

OPTIMAL SCALING LAWS IN LEARNING HIERARCHICAL MULTI-INDEX MODELS

Leonardo Defilippis¹, Florent Krzakala², Bruno Loureiro¹, and Antoine Maillard^{1,3}

¹Departement d'Informatique, École Normale Supérieure, PSL & CNRS

²Information, Learning and Physics Laboratory, École Polytechnique Fédérale de Lausanne (EPFL)

³INRIA Paris

ABSTRACT

In this work, we provide a sharp theory of scaling laws for two-layer neural networks trained on a class of *hierarchical multi-index* targets, in a genuinely representation-limited regime. We derive exact information-theoretic scaling laws for subspace recovery and prediction error, revealing how the hierarchical features of the target are sequentially learned through a cascade of phase transitions. We further show that these optimal rates are achieved by a simple, target-agnostic spectral estimator, which can be interpreted as the small learning-rate limit of gradient descent on the first-layer weights. Once an adapted representation is identified, the readout can be learned statistically optimally, using an efficient procedure. As a consequence, we provide a unified and rigorous explanation of scaling laws, plateau phenomena, and spectral structure in shallow neural networks trained on such hierarchical targets.

1 INTRODUCTION

Despite the staggering practical success of neural networks, we still lack a predictive theory answering a deceptively simple question: *given a structured learning problem, how does the network adapt to the task, in what order are the relevant features in the data learned, and how these translate in statistical efficiency?* This question sits at the intersection of three active lines of research. First, the empirical observation of *neural scaling laws* (Kaplan et al., 2020; Brown et al., 2020; Hoffmann et al., 2022) suggest that the performance of large models scale as a power-law in the training resources, yet—with few exceptions—our mathematical understanding of this relationship remains largely confined to linear models or networks in the lazy regime. Second, recent literature of the training dynamics of neural networks increasingly suggest that feature learning is not a smooth process, but are associated to long *plateaus and abrupt transitions* in the risk, with features (or *concepts* in this context) appearing sequentially rather than all at once Saxe et al. (2014); Wei et al. (2022); Schaeffer et al. (2023). Third, empirical analyses of trained networks have uncovered robust regularities in how the learned representations manifest in the trained network weights, such as in their spectral structure, but without a first-principles explanation of *why* particular features are preferred or *when* they should emerge (Martin & Mahoney, 2021; Wang et al., 2023; Thamm et al., 2024).

This paper provides an end-to-end answer to these questions in a mathematically tractable—yet genuinely feature learning—setting of a two-layer neural network trained on *hierarchical multi-index* data, a class of structured supervised learning tasks where the target function depends on a hierarchical combination of functions of the covariates Ren et al. (2025); Ben Arous et al. (2025); Defilippis et al. (2025b). Here, hierarchical denote the fact that the direction are *ordered*, in the sense that their relative importance decay as a power-law with their index number. This leads to a quasi-sparsity of the target representation, with ordered hierarchy of feature strengths as is classical in signal processing Mallat (1999); Donoho (2006). In this task, learning is fundamentally *representation-limited*: the labels $y = g(\mathbf{W}_* \mathbf{x})$ depend on a low-dimensional but unknown subspace $\text{span}(\mathbf{W}_*) \subset \mathbb{R}^d$ of the input space $\mathbf{x} \in \mathbb{R}^d$, and generalization hinges on discovering this subspace as well as possible from the data. Our goal is to turn this qualitative picture into sharp, quantitative predictions. More precisely, our **main contributions** are:

- (i) **Optimal Bayes rates for feature recovery.** We derive the exact information-theoretic limits for recovering the features subspace $\text{span}(\mathbf{W}_*)$ from $n = \Theta(d)$ samples. Further, we precisely characterize the associated optimal mean-squared error rates (a.k.a. *scaling laws*), unveiling the presence of cross-overs and plateaus regions that translate a fundamental trade-off between data-scarce and model-limited regions, depending on the hardness of the underlying hierarchical structure. Interestingly, these rates coincide with the minimax bounds for quasi-sparse recovery achieved by the LASSO Raskutti et al. (2011), and with previously conjectured results for shallow networks with quadratic activation Defilippis et al. (2024). Our analysis shows that these scalings are in fact universal for a broad class of hierarchical multi-index targets, well beyond the quadratic setting.
- (ii) **A matching spectral algorithm and sequential feature emergence.** We introduce a simple, *target-agnostic* spectral estimator that provably achieves the Bayes-optimal rates derived above. The recovery proceeds sequentially: the i -th direction in the hierarchy becomes detectable at a sample complexity $n_i = \Theta(i^{2\gamma} d)$ —where γ is the exponent controlling the hierarchy—leading to a cascade of sharp phase transitions.
- (iii) **Learnability of the target function by neural networks.** Once the features/concepts subspace is identified, we show that learning the second-layer (readout) weights incurs no additional statistical bottleneck: the resulting prediction error matches the Bayes-optimal rate for subspace recovery.

Together, these results show that representation learning proceeds through a sequence of sharp phase transitions as the number of samples increases. New directions in the signal subspace emerge one after another, leading to plateaus and abrupt drops in the prediction error. This sequential emergence (to borrow the terminology of Wei et al. (2022); Schaeffer et al. (2023)) of features provides a precise theoretical underpinning for the empirically observed phenomenon of progressive concept learning in neural networks. Strikingly, the resulting phenomenology closely mirrors that recently uncovered for diagonal and quadratic neural networks using heuristic tools from statistical physics, where progressive feature emergence and distinct scaling regimes were observed (Defilippis et al., 2025b). Our contribution is to place this picture on firm theoretical ground and to substantially generalize it to a broader and more realistic class of models and targets.

2 SETTING

Consider a supervised regression problem with training data $\mathcal{D} = \{(\mathbf{x}_i, y_i) \in \mathbb{R}^{d+1} : i \in [n]\}$, which we assumed were drawn i.i.d. from a joint distribution over \mathbb{R}^{d+1} . Recall that the goal in regression is to learn a target function $f_*(\mathbf{x}) = \mathbb{E}[y|\mathbf{x}]$ from the training data \mathcal{D} . In the following, we will be interested in particular structured tasks where the dependence of the target f_* on the covariates $\mathbf{x} \in \mathbb{R}^d$ is given by a hierarchical combination of low-dimensional subtasks. Mathematically, this is formalized by *hierarchical multi-index models*.

Definition 2.1 (Hierarchical multi-index model) Let $\mathbf{W}_* = (\mathbf{w}_k^* \in \mathbb{R}^d)_{k \in [m_*]}$ denote a family of m_* orthogonal vectors of norm $\|\mathbf{w}_k^*\|^2 = d$. A hierarchical multi-index target is defined as

$$f_*(\mathbf{x}) = \sum_{k=1}^{m_*} a_k^* g_k(\langle \mathbf{w}_k^*, \mathbf{x} \rangle), \quad (1)$$

where $a_1^* > a_2^* > \dots > a_{m_*}^*$ satisfy $\sum_{k=1}^{m_*} (a_k^*)^2 = 1$ and $g_k : \mathbb{R} \rightarrow \mathbb{R}$. Moreover, we say f_* is a scale-free hierarchical multi-index model if $a_k^* = \Theta(k^{-\gamma})$ for some $\gamma > 0$.

A few remarks are in order.

- As the name suggests, hierarchical multi-index models can be seen as a sum of m_* effectively one-dimensional tasks $g_k(z_k)$, with decreasing weight a_k^* .
- More generally, multi-index models are functions of the type $f_*(\mathbf{x}) = g(\mathbf{W}_* \mathbf{x})$, where $g : \mathbb{R}^{m_*} \rightarrow \mathbb{R}$ is known as the *link function* and

$$\mathbf{z} = \mathbf{W}_* \mathbf{x} \in \mathbb{R}^{m_*} \quad (2)$$

are known as the *indices*. They have been widely studied as statistical models for regression in the statistics literature Li (1991); Yuan (2011); Babichev & Bach (2018).

- More recently, multi-index models have gained in popularity in the machine learning theory literature as generative models for data, where they have been used to prove feature learning separation results for neural networks Damian et al. (2022); Abbe et al. (2023); Dandi et al. (2024).

- Scale-free hierarchical multi-index models were considered recently in Ren et al. (2025); Ben Arous et al. (2025); Defilippis et al. (2025b). Note that when $\gamma > 1/2$, the target is *quasi-sparse* in the index basis (Mallat, 1999; Donoho, 2006). This makes the model a suitable framework for investigating feature learning in neural networks, which are known to exhibit an implicit bias towards sparse estimators Gunasekar et al. (2017); Soudry et al. (2018); Andriushchenko et al. (2023).
- In scale-free hierarchical multi-index models with $\gamma < 1/2$, the coefficients must scale as $a_k^* = \Theta(k^{-\gamma} m_*^{\gamma-1/2})$ to guarantee the boundedness of the labels.

Definition 2.2 *The training data $\mathcal{D} = \{(\mathbf{x}_i, y_i) \in \mathbb{R}^{d+1} : i \in [n]\}$ is drawn i.i.d. from a scale-free hierarchical multi-index model $f_*(\mathbf{x})$ as in Def. 2.1, in particular*

$$y_i = f_*(\mathbf{x}_i) + \sqrt{\Delta} \xi_i \quad (3)$$

where $\mathbf{x}_i \sim \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$, $\Delta > 0$ is the noise variance and $\xi_i \sim \mathcal{N}(0, 1)$.

Our results depend on mild assumptions on the link function, stated formally in Assumption 5.1.

In order to quantify the recovery of each individual index, or concept, we define following metric.

Definition 2.3 (Matrix-MSE) *Let $\mathbf{w}_k^* \in \mathbb{S}^{d-1}(\sqrt{d})$ denote one of the target indices $k \in [d]$. We define the matrix-MSE associated to a predictor $\mathbf{w} \in \mathbb{R}^d$ as $\text{mse}_k(\mathbf{w}) := \frac{1}{d^2} \mathbb{E} [\|\mathbf{w} \mathbf{w}^\top - \mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_F^2]$.*

Additionally, we characterize the *weak recovery* transition for the index k , i.e. the minimum number of samples required by an estimator to correlate non-trivially with a specific concept.

Definition 2.4 (k -critical threshold) *Given an estimator $\hat{\mathbf{w}}_k \in \mathbb{R}^d$ of the signal direction \mathbf{w}_k^* (i.e. a measurable function of the training data \mathbf{X}, \mathbf{y}), the k -critical threshold is defined as the minimum sample complexity such that the estimator exhibits a finite overlap with $\mathbf{w}_{*,k}$, namely $\inf \{\alpha > 0 : d^{-1} |\langle \hat{\mathbf{w}}_k, \mathbf{w}_{*,k} \rangle| = \Theta_d(1)\}$.*

The recovery of $\text{span}(\mathbf{W}^*)$ is assessed by the sum of the squared errors for each individual direction, weighted by the contribution of each concept to the target variance.

Definition 2.5 (Weighted MSE) *Let $\mathbf{W} \in \mathbb{R}^{m_* \times d}$. We define the weighted mean-squared error as*

$$\text{MSE}_\gamma(\mathbf{W}) := \sum_{k \in K} (a_k^*)^2 \text{mse}_k(\mathbf{w}_k). \quad (4)$$

For the second goal — showing that neural networks can efficiently learn f_* — we will focus on two-layer neural networks:

$$f(\mathbf{x}; \Theta) := \mathbf{a}^\top \sigma(\mathbf{W} \mathbf{x} + \mathbf{b}), \quad (5)$$

with weights $\Theta := (\mathbf{a} \in \mathbb{R}^p, \mathbf{W} \in \mathbb{R}^{p \times d})$ and a generic activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$. As we shall see, two-layer neural networks can agnostically learn with respect to the model, i.e. without knowledge of the individual tasks g_k nor the details in Def. 2.1. We quantify the generalization capacity of this model through its excess risk

$$R(\Theta) = \mathbb{E} \left[(f_*(\mathbf{x}) - f(\mathbf{x}; \Theta))^2 \right]. \quad (6)$$

3 MAIN RESULTS

3.1 BAYES-OPTIMAL RATES FOR FEATURE RECOVERY

As a first result, we characterize the information-theoretic limits for the recovery of $\text{span}(\mathbf{W}_*)$ in terms of the sample complexity α and the subspace dimension m_* .

Definition 3.1 (Optimal MSE) *Let $\mathcal{D} = (\mathbf{X}, \mathbf{y})$ be drawn as in def. 2.2. Then, the optimal (weighted) mean-square error is achieved by the posterior average $\text{MMSE}_\gamma := \sum_{k=1}^{m_*} \frac{(a_k^*)^2}{d^2} \mathbb{E} [\|\mathbb{E}[\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D}] - \mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_F^2]$.*

By definition, MMSE_γ is a lower-bound for the smallest achievable weighted mean-squared error $\text{MSE}_\gamma(\hat{\mathbf{W}})$ by any estimator $\hat{\mathbf{W}}$ that is a function of the dataset \mathcal{D} . Our first main result quantifies precisely the rates of MMSE_γ as $\alpha \gg 1$, which by construction define the optimal scaling laws for subspace reconstruction in the class of scale-free hierarchical multi-index models. A proof of this result is discussed in Appendix 5.

Theorem 3.2 (Optimal scaling-laws) *In the setting of Definitions 2.1, 2.2, under Assumption 5.1, for $\alpha, m_* \gg 1$, the Bayes-optimal mean-squared error satisfies*

$$\text{MMSE}_\gamma = \Theta_{\alpha, m_*} \begin{cases} \min(\alpha^{-1+\frac{1}{2\gamma}}, \frac{m_*}{\alpha}) & \gamma > 1/2, \\ \min(1, \frac{m_*}{\alpha}) & \gamma < 1/2 \end{cases} \quad (7)$$

Moreover, the k -critical threshold of the Bayes estimator satisfies $\alpha_k^{\text{Bayes}} = \Theta_k(k^{2\gamma} m_*^{\max((1-2\gamma), 0)})$

3.2 OPTIMAL AGNOSTIC SUBSPACE RECOVERY

While Theorem 3.2 provides a fundamental benchmark, an equally important question is whether the optimal reconstruction rates can be *efficiently* achieved by an algorithm which is agnostic of the underlying data distribution.

Definition 3.3 (Spectral estimator) *Given a pre-processing function $\mathcal{T} : \mathbb{R} \rightarrow \mathbb{R}$ on the labels, consider the symmetric random matrix*

$$\mathbf{T} = \sum_{i=1}^n \mathcal{T}(y_i) \mathbf{x}_i \mathbf{x}_i^\top, \quad (8)$$

with spectrum $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$. Define $r := \min\{i \in \mathbb{N} \mid \lambda_{i+1} - \lambda_{i+2} < C/\sqrt{d}\}$, for some arbitrary constant C . We define the spectral estimator $\hat{\mathbf{W}}^{\text{sp}} = (\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_r)^\top \in \mathbb{R}^{r \times d}$ where $\hat{\mathbf{w}}_k$ is the $k \in [r]$ eigenvector of \mathbf{T} corresponding to λ_k , normalized such that $\|\hat{\mathbf{w}}_i\|^2 = d$.

In Appendix 5.2.1 we discuss an interpretation of the spectral estimator in terms of a gradient descent algorithm. Crucially, we consider \mathcal{T} to be data-agnostic, satisfying only the following condition.

Assumption 3.4 *Given $Y \sim \mathcal{N}(\sum_{k=1}^{m_*} a_k^* g_k(z_k), \Delta)$ with $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{m_*})$, assume that $\mathbb{P}(\mathcal{T}(Y) = 0) < 1$ and \mathcal{T} is bounded and there exists $\tau > 0$ such that $\tau = \inf\{c : \mathbb{P}(\mathcal{T}(Y) < c) = 1\}$.*

Our second main theoretical result shows that the data-agnostic spectral method defined above indeed achieves the optimal reconstruction rates of Theorem 3.2.

Theorem 3.5 *In the setting of Definitions 2.1, 2.2, under Assumption 5.1, denote $\hat{\mathbf{W}} = (\hat{\mathbf{W}}^{\text{sp}}, \mathbf{0}_{(m_*-r) \times d}) \in \mathbb{R}^{m_* \times d}$, where $\hat{\mathbf{W}}^{\text{sp}} \in \mathbb{R}^{r \times d}$, $r \leq m_*$, is the spectral estimator defined in 3.3, with \mathcal{T} satisfying Assumption 3.4. Then, for $\alpha, m_* \gg 1$, denoting the spectral estimator's k -critical threshold as α_k^{sp} , we have that $\text{MSE}_\gamma(\hat{\mathbf{W}}) = \Theta_{\alpha, m_*}(\text{MMSE}_\gamma)$, and $\alpha_k^{\text{sp}} = \Theta_k(\alpha_k^{\text{Bayes}})$.*

Figure 2 shows the decay of the weighted MSE for three scale-free hierarchical model variants as a function of the sample complexity α , based on finite-size experiments, using the spectral method introduced in Definition 3.3. Figure 3 illustrates the sequential emergence of new concepts as sample complexity grows, visualized as spikes detaching from the eigenvalue bulk of the matrix \mathbf{T} in eq. 8.

3.3 LEARNING MULTI-INDEX WITH NEURAL NETWORKS

Finally, we turn our attention to the problem of learning hierarchical multi-index targets function with a two-layer neural network and show that the excess risk associated with a suitable training procedure described in Algorithm 1 in Appendix 6, achieves the optimal rates in Theorem 3.2.

Theorem 3.6 *In the setting of Definitions 2.1, 2.2, under Assumption 5.1, with g_k^* , $k \in [m_*]$, polynomials of finite degree. There exist n_0 such that, for $n > n_0$, $p = \omega(n^{1/2})$, $\lambda = \Theta(n^{-1/2})$, the excess risk of eq. equation 6 for a two-layer neural network equation 5, with σ bounded and continuous, trained according to Algorithm 1 satisfies*

$$R_{\text{NN}} = \Theta(\text{MMSE}_\gamma) + O(n^{-1/2}). \quad (9)$$

REFERENCES

- Emmanuel Abbe, Enric Boix Adserà, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 2552–2623. PMLR, 12–15 Jul 2023. URL <https://proceedings.mlr.press/v195/abbe23a.html>.
- Maksym Andriushchenko, Aditya Vardhan Varre, Loucas Pillaud-Vivien, and Nicolas Flammarion. Sgd with large step sizes learns sparse features. In *International Conference on Machine Learning*, pp. 903–925. PMLR, 2023.
- Alexander Atanasov, Jacob A Zavatone-Veth, and Cengiz Pehlevan. Scaling and renormalization in high-dimensional regression. *arXiv preprint arXiv:2405.00592*, 2024.
- Benjamin Aubin, Antoine Maillard, Florent Krzakala, Nicolas Macris, Lenka Zdeborová, et al. The committee machine: Computational to statistical gaps in learning a two-layers neural network. *Advances in Neural Information Processing Systems*, 31, 2018.
- Dmitry Babichev and Francis Bach. Slice inverse regression with score functions. *Electronic Journal of Statistics*, 12(1):1507 – 1543, 2018. doi: 10.1214/18-EJS1428. URL <https://doi.org/10.1214/18-EJS1428>.
- Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural scaling laws. *Proceedings of the National Academy of Sciences*, 121(27):e2311878121, 2024.
- Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Phase transitions, optimal errors and optimality of message-passing in generalized linear models. *CoRR*, abs/1708.03395, 2017. URL <http://arxiv.org/abs/1708.03395>.
- Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019. doi: 10.1073/pnas.1802705116. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1802705116>.
- Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.
- Gérard Ben Arous, Murat A Erdogdu, N Mert Vural, and Denny Wu. Learning quadratic neural networks in high dimensions: SGD dynamics and scaling laws. *arXiv preprint arXiv:2508.03688*, 2025.
- Raphael Berthier, Andrea Montanari, and Phan-Minh Nguyen. State evolution for approximate message passing with non-separable functions. *Information and Inference: A Journal of the IMA*, 9(1):33–79, 2020.
- Tony Bonnaire, Giulio Biroli, and Chiara Cammarota. The role of the time-dependent hessian in high-dimensional optimization. *Journal of Statistical Mechanics: Theory and Experiment*, 2025(8):083401, 2025.
- Blake Bordelon, Abdulkadir Canatar, and Cengiz Pehlevan. Spectrum dependent learning curves in kernel regression and wide neural networks. In *International Conference on Machine Learning*, pp. 1024–1034. PMLR, 2020.
- Blake Bordelon, Alexander Atanasov, and Cengiz Pehlevan. A dynamical model of neural scaling laws. *Proceedings of the 41st International Conference on Machine Learning*, 2024. arXiv preprint arXiv:2402.01092.
- Blake Bordelon, Alexander Atanasov, and Cengiz Pehlevan. How feature learning can improve neural scaling laws. *Journal of Statistical Mechanics: Theory and Experiment*, 2025(8):084002, 2025.

- Guillaume Braun, Bruno Loureiro, Ha Quang Minh, and Masaaki Imaizumi. Fast escape, slow convergence: Learning dynamics of phase retrieval under power-law data. *arXiv preprint arXiv:2511.18661*, 2025.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- Andreas Christmann and Ingo Steinwart. Support vector machines. *Systems and Virtualization Management. Standards and New Technologies*, 2008.
- Hugo Cui, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime. *Advances in Neural Information Processing Systems*, 34:10131–10143, 2021.
- Hugo Cui, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Error scaling laws for kernel classification under source and capacity conditions. *Machine Learning: Science and Technology*, 4(3):035033, 2023.
- Alex Damian, Loucas Pillaud-Vivien, Jason D. Lee, and Joan Bruna. Computational-statistical gaps in gaussian single-index models. *arXiv preprint arXiv:2403.05529*, 2024.
- Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pp. 5413–5452. PMLR, 2022.
- Yatin Dandi, Florent Krzakala, Bruno Loureiro, Luca Pesce, and Ludovic Stephan. How two-layer neural networks learn, one (giant) step at a time. *Journal of Machine Learning Research*, 25(349): 1–65, 2024.
- Leonardo Defilippis, Bruno Loureiro, and Theodor Misiakiewicz. Dimension-free deterministic equivalents and scaling laws for random feature regression. *Advances in Neural Information Processing Systems*, 37:104630–104693, 2024.
- Leonardo Defilippis, Yatin Dandi, Pierre Mergny, Florent Krzakala, and Bruno Loureiro. Optimal spectral transitions in high-dimensional multi-index models, 2025a. URL <https://arxiv.org/abs/2502.02545>.
- Leonardo Defilippis, Yizhou Xu, Julius Girardin, Emanuele Troiani, Vittorio Erba, Lenka Zdeborová, Bruno Loureiro, and Florent Krzakala. Scaling laws and spectra of shallow neural networks in the feature learning regime. *arXiv preprint arXiv:2509.24882*, 2025b.
- David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- David L Donoho, Matan Gavish, and Andrea Montanari. The phase transition of matrix recovery from Gaussian measurements matches the minimax mse of matrix denoising. *Proceedings of the National Academy of Sciences*, 110(21):8405–8410, 2013.
- Rishabh Dudeja, Yue M. Lu, and Subhabrata Sen. Universality of approximate message passing with semirandom matrices. *The Annals of Probability*, 51(5):1616–1683, 2023.
- Vittorio Erba, Emanuele Troiani, Lenka Zdeborová, and Florent Krzakala. The nuclear route: Sharp asymptotics of ERM in overparameterized quadratic networks. *NeurIPS 2025*, 2025. arXiv preprint arXiv:2505.17958.
- Oliver Y. Feng, Ramji Venkataramanan, Cynthia Rush, and Richard J. Samworth. A unifying tutorial on approximate message passing. *Foundations and Trends® in Machine Learning*, 15(4):335–536, 2022. ISSN 1935-8237. doi: 10.1561/22000000092. URL <http://dx.doi.org/10.1561/22000000092>.

- Cédric Gerbelot and Raphaël Berthier. Graph-based approximate message passing iterations. *Information and Inference: A Journal of the IMA*, 12(4):2562–2628, 2023.
- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. *Advances in neural information processing systems*, 30, 2017.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. An empirical analysis of compute-optimal large language model training. *Advances in neural information processing systems*, 35:30016–30030, 2022.
- Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference: A Journal of the IMA*, 2(2):115–144, 2013.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Filip Kovačević, Yihan Zhang, and Marco Mondelli. Spectral estimators for multi-index models: Precise asymptotics and optimal weak recovery. *arXiv preprint arXiv:2502.01583*, 2025.
- Filip Kovačević, Zhang Yihan, and Marco Mondelli. Spectral estimators for multi-index models: Precise asymptotics and optimal weak recovery. In Nika Haghtalab and Ankur Moitra (eds.), *Proceedings of Thirty Eighth Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pp. 3354–3404. PMLR, 30 Jun–04 Jul 2025. URL <https://proceedings.mlr.press/v291/kovacevic25a.html>.
- Frederik Kunstner and Francis Bach. Scaling laws for gradient descent and sign descent for linear bigram models under Zipf’s law. *arXiv preprint arXiv:2505.19227*, 2025.
- Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- Yue M Lu and Gen Li. Phase transitions of spectral initialization for high-dimensional non-convex estimation. *Information and Inference: A Journal of the IMA*, 9(3):507–541, 2020.
- Antoine Maillard, Florent Krzakala, Yue M Lu, and Lenka Zdeborová. Construction of optimal spectral methods in phase retrieval. In *Mathematical and Scientific Machine Learning*, pp. 693–720. PMLR, 2022.
- Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.
- Alexander Maloney, Daniel A Roberts, and James Sully. A solvable model of neural scaling laws. *arXiv preprint arXiv:2210.16859*, 2022.
- Charles H Martin and Michael W Mahoney. Implicit self-regularization in deep neural networks: Evidence from random matrix theory and implications for learning. *Journal of Machine Learning Research*, 22(165):1–73, 2021.
- Marco Mondelli and Andrea Montanari. Fundamental limits of weak recovery with applications to phase retrieval. *Foundations of Computational Mathematics*, 19(3):703–773, Jun 2019. ISSN 1615-3383. doi: 10.1007/s10208-018-9395-y. URL <https://doi.org/10.1007/s10208-018-9395-y>.
- Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- Elliot Paquette, Courtney Paquette, Lechao Xiao, and Jeffrey Pennington. 4+3 phases of compute-optimal neural scaling laws. *Advances in Neural Information Processing Systems*, 37:16459–16537, 2024.

- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis (eds.), *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL https://proceedings.neurips.cc/paper_files/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf.
- Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE transactions on information theory*, 57(10): 6976–6994, 2011.
- Yunwei Ren, Eshaan Nichani, Denny Wu, and Jason D Lee. Emergence and scaling laws in SGD learning of shallow neural networks. *arXiv preprint arXiv:2504.19983*, 2025.
- Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/61b1fb3f59e28c67f3925f3c79be81a1-Paper.pdf.
- A Saxe, J McClelland, and S Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *Proceedings of the International Conference on Learning Representations 2014*. International Conference on Learning Representations 2014, 2014.
- Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. Are emergent abilities of large language models a mirage? *Advances in neural information processing systems*, 36:55565–55581, 2023.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70): 1–57, 2018.
- Stefano Spigler, Mario Geiger, and Matthieu Wyart. Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124001, 2020.
- Matthias Thamm, Max Staats, and Bernd Rosenow. Random matrix theory analysis of neural network weight matrices. In *High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning*, 2024.
- Emanuele Troiani, Yatin Dandi, Leonardo Defilippis, Lenka Zdeborova, Bruno Loureiro, and Florent Krzakala. Fundamental computational limits of weak learnability in high-dimensional multi-index models. In Yingzhen Li, Stephan Mandt, Shipra Agrawal, and Emtiyaz Khan (eds.), *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 2467–2475. PMLR, 03–05 May 2025. URL <https://proceedings.mlr.press/v258/troiani25a.html>.
- Zhichao Wang, Andrew Engel, Anand D Sarwate, Ioana Dumitriu, and Tony Chiang. Spectral evolution and invariance in linear-width neural networks. *Advances in neural information processing systems*, 36:20695–20728, 2023.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- Roman Worschch and Bernd Rosenow. Analyzing neural scaling laws in two-layer networks with power-law data spectra. *arXiv preprint arXiv:2410.09005*, 2024.
- Arie Wortsman and Bruno Loureiro. Kernel ridge regression under power-law data: spectrum and generalization. *arXiv preprint arXiv:2510.04780*, 2025.
- Ming Yuan. On the identifiability of additive index models. *Statistica Sinica*, pp. 1901–1911, 2011.
- Bohan Zhang, Zihao Wang, Hengyu Fu, and Jason D. Lee. Neural networks learn generic multi-index models near information-theoretic limit, 2025. URL <https://arxiv.org/abs/2511.15120>.

Qiuyun Zou and Hongwen Yang. A concise tutorial on approximate message passing. *arXiv preprint arXiv:2201.07487*, 2022.

4 FURTHER RELATED WORKS —

Most theoretical work on neural scaling laws focus on effectively linear models, such as kernel methods and random features, where the generalization behavior can be characterized through spectral properties of fixed representations (Caponnetto & De Vito, 2007; Spigler et al., 2020; Cui et al., 2021; 2023; Defilippis et al., 2024; Atanasov et al., 2024; Bahri et al., 2024; Maloney et al., 2022; Paquette et al., 2024; Kunstner & Bach, 2025; Bordelon et al., 2020). Exceptions have focused either in the joint-training of both layers for linear networks Bordelon et al. (2024; 2025); Worschech & Rosenow (2024) or non-linear settings with fixed-features Wortsman & Loureiro (2025); Worschech & Rosenow (2024); Braun et al. (2025). A central difference is that these works introduce scaling through the covariate distribution, while in our work it is induced by the hierarchical nature of the task. Moreover, our work departs from this literature by addressing scaling laws in a genuinely non-linear, feature-learning regime.

Closer to us are a recent line of work analyzing two-layer networks with structured first-layer weights and hierarchical multi-index target (Ren et al., 2025; Ben Arous et al., 2025; Defilippis et al., 2025b). In particular, Ren et al. (2025); Ben Arous et al. (2025) studied one-pass SGD in this setting, with Ben Arous et al. (2025) focusing on the case of a quadratic neural network architecture. The main difference with this work is that we analyze full-batch ERM, characterizing the optimal scaling laws and showing that they can be achieved computationally by a gradient-descent like algorithm, thus providing a fundamental benchmark for the SGD rates in these works.

Complementary, Defilippis et al. (2025b) derived rates for ERM in quadratic neural networks trained on quadratic targets. Our results show that the rates from Defilippis et al. (2025b) are universal a large class of target functions (any generative exponent two functions in the language of Damian et al. (2024)), closing a gap between the purely quadratic setting and general hierarchical multi-index targets. The key underlying these universal scaling laws is the combination *quasi-sparsity* of the target representation, encoded by a heavy-tailed spectrum that induces an ordered hierarchy of feature strengths (Mallat, 1999; Donoho, 2006), and an implicit *rank-sparsity* bias, leading to LASSO-like behavior (Raskutti et al., 2011).

On a technical level, our work builds on the toolbox of approximate message passing (AMP) and its associated state evolution (Bayati & Montanari, 2011; Donoho et al., 2013; Javanmard & Montanari, 2013; Berthier et al., 2020; Zou & Yang, 2022; Feng et al., 2022; Gerbelot & Berthier, 2023; Dudeja et al., 2023; Erba et al., 2025). In particular, our results leverage the connection between AMP and Bayes-optimal estimation for single- Barbier et al. (2019) and multi-index Aubin et al. (2018); Troiani et al. (2025) models. Similarly, we build up on the literature on optimal spectral methods derived from AMP for single- (Mondelli & Montanari, 2019; Lu & Li, 2020; Maillard et al., 2022; Damian et al., 2024) and multi-index Kovačević et al. (2025); Defilippis et al. (2025a) functions.

5 PROOF SKETCHES OF THEOREMS 3.2 AND 3.5

In this section, we prove the information-theoretic results stated in Theorem 3.2. In particular, we derive matching bounds for the MMSE $_{\gamma}$, defined 3.1, and the k -critical threshold α_k (Definition 2.4). Our results depend on the following additional assumption of the link functions.

Assumption 5.1 *For each index $k \in [m_{\star}]$, g_k is an even function. Moreover, for $z \sim \mathcal{N}(0, 1)$, for some constants $C, D > 0$, independent of $k, n, d, \alpha, m_{\star}$,*

$$D < \mathbb{E}_z[g_k^2(z)] < C, \quad D < |\mathbb{E}_z[g_k''(z)]| < C. \quad (10)$$

Remark 5.2 *The parity assumption on g_k ensures that learning $\text{span}(\mathbf{W}_{\star})$ is non-trivial, by ruling out linear correlations that would allow recovery at arbitrarily small sample complexity. Indeed, given \mathbf{z} as in equation 2 and y as in eq. 3, when $\mathbb{E}[\mathbf{z} | y] = 0$ a.s.—which holds here by parity—no efficient algorithm can even weakly recover $\text{span}(\mathbf{W}_{\star})$ below the critical threshold*

$$\alpha_c := \left(\sup_{\mathbf{M} \in \mathbb{S}_{m_{\star}}^+, \|\mathbf{M}\|_F = 1} \|\mathbb{E}_y \mathbf{G}(y) \mathbf{M} \mathbf{G}(y)\|_F \right)^{-1}, \quad (11)$$

with $\mathbf{G}(y) := \mathbb{E}[\mathbf{z}\mathbf{z}^{\top} - \mathbf{I}_{m_{\star}} | Y = y]$ Barbier et al. (2019); Mondelli & Montanari (2019); Lu & Li (2020); Damian et al. (2024); Troiani et al. (2025).

The condition $\mathbb{E}_z[g_k^2(z)] < C$ controls the label variance, while the lower bound on $\mathbb{E}_z[g_k''(z)]$ ensures detectability in the proportional regime $n = \Theta(d)$ (equivalently, a generative exponent equal to 2), so that the relative difficulty of each index is governed solely by the decay of a_k^* .¹

Hierarchical multi-index models with generative exponent equal to 2 includes the vast majority of cases of interest, while the class of models with exponent larger than 2 mostly includes fine-tuned examples Barbier et al. (2017); Damian et al. (2024).

In Appendix 5.1, we analyze an oracle estimator that exploits additional information to achieve a weighted mean-squared error smaller than MMSE_γ . Finally, in Appendix 5.2, we complete the proof by characterizing the scaling laws and k -critical thresholds of the spectral estimator defined in 3.3, which simultaneously establishes the results in Theorem 3.5.

5.1 LOWER BOUND

Definition 5.3 (Oracle Estimator) Consider a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i \in [n]}$, where the labels are generated by a multi-index model $y \sim \mathcal{P}(\cdot | \langle \mathbf{w}_1^*, \mathbf{x} \rangle, \dots, \langle \mathbf{w}_{m_*}^*, \mathbf{x} \rangle)$, with weights \mathbf{W}^* drawn from a distribution $\mathcal{P}_{\mathbf{W}}$. We define the Oracle Estimator as the matrix $\mathbf{W}^{\text{oracle}} \in \mathbb{R}^{m_* \times d}$ with rows

$$\mathbf{w}_k^{\text{oracle}} = \arg \min_{\mathbf{w}} \mathbb{E} [\|\mathbf{w}_k \mathbf{w}_k^\top - \mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_F^2 | \mathcal{D}, \{\mathbf{w}_h\}_{h \neq k}], k \in [m_*]. \quad (12)$$

Analogously to the Bayes-optimal estimator case, the weighed mean-squared error $\text{MSE}_\gamma(\mathbf{W}^{\text{oracle}})$ is lower bounded by

$$\text{MMSE}_\gamma^{\text{oracle}} = \frac{1}{d^2} \sum_{k=1}^{m_*} (a_k^*)^2 \mathbb{E} [\|\mathbb{E}[\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D}, \{\mathbf{w}_h^*\}_{h \neq k}] - \mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_F^2]. \quad (13)$$

As a consequence of the law of total variance

$$d^2 \text{mmse}_k := \mathbb{E} [\text{Cov}(\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D})] \quad (14)$$

$$= \mathbb{E} [\text{Cov}(\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D}, \{\mathbf{w}_h\}_{h \neq k})] + \underbrace{\mathbb{E} [\text{Cov}_{\{\mathbf{w}_h\}_{h \neq k}}(\mathbb{E}[\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D}, \{\mathbf{w}_h\}_{h \neq k}])]}_{\geq 0} \quad (15)$$

$$\geq d^2 \mathbb{E} [\|\mathbb{E}[\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D}, \{\mathbf{w}_h^*\}_{h \neq k}] - \mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_F^2]. \quad (16)$$

Therefore,

$$\text{MMSE}_\gamma^{\text{oracle}} \leq \text{MMSE}_\gamma. \quad (17)$$

We now focus on the specific setting of our interest defined in Section 2. For each index $k \in [m_*]$, conditioning on $\{\mathbf{w}_h\}_{h \neq k}$, the problem of the optimal estimation of \mathbf{w}_k^* becomes statistically equivalent to the Bayes-optimal estimation given a dataset $\mathcal{D}_k := \{(\mathbf{x}_\nu, \bar{y}_\nu)\}_{\nu \in [n]}$, where the labels are generated by the single-index model

$$\bar{y}_i = a_k^* g_k(\langle \mathbf{w}_k^*, \mathbf{x}_i \rangle) + \sqrt{\Delta} \xi_i, \quad (18)$$

where ξ_i is the same label noise as in the original dataset. We can now characterize the oracle estimator using the results from the literature on *single-index models*, in particular Barbier et al. (2019). In Appendix 7 we show that the information-theoretic weak recovery threshold for single-index models is a monotonic decreasing function of the signal-to-noise ratio. Combined with Assumption 5.1, this implies that the sequence of k -critical thresholds α_k is bounded by strictly increasing functions.² In particular, the first result in Corollary 7.4 implies that the sequence is bounded by

$$\alpha_k^{\text{oracle}} = \Theta((a_k^*)^{-2}) = \Theta(k^{2\gamma} m_*^{(1-2\gamma)_+}), \quad (19)$$

¹Our analysis can be easily adapted to the less restrictive assumption that there exists an integer $\beta \geq 1$ such that $D < \|\mathbb{E}_Z[g_k^\beta(Z)(Z^2 - 1)]\| < C$. This would result in a simple modification of the scaling laws and not affect the overall message of the present manuscript.

²If $g_k = g$, $\forall k$, the sequence α_k^{oracle} itself is strictly increasing.




Figure 1: Empirical spectrum of \mathbf{T} equation 8 with preprocessing $\mathcal{T}(y) = y/(1 + |y|)$, for a hierarchical multi-index model with $g_k(z) = \frac{1}{2}\text{He}_2(z) + \frac{1}{2 \cdot 4!}\text{He}_4(z)$ and $\gamma = 1.3$. The covariates dimension is $d = 1000$, while the feature space dimension is $m_* = 20$. The figure illustrates the change in scale of the eigenvalue gaps, transitioning from the informative spikes ($\Theta_d(1)$) to the uninformative bulk ($o_d(1)$). This behavior forms the basis of the selection method described in Def. 3.3.

with $(x)_+ = \max(0, x)$. Further, denoting by k_α the largest index k such that $\alpha > \alpha_k$, i.e. $k_\alpha = \max(m_*, \Theta(\alpha^{1/(2\gamma)}))$ for $\gamma > 1/2$ or $k_\alpha = \max(m_*, \Theta(m_*^{-1+1/(2\gamma)} \alpha^{1/(2\gamma)}))$ for $\gamma < 1/2$, and exploiting the second result in Corollary 7.4,

$$\text{MMSE}_\gamma^{\text{oracle}} = \left(\sum_{k=1}^{k_\alpha} + \sum_{k=k_\alpha+1}^{m_*} \right) (a_k^*)^2 \mathbb{E} [\|\mathbb{E}[\mathbf{w}_k \mathbf{w}_k^\top | \mathcal{D}, \{\mathbf{w}_h^*\}_{h \neq k}] - \mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_F^2] \quad (20)$$

$$\geq C \left(\sum_{k=1}^{k_\alpha} \frac{(a_k^*)^2}{(a_k^*)^2 \alpha} + \sum_{k=k_\alpha+1}^{m_*} (a_k^*)^2 \right) \quad (21)$$

$$= \begin{cases} \Theta(\alpha^{-1+1/(2\gamma)}), & \gamma > 1/2, \alpha \ll m_*^{2\gamma} \\ \Theta(m_*/\alpha), & \gamma > 1/2, \alpha \gg m_*^{2\gamma} \\ \Theta(1), & \gamma < 1/2, \alpha \ll m_* \\ \Theta(m_*/\alpha), & \gamma < 1/2, \alpha \gg m_* \end{cases} \quad (22)$$

Note that we have used

$$\sum_{k=k_\alpha+1}^{m_*} k^{-2\gamma} = \begin{cases} \Theta(k_\alpha^{1-2\gamma}), & \gamma > 1/2, \\ \Theta(m_*^{1-2\gamma}), & \gamma < 1/2. \end{cases} \quad (23)$$

In equation 21, the first term is a lower bound to the (weighted) mean-squared error of weakly recovered features, while the second corresponds to the underfitting contribution of the unlearned ones.

5.2 UPPER BOUND: SPECTRAL METHOD

By definition, any estimator that is a function of the dataset only has a weighted mean-squared error larger than MMSE_γ . In this section we consider the spectral method defined in 3.3. Note that, as the spectral method 3.3 retrieves $r \leq m_*$ directions, we construct the full estimator by filling the remaining columns with zeros to evaluate MSE_γ in Def. 2.5. This zero-padding is harmless: the added zeros do not impact the derived scaling laws or the estimator's agnostic nature. The following Theorem, proven in Kovačević et al. (2025), is valid for generic Gaussian multi-index models with m_* indices. We refer to the original work for further details. In the rest of the section we denote by \mathbb{E} the expected value with respect to $\mathbf{z} \sim \mathcal{N}(\mathbf{0}_{m_*}, \mathbf{I}_{m_*})$ and $\mathbf{y} \sim \mathcal{N}(\sum_{k=1}^{m_*} (a_k^*)^2 g_k(z_k), \Delta)$.

Theorem 5.4 (Theorem 4.1 in Kovačević et al. (2025)) Let $\mathcal{T} : \mathbb{R} \rightarrow \mathbb{R}$ be a preprocessing function subject to Assumption 3.4 and \mathbf{T} defined as in eq. 8. Let $t_1 \geq t_2 \geq \dots \geq t_r \geq \tau$, for some $r \in [m]$, be all the solutions to

$$\det \left(\alpha \mathbb{E} \left[\frac{(\mathbf{z}\mathbf{z}^\top - \mathbf{I}_m)\mathcal{T}(y)}{t - \mathcal{T}(y)} \right] - \mathbf{I}_p \right) = 0 \quad (24)$$

such that

$$t_k \geq \bar{t}_\alpha := \arg \min_{t \geq \tau} \zeta_\alpha(t), \quad \forall k \in [j], \quad (25)$$

where

$$\zeta_\alpha(t) := t \left(1 + \alpha \mathbb{E} \left[\frac{\mathcal{T}(y)}{t - \mathcal{T}(y)} \right] \right). \quad (26)$$

Then, denote $\lambda_1^{\mathbf{T}}, \dots, \lambda_{m_\star}^{\mathbf{T}}$ the largest m_\star eigenvalues of \mathbf{T} . For the top r eigenvalues it holds that

$$\lambda_1^{\mathbf{T}}, \dots, \lambda_r^{\mathbf{T}} \xrightarrow{\text{a.s.}} \zeta_\alpha(t_1), \dots, \zeta_\alpha(t_r), \quad (27)$$

and for the remaining $m_\star - r$ eigenvalues it holds that

$$\lambda_{r+1}^{\mathbf{T}}, \dots, \lambda_{m_\star}^{\mathbf{T}} \xrightarrow{\text{a.s.}} \zeta_\alpha(\bar{t}_\alpha). \quad (28)$$

As a first result, we derive a bound for α_k .

Consider the matrix in equation 24

$$\mathbf{G}(y) := \mathbb{E} \left[\frac{(\mathbf{z}\mathbf{z}^\top - \mathbf{I}_m)\mathcal{T}(y)}{t - \mathcal{T}(y)} \right] = \mathbb{E}_y \left[\frac{\mathbb{E}[\mathbf{z}\mathbf{z}^\top - \mathbf{I}_m | y] \mathcal{T}(y)}{t - \mathcal{T}(y)} \right]. \quad (29)$$

It is straightforward to show that it is a diagonal matrix, due to the parity of each function g_k . Indeed, given $k \neq h$,

$$G_{kh}(y) = \mathbb{E}[z_k z_h | y] \propto \int e^{-\|z\|^2/2} \exp \left(- \left(y - \sum_{k=1}^m a_k g_k(z_k) \right)^2 / (2\Delta) \right) z_h z_k \, dz = 0. \quad (30)$$

For simplicity, we denote $G_k(y) := G_{kk}(y)$. Therefore, the solutions of equation 24 coincides with the solutions of

$$\alpha^{-1} = \mathbb{E}_y \left[\frac{\mathcal{T}(y)}{t - \mathcal{T}(y)} G_k(y) \right], \quad k \in [m]. \quad (31)$$

Further, note that, by definition \bar{t}_α is the solution of

$$\alpha^{-1} = \mathbb{E}_y \left[\left(\frac{\mathcal{T}(y)}{t - \mathcal{T}(y)} \right)^2 \right]. \quad (32)$$

A spectral transition occurs if, for sample complexity α and index $k \in [m]$, the value \bar{t}_α is also solution of equation 31. Indeed by Theorem 4.2 in Kovačević et al. (2025) – which characterizes the overlap of the principal r eigenvalues, such sample complexity corresponds to the k -critical threshold α_k . This implies, for $\alpha = \alpha_k^{\text{SP}}$

$$\begin{cases} (\alpha_k^{\text{SP}})^{-1} = \mathbb{E}_y \left[\frac{\mathcal{T}(y)}{t - \mathcal{T}(y)} G_k(y) \right] \\ (\alpha_k^{\text{SP}})^{-1} = \mathbb{E}_y \left[\left(\frac{\mathcal{T}(y)}{t - \mathcal{T}(y)} \right)^2 \right]. \end{cases} \quad (33)$$

In order to characterize the threshold, we consider the following expansions for small SNR $a_{k^\star} \ll 1$. Define Z the marginal distribution of y and consider the Fourier representation $(2\pi)^{-1/2} e^{-x^2/2} =$

$(2\pi)^{-1} \int d\omega e^{i\omega x - \omega^2/2}$, then

$$Z(y) := \frac{1}{\sqrt{2\pi\Delta}} \mathbb{E}_z \left[\exp \left(- \left(y - \sum_{k=1}^{m_*} a_k^* g_k(z_k) \right)^2 / (2\Delta) \right) \right] \quad (34)$$

$$= \frac{1}{2\pi} \int d\omega e^{i\omega y - \omega^2 \Delta/2} \left(\prod_{h \neq k} \mathbb{E}_z \left[e^{-i\omega a_h^* g_h(z)} \right] \right) \left[e^{-i\omega a_k^* g_k(z)} \right] \quad (35)$$

$$= \frac{1}{2\pi} \int d\omega e^{i\omega y - \omega^2 \Delta/2} \left(\prod_{h \neq k} \mathbb{E}_z \left[e^{-i\omega a_h^* g_h(z)} \right] \right) \left[1 + \sum_{\beta \geq 1} \frac{(-i\omega a_k^*)^\beta}{\beta!} \mathbb{E}_z [g_k^\beta(z)] \right] \quad (36)$$

$$= \frac{1}{2\pi} \int d\omega e^{i\omega y - \omega^2 \Delta/2} \left(\prod_{h \neq k} \mathbb{E}_z \left[e^{-i\omega a_h^* g_h(z)} \right] \right) + O(a_k^* Z'(y)), \quad (37)$$

where we have used the identity $\int d\omega e^{i\omega y} \omega f(\omega) = \frac{\partial}{\partial y} \int d\omega e^{i\omega y} f(\omega)$. Similarly, one can bound the following quantity

$$\left| G_k(y) + a_k^* \mathbb{E}_z [g_k(z)(z^2 - 1)] \frac{Z'(y)}{Z(y)} \right| \quad (38)$$

$$= \frac{1}{Z(y)} \left| \mathbb{E}_{z \sim \mathcal{N}(0, I_p)} [\mathbf{P}(y|z)(z_k^2 - 1) + a_k^* \mathbb{E}_z [g_k(z)(z^2 - 1)] \partial_y \mathbf{P}(y|z)] \right| \quad (39)$$

$$= \frac{1}{2\pi Z(y)} \left| \int_{\mathbb{R}} d\omega e^{i\omega y - \omega^2 \Delta/2} \prod_{h \neq k} \mathbb{E}_z [e^{-i\omega a_h^* g_h(z)}] \left(\mathbb{E}_z [e^{-i\omega a_k^* g_k(z)}(z^2 - 1)] + i\omega a_k^* \mathbb{E}_z [g_k(z)(z^2 - 1)] \mathbb{E}_z [e^{-i\omega a_k^* g_k(z)}] \right) \right| \quad (40)$$

$$= \frac{1}{2\pi Z(y)} \left| \int_{\mathbb{R}} d\omega e^{i\omega y - \omega^2 \Delta/2} \prod_{h \neq k} \mathbb{E}_z [e^{-i\omega a_h^* g_h(z)}] O((a_k^*)^2 \omega^2) \right| \quad (41)$$

$$= \frac{1}{2\pi Z(y)} O \left((a_k^*)^2 \left| \frac{\partial^2}{\partial y^2} \int_{\mathbb{R}} d\omega e^{i\omega y - \omega^2 \Delta/2} \prod_{k \neq h} \mathbb{E}_z [e^{-i\omega a_k^* g_k(z)}] \right| \right) \quad (42)$$

$$= O \left((a_k^*)^2 \left| \frac{Z''(y)}{Z(y)} \right| \right). \quad (43)$$

Note that $\mathbb{E}_z [g_k(z)(z^2 - 1)] \stackrel{\text{parts}}{=} \mathbb{E}_z [g_k''(z)] \neq 0 \forall k \in [m_*]$ by Assumption 5.1. We now look for $t > \tau$ and α that are solutions of eqs. equation 33. In particular, the oracle k -critical threshold in Appendix 5.1 is a lower bound to the spectral one, therefore $\alpha_k^{\text{sp}} = \Omega(\alpha_k^{\text{oracle}})$ and diverges with k . The second equation of equation 33 – due to the monotonicity of the RHS, implies that also the solution t diverges with k . At leading order in a_k^* and t

$$\begin{cases} (\alpha_k^{\text{sp}})^{-1} \approx - \int dy Z(y) \frac{T(y)Z'(y)}{iZ(y)} a_k^* \mathbb{E}_z [g_k(z)(z^2 - 1)] \approx t^{-1} a_k^* \mathbb{E}_z [g_k(z)(z^2 - 1)] \mathbb{E}[T'(y)] \\ (\alpha_k^{\text{sp}})^{-1} \approx t^{-2} \mathbb{E}[T(y)^2]. \end{cases} \quad (44)$$

The system is thereby solved for

$$t^{-1} = \Theta(a_k^*), \quad \alpha_k^{\text{sp}} = \Theta(\alpha_k^{\text{oracle}}). \quad (45)$$

We can now derive the scaling laws for the spectral weighted mean-squared error. Taking $\alpha \gg \alpha_k^{\text{sp}}$, for k large enough, the solution t_k of equation 31 scales as $t_k = \Theta(\alpha a_k^*)$. As a consequence of Theorem 4.2 in Kovačević et al. (2025), specialized in our setting – where we can exploit the diagonality of $\mathbf{G}(y)$, we have that the overlap between the eigenvector \mathbf{v}_k corresponding to the eigenvalue λ_k^T , and the true signal \mathbf{w}_h^* , converges, in the high-dimensional limit, to an overlap

$$m_{hk}^2 := \frac{1}{d^2} \langle \mathbf{w}_h^*, \mathbf{v}_k \rangle^2 = \delta_{kh} \frac{\zeta'_\alpha(t_k)}{\zeta'_\alpha(t_k) + \frac{d}{dt} R_k(t_k)}, \quad (46)$$




Figure 2: Weighted mean square error MSE_γ – see Def. 2.5 – of the spectral estimator of Def. 3.3 with preprocessing function $\mathcal{T}(y) = y/(1 + |y|)$, averaged over 70 instances. The target is given by the hierarchical multi-index model 2.1, with $g_k(z) = g(z) \forall k$, stated in the legend, and $a_k^* \propto k^{-\gamma}$, $\gamma = 1.3$. The covariates dimension is $d = 1000$, the feature space dimension is $m_* = 10$.

where $R_k(t) = \mathbb{E}[z_k^2 \mathcal{T}(y)/(t - \mathcal{T}(y))]$. For $\alpha \gg \alpha_k^{\text{sp}}$ we have that $t_k \gg 1$ and equation 31 is solved by $t_k = \Theta(a_k^* \alpha)$. Putting all together

$$m_{hk}^2 = \delta_{kh} \left(1 - \Theta \left(\frac{(\alpha_k^*)^{-2}}{\alpha} \right) \right). \quad (47)$$

Recall that

$$\text{mse}_k(\mathbf{v}_k) = \frac{1}{d^2} \mathbb{E} [\|\mathbf{v}_k \mathbf{v}_k^\top\|_{\text{F}}^2 + \|\mathbf{w}_k^* \mathbf{w}_k^{*\top}\|_{\text{F}}^2 - 2\langle \mathbf{w}_k^*, \mathbf{v}_k \rangle^2] \rightarrow 2(1 - m_{kk}^2). \quad (48)$$

By repeating computations analogous to the ones in Appendix 5.1, one finds, setting $\hat{\mathbf{W}} = (\mathbf{v}_1, \dots, \mathbf{v}_r)^\top$,

$$\text{MMSE}_\gamma(\hat{\mathbf{W}} := (\hat{\mathbf{W}}^{\text{sp}}, \mathbf{0}_{(m_*-r) \times d})) = \Theta(\text{MMSE}_\gamma^{\text{oracle}}), \quad (49)$$

which proves Theorem 3.2 and 3.5.

5.2.1 GD INTERPRETATION OF THE SPECTRAL METHOD

Note that this spectral method also admits an interpretation in terms of a gradient descent algorithm. Indeed, considering a slightly modified loss than the pure square one, this matrix is nothing but the Hessian of the loss Bonnaire et al. (2025). The early dynamics of GD on the first layer of (5), for weights initialized on a sphere with small radius, will thus follow the proposed spectral method. This construction was recently used in Zhang et al. (2025) to show how two-layer neural networks identify features within the first few GD steps in the first layer weights. More precisely, they show that the early dynamics of a suitable gradient-based algorithm, reads as a power iteration

$$\mathbf{w}_k^{t+1} = \mathbf{w}_k^t - \nabla_{\mathbf{w}_k} \frac{1}{n} \sum_{i=1}^n \ell(y_i, \mathbf{a}^\top \sigma(\mathbf{W}^t \mathbf{x}_i)) \quad (50)$$

$$\approx (\mathbf{I}_d - a_k \sigma''(0) \mathbf{T}) \mathbf{w}_k^t, \quad (51)$$

where \mathbf{T} has the structure defined in equation 8 with pre-processing $\mathcal{T} = \ell'(\cdot, 0)$, and ℓ' denoting the derivative of ℓ with respect to its second argument.

6 PROOF SKETCH OF THEOREM 3.6

In this section we prove Theorem 3.6. This final result demonstrate that neural networks can efficiently learn the target function, achieving the exact same rates as those of the optimal feature




Figure 3: Empirical spectrum density of the matrix \mathbf{T} defined in equation 8, with preprocessing function $\mathcal{T}(y) = y/(1+|y|)$, at different sample complexities, highlighting the sequential emergence of concepts as the sample size increases. The target is given by a hierarchical multi-index model 2.1, with $g_k(z) = \frac{1}{2}\text{He}_2(z) + \frac{1}{2.4!}\text{He}_4(z) \forall k$ and $a_k \propto k^{-\gamma}$, $\gamma = 1.3$. The covariates dimension is $d = 1000$, the feature space dimension is $m_* = 20$. (**top left**) $\alpha = 5$, (**top right**) $\alpha = 164$, (**bottom left**) $\alpha = 611$, (**bottom right**) $\alpha = 1638$.

subspace recovery. We note that our results do not imply that these rates correspond to the decay of the Bayes risk, defined as $\mathbb{E} \left[(f_*(\mathbf{x}) - \mathbb{E}[f_*(\mathbf{x})|\mathbf{x}, \mathcal{D}])^2 \right]$, and which constitutes a lower bound for the risk $R(\Theta)$ in equation 6. Nevertheless, we notice that under the hypothesis of a *fully specialized* Bayes-optimal estimator—i.e., assuming $d^{-1}\hat{\mathbf{W}}^{\text{Bayes}}(\mathbf{W}^*)^\top$ converges to a diagonal matrix in the high-dimensional limit—the risk for a network with matching architecture and Bayes weights, $f(\mathbf{x}) = \sum_{k=1}^{m_*} a_k^* g_k(\hat{\mathbf{w}}_k^{\text{Bayes}})$, achieves the rates established in Theorem 3.2. We leave a more precise analysis of the optimal risk as a future direction.

Algorithm 1 Spectral Initialization and Ridge Training

Input: Dataset $\mathcal{D} = (\mathbf{X}, \mathbf{y}) \in \mathbb{R}^{2n \times d} \times \mathbb{R}^{2n}$, hidden layer width m , regularization λ .

Output: Network parameters $\mathbf{W} \in \mathbb{R}^{p \times d}$, $\mathbf{a} \in \mathbb{R}^p$.

1. Data Splitting

Partition the dataset \mathcal{D} into two disjoint sets \mathcal{D}_1 and \mathcal{D}_2 such that $|\mathcal{D}_1| = |\mathcal{D}_2| = n$.

2. Feature Learning (Spectral Initialization on \mathcal{D}_1)

Compute the spectral estimator $\hat{\mathbf{W}}^{\text{sp}} \in \mathbb{R}^{r \times d}$ Def. 3.3 using \mathcal{D}_1 .

Sample a random matrix $\mathbf{Z} \in \mathbb{R}^{p \times r}$ with i.i.d. entries drawn from $\mathcal{N}(0, r^{-1})$.

Set first-layer weights: $\mathbf{W} = \mathbf{Z}\hat{\mathbf{W}}^{\text{sp}}$.

3. Readout Training (Ridge Regression on \mathcal{D}_2)

Compute the feature matrix $\Psi \in \mathbb{R}^{n \times p}$ on \mathcal{D}_2 , where $\Psi_{ij} = \sigma(\langle \mathbf{w}_j, \mathbf{x}_i^{(2)} \rangle + b_j)$, $b_j \sim \mathcal{N}(0, 1)$.

Solve for the second-layer weights:

$$\mathbf{a} = (\Psi^\top \Psi + n\lambda \mathbf{I}_p)^{-1} \Psi^\top \mathbf{y}^{(2)}$$

return \mathbf{W}, \mathbf{a}

Denote $\hat{\mathbf{W}}^{\text{sp}} \in \mathbb{R}^{r \times d}$ the spectral estimator defined in 3.3. In Appendix 5.2 we have shown that each column k correlates with the signal component \mathbf{w}_k^* only, with squared overlap $m_k = 1 - \Theta((\alpha_k^*)^{-2}/\alpha)$. Denote \mathbf{v}_k as the k^{th} column of the spectral estimator and $\eta_k^2 := (\alpha_k^*)^2/\alpha \ll 1$. Up

to negligible corrections $O(\eta_k^2)$

$$\mathbf{w}_k^* = \mathbf{v}_k + \eta_k \boldsymbol{\xi}_k, \quad (52)$$

with $\boldsymbol{\xi}_k$ a unit vector orthogonal to all $\{\mathbf{v}_h\}_{h \in [r]}$. Given a covariate $\mathbf{x} \in \mathbb{R}^d$, define the projected inputs $\mathbf{s} \in \mathbb{R}^r$ such that $s_k = \langle \mathbf{v}_k, \mathbf{x} \rangle$, $k \in [r]$. The first layer pre-activations are given by

$$\mathbf{W}\mathbf{x} = \mathbf{Z}\hat{\mathbf{W}}^{\text{sp}}\mathbf{x} = \mathbf{Z}\mathbf{s}. \quad (53)$$

Similarly, up to negligible corrections $O(\eta_k^2)$, the target function

$$f_*(\mathbf{x}) = \sum_{k=1}^r a_k^* g_k(s_k + \eta_k \langle \boldsymbol{\xi}_k, \mathbf{x} \rangle) + \sum_{k=r+1}^{m_*} a_k^* g_k(\langle \mathbf{w}_k^*, \mathbf{x} \rangle) \quad (54)$$

$$= \sum_{k=1}^r a_k^* g_k(s_k) + \sum_{k=1}^r a_k^* g'_k(s_k) \eta_k \langle \boldsymbol{\xi}_k, \mathbf{x} \rangle + \sum_{k=r+1}^{m_*} a_k^* g_k(\langle \mathbf{w}_k^*, \mathbf{x} \rangle). \quad (55)$$

Due to the orthogonality between $\boldsymbol{\xi}_k$ and \mathbf{v}_h , for any $h, k \in [r]$, the variables $\langle \mathbf{x}, \boldsymbol{\xi}_k \rangle$ and s_h are independent centered Gaussian variables. With respect to the projected input space, the effective target function is

$$f_*^{\text{eff}}(\mathbf{s}) := \mathbb{E}[f_*(\mathbf{x}) | \mathbf{s}] = \sum_{i=1}^r a_i^* g_i(s_i). \quad (56)$$

Then, the readout training in Algorithm 1 corresponds to random feature ridge regression Rahimi & Recht (2007) with weights $\mathbf{Z} = (z_1, \dots, z_p)^\top$ on the projected covariates $\mathbf{s}_i = \hat{\mathbf{W}}^{\text{sp}} \mathbf{x}_i^{(2)}$, $i \in [n]$, where $\{\mathbf{x}_i^{(2)}\}_{i \in [n]}$ are the covariates in the dataset \mathcal{D}_2 :³

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a} \in \mathbb{R}^p} \frac{1}{2n} \sum_{i=1}^n \left(y_i - \sum_{j=1}^p a_j \sigma(\langle \mathbf{z}_j, \mathbf{s}_i \rangle) \right)^2 + \frac{\lambda}{2} \|\mathbf{a}\|^2 \quad (57)$$

Since σ is bounded and continuous, we can apply Theorem 1 in Rudi & Rosasco (2017)⁴, choosing $p = \omega(n^{1/2})$, $\lambda = \Theta(n^{-1/2})$, so that we obtain that the risk equation 6 $R(\hat{\mathbf{a}}, \mathbf{W})$ satisfies

$$R(\hat{\mathbf{a}}, \mathbf{W}) - R_{f_{\mathcal{H}}} = O(n^{-1/2}), \quad (58)$$

where, given the *reproducing kernel Hilbert space* (RKHS) \mathcal{H} associated to the kernel $K(\mathbf{s}, \mathbf{s}') = \mathbb{E}_{\mathbf{z}, b}[\sigma(\mathbf{s}^\top \mathbf{s} + b)\sigma(\mathbf{w}^\top \mathbf{s}' + b)]$,

$$R_{f_{\mathcal{H}}} = \min_{f \in \mathcal{H}} \mathbb{E}_{y, \mathbf{s}}[(f(\mathbf{s}) - y)^2]. \quad (59)$$

Such irreducible risk is lower-bounded by

$$R_* = \mathbb{E}[(f_*^{\text{eff}}(\mathbf{s}) - y)^2] = \Theta \left(\sum_{k=1}^r (a_k^*)^2 \eta_k^2 + \sum_{k=r+1}^{m_*} (a_k^*)^2 \right) = \Theta(\text{MMSE}_\gamma), \quad (60)$$

where, as for the spectral method mean-squared error, the two terms correspond to the approximation error of learned features and underfitting of unlearned ones. If $f_*^{\text{eff}} \in \mathcal{H}$, the argument is complete. By assumption, the function σ allows for a decomposition in Hermite polynomials

$$\sigma(z) = \sum_{\beta \geq 0} \frac{\sigma_\beta}{\beta!} \text{He}_\beta(z), \quad \sigma_\beta := \frac{1}{\beta!} \mathbb{E}_z[\text{He}_\beta(z)\sigma(z)]. \quad (61)$$

³In order to simplify the notation, we denote the labels in \mathcal{D}_2 as $y_i := y_i^{(2)}$.

⁴Note that the assumption of bounded outputs in Theorem 1 is relaxed in the Appendix of Rudi & Rosasco (2017), see Assumption 4. Further, the analysis in Rudi & Rosasco (2017) allows for a more refined error rate in our setting. However, the current result is already subleading with respect to the rates in Theorem 3.2.

Further, being bounded, it cannot be a polynomial of finite degree, as there is no β_0 such that $\sigma_\beta = 0$ for all $\beta \geq \beta_0$, excluding the trivial case of constant functions. Considering the Taylor expansion of Hermite polynomials

$$\text{He}_\beta(z+b) = \sum_{\zeta=0}^{\beta} \binom{\beta}{\zeta} b^{\beta-\zeta} \text{He}_\zeta(z), \quad (62)$$

with $\binom{\beta}{\zeta} = \frac{\beta!}{\zeta!(\beta-\zeta)!}$, we find that the kernel K is given by

$$K(\mathbf{s}, \mathbf{s}') = \mathbb{E}_{\mathbf{z}, b} [\sigma(\mathbf{z}^\top \mathbf{s} + b) \sigma(\mathbf{z}^\top \mathbf{s}' + b)] \quad (63)$$

$$= \mathbb{E}_{\mathbf{z}, b} \sum_{\beta, \beta' \geq 0} \sum_{\zeta=0}^{\beta} \sum_{\zeta'=0}^{\beta'} \frac{\sigma_\beta \sigma_{\beta'}}{\zeta! \zeta'! (\beta-\zeta)! (\beta'-\zeta')!} b^{\beta-\zeta} b^{\beta'-\zeta'} \text{He}_\zeta(\langle \mathbf{z}, \mathbf{s} \rangle) \text{He}_{\zeta'}(\langle \mathbf{z}, \mathbf{s}' \rangle) \quad (64)$$

By Proposition 11.31 in O'Donnell (2014),

$$K(\mathbf{s}, \mathbf{s}') = \mathbb{E}_b \sum_{\beta, \beta' \geq 0} \sum_{\zeta=0}^{\min(\beta, \beta')} \frac{\sigma_\beta \sigma_{\beta'}}{\zeta! (\beta-\zeta)! (\beta'-\zeta)!} b^{\beta+\beta'-2\zeta} \frac{\langle \mathbf{s}, \mathbf{s}' \rangle^\zeta}{r^\zeta} \quad (65)$$

$$= \mathbb{E}_b \sum_{\zeta=0}^{\infty} \sum_{\beta, \beta' \geq \zeta} \frac{\sigma_\beta \sigma_{\beta'}}{\zeta! (\beta-\zeta)! (\beta'-\zeta)!} b^{\beta+\beta'-2\zeta} \frac{\langle \mathbf{s}, \mathbf{s}' \rangle^\zeta}{r^\zeta} \quad (66)$$

$$= \sum_{\zeta=0}^{\infty} \mathbb{E}_b \left(\underbrace{\sum_{\beta \geq \zeta} \frac{\sigma_\beta}{\zeta! (\beta-\zeta)!} b^{\beta-\zeta}} \right)^2 \frac{\langle \mathbf{s}, \mathbf{s}' \rangle^\zeta}{r^\zeta} \quad (67)$$

The bracketed term is always strictly positive as $\sum_{\beta \geq \zeta} \frac{\sigma_\beta}{\zeta! (\beta-\zeta)!} b^{\beta-\zeta}$ in b cannot be identically zero. This would require $\sigma_\beta = 0$ for all $\beta > \zeta$, which contradicts the hypothesis of boundedness. Expanding the term $\langle \mathbf{s}, \mathbf{s}' \rangle$, it follows that the kernel admits a decomposition as $K(\mathbf{s}, \mathbf{s}') = \Phi(\mathbf{s})^\top \Phi(\mathbf{s}')$, where Φ is a feature map with components given by multivariate monomials $\lambda_\beta s_1^{\beta_1} \dots s_r^{\beta_r}$, for some $\lambda_\beta > 0$ for all degrees $|\beta|$. Theorem 4.21 in Christmann & Steinwart (2008) readily implies that the RKHS of K includes all finite degree polynomials, completing our argument.

7 USEFUL RESULTS ON SINGLE-INDEX MODELS

In this section we present various results applied in the derivation of the lower bound to MMSE_γ in Appendix 5.1. In particular, we focus on the following single-index model setting.

Definition 7.1 Let $\mathbf{w}^* \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, $g : \mathbb{R} \rightarrow \mathbb{R}$ satisfy Assumption 5.1. Consider the supervised learning problem of estimating \mathbf{w}_* from n i.i.d. observations $\mathcal{D} = \{(\mathbf{x}_i, y_i) \in \mathbb{R}^{d \times 1} : i \in [n]\}$ generated as

$$\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad y_i = \sqrt{\lambda} g(\langle \mathbf{x}_i, \mathbf{w}^* \rangle) + \xi_i, \quad (68)$$

where $\lambda \geq 0$ is the signal-to-noise ratio (SNR) and $\xi_i \sim \mathcal{N}(0, 1)$ is additive noise.

Given an estimator $\hat{\mathbf{w}}$ of \mathbf{w}^* that is a function of the dataset, we evaluate its estimation error using the following matrix-MMSE:

$$\text{mse}(\mathbf{w}) := \frac{1}{d^2} \mathbb{E} [\|\mathbf{w}\mathbf{w}^\top - \mathbf{w}^* \mathbf{w}^{*\top}\|_F^2], \quad (69)$$

where \mathbb{E} computes the expected value with respect to the joint distributions of \mathcal{D} and \mathbf{w}_* . A lower bound to this quantity is given by the following optimal matrix-MSE

$$\text{mmse} := \frac{1}{d^2} \mathbb{E} [\|\mathbb{E}[\mathbf{w}\mathbf{w}^\top | \mathcal{D}] - \mathbf{w}^* \mathbf{w}^{*\top}\|_F^2] = \arg \min_{\mathbf{Q} \in \mathbb{R}^{d \times d}} \mathbb{E} [\|\mathbf{Q} - \mathbf{w}^* \mathbf{w}^{*\top}\|_F^2]. \quad (70)$$

Intuitively, such optimal estimation error decreases with the SNR.

Lemma 7.2 *In the setting defined in 7.1, the optimal matrix-MSE satisfies*

$$\frac{\partial}{\partial \lambda} \text{mmse} < 0. \quad (71)$$

The proof follows from direct computation. We provide here a condensed derivation.

$$\text{mmse} = \mathbb{E} \left[\|\mathbf{w}\mathbf{w}^T\|_F^2 \right] - \mathbb{E}_{\mathbf{y}} \left[\|\mathbb{E}[\mathbf{w}\mathbf{w}^T \mid \mathbf{y}]\|_F^2 \right]. \quad (72)$$

Since the first term is independent of λ , it suffices to show that the second term is non-decreasing. Let us denote the posterior mean as $\langle \mathbf{w}\mathbf{w}^T \rangle_\lambda \equiv \mathbb{E}[\mathbf{w}\mathbf{w}^T \mid \mathbf{y}]$. We compute its derivative with respect to λ :

$$\frac{d}{d\lambda} \mathbb{E}_{\mathbf{y}} \left[\|\langle \mathbf{w}\mathbf{w}^T \rangle_\lambda\|_F^2 \right] = \int d\mathbf{y} \frac{\partial}{\partial \lambda} \left(P(\mathbf{y}) \|\langle \mathbf{w}\mathbf{w}^T \rangle_\lambda\|_F^2 \right). \quad (73)$$

Using the explicit form of $P(\mathbf{y}|\mathbf{w}) \propto \exp(-\frac{1}{2}\|\mathbf{y} - \sqrt{\lambda}g(\mathbf{X}\mathbf{w})\|^2)$, the derivative $\partial_\lambda \ln P(\mathbf{y})$ introduces terms involving $\mathbf{y}^T \langle g(\mathbf{X}\mathbf{w}) \rangle_\lambda$. We handle these terms via the multivariate Stein's Lemma (corresponding to integration by parts), which states that for the Gaussian measure, $\mathbb{E}_{\mathbf{y}}[\mathbf{y}^T \mathbf{f}(\mathbf{y})] = \mathbb{E}_{\mathbf{y}}[\text{Tr}(\nabla \mathbf{f}) + \sqrt{\lambda} \langle h(\mathbf{X}\mathbf{w}) \rangle_\lambda^T \mathbf{f}]$.

Applying this identity results in cancellations of the lower-order terms. The surviving term is proportional to the gradient of the estimator with respect to the observations $\nabla_{\mathbf{y}} \langle \mathbf{w}\mathbf{w}^T \rangle_\lambda = \sqrt{\lambda} \text{Cov}(\mathbf{w}\mathbf{w}^T, g(\mathbf{X}\mathbf{w}) \mid \mathbf{y})$. Therefore,

$$\frac{d}{d\lambda} \mathbb{E}_{\mathbf{y}} \left[\|\langle \mathbf{w}\mathbf{w}^T \rangle_\lambda\|_F^2 \right] = \mathbb{E}_{\mathbf{y}} \left[\|\text{Cov}(\mathbf{w}\mathbf{w}^T, g(\mathbf{X}\mathbf{w}) \mid \mathbf{y})\|_F^2 \right] \implies \frac{\partial}{\partial \lambda} \text{mmse} < 0. \quad (74)$$

We now consider the high-dimensional limit $n, d \rightarrow \infty$ with fixed ratio $\alpha = n/d$, referred to as the sample complexity. The following theorem specializes Theorems 1 and 2 of Barbier et al. (2019) to the setting of interest.

Theorem 7.3 (Barbier et al. (2019)) *Consider the setting defined in 7.1. Then, in the limit $n, d \rightarrow \infty$, with fixed ratio $n/d = \alpha$,*

$$\frac{1}{d^2} \mathbb{E} \|\mathbf{w}^* \mathbf{w}^{*\top} - \mathbb{E}[\mathbf{w}\mathbf{w}^\top \mid \mathcal{D}]\|_F^2 \rightarrow 1 - m^2, \quad (75)$$

and, given $\mathbf{w} \sim P(\cdot \mid \mathcal{D})$,

$$\frac{1}{d} |\mathbf{w}^\top \mathbf{w}^*| \xrightarrow{\mathbb{P}} m, \quad (76)$$

where $m = m(\alpha)$ is the maximizer of

$$\sup_{m \in [0,1]} f_{\text{RS}}(m), \quad f_{\text{RS}}(m) := \{m + \log(1 - m) + 2\alpha \Psi_{\text{out}}(m)\}, \quad (77)$$

$$\Psi_{\text{out}}(m) := \mathbb{E}_{W,V,Y} \log \mathbb{E}_{w \sim \mathcal{N}(0,1)} \left[P(Y \mid \sqrt{m}V + \sqrt{1 - m}w) \right], \quad (78)$$

with $V, W \sim \mathcal{N}(0, 1)$, $Y \sim P(\cdot \mid \sqrt{m}V + \sqrt{1 - m}W)$.

As a direct consequence of the theorem, the information-theoretic weak recovery threshold α^{IT} is the smallest sample complexity α such that the maximizer of the free entropy equation 77 is $m \neq 0$. Equivalently, α^{IT} is the smallest α such that $\text{mmse} < 1$. Lemma 7.2 readily implies that $\alpha^{\text{IT}} = \alpha^{\text{IT}}(\lambda)$ is decreasing with the SNR.

Finally, we characterize the IT weak recovery threshold in the limit of small SNR. By expanding equation 78 around $\lambda = 0$, we derive the following Corollary.

Corollary 7.4 (IT weak recovery threshold in the large noise regime) *Consider the setting of Theorem 7.3. Then, in the limit $\lambda \rightarrow 0$, the information theoretic weak-recovery threshold satisfies*

$$\alpha^{\text{IT}} = \Theta(\lambda^{-1}) \quad (79)$$

and, for $\alpha \rightarrow \infty$, and fixed λ the optimal matrix-MSE scales as

$$\text{mmse} = O\left(\frac{1}{\lambda\alpha}\right). \quad (80)$$

In order to simplify the notation, we introduce the shorthand $\phi_\beta(V; m) = \mathbb{E}_{\mathbf{w}}[g^\beta(\sqrt{m}V + \sqrt{1-m}W)]$.

$$\mathbb{E}_{\mathbf{w}}[\mathbb{P}(Y = y | \sqrt{m}V + \sqrt{1-m}W)] = \frac{1}{\sqrt{2\pi}} \mathbb{E}_{\mathbf{w}} \left[\exp \left(-\frac{(y - \sqrt{\lambda}(\sqrt{m}V + \sqrt{1-m}W))^2}{2} \right) \right] \quad (81)$$

$$= \frac{e^{-y^2/2}}{\sqrt{2\pi}} \left(1 + \sum_{\beta \geq 1} \frac{\lambda^{\beta/2}}{\beta!} \phi_\beta(V; m) \text{He}_\beta(y) \right) \quad (82)$$

and therefore, expanding $\mathbb{P}(y | \sqrt{m}V + \sqrt{1-m}W)$ in a similar fashion and leveraging the expansion $\log(1+x) = -\sum_{k \geq 0} (-1)^k x^k / k$, up to constant terms with respect to m ,

$$\Psi_{\text{out}}(m) = \mathbb{E}_{V,W,Y} \left[-\frac{y^2}{2} + \sum_{\beta \geq 1} \frac{\lambda^{\beta/2}}{\beta!} \phi_\beta(V; m) \text{He}_\beta(y) - \frac{1}{2} \left(\sum_{\beta \geq 1} \frac{\lambda^{\beta/2}}{\beta!} \phi_\beta(V; m) \text{He}_\beta(y) \right)^2 \right] + O(\lambda^{3/2}) \quad (83)$$

$$= \mathbb{E}_V \int \frac{e^{-y^2/2}}{\sqrt{2\pi}} \left(-\frac{y^2}{2} + \sum_{\beta \geq 1} \frac{\lambda^{\beta/2}}{\beta!} \phi_\beta(V; m) \text{He}_\beta(y) - \frac{\lambda}{2} \phi_1^2(V; m) y^2 \right) dy + \quad (84)$$

$$+ \mathbb{E}_V \int \frac{e^{-y^2/2}}{\sqrt{2\pi}} \left(\left(-\frac{y^2}{2} + \sum_{\beta \geq 1} \frac{\lambda^{\beta/2}}{\beta!} \phi_\beta(V; m) \text{He}_\beta(y) \right) \sum_{\beta \geq 1} \frac{\lambda^{\beta/2}}{\beta!} \phi_\beta(V; m) \text{He}_\beta(y) \right) dy + O(\lambda^{3/2})$$

$$= -\frac{1}{2} + \frac{\lambda}{2} \mathbb{E}_V[\phi_1^2(V; m)] - \frac{\lambda}{2} \mathbb{E}_V[\phi_2(V; m)] + O(\lambda^{3/2}) \quad (85)$$

$$= \frac{1}{2} (-1 - \lambda \mathbb{E}_{z \sim \mathcal{N}(0,1)}[g^2(z)] + \lambda \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}_2, \mathbf{C})}[g(z_1)g(z_2)]) + O(\lambda^{3/2}) \quad (86)$$

with

$$\mathbf{C} = \begin{pmatrix} 1 & m \\ m & 1 \end{pmatrix}. \quad (87)$$

In the above we used

$$\mathbb{E}_V[\phi_2(V; m)] = \mathbb{E}_{W,V}[g^2(\sqrt{m}V + \sqrt{1-m}W)] = \mathbb{E}_z[g^2(z)], \quad (88)$$

$$\mathbb{E}_V[\phi_1^2(V; m)] = \mathbb{E}_V[\mathbb{E}_W[g(\sqrt{m}V + \sqrt{1-m}W)] \mathbb{E}_{W'}[g(\sqrt{m}V + \sqrt{1-m}W')]] \quad (89)$$

$$= \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}_2, \mathbf{C})}[g(z_1)g(z_2)]. \quad (90)$$

Assumption 5.1 ensures that g can be decomposed in the Hermite basis as

$$g(z) = \sum_{k \geq 0} \frac{c_k}{k!} \text{He}_k(z), \quad c_k := \frac{1}{k!} \mathbb{E}_{z \sim \mathcal{N}(0,1)}[g(z) \text{He}_k(z)]. \quad (91)$$

Leveraging Proposition 11.31 in O'Donnell (2014),

$$\mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}_2, \mathbf{C})}[g(z_1)g(z_2)] = \sum_{k \geq 0} \frac{c_k^2}{k!} m^k \in [0, \mathbb{E}_z[g^2(z)]], \quad (92)$$

which is a non-decreasing function of $m \in [0, 1]$. Note that, since g has generative exponent 2, and $\mathbb{E}_z[g(z)] = 0$, we have that $c_0 = c_1 = 0$ necessarily. Moreover, by Assumption 5.1, $c_2 \neq 0$ and bounded. Therefore, we are interested in maximizing the quantity⁵

$$f_{\text{RS}}(m) := m + \log(1-m) + \alpha \lambda \sum_{k \geq 2} \frac{c_k^2}{k!} m^k + O(\alpha \lambda^{3/2}) \quad (93)$$

$$= \frac{1}{2} (\alpha \lambda c_2^2 - 1) m^2 + \sum_{k > 2} \left(\frac{\alpha \lambda}{k!} c_k^2 - \frac{1}{k} \right) m^k + O(\alpha \lambda^{3/2}), \quad (94)$$

⁵Note that we are neglecting constant terms with respect to m .

where we have expanded $\log(1 - m)$. For $\alpha > \lambda^{-1}c_2^{-2}$, $m = 0$ is a minimum of the free entropy, therefore $\alpha^{\text{IT}} \leq \lambda^{-1}c_2^{-2}$. Denote

$$D := \inf_{m \in (0,1]} \frac{-m - \log(1 - m)}{\sum_k \frac{c_k^2}{k!} m^k}, \quad (95)$$

which is strictly positive and well-defined. Note that

$$\lim_{m \rightarrow 0^+} \frac{-m - \log(1 - m)}{\sum_k \frac{c_k^2}{k!} m^k} = c_2^{-2} \implies \frac{1}{2\mathbb{E}_z[g^2(z)]} \leq D \leq c_2^{-2} = \frac{1}{2\mathbb{E}_z[g''(z)]}. \quad (96)$$

Then, for all $\alpha < D\lambda^{-1}$, $f(m \neq 0) < f(0) = 0$, i.e. $m = 0$ is the global maximizer and

$$D \leq \alpha^{\text{IT}} \lambda \leq c_2^{-2} \implies \alpha^{\text{IT}} = \Theta(\lambda^{-1}), \quad (97)$$

For the second result, we first consider the setting $\lambda \rightarrow 0$ and $\alpha \rightarrow \infty$, with $\alpha \gg \lambda$. There exists a non-zero maximizer m , which satisfies⁶

$$\frac{d}{dm} f_{\text{RS}}(m) \stackrel{!}{=} 0 \implies \frac{1}{1 - m} = \alpha \lambda \sum_{k \geq 2} \frac{c_k^2}{k!} m^{k-2} \implies \sum_{k \geq 2} \frac{c_k^2}{k!} (m^{k-2} - m^{k-1}) = \frac{1}{\alpha \lambda}, \quad (98)$$

The equation is solved by $m = 1 - \Theta(\frac{1}{\alpha \lambda})$. Theorem 7.3 implies that, in this regime, $\text{mmse}(\lambda) = \Theta((\alpha \lambda)^{-1})$. Together Lemma 7.2, the result for arbitrary λ is proved.

⁶In the following we retain the leading terms in the free entropy expansion.