Deep learning with missing data

Tianyi Ma^{*}, Tengyao Wang[†] and Richard J. Samworth^{*}

*Statistical Laboratory, University of Cambridge [†]Department of Statistics, London School of Economics

Abstract

In the context of multivariate nonparametric regression with missing covariates, we propose Pattern Embedded Neural Networks (PENNs), which can be applied in conjunction with any existing imputation technique. In addition to a neural network trained on the imputed data, PENNs pass the vectors of observation indicators through a second neural network to provide a compact representation. The outputs are then combined in a third neural network to produce final predictions. Our main theoretical result exploits an assumption that the observation patterns can be partitioned into cells on which the Bayes regression function behaves similarly, and belongs to a compositional Hölder class. It provides a finite-sample excess risk bound that holds for an arbitrary missingness mechanism, and in combination with a complementary minimax lower bound, demonstrates that our PENN estimator attains in typical cases the minimax rate of convergence as if the cells of the partition were known in advance, up to a poly-logarithmic factor in the sample size. Numerical experiments on simulated, semi-synthetic and real data confirm that the PENN estimator consistently improves, often dramatically, on standard neural networks without pattern embedding. Code to reproduce our experiments, as well as a tutorial on how to apply our method, is publicly available.

1 Introduction

Over the last decade or so, deep neural networks have achieved stunning empirical successes in learning tasks across diverse application areas, including image classification (Krizhevsky, Sutskever and Hinton, 2012; He et al., 2016), protein folding (Jumper et al., 2021), natural language processing (Vaswani et al., 2017; Devlin et al., 2019) and many others. While their complexity initially led to the temptation to regard them as a black box, recent theory does provide insights into the origins of their impressive practical performance (Jacot, Gabriel and Hongler, 2018; Schmidt-Hieber, 2020; Kohler and Langer, 2021; Mei and Montanari, 2022; Suh and Cheng, 2024).

The works described in the previous paragraph all concern settings where the (large) datasets involved are fully observed. Nevertheless, as argued by Zhu, Wang and Samworth (2022), missing data play an ever more significant role in high-dimensional learning problems, and new methods have now been introduced to tackle several different contemporary statistical challenges involving missing data, including sparse linear regression (Loh and Wainwright, 2012; Chandrasekher, Alaoui and Montanari, 2020), principal component analysis (Elsener and van de Geer, 2019; Zhu, Wang and Samworth, 2022), classification (Cai and Zhang, 2019; Sell, Berrett and Cannings, 2024) and changepoint estimation (Follain, Wang and Samworth, 2022). Most of these papers study the simplest, idealised case where the data and missingness indicators are independent, a setting known as Missing Completely At Random (MCAR). Indeed, it is now understood that for more general (dependent) missingness mechanisms, even consistent mean estimation may be impossible without further assumptions (Ma et al., 2024).

The goal of this paper is to study supervised deep learning problems with missing values among the covariates, where we do not rely on an MCAR assumption, or indeed any other restriction on the missingness mechanism. Motivated by many applications in which deep learning is applied, we focus on the problem of prediction (i.e. making a forecast of the response at a new covariate vector) as opposed to estimation (i.e. learning the underlying parameters of the model). Our primary methodological contribution is to introduce the idea of (observation) pattern embedding into the neural network framework. In other words, in addition to training a neural network on our original covariates (with missing values imputed via any existing technique), we pass the vectors of observation indicators, which we call *revelation vectors*, through another neural network to obtain a compact representation that summarises the information they contain. We can then train a third neural network that combines these two earlier ones to produce final predictions.

The benefits of our approach are illustrated in Figure 1, where our training data of size n = 1,000 and d = 1 are generated from the model described in Example 1 in Section 2.1 below. Our original covariates $(\boldsymbol{X}_i)_{i=1}^n$ are observed with missingness, and we apply zero imputation based on the corresponding revelation vectors $(\boldsymbol{\Omega}_i)_{i=1}^n$ to obtain imputed covariates $(\boldsymbol{Z}_i)_{i=1}^n$. Thus, together with the responses $(Y_i)_{i=1}^n$, which satisfy $Y_i | \boldsymbol{X}_i \sim N(\boldsymbol{X}_i^2, 0.1^2)$ independently for $i = 1, \ldots, n$, we have training data $(\boldsymbol{Z}_i, \boldsymbol{\Omega}_i, Y_i)_{i=1}^n$. Following the imputation, the standard approach would be to train a model on $(\boldsymbol{Z}_i, Y_i)_{i=1}^n$ (see Figure 1(a)), yielding fitted values on an independent test set of size 500 displayed in panel (c). On the other hand, our approach trains a neural network on the augmented data $(\boldsymbol{Z}_i, \boldsymbol{\Omega}_i, Y_i)_{i=1}^n$ in panel (b), leading to much more accurate predictions on the test set, as illustrated in panel (d). The main point to observe here is that the failure to include the revelation vectors $(\boldsymbol{\Omega}_i)_{i=1}^n$ in panel (c) corrupts the output in a neighbourhood of the origin, which is a point of discontinuity of the function $\boldsymbol{z} \mapsto \mathbb{E}(Y_1 | \boldsymbol{Z}_1 = \boldsymbol{z})$, whereas this is corrected in panel (d) by the inclusion of the additional covariate.

After a more formal description of our problem set-up and some background on ReLU neural networks, we introduce our *Pattern Embedded Neural Networks* (PENNs) in Section 2. These can be fitted using standard empirical risk minimisation algorithms such as Adam (Kingma and Ba, 2015), as implemented in PyTorch in Python, yielding an estimator \hat{f} of the Bayes regression function f^* , given by $f^*(\boldsymbol{z}, \boldsymbol{\omega}) \coloneqq \mathbb{E}(Y_1 | \boldsymbol{Z}_1 = \boldsymbol{z}, \boldsymbol{\Omega}_1 = \boldsymbol{\omega})$.

In Section 3 we turn our attention to the theoretical properties of our procedure. We begin with a general oracle inequality (Theorem 1), revealing that under a sub-Gaussian condition on the response, the excess risk of our estimator is controlled by the sum of optimisation error, approximation error and estimation error terms, with the latter two



(a) Test data (blue) and Bayes regression func- (b) Test data (blue) and Bayes regression function $\boldsymbol{z} \mapsto \mathbb{E}(Y_1 \mid \boldsymbol{Z}_1 = \boldsymbol{z})$ (red). tion $(\boldsymbol{z}, \boldsymbol{\omega}) \mapsto \mathbb{E}(Y_1 \mid \boldsymbol{Z}_1 = \boldsymbol{z}, \boldsymbol{\Omega}_1 = \boldsymbol{\omega})$ (red).



(c) Output of the neural network trained with- (d) Output of the neural network trained with out the revelation vectors.

Figure 1: An illustration of Example 1, and the outputs of neural networks trained without and with the revelation vectors.

reflecting a bias-variance trade-off in the complexity of the neural network class. Our main result (Theorem 3 in Section 3.2) specialises the general setting studied previously to give a more explicit excess risk bound. In particular, since there 2^d possible observation patterns, we introduce a new condition on the Bayes regression function f^* to encapsulate the notion that the set of all possible observation patterns may be partitioned into sets on which f^* behaves similarly as a function of the covariates. We further ask that on each cell of this partition, f^* belongs to a compositional Hölder smoothness class. This allows us to show that the sum of the approximation error and estimation error terms can be controlled by a weighted average of estimation rates for each cell of the partition, together with an additional term reflecting the complexity of the partition that will typically be negligible, up to a poly-logarithmic factor in the sample size. This theorem is complemented by a minimax lower bound in Theorem 4, which confirms that the weighted average of the estimation rates over different cells in Theorem 3 is optimal. All proofs are deferred to the Appendix.

A numerical study of the performance of our PENN estimator is presented in Section 4. Our simulated data settings consider both MCAR and Missing Not At Random (MNAR) scenarios, as well as three commonly-used imputation techniques, namely (columnwise) mean imputation, MissForest imputation (Stekhoven and Bühlmann, 2012) and Multiple Imputation by Chained Equations (MICE) (van Buuren and Groothuis-Oudshoorn, 2011). A consistent pattern emerges whereby the PENN estimator is able to improve, sometimes very significantly, on the corresponding neural network estimator that does not include pattern embedding. We also study two semi-synthetic and two real datasets with the same imputation methods. In the former cases, the original real data, on handwritten digits and bank loans, are observed without missingness, which allows us to introduce either MCAR or MNAR missingness via known mechanisms. On the other hand, in the latter real datasets, on credit scores and public procurement, missingness is already present in the original data. By splitting the data into training, validation and test sets, we again see consistent improvements from our pattern embedding approach. Python code to reproduce all of the experiments in this section, together with an accompanying tutorial on how to apply our method, is available at https://github.com/tianyima2000/DNN_ missing_data.

1.1 Related literature

The desire to understand and explain the impressive empirical performance of deep learning represents a major current research theme in statistics and machine learning. One line of work provides explicit rates of convergence through the lens of approximation theory and empirical process theory, assuming that the empirical risk can be minimised sufficiently well. Minimax optimality (in terms of the sample size) of neural networks has been studied, for example, in the context of nonparametric regression (Bauer and Kohler, 2019; Imaizumi and Fukumizu, 2019; Schmidt-Hieber, 2020; Kohler and Langer, 2021; Jiao et al., 2023; Fan and Gu, 2024), classification (Bos and Schmidt-Hieber, 2022; Zhang, Shi and Zhou, 2024), the partially linear Cox model (Zhong, Mueller and Wang, 2022) and density estimation (Bos and Schmidt-Hieber, 2023). In particular, it is now known that neural networks can avoid the curse of dimensionality in nonparametric regression under a compositional Hölder assumption (Schmidt-Hieber, 2020; Kohler and Langer, 2021). While most of the bounds in the work mentioned above have pre-factors depending exponentially on the covariate dimension or the number of variables on which the regression function depends, Jiao et al. (2023) provide upper bounds that depend only polynomially on the dimension, using the approximation scheme of Lu et al. (2021), while Imaizumi and Fukumizu (2019) allow discontinuities in the regression function. Other research directions focus on understanding the training dynamics of neural networks. These include the analysis of stochastic gradient descent for extremely (or infinitely) wide neural networks, known as the Neural Tangent Kernel regime (Jacot, Gabriel and Hongler, 2018; Du et al., 2019; Arora et al., 2019), and single hidden layer neural networks in the proportional asymptotics regime (Mei, Montanari and Nguyen, 2018; Mei and Montanari, 2022).

Missing data has been an active research topic in statistics for several decades, but is currently undergoing a renaissance for two main reasons. First, as already mentioned, missing values are ubiquitous in large datasets, and, unlike for other forms of data corruption, many statistical learning algorithms cannot be applied until the missingness is handled appropriately. Second, and somewhat related, the complexity of modern data demands new models for missingness and new inferential methods with appropriate guarantees on their performance (Berrett and Samworth, 2023; Ma et al., 2024). It has long been recognised that the MCAR assumption is unrealistic for many practical applications. On the other hand, modern data generating mechanisms can often only be adequately described by high- or infinite-dimensional parameter spaces. In such settings, missingness models such as Missing At Random (MAR) that are predicated on the correctness of lowdimensional parametric models fitted using maximum likelihood may be inappropriate.

For regression problems with missing covariates, one widely used general strategy is *impute-then-regress*, i.e. we first impute the missing data and then treat the imputed dataset as complete to train a regression algorithm. Many imputation algorithms have been proposed under MCAR or MAR assumptions, such as MissForest imputation (Stekhoven and Bühlmann, 2012), MICE (van Buuren and Groothuis-Oudshoorn, 2011), and methods based on deep learning and generative models (Li, Jiang and Marlin, 2019; Mattei and Frellsen, 2019; Nazabal et al., 2020; Zhang et al., 2025). If there exists a universally consistent estimator, then under some conditions that still allow MNAR, the impute-then-regress approach leads to asymptotically vanishing excess risk on new covariate vectors as the sample size diverges to infinity (Le Morvan et al., 2021; Josse et al., 2024). However, the Bayes regression function, which in this case is the conditional expectation of the response given the imputed covariate vector, is generally discontinuous and hard to learn (Le Morvan et al., 2021). Efformovich (2011) studies orthogonal series methods in univariate nonparametric regression with MAR missingness. Other strategies for regression with missing data include augmenting the covariate space with a distinguished point, reflecting a missing value, in relevant coordinates; this can be applied with regression trees or other (non-orthogonally equivariant) methods such as XGBoost (Chen and Guestrin, 2016). Twala, Jones and Hand (2008) propose regression trees with the observation patterns included as covariates, Śmieja et al. (2018) replace neurons in the first

hidden layer of a neural network by estimated expected values to handle missing data, while Le Morvan et al. (2021) train imputation and regression algorithms simultaneously using neural networks. Nevertheless, to the best of our knowledge, our work is the first to provide minimax optimality guarantees for multivariate nonparametric regression with missing data.

1.2 Notation

We conclude the introduction with some notation employed throughout the paper. For $n \in \mathbb{N}$, we define $[n] \coloneqq \{1, \ldots, n\}$, and for $a, b \in \mathbb{R}$, we let $a \wedge b \coloneqq \min\{a, b\}$ and $a \vee b \coloneqq \max\{a, b\}$. For $a, b \in \mathbb{R}^d$, we let $a \odot b \in \mathbb{R}^d$ denote the Hadamard product (i.e. coordinatewise product) of a and b. For $q \in [1, \infty)$ and $v = (v_1, \ldots, v_d)^\top \in \mathbb{R}^d$, we write $\|v\|_q \coloneqq (\sum_{j=1}^d |v_j|^q)^{1/q}$, as well as $\|v\|_0 \coloneqq \sum_{j=1}^d \mathbb{1}_{\{v_j \neq 0\}}$ and $\|v\|_\infty \coloneqq \max_{j \in [d]} |v_j|$. The all-ones vector is $\mathbf{1}_d \coloneqq (1, \ldots, 1)^\top \in \mathbb{R}^d$. If $(\mathcal{X}, \mathcal{A}, \mu)$ is a measure space and $f : \mathcal{X} \to \mathbb{R}$ is measurable, then we define $\|f\|_{L_q(\mu)} \coloneqq (\int_{\mathcal{X}} |f|^q d\mu)^{1/q}$, as well as $\|f\|_\infty \coloneqq \sup_{x \in \mathcal{X}} |f(x)|$. For function classes \mathcal{F} and \mathcal{G} , we define $\mathcal{F} \circ \mathcal{G} \coloneqq \{f \circ g : f \in \mathcal{F}, g \in \mathcal{G}\}$, and for $B \ge 0$, we define the truncation operator $T_B : \mathbb{R} \to [-B, B]$ by $T_B(y) \coloneqq (-B) \lor y \land B$. For a non-empty set \mathcal{S} , we say that $\{\mathcal{S}_1, \ldots, \mathcal{S}_K\}$ is a partition of \mathcal{S} if $\mathcal{S}_1, \ldots, \mathcal{S}_K$ are pairwise disjoint, non-empty subsets of \mathcal{S} whose union is \mathcal{S} . The sub-Gaussian norm of a realvalued random variable X is $\|X\|_{\psi_2} \coloneqq \inf\{t > 0 : \exp(X^2/t^2) \le 2\}$; its sub-exponential norm is $\|X\|_{\psi_1} \coloneqq \inf\{t > 0 : \exp(|X|/t) \le 2\}$. For positive sequences $(a_n), (b_n)$, we write $a_n \lesssim b_n$ if there exists a universal constant C > 0 such that $a_n \le Cb_n$ for all $n \in \mathbb{N}$; if we also have $b_n \lesssim a_n$, then we write $a_n \asymp b_n$.

2 Problem set-up and methodology

2.1 Problem set-up

Let $\mathcal{S} \subseteq \{0,1\}^d$, and let $(\mathbf{X}_i, \mathbf{\Omega}_i, Y_i)_{i=0}^n$ be independent and identically distributed random triples in $\mathbb{R}^d \times \mathcal{S} \times \mathbb{R}$. Without loss of generality, we assume that $\mathbb{P}(\mathbf{\Omega}_0 = \boldsymbol{\omega}) > 0$ for all $\boldsymbol{\omega} \in \mathcal{S}$. Define $\mathbb{R}^d_{\star} := (\mathbb{R} \cup \{\star\})^d$, equipped with the natural topology and σ -algebra described in Ma et al. (2024, Section 2.1). For $\boldsymbol{x} = (x_1, \ldots, x_d)^\top \in \mathbb{R}^d$ and $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_d)^\top \in \{0, 1\}^d$, define $\boldsymbol{x} \oplus \boldsymbol{\omega} \in \mathbb{R}^d_{\star}$ by $(\boldsymbol{x} \oplus \boldsymbol{\omega})_j := x_j$ if $\omega_j = 1$ and $(\boldsymbol{x} \oplus \boldsymbol{\omega})_j := \star$ if $\omega_j = 0$; i.e. we only observe the coordinates of \boldsymbol{x} for which the corresponding coordinates of the revelation vector $\boldsymbol{\omega}$ are one, and the missing entries are encoded by \star . For $i \in \{0, 1, \ldots, n\}$, let $\widetilde{\boldsymbol{X}}_i := \boldsymbol{X}_i \oplus \boldsymbol{\Omega}_i$ be the partially observed covariate vector; we emphasise that we do not assume independence between the covariate \boldsymbol{X}_i and observation pattern $\boldsymbol{\Omega}_i$, and indeed we allow the entries of \boldsymbol{X}_i to be MNAR. Let $\mathsf{Imp} : \mathbb{R}^d_\star \to \mathbb{R}^d$ be a potentially randomised imputation algorithm, where any randomness in the construction of Imp is independent of $(\boldsymbol{X}_i, \boldsymbol{\Omega}_i, Y_i)_{i=0}^n$, and let $\boldsymbol{Z}_i := \mathsf{Imp}(\widetilde{\boldsymbol{X}_i})$ for $i \in \{0, 1, \ldots, n\}$. For example, Imp can be zero (or mean) imputation, or other regression-based imputation algorithms trained on an independent dataset¹. We observe training data $(\boldsymbol{Z}_i, \boldsymbol{\Omega}_i, Y_i)_{i=1}^n$ and our goal is to

¹This is assumed for theoretical convenience. In practice (and indeed in our simulations in Section 4), the neural network and imputation algorithms would be trained on the same dataset.

predict Y_0 given a test point $(\mathbf{Z}_0, \mathbf{\Omega}_0)$. We remark that $(\mathbf{Z}_i, \mathbf{\Omega}_i, Y_i)_{i=0}^n$ are independent and identically distributed random triples in $\mathbb{R}^d \times \mathcal{S} \times \mathbb{R}$.

Writing \mathcal{G} for the set of Borel measurable functions from $\mathbb{R}^d \times \mathcal{S}$ to \mathbb{R} , we define the generalisation error of $\tilde{f} \in \mathcal{G}$ as

$$R(\widetilde{f}) \coloneqq \mathbb{E}_{(\mathbf{Z}_0, \mathbf{\Omega}_0, Y_0)} \left\{ \left(\widetilde{f}(\mathbf{Z}_0, \mathbf{\Omega}_0) - Y_0 \right)^2 \right\} \\ = \int_{\mathbb{R}^d \times \mathcal{S} \times \mathbb{R}} \left(\widetilde{f}(\mathbf{z}, \boldsymbol{\omega}) - y \right)^2 \mathrm{d}\mu_{\mathbf{Z}_0, \mathbf{\Omega}_0, Y_0}(\mathbf{z}, \boldsymbol{\omega}, y),$$

where $\mu_{\mathbf{Z}_0,\mathbf{\Omega}_0,Y_0}$ denotes the distribution of $(\mathbf{Z}_0,\mathbf{\Omega}_0,Y_0)$. Note that if \tilde{f} is an estimator depending on $(\mathbf{Z}_i,\mathbf{\Omega}_i,Y_i)_{i=1}^n$, then $R(\tilde{f})$ is random. The Bayes regression function f^* : $\mathbb{R}^d \times S \to \mathbb{R}$ is defined by

$$f^{\star}(\boldsymbol{z}, \boldsymbol{\omega}) \coloneqq \mathbb{E}(Y_0 \,|\, \boldsymbol{Z}_0 = \boldsymbol{z}, \boldsymbol{\Omega}_0 = \boldsymbol{\omega}).$$

This function, which satisfies $R(f^*) = \inf_{f \in \mathcal{G}} R(f)$, therefore depends on the imputation algorithm employed. For an estimator \tilde{f} , we measure its performance by its *excess risk* $\mathbb{E}\{R(\tilde{f})\} - R(f^*)$. Finally, we define the empirical risk of \tilde{f} as

$$\widehat{R}_n(\widetilde{f}) \coloneqq \frac{1}{n} \sum_{i=1}^n \left(\widetilde{f}(\boldsymbol{Z}_i, \boldsymbol{\Omega}_i) - Y_i \right)^2.$$

Example 1. Let $\mathbf{X}_0 = (X_{0,1}, \ldots, X_{0,d})^\top \sim \text{Unif}[-1, 1]^d$, and let $Y_0 = 3X_{0,1}^2 + \varepsilon_0$ where $\varepsilon_0 \sim N(0, 0.01)$ and $\varepsilon_0 \perp \mathbf{X}_0$. Suppose that $\mathbf{\Omega}_0 = (\Omega_{0,1}, \ldots, \Omega_{0,d})^\top \perp (\mathbf{X}_0, \varepsilon_0)$ satisfies $\Omega_{0,1}, \ldots, \Omega_{0,d} \stackrel{\text{iid}}{\sim} \text{Ber}(0.7)$, i.e. we have MCAR missingness where each coordinate is missing independently with probability 0.3. Let $\widetilde{\mathbf{X}}_0 \coloneqq \mathbf{X}_0 \otimes \mathbf{\Omega}_0$, and let $\text{Imp} : \mathbb{R}^d_* \to \mathbb{R}^d$ be the zero imputation algorithm (which is the same as mean imputation in this example) that replaces \star by zero, so $\mathbf{Z}_0 \coloneqq \text{Imp}(\widetilde{\mathbf{X}}_0) = \mathbf{X}_0 \odot \mathbf{\Omega}_0$. For $\mathbf{z} = (z_1, \ldots, z_d)^\top \in [-1, 1]^d$ and $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_d)^\top \in \{0, 1\}^d$, the Bayes regression function can be written as

$$f^{\star}(\boldsymbol{z}, \boldsymbol{\omega}) = \begin{cases} 3z_1^2 & \text{if } \omega_1 = 1\\ 1 & \text{if } \omega_1 = 0. \end{cases}$$

2.2 ReLU neural networks

In this section, we formally define classes of neural networks with ReLU activation function σ given by $\sigma(a) := a \lor 0$ for $a \in \mathbb{R}$. For a vector $\boldsymbol{x} = (x_1, \ldots, x_d)^\top \in \mathbb{R}^d$, we define

$$\boldsymbol{\sigma}(\boldsymbol{x}) \coloneqq \left(\sigma(x_1), \ldots, \sigma(x_d)\right)^{\top} \in [0, \infty)^d,$$

though we allow d to vary in different instances of this function without comment. Given $L \in \mathbb{N}$ and $\mathbf{p} = (p_0, p_1, \dots, p_{L+1}) \in \mathbb{N}^{L+2}$, the set of neural networks with architecture (L, \mathbf{p}) is defined as

$$\mathcal{F}(L,\boldsymbol{p}) \coloneqq \Big\{ \boldsymbol{f} : \mathbb{R}^{p_0} \to \mathbb{R}^{p_{L+1}} : \ \boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{A}_{L+1} \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_L \circ \boldsymbol{\sigma} \circ \cdots \circ \boldsymbol{A}_2 \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_1(\boldsymbol{x}), \text{ where } \boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{f}(\boldsymbol{x$$

$$\boldsymbol{A}_{\ell}(\boldsymbol{v}) = \boldsymbol{W}_{\ell}\boldsymbol{v} + \boldsymbol{b}_{\ell}, \ \boldsymbol{W}_{\ell} \in \mathbb{R}^{p_{\ell} \times p_{\ell-1}} \text{ and } \boldsymbol{b}_{\ell} \in \mathbb{R}^{p_{\ell}} \ \forall \ell \in [L+1] \Big\}.$$
(1)

Here, L is the number of hidden layers (or depth of the network), p_1, \ldots, p_L are the widths of the hidden layers, p_0 is the input dimension, p_{L+1} is the output dimension, W_1, \ldots, W_{L+1} are the weight matrices and b_1, \ldots, b_{L+1} are the bias vectors. For $f \in \mathcal{F}(L, p)$ with weight matrices $(W_\ell)_{\ell=1}^{L+1}$ and bias vectors $(b_\ell)_{\ell=1}^{L+1}$, we define $\Theta(f)$ to be the vector consisting of all the parameters of the neural network f, i.e. $\Theta(f) := (\operatorname{vec}(W_1)^\top, b_1^\top, \ldots, \operatorname{vec}(W_{L+1})^\top, b_{L+1}^\top)^\top \in \mathbb{R}^V$, where $V := \sum_{\ell=1}^{L+1} p_\ell(p_{\ell-1}+1)$ is the total number of parameters. For $s \in \mathbb{N}$, define

$$\mathcal{F}(L, \boldsymbol{p}, s) \coloneqq \left\{ \boldsymbol{f} \in \mathcal{F}(L, \boldsymbol{p}) : \|\boldsymbol{\Theta}(\boldsymbol{f})\|_0 \le s \right\}$$

to be the set of neural networks with architecture (L, \mathbf{p}) and sparsity s. Note that $\mathcal{F}(L, \mathbf{p}, s) = \mathcal{F}(L, \mathbf{p})$ for all $s \geq V$.

2.3 Pattern Embedded Neural Network estimators

Our aim is to estimate f^* by regressing $(Y_i)_{i=1}^n$ onto $(\mathbf{Z}_i, \mathbf{\Omega}_i)_{i=1}^n$ using neural networks, while borrowing strength across similar observation patterns. The number of observation patterns $|\mathcal{S}|$ may grow exponentially in the dimension d, so naive training of a neural network on $(\mathbf{Z}_i, \mathbf{\Omega}_i)_{i=1}^n$ may lead to overfitting. Motivated by ideas of categorical variable embedding and autoencoders in the deep learning literature (Hinton and Salakhutdinov, 2006; Mikolov et al., 2013), we propose a *pattern embedding* method to avoid this issue. The pattern embedding method maps $\boldsymbol{\omega} \in \mathcal{S} \subseteq \{0,1\}^d$ to a lower dimensional vector in \mathbb{R}^m for some $m \leq d$. Formally, for $r \in \{1,2,3\}$, let $L_r \in \mathbb{N}$, suppose that $\boldsymbol{p}_r =$ $(p_{r,0}, \ldots, p_{r,L_r+1}) \in \mathbb{N}^{L_r+2}$ satisfies $p_{3,0} = p_{1,L_1+1} + p_{2,L_2+1}$, and let $s \in \mathbb{N}$. Writing $p_{\text{in}} \coloneqq$ $p_{1,0} + p_{2,0}$ and $p_{\text{out}} \coloneqq p_{3,L_3+1}$, we then define our classes of *pattern embedded neural networks* (PENNs) by

$$\mathcal{F}_{\text{PENN}}\left(\begin{bmatrix} (L_1, \boldsymbol{p}_1) \\ (L_2, \boldsymbol{p}_2) \end{bmatrix}, s\right) \coloneqq \left\{\boldsymbol{f} : \mathbb{R}^{p_{\text{in}}} \to \mathbb{R}^{p_{\text{out}}} : \boldsymbol{f}(\cdot, \cdot) = \boldsymbol{f}_3\big(\boldsymbol{f}_1(\cdot), \boldsymbol{f}_2(\cdot)\big), \\ \text{where } \boldsymbol{f}_r \in \mathcal{F}(L_r, \boldsymbol{p}_r) \text{ for } r \in \{1, 2, 3\} \text{ and } \sum_{r=1}^3 \|\boldsymbol{\Theta}(\boldsymbol{f}_r)\|_0 \le s \right\}.$$
(2)

See Figure 2 for an illustration of a PENN. In the setting of Section 2.1, we seek an estimator \hat{f} that minimises the empirical risk $\hat{R}_n(f)$ over an appropriate PENN class as defined in (2). The optimisation involved in such a definition can be carried out easily using, for example, PyTorch (Paszke et al., 2019), though our theory in Section 3 allows for a residual optimisation error. In our applications, we will set $p_{1,0} = p_{2,0} = d$, $p_{2,L_2+1} = m$ where we refer to m as the *embedding dimension*, and $p_{3,L_3+1} = 1$. The function $f_2 = (f_{2,1}, \ldots, f_{2,m})^{\top} : \mathbb{R}^d \to \mathbb{R}^m$ is called the *embedding function* of f, and $f(z, \omega)$ denotes the output of $f \in \mathcal{F}_{\text{PENN}}$ evaluated at a test point $(z, \omega) \in \mathbb{R}^d \times S$.



Figure 2: An illustration of the class $\mathcal{F}_{\text{PENN}}\left(\begin{bmatrix} (L_1, \boldsymbol{p}_1) & (L_3, \boldsymbol{p}_3) \\ (L_2, \boldsymbol{p}_2) & (L_3, \boldsymbol{p}_3) \end{bmatrix}, s\right).$

3 Theoretical results

3.1 Oracle inequality

Our first main result is a general oracle inequality that controls the excess risk of a sparse neural network estimator in the missing data regression setting of Section 2.1. Let $L \in \mathbb{N}$ and $\boldsymbol{p} = (p_0, \ldots, p_{L+1})^\top \in \mathbb{N}^{L+2}$ with $p_0 = 2d$ and $p_{L+1} = 1$. By a *neural network estimator* in $\mathcal{F} \subseteq \mathcal{F}(L, \boldsymbol{p}, s)$, we mean a jointly measurable function $\tilde{f} : \mathbb{R}^d \times \mathcal{S} \times (\mathbb{R}^d \times \mathcal{S} \times \mathbb{R})^n \to \mathbb{R}$ such that for each $\mathcal{D} \coloneqq (\boldsymbol{z}_i, \boldsymbol{\omega}_i, y_i)_{i=1}^n \in (\mathbb{R}^d \times \mathcal{S} \times \mathbb{R})^n$, the function $\tilde{f}(\cdot, \cdot; \mathcal{D}) \in \mathcal{F}$. Where the data are clear from context, we often omit the final argument of \tilde{f} .

Theorem 1. In the setting of Section 2.1, assume that $||Y_0||_{\psi_2} \leq \xi$ for some $\xi \geq 1$. Let \tilde{f} be a neural network estimator in $\mathcal{F} \subseteq \mathcal{F}(L, \boldsymbol{p}, s)$ based on data $\mathcal{D} \coloneqq (\boldsymbol{Z}_i, \Omega_i, Y_i)_{i=1}^n$ and let $B_n \coloneqq \xi \sqrt{2 \log n}$. Then there exists a universal constant C > 0 such that for $n \geq 2$,

$$\mathbb{E}\left\{R(T_{B_n}\widetilde{f})\right\} - R(f^*) \le 2\mathbb{E}\left\{\widehat{R}_n(\widetilde{f}) - \inf_{f\in\mathcal{F}}\widehat{R}_n(f)\right\} + 2\inf_{f\in\mathcal{F}}\mathbb{E}\left\{\left(f(\mathbf{Z}_0, \mathbf{\Omega}_0) - f^*(\mathbf{Z}_0, \mathbf{\Omega}_0)\right)^2\right\} + \frac{C\xi^4\log(e\xi)\log^3 n \cdot \left(sL\log(es) + s\log(ed)\right)}{n}.$$

As a simple illustration of the bound on $||Y_0||_{\psi_2}$ in Theorem 1, if $Y_0 = f^0(\mathbf{X}_0) + \varepsilon_0$ where $||f^0||_{\infty} \leq \xi_1$ and $||\varepsilon_0||_{\psi_2} \leq \xi_2$, then $||Y_0||_{\psi_2} \leq \xi_1/\sqrt{\log 2} + \xi_2$. The upper bound on the excess risk in Theorem 1 is a sum of three terms, where the first term corresponds to the optimisation error, the second term represents the approximation error, and the last term corresponds to estimation error and reflects the complexity of the function class \mathcal{F} . If \tilde{f} is an empirical risk minimiser, then the optimisation error term vanishes. Importantly, we do not insist that $f^* \in \mathcal{F}$, though if it does, then we may take the approximation error term to be zero. The estimation error term is of order sL/n up to poly-logarithmic factors. In general, there is a trade-off between the approximation and estimation error terms that is akin to a bias-variance trade-off: more complex classes \mathcal{F} will have smaller approximation error, but the price to be paid is through larger values of s and L in the estimation error term. Related results to Theorem 1 include those of Schmidt-Hieber (2020), who has an additional boundedness assumption on the parameter space, and Kohler and Langer (2021), who work with fully dense neural networks. The key to our proof is a new bound on the Vapnik-Chervonenkis dimension and covering numbers of the $\mathcal{F}(L, \mathbf{p}, s)$ class given in Proposition 5.

3.2 Minimax rate under a piecewise smoothness assumption

The general theory of Section 3.1 allows us to provide an explicit upper bound on the excess risk of our PENN estimator under a piecewise smoothness assumption on the regression function; see Theorem 3 below. Our bound reveals that improved bounds are achievable in settings where the regression function is a composition of (smooth) functions that depend only on a subset of the variables, thereby providing a sense in which the estimator is able to evade the curse of dimensionality.

In addition to tail conditions on the covariates and response, Assumption 1 introduces a structural condition on the Bayes regression function, recognising that this regression function may be the same for certain different missingness patterns.

Assumption 1. Assume that each coordinate of \mathbb{Z}_0 is sub-exponential and that Y_0 is sub-Gaussian, i.e. there exist $\xi_1, \xi_2 > 0$ such that $||Z_{0,j}||_{\psi_1} \leq \xi_1$ for all $j \in [d]$ and $||Y_0||_{\psi_2} \leq \xi_2$. Further assume that there exist $K \leq |\mathcal{S}|$, a partition $\{\mathcal{S}_1, \ldots, \mathcal{S}_K\}$ of \mathcal{S} , and functions $f^{\mathcal{S}_1}, \ldots, f^{\mathcal{S}_K} : \mathbb{R}^d \to \mathbb{R}$ such that²

$$f^{\star}(\boldsymbol{z},\boldsymbol{\omega}) = \sum_{k=1}^{K} f^{\mathcal{S}_{k}}(\boldsymbol{z}) \mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}}$$
(3)

for all $\boldsymbol{z} \in \mathbb{R}^d$ and $\boldsymbol{\omega} \in \mathcal{S}$. For $k \in [K]$, we write $\pi_k := \mathbb{P}(\boldsymbol{\Omega}_0 \in \mathcal{S}_k)$ and $n_k := n\pi_k$.

The tail condition on \mathbb{Z}_0 allows our covariates to be unbounded, in contrast to much of the literature on nonparametric regression. Of course, any Bayes regression function f^* satisfies (3) with $K = |\mathcal{S}|$. However, K may be much smaller than $|\mathcal{S}|$; e.g. in Example 1 we have $|\mathcal{S}| = 2^d$ and K = 2, since we may take $\mathcal{S}_1 = \{ \boldsymbol{\omega} \in \{0, 1\}^d : \omega_1 = 1 \}$ and $\mathcal{S}_2 = \{ \boldsymbol{\omega} \in \{0, 1\}^d : \omega_1 = 0 \}$. In fact, given an arbitrary partition $\{\mathcal{S}_1, \ldots, \mathcal{S}_K\}$ of \mathcal{S} , there exists a Bayes regression function f^* that satisfies Assumption 1. Indeed, let $\boldsymbol{X} = (X_1, \ldots, X_d)^\top \sim \text{Unif}[0, 1]^d$ and let $Y = \sum_{k=1}^K \mathbbm{1}_{\{X_1 \ge (k-1)/K\}} + \varepsilon$ where $\varepsilon \sim N(0, 1)$

²More formally, since f^* is only defined up to sets having zero measure under the distribution of $(\mathbf{Z}_0, \mathbf{\Omega}_0)$, we ask that there exists a version of f^* for which the statement in Assumption 1 holds.

is independent of \mathbf{X} . Further define $\mathbf{\Omega}$ by $\mathbf{\Omega} | \mathbf{X} \sim \text{Unif}(\mathcal{S}_k)$ when $X_1 \in [(k-1)/K, k/K)$. Then $f^*(\mathbf{z}, \boldsymbol{\omega}) = k$ for all $\mathbf{z} \in [0, 1]^d$ whenever $\boldsymbol{\omega} \in \mathcal{S}_k$. It turns out that Assumption 1 may be weakened to only requiring that (3) holds approximately; see the discussion following Theorem 3 below.

We now introduce further sparsity and smoothness assumptions that will be imposed on each f^{S_k} in (3).

Definition 1. Let $D \subseteq \mathbb{R}^d$, $f : D \to \mathbb{R}$ and $\mathcal{J} \subseteq [d]$. We say that f depends only on the coordinates in \mathcal{J} if there exists $g : \mathbb{R}^{|\mathcal{J}|} \to \mathbb{R}$ such that for all $\boldsymbol{x} = (x_1, \ldots, x_d)^\top \in D$, we have $f(\boldsymbol{x}) = g(\boldsymbol{x}_{\mathcal{J}})$ where $\boldsymbol{x}_{\mathcal{J}} \coloneqq (x_j)_{j \in \mathcal{J}}$; when $\mathcal{J} = \emptyset$, this means that f is constant. For $t \in [d] \cup \{0\}$, we say that f depends only on t variables if there exists $\mathcal{J} \subseteq [d]$ with $|\mathcal{J}| = t$ such that f depends only on the coordinates in \mathcal{J} .

We use multi-index notation for partial derivatives, whereby for $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)^\top \in \mathbb{N}_0^d$ and an $\|\boldsymbol{\alpha}\|_1$ -times differentiable real-valued function f defined on a subset of \mathbb{R}^d , we set $\partial^{\boldsymbol{\alpha}} f := \partial^{\alpha_1} \cdots \partial^{\alpha_d} f$. It is also convenient to write $\boldsymbol{x}^{\boldsymbol{\alpha}} := \prod_{j=1}^d x_j^{\alpha_j}$ for $\boldsymbol{x} = (x_1, \ldots, x_d)^\top \in \mathbb{R}^d$.

Definition 2. For $\beta, \gamma > 0$, $d \in \mathbb{N}$, $t \in [d] \cup \{0\}$ and $D \subseteq \mathbb{R}^d$, we write $\beta_0 \coloneqq \lceil \beta \rceil - 1$ and define the class of (β, γ) -Hölder functions that depend only on t variables as

$$\mathcal{H}_{t}^{\beta}(D,\gamma) \coloneqq \left\{ f: D \to \mathbb{R} : f \text{ depends only on } t \text{ variables, } f \text{ is } \beta_{0}\text{-times differentiable,} \right. \\ \left. \max_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \le \beta_{0}} \|\partial^{\boldsymbol{\alpha}}f\|_{\infty} \le \gamma, \max_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} = \beta_{0}} \sup_{\boldsymbol{x} \neq \boldsymbol{y} \in D} \frac{|\partial^{\boldsymbol{\alpha}}f(\boldsymbol{x}) - \partial^{\boldsymbol{\alpha}}f(\boldsymbol{y})|}{\|\boldsymbol{x} - \boldsymbol{y}\|_{2}^{\beta - \beta_{0}}} \le \gamma \right\}.$$

Our functions f^{S_k} will be assumed to be compositions of vector-valued functions whose corresponding component functions belong to these Hölder classes.

Definition 3. Let $q \in \mathbb{N}$, $\boldsymbol{d} = (d_r)_{r=1}^{q+1} \in \mathbb{N}^{q+1}$ with $d_1 = d$ and $d_{q+1} = 1$, $\boldsymbol{t} = (t_r)_{r=1}^q \in \mathbb{N}_0^q$ with $t_r \in [d_r] \cup \{0\} \ \forall r \in [q], \ \boldsymbol{\beta} = (\beta_r)_{r=1}^q \in (0, \infty)^q$ and $\boldsymbol{\gamma} = (\gamma_r)_{r=1}^q \in (0, \infty)^q$. We define $\mathcal{H}_{\text{comp}}(q, \boldsymbol{d}, \boldsymbol{t}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ to be the class of functions $f : \mathbb{R}^d \to \mathbb{R}$ of the form

$$f(\boldsymbol{x}) = \boldsymbol{g}_q \circ \cdots \circ \boldsymbol{g}_1(\boldsymbol{x})$$

where³

$$\boldsymbol{g}_{r} = (g_{r,1}, \dots, g_{r,d_{r+1}})^{\top} : \mathbb{R}^{d_{r}} \to \mathbb{R}^{d_{r+1}} \text{ and}$$
$$g_{r,j} \in \mathcal{H}_{t_{r}}^{\beta_{r}} (\mathbb{R}^{d_{r}}, \gamma_{r}) \text{ for all } r \in [q], \ j \in [d_{r+1}].$$

Assumption 2. For $k \in [K]$, let $q_k \in \mathbb{N}$, let $d^{(k)} = (d_r^{(k)})_{r=1}^{q_k+1} \in \mathbb{N}^{q_k+1}$ with $d_1^{(k)} = d$ and $d_{q_k+1}^{(k)} = 1$, let $t^{(k)} = (t_r^{(k)})_{r=1}^{q_k} \in \mathbb{N}_0^{q_k}$ with $t_r^{(k)} \in [d_r^{(k)}] \cup \{0\}$ for all $r \in [q_k]$, let $\beta^{(k)} = (\beta_r^{(k)})_{r=1}^{q_k} \in (0, \infty)^{q_k}$ and let $\gamma^{(k)} = (\gamma_r^{(k)})_{r=1}^{q_k} \in (0, \infty)^{q_k}$. We assume that $f^{\mathcal{S}_k} \in \mathcal{H}_{\text{comp}}(q_k, d^{(k)}, t^{(k)}, \beta^{(k)}, \gamma^{(k)})$ for each $k \in [K]$.

Our final preliminary is to introduce the notion of \mathcal{F} -separability of a partition of \mathcal{S} .

³Although g_q is a real-valued function, we write it with a bold letter for convenience.

Definition 4. Let $S \subseteq \{0,1\}^d$, let $m \in \mathbb{N}$, let \mathcal{F} be a class of functions from $\mathbb{R}^d \to \mathbb{R}^m$ and let $\{S_1, \ldots, S_K\}$ be a partition of S. We say that $\{S_1, \ldots, S_K\}$ is \mathcal{F} -separable if there exist $\mathbf{f} \in \mathcal{F}, \mathbf{v}_1, \ldots, \mathbf{v}_K \in \mathbb{R}^m$ and $\epsilon > 0$ such that

(i)
$$\|\boldsymbol{f}(\boldsymbol{\omega}) - \boldsymbol{v}_k\|_{\infty} \leq \epsilon/2$$
 for all $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_k$;

(ii)
$$\|\boldsymbol{v}_k - \boldsymbol{v}_{k'}\|_{\infty} \geq 2\epsilon$$
 for all $k \neq k'$

In this case, we say that $\{S_1, \ldots, S_K\}$ is separated by f.

Thus $\{S_1, \ldots, S_K\}$ is separated by \boldsymbol{f} if the scale over which \boldsymbol{f} varies on each S_k is small by comparison with its variability across different cells of the partition. Proposition 2 below guarantees that an arbitrary partition $\{S_1, \ldots, S_K\}$ of S is $\mathcal{F}(2, \boldsymbol{p})$ -separable for a suitable choice of \boldsymbol{p} . Moreover, if the partition is defined by a small number of coordinates (as in part (a)) or a small number of halfspaces (as in part (b)), then the partition may be separated by a class of neural networks with fewer parameters.

Proposition 2. Let $S \subseteq \{0,1\}^d$ be such that $|S| \ge 2$, let $K \in \{2,\ldots,|S|\}$ and let $\{S_1,\ldots,S_K\}$ be a partition of S.

(a) Suppose that there exists $\mathcal{J} \subseteq [d]$ such that $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \mathcal{S}$ belong to the same cell of the partition whenever $\boldsymbol{\omega}_{\mathcal{J}} = \boldsymbol{\omega}'_{\mathcal{J}}$. Let $\mathcal{S}_{\mathcal{J}} \coloneqq \{\boldsymbol{\omega}_{\mathcal{J}} : \boldsymbol{\omega} \in \mathcal{S}\} \subseteq \{0,1\}^{|\mathcal{J}|}$. Define

$$p_* \coloneqq 2 [|\mathcal{S}_{\mathcal{J}}|^{1/2}] \quad and \quad \boldsymbol{p} \coloneqq (d, p_*, p_*, 1)^\top \in \mathbb{N}^4.$$

Then $\{S_1, \ldots, S_K\}$ is $\mathcal{F}(2, \mathbf{p})$ -separable. In particular, if S_1, \ldots, S_K is an arbitrary partition of S, then we may take $\mathcal{J} = [d]$.

(b) Suppose that for each $k \in [K]$, there exist $P_k \in \mathbb{N}$ and $(\boldsymbol{v}_{\ell}^{(k)}, b_{\ell}^{(k)})_{\ell \in [P_k]} \in (\mathbb{R}^d \times \mathbb{R})^{P_k}$ such that $\mathcal{S}_k = \{ \boldsymbol{\omega} \in \mathcal{S} : \boldsymbol{\omega}^\top \boldsymbol{v}_{\ell}^{(k)} \le b_{\ell}^{(k)} \text{ for all } \ell \in [P_k] \}$. Define

$$\boldsymbol{p} \coloneqq \left(d, 2\sum_{k=1}^{K} P_k, K, 1\right)^{\top} \in \mathbb{N}^4.$$

Then $\{S_1, \ldots, S_K\}$ is $\mathcal{F}(2, \boldsymbol{p})$ -separable.

For $k \in [K]$ and $r \in [q_k]$, we define

$$\bar{\beta}_r^{(k)} \coloneqq \beta_r^{(k)} \prod_{\ell=r+1}^{q_k} (\beta_\ell^{(k)} \wedge 1).$$

We also set⁴ $r_*^{(k)} \coloneqq \operatorname{sargmax}_{r \in [q_k]} t_r^{(k)} / \bar{\beta}_r^{(k)}$ and let $\bar{\beta}_*^{(k)} \coloneqq \bar{\beta}_{r_*^{(k)}}^{(k)}$, $\beta_*^{(k)} \coloneqq \beta_{r_*^{(k)}}^{(k)}$, $t_*^{(k)} \coloneqq t_{r_*^{(k)}}^{(k)}$ and $\gamma_*^{(k)} \coloneqq \gamma_{r_*^{(k)}}^{(k)}$. Finally, then, we can state our main result:

⁴Here, sargmax denotes the smallest element of the argmax set.

Theorem 3. Suppose that Assumptions 1 and 2 hold. Suppose further that $\{S_1, \ldots, S_K\}$ is $\mathcal{F}(L_2, \mathbf{p}_2, s_2)$ -separable, and let $B_n \coloneqq \xi_2 \sqrt{2 \log n}$. Then there exist $L_1, L_3 \in \mathbb{N}, \mathbf{p}_1 \in \mathbb{N}^{L_1+2}, \mathbf{p}_3 \in \mathbb{N}^{L_3+2}$ and $s \in \mathbb{N}$ such that, writing

$$\mathcal{F} \coloneqq \mathcal{F}_{\text{PENN}} \left(\begin{bmatrix} (L_1, \boldsymbol{p}_1) & (L_3, \boldsymbol{p}_3) \\ (L_2, \boldsymbol{p}_2) & (L_3, \boldsymbol{p}_3) \end{bmatrix}, s \right),$$

and letting \hat{f} denote any neural network estimator in \mathcal{F} based on data $\mathcal{D} \coloneqq (\mathbf{Z}_i, \mathbf{\Omega}_i, Y_i)_{i=1}^n$, we have for $n \geq 2$ that

$$\mathbb{E}\left\{R(T_{B_n}\widehat{f})\right\} - R(f^{\star}) \le C\left\{\sum_{k=1}^{K} \pi_k n_k^{-2\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})} + \frac{s_2 \log s_2}{n}\right\} \cdot (\log n)^{2\max_{k \in [K]} \bar{\beta}_1^{(k)} \vee 6} + 2\mathbb{E}\left\{\widehat{R}_n(\widehat{f}) - \inf_{f \in \mathcal{F}} \widehat{R}_n(f)\right\},$$

where C > 0 does not depend on n, $(\pi_k)_{k=1}^K$ or s_2 .

In order to understand the main messages of Theorem 3, first suppose that we are able to compute the empirical risk minimiser exactly, so that the optimisation error $\mathbb{E}\{\widehat{R}_n(\widehat{f}) - \inf_{f \in \mathcal{F}} \widehat{R}_n(f)\}$ is zero. Then the excess risk of a truncated PENN estimator is controlled by the sum of two interpretable terms. The first of these would be the minimax rate for estimating the Bayes regression function if the partition $\{S_1, \ldots, S_K\}$ of Assumption 1 were known, up to a poly-logarithmic factor in the sample size. It comprises a weighted average over $k \in [K]$ of the minimax rates of estimating the Bayes regression function $f^{\mathcal{S}_k}$ on the *k*th cell of the partition, with effective sample size n_k , effective smoothness $\overline{\beta}_*^{(k)}$ and effective dimension $t_*^{(k)}$; see Theorem 4 below. As an attraction of this weighted average form, we see that the *k*th summand

$$\pi_k n_k^{-2\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)}+t_*^{(k)})} = \pi_k^{t_*^{(k)}/(2\bar{\beta}_*^{(k)}+t_*^{(k)})} n^{-2\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)}+t_*^{(k)})} \to 0$$

as $\pi_k \to 0$, for fixed *n*; thus, rarely observed missingness patterns have negligible effect on the excess risk. The second term in the bound in Theorem **3** is the additional error incurred due to the fact that the $\mathcal{F}(L_2, \mathbf{p}_2, s_2)$ -separable partition $\{S_1, \ldots, S_K\}$ is unknown. In cases where the partition is defined by a small number of coordinates or a small number of halfspaces, we can expect this second term to be dominated by the first. In particular, this occurs in the setting of Proposition 2(a) with $|\mathcal{S}_{\mathcal{J}}| \lesssim \sum_{k=1}^{K} n_k^{t_k^{(k)}/(2\bar{\beta}_k^{(k)}+t_k^{(k)})}$ and in the setting of Proposition 2(b) with $K \sum_{k=1}^{K} P_k \lesssim \sum_{k=1}^{K} n_k^{t_k^{(k)}/(2\bar{\beta}_k^{(k)}+t_k^{(k)})}$. Theorem **3** also accounts for the setting where the empirical risk minimiser is not computed exactly, in which case we incur an additional optimisation error.

Writing $\boldsymbol{p}_1 = (d, p_{1,*}, \dots, p_{1,*}, p_{1,L_1+1})^\top \in \mathbb{N}^{L_1+2}$ and $\boldsymbol{p}_3 = (d, p_{3,*}, \dots, p_{3,*}, 1)^\top \in \mathbb{N}^{L_3+2}$, from the proof of Theorem 3, we can see that it suffices to choose L_1 and L_3 of constant order in the sample size, together with

$$p_{1,*} \asymp \sum_{k=1}^{K} n_k^{t_*^{(k)}/(4\bar{\beta}_*^{(k)} + 2t_*^{(k)})} \log n, \quad p_{3,*} \asymp \sum_{k=1}^{K} n_k^{t_*^{(k)}/\{(L_3 - 3)(2\bar{\beta}_*^{(k)} + t_*^{(k)})\}}$$

$$s \simeq s_2 + \sum_{k=1}^{K} n_k^{t_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})} \log^2 n.$$

Moreover, by (P1) in Appendix A and the fact that the bound in Theorem 1 does not depend on the widths of the hidden layers, we may increase $p_{1,*}$ and $p_{3,*}$ without affecting the bound in Theorem 3. This means that Theorem 3 applies to over-parametrised neural networks of appropriate sparsity.

From the proof, we see that the quantity C > 0 in Theorem 3 may be chosen as the maximum over $k \in [K]$ of polynomial functions of $d, d^{(k)}, t^{(k)}$, whose degrees depend only on $\boldsymbol{\beta}^{(k)}$ and whose coefficients depend only on $p_{2,L_2+1}, L_1, L_2, L_3, \xi_1, \xi_2, d, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)}$. In particular, if $\pi_k = 1/K$, $\beta_r^{(k)} = \beta \ge 1$ and $t_r^{(k)} = t$ for all $k \in [K]$ and $r \in [q_k]$, then the first term in the upper bound simplifies to $(K/n)^{2\beta/(2\beta+t)}$, up to a poly-logarithmic factor in n that does not depend on K. A further observation from the proof is that Assumption 1 can be weakened so that the Bayes regression function is not required to be identical on different elements within the same S_k ; in fact, it suffices to assume that $|f^*(\boldsymbol{z}, \boldsymbol{\omega}) - f^{S_k}(\boldsymbol{z})| \lesssim n_k^{-\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})}$ for $\boldsymbol{z} \in \mathbb{R}^d$, $k \in [K]$ and $\boldsymbol{\omega} \in S_k$.

We complement the upper bound of Theorem 3 with a corresponding minimax lower bound, that uses the notation of Assumptions 1 and 2, as well as the definitions that immediately follow Proposition 2. Further for $k \in [K]$, define $\mathcal{J}^{(k)} := \{j \in [d] : \omega_j =$ 1 for all $\omega \in \mathcal{S}_k\}$.

Theorem 4. Let \mathcal{P} be the set of all distributions of $(\mathbf{Z}_0, \Omega_0, Y_0)$ where $\mathbf{Z}_0 = \text{Imp}(\mathbf{X}_0 \oplus \Omega_0)$, and $(\mathbf{X}_0, \Omega_0, Y_0)$ satisfies Assumptions 1 and 2 with $\xi_1, \xi_2 \geq 1$. Suppose further that for some $j_* \in [d]$, we have

$$t_*^{(k)} \le \min\{d_1^{(k)}, \dots, d_{r_*^{(k)}}^{(k)}\} \land |\mathcal{J}^{(k)} \setminus \{j_*\}| \land \gamma_*^{(k)}$$

for all $k \in [K]$. For $P \in \mathcal{P}$, let $f^* \equiv f_P^* \coloneqq \mathbb{E}_P(Y_0 | \mathbf{Z}_0, \mathbf{\Omega}_0)$, and let $\widehat{\mathcal{F}}$ be the set of all estimators of f^* based on a sample of size n, i.e. the set of Borel measurable functions from $\mathbb{R}^d \times \mathcal{S} \times (\mathbb{R}^d \times \mathcal{S} \times \mathbb{R})^n$ to \mathbb{R} . Then there exists c > 0, depending only on ξ_2 and $(\overline{\beta}_*^{(k)}, \beta_*^{(k)}, t_*^{(k)})_{k=1}^K$, such that

$$\inf_{\widehat{f}\in\widehat{\mathcal{F}}} \sup_{P\in\mathcal{P}} \mathbb{E}_{P^{\otimes n}} \left\{ R(\widehat{f}) - R(f^{\star}) \right\} \ge c \sum_{k=1}^{K} \pi_k n_k^{-2\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})}$$

for all $n \in \mathbb{N}$.

Theorem 4 reveals that the main term in the upper bound in Theorem 3, namely the first estimation error term, is minimax optimal in n and π_1, \ldots, π_K , up to a polylogarithmic factor in n.

4 Simulations

In this section, we study the empirical performance of the Pattern Embedded Neural Network (PENN) estimator on simulated, semi-synthetic and real data. Since the PENN estimator can be used in conjunction with any imputation technique, we consider (columnwise) mean imputation (MI), MissForest imputation (MF) (Stekhoven and Bühlmann, 2012) and a python implementation of Multiple Imputation by Chained Equations (MICE) (van Buuren and Groothuis-Oudshoorn, 2011) called IterativeImputer (II) from scikit-learn (Pedregosa et al., 2011). In each case, we compare PENN with standard neural networks that do not incorporate the revelation vectors as covariates.

For all of our numerical experiments, we divide the data into training, validation and test sets. We fit the neural networks on the training data by first running the stochastic optimisation method Adam (Kingma and Ba, 2015) for 10 epochs, and then keeping a proportion $\lambda \in \{0.1, 0.2, 0.4, 0.8\}$ of weights of largest magnitude (in addition to all of the bias vectors) to obtain a sparse network. Following the recommendation of Liu et al. (2019), we then randomly reinitialise the non-zero parameters after this pruning step using the PyTorch function torch.nn.init.kaiming_uniform_ and retrain the sparse model via Adam. Early stopping (Prechelt, 2002) is incorporated, so that the training process for each value of λ is terminated once the validation loss fails to decrease by at least 0.001 in 10 epochs. The tuning parameter λ is then chosen to minimise the average loss on the validation set, and finally the performance of this selected estimator on the test set is reported. We remark that although our theory requires the data used for imputation to be independent of the training data, in our simulations we train the imputation algorithms on the whole dataset (training, validation and test sets) and then impute the missing entries.

4.1 Simulated data

For each of our experiments on simulated data, we take d = 20, with a training set of size n = 10,000, and validation and test sets each of size 5,000. As explained above, the choice of sparsity s is determined by the value of λ chosen on the validation set. In our implementation:

• PENN uses

$$\mathcal{F}_{\text{PENN}}\left(\begin{bmatrix} \left(3, \, (d, 70, 70, 70, 70)\right) \\ \left(2, \, (d, 30, 30, 3)\right) & \left(3, \, (73, 70, 70, 70, 1)\right) \end{bmatrix}, s \right);$$

• NN uses $\mathcal{F}(6, (d, 70, 70, 70, 70, 70, 70, 1), s)$.

Thus, by (P2), the PENN class above has six hidden layers in total, not including the embedding function f_2 ; see (2) and Figure 2. We therefore compare it with a standard neural network with six hidden layers of the same width. Our data generating mechanisms were chosen as follows:

Model 1: $X_0 \sim \text{Unif}[-1,1]^d$, $Y_0 = \exp(X_{0,1} + X_{0,2}) + 4X_{0,3}^2 + \varepsilon_0$, $\varepsilon_0 \sim N(0,0.25)$, $\varepsilon_0 \perp \perp X_0$, $\Omega_0 \perp \perp (X_0, \varepsilon_0)$ and $\Omega_{0,j} \stackrel{\text{iid}}{\sim} \text{Ber}(0.7)$ for $j \in [d]$. **Model 2:** $X_0 \sim \text{Unif}[-1,1]^d$, $Y_0 = 2\sin(2X_{0,1} + 2X_{0,2}) + 2X_{0,3} + \varepsilon_0$, $\varepsilon_0 \sim N(0,0.25)$, $\varepsilon_0 \perp X_0$, $\Omega_{0,j} \stackrel{\text{iid}}{\sim} \text{Ber}(0.7)$ independent of X_0 for $j \notin \{2,3\}$, and $\Omega_{0,j} = \mathbb{1}_{\{X_{0,j} \leq 0.4\}}$ for $j \in \{2,3\}$.

Model 3: As for Model 1 except that we set $X_{0,1} = \sqrt{X_{0,4} + 1} - 0.7 + \text{Unif}[-0.3, 0.3]$ and $X_{0,3} = 0.7X_{0,5} + \text{Unif}[-0.3, 0.3].$

Model 4: As for Model 2 except that we set $X_{0,1} = \sqrt{X_{0,4} + 1} - 0.7 + \text{Unif}[-0.3, 0.3]$ and $X_{0,3} = 0.7X_{0,5} + \text{Unif}[-0.3, 0.3].$

Thus, Models 1 and 3 are MCAR, while Models 2 and 4 are MNAR. In Models 1 and 2, the different coordinates of the covariates are independent, while in Models 3 and 4 we introduce positive correlation between $X_{0,1}$ and $X_{0,4}$, and between $X_{0,3}$ and $X_{0,5}$. In each case, the true regression function depends only on the first three components of X_0 ; this facilitates computation of the Bayes risk via Monte Carlo integration.

The results of our simulations over 100 repetitions are presented in Figure 3. We observe that for all four models, PENN improves the performance of each imputation technique, often dramatically. In particular, it is able to substantially remedy the strikingly poor performance of NN_MF for Models 2 and 4. It seems that, to some extent, more successful imputation techniques for the vanilla neural network estimator tend to yield more successful PENN estimators.

4.2 Semi-synthetic data

The aim of this subsection is to consider two semi-synthetic datasets. By this we mean that we take two real datasets without missingness, and artificially introduce missingness according to two different prescribed mechanisms that we articulate below. An attraction of this approach is that it allows us to study the effects of different types of missingness on real data.

Bank loan dataset: The bank loan dataset from https://www.kaggle.com/datasets/ udaymalviya/bank-loan-data is a complete dataset with d = 13 and total sample size 45,000. The response variable is binary with a one indicating that the loan was paid off, while the covariates consist of information on the credit history and other personal characteristics of the debtor, as well as information about the loan, including the interest rate loan_int_rate and the loan amount as a percentage of the individual's income loan_percent_income. Here we again consider MCAR missingness with homogeneous observation probability 0.7 in each coordinate. For MNAR missingness, the probability of observing loan_int_rate given that loan_int_rate = x is taken to be $\frac{1}{e^{x-13}+1}$, while the probability of observing loan_percent_income given that loan_percent_income = x is $\frac{1}{e^{6x-2}+1}$, and the other features are MCAR with homogeneous observation probability 0.7. MI imputes missing entries for numeric variables by their means, and categorical variables by their mode. MF has no issues with categorical variables, whereas II cannot handle categorical variables, so we do not include it for this dataset. After imputation,



Figure 3: Estimates of excess risks for simulated data Models 1–4. The PENN estimators are shown in red, with the vanilla neural networks in blue; on the x-axis, the abbreviation of the imputation technique appears after the underscore symbol.

we apply one-hot encoding for categorical variables, which inflates the dimensionality of the covariates to 27 (although the revelation vectors are still in 13 dimensions).

MNIST dataset: The MNIST dataset (LeCun, Cortes and Burges, 1998) in PyTorch consists of 70,000 greyscale images of handwritten digits from 0 to 9, with each image represented as a 28 × 28 pixel grid. We normalise the greyscale so that it takes values in [0, 1], and the goal is to predict the labels based on the images. To reflect the fact that missingness is likely to be correlated among nearby pixels, we partition each image uniformly into 49 blocks, each of 4×4 pixels. In our MCAR setting, each block is observed independently with probability 0.5; in our MNAR setting, for each $\ell \in \{0, \ldots, 9\}$, we draw a random vector $\mathbf{p}_{\ell} = (p_{\ell,1}, \ldots, p_{\ell,49})^{\top}$ uniformly from $[0.2, 0.8]^{49}$, and generate $\Omega_{i,j} \sim \text{Ber}(p_{\ell_{i,j}})$ independently for $j \in [49]$, where ℓ_i is the label for the *i*th observation. In other words, in our MNAR setting, the missingness probability varies across different labels, while the overall expected missingness probability is 0.5 for all observations. For the MNIST dataset, we adopt only mean imputation, since MissForest and MICE are not well-suited for image data. Figure 4 presents the first 16 images from the MNIST dataset with MCAR and MNAR missingness.



(a) MCAR missingness



Figure 4: First 16 images (after mean imputation) from the MNIST dataset with MCAR and MNAR missingness, together with the true labels above each panel.

For the bank loan dataset,

• PENN uses

$$\mathcal{F}_{\text{PENN}}\left(\begin{bmatrix} (3, (d, 100, 100, 100, 100)) \\ (2, (d, 30, 30, 3)) \end{bmatrix} (3, (103, 100, 100, 100, 1)) \end{bmatrix}, s \right).$$

• NN uses $\mathcal{F}(6, (d, 100, 100, 100, 100, 100, 100, 1), s)$.

On the other hand, in image classification tasks, convolutional neural networks (LeCun et al., 1998; Krizhevsky, Sutskever and Hinton, 2012) have achieved great successes; see also Fan, Ma and Zhong (2020, Section 3.1) for an introduction. Thus, for the MNIST dataset, we use a similar architecture as above, but replace each of the first two hidden layers in f_1 in the definition of PENN (2), and each of the first two hidden layers in NN, by a composition of a 2D convolutional layer, ReLU activation layer and max pooling layer. Details of the architectures of PENN and NN with convolutional layers used in our simulations can be found at https://github.com/tianyima2000/DNN_missing_data. Again, we compare a PENN with a neural network of the same architecture except that the embedding function has been removed. Finally, we use the cross-entropy loss, which is the negative log-likelihood of the relevant multinomial distribution, for these two classification tasks.

For both datasets, and for each of 50 repetitions, we randomly split the dataset into training, validation and test sets with sizes in the ratio 8:1:1 after introducing the missingness. We measure the performance of the algorithms via their misclassification error (MCE) on the test set. The results for the two different datasets are presented in Figures 5 and 6. For MCAR missingness, the PENN estimator is at least comparable with the vanilla NN estimator in terms of performance, while for MNAR missingness, it provides a significant improvement.



Figure 5: Misclassification error (MCE) for bank loan dataset with different missingness mechanisms.

4.3 Real data

We now turn to two further real datasets that already have some missing values.

Credit score prediction dataset: The credit score prediction dataset from https://www.kaggle.com/datasets/prasy46/credit-score-prediction has d = 304 and total



Figure 6: Misclassification error (MCE) for MNIST dataset with different missingness mechanisms.

sample size of 100,000. The response variable of interest is the credit rating, quantified as a positive integer ranging from 300 to 839 in the available data. There are 41 columns with missingness, and the observation probabilities for the columns with missingness are given in Figure 7(a).



Figure 7: Histograms of the observation probabilities for the columns with missingness in credit score prediction dataset and public procurement dataset.

Public procurement dataset: For the public procurement dataset from https://www.openml.org/search?type=data&status=active&id=42163, we use only the numeric variables, yielding d = 25 variables and a total sample size of 565,163. The original

response award_value_euro records the values of public procurement contracts (in Euros), ranging from -1 to over thirteen billion, so we use a log transformation and predict $log(award_value_euro + 2)$ instead. There are 14 columns with missingness, and the observation probabilities for the columns with missingness are given in Figure 7(b).

For our real data,

• PENN uses

$$\mathcal{F}_{\text{PENN}}\left(\begin{bmatrix} (3, (d, 100, 100, 100, 100)) \\ (2, (d, 30, 30, 3)) \end{bmatrix} (3, (103, 100, 100, 100, 1)) \end{bmatrix}, s \right)$$

• NN uses $\mathcal{F}(6, (d, 100, 100, 100, 100, 100, 100, 1), s)$.

Again, for each dataset and for each of 10 repetitions, we randomly split the dataset into training, validation and test sets with sizes in the ratio 8:1:1. The proportions of unexplained variance (PUV), defined as the ratio of the mean squared prediction error on the test set to the sample variance of the response on this same test set, are presented in Figure 8 for the credit score data and Figure 9 for the public procurement data. The improvements of PENN over the vanilla NN estimator are again very notable.



Figure 8: PUV for credit score prediction dataset.

Acknowledgements: The research of TM and RJS was supported by RJS's European Research Council Advanced Grant 101019498.



Figure 9: PUV for public procurement dataset.

References

- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R. and Wang, R. (2019) On exact computation with an infinitely wide neural net. Advances in Neural Information Processing Systems, 32.
- Bartlett, P. L., Harvey, N., Liaw, C. and Mehrabian, A. (2019) Nearly-tight VC-dimension and pseudodimension bounds for piecewise linear neural networks. *Journal of Machine Learning Research*, 20, 1–17.
- 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**, 2261–2285.
- Berrett, T. B. and Samworth, R. J. (2023) Optimal nonparametric testing of Missing Completely At Random, and its connections to compatibility. *The Annals of Statistics*, 51, 2170–2193.
- Bos, T. and Schmidt-Hieber, J. (2022) Convergence rates of deep ReLU networks for multiclass classification. *Electronic Journal of Statistics*, **16**, 2724–2773.
- Bos, T. and Schmidt-Hieber, J. (2023) A supervised deep learning method for nonparametric density estimation. arXiv preprint arXiv:2306.10471.
- Cai, T. T. and Zhang, L. (2019) High dimensional linear discriminant analysis: optimality, adaptive algorithm and missing data. *Journal of the Royal Statistical Society, Series* B: Statistical Methodology, 81, 675–705.

- Chandrasekher, K. A., Alaoui, A. E. and Montanari, A. (2020) Imputation for highdimensional linear regression. arXiv preprint arXiv:2001.09180.
- Chen, T. and Guestrin, C. (2016) XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794.
- Devlin, J., Chang, M.-W., Lee, K. and Toutanova, K. (2019) Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (long and short papers), 4171–4186.
- Du, S. S., Zhai, X., Poczos, B. and Singh, A. (2019) Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Repre*sentations.
- Efromovich, S. (2011) Nonparametric regression with predictors missing at random. *Journal of the American Statistical Association*, **106**, 306–319.
- Elsener, A. and van de Geer, S. (2019) Sparse spectral estimation with missing and corrupted measurements. *Stat*, **8**, e229.
- Fan, J. and Gu, Y. (2024) Factor augmented sparse throughput deep ReLU neural networks for high dimensional regression. *Journal of the American Statistical Association*, 119, 2680–2694.
- Fan, J., Ma, C. and Zhong, Y. (2020) A selective overview of deep learning. *Statistical Science*, 36, 264.
- Follain, B., Wang, T. and Samworth, R. J. (2022) High-dimensional changepoint estimation with heterogeneous missingness. Journal of the Royal Statistical Society Series B: Statistical Methodology, 84, 1023–1055.
- Györfi, L., Kohler, M., Krzyzak, A. and Walk, H. (2006) *A Distribution-free Theory of Nonparametric Regression*. Springer Science & Business Media.
- He, K., Zhang, X., Ren, S. and Sun, J. (2016) Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778.
- Hinton, G. E. and Salakhutdinov, R. R. (2006) Reducing the dimensionality of data with neural networks. *Science*, **313**, 504–507.
- Imaizumi, M. and Fukumizu, K. (2019) Deep neural networks learn non-smooth functions effectively. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 869–878, PMLR.
- Jacot, A., Gabriel, F. and Hongler, C. (2018) Neural tangent kernel: Convergence and generalization in neural networks. Advances in Neural Information Processing Systems, 31.

- 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**, 691–716.
- Josse, J., Chen, J. M., Prost, N., Varoquaux, G. and Scornet, E. (2024) On the consistency of supervised learning with missing values. *Statistical Papers*, **65**, 5447–5479.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P. and Hassabis, D. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature*, **596**, 583–589.
- Kingma, D. P. and Ba, J. (2015) Adam: A method for stochastic optimization. In *Inter*national Conference on Learning Representations.
- Kohler, M. and Langer, S. (2021) On the rate of convergence of fully connected deep neural network regression estimates. *The Annals of Statistics*, **49**, 2231–2249.
- Krizhevsky, A., Sutskever, I. and Hinton, G. E. (2012) Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25.
- Le Morvan, M., Josse, J., Scornet, E. and Varoquaux, G. (2021) What's a good imputation to predict with missing values? Advances in Neural Information Processing Systems, 34, 11530–11540.
- LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. (1998) Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, **86**, 2278–2324.
- LeCun, Y., Cortes, C. and Burges, C. J. (1998) The MNIST database of handwritten digits.
- Li, S. C.-X., Jiang, B. and Marlin, B. (2019) MisGAN: Learning from incomplete data with generative adversarial networks. In *International Conference on Learning Repre*sentations.
- Liu, Z., Sun, M., Zhou, T., Huang, G. and Darrell, T. (2019) Rethinking the value of network pruning. In *International Conference on Learning Representations*.
- Loh, P.-L. and Wainwright, M. J. (2012) High-dimensional regression with noisy and missing data: provable guarantees with nonconvexity. *The Annals of Statistics*, 40, 1637–1664.
- Lu, J., Shen, Z., Yang, H. and Zhang, S. (2021) Deep network approximation for smooth functions. SIAM Journal on Mathematical Analysis, 53, 5465–5506.

- Ma, T., Verchand, K. A., Berrett, T. B., Wang, T. and Samworth, R. J. (2024) Estimation beyond Missing (Completely) at Random. arXiv preprint arXiv:2410.10704.
- Mattei, P.-A. and Frellsen, J. (2019) MIWAE: Deep generative modelling and imputation of incomplete data sets. In *International Conference on Machine Learning*, 4413–4423, PMLR.
- Mei, S. and Montanari, A. (2022) The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, **75**, 667–766.
- Mei, S., Montanari, A. and Nguyen, P.-M. (2018) A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, **115**, E7665–E7671.
- Mikolov, T., Chen, K., Corrado, G. and Dean, J. (2013) Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781.
- Nazabal, A., Olmos, P. M., Ghahramani, Z. and Valera, I. (2020) Handling incomplete heterogeneous data using VAEs. *Pattern Recognition*, **107**, 107501.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. and Chintala, S. (2019) Pytorch: An imperative style, high-performance deep learning library. Advances in Neural Information Processing Systems, 32.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E. (2011) Scikit-learn: Machine learning in Python. the Journal of Machine Learning Research, 12, 2825–2830.
- Prechelt, L. (2002) Early stopping but when? In *Neural Networks: Tricks of the Trade* (Orr, G. B. and Müller, K.-R., eds.), 55–69, Springer.
- Royden, H. L. and Fitzpatrick, P. (2010) Real Analysis. Printice-Hall Inc, Boston.
- Samworth, R. J. and Shah, R. D. (2025+) Modern Statistical Methods and Theory. Cambridge University Press, Cambridge.
- Schmidt-Hieber, J. (2020) Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, **48**, 1875–1897.
- Sell, T., Berrett, T. B. and Cannings, T. I. (2024) Nonparametric classification with missing data. *The Annals of Statistics*, **52**, 1178–1200.
- Shen, Z., Yang, H. and Zhang, S. (2019) Nonlinear approximation via compositions. Neural Networks, 119, 74–84.

- Śmieja, M., Struski, Ł., Tabor, J., Zieliński, B. and Spurek, P. (2018) Processing of missing data by neural networks. Advances in Neural Information Processing Systems, 31.
- Stekhoven, D. J. and Bühlmann, P. (2012) MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, **28**, 112–118.
- Suh, N. and Cheng, G. (2024) A survey on statistical theory of deep learning: Approximation, training dynamics, and generative models. Annual Review of Statistics and Its Application, 12, 177–207.
- Twala, B. E., Jones, M. and Hand, D. J. (2008) Good methods for coping with missing data in decision trees. *Pattern Recognition Letters*, **29**, 950–956.
- van Buuren, S. and Groothuis-Oudshoorn, K. (2011) mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, **45**, 1–67.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. and Polosukhin, I. (2017) Attention is all you need. Advances in Neural Information Processing Systems, 30.
- Zhang, H., Fang, L., Wu, Q. and Yu, P. S. (2025) DiffPuter: An EM-driven diffusion model for missing data imputation. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, Z., Shi, L. and Zhou, D.-X. (2024) Classification with deep neural networks and logistic loss. Journal of Machine Learning Research, 25, 1–117.
- Zhong, Q., Mueller, J. and Wang, J.-L. (2022) Deep learning for the partially linear Cox model. The Annals of Statistics, 50, 1348–1375.
- Zhu, Z., Wang, T. and Samworth, R. J. (2022) High-dimensional principal component analysis with heterogeneous missingness. *Journal of the Royal Statistical Society Series* B: Statistical Methodology, 84, 2000–2031.

A Some properties of neural networks

We summarise some basic operations associated with neural networks:

- (P1) Neural network enlarging: Let $L \in \mathbb{N}$, let $p_1, p_2 \in \mathbb{N}^{L+2}$ be such that $p_1 \leq p_2$ coordinate-wise and let $s \in \mathbb{N}$. Then, from the definition, $\mathcal{F}(L, p_1, s) \subseteq \mathcal{F}(L, p_2, s)$.
- (P2) Neural network composition: $\boldsymbol{p}_1 = (p_{1,0}, \dots, p_{1,L_1+1})^\top \in \mathbb{N}^{L_1+2}$ and $\boldsymbol{p}_2 = (p_{2,0}, \dots, p_{2,L_2+1})^\top \in \mathbb{N}^{L_2+2}$ satisfy $p_{1,L_1+1} = p_{2,0}$; let $\boldsymbol{f}_1 \in \mathcal{F}(L_1, \boldsymbol{p}_1)$ and $\boldsymbol{f}_2 \in \mathcal{F}(L_2, \boldsymbol{p}_2)$ (if either \boldsymbol{f}_1 or \boldsymbol{f}_2 is an affine function, then we set the corresponding L_1 or L_2 to be zero). Then

$$\boldsymbol{f}_2 \circ \boldsymbol{f}_1 \in \mathcal{F}(L_1 + L_2, \boldsymbol{p}_3), \text{ where }$$

$$\boldsymbol{p}_3 \coloneqq (p_{1,0}, \ldots, p_{1,L_1}, p_{2,1}, \ldots, p_{2,L_2+1}) \in \mathbb{N}^{L_1+L_2+2}.$$

To see this, suppose that $\boldsymbol{f}_1(\cdot) = \boldsymbol{A}_{L_1+1}^{(1)} \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_{L_1}^{(1)} \circ \boldsymbol{\sigma} \circ \cdots \circ \boldsymbol{A}_{2}^{(1)} \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_{1}^{(1)}(\cdot)$ and $\boldsymbol{f}_2(\cdot) = \boldsymbol{A}_{L_2+1}^{(2)} \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_{L_2}^{(2)} \circ \boldsymbol{\sigma} \circ \cdots \circ \boldsymbol{A}_{2}^{(2)} \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_{1}^{(2)}(\cdot)$. Then $\boldsymbol{f}_2 \circ \boldsymbol{f}_1(\cdot) = \boldsymbol{A}_{L_2+1}^{(2)} \circ \boldsymbol{\sigma} \circ \cdots \circ \boldsymbol{A}_{2}^{(2)} \circ \boldsymbol{\sigma} \circ (\boldsymbol{A}_{1}^{(2)} \circ \boldsymbol{A}_{L_1+1}^{(1)}) \circ \boldsymbol{\sigma} \circ \cdots \circ \boldsymbol{A}_{2}^{(1)} \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_{1}^{(1)}(\cdot)$

where we note that $A_1^{(2)} \circ A_{L_1+1}^{(1)} : \mathbb{R}^{p_{1,L_1}} \to \mathbb{R}^{p_{2,1}}$ is an affine function.

- (P3) Neural network padding: Let $L_1, L_2 \in \mathbb{N}$ be such that $L_1 < L_2$, let $\boldsymbol{p}_1 = (p_{1,0}, \ldots, p_{1,L_1+1})^\top \in \mathbb{N}^{L_1+2}, \ \boldsymbol{p}_2 = (\boldsymbol{p}_1, 2p_{1,L_1+1}, \ldots, 2p_{1,L_1+1})^\top \in \mathbb{N}^{L_2+2}$ and $s \in \mathbb{N}$. Then $\mathcal{F}(L_1, \boldsymbol{p}_1, s) \subseteq \mathcal{F}(L_2, \boldsymbol{p}_2, 2s + 2p_{1,L_1+1}(L_2 L_1))$. To see this, let $\boldsymbol{\phi} \in \mathcal{F}(1, (p_{1,L_1+1}, 2p_{1,L_1+1}, p_{1,L_1+1}))$ be defined by $\boldsymbol{\phi}(\boldsymbol{y}) \coloneqq \boldsymbol{\sigma}(\boldsymbol{y}) \boldsymbol{\sigma}(-\boldsymbol{y})$ for $\boldsymbol{y} \in \mathbb{R}^{p_{1,L_1+1}},$ so that $\boldsymbol{\phi}(\boldsymbol{y}) = \boldsymbol{y}$. Let $\boldsymbol{f}_1 \in \mathcal{F}(L_1, \boldsymbol{p}_1, s)$, and define $\boldsymbol{f}_2 \coloneqq \boldsymbol{\phi} \circ \cdots \circ \boldsymbol{\phi} \circ \boldsymbol{f}_1$, where $\boldsymbol{\phi}$ is applied $(L_2 - L_1)$ -times. Then $\boldsymbol{f}_1(\boldsymbol{x}) = \boldsymbol{f}_2(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathbb{R}^{p_{1,0}}$, and $\boldsymbol{f}_2 \in \mathcal{F}(L_2, \boldsymbol{p}_2, 2s + 2p_{1,L_1+1}(L_2 - L_1))$ by (P2) and counting parameters.
- (P4) Neural network parallelisation: Let $N, L \in \mathbb{N}$ and for $i \in [N]$, let $f_i \in \mathcal{F}(L, p_i, s_i)$ where $p_i = (d, p_{i,1}, \dots, p_{i,L+1})^\top \in \mathbb{N}^{L+2}$ and $s_i \in \mathbb{N}$. Then, writing $f(\cdot) = (f_1(\cdot)^\top, \dots, f_N(\cdot)^\top)^\top$, we have from the definition that

$$\boldsymbol{f} \in \mathcal{F}\left(L, \left(d, \sum_{i=1}^{N} p_{i,1}, \dots, \sum_{i=1}^{N} p_{i,L+1}\right), \sum_{i=1}^{N} s_{i}\right).$$

B Covering number bounds for sparse neural networks

Our oracle inequality for the excess risk of sparse neural network estimators in Theorem 1 relies on bounds on the covering number of the relevant class that we develop in this subsection. In fact, our proof proceeds via a bound on the *pseudo-dimension* of this class, which is related to its Vapnik–Chervonenkis (VC) dimension, so we begin by defining the notions we require.

Definition 5. Let \mathcal{H} be a set of functions from $\mathcal{X} \subseteq \mathbb{R}^d$ to $\{0,1\}$. For $m \in \mathbb{N}$, the shattering coefficient is defined as

shat
$$(\mathcal{H}, m) \coloneqq \max_{\boldsymbol{x}_1, \dots, \boldsymbol{x}_m \in \mathcal{X}} |\{(h(\boldsymbol{x}_1), \dots, h(\boldsymbol{x}_m)) : h \in \mathcal{H}\}|.$$

The VC-dimension of \mathcal{H} is defined as

 $\operatorname{VCdim}(\mathcal{H}) \coloneqq \sup\{m \in \mathbb{N} : \operatorname{shat}(\mathcal{H}, m) = 2^m\}.$

We define the sign function sgn : $\mathbb{R} \to \{0,1\}$ as⁵ sgn $(x) \coloneqq \mathbb{1}_{\{x>0\}}$. If \mathcal{F} is a class of real-valued functions, then we define sgn $\circ \mathcal{F} \coloneqq \{\text{sgn } \circ f : f \in \mathcal{F}\}.$

⁵Although this is not the standard definition of the sign function, it is used elsewhere in the neural network literature (e.g. Bartlett et al., 2019), and it is convenient for our purposes here.

Definition 6. Let \mathcal{F} be a set of functions from $\mathcal{X} \subseteq \mathbb{R}^d$ to \mathbb{R} . The pseudo-dimension of \mathcal{F} is defined as

$$\operatorname{Pdim}(\mathcal{F}) \coloneqq \operatorname{VCdim}(\{(\boldsymbol{x}, y) \mapsto \operatorname{sgn}(f(\boldsymbol{x}) - y) : f \in \mathcal{F}\}).$$

For $\epsilon > 0$, we say that a collection of functions $g_1, \ldots, g_N : \mathcal{X} \to \mathbb{R}$ is an ϵ -cover of \mathcal{F} with respect to $\|\cdot\|_{L_q(\mu)}$ if for every $f \in \mathcal{F}$, there exists $j = j(f) \in [N]$ such that $\|f - g_j\|_{L_q(\mu)} \leq \epsilon$. The ϵ -covering number of \mathcal{F} with respect to $\|\cdot\|_{L_q(\mu)}$, written as $\mathcal{N}(\epsilon, \mathcal{F}, \|\cdot\|_{L_q(\mu)})$, is the cardinality of the smallest ϵ -cover of \mathcal{F} with respect to $\|\cdot\|_{L_q(\mu)}$ (if no finite ϵ -cover exists, then the ϵ -covering number is ∞).

The following proposition provides upper bounds on the VC-dimension and covering number of the class $\mathcal{F}(L, \boldsymbol{p}, s)$. The proof is based on Bartlett et al. (2019, Theorem 7), which gives an upper bound on the VC-dimension of the class $\mathcal{F}(L, \boldsymbol{p})$. We also remark that upper bounds on the covering number of the class $\{f \in \mathcal{F}(L, \boldsymbol{p}, s) : \|\Theta(f)\|_{\infty} \leq 1\}$ have been obtained by, e.g. Schmidt-Hieber (2020, Lemma 5); however, these do not imply an upper bound on its VC-dimension, and they do not generalise to the class $\mathcal{F}(L, \boldsymbol{p}, s)$.

Proposition 5. Let $L \in \mathbb{N}$, $p = (p_0, ..., p_{L+1})^{\top} \in \mathbb{N}^{L+2}$ with $p_{L+1} = 1$, let $V \coloneqq \sum_{\ell=1}^{L+1} p_\ell(p_{\ell-1}+1)$ and let $s \in [V]$.

(a) We have

$$\operatorname{VCdim}(\operatorname{sgn} \circ \mathcal{F}(L, \boldsymbol{p}, s)) \leq \operatorname{Pdim}(\mathcal{F}(L, \boldsymbol{p}, s))$$
$$\leq 6s(L+1)\log_2(3s) + 2s\log_2(2p_0)$$
$$\lesssim sL\log(es) + s\log(ep_0).$$

(b) For $q \in [1, \infty)$, B > 0, $\epsilon \in (0, B/2)$ and any probability measure μ on \mathbb{R}^d , we have

$$\log \mathcal{N}(\epsilon, T_B \circ \mathcal{F}(L, \boldsymbol{p}, s), \|\cdot\|_{L_q(\mu)}) \leq 2q \cdot \operatorname{Pdim}(\mathcal{F}(L, \boldsymbol{p}, s)) \log(10B/\epsilon) \\ \lesssim q \{ sL \log(es) + s \log(ep_0) \} \log(B/\epsilon).$$

By Bartlett et al. (2019, Theorem 3), there exists a universal constant $C_3 > 0$ such that if $s \ge C_3 L \ge C_3^2$, then there exists $\mathbf{p}' = (p'_0, \ldots, p'_{L+1})^\top \in \mathbb{N}^{L+2}$ with $V' \coloneqq \sum_{\ell=1}^{L+1} p'_{\ell}(p'_{\ell-1} + 1) \le s$, such that VCdim $(\text{sgn} \circ \mathcal{F}(L, \mathbf{p}')) \ge c_4 \cdot sL \log(s/L)$ for some universal constant $c_4 > 0$. By (P1), $\mathcal{F}(L, \mathbf{p}') \subseteq \mathcal{F}(L, \mathbf{p}, s)$ for all $\mathbf{p} \in \mathbb{N}^{L+2}$ such that $\mathbf{p} \ge \mathbf{p}'$ coordinate-wise, so the same lower bound then applies to VCdim $(\text{sgn} \circ \mathcal{F}(L, \mathbf{p}, s))$. Thus, under the above condition on s and L, our upper bound in Proposition 5(a) is tight up to a logarithmic factor in n when \mathbf{p} is sufficiently large.

Proof. (a) We first observe that

$$\operatorname{Pdim}(\mathcal{F}(L,\boldsymbol{p},s)) = \operatorname{VCdim}(\{(\boldsymbol{x},y) \mapsto \operatorname{sgn}(f(\boldsymbol{x}) - y) : f \in \mathcal{F}(L,\boldsymbol{p},s)\})$$
$$\geq \operatorname{VCdim}(\{\boldsymbol{x} \mapsto \operatorname{sgn}(f(\boldsymbol{x})) : f \in \mathcal{F}(L,\boldsymbol{p},s)\})$$
$$= \operatorname{VCdim}(\operatorname{sgn} \circ \mathcal{F}(L,\boldsymbol{p},s)).$$

Writing $\widetilde{\boldsymbol{p}} := (p_0, p_1 \wedge s, \dots, p_L \wedge s, p_{L+1})^{\top}$, we have by Schmidt-Hieber (2020, Equation (19)) that $\mathcal{F}(L, \boldsymbol{p}, s) = \mathcal{F}(L, \widetilde{\boldsymbol{p}}, s)$. Define a class of functions $\mathcal{H} := \{(\boldsymbol{x}, y) \mapsto \operatorname{sgn}(f(\boldsymbol{x}) - y) : f \in \mathcal{F}(L, \widetilde{\boldsymbol{p}}, s), \gamma \in \mathbb{R}\}$ from $\mathbb{R}^{p_0} \times \mathbb{R}$ to \mathbb{R} . Then $\operatorname{Pdim}(\mathcal{F}(L, \boldsymbol{p}, s)) = \operatorname{VCdim}(\mathcal{H})$, so it suffices to upper bound $m := \operatorname{VCdim}(\mathcal{H})$. If m < s, then part (a) of the proposition follows since the right-hand side is at least s. Therefore, we may assume without loss of generality that $m \geq s$. Suppose that $(\boldsymbol{x}_1, y_1), \dots, (\boldsymbol{x}_m, y_m) \in \mathbb{R}^{p_0} \times \mathbb{R}$ are shattered by \mathcal{H} . Let $\widetilde{V} := \sum_{\ell=1}^{L+1} \widetilde{p}_{\ell}(\widetilde{p}_{\ell-1}+1)$ denote the total number of parameters of a neural network in $\mathcal{F}(L, \widetilde{\boldsymbol{p}}, s)$. For $\boldsymbol{x} \in \mathbb{R}^{p_0}$ and $\boldsymbol{\theta} \in \{\boldsymbol{\theta}' \in \mathbb{R}^{\widetilde{V}} : \|\boldsymbol{\theta}'\|_0 \leq s\} =: \mathbb{B}_0^{\widetilde{V}}(s)$, let $g(\boldsymbol{x}, \boldsymbol{\theta}) := f(\boldsymbol{x})$ when $f \in \mathcal{F}(L, \widetilde{\boldsymbol{p}}, s)$ satisfies $\Theta(f) = \boldsymbol{\theta}$. In other words, $g(\cdot, \boldsymbol{\theta})$ is the neural network in $\mathcal{F}(L, \widetilde{\boldsymbol{p}}, s)$ with parameter vector $\boldsymbol{\theta}$, so $\mathcal{F}(L, \widetilde{\boldsymbol{p}}, s) = \{g(\cdot, \boldsymbol{\theta}) : \boldsymbol{\theta} \in \mathbb{B}_0^{\widetilde{V}}(s)\}$. We partition $\mathbb{B}_0^{\widetilde{V}}(s)$ into sets $B_1, \dots, B_{\binom{\widetilde{V}}{s}}$, where the elements in each set B_i are all supported on the same set of cardinality s. Then, by definition of m,

$$2^{m} = \left| \left\{ \left(\operatorname{sgn} \left(g(\boldsymbol{x}_{1}, \boldsymbol{\theta}) - y_{1} \right), \dots, \operatorname{sgn} \left(g(\boldsymbol{x}_{m}, \boldsymbol{\theta}) - y_{m} \right) \right)^{\top} : \boldsymbol{\theta} \in \mathbb{B}_{0}^{\widetilde{V}}(s) \right\} \right|$$

$$\leq \sum_{i=1}^{\binom{\widetilde{V}}{s}} \left| \left\{ \left(\operatorname{sgn} \left(g(\boldsymbol{x}_{1}, \boldsymbol{\theta}) - y_{1} \right), \dots, \operatorname{sgn} \left(g(\boldsymbol{x}_{m}, \boldsymbol{\theta}) - y_{m} \right) \right)^{\top} : \boldsymbol{\theta} \in B_{i} \right\} \right| \rightleftharpoons \sum_{i=1}^{\binom{\widetilde{V}}{s}} K_{i}. \quad (4)$$

We will prove upper bounds for $K_1, \ldots, K_{\binom{\widetilde{V}}{s}}$, which then imply an upper bound on m. To this end, without loss of generality, we upper bound K_1 . For $\ell \in [L+1]$, $\boldsymbol{x} \in \mathbb{R}^{p_0}$ and $\boldsymbol{\theta} \in B_1$, define $\boldsymbol{g}^{(\ell)}(\boldsymbol{x}, \boldsymbol{\theta}) \coloneqq \boldsymbol{A}_{\ell} \circ \boldsymbol{\sigma} \circ \cdots \circ \boldsymbol{\sigma} \circ \boldsymbol{A}_1(\boldsymbol{x})$, where $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_{L+1}$ are defined analogously to (1) (but with each p_{ℓ} there replaced with \widetilde{p}_{ℓ}) with weight matrices and bias vectors given by relevant components of $\boldsymbol{\theta}$. For $\ell \in [L+1]$ and $u \in [\widetilde{p}_{\ell}]$, let $g_u^{(\ell)}$ be the *u*th coordinate function of $\boldsymbol{g}^{(\ell)}$, let $\mathcal{U}_{\ell} \coloneqq \{u \in [\widetilde{p}_{\ell}] : g_u^{(\ell)}(\cdot, \boldsymbol{\theta}) \neq 0$ for some $\boldsymbol{\theta} \in B_1\}$ be the coordinates of active neurons in the ℓ th layer, and let $k_{\ell} \coloneqq |\mathcal{U}_{\ell}|$. If $k_{\ell} = 0$ for some $\ell \in [L+1]$, then $\boldsymbol{g}(\cdot, \boldsymbol{\theta}) = 0$ for all $\boldsymbol{\theta} \in B_1$, so $K_1 = 1$. Now assume that $k_{\ell} \ge 1$ for all $\ell \in [L+1]$. We will construct a finite sequence of partitions $\mathcal{P}_1, \ldots, \mathcal{P}_{L+1}$ of B_1 , each refining the previous one, satisfying the following properties:

(i) $N_{\ell} := |\mathcal{P}_{\ell}|$ satisfies $N_1 = 1$ and for $\ell \in [L]$, we have

$$\frac{N_{\ell+1}}{N_{\ell}} \le 2\left(\frac{2em\ell k_{\ell}}{s}\right)^s \eqqcolon \phi_{\ell}.$$

(ii) $\boldsymbol{\theta} \mapsto \boldsymbol{g}^{(\ell)}(\boldsymbol{x}_j, \boldsymbol{\theta})$ is a polynomial of degree at most ℓ on P_ℓ , for each $j \in [m], \ell \in [L+1]$ and $P_\ell \in \mathcal{P}_\ell$.

We start by defining $\mathcal{P}_1 \coloneqq \{B_1\}$. By the definition of A_1 in (1), we have that $\boldsymbol{\theta} \mapsto \boldsymbol{g}^{(1)}(\boldsymbol{x}_j, \boldsymbol{\theta}) = A_1(\boldsymbol{x}_j)$ is linear on B_1 . Now suppose that we have constructed $\mathcal{P}_1, \ldots, \mathcal{P}_\ell$ for some $\ell \in [L]$ satisfying both Properties (i) and (ii). By the induction hypothesis, $\mathcal{P}_\ell = \{P_{\ell,1}, \ldots, P_{\ell,N_\ell}\}$ is such that, for each $r \in [N_\ell], j \in [m]$ and $u \in \mathcal{U}_\ell$, the map $\boldsymbol{\theta} \mapsto g_u^{(\ell)}(\boldsymbol{x}_j, \boldsymbol{\theta})$ is a polynomial of degree at most ℓ on $P_{\ell,r}$, depending on at most s coordinates of $\boldsymbol{\theta}$. Then, for a fixed $r \in [N_\ell]$, by Bartlett et al. (2019, Lemma 17), since $m \geq s$, we have

$$\left|\left\{\left(\operatorname{sgn}\left(g_{u}^{(\ell)}(\boldsymbol{x}_{j},\boldsymbol{\theta})\right): u \in [\widetilde{p}_{\ell}], \ j \in [m]\right): \boldsymbol{\theta} \in P_{\ell,r}\right\}\right|$$

$$= \left| \left\{ \left(\operatorname{sgn} \left(g_u^{(\ell)}(\boldsymbol{x}_j, \boldsymbol{\theta}) \right) : u \in \mathcal{U}_{\ell}, \, j \in [m] \right) : \boldsymbol{\theta} \in P_{\ell, r} \right\} \right| \le \phi_{\ell}$$

Therefore, for each $r \in [N_{\ell}]$, we can partition $P_{\ell,r}$ into regions $R_{r,1}, \ldots, R_{r,\phi_{\ell}}$ such that for each $t \in [\phi_{\ell}]$, the map $\boldsymbol{\theta} \mapsto \left(\operatorname{sgn} \left(g_{u}^{(\ell)}(\boldsymbol{x}_{j}, \boldsymbol{\theta}) \right) : u \in [\widetilde{p}_{\ell}], j \in [m] \right)$ is constant on $R_{r,t}$. We then define $\mathcal{P}_{\ell+1} \coloneqq \left\{ R_{r,t} : r \in [N_{\ell}], t \in [\phi_{\ell}] \right\}$. By construction, $\mathcal{P}_{\ell+1}$ satisfies Property (i). Moreover, writing $\boldsymbol{\xi}_{\boldsymbol{x}_{j},\boldsymbol{\theta}} \coloneqq \left(\operatorname{sgn} \left(g_{u}^{(\ell)}(\boldsymbol{x}_{j}, \boldsymbol{\theta}) \right) : u \in [\widetilde{p}_{\ell}] \right)^{\top} \in \{0, 1\}^{\widetilde{p}_{\ell}}$, we have

$$\boldsymbol{g}^{(\ell+1)}(\boldsymbol{x}_j, \boldsymbol{\theta}) = \boldsymbol{A}_{\ell+1} \circ \boldsymbol{\sigma} \circ \boldsymbol{g}^{(\ell)}(\boldsymbol{x}_j, \boldsymbol{\theta}) = \boldsymbol{A}_{\ell+1} \big(\operatorname{diag}(\boldsymbol{\xi}_{\boldsymbol{x}_j, \boldsymbol{\theta}}) \boldsymbol{g}^{(\ell)}(\boldsymbol{x}_j, \boldsymbol{\theta}) \big)$$

Since for each $j \in [m]$, $r \in [N_{\ell}]$ and $t \in [\phi_{\ell}]$, the sign vector $\boldsymbol{\xi}_{\boldsymbol{x}_{j},\boldsymbol{\theta}}$ is constant for all $\boldsymbol{\theta} \in R_{r,t}$, the map $\boldsymbol{\theta} \mapsto \boldsymbol{g}^{(\ell+1)}(\boldsymbol{x}_{j},\boldsymbol{\theta})$ is a polynomial of degree at most $\ell + 1$ on $R_{r,t}$, which verifies Property (ii) for $\ell + 1$ and hence completes the induction.

Next, by Property (ii), for each $j \in [m]$ and $P_L \in \mathcal{P}_L$, we have that $\boldsymbol{\theta} \mapsto g(\boldsymbol{x}_j, \boldsymbol{\theta}) - y_j = \boldsymbol{g}^{(L+1)}(\boldsymbol{x}_j, \boldsymbol{\theta}) - y_j$ is a polynomial of degree at most L on P_L , depending on at most s coordinates of $\boldsymbol{\theta}$. Thus, by Bartlett et al. (2019, Lemma 17) again,

$$\left|\left\{\left(\operatorname{sgn}(g(\boldsymbol{x}_{1},\boldsymbol{\theta})-y_{1}),\ldots,\operatorname{sgn}(g(\boldsymbol{x}_{m},\boldsymbol{\theta})-y_{m})\right)^{\top}:\boldsymbol{\theta}\in P_{L}\right\}\right|\leq 2\left(\frac{2emL}{s}\right)^{s}.$$
 (5)

By Property (i), we also have

$$|\mathcal{P}_L| \le 2^L \prod_{\ell=1}^L \left(\frac{2em\ell k_\ell}{s}\right)^s.$$
(6)

Combining (5) and (6) yields that

$$K_{1} \leq 2^{L+1} \left(\frac{2emL}{s}\right)^{s} \prod_{\ell=1}^{L} \left(\frac{2em\ell k_{\ell}}{s}\right)^{s} = 2^{L+1} \left(\frac{2em}{s}\right)^{s(L+1)} \left(\prod_{\ell=1}^{L} \ell k_{\ell}\right)^{s}$$
$$\leq 2^{L+1} \left(\frac{2em}{s}\right)^{s(L+1)} \left(\frac{1}{L} \sum_{\ell=1}^{L} \ell k_{\ell}\right)^{sL} \leq 2^{L+1} (2em)^{s(L+1)},$$

where the second inequality is an application of the AM–GM inequality, and the final bound uses the fact that $\sum_{\ell=1}^{L} k_{\ell} \leq s$. Therefore, by (4), we deduce that

$$2^{m} \leq {\widetilde{V} \choose s} \cdot 2^{L+1} (2em)^{s(L+1)} \leq \left(4em\widetilde{V}^{1/(L+1)}\right)^{s(L+1)}.$$

By Lemma 12, since $4e\widetilde{V}^{1/(L+1)} \geq 4$, we have

$$m \le 2s(L+1)\log_2(4es(L+1)\widetilde{V}^{1/(L+1)})$$

$$\le 2s(L+1)\log_2(8es^3(2p_0)^{1/(L+1)}) \le 6s(L+1)\log_2(3s) + 2s\log_2(2p_0),$$

where the second inequality follows since $L + 1 \leq s$ by our assumption that $k_{\ell} \geq 1$ for all $\ell \in [L + 1]$ and

$$\widetilde{V} \le (L-1)s(s+1) + (p_0+1)s + s + 1 \le 2(L+1)s^2 + p_0s \le 2p_0(2s)^{L+1}.$$

(b) If $(\boldsymbol{x}_1, y_1), \ldots, (\boldsymbol{x}_m, y_m) \in \mathbb{R}^{p_0} \times \mathbb{R}$ are shattered by $\{(\boldsymbol{x}, y) \mapsto \operatorname{sgn}(T_B f(\boldsymbol{x}) - y) : f \in \mathcal{F}(L, \boldsymbol{p}, s)\}$, then we must have $y_j \in [-B, B)$ for all $j \in [m]$. Therefore, $(\boldsymbol{x}_1, y_1), \ldots, (\boldsymbol{x}_m, y_m)$ are also shattered by $\{(\boldsymbol{x}, y) \mapsto \operatorname{sgn}(f(\boldsymbol{x}) - y) : f \in \mathcal{F}(L, \boldsymbol{p}, s)\}$, so $\operatorname{Pdim}(T_B \circ \mathcal{F}(L, \boldsymbol{p}, s)) \leq \operatorname{Pdim}(\mathcal{F}(L, \boldsymbol{p}, s))$. Then, by Györfi et al. (2006, Theorem 9.4), since $\operatorname{Pdim}(T_B \circ \mathcal{F}(L, \boldsymbol{p}, s)) \geq 2$, we have

$$\mathcal{N}(\epsilon, T_B \circ \mathcal{F}(L, \boldsymbol{p}, s), \| \cdot \|_{L_q(\mu)}) \leq 3 \left\{ \frac{2e(2B)^q}{\epsilon^q} \log \left(\frac{3e(2B)^q}{\epsilon^q} \right) \right\}^{\operatorname{Pdim}(T_B \circ \mathcal{F}(L, \boldsymbol{p}, s))} \\ \leq \left(\frac{10B}{\epsilon} \right)^{2q \cdot \operatorname{Pdim}(\mathcal{F}(L, \boldsymbol{p}, s))}.$$

Taking logarithms and applying part (a) yields the desired result.

C Proof of Theorem 1

Proof of Theorem 1. For i = 0, 1, ..., n, let $\widetilde{Z}_i := (Z_i, \Omega_i)$ and let $f_{B_n}^{\star}(\widetilde{Z}_i) := \mathbb{E}(T_{B_n}Y_i | \widetilde{Z}_i)$. Then

$$\begin{split} \mathbb{E}\{R(T_{B_{n}}\widetilde{f}) - R(f^{\star})\} &= \mathbb{E}\{\left(T_{B_{n}}\widetilde{f}(\widetilde{Z}_{0}) - Y_{0}\right)^{2} - \left(T_{B_{n}}\widetilde{f}(\widetilde{Z}_{0}) - T_{B_{n}}Y_{0}\right)^{2}\} \\ &+ \mathbb{E}\{\left(f_{B_{n}}^{\star}(\widetilde{Z}_{0}) - T_{B_{n}}Y_{0}\right)^{2} - \left(f^{\star}(\widetilde{Z}_{0}) - Y_{0}\right)^{2}\} \\ &+ \mathbb{E}\left[\left(T_{B_{n}}\widetilde{f}(\widetilde{Z}_{0}) - T_{B_{n}}Y_{0}\right)^{2} - \left(f_{B_{n}}^{\star}(\widetilde{Z}_{0}) - T_{B_{n}}Y_{0}\right)^{2}\right] \\ &- \frac{2}{n}\sum_{i=1}^{n}\left\{\left(T_{B_{n}}\widetilde{f}(\widetilde{Z}_{i}) - T_{B_{n}}Y_{i}\right)^{2} - \left(f_{B_{n}}^{\star}(\widetilde{Z}_{i}) - T_{B_{n}}Y_{i}\right)^{2}\right\}\right] \\ &+ \mathbb{E}\left[\frac{2}{n}\sum_{i=1}^{n}\left\{\left(T_{B_{n}}\widetilde{f}(\widetilde{Z}_{i}) - T_{B_{n}}Y_{i}\right)^{2} - \left(f^{\star}(\widetilde{Z}_{i}) - Y_{i}\right)^{2}\right\}\right] \\ &+ \mathbb{E}\left[\frac{2}{n}\sum_{i=1}^{n}\left\{\left(f^{\star}(\widetilde{Z}_{i}) - Y_{i}\right)^{2} - \left(f^{\star}_{B_{n}}(\widetilde{Z}_{i}) - T_{B_{n}}Y_{i}\right)^{2}\right\}\right] \\ &=: E_{1} + E_{2} + E_{3} + E_{4} + E_{5}. \end{split}$$

Bounding E_1 : By Lemma 13,

$$E_1 \le \mathbb{E}\left\{Y_0^2 - (T_{B_n}Y_0)^2\right\} + 2B_n \mathbb{E}\left\{|Y_0 - T_{B_n}Y_0|\right\} \le \frac{5\xi^2}{n}.$$
(7)

Bounding $E_2 + E_5$: Note that $E_5 = -2E_2$, so by Lemma 13,

$$E_{2} + E_{5} = \mathbb{E}\left\{ \left(f^{*}(\widetilde{Z}_{0}) - Y_{0} \right)^{2} - \left(f^{*}_{B_{n}}(\widetilde{Z}_{0}) - T_{B_{n}}Y_{0} \right)^{2} \right\} \\ = \mathbb{E}\left\{ \left(f^{*}(\widetilde{Z}_{0}) \right)^{2} - \left(f^{*}_{B_{n}}(\widetilde{Z}_{0}) \right)^{2} \right\} + \mathbb{E}\left\{ Y_{0}^{2} - (T_{B_{n}}Y_{0})^{2} \right\} \\ + 2\mathbb{E}\left\{ f^{*}(\widetilde{Z}_{0}) \cdot Y_{0} - f^{*}_{B_{n}}(\widetilde{Z}_{0}) \cdot T_{B_{n}}Y_{0} \right\} \le \frac{13\xi^{2}}{n}.$$
(8)

Bounding E_3 : For $(\boldsymbol{u}_1, \boldsymbol{v}_1), \ldots, (\boldsymbol{u}_n, \boldsymbol{v}_n) \in \mathbb{R}^d \times \mathbb{R}^d$, let \mathbb{P}_n be their empirical measure and let

$$\mathcal{N}_{1,n} \coloneqq \sup_{(\boldsymbol{u}_i, \boldsymbol{v}_i)_{i=1}^n \in (\mathbb{R}^d \times \mathbb{R}^d)^n} \mathcal{N}\big((80B_n n)^{-1}, T_{B_n} \circ \mathcal{F}, \|\cdot\|_{L_1(\mathbb{P}_n)} \big).$$

For $\delta \coloneqq \frac{24 \cdot 214 B_n^4 \log(14\mathcal{N}_{1,n})}{n} \ge 1/n$, we have

$$\begin{split} E_{3} &= \mathbb{E} \left[\mathbb{E} \left\{ \left(T_{B_{n}} \widetilde{f}(\widetilde{Z}_{0}) - T_{B_{n}} Y_{0} \right)^{2} - \left(f_{B_{n}}^{*}(\widetilde{Z}_{0}) - T_{B_{n}} Y_{0} \right)^{2} \mid \mathcal{D} \right\} \\ &\quad - \frac{2}{n} \sum_{i=1}^{n} \left\{ \left(T_{B_{n}} \widetilde{f}(\widetilde{Z}_{i}) - T_{B_{n}} Y_{i} \right)^{2} - \left(f_{B_{n}}^{*}(\widetilde{Z}_{i}) - T_{B_{n}} Y_{i} \right)^{2} \right\} \right] \\ &\leq \mathbb{E} \left(\sup_{g \in T_{B_{n}} \circ \mathcal{F}(L, p, s)} \left[\mathbb{E} \left\{ \left(g(\widetilde{Z}_{0}) - T_{B_{n}} Y_{0} \right)^{2} - \left(f_{B_{n}}^{*}(\widetilde{Z}_{0}) - T_{B_{n}} Y_{i} \right)^{2} \right\} \right] \right) \\ &\leq \delta + \int_{\delta}^{\infty} \mathbb{P} \left(\sup_{g \in T_{B_{n}} \circ \mathcal{F}(L, p, s)} \left[\mathbb{E} \left\{ \left(g(\widetilde{Z}_{0}) - T_{B_{n}} Y_{i} \right)^{2} - \left(f_{B_{n}}^{*}(\widetilde{Z}_{0}) - T_{B_{n}} Y_{i} \right)^{2} \right\} \right] \right) \right] dt \\ &\leq \delta + \int_{\delta}^{\infty} \mathbb{P} \left(\sup_{g \in T_{B_{n}} \circ \mathcal{F}(L, p, s)} \left[\mathbb{E} \left\{ \left(g(\widetilde{Z}_{0}) - T_{B_{n}} Y_{0} \right)^{2} - \left(f_{B_{n}}^{*}(\widetilde{Z}_{0}) - T_{B_{n}} Y_{0} \right)^{2} \right\} \right] \right] dt \\ &\leq \delta + 14 \mathcal{N}_{1,n} \int_{\delta}^{\infty} \exp \left(- \frac{nt}{24 \cdot 214 B_{n}^{4}} \right) dt \\ &= \delta + 14 \mathcal{N}_{1,n} \cdot \frac{24 \cdot 214 B_{n}^{4}}{n} \exp \left(- \frac{n\delta}{24 \cdot 214 B_{n}^{4}} \right) \\ &\leq \frac{24 \cdot 214 B_{n}^{4} \log(14 \mathcal{N}_{1,n})}{n} + \frac{24 \cdot 214 B_{n}^{4}}{n} \right] \\ &\leq \frac{24 \cdot 214 B_{n}^{4} \log(14 \mathcal{N}_{1,n})}{n} + \frac{24 \cdot 214 B_{n}^{4}}{n} \right] \end{cases}$$

were (i) uses Györfi et al. (2006, Theorem 11.4) with $\epsilon = 1/2$ and $\alpha = \beta = t/2$ therein, and (ii) uses Proposition 5(b).

Bounding E_4 : The functions in \mathcal{F} are parametrised by a subset A of a Euclidean space, meaning that we can write $\mathcal{F} = \{f_{\theta} : \theta \in A\}$. Since A has a countable, dense subset \widetilde{A} (e.g. Royden and Fitzpatrick, 2010, Proposition 9.26) and the map $\theta \mapsto \widehat{R}_n(f_{\theta})$ is continuous, it follows that $\inf_{\theta \in A} \widehat{R}_n(f_{\theta}) = \inf_{\theta \in \widetilde{A}} \widehat{R}_n(f_{\theta})$. Thus $\inf_{f \in \mathcal{F}} \widehat{R}_n(f)$ is measurable, and

$$E_{4} \leq \mathbb{E}\left[\frac{2}{n}\sum_{i=1}^{n}\left\{\left(\widetilde{f}(\widetilde{Z}_{i})-Y_{i}\right)^{2}-\left(f^{\star}(\widetilde{Z}_{i})-Y_{i}\right)^{2}\right\}\right]$$
$$= 2\mathbb{E}\left\{\widehat{R}_{n}(\widetilde{f})-\inf_{f\in\mathcal{F}}\widehat{R}_{n}(f)\right\}+2\mathbb{E}\left\{\inf_{f\in\mathcal{F}}\widehat{R}_{n}(f)-\widehat{R}_{n}(f^{\star})\right\}$$
$$\leq 2\mathbb{E}\left\{\widehat{R}_{n}(\widetilde{f})-\inf_{f\in\mathcal{F}}\widehat{R}_{n}(f)\right\}+2\inf_{f\in\mathcal{F}}\mathbb{E}\left\{\widehat{R}_{n}(f)-\widehat{R}_{n}(f^{\star})\right\}$$
$$= 2\mathbb{E}\left\{\widehat{R}_{n}(\widetilde{f})-\inf_{f\in\mathcal{F}}\widehat{R}_{n}(f)\right\}+2\inf_{f\in\mathcal{F}}\mathbb{E}\left\{\left(f(\widetilde{Z}_{0})-Y_{0}\right)^{2}-\left(f^{\star}(\widetilde{Z}_{0})-Y_{0}\right)^{2}\right\}$$

$$= 2\mathbb{E}\Big\{\widehat{R}_n(\widetilde{f}) - \inf_{f \in \mathcal{F}}\widehat{R}_n(f)\Big\} + 2\inf_{f \in \mathcal{F}}\mathbb{E}\Big\{\big(f(\widetilde{Z}_0) - f^{\star}(\widetilde{Z}_0)\big)^2\Big\}.$$
(10)

Combining (7), (8), (9) and (10) yields the final result.

D Proof of Proposition 2

Proof of Proposition 2. (a) Let $J := |\mathcal{J}|$. Without loss of generality, suppose that $\mathcal{J} = [J]$, and let $\mathcal{S}_{\mathcal{J}} = \{u_1, \ldots, u_{|\mathcal{S}_{\mathcal{J}}|}\} \subseteq \{0, 1\}^J$. For $i \in [|\mathcal{S}_{\mathcal{J}}|]$, define $x_i := \sum_{j=1}^J 2^{j-1} u_{ij} \in \mathbb{N}_0$, where u_{ij} is the *j*th coordinate of u_i . Further define $y_i := k/K$ where $k \in [K]$ is such that there exists $\boldsymbol{\omega} \in \mathcal{S}_k$ such that $\boldsymbol{\omega}_{\mathcal{J}} = u_i$; note that *k* is uniquely defined by our assumption. By Shen, Yang and Zhang (2019, Lemma 2.2) and (P1), there exists $f_1 \in \mathcal{F}(2, (1, p_*, p_*, 1))$ such that $f_1(x_i) = y_i$ for all $i \in [|\mathcal{S}_{\mathcal{J}}|]$. Thus, by (P2), there exists $f \in \mathcal{F}(2, p)$ such that

$$f(\boldsymbol{v}) = f_1\left(\sum_{j=1}^J 2^{j-1} v_j\right)$$

for $\boldsymbol{v} = (v_1, \dots, v_d)^\top \in \mathbb{R}^d$. For any $k \in [K]$ and $\boldsymbol{\omega} = (\omega_1, \dots, \omega_d)^\top \in \mathcal{S}_k$, there exists $i \in |\mathcal{S}_{\mathcal{J}}|$ such that $\boldsymbol{\omega}_{\mathcal{J}} = \boldsymbol{u}_i$. Therefore,

$$f(\boldsymbol{\omega}) = f_1\left(\sum_{j=1}^J 2^{j-1}\omega_j\right) = f_1\left(\sum_{j=1}^J 2^{j-1}u_{ij}\right) = f_1(x_i) = \frac{k}{K}$$

Hence, S_1, \ldots, S_K are $\mathcal{F}(2, \boldsymbol{p})$ -separable with $\epsilon = 1/(2K)$.

(b) First define

$$\epsilon \coloneqq \min_{k \in [K]} \min_{\boldsymbol{\omega} \in \mathcal{S} \setminus \mathcal{S}_k} \max_{\ell \in [P_k]} (\boldsymbol{\omega}^\top \boldsymbol{v}_{\ell}^{(k)} - b_{\ell}^{(k)}).$$

For each $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S} \setminus \mathcal{S}_k$, we have $\max_{\ell \in [P_k]} (\boldsymbol{\omega}^\top \boldsymbol{v}_{\ell}^{(k)} - b_{\ell}^{(k)}) > 0$ by assumption, so $\epsilon > 0$. For $k \in [K]$ and $\ell \in [P_k]$, define $\phi_{(k,\ell)} \in \mathcal{F}(1, (d, 2, 1))$ by

$$\phi_{(k,\ell)}(\boldsymbol{x}) \coloneqq \sigma\left(-\frac{\boldsymbol{x}^{\top}\boldsymbol{v}_{\ell}^{(k)} - b_{\ell}^{(k)}}{\epsilon} + 1\right) - \sigma\left(-\frac{\boldsymbol{x}^{\top}\boldsymbol{v}_{\ell}^{(k)} - b_{\ell}^{(k)}}{\epsilon}\right)$$

for $\boldsymbol{x} \in \mathbb{R}^{d}$. If $\boldsymbol{x}^{\top} \boldsymbol{v}_{\ell}^{(k)} - b_{\ell}^{(k)} \leq 0$, then $\phi_{(k,\ell)}(\boldsymbol{x}) = 1$; if $\boldsymbol{x}^{\top} \boldsymbol{v}_{\ell}^{(k)} - b_{\ell}^{(k)} \geq \epsilon$, then $\phi_{(k,\ell)}(\boldsymbol{x}) = 0$. By (P4), the function $\boldsymbol{\phi} : \mathbb{R}^{d} \to \mathbb{R}^{\sum_{k=1}^{K} P_{k}}$ defined by $\boldsymbol{\phi}(\boldsymbol{x}) \coloneqq (\phi_{(k,\ell)}(\boldsymbol{x}))_{k\in[K],\ell\in[P_{k}]}$ belongs to $\mathcal{F}(1, (d, 2\sum_{k=1}^{K} P_{k}, \sum_{k=1}^{K} P_{k}))$. For $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_{k}$, we have $\sum_{\ell=1}^{P_{k}} \phi_{(k,\ell)}(\boldsymbol{\omega}) = P_{k}$, and for all $k' \neq k$, we have $\sum_{\ell'=1}^{P_{k'}} \phi_{(k',\ell')}(\boldsymbol{\omega}) \leq P_{k'} - 1$ since there exists $\ell^{*} \in [P_{k'}]$ such that $\boldsymbol{\omega}^{\top} \boldsymbol{v}_{\ell^{*}}^{(k')} - b_{\ell^{*}}^{(k')} \geq \epsilon$. Next define $\psi \in \mathcal{F}(1, (\sum_{k=1}^{K} P_{k}, K, 1))$ by

$$\psi(\boldsymbol{u}) \coloneqq \sum_{k=1}^{K} k \cdot \sigma \left(\sum_{\ell=1}^{P_k} u_{k,\ell} - P_k + 1 \right)$$

for $\boldsymbol{u} = (u_{k,\ell})_{k \in [K], \ell \in [P_k]} \in \mathbb{R}^{\sum_{k=1}^{K} P_k}$. Then, for $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_k$, we have $\psi(\boldsymbol{\phi}(\boldsymbol{\omega})) = k$. Finally, $\psi \circ \boldsymbol{\phi} \in \mathcal{F}(2, \boldsymbol{p})$ by (P2), so $\{\mathcal{S}_1, \ldots, \mathcal{S}_K\}$ is $\mathcal{F}(2, \boldsymbol{p})$ -separable.

E Approximation theory for neural networks

The following lemma follows the arguments of Lu et al. (2021, Theorem 2.2) and Jiao et al. (2023, Theorem 3.3), with very minor changes.

Lemma 6. Let $\beta, \gamma > 0$, $\beta_0 \coloneqq \lceil \beta \rceil - 1$, $d \in \mathbb{N}$ and $g \in \mathcal{H}^{\beta}_d([0,1]^d, \gamma)$. For any $M, N \in \mathbb{N}$, let

$$R \coloneqq \lfloor N^{1/d} \rfloor^2 \lfloor M^{2/d} \rfloor, \quad L \coloneqq 12(\beta_0 + 1)^2 (M + 2) \lceil \log_2(4M) \rceil,$$
$$p_* \coloneqq 30(\beta_0 + 1)^2 d^{\beta_0 + 1} (N + 1) \lceil \log_2(8N) \rceil \quad and \quad \boldsymbol{p} \coloneqq (d, p_*, \dots, p_*, 1)^\top \in \mathbb{N}^{L+2}.$$

For $\delta > 0$, let

$$\Gamma([0,1]^d, R, \delta) \coloneqq \bigcup_{j=1}^d \left\{ \boldsymbol{x} = (x_1, \dots, x_d)^\top \in [0,1]^d : x_j \in \bigcup_{r=1}^{R-1} \left(\frac{r}{R} - \delta, \frac{r}{R} \right) \right\}.$$

Then, for any $\delta \in (0, \frac{1}{3R}]$, there exists $f \in \mathcal{F}(L, \mathbf{p})$ such that $|f(\mathbf{x})| \leq \gamma$ for all $\mathbf{x} \in \mathbb{R}^d$, and

$$|f(\boldsymbol{x}) - g(\boldsymbol{x})| \le 9\gamma(\beta_0 + 3)^2 8^{\beta} d^{\beta_0 + \beta/2} (NM)^{-2\beta/d},$$

for all $\boldsymbol{x} \in [0,1]^d \setminus \Gamma([0,1]^d, R, \delta)$.

Proof. Step 1 (discretisation): By Lu et al. (2021, Proposition 4.3), there exists $\phi_1 \in \mathcal{F}(4M + 5, (1, p_{*1}, \ldots, p_{*1}, 1))$, where $p_{*1} \coloneqq 4\lfloor N^{1/d} \rfloor + 3$ such that

$$\phi_1(x) = \frac{r}{R}$$
 if $x \in \left[\frac{r}{R}, \frac{r+1}{R} - \delta \cdot \mathbb{1}_{\{r \le R-2\}}\right]$ for $r \in \{0, 1, \dots, R-1\}$.

For $\boldsymbol{v} = (v_1, \dots, v_d)^{\top} \in \{0, 1, \dots, R-1\}^d$, define

$$Q_{\boldsymbol{v}} \coloneqq \left\{ \boldsymbol{x} = (x_1, \dots, x_d)^\top \in [0, 1]^d : x_j \in \left[\frac{v_j}{R}, \frac{v_j + 1}{R} - \delta \cdot \mathbb{1}_{\{v_j \le R - 2\}} \right] \text{ for all } j \in [d] \right\},\$$

so that $[0,1]^d \setminus \Gamma([0,1]^d, R, \delta) = \bigcup_{\boldsymbol{v} \in \{0,\dots,R-1\}^d} Q_{\boldsymbol{v}}$. By Property (P4), the function $\boldsymbol{\phi}_{dsc}$: $\mathbb{R}^d \to \mathbb{R}^d$ given by $\boldsymbol{\phi}_{dsc}(\boldsymbol{x}) \coloneqq (\boldsymbol{\phi}_1(x_1),\dots,\boldsymbol{\phi}_1(x_d))^\top$ for $\boldsymbol{x} = (x_1,\dots,x_d)^\top$ belongs to $\mathcal{F}(4M+5, (d, dp_{*1},\dots, dp_{*1}, d))$. Thus,

$$\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}) = \left(\frac{v_1}{R}, \dots, \frac{v_d}{R}\right)^{\top}$$
 for all $\boldsymbol{x} \in Q_{\boldsymbol{v}}, \, \boldsymbol{v} \in \{0, 1, \dots, R-1\}^d$.

Step 2 (approximation of Taylor coefficients): The function $\phi_2 : \mathbb{R}^d \to \mathbb{R}$ given by $\phi_2(\boldsymbol{x}) \coloneqq \sum_{j=1}^d R^j \phi_1(x_j)$ is a composition of $\boldsymbol{\phi}_{dsc}$ with a projection along the vector $(R, R^2, \ldots, R^d)^\top$ and hence by (P2), $\phi_2 \in \mathcal{F}(4M + 5, (d, dp_{*1}, \ldots, dp_{*1}, 1))$. Moreover, for all $\boldsymbol{v} \in \{0, 1, \ldots, R-1\}^d$ and $\boldsymbol{x} \in Q_{\boldsymbol{v}}$, we have

$$\phi_2(\boldsymbol{x}) = \sum_{j=1}^d R^{j-1} v_j \eqqcolon I_{\boldsymbol{v}} \in \{0, 1, \dots, R^d - 1\}.$$
(11)

For each $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ with $\|\boldsymbol{\alpha}\|_1 \leq \beta_0$ we have $\frac{\partial^{\boldsymbol{\alpha}} g(\boldsymbol{x}) + \gamma}{2\gamma} \in [0, 1]$ for $\boldsymbol{x} \in [0, 1]^d$. Hence, since $R^d \leq N^2 M^2$ and since $\boldsymbol{v} \mapsto I_{\boldsymbol{v}}$ is a bijection, by Lu et al. (2021, Proposition 4.4) there exists $\psi_{\boldsymbol{\alpha}}^{\circ} \in \mathcal{F}(5(M+2)\lceil \log_2(4M) \rceil, (1, p_{*2}, \dots, p_{*2}, 1))$, where $p_{*2} \coloneqq 16(\beta_0 + 1)(N + 1)\lceil \log_2(8N) \rceil$, such that

$$\left|\psi_{\boldsymbol{\alpha}}^{\circ}(I_{\boldsymbol{v}}) - \frac{\partial^{\boldsymbol{\alpha}}g(\boldsymbol{v}/R) + \gamma}{2\gamma}\right| \le (NM)^{-2(\beta_0+1)} \quad \text{for all } \boldsymbol{v} \in \{0, 1, \dots, R-1\}^d.$$
(12)

Now

$$dp_{*1} \lor p_{*2} \le 16d(\beta_0 + 1)(N + 1)\lceil \log_2(8N) \rceil =: p_{*3}$$

and

$$4M + 5 + 5(M + 2)\lceil \log_2(4M) \rceil \le 7(M + 2)\lceil \log_2(4M) \rceil =: L_0$$

Hence, by Properties (P1), (P2) and (P3), there exists $\psi_{\alpha} \in \mathcal{F}(L_0, (d, p_{*3}, \ldots, p_{*3}, 1))$ such that $\psi_{\alpha}(\boldsymbol{x}) = 2\psi_{\alpha}^{\circ}(\phi_2(\boldsymbol{x})) - 1 \in [-1, 1]$ for all $\boldsymbol{x} \in [0, 1]^d$. By (11) and (12), we deduce that for all $\boldsymbol{v} \in \{0, 1, \ldots, R-1\}^d$ and $\boldsymbol{x} \in Q_{\boldsymbol{v}}$, we have

$$\left|\psi_{\boldsymbol{\alpha}}(\boldsymbol{x}) - \frac{1}{\gamma} \cdot \partial^{\boldsymbol{\alpha}} g(\boldsymbol{\phi}_{\rm dsc}(\boldsymbol{x}))\right| \leq 2(NM)^{-2(\beta_0+1)}.$$
(13)

Step 3 (local Taylor expansion): By Lu et al. (2021, Lemma 4.2), there exists $\phi_{\times} \in \mathcal{F}(2(\beta_0 + 1)(M + 1), (1, 9(N + 1) + 1, \dots, 9(N + 1) + 1, 1))$ such that for all $x, y \in [-1, 1.1]$, we have

$$|\phi_{\times}(x,y) - xy| \le 27(N+1)^{-2(\beta_0+1)(M+1)}.$$
(14)

Moreover, by Lu et al. (2021, Proposition 4.1), for $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ with $\|\boldsymbol{\alpha}\|_1 \leq \beta_0$, there exists $\operatorname{Poly}_{\boldsymbol{\alpha}} \in \mathcal{F}(7(\beta_0+1)^2 M, (d, 9(N+1) + \beta_0, \dots, 9(N+1) + \beta_0, 1))$ such that

$$|\operatorname{Poly}_{\alpha}(\boldsymbol{x}) - \boldsymbol{x}^{\alpha}| \le 9(\beta_0 + 1)(N + 1)^{-7(\beta_0 + 1)M},$$
 (15)

for all $\boldsymbol{x} \in [0, 1]^d$. Moreover, since $9(\beta_0 + 1)(N + 1)^{-7(\beta_0 + 1)M} \leq 9(\beta_0 + 1)2^{-7(\beta_0 + 1)} < 0.1$, we have $\text{Poly}_{\boldsymbol{\alpha}}(\boldsymbol{x}) \in [-1, 1.1]$ for all $\boldsymbol{x} \in [0, 1]^d$. Now define $f^\circ : [0, 1]^d \to \mathbb{R}$ by

$$f^{\circ}(\boldsymbol{x}) \coloneqq \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\gamma}{\boldsymbol{\alpha}!} \cdot \phi_{\times} \Big(\psi_{\boldsymbol{\alpha}}(\boldsymbol{x}), \operatorname{Poly}_{\boldsymbol{\alpha}} \big(\boldsymbol{x} - \boldsymbol{\phi}_{\operatorname{dsc}}(\boldsymbol{x}) \big) \Big).$$

Since $|\{\boldsymbol{\alpha} \in \mathbb{N}_0^d : \|\boldsymbol{\alpha}\|_1 \leq \beta_0\}| = \binom{d+\beta_0}{\beta_0} \leq (\beta_0+1)d^{\beta_0}$, we have by Properties (P1)– (P4) that $f^{\circ} \in \mathcal{F}(L-1, \boldsymbol{p}^{\circ})$ where $\boldsymbol{p}^{\circ} \coloneqq (d, p_*, \dots, p_*, 1)^{\top} \in \mathbb{N}^{L+1}$. Moreover, for $\boldsymbol{v} \in \{0, \dots, R-1\}^d$ and $\boldsymbol{x} \in Q_{\boldsymbol{v}}$, we have

$$\begin{split} \left| f^{\circ}(\boldsymbol{x}) - \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\partial^{\boldsymbol{\alpha}} g(\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))}{\boldsymbol{\alpha}!} (\boldsymbol{x} - \boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))^{\boldsymbol{\alpha}} \right| \\ &\leq \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\gamma}{\boldsymbol{\alpha}!} \cdot \left| \boldsymbol{\phi}_{\times} \Big(\psi_{\boldsymbol{\alpha}}(\boldsymbol{x}), \operatorname{Poly}_{\boldsymbol{\alpha}} \big(\boldsymbol{x} - \boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}) \big) \Big) - \psi_{\boldsymbol{\alpha}}(\boldsymbol{x}) \cdot \operatorname{Poly}_{\boldsymbol{\alpha}} \big(\boldsymbol{x} - \boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}) \big) \right| \\ &+ \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\gamma}{\boldsymbol{\alpha}!} \cdot \left| \psi_{\boldsymbol{\alpha}}(\boldsymbol{x}) - \frac{1}{\gamma} \cdot \partial^{\boldsymbol{\alpha}} g\big(\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}) \big) \right| \cdot \operatorname{Poly}_{\boldsymbol{\alpha}} \big(\boldsymbol{x} - \boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}) \big) \end{split}$$

$$+\sum_{\boldsymbol{\alpha}\in\mathbb{N}_{0}^{d}:\|\boldsymbol{\alpha}\|_{1}\leq\beta_{0}}\frac{\partial^{\boldsymbol{\alpha}}g(\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))}{\boldsymbol{\alpha}!}\Big|\operatorname{Poly}_{\boldsymbol{\alpha}}\big(\boldsymbol{x}-\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x})\big)-\big(\boldsymbol{x}-\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x})\big)^{\boldsymbol{\alpha}}\Big|$$

$$\leq\gamma(\beta_{0}+1)d^{\beta_{0}}\big\{27(N+1)^{-2(\beta_{0}+1)(M+1)}+3(NM)^{-2(\beta_{0}+1)}+9(\beta_{0}+1)(N+1)^{-7(\beta_{0}+1)M}\big\}$$

$$\leq\gamma(\beta_{0}+1)(9\beta_{0}+39)d^{\beta_{0}}(NM)^{-2(\beta_{0}+1)},$$
(16)

where the second inequality follows from (14), (13) and (15). Therefore, we deduce that, for $\boldsymbol{v} \in \{0, \ldots, R-1\}^d$ and $\boldsymbol{x} \in Q_{\boldsymbol{v}}$,

$$\begin{split} |f^{\circ}(\boldsymbol{x}) - g(\boldsymbol{x})| &\leq \left| f^{\circ}(\boldsymbol{x}) - \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\partial^{\boldsymbol{\alpha}} g(\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))}{\boldsymbol{\alpha}!} (\boldsymbol{x} - \boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))^{\boldsymbol{\alpha}} \right| \\ &+ \left| g(\boldsymbol{x}) - \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\partial^{\boldsymbol{\alpha}} g(\boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))}{\boldsymbol{\alpha}!} (\boldsymbol{x} - \boldsymbol{\phi}_{\mathrm{dsc}}(\boldsymbol{x}))^{\boldsymbol{\alpha}} \right| \\ &\leq \gamma(\beta_{0} + 1)(9\beta_{0} + 39)d^{\beta_{0}}(NM)^{-2(\beta_{0}+1)} + \gamma d^{\beta_{0}}(\sqrt{d}/R)^{\beta} \\ &\leq \gamma(\beta_{0} + 1)(9\beta_{0} + 39)d^{\beta_{0}}(NM)^{-2(\beta_{0}+1)} + \gamma d^{\beta_{0}+\beta/2}8^{\beta}(NM)^{-2\beta/d} \\ &\leq 9\gamma(\beta_{0} + 3)^{2}8^{\beta}d^{\beta_{0}+\beta/2}(NM)^{-2\beta/d}, \end{split}$$

where the second inequality follows from (16) and Lemma 14, and the third inequality follows since $R \geq \frac{1}{8}(NM)^{2/d}$. Finally, we define $f : \mathbb{R}^d \to [-\gamma, \gamma]$ by

$$f(\boldsymbol{x}) \coloneqq \sigma \big(f^{\circ}(\boldsymbol{x}) + \gamma \big) - \sigma \big(f^{\circ}(\boldsymbol{x}) - \gamma \big) - \gamma = T_{\gamma} f^{\circ}(\boldsymbol{x}).$$

By (P1) and (P2), we have $f \in \mathcal{F}(L, \boldsymbol{p})$. Moreover, since $g(\boldsymbol{x}) \in [-\gamma, \gamma]$ for all $\boldsymbol{x} \in [0, 1]^d$, we have $|f(\boldsymbol{x}) - g(\boldsymbol{x})| \leq |f^{\circ}(\boldsymbol{x}) - g(\boldsymbol{x})| \leq 9\gamma(\beta_0 + 3)^{2}8^{\beta}d^{\beta_0 + \beta/2}(NM)^{-2\beta/d}$ for $\boldsymbol{v} \in \{0, \ldots, R-1\}^d$ and $\boldsymbol{x} \in Q_{\boldsymbol{v}}$.

Our next lemma quantifies the extent to which a Hölder function on a bounded hypercube that depends only on a subset of the variables can be approximated by a neural network once we excise a finite set of strips in each coordinate.

Lemma 7. Let $-\infty < a < c < \infty$, $\beta, \gamma > 0$, $\beta_0 \coloneqq \lceil \beta \rceil - 1$, $d \in \mathbb{N}$, $t \in [d] \cup \{0\}$ and $g \in \mathcal{H}_t^{\beta}([a,c]^d,\gamma)$. For any $M, N \in \mathbb{N}$, let

$$R \coloneqq \lfloor N^{1/t} \rfloor^2 \lfloor M^{2/t} \rfloor, \quad L \coloneqq 12(\beta_0 + 1)^2 (M + 2) \lceil \log_2(4M) \rceil,$$
$$p_* \coloneqq 30(\beta_0 + 1)^2 t^{\beta_0 + 1} (N + 1) \lceil \log_2(8N) \rceil \lor 1 \quad and \quad \boldsymbol{p} \coloneqq (d, p_*, \dots, p_*, 1)^\top \in \mathbb{N}^{L+2}.$$

Then, for any $\delta \in (0, \frac{1}{3R}]$, there exists $f \in \mathcal{F}(L, \mathbf{p})$ such that $|f(\mathbf{x})| \leq \gamma \vee (c-a)^{\beta_0} \gamma$ for all $\mathbf{x} \in \mathbb{R}^d$, and

$$|f(\boldsymbol{x}) - g(\boldsymbol{x})| \le 9 (1 \lor (c-a)^{\beta_0}) \gamma (\beta_0 + 3)^2 8^{\beta} t^{\beta_0 + \beta/2} (NM)^{-2\beta/t},$$

for all $\boldsymbol{x} \in [a, c]^d \setminus \Gamma([a, c]^d, R, \delta)$, where

$$\Gamma\left([a,c]^d,R,\delta\right) \coloneqq \bigcup_{j=1}^d \left\{ \boldsymbol{x} = (x_1,\dots,x_d)^\top \in [a,c]^d : \frac{x_j-a}{c-a} \in \bigcup_{r=1}^{R-1} \left(\frac{r}{R}-\delta, \frac{r}{R}\right) \right\}$$

if $t \in [d]$, and $\Gamma([a, c]^d, R, \delta) \coloneqq \emptyset$ if t = 0.

Proof. In the case where t = 0, the function g is constant, so belongs to $\mathcal{F}(L, \mathbf{p})$ and we have zero approximation error. We next consider the case where t = d. Define $h: [0, 1]^d \to \mathbb{R}$ by

$$h(\boldsymbol{x}) \coloneqq g(a\boldsymbol{1}_d + (c-a)\boldsymbol{x}),$$

so that $h \in \mathcal{H}_d^\beta([0,1]^d, \gamma \lor (c-a)^{\beta_0}\gamma)$. By Lemma 6, there exists $f^\circ \in \mathcal{F}(L, p)$ such that $|f^\circ(\boldsymbol{x})| \leq \gamma \lor (c-a)^{\beta_0}\gamma$ for all $\boldsymbol{x} \in \mathbb{R}^d$, and

$$|f^{\circ}(\boldsymbol{x}) - h(\boldsymbol{x})| \le 9 (\gamma \vee (c-a)^{\beta_0} \gamma) (\beta_0 + 3)^2 8^{\beta} t^{\beta_0 + \beta/2} (NM)^{-2\beta/t}$$
(17)

for all $\boldsymbol{x} \in [0,1]^d \setminus \Gamma([0,1]^d, R, \delta)$. By (P2), there exists $f \in \mathcal{F}(L, \boldsymbol{p})$ such that

$$f(\boldsymbol{x}) = f^{\circ}\left(\frac{\boldsymbol{x} - a\boldsymbol{1}_d}{c - a}\right),$$

and $|f(\boldsymbol{x})| \leq \gamma \vee (c-a)^{\beta_0} \gamma$ for all $\boldsymbol{x} \in \mathbb{R}^d$. Thus, by (17),

$$|f(\boldsymbol{x}) - g(\boldsymbol{x})| \le 9(\gamma \lor (c-a)^{\beta_0} \gamma)(\beta_0 + 3)^2 8^{\beta} t^{\beta_0 + \beta/2} (NM)^{-2\beta/t}$$

for all $\boldsymbol{x} \in [a,c]^d \setminus \Gamma([a,c]^d, R, \delta)$. This proves the claim when t = d.

Now assume that $t \in [d-1]$. Without loss of generality, assume that g depends only on the coordinates in [t]. Define $g^{[t]} : [a,c]^t \to \mathbb{R}$ by $g^{[t]}(\boldsymbol{y}) \coloneqq g(\boldsymbol{y}, \boldsymbol{0}_{d-t})$, so that $g^{[t]} \in \mathcal{H}_t^\beta([a,c]^t,\gamma)$. Then, by the case where t = d, there exists $f^{[t]} \in \mathcal{F}(L, \boldsymbol{p}^{[t]})$ where $\boldsymbol{p}^{[t]} \coloneqq (t, p_*, \dots, p_*, 1)^\top \in \mathbb{N}^{L+2}$ such that $|f^{[t]}(\boldsymbol{y})| \leq \gamma \lor (c-a)^{\beta_0} \gamma$ for all $\boldsymbol{y} \in \mathbb{R}^t$, and

$$\left| f^{[t]}(\boldsymbol{y}) - g^{[t]}(\boldsymbol{y}) \right| \le 9 \left(\gamma \lor (c-a)^{\beta_0} \gamma \right) (\beta_0 + 3)^2 8^{\beta} t^{\beta_0 + \beta/2} (NM)^{-2\beta/2}$$

for all $\boldsymbol{y} \in [a,c]^t \setminus \Gamma([a,c]^t, R, \delta)$. By (P2), we can define $f \in \mathcal{F}(L, \boldsymbol{p})$ by $f(\boldsymbol{x}) \coloneqq f^{[t]}(\boldsymbol{M}\boldsymbol{x})$, where $\boldsymbol{M} \coloneqq (\boldsymbol{I}_t \quad \boldsymbol{0}) \in \mathbb{R}^{t \times d}$, and f satisfies the requirements in the statement. \Box

By Assumption 2, for each $k \in [K]$, there exist $\boldsymbol{g}_1^{(k)}, \ldots, \boldsymbol{g}_{q_k}^{(k)}$ such that

$$\boldsymbol{g}_{r}^{(k)} = \left(g_{r,1}^{(k)}, \dots, g_{r,d_{r+1}}^{(k)}\right)^{\top} : \mathbb{R}^{d_{r}^{(k)}} \to \mathbb{R}^{d_{r+1}^{(k)}},$$
$$g_{r,j}^{(k)} \in \mathcal{H}_{t_{r}^{(k)}}^{\beta_{r}^{(k)}} \left(\mathbb{R}^{d_{r}^{(k)}}, \gamma_{r}^{(k)}\right) \text{ for all } r \in [q_{k}], j \in [d_{r+1}^{(k)}]$$

and

$$f^{\mathcal{S}_k}(oldsymbol{z}) = oldsymbol{g}_{q_k}^{(k)} \circ oldsymbol{g}_{q_k-1}^{(k)} \circ \cdots \circ oldsymbol{g}_1^{(k)}(oldsymbol{z})$$

for all $\boldsymbol{z} \in \mathbb{R}^d$. Now, for $k \in [K]$, define

$$a_1^{(k)} \coloneqq -\xi_1 \log(2dn), \quad c_1^{(k)} \coloneqq \xi_1 \log(2dn)$$

and $a_r^{(k)} \coloneqq -\gamma_{r-1}^{(k)}, \quad c_r^{(k)} \coloneqq \gamma_{r-1}^{(k)} \text{ for } r \in \{2, \dots, q_k + 1\}.$ (18)

For $r \in \{2, \ldots, q_k + 1\}$, we have $\boldsymbol{g}_{r-1}^{(k)}(\boldsymbol{z}) \in [a_r^{(k)}, c_r^{(k)}]^{d_r^{(k)}}$ for all $\boldsymbol{z} \in \mathbb{R}^{d_{r-1}^{(k)}}$ by the Hölder property of these functions. Thus, it is sufficient to restrict the domain of $\boldsymbol{g}_r^{(k)}$ to $[a_r^{(k)}, c_r^{(k)}]^{d_r^{(k)}}$, for $r \in \{2, \ldots, q_k\}$. Then, Assumption 2 yields that

$$f^{\mathcal{S}_k}(\boldsymbol{z}) = \boldsymbol{g}_{q_k}^{(k)} \circ \boldsymbol{g}_{q_k-1}^{(k)} \cdots \circ \boldsymbol{g}_1^{(k)}(\boldsymbol{z})$$

for all $\boldsymbol{z} \in [a_1^{(k)}, c_1^{(k)}]^d = \left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d$, where

$$\boldsymbol{g}_{r}^{(k)} = \left(g_{r,1}^{(k)}, \dots, g_{r,d_{i+1}^{(k)}}^{(k)}\right)^{\top} : \left[a_{r}^{(k)}, c_{r}^{(k)}\right]^{d_{r}^{(k)}} \to \left[a_{r+1}^{(k)}, c_{r+1}^{(k)}\right]^{d_{r+1}^{(k)}}, \text{ and}$$
$$g_{r,j}^{(k)} \in \mathcal{H}_{t_{r}^{(k)}}^{\beta_{r}^{(k)}}\left(\left[a_{r}^{(k)}, c_{r}^{(k)}\right]^{d_{r}^{(k)}}, \gamma_{r}^{(k)}\right) \text{ for all } r \in [q_{k}], j \in [d_{r+1}^{(k)}].$$
(19)

The reason that we are restricting the domain of $f^{\mathcal{S}_k}$ to $\left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d$ is that by Assumption 1, each coordinate of \mathbb{Z}_0 is sub-exponential, so $\mathbb{P}\left(|Z_{0,j}| \ge \xi_1 \log(2dn)\right) \le 1/(dn)$ for all $j \in [d]$. Thus, by a union bound,

$$\mathbb{P}\left\{\boldsymbol{Z}_{0} \notin \left[-\xi_{1}\log(2dn), \,\xi_{1}\log(2dn)\right]^{d}\right\} \leq \frac{1}{n},\tag{20}$$

so it will suffice to approximate $f^{\mathcal{S}_k}$ on $\left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d$. To this end, we begin by approximating each function in the composition that defines $f^{\mathcal{S}_k}$.

Lemma 8. Suppose that Assumptions 1 and 2 hold. With the notation in (18) and (19), for $M, N \in \mathbb{N}$ and $k \in [K]$, let

$$L^{(k)} \coloneqq \max_{r \in [q_k]} 12\lceil \beta_r^{(k)} \rceil^2 (M+2) \lceil \log_2(4M) \rceil,$$

$$p_*^{(k)} \coloneqq \max_{r \in [q_k]} \left\{ 30d_{r+1}^{(k)} \lceil \beta_r^{(k)} \rceil^2 (t_r^{(k)})^{\lceil \beta_r^{(k)} \rceil} (N+1) \lceil \log_2(8N) \rceil \lor 2d_{r+1}^{(k)} \right\} and$$

$$\boldsymbol{p}_r^{(k)} \coloneqq (d_r^{(k)}, p_*^{(k)}, \dots, p_*^{(k)}, d_{r+1}^{(k)})^\top \in \mathbb{N}^{L^{(k)}+2} \text{ for } r \in [q_k].$$

For each $k \in [K]$, $r \in [q_k]$ and $n \in \mathbb{N}$, there exist $\mathbf{f}_r^{(k)} \in \mathcal{F}(L^{(k)}, \mathbf{p}_r^{(k)})$ and $E_r^{(k)} \subseteq [a_r^{(k)}, c_r^{(k)}]^{d_r^{(k)}}$ satisfying:

$$\begin{aligned} (i) \ \| \boldsymbol{f}_{r}^{(k)}(\boldsymbol{x}) \|_{\infty} &\leq \gamma_{r}^{(k)} \vee (c_{r}^{(k)} - a_{r}^{(k)})^{\lceil \beta_{r}^{(k)} \rceil - 1} \gamma_{r}^{(k)} \text{ for all } \boldsymbol{x} \in \mathbb{R}^{d_{r}^{(k)}} \text{ and} \\ \| \boldsymbol{f}_{r}^{(k)}(\boldsymbol{x}) - \boldsymbol{g}_{r}^{(k)}(\boldsymbol{x}) \|_{\infty} \\ &\leq 9 \big(1 \vee (c_{r}^{(k)} - a_{r}^{(k)})^{\lceil \beta_{r}^{(k)} \rceil - 1} \big) \gamma_{r}^{(k)} (\lceil \beta_{r}^{(k)} \rceil + 2)^{2} 8^{\beta_{r}^{(k)}} (t_{r}^{(k)})^{\lceil \beta_{r}^{(k)} \rceil - 1 + \beta_{r}^{(k)}/2} (NM)^{-2\beta_{r}^{(k)}/t_{r}^{(k)}}, \\ &\text{for all } \boldsymbol{x} \in [a_{r}^{(k)}, c_{r}^{(k)}]^{d_{r}^{(k)}} \setminus E_{r}^{(k)}. \end{aligned}$$

(ii)
$$\mu_{\mathbf{Z}_0}(E_1^{(k)}) \leq \frac{1}{Knq_k} \text{ and } \mu_{\mathbf{Z}_0}((\mathbf{F}_{r-1}^{(k)})^{-1}(E_r^{(k)})) \leq \frac{1}{Knq_k} \text{ for all } r \in \{2, \dots, q_k\}, \text{ where } \mathbf{F}_r^{(k)} \coloneqq \mathbf{f}_r^{(k)} \circ \cdots \circ \mathbf{f}_1^{(k)}.$$

Proof. Fixing $k \in [K]$, we construct $(f_r^{(k)}, E_r^{(k)})_{r=1}^{q_k}$ inductively. For $r \in [q_k]$, let $R_r^{(k)} := \lfloor N^{1/t_r^{(k)}} \rfloor^2 \lfloor M^{2/t_r^{(k)}} \rfloor$ for $r \in [q_k]$. Let $\delta_{\max} := \frac{1}{3 \max_{r \in [q_k]} R_r^{(k)}}$. By Lemma 7, (P1), (P3) and (P4), for any $\delta \in (0, \delta_{\max}]$, there exists $\boldsymbol{f}_{1,\delta}^{(k)} \in \mathcal{F}(L^{(k)}, \boldsymbol{p}_1^{(k)})$ such that $\|\boldsymbol{f}_{1,\delta}^{(k)}(\boldsymbol{x})\|_{\infty} \leq (1 \vee (c_1^{(k)} - a_1^{(k)})^{\lceil \beta_1^{(k)} \rceil - 1}) \gamma_1^{(k)}$ for all $\boldsymbol{x} \in \mathbb{R}^{d_1^{(k)}}$ and

$$\left\| \boldsymbol{f}_{1,\delta}^{(k)}(\boldsymbol{x}) - \boldsymbol{g}_{1}^{(k)}(\boldsymbol{x}) \right\|_{\infty} \\ \leq 9 \left(1 \vee (c_{1}^{(k)} - a_{1}^{(k)})^{\lceil \beta_{1}^{(k)} \rceil - 1} \right) \gamma_{1}^{(k)} (\lceil \beta_{1}^{(k)} \rceil + 2)^{2} 8^{\beta_{1}^{(k)}} (t_{1}^{(k)})^{\lceil \beta_{1}^{(k)} \rceil - 1 + \beta_{1}^{(k)}/2} (NM)^{-2\beta_{1}^{(k)}/t_{1}^{(k)}},$$

for all $\boldsymbol{x} \in [a_1^{(k)}, c_1^{(k)}]^{d_1^{(k)}} \setminus E_{1,\delta}^{(k)}$, where $E_{1,\delta}^{(k)} \coloneqq \Gamma([a_1^{(k)}, c_1^{(k)}]^{d_1^{(k)}}, R_1^{(k)}, \delta)$ as defined in Lemma 7. Further note that $(E_{1,\delta}^{(k)})_{\delta \in (0,\delta_{\max}]}$ are nested and have empty intersection.

Thus, there exists $\delta^* \in (0, \delta_{\max}]$ such that $\mu_{Z_0}(E_{1,\delta^*}^{(k)}) \leq \frac{1}{Knq_k}$. We set $\boldsymbol{f}_1^{(k)} \coloneqq \boldsymbol{f}_{1,\delta^*}^{(k)}$ and $E_1^{(k)} \coloneqq E_{1,\delta^*}^{(k)}$, so that $\boldsymbol{f}_1^{(k)}$ and $E_1^{(k)}$ satisfy conditions (i) and (ii).

Now suppose that $\boldsymbol{f}_{1}^{(k)}, \ldots, \boldsymbol{f}_{r-1}^{(k)}$ and $E_{1}^{(k)}, \ldots, E_{r-1}^{(k)}$ satisfying conditions (i) and (ii) have been constructed for some $r \in \{2, \ldots, q_k\}$. We apply the same argument as above to construct $\boldsymbol{f}_r^{(k)}$ and $E_r^{(k)}$. By Lemma 7, (P1), (P3) and (P4), for any $\delta \in (0, \delta_{\max}]$, there exists $\boldsymbol{f}_{r,\delta}^{(k)} \in \mathcal{F}(L^{(k)}, \boldsymbol{p}_r^{(k)})$ such that $\|\boldsymbol{f}_{r,\delta}^{(k)}(\boldsymbol{x})\|_{\infty} \leq (1 \vee (c_r^{(k)} - a_r^{(k)})^{\lceil \beta_r^{(k)} \rceil - 1}) \gamma_r^{(k)}$ for all $\boldsymbol{x} \in \mathbb{R}^{d_r^{(k)}}$ and

$$\left\| \boldsymbol{f}_{r,\delta}^{(k)}(\boldsymbol{x}) - \boldsymbol{g}_{r}^{(k)}(\boldsymbol{x}) \right\|_{\infty}$$

$$\leq 9 \left(1 \vee (c_{r}^{(k)} - a_{r}^{(k)})^{\lceil \beta_{r}^{(k)} \rceil - 1} \right) \gamma_{r}^{(k)} (\lceil \beta_{r}^{(k)} \rceil + 2)^{2} 8^{\beta_{r}^{(k)}} (t_{r}^{(k)})^{\lceil \beta_{r}^{(k)} \rceil - 1 + \beta_{r}^{(k)}/2} (NM)^{-2\beta_{r}^{(k)}/t_{r}^{(k)}}$$

for all $\boldsymbol{x} \in [a_r^{(k)}, c_r^{(k)}]^{d_r^{(k)}} \setminus E_{r,\delta}^{(k)}$, where $E_{r,\delta}^{(k)} \coloneqq \Gamma([a_r^{(k)}, c_r^{(k)}]^{d_r^{(k)}}, R_r^{(k)}, \delta)$. Again, since $((\boldsymbol{F}_{r-1}^{(k)})^{-1}(E_{r,\delta}^{(k)}))_{\delta \in (0,\delta_{\max}]}$ are nested and have empty intersection, there exists $\delta^* \in (0, \delta_{\max}]$ such that $\mu_{\boldsymbol{Z}_0}((\boldsymbol{F}_{r-1}^{(k)})^{-1}(E_{r,\delta^*}^{(k)})) \leq \frac{1}{Knq_k}$. We set $\boldsymbol{f}_r^{(k)} \coloneqq \boldsymbol{f}_{r,\delta^*}^{(k)}$ and $E_r^{(k)} \coloneqq E_{r,\delta^*}^{(k)}$, so that $\boldsymbol{f}_r^{(k)}$ and $E_r^{(k)}$ satisfy conditions (i) and (ii).

We are now in a position to approximate each piece of the Bayes regression function by a neural network on a hypercube, once we have excised certain strips.

Lemma 9. Suppose that Assumptions 1 and 2 hold and let $n \ge 2$. For $M, N \in \mathbb{N}$ and $k \in [K]$, let

$$L^{(k)} \coloneqq \max_{r \in [q_k]} 12\lceil \beta_r^{(k)} \rceil^2 (M+2) \lceil \log_2(4M) \rceil \text{ and}$$
$$p_*^{(k)} \coloneqq \max_{r \in [q_k]} \Big\{ 30d_{r+1}^{(k)} \lceil \beta_r^{(k)} \rceil^2 (t_r^{(k)})^{\lceil \beta_r^{(k)} \rceil} (N+1) \lceil \log_2(8N) \rceil \lor 2d_{r+1}^{(k)} \Big\}.$$

For each $k \in [K]$, there exist $E^{(k)} \subseteq \left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d$ and $f^{(k)} \in \mathcal{F}(q_k L^{(k)}, \boldsymbol{p}^{(k)})$ where $\boldsymbol{p}^{(k)} \coloneqq (d, p_*^{(k)}, \dots, p_*^{(k)}, 1) \in \mathbb{N}^{q_k L^{(k)} + 2}$ such that $\mu_{\mathbf{Z}_0}(E^{(k)}) \leq \frac{1}{Kn}, |f^{(k)}(\boldsymbol{z})| \leq \gamma_{q_k}^{(k)} \vee (2\gamma_{q_k-1}^{(k)})^{\lceil \beta_{q_k}^{(k)} \rceil - 1} \gamma_{q_k}^{(k)}$ for all $\boldsymbol{z} \in \mathbb{R}^d$, and

$$\left| f^{(k)}(\boldsymbol{z}) - f^{\mathcal{S}_{k}}(\boldsymbol{z}) \right| \leq C_{1}(\xi_{1}, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)}) \cdot \log^{\bar{\beta}_{1}^{(k)}}(2dn) \sum_{r=1}^{q_{k}} \frac{1}{(NM)^{2\bar{\beta}_{r}^{(k)}/t_{r}^{(k)}}}$$

for all $z \in [-\xi_1 \log(2dn), \xi_1 \log(2dn)]^d \setminus E^{(k)}$, where $C_1(\xi_1, t^{(k)}, \beta^{(k)}, \gamma^{(k)}) > 0$ depends only on $(\xi_1, t^{(k)}, \beta^{(k)}, \gamma^{(k)})$, and $C_1(\xi_1, t^{(k)}, \beta^{(k)}, \gamma^{(k)}) \leq A(\xi_1, \beta^{(k)}, \gamma^{(k)}) \|t^{(k)}\|_{\infty}^{B(\beta^{(k)})}$ for some $A(\xi_1, \beta^{(k)}, \gamma^{(k)}), B(\beta^{(k)}) > 0$.

Proof. We use the notation in (18) and (19). Fix $k \in [K]$, and let $(\mathbf{f}_r^{(k)})_{r=1}^{q_k}$ and $(E_r^{(k)})_{r=1}^{q_k}$ be defined as in Lemma 8. By (P2), the function $f^{(k)} \coloneqq \mathbf{f}_{q_k}^{(k)} \circ \cdots \circ \mathbf{f}_1^{(k)}$ belongs to $\mathcal{F}(q_k L^{(k)}, \mathbf{p}^{(k)})$. For $r \in [q_k]$, recall the definition of $\mathbf{F}_r^{(k)}$ from Lemma 8 and let $\mathbf{G}_r^{(k)} \coloneqq$ $\mathbf{g}_r^{(k)} \circ \cdots \circ \mathbf{g}_1^{(k)}$. Define $E^{(k)} \coloneqq \bigcup_{r \in [q_k]} (\mathbf{F}_{r-1}^{(k)})^{-1} (E_r^{(k)})$ with $\mathbf{F}_0^{(k)}$ being the identity function. By Lemma 8 and a union bound, $\mu_{\mathbf{Z}_0}(E^{(k)}) \leq \frac{1}{K_n}$. For $\mathbf{f} : \mathbb{R}^d \to \mathbb{R}^m$ and $D \subseteq \mathbb{R}^d$, we define $\|\mathbf{f}\|_{L^{\infty}(D)} \coloneqq \sup_{\mathbf{x} \in D} \|\mathbf{f}(\mathbf{x})\|_{\infty}$. For $r \in [q_k]$, let $D_r^{(k)} \coloneqq [a_r^{(k)}, c_r^{(k)}]^{d_r^{(k)}} \setminus E_r^{(k)}$ and let $D^{(k)} \coloneqq [-\xi_1 \log(2dn), \xi_1 \log(2dn)]^d \setminus E^{(k)}$. Then, for $r \in \{2, \dots, q_k\}$,

$$\left\| \boldsymbol{F}_{r}^{(k)} - \boldsymbol{G}_{r}^{(k)} \right\|_{L^{\infty}(D^{(k)})} \leq \left\| \boldsymbol{f}_{r}^{(k)} \circ \boldsymbol{f}_{r-1}^{(k)} \circ \cdots \circ \boldsymbol{f}_{1}^{(k)} - \boldsymbol{g}_{r}^{(k)} \circ \boldsymbol{f}_{r-1}^{(k)} \circ \cdots \circ \boldsymbol{f}_{1}^{(k)} \right\|_{L^{\infty}(D^{(k)})}$$

$$+ \left\| \boldsymbol{g}_{r}^{(k)} \circ \boldsymbol{f}_{r-1}^{(k)} \circ \cdots \circ \boldsymbol{f}_{1}^{(k)} - \boldsymbol{g}_{r}^{(k)} \circ \boldsymbol{g}_{r-1}^{(k)} \circ \cdots \circ \boldsymbol{g}_{1}^{(k)} \right\|_{L^{\infty}(D^{(k)})}$$

$$\le \left\| \boldsymbol{f}_{r}^{(k)} - \boldsymbol{g}_{r}^{(k)} \right\|_{L^{\infty}(\boldsymbol{F}_{r-1}^{(k)}(D^{(k)}))} + \gamma_{r}^{(k)} \left\{ \left(t_{r}^{(k)} \right)^{1/2} \left\| \boldsymbol{F}_{r-1}^{(k)} - \boldsymbol{G}_{r-1}^{(k)} \right\|_{L^{\infty}(D^{(k)})} \right\}^{\beta_{r}^{(k)} \wedge 1}$$

$$\le \left\| \boldsymbol{f}_{r}^{(k)} - \boldsymbol{g}_{r}^{(k)} \right\|_{L^{\infty}(D^{(k)}_{r})} + \gamma_{r}^{(k)} \left(t_{r}^{(k)} \right)^{(\beta_{r}^{(k)} \wedge 1)/2} \left\| \boldsymbol{F}_{r-1}^{(k)} - \boldsymbol{G}_{r-1}^{(k)} \right\|_{L^{\infty}(D^{(k)})}^{\beta_{r}^{(k)} \wedge 1},$$

$$(21)$$

where the second inequality follows since each coordinate of $\boldsymbol{g}_{r}^{(k)}$ belongs to the class $\mathcal{H}_{t_{r}^{(k)}}^{\beta_{r}^{(k)}}([a_{r}^{(k)},c_{r}^{(k)}]^{d_{r}^{(k)}},\gamma_{r}^{(k)})$. Thus,

where $C_1(\xi_1, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)})$ has the properties claimed in the statement of the result. Here, the first inequality follows by applying (21) iteratively and using the fact that $(a+b)^t \leq a^t + b^t$ for $a, b \geq 0$ and $t \in [0,1]$, the second inequality follows from Lemma 8, and the third inequality follows by substituting the definitions of $(a_r^{(k)}, c_r^{(k)})_{r=1}^{q_k}$ in (18). Moreover, $\|f^{(k)}\|_{L^{\infty}(\mathbb{R}^d)} \leq \|\boldsymbol{f}_{q_k}^{(k)}\|_{L^{\infty}(\mathbb{R}^d)} \leq \gamma_{q_k}^{(k)} \vee (2\gamma_{q_{k-1}}^{(k)})^{\lceil \beta_{q_k}^{(k)}\rceil - 1}\gamma_{q_k}^{(k)}$ by Lemma 8.

The following lemma is used in the proof of Lemma 11 below, which quantifies the extent to which we can extract coordinates, based on a function f_2 that separates $\{S_1, \ldots, S_K\}$, using a neural network.

Lemma 10. For any B > 0 and $N, L \in \mathbb{N}$, there exists $\phi \in \mathcal{F}(L, (2, 9N+1, \dots, 9N+1, 1))$ such that for all $x \in [-B, B]$, we have

$$\phi(x,0) = 0$$
 and $|\phi(x,1) - x| \le \frac{12B}{N^L}$

Proof. By Lu et al. (2021, Lemma 5.2), there exists $\phi_1 \in \mathcal{F}(L, (2, 9N, \dots, 9N, 1))$ such that $|\phi_1(a, b) - ab| \leq 6N^{-L}$ for all $a, b \in [0, 1]$. Moreover, by Lu et al. (2021, Eq (5.2)), the function ϕ_1 is defined by $\phi_1(a, b) = 2(\psi(\frac{a+b}{2}) - \psi(\frac{a}{2}) - \psi(\frac{b}{2}))$, where ψ is the neural network constructed in Lu et al. (2021, Lemma 5.1), which satisfies $\psi(0) = 0$. Therefore, $\phi_1(a, 0) = 0$ for all $a \in [0, 1]$. By (P2) and (P4), there exists $\phi \in \mathcal{F}(L, (2, 9N + 1, \dots, 9N + 1, 1))$ such that $\phi(x, y) = 2B\phi_1(\frac{x+B}{2B}, y) - By$ for all $(x, y) \in \mathbb{R} \times [0, \infty)$. Moreover, for all $x \in [-B, B]$, we have $\phi(x, 0) = 2B\phi_1(\frac{x+B}{2B}, 0) = 0$ and

$$|\phi(x,1) - x| = 2B \left| \phi_1\left(\frac{x+B}{2B}, 1\right) - \frac{x+B}{2B} \right| \le \frac{12B}{N^L},$$

as required.

Lemma 11. Suppose that $\{S_1, \ldots, S_K\}$ is separated by $f_2 \in \mathcal{F}(L_2, p_2)$. Let $m \coloneqq p_{2,L_2+1}$ and B > 0. For any $M, N_1, \ldots, N_K \in \mathbb{N}$, let

$$s \coloneqq (13m+7)K + (2M+2)\sum_{k=1}^{K} (9N_k+1)^2 + 7,$$
$$p_* \coloneqq (4m+1)K \lor \sum_{k=1}^{K} (9N_k+1) \quad and \quad \mathbf{p} \coloneqq (K+m, p_*, \dots, p_*, 1) \in \mathbb{N}^{M+5}$$

Then there exists $f_3 \in \mathcal{F}(M+3, \boldsymbol{p}, s)$ such that $|f_3(\boldsymbol{u}, \boldsymbol{v})| \leq B$ for all $(\boldsymbol{u}, \boldsymbol{v}) \in \mathbb{R}^K \times \mathbb{R}^m$, and for all $\boldsymbol{u} = (u_1, \dots, u_K)^\top \in [-B, B]^K$, $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_k$, we have

$$\left|f_3(\boldsymbol{u}, \boldsymbol{f}_2(\boldsymbol{\omega})) - u_k\right| \le \frac{12B}{N_k^M}.$$

Proof. Suppose that $\{S_1, \ldots, S_K\}$ and $f_2 \in \mathcal{F}(L_2, p_2)$ satisfy Definition 4 for some $\epsilon > 0$ and $v_1, \ldots, v_K \in \mathbb{R}^m$. For $a \in \mathbb{R}$, define $h_a : \mathbb{R} \to [0, 1]$ by

$$h_a(x) \coloneqq \sigma\left(\frac{2x - 2a + 2\epsilon}{\epsilon}\right) - \sigma\left(\frac{2x - 2a + \epsilon}{\epsilon}\right) - \sigma\left(\frac{2x - 2a - \epsilon}{\epsilon}\right) + \sigma\left(\frac{2x - 2a - 2\epsilon}{\epsilon}\right).$$

Then $h_a \in \mathcal{F}(1, (1, 4, 1), 12)$, with $h_a(x) = 1$ for all $x \in \lfloor a - \frac{\epsilon}{2}, a + \frac{\epsilon}{2} \rfloor$ and $h_a(x) = 0$ for all $x \in (-\infty, a - \epsilon] \cup [a + \epsilon, \infty)$. Further, for $k \in [K]$, write $\boldsymbol{v}_k = (v_{k,1}, \dots, v_{k,m})^\top \in \mathbb{R}^m$ and define $g_k \in \mathcal{F}(2, (m, 4m, 1, 1), 13m + 2)$ by

$$g_k(\boldsymbol{x}) \coloneqq \sigma \left(\sum_{j=1}^m h_{v_{k,j}}(x_j) - m + 1 \right).$$

Then, for $\boldsymbol{\omega} \in \mathcal{S}_k$, we have $g_k(\boldsymbol{f}_2(\boldsymbol{\omega})) = \sigma(m-m+1) = 1$; for $\boldsymbol{\omega} \notin \mathcal{S}_k$, we have $\sum_{j=1}^m h_{v_{k,j}}(f_{2,j}(\boldsymbol{\omega})) \leq m-1$, so $g_k(\boldsymbol{f}_2(\boldsymbol{\omega})) = 0$. By Lemma 10, for $k \in [K]$, there exists $\phi_k \in \mathcal{F}(M, (2, 9N_k + 1, \ldots, 9N_k + 1, 1))$ such that $\phi_k(x, 0) = 0$ and $|\phi_k(x, 1) - x| \leq 12BN_k^{-M}$ for all $x \in [-B, B]$. Further, by (P1), (P2) and (P4), there exists $f_3^\circ \in \mathcal{F}(M+2, (K+m, p_*, \ldots, p_*, 1), s-7)$ such that $f_3^\circ(\boldsymbol{u}, \boldsymbol{v}) = \sum_{k=1}^K \phi_k(\boldsymbol{u}_k, g_k(\boldsymbol{v}))$ for all $\boldsymbol{u} \in [-B, B]^K$ and $\boldsymbol{v} \in \mathbb{R}^m$. Moreover, for $\boldsymbol{\omega} \in \mathcal{S}_k$ and $\boldsymbol{u} = (u_1, \ldots, u_K)^\top \in [-B, B]^K$, we have

$$\begin{split} \left| f_3^{\circ} \big(\boldsymbol{u}, \boldsymbol{f}_2(\boldsymbol{\omega}) \big) - u_k \right| &= \left| \sum_{\ell=1}^K \Big\{ \phi_k \Big(u_\ell, g_\ell \big(\boldsymbol{f}_2(\boldsymbol{\omega}) \big) \Big) - u_\ell g_\ell \big(\boldsymbol{f}_2(\boldsymbol{\omega}) \big) \Big\} \right| \\ &\leq \sum_{\ell=1}^K \Big| \phi_k \Big(u_\ell, g_\ell \big(\boldsymbol{f}_2(\boldsymbol{\omega}) \big) \Big) - u_\ell g_\ell \big(\boldsymbol{f}_2(\boldsymbol{\omega}) \big) \Big| \\ &= \left| \phi_k(u_k, 1) - u_k \right| + \sum_{\ell \neq k} \left| \phi_\ell(u_\ell, 0) - 0 \right| \leq \frac{12B}{N_k^M} \end{split}$$

Finally, define $f_3 : \mathbb{R}^{K+m} \to [-B, B]$ by

$$f_3 \coloneqq \sigma(f_3^\circ + B) - \sigma(f_3^\circ - B) - B = T_B f_3^\circ.$$

By (P1) and (P2), we have $f_3 \in \mathcal{F}(M+3, \boldsymbol{p}, s)$ and since $u_k \in [-B, B]$, we deduce that $\left|f_3(\boldsymbol{u}, \boldsymbol{f}_2(\boldsymbol{\omega})) - u_k\right| \leq \left|f_3^{\circ}(\boldsymbol{u}, \boldsymbol{f}_2(\boldsymbol{\omega})) - u_k\right|$ for all $\boldsymbol{u} = (u_1, \dots, u_K)^{\top} \in [-B, B]^K$, $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_k$.

F Proof of Theorem 3

Proof of Theorem 3. For any $M_1 \in \mathbb{N}$, let

$$L_1 \coloneqq \max_{k \in [K], r \in [q_k]} 12q_k \lceil \beta_r^{(k)} \rceil^2 (M_1 + 2) \lceil \log_2(4M_1) \rceil \in \mathbb{N}$$

and for $k \in [K]$, let

$$N_{1}^{(k)} \coloneqq \left\lceil M_{1}^{-1} n_{k}^{t_{*}^{(k)}/(4\bar{\beta}_{*}^{(k)}+2t_{*}^{(k)})} \right\rceil,$$
$$p_{1,*}^{(k)} \coloneqq \max_{r \in [q_{k}]} \left\{ 30d_{r+1}^{(k)} \lceil \beta_{r}^{(k)} \rceil^{2} (t_{r}^{(k)})^{\lceil \beta_{r}^{(k)} \rceil} (N_{1}^{(k)}+1) \lceil \log_{2}(8N_{1}^{(k)}) \rceil \lor 2d_{r+1}^{(k)} \right\},$$
$$\boldsymbol{p}_{1}^{(k)} \coloneqq (d, p_{1,*}^{(k)}, \dots, p_{1,*}^{(k)}, 1) \in \mathbb{N}^{L_{1}+2}.$$

By Lemma 9 and (P3), for $k \in [K]$ we can find $E^{(k)} \subseteq \left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d$ and $f_1^{(k)} \in \mathcal{F}(L_1, \boldsymbol{p}_1^{(k)})$ such that $\mu_{\boldsymbol{Z}_0}(E^{(k)}) \leq \frac{1}{Kn}$,

$$|f_1^{(k)}(\boldsymbol{z})| \le \max_{k \in [K]} \{ \gamma_{q_k}^{(k)} \lor (2\gamma_{q_k-1}^{(k)})^{\lceil \beta_{q_k}^{(k)} \rceil - 1} \gamma_{q_k}^{(k)} \} =: B$$

for all $\boldsymbol{z} \in \mathbb{R}^d$ and

$$\left| f_{1}^{(k)}(\boldsymbol{z}) - f^{\mathcal{S}_{k}}(\boldsymbol{z}) \right| \leq C_{1}(\xi_{1}, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)}) \cdot \log^{\bar{\beta}_{1}^{(k)}}(2dn) \cdot q_{k} \cdot n_{k}^{-\bar{\beta}_{k}^{(k)}/(2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)})} \\ \leq C_{2}(\xi_{1}, d, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)}) \cdot \log^{\bar{\beta}_{1}^{(k)}}(n) \cdot n_{k}^{-\bar{\beta}_{*}^{(k)}/(2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)})}$$
(22)

for all $\mathbf{z} \in \left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d \setminus E^{(k)}$. Furthermore, each $f_1^{(k)}$ has $V_1^{(k)} \coloneqq (d+L_1+1)p_{1,*}^{(k)} + (L_1-1)(p_{1,*}^{(k)})^2 + 1$ parameters. Thus, by (P4), the function $\mathbf{f}_1 \coloneqq (f_1^{(1)}, \ldots, f_1^{(K)})^\top$: $\mathbb{R}^d \to [-B, B]^K$ belongs to $\mathcal{F}(L_1, \mathbf{p}_1, s_1)$, where $\mathbf{p}_1 \coloneqq (d, \sum_{k=1}^K p_{1,*}^{(k)}, \ldots, \sum_{k=1}^K p_{1,*}^{(k)}, K) \in \mathbb{N}^{L_1+2}$ and $s_1 \coloneqq \sum_{k=1}^K V_1^{(k)}$.

By assumption, there exists $f_2 \in \mathcal{F}(L_2, p_2, s_2)$ such that $\{S_1, \ldots, S_K\}$ is separated by f_2 . For any $M_3 \in \mathbb{N}$, let $m \coloneqq p_{2,L_2+1}$, and let

$$L_{3} \coloneqq \left[2M_{3} \max_{k \in [K]} \frac{\bar{\beta}_{*}^{(k)}}{t_{*}^{(k)}} \right] + 3, \quad N_{3,k} \coloneqq \left[n_{k}^{\bar{\beta}_{*}^{(k)} / \{(L_{3}-3)(2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)})\}} \right] \text{ for } k \in [K],$$

$$p_{3,*} \coloneqq (4m+1)K \lor \sum_{k=1}^{K} (9N_{3,k}+1), \quad \boldsymbol{p}_{3} \coloneqq (K+m, p_{3,*}, \dots, p_{3,*}, 1) \in \mathbb{N}^{L_{3}+2} \text{ and}$$

$$s_{3} \coloneqq (13m+7)K + (2M_{3}+2)\sum_{k=1}^{K} (9N_{3,k}+1)^{2} + 7.$$

By Lemma 11, there exists $f_3 \in \mathcal{F}(L_3, \boldsymbol{p}_3, s)$ such that $|f_3(\boldsymbol{u}, \boldsymbol{v})| \leq B$ for all $(\boldsymbol{u}, \boldsymbol{v}) \in \mathbb{R}^K \times \mathbb{R}^m$, and that for all $\boldsymbol{z} \in \mathbb{R}^d$, $k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_k$, we have

$$\left| f_3(\boldsymbol{f}_1(\boldsymbol{z}), \boldsymbol{f}_2(\boldsymbol{\omega})) - f_1^{(k)}(\boldsymbol{z}) \right| \le 12B \cdot n_k^{-\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})}.$$
 (23)

Further note that for $k \in [K]$,

$$N_{3,k} \le 2n_k^{t_*^{(k)}/(4\bar{\beta}_*^{((k)}+2t_*^{(k)}))}$$

Thus,

$$s_3 \le C_3(m, M_3) \sum_{k=1}^K n_k^{t_*^{(k)}/(2\bar{\beta}_*^{((k)} + t_*^{(k)}))}.$$

Now, define $\bar{f} \in \mathcal{F}$ by

$$ar{f}(oldsymbol{z},oldsymbol{\omega})\coloneqq f_3ig(oldsymbol{f}_1(oldsymbol{z}),\,oldsymbol{f}_2(oldsymbol{\omega})ig)$$

so that $|\bar{f}(\boldsymbol{z},\boldsymbol{\omega})| \leq B$ for all $\boldsymbol{z} \in \mathbb{R}^d$ and $\boldsymbol{\omega} \in \mathcal{S}$. Let $D \coloneqq \left[-\xi_1 \log(2dn), \xi_1 \log(2dn)\right]^d \setminus \left(\bigcup_{k \in [K]} E^{(k)}\right)$, so that by Lemma 9 and (20), we have

$$\mathbb{P}(\boldsymbol{Z}_{0} \notin D) \leq \mathbb{P}(\boldsymbol{Z}_{0} \notin \left[-\xi_{1} \log(2dn), \xi_{1} \log(2dn)\right]^{d}) + \sum_{k=1}^{K} \mathbb{P}(\boldsymbol{Z}_{0} \in E^{(k)}) \leq \frac{2}{n}.$$
 (24)

By (22) and (23), for all $\boldsymbol{z} \in D, k \in [K]$ and $\boldsymbol{\omega} \in \mathcal{S}_k$, we have

$$\left| \bar{f}(\boldsymbol{z}, \boldsymbol{\omega}) - f^{\star}(\boldsymbol{z}, \boldsymbol{\omega}) \right| \leq \left| f_{3}(\boldsymbol{f}_{1}(\boldsymbol{z}), \, \boldsymbol{f}_{2}(\boldsymbol{\omega})) - f_{1}^{(k)}(\boldsymbol{z}) \right| + \left| f_{1}^{(k)}(\boldsymbol{z}) - f^{\mathcal{S}_{k}}(\boldsymbol{z}) \right| \\ \leq \left\{ C_{2}(\xi_{1}, d, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)}) + 12B \right\} \cdot \log^{\bar{\beta}_{1}^{(k)}}(n) \cdot n_{k}^{-\bar{\beta}_{*}^{(k)}/(2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)})}.$$
(25)

Hence there exists $C_4 > 0$, depending only on ξ_1, d, m and $(\boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)})_{k=1}^K$ such that

$$\inf_{f \in \mathcal{F}} \mathbb{E} \left\{ \left(f(\mathbf{Z}_{0}, \mathbf{\Omega}_{0}) - f^{*}(\mathbf{Z}_{0}, \mathbf{\Omega}_{0}) \right)^{2} \right\} \\
\leq \mathbb{E} \left\{ \left(\bar{f}(\mathbf{Z}_{0}, \mathbf{\Omega}_{0}) - f^{*}(\mathbf{Z}_{0}, \mathbf{\Omega}_{0}) \right)^{2} \mathbb{1}_{\{\mathbf{Z}_{0} \in D\}} \right\} + \frac{8B^{2}}{n} \\
= \sum_{k=1}^{K} \pi_{k} \mathbb{E} \left\{ \left(\bar{f}(\mathbf{Z}_{0}, \mathbf{\Omega}_{0}) - f^{*}(\mathbf{Z}_{0}, \mathbf{\Omega}_{0}) \right)^{2} \mathbb{1}_{\{\mathbf{Z}_{0} \in D\}} \mid \mathbf{\Omega}_{0} \in \mathcal{S}_{k} \right\} + \frac{8B^{2}}{n} \\
\leq \sum_{k=1}^{K} \pi_{k} \sup_{\mathbf{z} \in D, \boldsymbol{\omega} \in \mathcal{S}_{k}} \left\{ \bar{f}(\mathbf{z}, \boldsymbol{\omega}) - f^{*}(\mathbf{z}, \boldsymbol{\omega}) \right\}^{2} + \frac{8B^{2}}{n} \\
\leq C_{4} (\log n)^{2 \max_{k \in [K]} \bar{\beta}_{1}^{(k)}} \cdot \sum_{k=1}^{K} \pi_{k} n_{k}^{-2\bar{\beta}_{*}^{(k)}/(2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)})},$$
(26)

where the first inequality follows from (24) and $|\bar{f}(\boldsymbol{z},\boldsymbol{\omega}) - f^{\star}(\boldsymbol{z},\boldsymbol{\omega})| \leq 2B$ for all $\boldsymbol{z} \in \mathbb{R}^d$ and $\boldsymbol{\omega} \in \mathcal{S}$, and the final inequality follows from (25). Let $L_0 \coloneqq L_3 + (L_1 \vee L_2)$, $p_0 \coloneqq (2d, 2\|\boldsymbol{p}_1\|_{\infty} + 2\|\boldsymbol{p}_2\|_{\infty}, \dots, 2\|\boldsymbol{p}_1\|_{\infty} + 2\|\boldsymbol{p}_2\|_{\infty}, \boldsymbol{p}_3) \in \mathbb{N}^{L_0+2}$ and $s_0 \coloneqq 2(s_1 + s_2 + s_3) + 2(K \vee m)|L_1 - L_2|$. By (P1)–(P4), we have $\mathcal{F} \subseteq \mathcal{F}(L_0, \boldsymbol{p}_0, s_0)$. Moreover,

$$\frac{s_0 L_0 \log(es_0) + s_0 \log(ed)}{n} \le C_5 \log^3 n \cdot \frac{\sum_{k=1}^K n_k^{t_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})} + s_2 \log s_2}{n}$$
$$= C_5 \log^3 n \cdot \left\{ \sum_{k=1}^K \pi_k n_k^{-2\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})} + \frac{s_2 \log s_2}{n} \right\}, \quad (27)$$

where $C_5 > 0$ depends only on m, L_2, M_1, M_3, d and $(\boldsymbol{d}^{(k)}, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)})_{k=1}^K$. The final result then follows by applying Theorem 1 in conjunction with (26) and (27).

G Proof of Theorem 4

Proof of Theorem 4. Without loss of generality, we may assume that $j_* = d$. Let \mathbf{X}_0 be uniformly distributed on $[0,1]^d$, so that $\|\mathbf{X}_{0,j}\|_{\psi_1} \leq 0.8 \leq \xi_1$ for $j \in [d]$. For $k \in [K-1]$, let $A_k \coloneqq \left[\sum_{\ell=1}^{k-1} \pi_\ell, \sum_{\ell=1}^k \pi_\ell\right)$ and let $A_K \coloneqq \left[\sum_{\ell=1}^{K-1} \pi_\ell, 1\right]$. Given $f^{(1)}, \ldots, f^{(K)} \colon [0,1]^d \to [0,1/2]$, we define $g_{(f^{(1)},\ldots,f^{(K)})} \colon [0,1]^d \to [0,1/2]$ by $g_{(f^{(1)},\ldots,f^{(K)})}(\mathbf{x}) \coloneqq \sum_{k=1}^K \mathbbm{1}_{\{x_d \in A_k\}} \cdot f^{(k)}(\mathbf{x})$ where $\mathbf{x} = (x_1,\ldots,x_d)^\top \in [0,1]^d$. Further, let $P_{(f^{(1)},\ldots,f^{(K)})}$ be the distribution of $(\mathbf{Z}_0, \mathbf{\Omega}_0, Y_0)$, where $\mathbf{\Omega}_0 \mid \mathbf{X}_0 \sim \text{Unif}(\mathcal{S}_k)$ if $X_{0,d} \in A_k$, $\mathbf{Z}_0 = \text{Imp}(\mathbf{X}_0 \otimes \mathbf{\Omega}_0)$ and $Y_0 = g_{(f^{(1)},\ldots,f^{(K)})}(\mathbf{X}_0) + \varepsilon_0$, where $\varepsilon_0 \sim N(0, \xi_2^2/25)$ is independent of $(\mathbf{X}_0, \mathbf{\Omega}_0)$. Thus

$$\|Y_0\|_{\psi_2} \le \|g_{(f^{(1)},\dots,f^{(K)})}(\boldsymbol{X}_0)\|_{\psi_2} + \|\varepsilon_0\|_{\psi_2} \le \frac{1}{2\sqrt{\log 2}} + \sqrt{\frac{8}{75}}\xi_2 \le \xi_2$$

Moreover, observe that when $\boldsymbol{\omega} \in \mathcal{S}_k$, we must have $X_{0,d} \in A_k$ and hence $g_{(f^{(1)},\ldots,f^{(K)})}(\boldsymbol{X}_0) = f^{(k)}(\boldsymbol{X}_0)$. Hence, if for each $k \in [K]$, the function $f^{(k)}$ depends only on the coordinates in $\widetilde{\mathcal{J}}^{(k)} \coloneqq \mathcal{J}^{(k)} \setminus \{d\}$, then when $\boldsymbol{\omega} \in \mathcal{S}_k$,

$$egin{aligned} f^{\star}(oldsymbol{z},oldsymbol{\omega}) &= \mathbb{E}ig(oldsymbol{Y}_0 \,|\, oldsymbol{Z}_0 = oldsymbol{z}, oldsymbol{\Omega}_0 = oldsymbol{\omega}ig) = \mathbb{E}ig(f^{(1)}, ..., f^{(K)})(oldsymbol{X}_0) \,|\, oldsymbol{Z}_0 = oldsymbol{z}, oldsymbol{\Omega}_0 = oldsymbol{\omega}ig) = oldsymbol{f}^{(k)}(oldsymbol{z}). \end{aligned}$$

In general then, $f^{\star}(\boldsymbol{z}, \boldsymbol{\omega}) = \sum_{k=1}^{K} f^{(k)}(\boldsymbol{z}) \mathbb{1}_{\{\boldsymbol{\omega} \in S_k\}}$. For $k \in [K]$, let

$$\mathcal{F}^{(k)} \coloneqq \Big\{ f \in \mathcal{H}_{\text{comp}}(q_k, \boldsymbol{d}^{(k)}, \boldsymbol{t}^{(k)}, \boldsymbol{\beta}^{(k)}, \boldsymbol{\gamma}^{(k)}) : f \text{ takes values in } [0, 1/2] \text{ and} \\ \text{depends only on the coordinates in } \widetilde{\mathcal{J}}^{(k)} \Big\}.$$

We have established that if $f^{(k)} \in \mathcal{F}^{(k)}$ for all $k \in [K]$, then $P_{(f^{(1)},\dots,f^{(K)})} \in \mathcal{P}$. Let $\mu_{\mathbb{Z}_0,\Omega_0}$ be the joint distribution of (\mathbb{Z}_0,Ω_0) when $(\mathbb{Z}_0,\Omega_0,Y_0) \sim P_{(f^{(1)},\dots,f^{(K)})} \in \mathcal{P}$, and for $k \in [K]$ and $f^{(k)} \in \mathcal{F}^{(k)}$, let $P_{f^{(k)}} \coloneqq P_{(0,\dots,0,f^{(k)},0,\dots,0)}$. Then

$$\inf_{\hat{f}\in\hat{\mathcal{F}}} \sup_{P\in\mathcal{P}} \mathbb{E}_{P\otimes^{n}} \left\{ R(\hat{f}) - R(f^{\star}) \right\} \geq \inf_{\hat{f}\in\hat{\mathcal{F}}} \sup_{\substack{f^{(1)},\dots,f^{(K)}:\\f^{(\ell)}\in\mathcal{F}^{(\ell)}\forall\ell\in[K]}} \mathbb{E}_{P^{\otimes n}_{(f^{(1)},\dots,f^{(K)})}} \left\{ \|\hat{f} - f^{\star}\|_{L_{2}(\mu_{Z_{0},\Omega_{0}})}^{2} \right\} \\
= \inf_{\hat{f}\in\hat{\mathcal{F}}} \sup_{\substack{f^{(1)},\dots,f^{(K)}:\\f^{(\ell)}\in\mathcal{F}^{(\ell)}\forall\ell\in[K]}} \sum_{k=1}^{K} \mathbb{E}_{P^{\otimes n}_{(f^{(1)},\dots,f^{(K)})}} \int_{[0,1]^{d}\times\mathcal{S}} \left\{ \hat{f}(\boldsymbol{z},\boldsymbol{\omega}) - f^{(k)}(\boldsymbol{z}) \right\}^{2} \mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}} d\mu_{Z_{0},\Omega_{0}}(\boldsymbol{z},\boldsymbol{\omega}) \\
\geq \sum_{k=1}^{K} \inf_{\hat{f}\in\hat{\mathcal{F}}} \sup_{\substack{f^{(1)},\dots,f^{(K)}:\\f^{(\ell)}\in\mathcal{F}^{(\ell)}\forall\ell\in[K]}} \mathbb{E}_{P^{\otimes n}_{(f^{(1)},\dots,f^{(K)})}} \int_{[0,1]^{d}\times\mathcal{S}} \left\{ \hat{f}(\boldsymbol{z},\boldsymbol{\omega}) - f^{(k)}(\boldsymbol{z}) \right\}^{2} \mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}} d\mu_{Z_{0},\Omega_{0}}(\boldsymbol{z},\boldsymbol{\omega}) \\
\geq \sum_{k=1}^{K} \inf_{\hat{f}\in\hat{\mathcal{F}}} \sup_{f^{(k)}\in\mathcal{F}^{(k)}} \mathbb{E}_{P^{\otimes n}_{f^{(k)}}} \int_{[0,1]^{d}\times\mathcal{S}} \left\{ \hat{f}(\boldsymbol{z},\boldsymbol{\omega}) - f^{(k)}(\boldsymbol{z}) \right\}^{2} \mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}} d\mu_{Z_{0},\Omega_{0}}(\boldsymbol{z},\boldsymbol{\omega}). \tag{28}$$

Now fix $k \in [K]$ and take $f^{(\ell)} = 0$ for all $\ell \neq k$. Since any $f^{(k)} \in \mathcal{F}^{(k)}$ depends only on the coordinates in $\widetilde{\mathcal{J}}^{(k)}$, there exists $\widetilde{f}^{(k)} : [0,1]^{|\widetilde{\mathcal{J}}^{(k)}|} \to [0,1/2]$ such that $f^{(k)}(\boldsymbol{x}) = \widetilde{f}^{(k)}(\boldsymbol{x}_{\widetilde{\mathcal{J}}^{(k)}})$ for all $\boldsymbol{x} = (x_1, \ldots, x_d)^{\top} \in [0,1]^d$, where $\boldsymbol{x}_{\widetilde{\mathcal{J}}^{(k)}} \coloneqq (x_j)_{j \in \widetilde{\mathcal{J}}^{(k)}} \in [0,1]^{|\widetilde{\mathcal{J}}^{(k)}|}$. Thus, writing $||f||_{L_2} := \left(\int_{\boldsymbol{x} \in [0,1]^m} f(\boldsymbol{x})^2 \, \mathrm{d}\boldsymbol{x}\right)^{1/2}$ for square-integrable $f : [0,1]^m \to \mathbb{R}$ and $m \in \mathbb{N}$, we deduce that for $f_1^{(k)}, f_2^{(k)} \in \mathcal{F}^{(k)}$,

$$\int_{[0,1]^{d}\times\mathcal{S}} \left\{ f_{1}^{(k)}(\boldsymbol{z}) - f_{2}^{(k)}(\boldsymbol{z}) \right\}^{2} \mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}} d\mu_{\boldsymbol{Z}_{0},\boldsymbol{\Omega}_{0}}(\boldsymbol{z},\boldsymbol{\omega}) = \mathbb{E} \left[\left\{ f_{1}^{(k)}(\boldsymbol{Z}_{0}) - f_{2}^{(k)}(\boldsymbol{Z}_{0}) \right\}^{2} \mathbb{1}_{\{\boldsymbol{\Omega}_{0}\in\mathcal{S}_{k}\}} \right]
= \mathbb{E} \left[\left\{ \tilde{f}_{1}^{(k)}(\boldsymbol{X}_{0,\tilde{\mathcal{J}}^{(k)}}) - \tilde{f}_{2}^{(k)}(\boldsymbol{X}_{0,\tilde{\mathcal{J}}^{(k)}}) \right\}^{2} \mathbb{1}_{\{X_{0,d}\in\mathcal{A}_{k}\}} \right]
= \pi_{k} \mathbb{E} \left[\left\{ \tilde{f}_{1}^{(k)}(\boldsymbol{X}_{0,\tilde{\mathcal{J}}^{(k)}}) - \tilde{f}_{2}^{(k)}(\boldsymbol{X}_{0,\tilde{\mathcal{J}}^{(k)}}) \right\}^{2} \middle| X_{0,d}\in\mathcal{A}_{k} \right]
= \pi_{k} \left\| f_{1}^{(k)} - f_{2}^{(k)} \right\|_{L_{2}}^{2}.$$
(29)

Next, we use a construction similar to Schmidt-Hieber (2020, Theorem 3) to prove the lower bound. Define $R \coloneqq \lfloor \lambda n_k^{1/(2\bar{\beta}_*^{(k)} + t_*^{(k)})} \rfloor \in \mathbb{N}$ where $\lambda \ge 16$ will be chosen later, let $\rho \coloneqq 1/R$ and let $\mathcal{U} \coloneqq \{0, \rho, 2\rho, \dots, (R-1)\rho\}^{t_*^{(k)}}$. Define $h : \mathbb{R} \to [0, 1]$ by $h(x) \coloneqq c_1 e^{-1/\{x(1-x)\}} \mathbb{1}_{\{x \in (0,1)\}}$, where $c_1 > 0$ depends only on $\beta_*^{(k)}$ and is chosen such that $h \in \mathcal{H}_1^{\beta_*^{(k)}}(\mathbb{R}, 1)$. Further, without loss of generality, suppose that $\widetilde{\mathcal{J}}^{(k)} = \{1, \dots, |\widetilde{\mathcal{J}}^{(k)}|\}$. For $\boldsymbol{u} = (u_1, \dots, u_{t_*^{(k)}})^\top \in \mathcal{U}$, define $\psi_{\boldsymbol{u}} : [0, 1]^{t_*^{(k)}} \to [0, 1]$ by

$$\psi_{\boldsymbol{u}}(x_1, \dots, x_{t_*^{(k)}}) \coloneqq \rho^{\beta_*^{(k)}} \prod_{j=1}^{t_*^{(k)}} h\left(\frac{x_j - u_j}{\rho}\right)$$

For $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ with $\|\boldsymbol{\alpha}\|_1 \leq \lceil \beta_*^{(k)} \rceil - 1$, we have $\|\partial^{\boldsymbol{\alpha}} \psi_{\boldsymbol{u}}\|_{\infty} \leq 1$ since $h \in \mathcal{H}_1^{\beta_*^{(k)}}(\mathbb{R}, 1)$. Moreover, for $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ with $\|\boldsymbol{\alpha}\|_1 = \lceil \beta_*^{(k)} \rceil - 1$ and for all $\boldsymbol{x}, \boldsymbol{y} \in [0, 1]^{t_*^{(k)}}$, we have

$$\begin{aligned} \frac{|\partial^{\alpha}\psi_{\boldsymbol{u}}(\boldsymbol{x}) - \partial^{\alpha}\psi_{\boldsymbol{u}}(\boldsymbol{y})|}{\|\boldsymbol{x} - \boldsymbol{y}\|_{2}^{\beta_{*}^{(k)} + 1 - \lceil \beta_{*}^{(k)} \rceil}} &= \rho^{\beta_{*}^{(k)} + 1 - \lceil \beta_{*}^{(k)} \rceil} \frac{\left|\prod_{j=1}^{t_{*}^{(k)}} h^{(\alpha_{j})}\left(\frac{x_{j} - u_{j}}{\rho}\right) - \prod_{j=1}^{t_{*}^{(k)}} h^{(\alpha_{j})}\left(\frac{y_{j} - u_{j}}{\rho}\right)\right|}{\|\boldsymbol{x} - \boldsymbol{y}\|_{2}^{\beta_{*}^{(k)} + 1 - \lceil \beta_{*}^{(k)} \rceil}} \\ &\leq \frac{\rho^{\beta_{*}^{(k)} + 1 - \lceil \beta_{*}^{(k)} \rceil}}{\|\boldsymbol{x} - \boldsymbol{y}\|_{2}^{\beta_{*}^{(k)} + 1 - \lceil \beta_{*}^{(k)} \rceil}} \sum_{\ell=1}^{t_{*}^{(k)}} \left|\prod_{j=\ell}^{t_{*}^{(k)}} h^{(\alpha_{j})}\left(\frac{x_{j} - u_{j}}{\rho}\right)\prod_{j=1}^{\ell-1} h^{(\alpha_{j})}\left(\frac{y_{j} - u_{j}}{\rho}\right)\right| \\ &- \prod_{j=\ell+1}^{t_{*}^{(k)}} h^{(\alpha_{j})}\left(\frac{x_{j} - u_{j}}{\rho}\right)\prod_{j=1}^{\ell} h^{(\alpha_{j})}\left(\frac{y_{j} - u_{j}}{\rho}\right)\right| \\ &\leq \sum_{\ell=1}^{t_{*}^{(k)}} \frac{\left|h^{(\alpha_{\ell})}\left(\frac{x_{\ell} - u_{\ell}}{\rho}\right) - h^{(\alpha_{\ell})}\left(\frac{y_{\ell} - u_{\ell}}{\rho}\right)\right|}{\left|(x_{\ell} - y_{\ell})/\rho|^{\beta_{*}^{(k)} + 1 - \lceil \beta_{*}^{(k)} \rceil}} \leq t_{*}^{(k)}, \end{aligned}$$

where the first inequality follows from the triangle inequality, and the second and third inequalities follow since $h \in \mathcal{H}_{1}^{\beta_{*}^{(k)}}(\mathbb{R}, 1)$. Thus, we have shown that $\psi_{\boldsymbol{u}} \in \mathcal{H}_{t_{*}^{(k)}}^{\beta_{*}^{(k)}}([0, 1]^{t_{*}^{(k)}}, t_{*}^{(k)})$. For $\boldsymbol{v} = (v_{\boldsymbol{u}})_{\boldsymbol{u} \in \mathcal{U}} \in \{0, 1\}^{|\mathcal{U}|}$, define $\phi_{\boldsymbol{v}} : [0, 1]^{t_{*}^{(k)}} \to [0, 1]$ by

$$\phi_{\boldsymbol{v}} \coloneqq \sum_{\boldsymbol{u} \in \mathcal{U}} v_{\boldsymbol{u}} \psi_{\boldsymbol{u}}.$$

Since $\psi_{\boldsymbol{u}}$ and $\psi_{\boldsymbol{u}'}$ have disjoint support for $\boldsymbol{u} \neq \boldsymbol{u}'$, we have $\phi_{\boldsymbol{v}} \in \mathcal{H}_{t_*^{(k)}}^{\beta_*^{(k)}}([0,1]^{t_*^{(k)}}, t_*^{(k)})$. For $r < r_*^{(k)}$, if $d_{r+1}^{(k)} \leq d_r^{(k)}$, then we define $\boldsymbol{g}_r^{(k)} : [0,1]^{d_r^{(k)}} \to [0,1]^{d_{r+1}^{(k)}}$ by $\boldsymbol{g}_r^{(k)}(\boldsymbol{x}) :=$
$$\begin{split} & (x_1, \dots, x_{d_{r+1}^{(k)}})^\top; \text{otherwise, we define } \boldsymbol{g}_r^{(k)} : [0, 1]^{d_r^{(k)}} \to [0, 1]^{d_{r+1}^{(k)}} \text{ by } \boldsymbol{g}_r^{(k)}(\boldsymbol{x}) \coloneqq (\boldsymbol{x}, 0, \dots, 0)^\top. \\ \text{For } r > r_*^{(k)}, \text{ define } \boldsymbol{g}_r^{(k)} : [0, 1]^{d_r^{(k)}} \to [0, 1]^{d_{r+1}^{(k)}} \text{ by } \boldsymbol{g}_r^{(k)}(\boldsymbol{x}) \coloneqq (x_1^{\beta_r^{(k)} \wedge 1}, 0, \dots, 0)^\top. \\ \text{For } r > r_*^{(k)}, \text{ define } \boldsymbol{g}_{r_*^{(k)}, \boldsymbol{v}}^{(k)} : [0, 1]^{d_{r+1}^{(k)}} \to [0, 1]^{d_{r+1}^{(k)}} \text{ by } \boldsymbol{g}_r^{(k)}(\boldsymbol{x}) \coloneqq (x_1^{\beta_r^{(k)} \wedge 1}, 0, \dots, 0)^\top. \\ \text{For } \boldsymbol{v} \in \{0, 1\}^{|\mathcal{U}|}, \text{ define } \boldsymbol{g}_{r_*^{(k)}, \boldsymbol{v}}^{(k)} : [0, 1]^{d_{r+1}^{(k)}} \to [0, 1]^{d_{r+1}^{(k)}} \text{ by } \boldsymbol{g}_{r_*^{(k)}, \boldsymbol{v}}^{(k)}(\boldsymbol{x}) \coloneqq (\phi_{\boldsymbol{v}}(x_1, \dots, x_{t_*^{(k)}}), 0, \dots, 0)^\top. \\ \text{Then each component of } \boldsymbol{g}_r^{(k)} \text{ belongs to } \mathcal{H}_{t_r^{(k)}}^{\beta_r^{(k)}}([0, 1]^{d_r^{(k)}}, \gamma_r^{(k)}) \text{ for } r \in [q_k] \setminus \{r_*^{(k)}\}, \text{ and} \\ \text{ each component of } \boldsymbol{g}_{r_*^{(k)}, \boldsymbol{v}}^{(k)} \text{ belongs to } \mathcal{H}_{t_*^{(k)}}^{\beta_r^{(k)}}([0, 1]^{d_{r*}^{(k)}}, \gamma_*^{(k)}) \text{ for } \boldsymbol{v} \in \{0, 1\}^{|\mathcal{U}|}. \\ \text{ Let } B \coloneqq \prod_{r=r_*^{(k)}+1}^{q_k}(\beta_r^{(k)} \wedge 1) = \overline{\beta}_*^{(k)}/\beta_*^{(k)}. \\ \text{ For } \boldsymbol{v} = (v_{\boldsymbol{u}})_{\boldsymbol{u}\in\mathcal{U}} \in \{0, 1\}^{|\mathcal{U}|}, \text{ define } f_{\boldsymbol{v}}^{(k)} : [0, 1]^d \to [0, 1/2] \\ \text{ by} \end{split}$$

$$f_{\boldsymbol{v}}^{(k)}(\boldsymbol{x}) \coloneqq \frac{1}{2} \boldsymbol{g}_{q_{k}}^{(k)} \circ \cdots \circ \boldsymbol{g}_{r_{*}^{(k)}+1}^{(k)} \circ \boldsymbol{g}_{r_{*}^{(k)},\boldsymbol{v}}^{(k)} \circ \boldsymbol{g}_{r_{*}^{(k)}-1}^{(k)} \circ \cdots \circ \boldsymbol{g}_{1}^{(k)}(\boldsymbol{x}) \\ = \frac{1}{2} \phi_{\boldsymbol{v}}(\boldsymbol{x})^{B} = \frac{1}{2} \sum_{\boldsymbol{u} \in \mathcal{U}} v_{\boldsymbol{u}} \psi_{\boldsymbol{u}}(x_{1}, \dots, x_{t_{*}^{(k)}})^{B},$$

which satisfies $f_{\boldsymbol{v}}^{(k)} \in \mathcal{F}^{(k)}$ by construction. For $\boldsymbol{u} \in \mathcal{U}$, we have $\|\psi_{\boldsymbol{u}}^B\|_{L_2}^2 = \rho^{2\bar{\beta}_*^{(k)} + t_*^{(k)}} \|h^B\|_{L_2}^{2t_*^{(k)}}$. Therefore, for $\boldsymbol{v}, \boldsymbol{v}' \in \{0,1\}^{|\mathcal{U}|}$,

$$\|f_{\boldsymbol{v}}^{(k)} - f_{\boldsymbol{v}'}^{(k)}\|_{L_2}^2 = \frac{1}{4} \|\boldsymbol{v} - \boldsymbol{v}'\|_1 \cdot \rho^{2\bar{\beta}_*^{(k)} + t_*^{(k)}} \|h^B\|_{L_2}^{2t_*^{(k)}}.$$
(30)

By the Gilbert–Varshamov lemma Samworth and Shah (2025+, Exercise 8.9), there exists $\mathcal{V} \subseteq \{0,1\}^{|\mathcal{U}|}$ such that $|\mathcal{V}| \geq e^{|\mathcal{U}|/8}$ and $||\boldsymbol{v} - \boldsymbol{v}'||_1 > |\mathcal{U}|/4$ for all $\boldsymbol{v}, \boldsymbol{v}' \in \mathcal{V}$ with $\boldsymbol{v} \neq \boldsymbol{v}'$. Thus, for $\boldsymbol{v}, \boldsymbol{v}' \in \mathcal{V}$ with $\boldsymbol{v} \neq \boldsymbol{v}'$, we have by (29) and (30) that

$$\int_{[0,1]^{d}\times\mathcal{S}} \left\{ f_{\boldsymbol{v}}^{(k)}(\boldsymbol{z}) - f_{\boldsymbol{v}'}^{(k)}(\boldsymbol{z}) \right\}^{2} \mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}} d\mu_{\boldsymbol{Z}_{0},\boldsymbol{\Omega}_{0}}(\boldsymbol{z},\boldsymbol{\omega}) = \pi_{k} \|f_{\boldsymbol{v}}^{(k)} - f_{\boldsymbol{v}'}^{(k)}\|_{L_{2}}^{2} \\
> \frac{\pi_{k}}{4} \cdot \frac{\rho^{-t_{*}^{(k)}}}{4} \cdot \rho^{2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)}} \|h^{B}\|_{L_{2}}^{2t_{*}^{(k)}} \ge \frac{\|h^{B}\|_{L_{2}}^{2t_{*}^{(k)}}}{16\lambda^{2\bar{\beta}_{*}^{(k)}}} \cdot \pi_{k} n_{k}^{-2\bar{\beta}_{*}^{(k)}/(2\bar{\beta}_{*}^{(k)} + t_{*}^{(k)})}. \quad (31)$$

Moreover, with $\lambda \coloneqq 2 \left(\frac{100 \|h^B\|_{L_2}^{2t_*^{(k)}}}{\xi_2^2} \right)^{1/(2\bar{\beta}_*^{(k)} + t_*^{(k)})} \vee 16$,

$$\operatorname{KL}\left(P_{f_{\boldsymbol{v}}^{(k)}}^{\otimes n}, P_{f_{\boldsymbol{v}'}^{(k)}}^{\otimes n}\right) = n\operatorname{KL}\left(P_{f_{\boldsymbol{v}}^{(k)}}, P_{f_{\boldsymbol{v}'}^{(k)}}\right)
= \frac{25n}{2\xi_{2}^{2}} \cdot \int_{[0,1]^{d} \times \mathcal{S}} \left\{f_{\boldsymbol{v}}^{(k)}(\boldsymbol{z}) - f_{\boldsymbol{v}'}^{(k)}(\boldsymbol{z})\right\}^{2} \mathbb{1}_{\{\boldsymbol{\omega} \in \mathcal{S}_{k}\}} d\mu_{\boldsymbol{Z}_{0},\boldsymbol{\Omega}_{0}}(\boldsymbol{z}, \boldsymbol{\omega})
= \frac{25n\pi_{k}}{2\xi_{2}^{2}} \|f_{\boldsymbol{v}}^{(k)} - f_{\boldsymbol{v}'}^{(k)}\|_{L_{2}}^{2} \leq \frac{25\|h^{B}\|_{L_{2}}^{2t_{*}^{(k)}}}{8\xi_{2}^{2}} \cdot n_{k}\rho^{2\bar{\beta}_{*}^{(k)}}
\leq \frac{1}{32\rho^{t_{*}^{(k)}}} \leq \frac{\log(|\mathcal{V}|)}{4}.$$
(32)

By applying Fano's lemma Samworth and Shah (2025+, Corollary 8.12) in conjunction with (31) and (32), we conclude that there exists $c^{(k)} > 0$, depending only on $(\xi_2, \bar{\beta}_*^{(k)}, \beta_*^{(k)}, t_*^{(k)})$, such that

$$\inf_{\widehat{f}\in\widehat{\mathcal{F}}}\sup_{f^{(k)}\in\mathcal{F}^{(k)}}\mathbb{E}_{P_{f^{(k)}}^{\otimes n}}\int_{[0,1]^{d}\times\mathcal{S}}\left\{\widehat{f}(\boldsymbol{z},\boldsymbol{\omega})-f^{(k)}(\boldsymbol{z})\right\}^{2}\mathbb{1}_{\{\boldsymbol{\omega}\in\mathcal{S}_{k}\}}\,\mathrm{d}\mu_{\boldsymbol{Z}_{0},\boldsymbol{\Omega}_{0}}(\boldsymbol{z},\boldsymbol{\omega})$$

$$\geq c^{(k)} \pi_k n_k^{-2\bar{\beta}_*^{(k)}/(2\bar{\beta}_*^{(k)} + t_*^{(k)})}.$$
 (33)

The final result follows from (28) and (33).

H Auxiliary lemmas

Lemma 12. For $c \ge 4$ and $d \ge 1$, if $m \ge 0$ satisfies $2^m \le (cm)^d$, then $m \le 2d \log_2(cd)$.

Proof. First, in the case d = 1, we have that $x \ge 0$ satisfies $2^x \le cx$ if and only if $f(x) \coloneqq x - \log_2 x - \log_2 c \le 0$. Since $f'(x) \ge 0$ for $x \ge \log_2 e$, any $m_0 \ge \log_2 e$ satisfying $f(m_0) > 0$ is an upper bound of m. But for $m_0 = 2 \log_2 c$,

$$f(m_0) = \log_2 c - \log_2 \log_2 c - 1 \ge 0$$

for all $c \ge 4$. Hence $m \le m_0$ as desired. For general $d \ge 1$, defining y := m/d, we have $2^y \le cdy$, so the result follows from the case d = 1.

Lemma 13. Let Y be a random variable such that $||Y||_{\psi_2} \leq \xi$ for some $\xi > 0$. Let Z be a random variable taking values in a measurable space \mathcal{Z} . Writing $B_n \coloneqq \xi \sqrt{2 \log n}$, we have

$$\mathbb{E}\{|Y - T_{B_n}Y|\} \le \frac{\sqrt{\pi}\xi}{n^2}, \quad \mathbb{E}\{Y^2 - (T_{B_n}Y)^2\} \le \frac{2\xi^2}{n^2}$$

and

$$\mathbb{E}\left\{\left(\mathbb{E}(Y \mid Z)\right)^{2} - \left(\mathbb{E}(T_{B_{n}}Y \mid Z)\right)^{2}\right\} \leq \frac{4\xi^{2}}{n}, \ \mathbb{E}\left\{\mathbb{E}(Y \mid Z)Y - \mathbb{E}(T_{B_{n}}Y \mid Z) \cdot T_{B_{n}}Y\right\} \leq \frac{4\xi^{2}}{n}$$

Proof. By Markov's inequality,

$$\mathbb{P}(|Y| \ge t) \le \mathbb{E}(e^{|Y|^2/\xi^2}) \cdot e^{-t^2/\xi^2} \le 2e^{-t^2/\xi^2},$$

for all $t \ge 0$. For the first inequality,

$$\mathbb{E}\{|Y - T_{B_n}Y|\} = \int_0^\infty \mathbb{P}(|Y - T_{B_n}Y| \ge t) \,\mathrm{d}t$$

$$\leq \int_0^\infty \mathbb{P}(|Y| \ge B_n + t) \,\mathrm{d}t$$

$$\leq 2\int_0^\infty e^{-(B_n^2 + t^2)/\xi^2} \,\mathrm{d}t = \frac{2}{n^2}\int_0^\infty e^{-t^2/\xi^2} \,\mathrm{d}t = \frac{\sqrt{\pi\xi}}{n^2}.$$

For the second inequality,

$$\mathbb{E}\left\{Y^{2} - (T_{B_{n}}Y)^{2}\right\} = \int_{0}^{\infty} \mathbb{P}\left(Y^{2} - (T_{B_{n}}Y)^{2} \ge t\right) dt$$
$$\leq \int_{0}^{\infty} \mathbb{P}\left(Y^{2} \ge B_{n}^{2} + t\right) dt \le 2\int_{0}^{\infty} e^{-(B_{n}^{2} + t)/\xi^{2}} dt = \frac{2\xi^{2}}{n^{2}}.$$

For the third inequality, by the Cauchy–Schwarz inequality,

$$\mathbb{E}\left\{\left(\mathbb{E}(Y \mid Z)\right)^{2} - \left(\mathbb{E}(T_{B_{n}}Y \mid Z)\right)^{2}\right\}$$

$$= \mathbb{E}\left\{\left(\mathbb{E}(Y \mid Z) - \mathbb{E}(T_{B_n}Y \mid Z)\right) \cdot \left(\mathbb{E}(Y \mid Z) + \mathbb{E}(T_{B_n}Y \mid Z)\right)\right\}$$

$$\leq \sqrt{\mathbb{E}\left\{\left(\mathbb{E}(Y - T_{B_n}Y \mid Z)\right)^2\right\} \cdot \mathbb{E}\left\{2\left(\mathbb{E}(Y \mid Z)\right)^2 + 2\left(\mathbb{E}(T_{B_n}Y \mid Z)\right)^2\right\}}.$$
 (34)

Moreover, by the conditional version of Jensen's inequality,

$$\mathbb{E}\left\{\left(\mathbb{E}(Y - T_{B_n}Y \mid Z)\right)^2\right\} \le \mathbb{E}\left\{\mathbb{E}\left((Y - T_{B_n}Y)^2 \mid Z\right)\right\} \le \mathbb{E}\left\{Y^2 - (T_{B_n}Y)^2\right\} \le \frac{2\xi^2}{n^2}, \quad (35)$$

and

$$\mathbb{E}\left\{2\left(\mathbb{E}(Y|Z)\right)^{2} + 2\left(\mathbb{E}(T_{B_{n}}Y|Z)\right)^{2}\right\} \leq \mathbb{E}\left\{2\mathbb{E}(Y^{2}|Z) + 2\mathbb{E}\left((T_{B_{n}}Y)^{2}|Z\right)\right\} \\
\leq 4\mathbb{E}(Y^{2}) = 4\int_{0}^{\infty}\mathbb{P}(Y^{2}\geq t)\,\mathrm{d}t \leq 8\int_{0}^{\infty}e^{-t/\xi^{2}}\,\mathrm{d}t = 8\xi^{2}. \quad (36)$$

The third inequality then follows by combining (34), (35) and (36).

Finally, for the fourth inequality, by Cauchy–Schwarz again,

$$\mathbb{E}\left\{\mathbb{E}(Y \mid Z)Y - \mathbb{E}(T_{B_n}Y \mid Z) \cdot T_{B_n}Y\right\} \\
= \mathbb{E}\left\{\mathbb{E}(Y \mid Z)(Y - T_{B_n}Y)\right\} + \mathbb{E}\left\{\left(\mathbb{E}(Y \mid Z) - \mathbb{E}(T_{B_n}Y \mid Z)\right) \cdot T_{B_n}Y\right\} \\
\leq \sqrt{\mathbb{E}\left\{\left(\mathbb{E}(Y \mid Z)\right)^2\right\} \cdot \mathbb{E}\left\{(Y - T_{B_n}Y)^2\right\}} + B_n\mathbb{E}\left\{\left|\mathbb{E}(Y - T_{B_n}Y \mid Z)\right|\right\}.$$
(37)

Now, by (36),

$$\mathbb{E}\left\{\left(\mathbb{E}(Y \mid Z)\right)^{2}\right\} \leq \mathbb{E}(Y^{2}) \leq 2\xi^{2},\tag{38}$$

and

$$\mathbb{E}\left\{ (Y - T_{B_n} Y)^2 \right\} \le \mathbb{E}\left\{ Y^2 - (T_{B_n} Y)^2 \right\} \le \frac{2\xi^2}{n^2}.$$
(39)

Lastly,

$$\mathbb{E}\left\{\left|\mathbb{E}(Y - T_{B_n}Y \mid Z)\right|\right\} \le \mathbb{E}\left(|Y - T_{B_n}Y|\right) \le \frac{\sqrt{\pi\xi}}{n^2}.$$
(40)

Combining (37), (38), (39) and (40) yields that

$$\mathbb{E}\Big\{\mathbb{E}(Y \mid Z)Y - \mathbb{E}(T_{B_n}Y \mid Z) \cdot T_{B_n}Y\Big\} \le \frac{4\xi^2}{n},$$

as desired.

Lemma 14. Let $\beta, \gamma > 0, \ \beta_0 \coloneqq \lceil \beta \rceil - 1, \ d \in \mathbb{N} \ and \ g \in \mathcal{H}^{\beta}_d([0, 1]^d, \gamma)$. Then

$$\left|g(\boldsymbol{y}) - \sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d: \|\boldsymbol{\alpha}\|_1 \leq \beta_0} \frac{\partial^{\boldsymbol{\alpha}} g(\boldsymbol{x})}{\boldsymbol{\alpha}!} (\boldsymbol{y} - \boldsymbol{x})^{\boldsymbol{\alpha}}\right| \leq \gamma d^{\beta_0} \|\boldsymbol{y} - \boldsymbol{x}\|_2^{\beta_0}$$

for all $\boldsymbol{x}, \boldsymbol{y} \in [0, 1]^d$.

Proof. Let $\boldsymbol{u} \coloneqq \boldsymbol{y} - \boldsymbol{x}$. By Taylor's theorem, there exists $\theta \in (0, 1)$ such that

$$g(oldsymbol{y}) = \sum_{oldsymbol{lpha} \in \mathbb{N}_0^d: \|oldsymbol{lpha}\|_1 \leq eta_0 - 1} rac{\partial^{oldsymbol{lpha}} g(oldsymbol{x})}{oldsymbol{lpha}!} oldsymbol{u}^{oldsymbol{lpha}} + \sum_{oldsymbol{lpha} \in \mathbb{N}_0^d: \|oldsymbol{lpha}\|_1 = eta_0} rac{\partial^{oldsymbol{lpha}} g(oldsymbol{x} + oldsymbol{ heta} oldsymbol{u})}{oldsymbol{lpha}!} oldsymbol{u}^{oldsymbol{lpha}}.$$

Thus,

$$\begin{split} \left| g(\boldsymbol{y}) - \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} \leq \beta_{0}} \frac{\partial^{\boldsymbol{\alpha}} g(\boldsymbol{x})}{\boldsymbol{\alpha}!} \boldsymbol{u}^{\boldsymbol{\alpha}} \right| &\leq \sum_{\boldsymbol{\alpha} \in \mathbb{N}_{0}^{d}: \|\boldsymbol{\alpha}\|_{1} = \beta_{0}} \left| \partial^{\boldsymbol{\alpha}} g(\boldsymbol{x} + \theta \boldsymbol{u}) - \partial^{\boldsymbol{\alpha}} g(\boldsymbol{x}) \right| \boldsymbol{u}^{\boldsymbol{\alpha}} \\ &\leq d^{\beta_{0}} \gamma \|\boldsymbol{u}\|_{2}^{\beta - \beta_{0}} \cdot \|\boldsymbol{u}\|_{2}^{\beta_{0}} \end{split}$$

where the final inequality uses the facts that $|\{\boldsymbol{\alpha} \in \mathbb{N}_0^d : \|\boldsymbol{\alpha}\|_1 = \beta_0\}| = \binom{d+\beta_0-1}{\beta_0} \leq d^{\beta_0},$ that $g \in \mathcal{H}_d^\beta([0,1]^d,\gamma)$ and that $\boldsymbol{u}^{\boldsymbol{\alpha}} = \prod_{j=1}^d u_j^{\alpha_j} \leq \|\boldsymbol{u}\|_{\infty}^{\beta_0} \leq \|\boldsymbol{u}\|_2^{\beta_0}.$