

# FEATURE LEARNING FOR THE HIGH DIMENSIONAL STATIONARY SCHÖDINGER EQUATION WITH DEEP RITZ METHOD

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

This paper investigates feature learning within the framework of the deep Ritz method for solving the stationary Schrödinger equation with Neumann boundary conditions. We first analyze the convergence of Riemannian gradient descent in an agnostic setting, where the hypothesis function is restricted to a single-index model while the PDE solution is arbitrary. We prove that gradient descent reaches an approximate global minimum: after  $T = O(\log(1/\epsilon))$  iterations, the loss is within  $\epsilon$  of a constant multiple of the optimal loss. We then examine the loss landscape when the source term of the PDE itself follows a single-index model, considering hypothesis functions given by either a single-index model or a two-neuron multi-index model. In the single-index case, we show that the minimum Ritz energy is attained at the feature vector aligned with that of the source term. In the two-neuron case, we study the landscape of regularized Ritz losses and characterize how a second feature emerges, given that the first feature is aligned with the source, as the regularization parameter varies. Finally, numerical experiments are presented to validate the feature emergence theory in the two-neuron setting.

## 1 INTRODUCTION

The past decade has witnessed remarkable progress in applying deep learning techniques to scientific computing, and in particular to the numerical solution of partial differential equations (PDEs). Classical numerical methods, such as finite difference, finite element, and spectral methods, often face challenges when addressing high-dimensional problems, irregular domains, or when solutions exhibit complex structures. Neural network-based approaches have emerged as promising alternatives, offering mesh-free approximations, strong expressive power, and flexibility in incorporating physical constraints. Several notable frameworks have been proposed for solving PDEs with neural networks. The deep Ritz method (E & Yu, 2018) formulates the variational problem associated with elliptic PDEs as an energy minimization task and leverages neural networks as trial functions. Physics-informed neural networks (PINNs) (Raissi et al., 2019) instead enforce PDE constraints through the residual of the governing equations, incorporating them into the training loss. The Deep Galerkin Method (Sirignano & Spiliopoulos, 2018) generalizes this perspective by introducing stochastic collocation strategies to enforce weak PDE formulations. More recently, Weak Adversarial Networks (Zang et al., 2020) have explored adversarial training principles to approximate PDE solutions in a weak sense.

Despite these advances, the theoretical foundations of deep learning for PDEs lag behind their empirical success. While substantial progress has been made in understanding approximation power and generalization error, much less is known about **how neural networks learn effective features or representations of PDE solutions**. *Feature learning* — the process of automatically discovering meaningful low-dimensional structure from high-dimensional data, is a key ingredient to the success of neural networks in a broad range of machine learning tasks. Many recent theoretical works have demonstrated that neural networks trained with gradient-based algorithms are able to learn certain high dimensional target functions with low-dimensional structures, e.g. single- or multi-index models, in regression settings. However, to the best of our knowledge, the emergence of feature learning has not been systematically studied in the context of neural network-based methods for PDEs.

In this work, we investigate the feature learning mechanism of simple neural network models applied to high dimensional stationary Schrödinger equation. Specifically, we consider fitting the solution of the Schrödinger equation in the framework of deep Ritz method where the hypothesis function is defined by either a single index model or a two-neuron network. Specifically, we study solutions within the framework of the deep Ritz method, where the hypothesis function is either a single-index model or a two-neuron network. Our overarching goal is to characterize the feature directions that minimize the Ritz loss and to identify which of these directions can be effectively captured by gradient descent.

## 1.1 OUR CONTRIBUTIONS

We highlight the major contributions of the paper as follows:

- We first investigate the convergence guarantees of gradient descent (GD) in the agnostic setting, where the hypothesis function is restricted to a single-index model while the PDE solution is generic. We show that GD achieves an approximate global minimum in the sense that the loss value after  $T = O(\log 1/\epsilon)$  iterations of GD is within  $\epsilon$  of a constant multiplier of the minimum loss.
- Next, we focus on the loss landscape when the source term of the PDE is itself a single-index model. We consider two cases for the hypothesis function: a single-index model or a multi-index model with two neurons. In the first case, we prove that the minimum Ritz energy is attained at the same feature vector as the source term. In the latter case, we analyze the landscape of a regularized Ritz loss, where the regularization is applied to the outer-layer weights, in the high-dimensional regime. We characterize the sets of local and global minimizers of the high-dimensional limit of the Ritz loss as the regularization parameter varies. Our main results demonstrate that, when one feature vector of the hypothesis function aligns with that of the source term, the behavior of the second feature vector depends on the strength of the regularization: the regularized Ritz loss may admit a local or global minimizer that deviates from the feature vector of the source term.
- Finally, we provide numerical experiments to validate the theory we established for the loss landscape of the two-neuron multi-index model.

## 1.2 RELATED WORK

**Analysis of neural networks for PDEs.** Many recent theoretical studies have investigated the approximation power (Marwah et al., 2021; 2023; Grohs et al., 2022; 2023; Grohs & Herrmann, 2022; De Ryck & Mishra, 2024) of neural networks for representing solutions of PDEs. A number of works have also analyzed generalization error estimates (Mishra & Molinaro, 2023; De Ryck & Mishra, 2022; Shin et al., 2023; Lu et al., 2022b) within variational frameworks such as PINNs and the deep Ritz method. These analyses typically assume that PDE solutions lie in Sobolev or Hölder spaces, which leads to convergence rates that suffer from the curse of dimensionality: achieving an  $\epsilon$ -accurate approximation of a  $d$ -dimensional solution requires  $n = O(\epsilon^{-cd})$  network parameters and training samples. In contrast, another line of research (Chen et al., 2021; 2023; Feng & Lu, 2025; Weinan & Wojtowytsch, 2022) has established dimension-free approximation rates for certain high-dimensional PDEs by developing new regularity theory in Barron spaces. More recently, dimension-free generalization error bounds (Lu et al., 2021a; Lu & Lu, 2022) have also been derived for specific classes of elliptic PDEs. Also, there is a series of stochastic differential equation (SDE)-based neural PDE solvers for elliptic equations (Nüsken & Richter, 2021; Han et al., 2020; Nam et al., 2024).

Compared to approximation-theoretic and generalization analyses, the optimization aspect of neural networks for PDEs remains significantly more challenging and less well understood, primarily due to the highly non-convex nature of the associated loss landscapes. Several recent works (Luo & Yang, 2024; Bonfanti et al., 2024; Zhao & Luo, 2025; Xu et al., 2024; Gao et al., 2023; Jiao et al., 2024) have investigated the convergence of gradient-based algorithms for training highly over-parameterized neural networks within the Neural Tangent Kernel (NTK) framework (Jacot et al., 2018). However, NTK-based analyses are intrinsically tied to the lazy training regime and thus fail to capture feature learning. An alternative line of research considers the mean-field regime (Mei et al., 2018; Sirignano & Spiliopoulos, 2020; Rotskoff & Vanden-Eijnden, 2018), where over-parameterized neural networks are studied via their distributional dynamics, described by Wasser-

stein gradient flows. In this setting, Dus & Virginie (2024); Dus & Ehlacher (2025) established convergence results for Wasserstein gradient flows of the Ritz energy with infinite-width two-layer neural networks, applied to problems such as the Poisson equation and the Schrödinger eigenvalue problem. Nonetheless, these mean-field results are limited to the infinite-width setting and do not directly extend to finite-width networks, where feature learning and optimization dynamics remain far less understood.

**Feature learning of neural networks with gradient-based algorithms.** Beyond the NTK regime, a growing body of work seeks to understand how neural networks trained with gradient descent (GD) can recover low-dimensional structures (feature directions) of relatively simple target functions in regression problems. Examples include polynomials Yehudai & Shamir (2019); Damian et al. (2022), single-index models Soltanolkotabi (2017); Dudeja & Hsu (2018); Damian et al. (2023); Bietti et al. (2022), multi-index models Ba et al. (2022); Dandi et al. (2024); Cui et al. (2024); Moniri et al. (2024); Bruna & Hsu (2025), and sparse Boolean functions Abbe et al. (2022; 2023), among others. A central challenge in extending these analyses from regression to PDE problems lies in the intrinsic mismatch between the PDE solutions and the neural network approximators. Although PDE solutions may exhibit low-dimensional structures, they are typically far more complex than simple neural network models such as multi-index functions. This places the learning problem into the so-called *agnostic* setting, where the hypothesis class does not perfectly capture the target function. While a line of research investigated the first-order methods to learn agnostically the single-index Frei et al. (2020); Wu (2022); Awasthi et al. (2023); Wang et al. (2023); Gollakota et al. (2023); Zarifis et al. (2024) and multi-index models Diakonikolas et al. (2024), these results do not directly generalize to PDEs because the loss functions in PDE learning involve complex differential operators. A rigorous understanding of feature learning in neural network-based PDE solvers remains largely open, even for simple architectures.

### 1.3 NOTATION

We use bold uppercase and lowercase letters to denote matrices and column vectors, respectively. For  $\mathbf{x} \in \mathbb{R}^d$ , we denote its  $p$ -norm by  $|\mathbf{x}|_p$ . When  $p = 2$ , we write  $|\mathbf{x}| = |\mathbf{x}|_2$ . Given a function  $f$ , we use  $\|f\|_\infty$  to denote the sup norm of  $f$ . Let  $\Omega = \mathcal{B}^d$  be the unit ball on  $\mathbb{R}^d$  and  $\partial\Omega$  be the boundary of  $\Omega$ . Also,  $\mathcal{S}^{d-1}$  denotes the sphere in  $d$ -dimension,  $\mathcal{S}^{d-1} := \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{x}| = 1\}$ . We denote by  $H^1(\Omega)$  the Sobolev space of square-integrable functions with square-integrable first derivatives. Given a unit vector  $\mathbf{w} \in \mathcal{S}^{d-1}$ ,  $\mathbf{P}_{\mathbf{w}^\perp} := \mathbf{I} - \mathbf{w}\mathbf{w}^\top$  denotes the orthogonal projector onto the hyperplane perpendicular to  $\mathbf{w}$ , i.e., for  $\mathbf{v} \in \mathbb{R}^d$ ,  $\mathbf{P}_{\mathbf{w}^\perp} \mathbf{v} = \mathbf{v}^\perp \mathbf{w}$ .

## 2 SET-UP AND MAIN RESULTS

We consider the following stationary Schrödinger equation

$$-\Delta u + u = f \text{ on } \Omega, \quad \frac{\partial u}{\partial \nu} = 0 \text{ on } \partial\Omega. \quad (1)$$

We assume that  $f$  is known and  $f \in L^2(\Omega)$ . Thanks to the standard well-posedness of PDEs, there exists a unique weak solution  $u^* \in H^1(\Omega)$  to the equation (1) and

$$u^* = \arg \min_{u \in H^1(\Omega)} \mathcal{E}(u) := \arg \min_{u \in H^1(\Omega)} \frac{1}{2} \int_{\Omega} |\nabla u|^2 + u^2 - 2f u dx. \quad (2)$$

It is also useful to note that for any  $u \in H^1(\Omega)$ ,

$$\begin{aligned} \mathcal{E}(u) &= \frac{1}{2} \int_{\Omega} |\nabla u|^2 + u^2 - 2(-\Delta u^* + u^*) u dx \\ &= \frac{1}{2} \int_{\Omega} |\nabla u - \nabla u^*|^2 + |u - u^*|^2 dx - \frac{1}{2} \int_{\Omega} |\nabla u^*|^2 + |u^*|^2 dx \\ &= \frac{1}{2} \int_{\Omega} |\nabla u - \nabla u^*|^2 + |u - u^*|^2 dx + \mathcal{E}(u^*), \end{aligned}$$

where we have used integration by parts in the second equality. As a consequence, minimizing the energy  $\mathcal{E}$  is equivalent to minimizing the shifted loss function  $\mathcal{L}$  defined as follows:

$$\begin{aligned}\mathcal{L}(u) &:= \mathcal{E}(u) - \mathcal{E}(u^*) \\ &= \frac{1}{2} \int_{\Omega} |\nabla u - \nabla u^*|^2 + |u - u^*|^2 dx \\ &= \frac{|\Omega|}{2} \mathbb{E}_{x \sim \mathcal{P}_{\Omega}} [|\nabla u(x) - \nabla u^*(x)|^2 + |u(x) - u^*(x)|^2],\end{aligned}\tag{3}$$

with  $\mathcal{P}_{\Omega}$  denoting the uniform probability distributions on the domain  $\Omega$ .

The deep Ritz method (DRM) E & Yu (2018) seeks an approximate solution of (2) by minimizing the energy  $\mathcal{E}$  (or its empirical version) within a hypothesis function class  $\mathcal{F}$  which is parameterized by neural networks. The resulting optimization problem over the neural network parameters is non-convex in general and it remains a challenging open problem whether standard gradient-based algorithms can find the globally optimal solution. The purpose of this paper is to approach this problem by understanding the feature learning mechanism of gradient descent in the minimization of the DRM loss. Specifically, we perform the analysis in three settings as discussed in what follows.

## 2.1 SINGLE-INDEX HYPOTHESIS IN THE FULLY AGNOSTIC SETTING

We first consider the case where the hypothesis function  $u$  is defined by a single-index model parameterized with a unit vector  $\mathbf{w} \in \mathcal{S}^{d-1}$ . More precisely, we assume that

$$u(\mathbf{x}) \equiv u_{\mathbf{w}}(\mathbf{x}) = \sigma^2(\mathbf{w} \cdot \mathbf{x}), \mathbf{x} \in \Omega\tag{4}$$

where the activation function  $\sigma(\cdot)$  is the ReLU function (i.e.,  $\sigma(\cdot) = \max\{0, \cdot\}$ ). We consider the squared ReLU activation instead of ReLU its own as it yields a differentiable (Lipschitz continuous) loss function (see equation (5)) leading to a well-defined gradient descent algorithm. Also, it has been shown that shallow neural networks with squared ReLU or high-order power of ReLU activation enjoy quantitative approximation rate in Sobolev spaces (Lu et al., 2022a; Mao et al., 2024). Moreover, we make the following a-priori assumption on the right hand side  $f$  and the exact solution  $u^*$ . It is important to note that the ground-truth solution  $u^*$  still belongs to a large function class and does not necessarily admit the same single-index form (4) as  $u$ . This places the problem in the fully *agnostic learning* setting.

**Assumption 2.1.** *The following statements hold for the equation (1).*

1. *There exist a constant  $C_{u^*}$  such that  $\|\nabla u^*\|_{\infty} \leq C_{u^*}$  and  $\|u^*\|_{\infty} \leq C_{u^*}$ .*
2. *There exists a constant  $C_f$  such that  $\|f\|_{\infty} \leq C_f$ .*

Ignoring the constant  $\frac{|\Omega|}{2}$  and noting that  $\sigma(\cdot)\sigma'(\cdot) = \sigma(\cdot)$  in the subgradient sense, one redefines the equivalent loss functions as follows

$$\begin{aligned}\mathcal{L}(\mathbf{w}) &:= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{\Omega}} [|\nabla_{\mathbf{x}} u_{\mathbf{w}}(\mathbf{x}) - \nabla_{\mathbf{x}} u^*(\mathbf{x})|^2 + |u_{\mathbf{w}}(\mathbf{x}) - u^*(\mathbf{x})|^2] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{\Omega}} [|\nabla_{\mathbf{x}} u_{\mathbf{w}}(\mathbf{x}) - \nabla_{\mathbf{x}} u^*(\mathbf{x})|^2 + |\sigma^2(\mathbf{w} \cdot \mathbf{x}) - u^*(\mathbf{x})|^2]\end{aligned}\tag{5}$$

and

$$\begin{aligned}\mathcal{E}(\mathbf{w}) &:= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{\Omega}} [|\nabla_{\mathbf{x}} u_{\mathbf{w}}(\mathbf{x})|^2 + |u_{\mathbf{w}}(\mathbf{x})|^2 - 2f(\mathbf{x})u_{\mathbf{w}}(\mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{\Omega}} [|\nabla_{\mathbf{x}} u_{\mathbf{w}}(\mathbf{x})|^2 + |u_{\mathbf{w}}(\mathbf{x})|^2 - 2f(\mathbf{x})\sigma^2(\mathbf{w} \cdot \mathbf{x})].\end{aligned}$$

Moreover, since  $\mathcal{L}(\mathbf{w})$  and  $\mathcal{E}(\mathbf{w})$  only differs by a constant (see 3), one has that  $\nabla \mathcal{L}(\mathbf{w}) = \nabla \mathcal{E}(\mathbf{w})$ .

### 2.1.1 OPTIMIZATION VIA RIEMANNIAN GRADIENT DESCENT

To optimize the loss  $\mathcal{L}$  defined by 5, we adopt the classic Riemannian gradient descent (GD) method. First, observe that the gradient of the loss function  $\mathcal{L}(\mathbf{w})$  with respect to  $\mathbf{w}$  is given by

$$\begin{aligned}\nabla \mathcal{L}(\mathbf{w}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{\Omega}} [4\sigma'(\mathbf{w} \cdot \mathbf{x})\mathbf{x}\mathbf{w}^{\top} (2\sigma(\mathbf{w} \cdot \mathbf{x})\mathbf{w} - \nabla u^*) + 4(2\sigma(\mathbf{w} \cdot \mathbf{x})\mathbf{w} - \nabla u^*)\sigma(\mathbf{w} \cdot \mathbf{x})] \\ &\quad + \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{\Omega}} [4(\sigma^2(\mathbf{w} \cdot \mathbf{x}) - u^*(\mathbf{x}))\sigma(\mathbf{w} \cdot \mathbf{x})\mathbf{x}].\end{aligned}$$

However, the expression above is intractable since the ground-truth  $u^*$  is unknown. Instead, noting that  $\nabla \mathcal{L}(\mathbf{w}) = \nabla \mathcal{E}(\mathbf{w})$ , we use the tractable gradient  $\nabla \mathcal{E}(\mathbf{w})$  of the Ritz energy  $\mathcal{E}$  given by

$$\nabla \mathcal{E}(\mathbf{w}) = 4\mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [2|\mathbf{w}|^2 \sigma(\mathbf{w} \cdot \mathbf{x}) \mathbf{x} + 2\sigma^2(\mathbf{w} \cdot \mathbf{x}) \mathbf{w} + \sigma^3(\mathbf{w} \cdot \mathbf{x}) \mathbf{x} - f(\mathbf{x}) \sigma(\mathbf{w} \cdot \mathbf{x}) \mathbf{x}].$$

Recall that  $\mathbf{P}_{\mathbf{w}^\perp} := \mathbf{I} - \mathbf{w}\mathbf{w}^\top$ . Then the Riemannian gradients of the loss function  $\mathcal{L}(\mathbf{w})$  and the energy function  $\mathcal{E}(\mathbf{w})$ , denoted by  $g_{\mathcal{L}}(\mathbf{w})$  and  $g_{\mathcal{E}}(\mathbf{w})$  respectively are  $g_{\mathcal{L}}(\mathbf{w}) := \mathbf{P}_{\mathbf{w}^\perp} \nabla \mathcal{L}(\mathbf{w})$  and  $g_{\mathcal{E}}(\mathbf{w}) := \mathbf{P}_{\mathbf{w}^\perp} \nabla \mathcal{E}(\mathbf{w})$ . Specifically, the gradient of the Ritz energy can be evaluated explicitly as

$$g_{\mathcal{E}}(\mathbf{w}) = 4\mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [2|\mathbf{w}|^2 \sigma(\mathbf{w} \cdot \mathbf{x}) \mathbf{P}_{\mathbf{w}^\perp} \mathbf{x} + \sigma^3(\mathbf{w} \cdot \mathbf{x}) \mathbf{P}_{\mathbf{w}^\perp} \mathbf{x} - f(\mathbf{x}) \sigma(\mathbf{w} \cdot \mathbf{x}) \mathbf{P}_{\mathbf{w}^\perp} \mathbf{x}].$$

In practice, the expectation above is approximated by an empirical estimator  $\hat{g}_{\mathcal{E}}$  computed with  $n$  iid uniform samples  $\{\mathbf{x}_i\}_{i=1}^n$  on  $\Omega$  by

$$\hat{g}_{\mathcal{E}}(\mathbf{w}^t) := \frac{1}{n} \sum_{i=1}^n 4 (2|\mathbf{w}^t|^2 \sigma(\mathbf{w}^t \cdot \mathbf{x}_i) \mathbf{P}_{(\mathbf{w}^t)^\perp} \mathbf{x}_i + \sigma^3(\mathbf{w}^t \cdot \mathbf{x}_i) \mathbf{P}_{(\mathbf{w}^t)^\perp} \mathbf{x}_i - f(\mathbf{x}_i) \sigma(\mathbf{w}^t \cdot \mathbf{x}_i) \mathbf{P}_{(\mathbf{w}^t)^\perp} \mathbf{x}_i).$$

We summarize the gradient descent procedure in Algorithm 1. It takes as input an initial guess  $\mathbf{w}^0$  for the desired parameter, accuracy tolerance parameter  $\epsilon$  (see Theorem 2.2), number of iterations  $T$ , step size  $\eta$ , and the sample distribution  $\mathcal{P}_\Omega$ . It outputs the estimated parameter  $\mathbf{w}^T$  obtained by the Riemannian GD.

---

#### Algorithm 1 Riemannian GD

---

- 1: **Input:**  $\mathbf{w}^0, \epsilon, T, \eta$ ; Sample distribution  $\mathcal{P}_\Omega$ .
- 2: Draw  $n = \Theta(d(d+2)/\epsilon)$  samples  $\{\mathbf{x}_i\}_{i=1}^n$  from  $\mathcal{P}_\Omega$ .
- 3: **for**  $t = 0, \dots, T-1$  **do**
- 4:     Compute the empirical estimate of  $g_{\mathcal{E}}(\mathbf{w}^t)$ :

$$\hat{g}_{\mathcal{E}}(\mathbf{w}^t) := \frac{1}{n} \sum_{i=1}^n 4 (2|\mathbf{w}^t|^2 \sigma(\mathbf{w}^t \cdot \mathbf{x}_i) \mathbf{P}_{(\mathbf{w}^t)^\perp} \mathbf{x}_i + \sigma^3(\mathbf{w}^t \cdot \mathbf{x}_i) \mathbf{P}_{(\mathbf{w}^t)^\perp} \mathbf{x}_i - f(\mathbf{x}_i) \sigma(\mathbf{w}^t \cdot \mathbf{x}_i) \mathbf{P}_{(\mathbf{w}^t)^\perp} \mathbf{x}_i).$$

- 5:     Gradient descent and normalize:  $\mathbf{w}^{t+1} = (\mathbf{w}^t - \eta \hat{g}_{\mathcal{E}}(\mathbf{w}^t)) / \|\mathbf{w}^t - \eta \hat{g}_{\mathcal{E}}(\mathbf{w}^t)\|_2$ .
  - 6: **end for**
  - 7: **return**  $\mathbf{w}^T$
- 

We measure the performance of  $\mathbf{w}^T$  by comparing the loss value  $\mathcal{L}(\mathbf{w}^T)$  with the minimum loss value OPT defined by

$$\text{OPT} := \mathcal{L}(\mathbf{w}^*) = \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [|\nabla_{\mathbf{x}} u_{\mathbf{w}^*}(\mathbf{x}) - \nabla_{\mathbf{x}} u^*(\mathbf{x})|^2 + |u_{\mathbf{w}^*}(\mathbf{x}) - u^*(\mathbf{x})|^2],$$

where  $\mathbf{w}^* := \arg \min_{\mathbf{w} \in \mathcal{S}^{d-1}} \mathcal{L}(\mathbf{w})$ . Our first main theorem below shows that Algorithm 1 produces an approximate solution  $u_{\mathbf{w}^T}$  to the problem (5) after  $T = O(\log(1/\epsilon))$  iterations in the sense that the loss value  $\mathcal{L}(\mathbf{w}^T)$  is within  $\epsilon$  of a constant multiplier of the OPT.

**Theorem 2.2.** *Suppose that Assumption 2.1 holds. Consider Algorithm 1 with initial condition  $\mathbf{w}^0$  satisfying  $\angle(\mathbf{w}^0, \mathbf{w}^*) \in [0, \frac{\pi}{2}]$ . Set the sample size  $n = \Theta\left(\frac{d(d+2)}{\epsilon}\right)$  and the step size  $\eta = \frac{d+2}{2048\pi}$ . Then after  $T = O(\log(1/\epsilon))$  iterations, with probability at least  $1 - \delta$ , the output of Algorithm 1,  $\mathbf{w}^T$ , satisfies  $\mathcal{L}(\mathbf{w}^T) < \gamma \cdot \text{OPT} + \frac{128}{d+2}\epsilon$ , where the absolute constant  $0 < \gamma < 2048\pi^2 + 2$ .*

The proof of Theorem 2.2 can be found in Appendix A. We remark that the linear dependence of the step size  $\eta$  on the dimension  $d$  arises from the fact that the Riemannian gradient satisfies  $g_{\mathcal{E}}(\mathbf{w}) = O(1/d)$ . Consequently, scaling the step size with  $d$  ensures that the effective update of magnitude remains balanced and prevents vanishingly small progress in high dimensions.

## 2.2 SINGLE-INDEX HYPOTHESIS

In this subsection, we focus on the setting where  $f$  is given by a single-index model  $f(\mathbf{x}) = \sigma^2(\mathbf{w}^* \cdot \mathbf{x})$  for some  $\mathbf{w}^* \in \mathcal{S}^{d-1}$ . We also consider the hypothesis function  $u$  is defined by a single index model  $u(\mathbf{x}) = \sigma^2(\mathbf{w} \cdot \mathbf{x})$ . Then the problem of minimizing the Ritz energy becomes

$$\min_{\mathbf{w} \in \mathcal{S}^{d-1}} \mathcal{E}(\mathbf{w}) := \frac{|\Omega|}{2} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [2\sigma(\mathbf{w} \cdot \mathbf{x}) \mathbf{w}^2 + |\sigma^2(\mathbf{w} \cdot \mathbf{x})|^2 - 2\sigma^2(\mathbf{w}^* \cdot \mathbf{x}) \sigma^2(\mathbf{w} \cdot \mathbf{x})]. \quad (6)$$

Similar to the previous section, we remark that the exact solution  $u^*$  associated to the source term  $f$  is not a single index model. However, the proposition below shows that the minimum of the Ritz energy function  $\mathcal{E}(\mathbf{w})$ , defined in (6), attains its minimum in the same direction  $\mathbf{w}^*$  as the source term  $f$ . We defer the proof of Proposition 2.3 to Appendix B.

**Proposition 2.3.** *The minimum value of the Ritz energy function  $\mathcal{E}(\mathbf{w})$  is achieved when  $\mathbf{w} = \mathbf{w}^*$ .*

### 2.3 MULTI-INDEX HYPOTHESIS

In this subsection, we assume that  $f$  is given by a single-index model  $f(\mathbf{x}) = \sigma^2(\mathbf{w}^* \cdot \mathbf{x})$ , but consider a more complicated hypothesis model — two-neuron neural network, which is a special case of *multi-index model*. More concretely, we assume that

$$u(\mathbf{x}) = \sum_{i=1}^2 a_i \sigma^2(\mathbf{w}_i \cdot \mathbf{x})$$

with  $\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{S}^{d-1}$  and  $a_1, a_2 \in \mathbb{R}$ . Motivated by Proposition 2.3 for the single-index hypothesis function, we fix the first feature vector  $\mathbf{w}_1 = \mathbf{w}^*$ . In fact, we can also show that if the hypothesis function is  $a_1 \sigma^2(\mathbf{w}_1 \cdot \mathbf{x}) + a_2 \sigma^2(\mathbf{w}_2 \cdot \mathbf{x})$  with  $a_i > 0$ , then one of the features also aligns with  $\mathbf{w}^*$ . For simplicity, we investigate the mechanism by which the second feature  $\mathbf{w}_2$  emerges under this constraint. Under this setting, the variational problem of minimizing the Ritz energy becomes

$$\begin{aligned} \mathcal{L}(\mathbf{w}, \mathbf{a}) &:= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [|\nabla_{\mathbf{x}} u|^2 + |u|^2 - 2fu] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} \left[ \left| 2 \sum_{i=1}^2 a_i \sigma(\mathbf{w}_i \cdot \mathbf{x}) \mathbf{w}_i \right|^2 + \left| \sum_{i=1}^2 a_i \sigma^2(\mathbf{w}_i \cdot \mathbf{x}) \right|^2 - 2\sigma^2(\mathbf{w}^* \cdot \mathbf{x}) \sum_{i=1}^2 a_i \sigma^2(\mathbf{w}_i \cdot \mathbf{x}) \right]. \end{aligned} \quad (7)$$

To write the loss function in a more compact form, we apply the similar techniques as in Cho & Saul (2009) and define the following kernels

$$\begin{aligned} K_1(\mathbf{w}_i, \mathbf{w}_j) &:= 4\mathbf{w}_i^\top \mathbf{w}_j \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [\sigma(\mathbf{w}_i \cdot \mathbf{x}) \sigma(\mathbf{w}_j \cdot \mathbf{x})] = \frac{2}{\pi(d+2)} h_1(\theta_{ij}), \\ K_2(\mathbf{w}_i, \mathbf{w}_j) &:= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_\Omega} [\sigma^2(\mathbf{w}_i \cdot \mathbf{x}) \sigma^2(\mathbf{w}_j \cdot \mathbf{x})] = \frac{1}{2\pi(d+2)(d+4)} h_2(\theta_{ij}), \end{aligned}$$

where  $\theta_{ij} = \angle(\mathbf{w}_i, \mathbf{w}_j)$  is the angle between  $\mathbf{w}_i$  and  $\mathbf{w}_j$  and

$$\begin{aligned} h_1(\theta_{ij}) &:= (\sin \theta_{ij} \cos \theta_{ij} + (\pi - \theta_{ij}) \cos^2 \theta_{ij}), \\ h_2(\theta_{ij}) &:= 3 \sin \theta_{ij} \cos \theta_{ij} + (\pi - \theta_{ij})(1 + 2 \cos^2 \theta_{ij}), \end{aligned}$$

Note that we have also used  $\mathbf{w}_i^\top \mathbf{w}_j = \cos \theta_{ij}$ . Therefore, the loss function (7) can be written as

$$\mathcal{L}(\mathbf{w}, \mathbf{a}) = \mathbf{a}^\top \mathbf{K}_1 \mathbf{a} + \mathbf{a}^\top \mathbf{K}_2 \mathbf{a} - 2\mathbf{a}^\top \mathbf{K}_*, \quad (8)$$

where  $\mathbf{K}_1 \in \mathbb{R}^{2 \times 2}$ ,  $\mathbf{K}_2 \in \mathbb{R}^{2 \times 2}$  and  $\mathbf{K}_* \in \mathbb{R}^{2 \times 1}$  are defined as follows:

$$\mathbf{K}_1 = \begin{bmatrix} K_1(\mathbf{w}_1, \mathbf{w}_1) & K_1(\mathbf{w}_1, \mathbf{w}_2) \\ K_1(\mathbf{w}_2, \mathbf{w}_1) & K_1(\mathbf{w}_2, \mathbf{w}_2) \end{bmatrix}, \mathbf{K}_2 = \begin{bmatrix} K_2(\mathbf{w}_1, \mathbf{w}_1) & K_2(\mathbf{w}_1, \mathbf{w}_2) \\ K_2(\mathbf{w}_2, \mathbf{w}_1) & K_2(\mathbf{w}_2, \mathbf{w}_2) \end{bmatrix}, \mathbf{K}_* = \begin{bmatrix} K_2(\mathbf{w}^*, \mathbf{w}_1) \\ K_2(\mathbf{w}^*, \mathbf{w}_2) \end{bmatrix}.$$

In practice, an  $\ell_2$ -regularization is commonly employed to promote weight decay. Here we consider two variants: one applied to  $\mathbf{a} = (a_1, a_2)$  and another applied solely to  $a_2$ . In both regularization settings, our primary interest is to understand how the other feature  $\mathbf{w}_2$  emerges in the high dimensional regime  $d \rightarrow \infty$  given that the first feature is fixed at  $\mathbf{w}_1 = \mathbf{w}^*$  as the regularization parameter  $\lambda$  varies.

#### 2.3.1 REGULARIZATION APPLIED TO $\mathbf{a}$

We first consider the following regularized loss function:

$$\mathcal{L}_\lambda(\mathbf{w}, \mathbf{a}) = \mathbf{a}^\top \mathbf{K}_1 \mathbf{a} + \mathbf{a}^\top \mathbf{K}_2 \mathbf{a} - 2\mathbf{a}^\top \mathbf{K}_* + \lambda \mathbf{a}^\top \mathbf{a}, \quad (9)$$

where  $\lambda > 0$  is the regularization parameter. To minimize the regularized loss function with respect to the parameters  $\mathbf{w}$  and  $\mathbf{a}$ , noting first that for any fixed  $\mathbf{w}$ , the function  $\mathbf{a} \mapsto \mathcal{L}_\lambda(\mathbf{w}, \mathbf{a})$  is quadratic and the minimizer  $\mathbf{a}^*$  has a closed form solution and the regularized loss function can thus be reduced to a loss function depending only on  $\mathbf{w}$ . The details are presented in the following lemma.

**Lemma 2.4.** For any fixed  $\mathbf{w}$ , the minimizer  $\mathbf{a}^*$  of the regularized loss function (9) has a closed form  $\mathbf{a}^* = (\mathbf{K}_1 + \mathbf{K}_2 + \lambda \mathbf{I}_2)^{-1} \mathbf{K}_*$  and  $\mathcal{L}_\lambda(\mathbf{w}) = -\mathbf{K}_*^\top (\mathbf{K}_1 + \mathbf{K}_2 + \lambda \mathbf{I}_2)^{-1} \mathbf{K}_*$ .

Since by assumption  $\mathbf{w}_1$  is aligned with  $\mathbf{w}^*$ , the loss function  $\mathcal{L}_\lambda(\mathbf{w})$  can be rewritten as a loss of the angle  $\theta := \angle(\mathbf{w}^*, \mathbf{w}_2) \in [0, \pi]$ . More precisely, letting  $c = \frac{2}{d+2}$  and defining  $\lambda = \xi c$  with a rescaled regularization parameter  $\xi$ , it can be shown that the loss  $\mathcal{L}_\lambda(\mathbf{w}) = \mathcal{L}_\xi(\theta)$  where

$$\mathcal{L}_\xi(\theta) = -\frac{c^2}{16\pi^2(d+4)^2} \begin{bmatrix} h_2(0) \\ h_2(\theta) \end{bmatrix}^\top \begin{bmatrix} \frac{c}{\pi} h_1(0) + \frac{c}{4\pi(d+4)} h_2(0) + \xi c & \frac{c}{\pi} h_1(\theta) + \frac{c}{4\pi(d+4)} h_2(\theta) \\ \frac{c}{\pi} h_1(\theta) + \frac{c}{4\pi(d+4)} h_2(\theta) & \frac{c}{\pi} h_1(0) + \frac{c}{4\pi(d+4)} h_2(0) + \xi c \end{bmatrix}^{-1} \begin{bmatrix} h_2(0) \\ h_2(\theta) \end{bmatrix}.$$

Since our focus is on feature learning in the high-dimensional regime and noting that  $h_1(0) = \pi$  and  $h_2(0) = 3\pi$ , we define the following limiting function as  $d \rightarrow +\infty$ :

$$\begin{aligned} \tilde{\mathcal{L}}_\xi(\theta) &:= \lim_{d \rightarrow +\infty} \frac{16(d+4)^2}{c} \mathcal{L}_\xi(\theta) \\ &= -\frac{(9+9\xi) + \frac{1}{\pi^2} [(1+\xi)h_2^2(\theta) - 6h_1(\theta)h_2(\theta)]}{(1+\xi)^2 - \frac{1}{\pi^2} h_1^2(\theta)}. \end{aligned} \quad (10)$$

The theorem below characterizes the set of minimizers of the limiting function  $\tilde{\mathcal{L}}_\xi(\theta)$  as  $\xi$  varies.

**Theorem 2.5.** Consider the minimization of the limiting loss function  $\tilde{\mathcal{L}}_\xi$  defined by (10).

1. When  $\xi \geq \frac{1}{2}$ ,  $\tilde{\mathcal{L}}_\xi(\theta)$  has a unique global minimizer at  $\theta = 0$  for  $\theta$  on  $[0, \frac{5\pi}{6}]$ .
2. When  $\xi \leq \xi_0$  with some  $\xi_0 < 1/2$ , besides the local minimizer  $\theta = 0$ , there exists at least one additional local minimizer of  $\tilde{\mathcal{L}}_\xi(\theta)$  in the interval  $(\frac{\pi}{4}, \frac{\pi}{2})$ .

We defer the proofs of Lemma 2.4 and Theorem 2.5 to Appendix C.

**Remark.** Although we can only rigorously prove the above theorem, our numerical experiments indicate a phase transition between a local minimizer (unique global minimizer) and two local minimizers, one at 0 and the other lying in  $(\pi/4, \pi/2)$ . This phase transition occurs approximately at the value  $\xi_0 = 0.13$ . More details can be found in Figure 1.

### 2.3.2 REGULARIZATION APPLIED TO $a_2$ SOLELY

We now consider the regularized loss function in which only the coefficient  $a_2$  is penalized:

$$\mathcal{L}_\lambda(\mathbf{w}, \mathbf{a}) = \mathbf{a}^\top \mathbf{K}_1 \mathbf{a} + \mathbf{a}^\top \mathbf{K}_2 \mathbf{a} - 2\mathbf{a}^\top \mathbf{K}_* + \lambda a_2^2. \quad (11)$$

To minimize the above regularized loss function with respect to parameters  $\mathbf{w}$  and  $\mathbf{a}$ , we first note that for any fixed  $\mathbf{w}$ , the minimizer  $\mathbf{a}^*$  has a closed form solution and the regularized loss function can be converted to a loss function that is only about  $\mathbf{w}$ . The details are presented in the following lemma.

**Lemma 2.6.** For any fixed  $\mathbf{w}$ , the minimizer  $\mathbf{a}^*$  of the regularized loss function (11) has a closed form  $\mathbf{a}^* = \left( \mathbf{K}_1 + \mathbf{K}_2 + \begin{bmatrix} 0 & 0 \\ 0 & \lambda \end{bmatrix} \right)^{-1} \mathbf{K}_*$  and  $\mathcal{L}_\lambda(\mathbf{w}) = -\mathbf{K}_*^\top \left( \mathbf{K}_1 + \mathbf{K}_2 + \begin{bmatrix} 0 & 0 \\ 0 & \lambda \end{bmatrix} \right)^{-1} \mathbf{K}_*$ .

Given that  $\mathbf{w}_1$  is aligned with  $\mathbf{w}^*$ , the loss function  $\mathcal{L}_\lambda(\mathbf{w})$  is equivalent to the following loss function  $\mathcal{L}_\xi(\theta)$ , upon substituting the definitions of  $\mathbf{K}_1$  and  $\mathbf{K}_2$ , where  $\theta := \angle(\mathbf{w}^*, \mathbf{w}_2)$ :

$$\mathcal{L}_\xi(\theta) = -\frac{c^2}{16\pi^2(d+4)^2} \begin{bmatrix} h_2(0) \\ h_2(\theta) \end{bmatrix}^\top \begin{bmatrix} \frac{c}{\pi} h_1(0) + \frac{c}{4\pi(d+4)} h_2(0) & \frac{c}{\pi} h_1(\theta) + \frac{c}{4\pi(d+4)} h_2(\theta) \\ \frac{c}{\pi} h_1(\theta) + \frac{c}{4\pi(d+4)} h_2(\theta) & \frac{c}{\pi} h_1(0) + \frac{c}{4\pi(d+4)} h_2(0) + \xi c \end{bmatrix}^{-1} \begin{bmatrix} h_2(0) \\ h_2(\theta) \end{bmatrix},$$

where  $c = \frac{2}{d+2}$  and  $\lambda = \xi c$  for a reparameterized regularization parameter  $\xi$ . Similar to the previous section, we consider the following limiting function (as  $d \rightarrow +\infty$ ):

$$\begin{aligned} \bar{\mathcal{L}}_\xi(\theta) &:= \lim_{d \rightarrow +\infty} \frac{16(d+4)^2}{c} \mathcal{L}_\xi(\theta) \\ &= -\frac{(9+9\xi) + \frac{1}{\pi^2} [h_2^2(\theta) - 6h_1(\theta)h_2(\theta)]}{1+\xi - \frac{1}{\pi^2} h_1^2(\theta)}. \end{aligned} \quad (12)$$

We are now ready to present the main result of this subsection and defer the proof of Lemma 2.6 and Theorem 2.7 to Appendix D.

**Theorem 2.7.** Consider the minimization of the limiting loss function  $\bar{\mathcal{L}}_\xi$  defined by (12). For any  $\xi > 0$ , the function  $\bar{\mathcal{L}}_\xi(\theta)$  has a **unique** global minimizer  $\theta^* \in (\frac{\pi}{3}, \frac{\pi}{2})$ .

Comparing Theorem 2.7 with Theorem 2.5, we observe that different forms of regularization on  $\mathbf{a}$  lead to distinct behaviors in feature emergence. In particular, penalizing only  $a_2$  consistently yields the emergence of an additional feature, whereas strong penalization on both outer-layer weights results in feature collapse.

### 3 NUMERICAL EXPERIMENTS

In this section, we provide numerical results to validate our theory established in Section 2.3 for a multi-index model with two neurons.

First, we show the limiting function  $\tilde{\mathcal{L}}_\xi(\theta)$  in Figure 1 that corresponds to the regularized loss function with regularization applied on the whole vector  $\mathbf{a}$  (Subsection 2.3.1). The left plot of Figure 1 shows the 3D surface defined by the function  $(\xi, \theta) \mapsto \tilde{\mathcal{L}}_\xi(\theta)$  with  $\xi \in [0, 0.2]$  and  $\theta \in [0, \pi]$ . The right plot of the same figure shows the function  $\tilde{\mathcal{L}}_\xi(\theta)$  with some specifically chosen values of  $\xi$ . We also mark the global minimizers of  $\tilde{\mathcal{L}}_\xi(\theta)$  for  $\xi = 0.01, 0.08, 0.09$ . From the figure, we can observe the landscape of  $\tilde{\mathcal{L}}_\xi(\theta)$  exhibits the following phase transitions as  $\xi$  varies:

- When  $0 < \xi \leq \xi_0 \approx 0.08$ ,  $\tilde{\mathcal{L}}_\xi(\theta)$  has two local minimizers, one at  $\theta = 0$  and the other is in the interval  $(\frac{\pi}{3}, \frac{\pi}{2})$ . Moreover, the global minimizer is in the interval  $(\frac{\pi}{3}, \frac{\pi}{2})$ .
- When  $\xi_0 \approx 0.08 < \xi \leq \xi_1 \approx 0.13$ ,  $\tilde{\mathcal{L}}_\xi(\theta)$  has two local minimizers, one at  $\theta = 0$  and the other is in the interval  $(\frac{\pi}{4}, \frac{\pi}{2})$ . Moreover,  $\theta = 0$  is the global minimizer.
- When  $\xi > \xi_1 \approx 0.13$ ,  $\tilde{\mathcal{L}}_\xi(\theta)$  has only one local (and hence global) minimizer at  $\theta = 0$ .

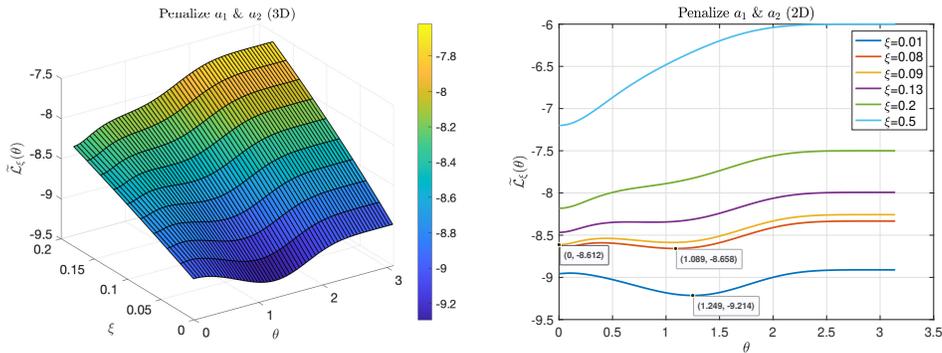


Figure 1: Graphs of the limiting function  $\tilde{\mathcal{L}}_\xi(\theta)$ .

Next, we plot the limiting function  $\bar{\mathcal{L}}_\xi(\theta)$  in Figure 2, which corresponds to the limiting regularized loss function with regularization applied only to  $a_2$  defined by (12). In Figure 2, the left graph shows the 3D plot of the function  $(\xi, \theta) \mapsto \bar{\mathcal{L}}_\xi(\theta)$  with  $\xi \in [0.01, 0.5]$  and  $\theta \in [0, \pi]$ . The right graph shows plots of the function  $\theta \mapsto \bar{\mathcal{L}}_\xi(\theta)$  with  $\xi$  taking a broader range of values. We also mark the global minimizers of  $\bar{\mathcal{L}}_\xi(\theta)$  for  $\xi = 0.01, 1, 5, 100$ . Observe that the loss function  $\bar{\mathcal{L}}_\xi(\theta)$  has a unique global minimizer in the interval  $(\frac{\pi}{3}, \frac{\pi}{2})$  for  $\xi > 0$ .

Moreover, we plot the loss function  $\mathcal{L}_\xi(\theta)$  with  $d = 2$  when the regularization is applied to  $\mathbf{a}$  and  $a_2$  solely in Figure 3. We observe that the loss function has the same shape as the corresponding limiting functions.

We also perform numerical experiments on the landscapes for sigmoid and GELU activation functions with  $d = 2$  and observe that they have different angle landscapes (Figure 4). For sigmoid

432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444

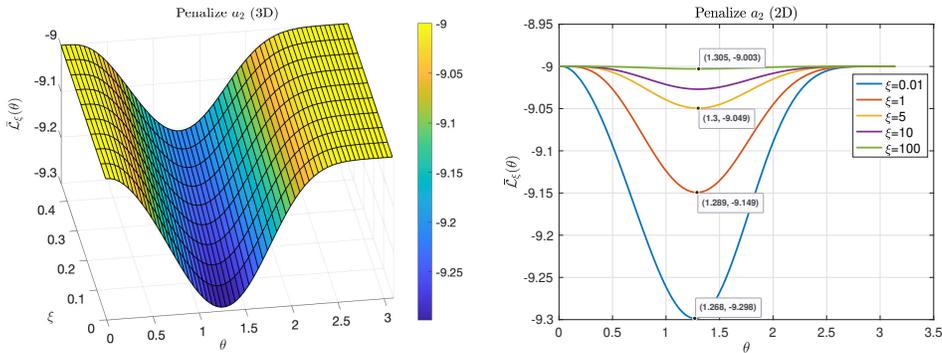


Figure 2: Graphs of the limiting function  $\bar{\mathcal{L}}_{\xi}(\theta)$ .

445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458

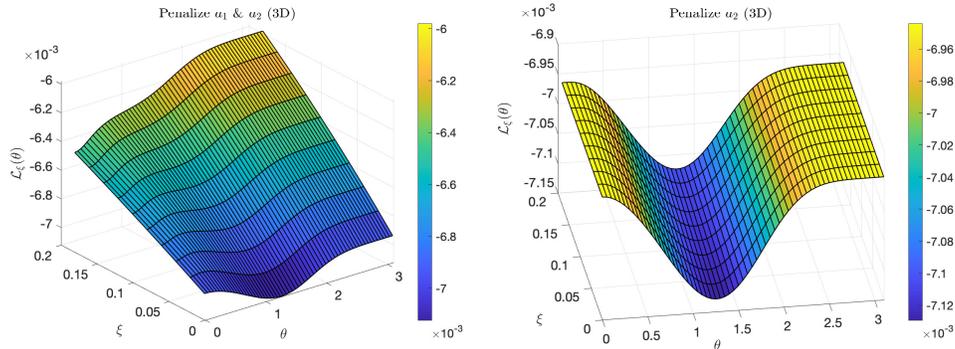


Figure 3: Graphs of the loss function  $\mathcal{L}_{\xi}(\theta)$  with  $d = 2$ .

459  
460  
461  
462  
463  
464  
465  
466

activation, the unique global minimum is achieved at  $\theta = \pi$  when the regularization is applied to  $\mathbf{a}$  and  $a_2$  solely. For GELU, when the regularization is applied to  $\mathbf{a}$ , there are two local minimizers at  $\theta = 0$  and  $\theta = \pi$  and the global minimizer changes from  $\theta = \pi$  to  $\theta = 0$  as  $\lambda$  increases. When the regularization is applied to  $a_2$  solely, there are two local minimizers, one is in  $(1, 1.5)$  and the other is at  $\theta = \pi$  and the global minimizer is at  $\theta = \pi$ .

467

#### 4 CONCLUSION AND DISCUSSION

468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482

In this work, we present a systematic study of feature learning in the context of stationary Schrödinger equation using the deep Ritz method. Unlike prior analyses that rely on infinite-width limits or strong over-parameterization assumptions, our study focuses on the behavior of finite-width models and examines how low-dimensional features of PDE solutions emerge as a result of an optimization procedure. Specifically, in an agnostic setting where the PDE solution is generic, we established convergence guarantees for gradient descent under a single-index hypothesis class, showing that approximate global optimality can be achieved in logarithmic iteration complexity. Further, when the source term of the PDE follows a single-index structure, we characterize the loss landscape for both single- and two-neuron models. Our analysis of the regularized Ritz losses revealed how feature emergence depends critically on the regularization strength: penalizing a single outer-layer weight consistently produces an additional feature, whereas strong joint penalization can result in feature collapse. Collectively, these results provide a first step toward a rigorous theory of feature learning in finite-width neural networks for PDEs.

483  
484  
485

Our work opens several directions for future research. While our analysis focused on single-index and two-neuron models, it remains an open question to investigate multi-index models with more neurons and to determine whether feature emergence exhibits hierarchical or sequential patterns in such settings. Another important direction is to extend the analysis to broader classes of PDEs

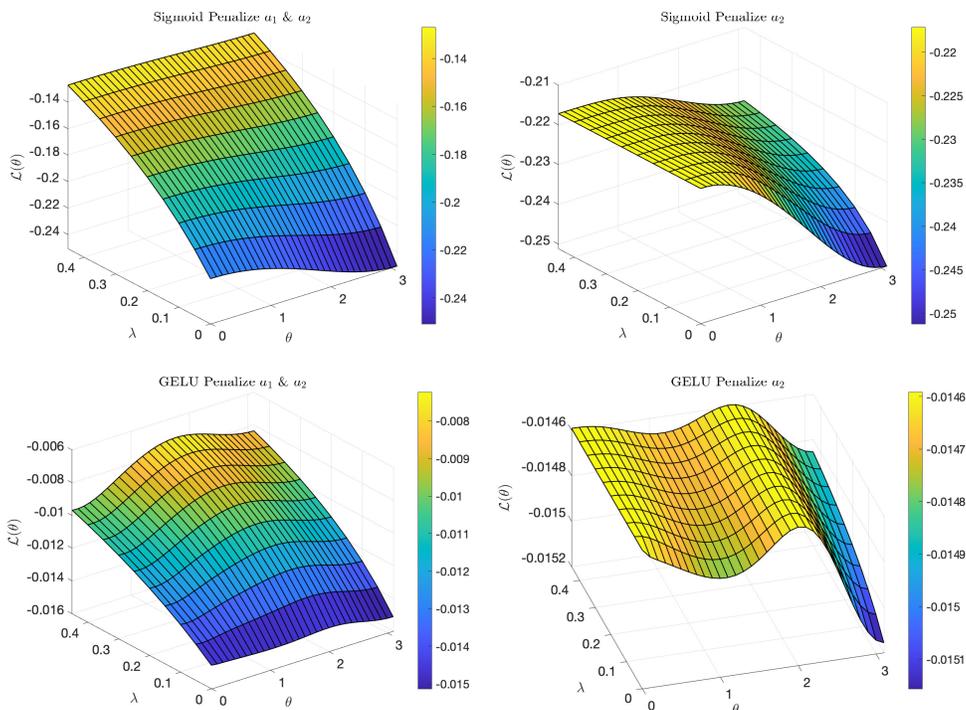


Figure 4: Graphs of the landscapes with  $d = 2$  for sigmoid and GELU activation functions.

and boundary conditions, including higher-order or nonlinear equations, where feature structures may be more intricate. Beyond the deep Ritz method, it would also be valuable to explore whether optimizing losses defined by physics-informed neural networks exhibit analogous mechanisms of feature emergence. We plan to study questions along these directions in future work.

## REFERENCES

- Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, pp. 4782–4887. PMLR, 2022.
- Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. In *The Thirty Sixth Annual Conference on Learning Theory*, pp. 2552–2623. PMLR, 2023.
- Pranjal Awasthi, Alex Tang, and Aravindan Vijayaraghavan. Agnostic learning of general relu activation using gradient descent. In *International Conference on Learning Representations (ICLR)*, 2023. poster.
- Jimmy Ba, Murat A. Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In *Advances in Neural Information Processing Systems*, volume 35, 2022. NeurIPS 2022.
- Alberto Bietti, Joan Bruna, Clayton Sanford, and Min Jae Song. Learning single-index models with shallow neural networks. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS 2022)*, 2022.
- Andrea Bonfanti, Giuseppe Bruno, and Cristina Cipriani. The challenges of the nonlinear regime for physics-informed neural networks. *Advances in Neural Information Processing Systems*, 37: 41852–41881, 2024.

- 540 Joan Bruna and Daniel Hsu. Survey on algorithms for multi-index models. *arXiv preprint*,  
541 2504.05426, 2025.
- 542
- 543 Ziang Chen, Jianfeng Lu, and Yulong Lu. On the representation of solutions to elliptic pdes in  
544 barron spaces. *Advances in neural information processing systems*, 34:6454–6465, 2021.
- 545
- 546 Ziang Chen, Jianfeng Lu, Yulong Lu, and Shengxuan Zhou. A regularity theory for static  
547 schrödinger equations on  $d$  in spectral barron spaces. *SIAM Journal on Mathematical Analysis*,  
548 55(1):557–570, 2023.
- 549 Youngmin Cho and Lawrence K. Saul. Kernel methods for deep learning. In *Advances in Neural*  
550 *Information Processing Systems*, volume 22, pp. 342–350, 2009.
- 551
- 552 Hugo Cui, Luca Pesce, Yatin Dandi, Florent Krzakala, Yue Lu, Lenka Zdeborová, and Bruno  
553 Loureiro. Asymptotics of feature learning in two-layer networks after one gradient-step. In  
554 Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scar-  
555 lett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine*  
556 *Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 9662–9695. PMLR,  
557 July 2024.
- 558 Alex Damian, Eshaan Nichani, Rong Ge, and Jason D. Lee. Smoothing the landscape boosts the  
559 signal for sgd: Optimal sample complexity for learning single index models. In *Advances in*  
560 *Neural Information Processing Systems*, volume 36, 2023. Oral presentation at NeurIPS 2023.
- 561
- 562 Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations  
563 with gradient descent. In *Conference on Learning Theory*, pp. 5413–5452. PMLR, 2022.
- 564
- 565 Yatin Dandi, Florent Krzakala, Bruno Loureiro, Luca Pesce, and Ludovic Stephan. How two-layer  
566 neural networks learn, one (giant) step at a time. *Journal of Machine Learning Research*, 25(349):  
1–65, 2024.
- 567
- 568 Tim De Ryck and Siddhartha Mishra. Error analysis for physics-informed neural networks (pinns)  
569 approximating kolmogorov pdes. *Advances in Computational Mathematics*, 48(6):79, 2022.
- 570
- 571 Tim De Ryck and Siddhartha Mishra. Numerical analysis of physics-informed neural networks and  
572 related models in physics-informed machine learning. *Acta Numerica*, 33:633–713, 2024.
- 573
- 574 Ilias Diakonikolas, Daniel M Kane, Vasilis Kontonis, Christos Tzamos, and Nikos Zarifis. Ag-  
575 nostically learning multi-index models with queries. In *2024 IEEE 65th Annual Symposium on*  
*Foundations of Computer Science (FOCS)*, pp. 1931–1952. IEEE, 2024.
- 576
- 577 Rishabh Dudeja and Daniel Hsu. Learning single-index models in gaussian space. In Sébastien  
578 Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference on*  
579 *Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 1887–1930.  
PMLR, Jul 2018.
- 580
- 581 Mathias Dus and Virginie Ehrlacher. Numerical solution of poisson partial differential equation in  
582 high dimension using two-layer neural networks. *Mathematics of computation*, 94(351):159–208,  
583 2025. ISSN 0025-5718.
- 584
- 585 Mathias Dus and Ehrlacher Virginie. Two-layers neural networks for schrödinger eigenvalue prob-  
586 lems, 2024. URL <https://arxiv.org/abs/2409.01640>.
- 587
- 588 Weinan E and Bing Yu. The deep ritz method: A deep learning-based numerical algorithm for  
589 solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, Mar  
2018. ISSN 2194-671X. doi: 10.1007/s40304-018-0127-z.
- 590
- 591 Ye Feng and Jianfeng Lu. Solution theory of hamilton-jacobi-bellman equations in spectral barron  
592 spaces. *arXiv preprint arXiv:2503.18656*, 2025.
- 593
- 594 Spencer Frei, Yuan Cao, and Quanquan Gu. Agnostic learning of a single neuron with gradient  
descent. *Advances in neural information processing systems*, 33:5417–5428, 2020.

- 594 Yihang Gao, Yiqi Gu, and Michael Ng. Gradient descent finds the global optima of two-layer  
595 physics-informed neural networks. In *International Conference on Machine Learning*, pp. 10676–  
596 10707. PMLR, 2023.
- 597 Aravind Gollakota, Parikshit Gopalan, Adam Klivans, and Konstantinos Stavropoulos. Agnostically  
598 learning single-index models using omnipredictors. In *Advances in Neural Information Process-*  
599 *ing Systems (NeurIPS) 36*, 2023. Poster Presentation.
- 600 Philipp Grohs and Lukas Herrmann. Deep neural network approximation for high-dimensional  
601 elliptic pdes with boundary conditions. *IMA Journal of Numerical Analysis*, 42(3):2055–2082,  
602 2022.
- 603 Philipp Grohs, Arnulf Jentzen, and Diyora Salimova. Deep neural network approximations for so-  
604 lutions of pdes based on monte carlo algorithms. *Partial Differential Equations and Applications*,  
605 3(4):45, 2022.
- 606 Philipp Grohs, Fabian Hornung, Arnulf Jentzen, and Philippe Von Wurstemberger. *A proof that*  
607 *artificial neural networks overcome the curse of dimensionality in the numerical approximation*  
608 *of Black–Scholes partial differential equations*, volume 284. American Mathematical Society,  
609 2023.
- 610 Jihun Han, Mihai Nica, and Adam R. Stinchcombe. A derivative-free method for solving elliptic  
611 partial differential equations with deep neural networks. *Journal of Computational Physics*, 419:  
612 109672:1–109672:18, 2020. doi: 10.1016/j.jcp.2020.109672.
- 613 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and gen-  
614 eralization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- 615 Yuling Jiao, Ruoxuan Li, Peiying Wu, Jerry Zhijian Yang, and Pingwen Zhang. Drm revisited: A  
616 complete error analysis. 2024.
- 617 Jianfeng Lu and Yulong Lu. A priori generalization error analysis of two-layer neural networks for  
618 solving high dimensional schrödinger eigenvalue problems. *Communications of the American*  
619 *Mathematical Society*, 2(1):1–21, 2022.
- 620 Yiping Lu, Haoxuan Chen, Jianfeng Lu, Lexing Ying, and Jose Blanchet. Machine learning for  
621 elliptic pdes: Fast rate generalization bound, neural scaling law and minimax optimality. In *Pro-*  
622 *ceedings of the 10th International Conference on Learning Representations (ICLR 2022)*, 2022a.  
623 arXiv preprint arXiv:2110.06897 (2021), accepted ICLR 2022.
- 624 Yiping Lu, Haoxuan Chen, Jianfeng Lu, Lexing Ying, and Jose Blanchet. Machine learning for  
625 elliptic pdes: Fast rate generalization bound, neural scaling law and minimax optimality. In *10th*  
626 *International Conference on Learning Representations, ICLR 2022*, 2022b.
- 627 Yulong Lu, Jianfeng Lu, and Min Wang. A priori generalization analysis of the deep ritz method for  
628 solving high dimensional elliptic partial differential equations. In *Conference on learning theory*,  
629 pp. 3196–3241. PMLR, 2021a.
- 630 Yulong Lu, Jianfeng Lu, and Min Wang. A priori generalization analysis of the deep ritz method  
631 for solving high dimensional elliptic partial differential equations. In Mikhail Belkin and Samory  
632 Kpotufe (eds.), *Proceedings of the Thirty-Fourth Conference on Learning Theory*, volume 134 of  
633 *Proceedings of Machine Learning Research*, pp. 3196–3241. PMLR, August 15–19 2021b.
- 634 Tao Luo and Haizhao Yang. Two-layer neural networks for partial differential equations: optimiza-  
635 tion and generalization theory. In *Handbook of Numerical Analysis*, volume 25, pp. 515–554,  
636 2024. ISBN 9780443239847.
- 637 Tong Mao, Jonathan W. Siegel, and Jinchao Xu. Approximation rates for shallow  $\text{relu}^k$  neural  
638 networks on sobolev spaces via the radon transform. *arXiv preprint*, 2024.
- 639 Tanya Marwah, Zachary Lipton, and Andrej Risteski. Parametric complexity bounds for approx-  
640 imating pdes with neural networks. *Advances in Neural Information Processing Systems*, 34:  
641 15044–15055, 2021.

- 648 Tanya Marwah, Zachary Chase Lipton, Jianfeng Lu, and Andrej Risteski. Neural network approxi-  
649 mations of pdes beyond linearity: A representational perspective. In *International Conference on*  
650 *Machine Learning*, pp. 24139–24172. PMLR, 2023.
- 651 Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-  
652 layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671,  
653 2018.
- 654 Siddhartha Mishra and Roberto Molinaro. Estimates on the generalization error of physics-informed  
655 neural networks for approximating pdes. *IMA Journal of Numerical Analysis*, 43(1):1–43, 2023.
- 656 Behrad Moniri, Donghwan Lee, Hamed Hassani, and Edgar Dobriban. A theory of non-linear feature  
657 learning with one gradient step in two-layer neural networks. In Ruslan Salakhutdinov, Zico  
658 Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp  
659 (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of  
660 *Proceedings of Machine Learning Research*, pp. 36106–36159. PMLR, July 2024.
- 661 Hong Chul Nam, Julius Berner, and Anima Anandkumar. Solving poisson equations using neural  
662 walk-on-spheres. In *Proceedings of the 41st International Conference on Machine Learning*  
663 *(ICML 2024)*, volume 235 of *Proceedings of Machine Learning Research*, pp. 37277–37292.  
664 PMLR, 2024.
- 665 Nikolaus Nüsken and Lorenz Richter. Interpolating between BSDEs and PINNs: deep learning for  
666 elliptic and parabolic boundary value problems. *arXiv preprint arXiv:2112.03749*, 2021.
- 667 M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learn-  
668 ing framework for solving forward and inverse problems involving nonlinear partial differen-  
669 tial equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. doi:  
670 <https://doi.org/10.1016/j.jcp.2018.10.045>.
- 671 Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence  
672 and asymptotic error scaling of neural networks. *Advances in neural information processing*  
673 *systems*, 31, 2018.
- 674 Yeonjong Shin, Zhongqiang Zhang, and George Em Karniadakis. Error estimates of residual mini-  
675 mization using neural networks for linear pdes. *Journal of Machine Learning for Modeling and*  
676 *Computing*, 4(4), 2023.
- 677 Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial  
678 differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- 679 Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of  
680 large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.
- 681 Mahdi Soltanolkotabi. Learning relus via gradient descent. In *Advances in Neural Information*  
682 *Processing Systems*, volume 30, 2017.
- 683 Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48 of  
684 *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press,  
685 2019. doi: 10.1017/9781108627771.
- 686 Puqian Wang, Nikos Zarifis, Ilias Diakonikolas, and Jelena Diakonikolas. Robustly learning a single  
687 neuron via sharpness. In *Proceedings of the 40th International Conference on Machine Learning*  
688 *(ICML)*, volume 202 of *Proceedings of Machine Learning Research*, pp. 36541–36577. PMLR,  
689 2023.
- 690 E Weinan and Stephan Wojtowytsch. Some observations on high-dimensional partial differential  
691 equations with barron data. In *Mathematical and Scientific Machine Learning*, pp. 253–269.  
692 PMLR, 2022.
- 693 Michael M. Wolf. Mathematical foundations of supervised learning. URL [https://mediatum.](https://mediatum.ub.tum.de/doc/1723378/1723378.pdf)  
694 [ub.tum.de/doc/1723378/1723378.pdf](https://mediatum.ub.tum.de/doc/1723378/1723378.pdf).

702 Lei Wu. Learning a single neuron for non-monotonic activation functions. In *International confer-*  
703 *ence on artificial intelligence and statistics*, pp. 4178–4197. PMLR, 2022.

704

705 Xianliang Xu, Ting Du, Wang Kong, Bin Shan, Ye Li, and Zhongyi Huang. Convergence analysis of  
706 natural gradient descent for over-parameterized physics-informed neural networks. *arXiv preprint*  
707 *arXiv:2408.00573*, 2024.

708 Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding  
709 neural networks. *Advances in neural information processing systems*, 32, 2019.

710

711 Yaohua Zang, Gang Bao, Xiaojing Ye, and Haomin Zhou. Weak adversarial networks for high-  
712 dimensional partial differential equations. *Journal of Computational Physics*, 411:109409, 2020.

713 Nikos Zarifis, Puqian Wang, Ilias Diakonikolas, and Jelena Diakonikolas. Robustly learning single-  
714 index models via alignment sharpness. In *Proceedings of the 41st International Conference on*  
715 *Machine Learning*, volume 235, pp. 58197–58243. PMLR, 2024.

716

717 Wei Zhao and Tao Luo. Convergence guarantees for gradient-based training of neural pde solvers:  
718 From linear to nonlinear pdes. *arXiv preprint arXiv:2505.14002*, 2025.

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755