312	0.271	0.166	0.076	0.126
	$\sigma = 0.2$	mean	0.628	0.899	1.155	0.705	0.338	1.526
		std	0.269	0.329	0.237	0.149	0.074	0.138
	$\sigma = 0.5$	mean	1.197	1.187	1.278	0.830	0.367	1.567
		std	0.213	0.204	0.200	0.169	0.072	0.131
Setting 3	$n = 200$	mean	0.639	1.246	1.759	1.669	1.728	1.740
		std	0.418	0.289	0.149	0.235	0.136	0.221
	$n = 500$	mean	0.248	1.076	1.752	1.650	1.683	1.721
		std	0.242	0.331	0.125	0.271	0.201	0.228
	$n = 1000$	mean	0.075	0.924	0.554	1.652	1.678	1.737
		std	0.081	0.382	0.371	0.259	0.245	0.191
Setting 4	$n = 1000, d = 4$	mean	0.127	0.293	0.368	1.363	0.429	0.960
		std	0.161	0.050	0.079	0.167	0.079	0.178
	$n = 1500, d = 6$	mean	0.144	0.415	0.386	1.530	0.467	0.902
		std	0.067	0.076	0.077	0.153	0.082	0.206
	$n = 2000, d = 8$	mean	0.140	0.360	0.344	1.410	0.364	0.735
		std	0.055	0.084	0.010	0.163	0.065	0.175

Table 1: The results of average and standard deviation of $\|\pi_{\hat{B}} - \pi_{B_0}\|_F$ on 100 replicates across different methods. NN represents the neural network-based method.

We considered the following four scenarios, with two of them attaining small $d = 1, 2$ and the rest of $d \geq 3$.

Setting 1: $y = x_1^4 + \epsilon$ where $x \in \text{Normal}(0, I_p)$ and $\epsilon \in \mathfrak{t}_5$.

Setting 2: $y = \log(x_1 + x_1 x_2) + \epsilon$ where $x \in \text{Uniform}([0, 1]^{10})$ and $\epsilon \in \text{Normal}(0, \sigma^2)$.

Setting 3: $y = (1 + \beta_1^\top x)^2 \exp(\beta_2^\top x) + 5 \cdot 1(\beta_3^\top x > 0) + \epsilon$ where $x \in \text{Uniform}([-1, 1]^6)$, $\epsilon \sim \chi_2^2 - 2$ and

$$(\beta_1, \beta_2, \beta_3) = \begin{pmatrix} -2 & -1 & 0 & 1 & 2 & 3 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 \end{pmatrix}^\top.$$

Setting 4:

$$y = \sum_{k=1}^d \frac{e^{x_k}}{2 + \sin(\pi x_k)} + \epsilon,$$

where $\epsilon \sim \text{Normal}(0, 0.1^2)$,

$$x \sim \text{Normal}(1_{10}, I_{10} - 1_{10}1_{10}^\top/20).$$

In setting 1, we tested the combinations of $(n, p) = (100, 10), (200, 30)$ and $(300, 30)$, where n was the sample size. In setting 2, we fixed $n = 200, p = 10$ and set $\sigma = 0.1, 0.2, 0.5$. In setting 3, we fixed $p = 6$ and tested $n = 200, 500, 1000$. In setting 4, we fixed $p = 10$ and tested $(n, d) = (1000, 4), (1500, 6)$ and $(2000, 8)$. Settings 1, 2, 4 were equipped with continuous regression functions, and

setting 3 involved discontinuity. We set the number of iterations of our method to 1000 for settings 1 and 2, and increased the iterations to 2000 and 4000 for settings 3 and 4, due to their high complexity.

To evaluate the performance of each method, we calculated the distance metric $\|\pi_{\hat{B}} - \pi_{B_0}\|_F$ where \hat{B} represented the estimate of B_0 . Particularly, for our neural network-based method, $\hat{B} = \mathcal{T}(\hat{f}_n)$. Smaller value of this metric indicated better performance. We ran 100 replicates for each setting. The results are displayed in Table 1 and Figure 2.

In simple cases (settings 1 and 2), the neural network-based method showed similar performance to classical methods, with sliced average variance estimation being the most effective method in setting 2. However, complex settings demonstrated the significant superiority of neural network-based method compared to other methods. Specifically, in setting 3, as the sample size n increased, the metric $\|\pi_{\hat{B}} - \pi_{B_0}\|_F$ decreased rapidly, indicating the convergence of $\Pi_{\hat{B}}$ to Π_{B_0} . According to the results in setting 4, the neural network-based method was capable of handling higher-dimensional scenarios. In summary, numerical studies advocated neural networks as a powerful tool for SDR.

Real Data Analysis

We further applied the neural network-based method involving rank regularization to a real regression task. In practice,

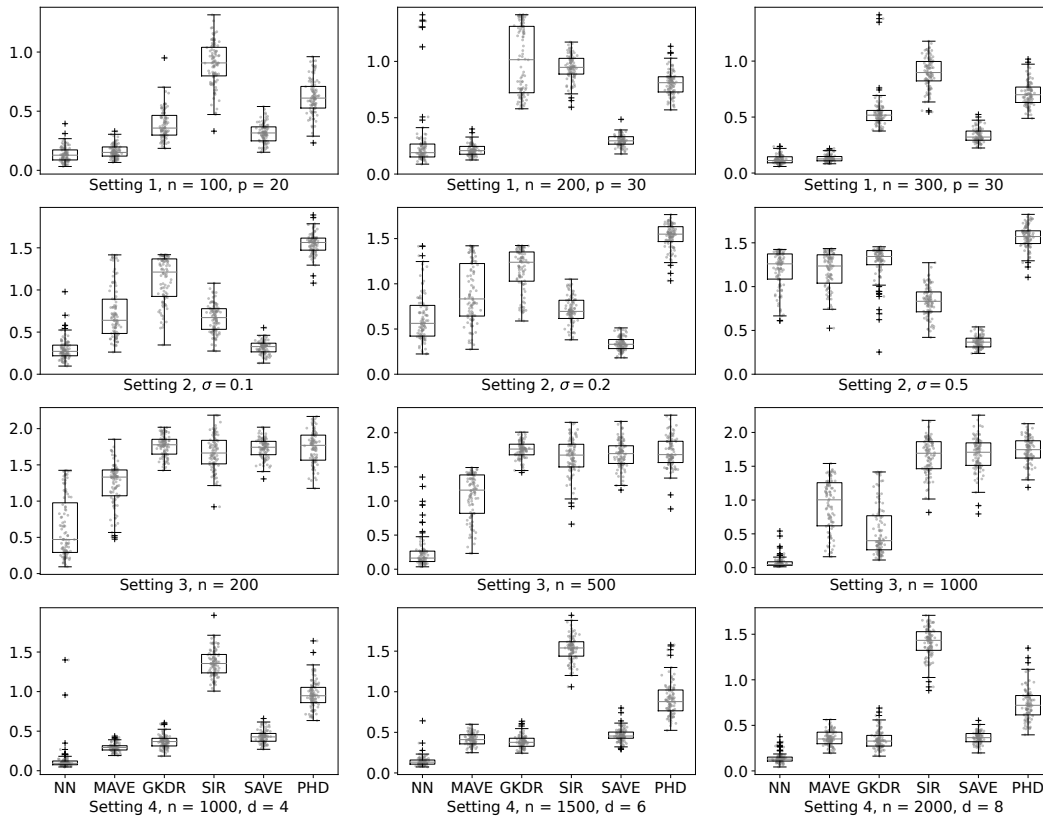


Figure 2: Boxplots of $\|\pi_{\hat{B}} - \pi_{B_0}\|_F$ on 100 replicates across different methods. In each panel, six methods for SDR are included, among which NN represents the neural network-based method.

the precise intrinsic dimension d is unknown, and it is uncertain whether a low-dimensional structure exists. To address this issue, we used cross validation to determine an appropriate d from the range of values $\{1, \dots, p\}$. More specifically, the raw data set was divided into 80% training data and 20% testing data. The optimal value of d was determined through cross validation on the training data. Subsequently, the final model was fit using the selected d on the training data, and the mean squared error on the testing data was evaluated.

In order to reduce randomness, the aforementioned process was repeated 20 times and the resulting testing errors were recorded. Additionally, we conducted a comparative analysis between the neural network-based approach and alternative methods including vanilla neural network without rank regularization, SIR-based regression, SAVE-based regression, and MAVE-based regression. For the latter three techniques, we initially executed SDR to acquire the embedded data, followed by the utilization of a fully-connected neural network for predictive purpose. The optimal value for d was also determined through cross validation.

We utilized a data set of weather records from Seoul, South Korea during the summer months from 2013 to 2017 (Cho et al. 2020), available at the UCI data set repository (bias correction of numerical prediction model temperature forecast). This data set contained 7750 observations with 23 predictors and 2 responses, specifically the next-day max-

	NN-RR	NN-VN	MAVE	SIR	SAVE
mean	0.602	0.774	0.743	1.324	0.772
std	0.043	0.116	0.159	0.190	0.161
\bar{d}	19.6	—	19.25	19.6	20.6

Table 2: The results of testing errors on 20 replicates across different methods. We report the average and standard deviation of testing errors, along with averaged optimal d determined through cross validation. NN-RR, NN-VN, MAVE, SIR, SAVE represent the neural network-based method with rank regularization, vanilla neural network regression, MAVE-based regression, SIR-based regression, and SAVE-based regression, respectively.

imum air temperature and next-day minimum air temperature. After excluding the variables for station and date, the data set was reduced to 21 predictors, which were further standardized using StandardScaler in scikit-learn package.

It is evident from Table 2 that the neural network-based method with rank regularization outperformed other methods, demonstrating the effectiveness of the modification compared to the vanilla neural network, and the sound performance in reducing dimensions compared to other SDR methods. The presence of latent structures was partially sup-

ported by the averaged optimal d . It was possible that 19 or 20 combinations of raw predictors might be sufficient, as opposed to the original 21 predictors.

Proofs

Proof of Theorem 1

Define the vanilla neural network function class without rank regularization as

$$\begin{aligned} \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}^* &= \left\{ \right. \\ f(x) &= \phi_L \circ \sigma_{L-1} \circ \phi_{L-1} \circ \cdots \circ \sigma_1 \circ \phi_1(x) : \\ \phi_i(z) &= W_i^\top z + b_i, W_i \in \mathbb{R}^{d_i \times d_{i+1}}, i = 1, \dots, L, \\ \text{with } L &= \mathcal{L}, \max_{i=1, \dots, L+1} d_i = \mathcal{M}, \sum_{i=1}^L (d_i + 1) d_{i+1} = \mathcal{S}, \\ &\left. \|f\|_{L^\infty([0,1]^d)} \leq \mathcal{R} \right\}. \end{aligned}$$

Lemma 1. *Suppose that h is Hölder continuous with $\alpha \in (0, 1]$ and $\lambda > 0$. Then, for any $\zeta > 0$, there exists a function g in neural network function class $\mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}^*$, with rectified linear unit activation and $\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}$ large enough, such that*

$$\|h - g\|_{L^\infty([0,1]^d)} < \zeta.$$

We note that Lemma 1 is a simplified version of Theorem 1.1 in (Shen, Yang, and Zhang 2020).

Proof of Theorem 1. For $f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}$ such that $\pi_{\mathcal{T}(f)} \neq \pi_{B_0}$, under Assumption 2, there exists a scalar $t > 0$ satisfying

$$\mathbb{E} \{ \text{Var} [f_0(B_0^\top x) | \mathcal{T}(f)^\top x] \} > t.$$

Then, we have

$$\begin{aligned} \mathbb{E}[y - f(x)]^2 &= \mathbb{E}[\epsilon + f_0(B_0^\top x) - f(x)]^2 \\ &= \mathbb{E}(\epsilon^2) + \mathbb{E}[f_0(B_0^\top x) - f(x)]^2 \\ &= \nu^2 + \text{Var}[f_0(B_0^\top x) - f(x)] \\ &\quad + \mathbb{E}^2[f_0(B_0^\top x) - f(x)] \\ &\geq \nu^2 + \text{Var}[f_0(B_0^\top x) - f(x)] \\ &\geq \nu^2 + \mathbb{E}\{\text{Var}[f_0(B_0^\top x) | \mathcal{T}(f)^\top x]\} \\ &> \nu^2 + t. \end{aligned}$$

On the other hand, for $\tilde{f} \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}$ such that $\pi_{\mathcal{T}(\tilde{f})} = \pi_{B_0}$, there exists an orthogonal matrix $Q \in \mathbb{R}^{d \times d}$ such that $B_0 = \mathcal{T}(\tilde{f})Q$. Hence, $\tilde{f}(x) = \tilde{f} \circ \mathcal{T}(\tilde{f}) \circ Q(B_0^\top x)$ and $\tilde{f} \circ \mathcal{T}(\tilde{f}) \circ Q$ is still a neural network function. Assumption 1 and Lemma 1 imply that there is a neural network function g satisfying

$$\|f_0(B_0^\top x) - g(B_0^\top x)\|_{L^\infty(\mathcal{B}([0,1]^p))} < t^{1/2}/2,$$

for sufficiently large $\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}$. As a result, select \tilde{f} such that $\tilde{f} \circ \mathcal{T}(\tilde{f}) \circ Q = g$, and it follows that

$$\mathbb{E}[y - \tilde{f}(x)]^2 = \nu^2 + \mathbb{E}[f_0(B_0^\top x) - g(B_0^\top x)]^2 < \nu^2 + t/4.$$

Therefore, for any $f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}$ such that $\pi_{\mathcal{T}(f)} \neq \pi_{B_0}$, there exists a neural network function \tilde{f} , with the same rank regularization, satisfying

$$\mathbb{E}[y - \tilde{f}(x)]^2 < \mathbb{E}[y - f(x)]^2.$$

To conclude, $\pi_{\mathcal{T}(f^*)} = \pi_{B_0}$, which entails that $\Pi_{\mathcal{T}(f^*)} = \Pi_{B_0}$. \square

Proof of Theorem 2

Let $L(f) = \mathbb{E}[y - f(x)]^2$. Without loss of generality, suppose $\mathcal{R} = \mathcal{R}_0$. We first present some useful lemmas.

Lemma 2. *For sufficiently large n , it follows that*

$$\sup_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} |L(f) - L_n(f)| = O_p \left(\sqrt{\frac{\mathcal{S}\mathcal{L} \log(\mathcal{S})(\log n)^5}{n}} \right).$$

The proof of Lemma 2 is presented in the Supplementary Material. We note this convergence rate may not be optimal, but is sufficient for deducing consistency.

Lemma 3. *Under Assumptions 1 and 3, for sufficiently large n , we have*

$$\mathbb{E}[y - \hat{f}_n(x)]^2 \leq \nu^2 + Cn^{-2\alpha/(d+4\alpha)} (\log n)^3,$$

where C is a constant depending on \mathcal{R}_0 and d .

Proof of Lemma 3. We begin with the decomposition that

$$\begin{aligned} L(\hat{f}_n) &\leq \mathbb{E} \left[L_n(\hat{f}_n) + \sup_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} |L(f) - L_n(f)| \right] \\ &\leq \mathbb{E} \left[\inf_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} L_n(f \circ \mathcal{T}(f) \circ B_0^\top) \right. \\ &\quad \left. + \sup_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} |L(f) - L_n(f)| \right] \\ &\leq \mathbb{E} \left[\inf_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} L(f \circ \mathcal{T}(f) \circ B_0^\top) \right. \\ &\quad \left. + 2 \sup_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} |L(f) - L_n(f)| \right] \\ &\leq \inf_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} L(f \circ \mathcal{T}(f) \circ B_0^\top) \\ &\quad + C_1 \sqrt{\mathcal{S}\mathcal{L} \log(\mathcal{S})(\log n)^5 n^{-1}}, \end{aligned}$$

where C_1 is a constant, for sufficiently large n . Hence,

$$\begin{aligned} &L(\hat{f}_n) - \nu^2 \\ &= L(\hat{f}_n) - L(f_0 \circ B_0^\top) \\ &= L(\hat{f}_n) - \inf_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} L(f \circ \mathcal{T}(f) \circ B_0^\top) \\ &\quad + \inf_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} L(f \circ \mathcal{T}(f) \circ B_0^\top) - L(f_0 \circ B_0^\top) \\ &\leq C_1 \sqrt{\mathcal{S}\mathcal{L} \log(\mathcal{S})(\log n)^5 n^{-1}} \\ &\quad + \inf_{f \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}} \mathbb{E}[f_0(B_0^\top x) - f \circ \mathcal{T}(f)(B_0^\top x)]^2 \\ &\leq C_1 \sqrt{\mathcal{S}\mathcal{L} \log(\mathcal{S})(\log n)^5 n^{-1}} \\ &\quad + \inf_{g \in \mathcal{F}_{\mathcal{L},\mathcal{M},\mathcal{S},\mathcal{R}}^*} \|f_0 - g\|_{L^\infty([0,1]^d)}^2. \end{aligned}$$

By Theorem 1.1 in Shen, Yang, and Zhang (2020), there exists g^* with $\mathcal{M} = 3^{d+3} \max(d \lfloor K^{1/d} \rfloor, K + 1)$, $\mathcal{L} = 12D + 14 + 2d$ for some constant $K, D > 0$ such that

$$\|f_0 - g^*\|_{L^\infty([0,1]^d)} \leq 19\sqrt{d}\lambda(KD)^{-2\alpha/d}.$$

Let

$$D = O\left(n^{d/(2d+8\alpha)}\right), \quad K = O(1),$$

and observe that $\mathcal{S} = O(\mathcal{M}^2\mathcal{L})$. Then, it follows that

$$\begin{aligned} L(\hat{f}_n) &= \mathbb{E}[y - \hat{f}_n(x)]^2 \\ &\leq \nu^2 + Cn^{-2\alpha/(d+4\alpha)}(\log n)^3, \end{aligned}$$

where C is a constant depending on \mathcal{R}_0 and d . \square

With Lemma 3, the proof of Theorem 2 is a direct application of Theorem 5.9 in Van der Vaart (2000).

Proof of Theorem 2. Recall that $m_0 = f_0 \circ B_0^\top$. Let $R(f) = L(f) - L(m_0) = \mathbb{E}[f_0(B_0^\top x) - f(x)]^2 \geq 0$ and $R_n(f) = L_n(f) - L_n(m_0)$. Then, Lemma 2 indicates that

$$\sup_{f \in \mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}} |R_n(f) - R(f)| \rightarrow 0, \quad \text{in probability,}$$

yielding $R_n(\hat{f}_n) = o_p(1)$ by incorporating with Lemma 3.

We denote the metric

$$d_1(f, m_0) = \min_{Q \in \mathcal{Q}} \|B_0 - \mathcal{T}(f)Q\|_2 \vee \|f_0 \circ Q - f \circ \mathcal{T}(f)\|_{L^2(\mu)}.$$

Here, $f \in \mathcal{F}_{\mathcal{L}, \mathcal{M}, \mathcal{S}, \mathcal{R}}$, $m_0 = f_0 \circ B_0^\top$, $a \vee b$ means $\max(a, b)$, and μ is the probability distribution of $\mathcal{T}(f)^\top x$. Recall that $d(B, B_0) = \min_{Q \in \mathcal{Q}} \|B_0 - BQ\|_2$, where $B \in \Psi_d$.

For f and $\delta > 0$ such that $d_1(f, m_0) \geq \delta$, if

$$d(\mathcal{T}(f), B_0) \leq \tilde{\delta} = \min\{\delta^2/(8\mathcal{R}_0\lambda)\}^{1/\alpha}, \delta/2\},$$

we have

$$\begin{aligned} |f_0(B_0^\top x) - f_0(Q^\top \mathcal{T}(f)^\top x)| &\leq \lambda \|B_0^\top x - Q^\top \mathcal{T}(f)^\top x\|_2^\alpha \\ &\leq \lambda \|B_0^\top - Q^\top \mathcal{T}(f)^\top\|_2^\alpha \\ &\leq \frac{\delta^2}{8\mathcal{R}_0}, \end{aligned}$$

for some orthogonal matrix Q . Hence,

$$\begin{aligned} R(f) &= \mathbb{E}[f_0(B_0^\top x) - f(x)]^2 \\ &\geq \mathbb{E}[f_0(Q^\top \mathcal{T}(f)^\top x) - f(x)]^2 \\ &\quad + 2\mathbb{E}[f_0(B_0^\top x) - F_0][F_0 - f(x)] \\ &\geq -\frac{\delta^2}{2} + d_1(f, m_0)^2 \geq \frac{\delta^2}{2} > 0, \end{aligned}$$

where $F_0 = f_0(Q^\top \mathcal{T}(f)^\top x)$.

For the case that $d(\mathcal{T}(f), B_0) > \tilde{\delta}$, by Assumption 4, we have $R(f) \geq \xi > 0$ for some positive constant ξ . To conclude, we obtain

$$\inf_{f: d_1(f, m_0) \geq \delta} R(f) > 0,$$

for any $\delta > 0$. Finally, applying Theorem 5.9 in Van der Vaart (2000), $R_n(\hat{f}_n) = o_p(1)$ implies that $d_1(\hat{f}_n, m_0) \rightarrow 0$ in probability. Hence, $d(\mathcal{T}(\hat{f}_n), B_0) \rightarrow 0$ in probability. \square

Discussion

We have demonstrated that neural networks attain the capability of detecting the underlying low-dimensional structure that preserves all information in x about the conditional mean function $\mathbb{E}(y|x)$. As a result, neural networks are suitable to be utilized for the estimation of central mean subspace Π_{B_0} . The theoretical investigations sharpen our understanding of neural networks, while broadening the scope of SDR as well.

In the context of SDR, a more general scenario than sufficient mean dimension reduction emerges when considering

$$y \perp\!\!\!\perp x | B_0^\top x. \quad (4)$$

And the column space spanned by B_0 corresponds to the central subspace (Cook 1998a; Li 2018). It is clear that relation (4) is equivalent to that $p(y|x) = p(y|B_0^\top x)$, where $p(\cdot|\cdot)$ represents the conditional probability density function. Following the work of Xia (2007), we can further adapt neural networks to estimate the central subspace by modifying the loss function.

Under mild conditions, Xia (2007) showed that

$$\mathbb{E}[K_h(y - y_0)|x = x_0] \rightarrow p(y_0|B_0^\top x_0), \quad \text{as } h \rightarrow 0^+.$$

Here, (x_0, y_0) is a fixed point, and $K_h(\cdot)$ is a suitable kernel function with a bandwidth h . Such finding then implies that

$$\begin{aligned} K_h(y - y_0) &= \mathbb{E}[K_h(y - y_0)|x] + \text{remainder term} \\ &\approx p(y_0|B_0^\top x) + \text{remainder term}. \end{aligned}$$

Based on this discovery, it is natural to employ a neural network function $f(B^\top x, y_0) : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ to approximate $p(y_0|B_0^\top x)$, and obtain the estimate of B_0 at the population level by solving the following problem

$$(B^*, f^*) = \underset{B, f}{\operatorname{argmin}} \mathbb{E}[K_h(y - \tilde{y}) - f(B^\top x, \tilde{y})]^2,$$

where \tilde{y} is an independent copy of y . Empirically, we define the loss function as

$$\tilde{L}_n(B, f) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n [K_h(y_j - y_i) - f(B^\top x_j, y_i)]^2.$$

We provide some additional simulation results to verify the feasibility utilizing neural networks for the estimation of central subspace; see the Supplementary Material. Theoretical analysis of neural networks for the estimation of the central subspace, including unbiasedness and consistency, deserves further studies.

References

- Abbe, E.; Adsera, E. B.; and Misiakiewicz, T. 2022. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, 4782–4887. PMLR.
- Abdeljawad, A.; and Grohs, P. 2022. Approximations with deep neural networks in Sobolev time-space. *Analysis and Applications*, 20(03): 499–541.

- Bach, F. 2017. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(19): 1–53.
- Barron, A. R. 1993. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3): 930–945.
- Bauer, B.; and Kohler, M. 2019. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47(4): 2261 – 2285.
- Brandes, N.; Ofer, D.; Peleg, Y.; Rappoport, N.; and Linal, M. 2022. ProteinBERT: A universal deep-learning model of protein sequence and function. *Bioinformatics*, 38(8): 2102–2110.
- Cho, D.; Yoo, C.; Im, J.; and Cha, D.-H. 2020. Comparative assessment of various machine learning-based bias correction methods for numerical weather prediction model forecasts of extreme air temperatures in urban areas. *Earth and Space Science*, 7(4): e2019EA000740.
- Cook, R. D. 1996. Graphics for regressions with a binary response. *Journal of the American Statistical Association*, 91(435): 983–992.
- Cook, R. D. 1998a. *Regression graphics*. New York: Wiley.
- Cook, R. D.; and Li, B. 2002. Dimension reduction for conditional mean in regression. *The Annals of Statistics*, 30(2): 455–474.
- Cook, R. D.; and Weisberg, S. 1991. Sliced inverse regression for dimension reduction: Comment. *Journal of the American Statistical Association*, 86(414): 328–332.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. arXiv:1702.08608.
- Fukumizu, K.; and Leng, C. 2014. Gradient-based kernel dimension reduction for regression. *Journal of the American Statistical Association*, 109(505): 359–370.
- Ghorbani, A.; Abid, A.; and Zou, J. 2019. Interpretation of neural networks is fragile. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 3681–3688.
- Ghosh, D. 2022. Sufficient dimension reduction: An information-theoretic viewpoint. *Entropy*, 24(2): 167.
- Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; and Pedreschi, D. 2018. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5): 1–42.
- Hornik, K.; Stinchcombe, M.; and White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5): 359–366.
- Jiao, Y.; Shen, G.; Lin, Y.; and Huang, J. 2023. Deep nonparametric regression on approximate manifolds: Nonasymptotic error bounds with polynomial prefactors. *The Annals of Statistics*, 51(2): 691–716.
- Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Židek, A.; Potapenko, A.; et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873): 583–589.
- Kapla, D.; Fertl, L.; and Bura, E. 2022. Fusing sufficient dimension reduction with neural networks. *Computational Statistics & Data Analysis*, 168: 107390.
- Lee, J.; and Abu-El-Haija, S. 2017. Large-scale content-only video recommendation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 987–995.
- Li, B. 2018. *Sufficient dimension reduction: Methods and applications with R*. CRC Press.
- Li, B.; and Dong, Y. 2009. Dimension reduction for nonelliptically distributed predictors. *The Annals of Statistics*, 37(3): 1272 – 1298.
- Li, K.-C. 1991. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414): 316–327.
- Li, K.-C. 1992. On principal Hessian directions for data visualization and dimension reduction: Another application of Stein’s lemma. *Journal of the American Statistical Association*, 87(420): 1025–1039.
- Li, K.-C.; and Duan, N. 1989. Regression analysis under link violation. *The Annals of Statistics*, 17(3): 1009–1052.
- Liang, S.; Sun, Y.; and Liang, F. 2022. Nonlinear sufficient dimension reduction with a stochastic neural network. *Advances in Neural Information Processing Systems*, 35: 27360–27373.
- Lipton, Z. C. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3): 31–57.
- Ma, Y.; and Zhu, L. 2013. Efficient estimation in sufficient dimension reduction. *Annals of statistics*, 41(1): 250.
- Saxe, A. M.; Bansal, Y.; Dapello, J.; Advani, M.; Kolchinsky, A.; Tracey, B. D.; and Cox, D. D. 2019. On the information bottleneck theory of deep learning. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12): 124020.
- Schmidt-Hieber, J. 2020. Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4): 1875–1897.
- Shen, G.; Jiao, Y.; Lin, Y.; and Huang, J. 2022. Approximation with cnns in Sobolev space: With applications to classification. *Advances in Neural Information Processing Systems*, 35: 2876–2888.
- Shen, Z.; Yang, H.; and Zhang, S. 2020. Deep network approximation characterized by number of neurons. *Communications in Computational Physics*, 28(5): 1768–1811.
- Shen, Z.; Yang, H.; and Zhang, S. 2021. Neural network approximation: Three hidden layers are enough. *Neural Networks*, 141: 160–173.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419): 1140–1144.
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nature medicine*, 29(8): 1930–1940.

- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, 1–5. IEEE.
- Troiani, E.; Dandi, Y.; Defilippis, L.; Zdeborova, L.; Loureiro, B.; and Krzakala, F. 2024. Fundamental limits of weak learnability in high-dimensional multi-index models. In *High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning*.
- Van der Vaart, A. W. 2000. *Asymptotic statistics*, volume 3. Cambridge university press.
- Xia, Y. 2007. A constructive approach to the estimation of dimension reduction directions. *The Annals of Statistics*, 35(6): 2654 – 2690.
- Xia, Y.; Tong, H.; Li, W. K.; and Zhu, L.-X. 2002. An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(3): 363–410.
- Yang, Y.; and Barron, A. 1999. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, 1564–1599.
- Yarotsky, D. 2017. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94: 103–114.
- Zhang, Y.; Tiño, P.; Leonardis, A.; and Tang, K. 2021. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(5): 726–742.
- Zhu, B.; Jiao, J.; and Jordan, M. I. 2022. Robust estimation for non-parametric families via generative adversarial networks. In *2022 IEEE International Symposium on Information Theory (ISIT)*, 1100–1105. IEEE.