

HOW 2-LAYER NETWORKS LEARN, ONE (GIANT) STEP AT A TIME

Anonymous authors
Paper under double-blind review

ABSTRACT

We investigate theoretically how the features of a 2-layer neural network adapt to the structure of the target function through a few large batch gradient descent steps, leading to improvement in the approximation capacity with respect to the initialization. We compare the influence of batch size and that of multiple (but finitely many) steps. For a single gradient step, a batch of size $n = \mathcal{O}(d)$ is both necessary and sufficient to align with the target function, although only a single direction can be learned. In contrast, $n = \mathcal{O}(d^2)$ is essential for neurons to specialize to multiple relevant directions of the target with a single gradient step. Even in this case, we show there might exist “hard” directions requiring $n = \mathcal{O}(d^\ell)$ samples to be learned, where ℓ is known as the leap index of the target. The picture drastically improves over multiple gradient steps: we show that a batch-size of $n = \mathcal{O}(d)$ is indeed enough to learn multiple target directions satisfying a staircase property, where more and more directions can be learned over time. Finally, we discuss how these directions allow to drastically improve the approximation capacity and generalization error over the initialization, illustrating a separation of scale between the random features/lazy regime, and the feature learning regime. Our technical analysis leverages a combination of techniques related to concentration, projection-based conditioning, and Gaussian equivalence which we believe are of independent interest. By pinning down the conditions necessary for specialization and learning, our results highlight the interaction between batch size and number of iterations, and lead to a hierarchical depiction where learning performance exhibits a stairway to accuracy over time and batch size, shedding new light on how neural nets adapt to features of the data.

A central property behind the success of neural networks is their capacity to adapt to the features in the training data. Indeed, many of the classical machine learning methods, e.g. linear or logistic regression, are specifically designed to a restricted class of functions (e.g. generalized linear functions). Others, such as kernel methods, can adapt to larger function classes (e.g. square-integrable functions), but sometimes at prohibitively many samples. Despite the limitations, these methods enjoy well-understood theoretical guarantees: they are convex (hence easy to train) and given a target function, it is well-understood how many samples are needed to achieve a target accuracy. The situation is dramatically different for neural networks: despite being universal approximators, little is known on how to optimally train them or how many hidden units and/or samples are required to learn a given class of functions. Nevertheless, they have proven to be flexible, efficient and easy to optimize in practice, properties which are often attributed to their capacity to adapt to features in the data. Curiously, most of our current theoretical understanding of neural networks stems from the investigation of their lazy regime where features are *not* learned during training. In this work, we take some (giant) steps forward from the lazy regime.

Our central goal is to paint a complete picture of how two-layer neural networks adapt to the features of training data $(z^\nu, y^\nu)_{\nu=1}^n \in \mathbb{R}^{d+1}$ in the *early phase* of training after the first few steps of gradient descent. We recall the reader that for data is supported in a high-dimensional space, the curse of dimensionality prevents efficient learning even under standard regularity assumptions on the target function such as

Lipschitzness (Devroye et al., 2013). Hence, understanding the efficient learning performance of neural networks observed in practice requires additional assumptions on the data distribution. In this work, we focus on a popular synthetic data model consisting of: a) independently drawn standard Gaussian covariates $\mathbf{z}^\nu \sim \mathcal{N}(0, I_d)$; b) a target function $y^\nu = f^*(\mathbf{z}^\nu)$ depending only on a finite number of relevant directions, also known as a *multi-index model*. In other words, there exists a finite number of orthogonal *teacher weights* $(\mathbf{w}_1^*, \dots, \mathbf{w}_r^*)$ and a function $g^* : \mathbb{R}^r \rightarrow \mathbb{R}$ such that:

$$y = f^*(\mathbf{z}) := g^*(\langle \mathbf{w}_1^*, \mathbf{z} \rangle, \dots, \langle \mathbf{w}_r^*, \mathbf{z} \rangle). \quad (1)$$

Note that in this model the features are isotropic, with all the structure in the data being in the target. In contrast to other popular models for structured data, such as low-dimensional support of the inputs or smoothness of the target function, kernel methods do not adapt to target functions depending on low-dimensional projections (Bach, 2017). This makes it an ideal playground for quantifying the adaptativity of neural networks in the feature-learning regime. Given this class of structured targets, we consider supervised learning with the simplest universal approximator neural network: a fully-connected two-layer network with first and second layer weights $W \in \mathbb{R}^{p \times d}$ and $\mathbf{a} \in \mathbb{R}^p$ and activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$:

$$\hat{f}(\mathbf{z}; W, \mathbf{a}) = \frac{1}{\sqrt{p}} \sum_{i=1}^p a_i \sigma(\langle \mathbf{w}_i, \mathbf{z} \rangle). \quad (2)$$

The central question in this work is to uncover how this two-layer neural net adapts to the low-dimensional target structure during training. In particular, we will be interested in quantifying how much data is required for the relevant directions of the target to be learned, and how feature learning translates into the approximation capacity of the network with respect to kernel methods.

1 SUMMARY OF MAIN RESULTS

To outperform the network at initialization (which can be regarded as a kernel method) the network must adapt to the data distribution. Mathematically, this translates to developing correlation in the first layer weights with the target subspace

$$V^* := \text{span}(\mathbf{w}_1^*, \dots, \mathbf{w}_r^*).$$

First we focus on **feature learning**, i.e. how the subspace $V^* \subset \mathbb{R}^d$ is learned during training.

First GD step — First, we discuss the case of a **single gradient step of full-batch gradient descent**, which turns out to be already non-trivial (Ba et al., 2022; Damian et al., 2022) and consider the update:

$$\mathbf{w}_i^1 = \mathbf{w}_i^0 - \frac{\eta}{2n} \sum_{\nu=1}^n \nabla_{\mathbf{w}_i} \left(y^\nu - \hat{f}(\mathbf{z}^\nu; W^0, \mathbf{a}^0) \right)^2. \quad (3)$$

Theorems 1 and 2 identify a fundamental interplay between batch size and the complexity of the underlying target function in what can be learned. They are summarized in Fig. 1 (and particular numerical examples are shown in the Appendix, Figs. 3 and 5). More precisely:

- We show that developing meaningful correlation with the target function requires a large batch size $n = \mathcal{O}(d)$ and learning rate $\eta = \Theta(p)$ when p, d, n are large. However, feature learning remains limited in this regime since we prove only a *single* direction can be learned. Thus, if the target depends on several relevant directions, only a "single neuron" approximation can be learned.




Figure 1: **Learning with a single gradient step.** Illustration of the relationship between batch size and target function complexity for learning multi-index functions with a single giant step in the $n = \Theta(d^k)$ regime (Thms. 1 & 2).

- Surpassing the single direction approximation with a single step *requires* a larger batch size with *at least* $n = \mathcal{O}(d^2)$ samples. This allows for the network weights to *specialize* to multiple target directions.
- Nonetheless, we show that there might be *hard* directions in the target which cannot be learned with $n = \mathcal{O}(d^2)$. Learning these directions necessitates a batch size of at least $n = \mathcal{O}(d^\ell)$, as well as suppressing the directions learned at $n = \mathcal{O}(d^{\ell-1})$, where ℓ is the *leap index* of the target (precisely defined in Def. 1) that thoroughly speaking corresponds to the lowest non-zero degree of the Hermite polynomials in the expansion of the target in this directions.

For a single-index function with leap index ℓ , one-pass SGD has a sample complexity of $\mathcal{O}(d^{\ell-1})$ (Ben Arous et al., 2021), and it has been recently shown that a smoothed version of SGD achieves a sample complexity of $\mathcal{O}(d^{\ell/2})$ (Damian et al., 2023). This matches a lower bound from the correlation statistical query family, which encompasses all SGD-like methods. In our single step setting, the sample complexity for large batch learning is $\mathcal{O}(d^\ell)$, which is worse than both the aforementioned methods. However, the time complexity of each algorithm paints a different picture. Both SGD algorithms are sequential, and require $\mathcal{O}(d)$ operations per step, which leads to a total time complexity of $\mathcal{O}(d^{\ell/2+1})$ at least. On the other hand, the computation of the update in Eq. 1 is simply an average of independent terms, which is easy to parallelize. Including the time to compute the average of each term e.g. using a *Gossip algorithm* (Boyd et al., 2006), this sums up to a time complexity of $\mathcal{O}(d + \log(n))$. Such an algorithm is also amenable to decentralized learning schemes, where each agent only has access to a fraction of the overall data.

Learning over many GD iterations — The situation drastically improves when taking for **multiple gradient steps**. In this case, focusing on the linear batch size $n = \mathcal{O}(d)$ regime, and using a fresh batch of data at each GD iteration:

$$\mathbf{w}_i^{t+1} = \mathbf{w}_i^t - \frac{\eta}{2n} \sum_{\nu=1}^n \nabla_{\mathbf{w}_i} \left(y^\nu - \hat{f}(\mathbf{z}^\nu; W^t, \mathbf{a}^0) \right)^2, \quad (4)$$

Note that splitting the training of the first and second layers and using a fresh batch of data at each iteration is a common approximation in this literature (Damian et al., 2022; Ba et al., 2022; Bietti et al., 2022). In contrast to the recent works considering the population limit, we stress that here we take the batch size n to scale with the dimension d . In the realm of distributed and federated learning, scenarios with large batches, a single pass, and few iterations are often the norm (Goyal et al., 2017; Li et al., 2020) (for instance this is the case when training large language models), further underlining the relevance of this scenario. In this case, Theorem 3 shows that more complex subspaces of the target directions *may* be progressively learned at each iteration, as we illustrate in Fig. 2 (and provide many numerical examples in App.A). More precisely:

- Each additional gradient step allows for learning new perpendicular directions **upon the important condition that they are linearly connected to the previously learned directions** (see Sec. 2.2 for precise definitions of this staircase property). Therefore, in contrast with a single step, taking multiple steps allows to learn a multiple-index target with only $n = \mathcal{O}(d)$ samples.
- Nonetheless, directions that are not coupled through the staircase property and with zero first Hermite coefficient cannot be learned in any finite number of steps. In fact, as discussed previously, they require a batch size of at least $n = \mathcal{O}(d^2)$. In other words, while multiple steps help specialization, it cannot help learning “hard” target directions.
- Similar to the *saddle-to-saddle* dynamics under gradient flow (Jacot et al., 2021; Abbe et al., 2023), the dynamics described through Theorem 3 involves sequential learning of directions. However, in the one-sample SGD regime (Ben Arous et al., 2021; Abbe et al., 2023) the gradient has vanishing correlation with new directions, thus requiring a polynomial number of updates to escape the saddles. In contrast, the large-batch gradient updates contain a finite fraction of components along the new directions, allowing their learning with a single step. Moreover, we show that each update leads to a $\mathcal{O}(1)$ change in the components along direction in V^* , obviating the need for coordinate-wise projections in Abbe et al. (2023).




Figure 2: **Learning with multiple gradient steps.** An illustration of how neural networks trained using $n = \mathcal{O}(d)$ batches learn relevant directions after multiple GD training steps, allowing to learn more complex functions over time (see Thm. 3). Here we represent the space V^* of the relevant direction of the target function $f^*(\cdot)$ and the (normalized) projection of learned direction W^t by the networks for different task $f^*(\cdot)$ - the function $g(\cdot)$ and $h(\cdot)$ are assumed to have zero first two Hermite coefficients. During the early stage of training, neural networks first learn first a single direction associated to the linear part of the target, and then can learn over time other directions that are linear conditioned on the previous learned ones. Let $\{e_i\}_{i \in [d]}$ be the standard basis of \mathbb{R}^d , the four examples show cases where: **Top left:** No directions is learned at all. **Top right:** The network can only learn a single direction e_1 (single index regime). **Bottom left:** The network learns a new direction each time, e_1 , then e_2 and finally e_3 . **Bottom right:** The network learns e_1 at the first step and both e_2 and e_3 at the second steps.

This set of results provides a mathematical theory on how neural networks learn representations over training and corroborate, among others, the findings of Kalimeris et al. (2019), who observed that neural networks first fits the best linear classifier and subsequently learns functions of increasing complexity.

From features to generalization — Our last set of results connects feature learning with generalization, and how feature learning improves the approximation capacity of the network.

- We show that a two-layer network with a finite second layer can *only* learn the component of the target function in the learned subspace (see Proposition 1). In fact, we conjecture that with p large enough (but still finite), it should be possible to approximate this part of f^* up to arbitrary precision (Conj. 1).
- As (possibly) universal kernels, large width two-layer networks at initialization enjoy of better approximation capacity. Indeed, Mei and Montanari (2022) proved that at initialization $W^{t=0}$, with $n = \Theta(d^k)$ only a degree k approximation of the target function can be learned. Our results characterizes how feature learning allows to improve over this sample complexity. In particular, we show that in the directions learned by the first layer, the target function f^* can be learned with less samples, while the component of f^* orthogonal to the learned features still requires $n = \mathcal{O}(d^k)$. While a complete mathematical control of the generalization error rates remains a difficult problem (see Conj. 2), our results provide a clear separation on scales between two-layer networks and NTK-like methods, improving over the state-of-the art in the literature.
- In particular, in Corollary 1 we prove that with a single gradient step and $n, p = \mathcal{O}(d)$, one can *only learn* features in this one-dimensional subspace, doing no better than kernels in the orthogonal direction. To do so, we provide a *conditional* version of the Gaussian equivalence theorem (Mei and Montanari, 2022; Goldt et al., 2022; Hu and Lu, 2022), which is the backbone of Thm. 4, and we believe is of independent interest.

We refer to App. A for a variety of numerical illustrations of our results. The code will be available on [GitHub](#) after de-anonymization. Mathematical proofs are detailed in Appendices B, C and D.

Further related works — The analysis of high-dimensional asymptotics of kernel regression has provided valuable insights into the advantages and limitations of kernel methods (Dietrich et al., 1999; Opper and Urbanczik, 2001; Ghorbani et al., 2019; 2020; Donhauser et al., 2021; Mei et al., 2022; Spigler et al., 2020; Bordelon et al., 2020; Canatar et al., 2021; Simon et al., 2022; Cui et al., 2021; 2022; Xiao et al., 2022). In particular, a similar stairway picture as in Fig. 2 emerged from these works (Opper and Urbanczik, 2001; Xiao et al., 2022). The key difference, however, is that at each regime $n = \mathcal{O}(d^k)$, kernels can only learn the k -th Hermite polynomial of the target. This should be contrasted with our feature learning regime where, once a direction is learned, *all* its Hermite coefficients are learned. Feature learning corrections to kernel methods have been investigated in Lewkowycz et al. (2021); Naveh and Ringel (2021); Seroussi et al. (2023); Atanasov et al. (2022); Bietti et al. (2022); Bordelon and Pehlevan (2023); Petrini et al. (2022). On a complementary line of work, exact asymptotic results for the the random features models have been derived in the literature (Mei and Montanari, 2022; Gerace et al., 2020; Dhifallah and Lu, 2020; Loureiro et al., 2021; 2022; Schröder et al., 2023; Bosch et al., 2023). A large part of these results are enabled by the Gaussian equivalence property (Goldt et al., 2022; Hu and Lu, 2022; Montanari and Saeed, 2022; Dandi et al., 2023).

Closer to us are (Ba et al., 2022; Damian et al., 2022). Ba et al. (2022) showed that a single gradient step yields an approximately rank-one change to the weights, and that a sufficiently large step allows to beat kernel methods. They did not, however, characterize the impact on the generalization error. Here, we prove that with a single gradient step and $n, p = \mathcal{O}(d)$, one can *only* learn features in this one-dimensional subspace, doing no better than kernels in the orthogonal direction. Their results are also limited to a single-index target and to a single gradient step. In contrast, Damian et al. (2022) showed that with $n = \omega(d^2)$ samples, 2-layer nets can specialize to more than one direction of a multi-index target function with zero first Hermite coefficient ($\ell=2$), but were again limited to a single step. Our work extends their sufficient conditions on the sample complexity to general $\ell \geq 1$. We also show they are necessary: $\forall \epsilon > 0$, with less data $n = \Theta(d^{\ell-\epsilon})$ one *cannot* do better than random features. Thus, our results prove a clear separation between the class of functions learned within the $\Theta(d)$ batch-size setting of Ba et al. (2022) and the $\Theta(d^2)$ batch-size setting of Damian et al. (2023) and establish a general hierarchy of functions requiring increasing batch-size to be learned with a gradient step (as well as which class of targets can be learned with *multiple* steps).

Abbe et al. (2021; 2022) show how a “staircase” property of target functions characterizes the sample complexity for infinite width two-layer networks trained by SGD, a.k.a. the mean-field regime (Mei et al., 2018; Chizat and Bach, 2018; Rotskoff and Vanden-Eijnden, 2022; Sirignano and Spiliopoulos, 2020). Their focus, however, was on the peculiar case of *sparse boolean functions* depending *solely* on linear coefficients, without considering direction-specific specialization. Abbe et al. (2023) provided further partial results for isotropic Gaussian data, always in the regime of one-pass SGD, involving basis-dependent projections. We consider the more generic case of multi-index target functions with Gaussian data and large-batch SGD, without assuming a knowledge of the basis, and show that a “directional staircase” behavior arises when iterating a few giant gradient steps, while a related, but definitely different, picture arises with single steps depending on the batch size. We also provide a sharp characterization of when this appears for multi-index targets and networks trained under large batch SGD, as well as a bound on the generalization error.

2 MAIN MATHEMATICAL RESULTS

2.1 SINGLE GRADIENT STEP

Our first goal is to determine under which conditions the relevant directions V^* of the target function f^* can be learned with a single step of gradient descent (see App. B and Assumption 3). First, we define the leap index, a fundamental quantity that precisely parametrizes what are the first directions to be learned.

Definition 1 (Leap index). *Since the input data is Gaussian $z \sim \mathcal{N}(0, I_d)$, the target function admits a decomposition in terms of the Hermite polynomials (see Definition 4 in the appendix). We define the leap index*

of the target function f^* as the first integer $\ell > 0$ such that $C_\ell^* \neq 0$:

$$\ell = \min\{j \in \mathbb{N} : \langle f^*, \mathcal{H}_j \rangle_\gamma \neq 0\} \quad (5)$$

Another crucial object in our analysis is given by the projection of the network weights in the space spanned by the target relevant directions (where Π^* is the projection on V^*):

$$\boldsymbol{\pi}_i^t = \Pi^* \boldsymbol{w}_i^t; \quad (6)$$

Our first result is of a negative nature, showing the impossibility of learning in the data-scarce regime:

Theorem 1. *Let ℓ be the leap index of f^* (1), and assume that $n = \mathcal{O}(d^{\ell-\delta})$ for some $\delta > 0$. Then, with probability at least $1 - cpe^{-c(\delta)\log(d)^2}$, there exists a universal constant c such that for any $i \in [p]$,*

$$\frac{\|\boldsymbol{\pi}_i^1\|^2}{\|\boldsymbol{w}_i^1\|^2} \leq c \frac{\sqrt{r} \text{polylog}(d)}{d^{(1 \wedge \delta)/2}}. \quad (7)$$

In other words, for every neuron i , only a vanishing fraction of the weight \boldsymbol{w}_i^1 lies in the target subspace V^* . In particular, if $\delta > 1$, this large gradient step does not improve over the initial random feature weights.

On the other hand, when $n = \Omega(d^\ell)$, we are able to characterize exactly what is being learned:

Theorem 2. *Assume that the ℓ -th Hermite coefficient μ_ℓ of σ is nonzero, and set the learning rate $\eta = pd^{(\ell-1)/2}$. Then, with probability at least $1 - ce^{-c\log(d)^2}$, there exists a random variable X independent of d with positive expectation such that*

$$\frac{\|\boldsymbol{\pi}_i^1\|^2}{\|\boldsymbol{w}_i^1\|^2} \geq X_i, \quad (8)$$

where X_1, \dots, X_p are i.i.d copies of X . Further, let $\boldsymbol{u}_1^*, \dots, \boldsymbol{u}_{r_\ell}^*$ be the higher-order singular vectors of C_ℓ^* , and define $V_\ell^* = \text{span}(\boldsymbol{u}_1^*, \dots, \boldsymbol{u}_{r_\ell}^*)$. Then, the projections $\boldsymbol{\pi}_i^1$ asymptotically belong to V_ℓ^* , in the sense that there exists a constant c such that

$$\|(I - \Pi_{V_\ell^*})\boldsymbol{\pi}_i^1\| \leq c \frac{\sqrt{r} \text{polylog}(d)}{\sqrt{d}}, \quad (9)$$

and they span the space V_ℓ^* as soon as $p \geq r$.

Note that in the case $\ell = 1$ the learned subspace V_ℓ^* is one-dimensional and is identified by the first Hermite coefficient of the target $C_1(f^*)$: this corresponds to the ‘‘linear subspace learning’’ regime exemplified in Fig. 1. Some aspects of the results above were already present in previous works, with key differences: Ba et al. (2022) shows the existence of a rank-one property of the gradient at initialization for $n = \Theta(d)$, and Damian et al. (2022) implies the positive part of our result for $n = \Theta(d^2)$, provided that $V_2^* = V^*$ (which corresponds to their Assumption 2). Our theorem allow us to obtain the matching lower bounds, demonstrating their tightness, and provides the generic picture for any higher powers of d . In particular, our results prove a clear separation between the class of functions learned within the $\Theta(d)$ batch-size setting of Ba et al. (2022) and the $\Theta(d^2)$ batch-size setting of Damian et al. (2022) and establish a general hierarchy of functions requiring increasing batch-size to be learned with a single gradient step. We refer to Appendix C for the proofs of the theorems and a more detailed technical discussion.

2.2 LEARNING WITH MANY STEPS, AND THE STAIRCASE PHENOMENON

For simplicity, we restrict ourselves to the case $n = \mathcal{O}(d)$, although we expect the findings to hold in the other regimes. We unveil the intertwined dependence between representation learning efficiency and a ‘‘directional staircase’’ condition. Informally: once a direction is learned by the first-layer weights, the next gradient step uses it as a ladder to learn the next ones, on the assumption that they are linearly connected. We formalize the hierarchical learning by introducing the notion of *subspace conditioning*:

Definition 2 (Subspace conditioning). *Let V be a vector space, and $U \subseteq V$ a subspace. For any function $f : V \rightarrow \mathbb{R}$, and $\mathbf{x} \in U$, we define the conditional function $f_{U,\mathbf{x}} : U^\perp \rightarrow \mathbb{R}$ as*

$$f_{U,\mathbf{x}}(\mathbf{x}^\perp) = f(\mathbf{x} + \mathbf{x}^\perp) \quad (10)$$

In short, the function $f_{U,\mathbf{x}}$ corresponds to f “conditioned” on the projection of its argument in U . Its first Hermite coefficient will be denoted as

$$\mu_{U,\mathbf{x}}(f) = \mathbb{E}_{\mathbf{x}^\perp \sim \mathcal{N}(0,I)} [\nabla_{\mathbf{x}^\perp} f_{U,\mathbf{x}}] \quad (11)$$

We are now equipped to state our main result describing sequential learning of directions under multiple gradient steps with batch size $n = \Theta(d)$. We denote by $W^* \in \mathbb{R}^{r \times d}$ the matrix with rows $\mathbf{w}_1^*, \dots, \mathbf{w}_r^*$. To avoid degeneracy issues, we restrict ourselves to polynomial activations and target functions.

Assumption 1. *Both the activation $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ and the target function $g^* : \mathbb{R}^r \rightarrow \mathbb{R}$ are fixed polynomials with degrees independent of d, n , such that $\deg(\sigma) \geq \deg(g^*)$.*

Theorem 3. *Assume that $n = \Theta(d)$, and η, p are fixed. Define a sequence of nested subspaces $U_0^* \subseteq U_1^* \subseteq \dots \subseteq U_t^* \subseteq \dots$ as a) $U_0^* = \{0\}$, and b) for any $t \geq 0$, $U_{t+1}^* = U_t^* \oplus \text{span}(\{\mu_{U_t^*,\mathbf{x}}(f^*) : \mathbf{x} \in U_t^*\})$.*

Then, under assumptions 3 and 1, after t gradient steps of the form (4), W^t satisfies the following with high-probability, for all $i \in [p]$:

(i) *Let $\mathbf{v} \in U_t^*$, with $\|\mathbf{v}\| = 1$ be arbitrary. Let $\mathbf{v}^m = W^* \mathbf{v}$ denote the components of \mathbf{v} along the target basis. Then there exists an almost surely positive random variable $X_{t,\mathbf{v},\mathbf{a}}$, independent of d, n such that $|\langle \mathbf{w}_i, \mathbf{v} \rangle| \geq X_{t,\mathbf{v},\mathbf{a}}$.*

(ii) *For any $\mathbf{v} \in U_t^{*\perp} \cap V^*$, $|\langle \mathbf{w}_i, \mathbf{v} \rangle| = O(\frac{\text{polylog}(d)}{\sqrt{d}})$.*

Informally, the above statements imply that almost surely over \mathbf{a} , the first layer learns exactly U_t^ .*

This is perfectly consistent with Thms. 1 and 2 since $f_{\{0\},\mathbf{0}} = f$, and \mathbf{v}^* is exactly the first Hermite coefficient of f^* . As an example, below we explicitly work out the subspaces U_t^* for two target functions:

• $f^*(\mathbf{z}) = z_1 + z_2 + z_1^2 - z_2^2$: the normalized first Hermite coefficient of f^* is $\mathbf{v}^* = (\mathbf{e}_2 + \mathbf{e}_1)/\sqrt{2}$, so $U_1^* = \text{span}(\mathbf{e}_1 + \mathbf{e}_2)$. Then, we can rewrite f^* in the new basis $(\mathbf{v}^*, \mathbf{v}^\perp)$, to get

$$f^*(\mathbf{z}) = \sqrt{2} \langle \mathbf{v}^*, \mathbf{z} \rangle + 2 \langle \mathbf{v}^*, \mathbf{z} \rangle \langle \mathbf{v}^\perp, \mathbf{z} \rangle.$$

Hence, if $\mathbf{x} = \lambda \mathbf{v}^*$, we have $\mu_{U_1^*,\mathbf{x}}(f^*) = 2\lambda \mathbf{v}^\perp$, and hence $U_2^* = V^*$.

• $f^*(\mathbf{z}) = z_1 + z_2 + z_1^2 + z_2^2$: just as the above example, $U_1^* = \text{span}(\mathbf{e}_1 + \mathbf{e}_2)$, and hence we perform the same change of basis. However, this time,

$$f^*(\mathbf{z}) = \sqrt{2} \langle \mathbf{v}^*, \mathbf{z} \rangle + \langle \mathbf{v}^*, \mathbf{z} \rangle^2 + \langle \mathbf{v}^\perp, \mathbf{z} \rangle^2.$$

This implies that for any $\mathbf{x} \in U_1^*$, $\mu_{U_1^*,\mathbf{x}}(f^*) = \mathbf{0}$, and the direction of \mathbf{v}^\perp is never learned.

We provide a proof of Theorem 3 in App. C. The notion of subspace conditioning (Def. 2) arises through an inductive decomposition of the projections of the gradient along the target subspace. Our theory leads to a precise prediction of the orientation of the gradient in the target subspace after a finite number of steps (see Fig. 6 in App. A). The non-linearity of the activation function allows different neurons to simultaneously specialize along different directions in contrast to the rank-one increase at each saddle in deep linear networks (Jacot et al., 2021) under vanishing initialization. We refer to App. A for additional discussion.

This picture is akin to the work of Abbe et al. (2023) for online SGD. However, their analysis relies on coordinate-dependent projections, which requires that V^* is aligned with the canonical basis. In contrast, our proof as well as our Definition 2 are both basis independent.

2.3 FROM FEATURE LEARNING TO GENERALIZATION BOUNDS

We now investigate the consequence of Thm. 2 to the capacity of the network to generalize. At a high-level, one expects that as the subspace U is learned by the first-layer, the network effectively reduces the input dimension to that of U , allowing it to learn functions of U with significantly fewer examples and hidden-units. On the other hand, one expects that learning functions depending on directions which were not learned require as many samples and hidden-units as with random first-layer weights. While a complete mathematical control of generalization rates is difficult, we shall make this intuition precise by proving that the generalization error of training the second layer on the top of the learned features is lower bounded precisely by the directions which were not learned by the first-layer. Our results incorporate the dependence on both the number of samples and the number of hidden-units.

The simplest case is when the **width p is finite**: then there is simply no way to fit anything beyond the learned directions, even with an infinite amount of data. This is formalized by the following proposition:

Proposition 1. *Assume that p is bounded as $n, d \rightarrow \infty$, and that the first layer W only learns a subspace $U \subseteq V^*$. For any choice of second layer \mathbf{a} such that $\|\mathbf{a}\|_\infty = \mathcal{O}(1)$, we have*

$$\mathbb{E} \left[\left(f^*(\mathbf{z}) - \hat{f}(\mathbf{z}; W, \mathbf{a}) \right)^2 \right] \geq \mathbb{E} [\text{Var}(f^*(\mathbf{z}) \mid P_U \mathbf{z})] - o(1) \quad (12)$$

where P_U is the orthogonal projection on U .

We refer to App. D for the proof of the proposition. Intuitively, the RHS of the above inequality corresponds to the predictor $\hat{f}(\mathbf{z}) = \mathbb{E}[f^*(\mathbf{z}) \mid P_U \mathbf{z}]$, which is the best predictor depending only on $P_U \mathbf{z}$. Conversely, we expect that with enough neurons, it is possible to approximate this predictor:

Conjecture 1. *For any $\delta > 0$, there exists a $p_0 \in \mathbb{N}$ such that if $p \geq p_0$, there exists a choice of second layer \mathbf{a} satisfying $\|\mathbf{a}\|_\infty \leq \mathcal{O}(1)$, such that as $n, p \rightarrow \infty$*

$$\mathbb{E} \left[\left(f^*(\mathbf{z}) - \hat{f}(\mathbf{z}; W, \mathbf{a}) \right)^2 \right] \leq \mathbb{E} [\text{Var}(f^*(\mathbf{z}) \mid P_U \mathbf{z})] + \delta + o(1) \quad (13)$$

Proving the above conjecture for generic target and activation functions requires proving a sufficient spread of W along the learned subspace U together with an approximation result for finite-dimensional neural networks having activation function σ . While this is done in Damian et al. (2022) for the *specific* case of $U = V^*$ (in which case the first term in the RHS is zero) and $\sigma = \text{relu}$ (with an additional bias term) such approximation results are scarcer for generic σ . In particular, since the projection of the weights W along U cannot be chosen arbitrarily, one cannot directly invoke classical approximation results such as the ones based on Barron spaces. However, the components of W along U are still randomly distributed across neurons. One possible (but delicate) approach could be to use results on the generalization of random feature models on finite-dimensional inputs such as Rudi and Rosasco (2017). We leave this for future work.

The situation is more complex when **the number of neurons p is allowed to diverge**. Indeed, Mei et al. (2022) shows that even in the absence of feature learning (i.e. when the first layer is random), it is possible to learn a polynomial approximation of f^* of degree k as long as $n, p = \Omega(d^{k+\delta})$ for some $\delta > 0$. Intuitively, we expect the following behavior: (i) On the directions present in W , with enough variety in the neurons, the model behaves similarly to a random features model in this *finite*-dimensional space, and is thus able to learn “everything”. (ii) On the orthogonal directions, however, the model still behave as a high-dimensional random feature model, and the polynomial limitations of Mei et al. (2022) still apply. To formalize this conjecture, we introduce the following definition:

Definition 3. *Let V be a vector space, and $U \subseteq V$ a subspace. For any $k \geq 0$, we define the space $\mathcal{P}_{U,k}$ of functions $f : V \rightarrow \mathbb{R}$ such that for any $\mathbf{x} \in U$, the function $f_{U,\mathbf{x}}$ introduced in Definition 2 is a polynomial of degree at most k .*

Simply put, the space $\mathcal{P}_{U,k}$ consists of polynomials in \mathbf{x}^\perp of degree at most k , whose coefficients can be functions of \mathbf{x} . We will denote by $P_{U,\leq k}$ and $P_{U,>k}$ the projections on $\mathcal{P}_{U,k}$ and $\mathcal{P}_{U,k}^\perp$ in $\ell^2(\mathbb{R}, \gamma)$, respectively, where γ is the Gaussian measure. This allows us to formalize the above intuition precisely as:

Conjecture 2. *Assume that $p = \mathcal{O}(d^{\kappa_1})$ and $n = \mathcal{O}(d^{\kappa_2})$, and that the first layer W only learns a subspace $U^* \subseteq V^*$. Then, if $\hat{\mathbf{a}}$ is obtained as in Eq. (29) for any value of λ ,*

$$\mathbb{E} \left[\left(f^*(\mathbf{z}) - \hat{f}(\mathbf{z}; W, \hat{\mathbf{a}}) \right)^2 \right] \geq \|P_{U^*,>\kappa} f^*\|_\gamma^2 - o(1), \quad (14)$$

where $\kappa = \min(\kappa_1, \kappa_2)$ and $\|\cdot\|_\gamma$ is the norm in $\ell^2(\mathbb{R}, \gamma)$.

While a general proof for this conjecture would require significant additional work we can readily prove part of it. First the case $\kappa = \kappa_1 = 0$ corresponds to our rigorous Prop. 1. Second, we now prove the already fairly delicate $\kappa_1 = \kappa_2$ case, also known as the *proportional regime*.

(Mei and Montanari, 2022; Gerace et al., 2020; Goldt et al., 2022; Hu and Lu, 2022) have shown that for the random features W^0 case, the generalization error of the minimizer for the second layer (29) is equivalent to the generalization error of an equivalent linear model; this result was extended in Ba et al. (2022) to a trained first layer W^1 for small, “lazy” stepsizes $\eta = \mathcal{O}(1)$. Ba et al. (2022) has also shown that when $\eta = \Theta(p)$, the equivalent linear description breaks down, due to the appearance of a spike term proportional to η in the trained weight matrix W^1 . We now show that this spike *only allows* the learning of non-linear functions *parallel to the spike*. Notice that the optimization problem (29) is equivalent to

$$\min_{\mathbf{a} \in \mathbb{R}^p} \frac{1}{n} \sum_{\nu=1}^n \left(\langle \mathbf{a}, \phi_{\text{CK}}(\mathbf{z}^\nu) \rangle - f^*(\mathbf{z}^\nu) \right)^2 + \lambda \|\mathbf{a}\|^2 \quad (15)$$

where $\phi_{\text{CK}}(\mathbf{z}) = \sigma(W\mathbf{z})$ is the feature map corresponding to the *conjugate kernel* of the neural network. For a given direction \mathbf{v} , we define the *conditional linear* equivalent map $\phi_{\text{CL}}(\mathbf{z}; \mathbf{v})$ as follows: given the decomposition $\mathbf{z} = z_{\mathbf{v}}\mathbf{v} + \mathbf{z}^\perp$,

$$\phi_{\text{CL}}(\mathbf{z}; \mathbf{v}) = \mu(z_{\mathbf{v}}) + \Psi(z_{\mathbf{v}})\mathbf{z}^\perp + \Phi(z_{\mathbf{v}})\boldsymbol{\xi} \quad (16)$$

where $\boldsymbol{\xi} \sim \mathcal{N}(0, I_p - \mathbf{v}\mathbf{v}^\top)$, and $\mu \in \mathbb{R}^p$, $\Psi \in \mathbb{R}^{p \times d}$, $\Phi \in \mathbb{R}^{p \times p}$ are chosen to match the first two conditional moments of ϕ_{CK} :

$$\mu(z_{\mathbf{v}}) = \mathbb{E}[\phi_{\text{CK}}(\mathbf{z}) | z_{\mathbf{v}}], \quad \Psi(z_{\mathbf{v}}) = \mathbb{E}[\phi_{\text{CK}}(\mathbf{z})(\mathbf{z}^\perp)^\top | z_{\mathbf{v}}], \quad \Phi(z_{\mathbf{v}}) = \text{Cov}[\phi_{\text{CK}}(\mathbf{z}) | z_{\mathbf{v}}] - \Psi(z_{\mathbf{v}})\Psi(z_{\mathbf{v}})^\top$$

Then, the following theorem holds:

Theorem 4. (informal) *Assume that $n, p = \Theta(d)$, and the vector $\mathbf{v}^* = V_1^*$ defined in Theorem 2 to be nonzero. Then, after one gradient step, there exists a vector $\mathbf{v} \in \mathbb{R}^d$ such that: (i) the projection of \mathbf{v} on V^* is proportional to \mathbf{v}^* , (ii) the solutions $\hat{\mathbf{a}}, \tilde{\mathbf{a}}$ to the optimization problem in (15) with the feature maps $\phi_{\text{CK}}(\mathbf{z})$ and $\phi_{\text{CL}}(\mathbf{z}; \mathbf{v})$ yield the same generalization error.*

The full statement of Theorem 4 with the relevant hypotheses is discussed App. D. In particular, as we show in App. D, Theorem 4 provides a proof of the case $\kappa_1 = \kappa_2 = 1$ of Conjecture 2:

Corollary 1. *Conjecture 2 holds for $\kappa_1 = \kappa_2 = 1$.*

This concludes the mathematical presentation of our results.

REFERENCES

- E. Abbe, E. Boix-Adsera, M. S. Brennan, G. Bresler, and D. Nagaraj. The staircase property: How hierarchical structure can guide deep learning. *Advances in Neural Information Processing Systems*, 34:26989–27002, 2021.
- E. Abbe, E. B. Adsera, and T. 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*, pages 4782–4887. PMLR, 2022.
- E. Abbe, E. Boix-Adsera, and T. Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics, 2023.
- L. Arnaboldi, L. Stephan, F. Krzakala, and B. Loureiro. From high-dimensional & mean-field dynamics to dimensionless ODEs: A unifying approach to SGD in two-layers networks, Feb. 2023. arXiv:2302.05882 [cond-mat, stat] type: article.
- A. Atanasov, B. Bordelon, and C. Pehlevan. Neural networks as kernel learners: The silent alignment effect. In *International Conference on Learning Representations*, 2022.
- J. Ba, M. A. Erdogdu, T. Suzuki, Z. Wang, D. Wu, and G. Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 37932–37946. Curran Associates, Inc., 2022.
- F. Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- G. Ben Arous, R. Gheissari, and A. Jagannath. Online stochastic gradient descent on non-convex losses from high-dimensional inference. *Journal of Machine Learning Research*, 22(106):1–51, 2021.
- R. Berthier, A. Montanari, and K. Zhou. Learning time-scales in two-layers neural networks, 2023.
- A. Bietti, J. Bruna, C. Sanford, and M. J. Song. Learning single-index models with shallow neural networks. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 9768–9783. Curran Associates, Inc., 2022.
- B. Bordelon and C. Pehlevan. Dynamics of finite width kernel and prediction fluctuations in mean field neural networks, 2023.
- B. Bordelon, A. Canatar, and C. Pehlevan. Spectrum dependent learning curves in kernel regression and wide neural networks. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1024–1034. PMLR, 13–18 Jul 2020.
- D. Bosch, A. Panahi, and B. Hassibi. Precise asymptotic analysis of deep random feature models, 2023.
- S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, June 2006. ISSN 1557-9654. doi: 10.1109/TIT.2006.874516.
- A. Canatar, B. Bordelon, and C. Pehlevan. Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks. *Nature Communications*, 12(1):2914, May 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-23103-1.
- L. Chizat and F. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.

- L. Chizat, E. Oyallon, and F. Bach. On Lazy Training in Differentiable Programming. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- H. Cui, B. Loureiro, F. Krzakala, and L. Zdeborová. Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 10131–10143. Curran Associates, Inc., 2021.
- H. Cui, B. Loureiro, F. Krzakala, and L. Zdeborová. Error rates for kernel classification under source and capacity conditions, 2022.
- A. Damian, J. Lee, and M. Soltanolkotabi. Neural networks can learn representations with gradient descent. In P.-L. Loh and M. Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 5413–5452. PMLR, 02–05 Jul 2022.
- A. Damian, E. Nichani, R. Ge, and J. D. Lee. Smoothing the Landscape Boosts the Signal for SGD: Optimal Sample Complexity for Learning Single Index Models. Technical report, May 2023. arXiv:2305.10633 [cs, math, stat] type: article.
- Y. Dandi, L. Stephan, F. Krzakala, B. Loureiro, and L. Zdeborová. Universality laws for gaussian mixtures in generalized linear models, 2023.
- L. Devroye, L. Györfi, and G. Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- O. Dhifallah and Y. M. Lu. A precise performance analysis of learning with random features, 2020.
- R. Dietrich, M. Opper, and H. Sompolinsky. Statistical mechanics of support vector networks. *Phys. Rev. Lett.*, 82:2975–2978, Apr 1999. doi: 10.1103/PhysRevLett.82.2975.
- K. Donhauser, M. Wu, and F. Yang. How rotational invariance of common kernels prevents generalization in high dimensions. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 2804–2814. PMLR, 18–24 Jul 2021.
- F. Gerace, B. Loureiro, F. Krzakala, M. Mezard, and L. Zdeborova. Generalisation error in learning with random features and the hidden manifold model. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3452–3462. PMLR, 13–18 Jul 2020.
- B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari. Limitations of lazy training of two-layers neural network. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari. When do neural networks outperform kernel methods? In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 14820–14830. Curran Associates, Inc., 2020.
- S. Goldt, B. Loureiro, G. Reeves, F. Krzakala, M. Mezard, and L. Zdeborova. The gaussian equivalence of generative models for learning with shallow neural networks. In J. Bruna, J. Hesthaven, and L. Zdeborova, editors, *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, volume 145 of *Proceedings of Machine Learning Research*, pages 426–471. PMLR, 16–19 Aug 2022.
- F. Gotze, H. Sambale, and A. Sinulis. Concentration inequalities for polynomials in α -sub-exponential random variables. *Electronic Journal of Probability*, 2019.

- P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- H. Grad. Note on N-dimensional hermite polynomials. *Communications on Pure and Applied Mathematics*, 2(4):325–330, 1949. ISSN 1097-0312. doi: 10.1002/cpa.3160020402.
- W. Greub. *Multilinear Algebra*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2012. ISBN 9783662007952.
- H. Hu and Y. M. Lu. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, 2022.
- A. Jacot, F. Ged, B. Şimşek, C. Hongler, and F. Gabriel. Saddle-to-saddle dynamics in deep linear networks: Small initialization training, symmetry, and sparsity. *arXiv preprint arXiv:2106.15933*, 2021.
- D. Kalimeris, G. Kaplun, P. Nakkiran, B. Edelman, T. Yang, B. Barak, and H. Zhang. Sgd on neural networks learns functions of increasing complexity. *Advances in neural information processing systems*, 32, 2019.
- A. Klenke. *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.
- M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer-Verlag, 1991. ISBN 9780387520131. Google-Books-ID: juC1QgAACAAJ.
- A. Lewkowycz, Y. Bahri, E. Dyer, J. Sohl-Dickstein, and G. Gur-Ari. The large learning rate phase of deep learning, 2021.
- L. Li, Y. Fan, M. Tse, and K.-Y. Lin. A review of applications in federated learning. *Computers & Industrial Engineering*, 149:106854, 2020.
- B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mezard, and L. Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 18137–18151. Curran Associates, Inc., 2021.
- B. Loureiro, C. Gerbelot, M. Refinetti, G. Sicuro, and F. Krzakala. Fluctuations, bias, variance & ensemble of learners: Exact asymptotics for convex losses in high-dimension. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 14283–14314. PMLR, 17–23 Jul 2022.
- S. Mei and A. Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766, 2022.
- S. Mei, A. Montanari, and P.-M. Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- S. Mei, T. Misiakiewicz, and A. Montanari. Generalization error of random feature and kernel methods: Hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 59: 3–84, 2022. ISSN 1063-5203. doi: <https://doi.org/10.1016/j.acha.2021.12.003>. Special Issue on Harmonic Analysis and Machine Learning.
- B. S. Mityagin. The zero set of a real analytic function. *Mathematical Notes*, 107(3-4):529–530, 2020.
- A. Montanari and B. N. Saeed. Universality of empirical risk minimization. In P.-L. Loh and M. Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 4310–4312. PMLR, 02–05 Jul 2022.

- G. Naveh and Z. Ringel. A self consistent theory of gaussian processes captures feature learning effects in finite cnns. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 21352–21364. Curran Associates, Inc., 2021.
- M. Opper and R. Urbanczik. Universal learning curves of support vector machines. *Phys. Rev. Lett.*, 86: 4410–4413, May 2001. doi: 10.1103/PhysRevLett.86.4410.
- V. H. d. l. Pena and S. J. Montgomery-Smith. Decoupling Inequalities for the Tail Probabilities of Multivariate \mathbb{U}^d -Statistics. *The Annals of Probability*, 23(2):806–816, Apr. 1995. ISSN 0091-1798, 2168-894X. doi: 10.1214/aop/1176988291.
- L. Petrini, F. Cagnetta, E. Vanden-Eijnden, and M. Wyart. Learning sparse features can lead to overfitting in neural networks, 2022.
- G. Rotskoff and E. Vanden-Eijnden. Trainability and accuracy of artificial neural networks: An interacting particle system approach. *Communications on Pure and Applied Mathematics*, 75(9):1889–1935, 2022. doi: <https://doi.org/10.1002/cpa.22074>.
- A. Rudi and L. Rosasco. Generalization properties of learning with random features. *Advances in neural information processing systems*, 30, 2017.
- D. Schröder, H. Cui, D. Dmitriev, and B. Loureiro. Deterministic equivalent and error universality of deep random features learning, 2023.
- I. Seroussi, G. Naveh, and Z. Ringel. Separation of scales and a thermodynamic description of feature learning in some cnns. *Nature Communications*, 14(1):908, Feb 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-36361-y.
- J. B. Simon, M. Dickens, D. Karkada, and M. R. DeWeese. The eigenlearning framework: A conservation law perspective on kernel regression and wide neural networks, 2022.
- J. Sirignano and K. Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020.
- S. Spigler, M. Geiger, and M. Wyart. Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124001, dec 2020. doi: 10.1088/1742-5468/abc61d.
- A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Science & Business Media, Mar. 1996. ISBN 9780387946405. Google-Books-ID: OCenCW9qmp4C.
- R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- L. Xiao, H. Hu, T. Misiakiewicz, Y. Lu, and J. Pennington. Precise learning curves and higher-order scalings for dot-product kernel regression. *Advances in Neural Information Processing Systems*, 35:4558–4570, 2022.

A NUMERICAL INVESTIGATION

In this section, we provide many plots corroborating the theoretical results presented in the main manuscript. We also give information on the numerical procedures needed to reproduce the different figures, along with the details behind the numerical experiments.

The code will be available on [GitHub](#) after de-anonymization.

Description of the training algorithm: First, we describe the training protocol reported in Alg. 1: we separately update the first layer with T -GD steps of learning rate η , followed by training with standard ridge regression for the second layer with fixed regularization strength λ . We vary adaptively the learning rate to satisfy the hypothesis of Thm. 2, i.e. $\eta = \mathcal{O}(p\sqrt{\frac{\eta}{d}})$, the regularization parameter is fixed to $\lambda = 1$, and we take noiseless labels. We average over different seeds to get the mean performance, and we use standard deviation for giving confidence intervals.

A.1 LEARNING WITH A SINGLE GIANT STEP

A sizable part of our results concerns the feature learning efficiency of two layer neural networks after one giant step of GD. We provide a toy illustration of the phenomenology in Fig. 1, and rigorously characterize this in Theorems 1 and 2. Moreover, in section 2.3 we provide a plethora of results analyzing the consequences of the theorems above in the actual learning performance of the network. Here, we perform a detailed numerical investigation of the different claims in these results.

Investigating the generalization performance: We start by analyzing the generalization performance of different networks after one giant GD step in Fig. 3. We compare the generalization performances of linear and quadratic kernel methods - horizontal lines marked by different colors computed at $n = n_{max} \sim d^{2.25}$ - with three networks: a) *random features* (red points) random network with fixed weight matrix W_0 at initialization; b) *1 GD step* (green points) two-layer network trained using one step in the protocol of Alg. 1; c) *1 GD step with preprocessing* (blue points) two-layer network trained using a preprocessing step in Alg. 1. The introduction of a preprocessed algorithm is linked with the theoretical results of Theorem 2 and we provide a detailed analysis in the next paragraph.

The importance of preprocessing: As Theorem. 2 provably states, it is not possible to get fully specialized hidden units with one giant step of GD in the $n = \mathcal{O}(d^k)$ regime (with $k > 1$), unless the directions associated to teacher Hermite coefficients lower than k are suppressed, or equivalently, if the leap index ℓ (Def. 1) of the target is equal to k - see Fig. 1 for an illustration. We can circumvent this issue by using a preprocessing step. Given a batch of size $n = \mathcal{O}(d^k)$, we preprocess the labels in Alg. 1 using a method introduced in [Damian et al. \(2022\)](#) for the case $k = 1$:

$$\hat{c}_{j_1, \dots, j_d} \leftarrow \frac{1}{n} \sum_{\nu=1}^n y_\nu \text{He}_{j_1}(\langle e_1, z_\nu \rangle) \cdots \text{He}_{j_d}(\langle e_d, z_\nu \rangle) \quad (17)$$

$$y_\nu \leftarrow y_\nu - \sum_{j_1, \dots, j_d: j_1 + \dots + j_d < k} \frac{\hat{c}_{j_1, \dots, j_d}}{j_1! \cdots j_d!} \text{He}_{j_1}(\langle e_1, z_\nu \rangle) \cdots \text{He}_{j_d}(\langle e_d, z_\nu \rangle) \quad (18)$$

where we denoted with $\hat{c}_{j_1, \dots, j_d}$ the plug-in estimates from data of the teacher Hermite coefficients, and with $\{e_i\}_{i \in [d]}$ the canonical basis in \mathbb{R}^d . By standard concentration arguments [Gotze et al. \(2019\)](#) the plug-in estimation of the coefficients is accurate only in the $n = \omega(d \text{polylog}(d))$ regime. Indeed, in Fig. 3 the inefficient estimation of eq. (17) in the $n = o(d)$ sample regime generates a noisy learning curve for the preprocessed algorithm (blue points). The ridge estimator \hat{a} is consequently found by training on the processed labels defined in eq. (18) and the suppressed part is injected back in the predictor only at test




Figure 3: Learning with training of the second layer. Simulation illustrating the different regimes in Fig. 1, using $d=512, p=1024$, a symmetric two-index target function $f^*(z) = \sigma^*(\langle w_1^*, z \rangle) + \sigma^*(\langle w_2^*, z \rangle)$ with activation $\sigma^*(z) = \text{He}_1(x) + \text{He}_2(x)/2! + \text{He}_4(x)/4!$, and a relu student. (a) The first algorithm (green) applies a giant step and then learns the second layer. When $n \gg d$, its performance goes beyond the linear predictor that would be obtained with a kernel method and reach the “linear subspace learning” regime in Fig. 1. (b) To go beyond this regime, the second algorithm (blue) preprocesses the data to remove a plug-in estimate of the first Hermite coefficient. It reaches a lower plateau as $n \approx d^2$, now beating the quadratic kernel. We contrast this behavior with the one of the random feature model (red).




Figure 4: Learning as a function of the number of hidden neurons. Simulations illustrating the influence of the number of hidden neurons in Fig. 3. We change the value of $p \in (128, 256, 512, 1024)$ from left to right.

time:

$$\hat{f}(z_\nu) = \frac{1}{\sqrt{p}} \hat{a}^\top \sigma(W z_\nu) + \sum_{j_1, \dots, j_d: j_1 + \dots + j_d < k} \frac{\hat{c}_{j_1, \dots, j_d}}{j_1! \dots j_d!} \text{He}_{j_1}(\langle e_1, z_\nu \rangle) \dots \text{He}_{j_d}(\langle e_d, z_\nu \rangle) \quad (19)$$

Comparison of different methods: The results presented in Fig. 3 clearly illustrate the theoretical predictions of Thm. 2: in the $n = \mathcal{O}(d)$ regime vanilla Alg. 1 attains the “linear subspace learning” regime (see Fig. 1) and beats the linear kernel, while the preprocessed version cannot. However, implementing preprocessing turns out definitely beneficial in the $n = \mathcal{O}(d^2)$ region. Indeed, while the vanilla Alg. 1 remains stuck on the linear subspace learning plateau, the preprocessed Alg. 1 reaches a lower test error than the quadratic kernel. This is achieved by effectively raising the leap index of the target function. More precisely, given a target with leap index $\ell = 1$ as in Fig. 3, the manipulation in eq. (18) aims exactly at the removal of the first Hermite coefficient of the target by estimating it from the data, allowing feature learning in the $n = \mathcal{O}(d^2)$ regime in accordance with Thm. 2. We complement the above picture by analyzing the influence of the number of hidden neurons p on the generalization performance in Fig. 4: by increasing the expressive power of the network, we attain the single-index regime by using a single giant step of Alg. 1 (in accordance with Conj. 1). Moreover, we note that it is necessary to use $p = 2d$ in order to be able to beat the performance of the quadratic kernel in this learning task (rightmost section).

Investigating representation learning efficiency: We move to an additional numerical investigation of feature learning efficiency, as characterized by Theorems 1 and 2. In Fig. 5 we again consider a single

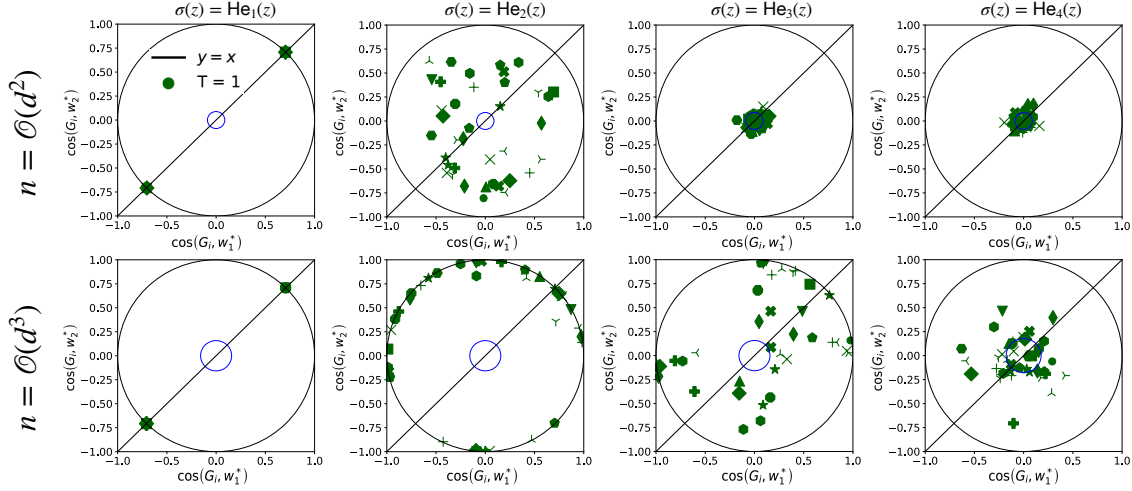


Figure 5: **Feature learning after a single step.** Specialization of hidden units in the $n = \mathcal{O}(d^k)$ regime ($k = 2, 3$). The plots show the cosine similarity of the gradient with respect to the target vectors (w_1^*, w_2^*) for $p = 40$ different neurons, identified by different markers. The bisectrix of the first quadrant is shown as a continuous black line, the circle of unitary radius in black, and the circle of radius $2/\sqrt{2}$ in blue. In the upper panel, $(n, d) = (2^{18}, 2^9)$, and $(n, d) = (2^{21}, 2^7)$ in the lower one.

We use a 2-index target $f^*(z) = \sigma^*(\langle w_1^*, z \rangle) + \sigma^*(\langle w_2^*, z \rangle)$, with matching student: $\sigma(z) = \sigma^*(z)$. **Left:** $\sigma(z) = \text{He}_1(z)$. **Center-Left:** $\sigma(z) = \text{He}_2(z)$. **Center-Right:** $\sigma(z) = \text{He}_3(z)$. **Right:** $\sigma(z) = \text{He}_4(z)$. We observe that if the leap index $\ell = 1$, we only learn a single direction, no matter the data quantity, while for $\ell > 1$ we learn every direction as soon as we reach $n = \mathcal{O}(d^\ell)$. The small spread observed for $\sigma(z) = \text{He}_4(z)$ and $n = \mathcal{O}(d^3)$ is due to the small value of d used for the experiments.

step of Alg. 1, focusing now on the analysis of the gradient. We compute the gradient matrix $G \in \mathbb{R}^{p \times d}$ and plot the cosine similarities of all the rows $\{G_i \in \mathbb{R}^d\}_{i=1}^p$ with the teacher vectors (w_1^*, w_2^*) . The figure clearly illustrates the claims of Thm. 2: in the $n = \mathcal{O}(d^k)$ (with $k > 1$) regime is necessary to analyze targets with leap index k in order to obtain specialized hidden units. Moreover, the leftmost section of Fig. 5 completes the picture offered by Figs. 3&4 about the lack of specialization in presence of teacher functions with non-zero first Hermite coefficient ($\ell = 1$): the gradient is stuck in the linear subspace learning regime theoretically predicted by Thm. 2, regardless of the sample regime considered, preventing feature learning.

A.2 LEARNING WITH MULTIPLE STEPS

We move now the numerical investigation of the learning behavior after multiple gradient steps. The general picture of the phenomenology is offered in Fig. 2, following the theoretical characterization of Theorem 3.

Investigating the generalization performance: First, we investigate the generalization behavior in the upper panel of Fig. 6. We modify slightly the training procedure in Alg. 1 to perform the numerical experiments: at every gradient step on the first layer weights we train the second layer sequentially with ridge regression. The analysis of the test error behavior in the upper panel of Fig. 6 sheds light on the consequences of Thm. 3 on the generalization performance of two-layer networks. Indeed, we observe a clear benefit in performing multiple gradient steps if the teacher function has a direction linearly connected to the rank-one spike in the gradient identified by $C_1(f^*)$ (right panel), while if such linearly connected direction does not exist (left panel) the generalization performance does not improve relevantly over time,




Figure 6: **Feature learning with multiple gradient steps.** **Top:** Generalization error as a function of n ($d = 512, p = 256$) after iterating the training procedure for six steps. **Bottom:** Cosine similarity of the projected gradient matrix G^p inside the target subspace for all the p neurons at a fixed ratio $n/d = 4$, plotted at different stages of the training. The blue and purple lines are the theoretical predictions for the orientation of the gradient in the second step. We fix a relu student and consider two different 2-index target functions $f^*(z) = \sigma_1^*(\langle w_1^*, z \rangle) + \sigma_2^*(\langle w_2^*, z \rangle)$. **Left:** $\sigma_1^*(z) = \sigma_2^*(z) = \text{He}_1(z) + \text{He}_2(z)/2 + \text{He}_4(z)/4!$. **Right:** $\sigma_1^*(z) = z - z^2$ and $\sigma_2^*(z) = z + z^2$. In accordance with Theorem 3, the difference between the two cases is clear already after the first GD step: while on the left the gradient is stuck around the predicted rank-one spike after the first step (black line), on the right the gradient changes orientation in the second step, allowing to learn multiple features.

and the network is stuck on the “linear subspace learning” (see the upper right plot of Fig. 2). These results are in perfect agreement with Thm. 3.

Investigating representation learning efficiency: In this paragraph we further analyze the claims of Thm. 3 in the context of feature learning. The experiments done in the lower panel of Fig. 6 are closely related to the ones of Fig. 5. However, contrary to the previous setting, we study the cosine similarity of the *projected gradient* $G^p = G\Pi^*$ in the teacher subspace. This quantity differs from the cosine similarity of the full gradient, plotted in Fig. 5, as we lose completely the information about the share of the gradient lying in the subspace orthogonal to the teacher one. This divergence in the choices is due to the different illustrative goals of the figures: while in Fig. 6 we highlight the change in orientation of the gradient inside the teacher subspace after a few steps, hence not caring about the relative magnitude, in Fig. 5 we contrast the magnitude of the true gradient with the one of a random object (blue circles) in order to claim the presence (or lack) of feature learning after a single step. The results in the lower panel of Fig. 6 are obtained iterating 2 steps of the training procedure in Alg. 1: in accordance with Thm. 3 we observe delocalization of the projected gradient only if there are linearly connected directions that can be exploited to escape the spike given by the first Hermite coefficient $C_1(f^*)$ (right panel). Moreover, we are able to theoretically predict the orientation of the gradient at the second step as well (see Appendix. C). On the contrary, when such linearly connected directions do not exist, the gradient is stuck on the spike $C_1(f^*)$ (Left panel). We elaborate on this last observation by checking that the lack of specialization persists iterating for more than two GD steps. We present the results in Fig. 7: the gradient is stuck in the linear subspace learning regime even as the training proceeds, again in agreement with Thm. 3. Moreover, we illustrate by changing the




Figure 7: **Lack of feature learning after few GD steps.** The plots show the cosine similarity with respect to the teacher vectors (w_1^*, w_2^*) for the gradient at different stages of the training. The predicted orientation (Thm. 2) of the gradient is shown as a continuous black line, the circle of unitary radius in black, and the circle of radius $2/\sqrt{a}$ in blue. We fix $n = d = 2^{13}$, the learning rate $\eta = p$, and we use a relu student. We vary the teacher functions: **Left:** $\sigma_1^*(z) = 4z^2 + z$, $\sigma_2^*(z) = z$ **Center:** $\sigma_1^*(z) = \sigma_2^*(z) = 4z^2 + z$, **Right:** $\sigma_2^*(z) = 4z^2 + z$, $\sigma_1^*(z) = z$. The orientation of the gradient does not change after $T = 6$ steps preventing specialization, in agreement with Thm. 3.

teacher functions, that the theoretical prediction of Thm. 2 on the gradient orientation, are valid beyond the symmetric teachers.

Multiple stairs: We complement the picture offered by Fig. 6 studying functions that have multiple linearly connected directions to the previously learned one, or informally, “multiple-stairs function”. The results are presented in Fig. 8 by considering the function $f_*(z) = z_1/3 + 2z_1z_2/3 + z_2z_3$; we consider 3 steps in the training of Alg. 1, the network is able to learn respectively e_1, e_2, e_3 after the first three steps of training in the proportional sample regime. This is clearly appreciable by studying the cosine similarity of the projected gradient on the teacher subspace: after the first step it is localized around e_1 , proceeding with training it has projections along e_2 while e_3 remains hidden, and only at the third step we obtain delocalization of the gradient along e_3 . These results are in perfect agreement with Thm. 3. Note that the hierarchical learning framework of Thm. 3 allows neurons to simultaneously specialize along different directions, as exemplified in Fig. 2 (see the bottom right plot). We observe one instance of this multidirectional staircase learning in Fig. 9 by considering the target $f_*(z) = z_1/3 + 2\text{He}_2(z_1)z_2 + z_1z_3$: while the results are unchanged in the first step with respect to Fig. 8 (with only the e_1 direction being learned), we observe that both directions e_2 & e_3 are learned at the second step.

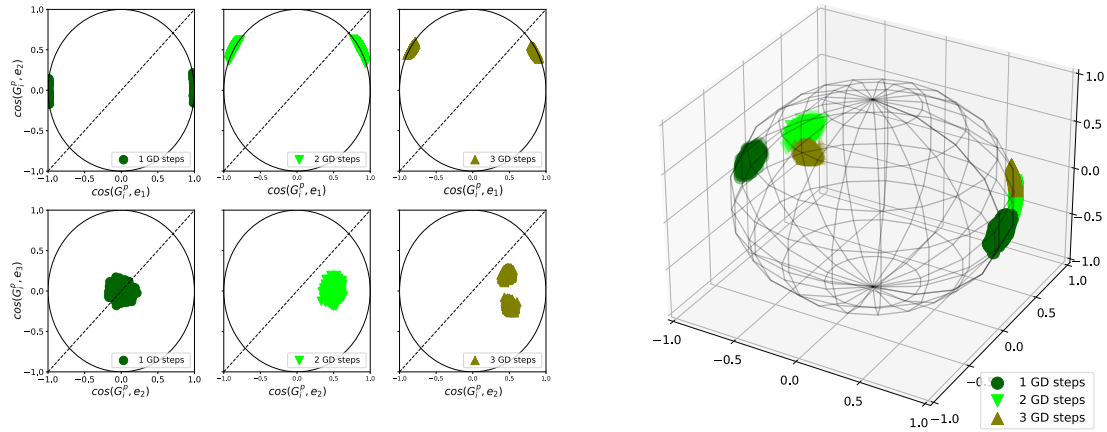


Figure 8: **Climbing multiple stairs.** Fix the teacher function $f_*(\mathbf{z}) = z_1/3 + 2z_1z_2/3 + z_2z_3$ and a relu student. The plots show the cosine similarity of the projected gradient matrix G^p inside the teacher subspace for all the p neurons at a fixed ratio $n/d = 4$, plotted at different stages of the training following Alg. 1. The plot shows the similarity in the 3D teacher subspace on the right, and two sections of it on the left: **Up:** (e_1, e_2) plane. **Bottom:** (e_2, e_3) plane. In accordance with Thm. 3, the gradient is first localized around e_1 , then sequentially learns e_2 , and only at the third step has components along e_3 .

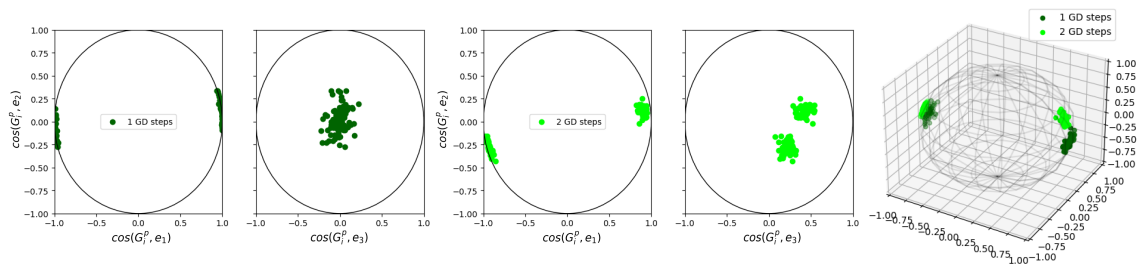


Figure 9: **Learning multiple directions at a time.** Fix the teacher function $f^*(\mathbf{z}) = z_1/3 + 2\text{He}_2(z_1)z_2 + z_1z_3$ and a relu student. The plots show the cosine similarity of the projected gradient matrix G^p inside the teacher subspace for all the p neurons at a fixed ratio $n/d = 4$, plotted at different stages of the training following Alg. 1. The plot shows the similarity measure in different cases. **Left:** (e_1, e_2) cross section. **Center:** (e_3, e_2) cross section. **Right:** 3D teacher subspace (e_1, e_2, e_3) . In accordance with Thm. 3, the gradient is first localized around the direction e_1 , and then learns both directions (e_3, e_2) at the second gradient step.

Algorithm 1 Training procedure

Choice of parameters Fix the data dimension and the width of the second layer (d, p) and sample (W_0, \mathbf{a}_0) obeying eq. (26). Fix a regularization parameter λ , and a number of GD steps T_{max} .

for n in a given range **do**

Learning rate tuning Fix the learning rate $\eta = \mathcal{O}(p\sqrt{\frac{n}{d}})$.

for $t < T_{max}$ **do**

Data generation Sample the data matrix $Z \sim \mathcal{N}(0, I_{n \times d})$ and get the labels $Y = f_*(Z) \in \mathbb{R}^n$

Update first layer Compute the gradient matrix $G_t = \{\mathbf{G}_i^{(t)}\}_{i \in [p]} \in \mathbb{R}^{p \times d}$ and update W :

$$\mathbf{G}_i^{(t)} \leftarrow \frac{a_{0,i}}{\sqrt{p}} \cdot \frac{1}{n} \sum_{\nu=1}^n \mathbf{x}^\nu \sigma'(\langle \mathbf{w}_i^{(t)}, \mathbf{z}^\nu \rangle) \left(\hat{f}(\mathbf{z}^\nu, W_t, \mathbf{a}_0) - f^*(\mathbf{z}^\nu) \right) \quad (20)$$

$$W_{t+1} \leftarrow W_t - \eta G_t \quad (21)$$

if $t == T_{max}$ **then**

Train second layer Get the feature matrix $X_t \leftarrow \sigma(W_t Z)$, and compute estimator:

$$\hat{\mathbf{a}} \leftarrow \begin{cases} X_t^\top (X_t X_t^\top + \lambda I_n)^{-1} Y & n < p \\ (X_t^\top X_t + \lambda I_p)^{-1} X_t^\top Y & n > p \end{cases}$$

end if

end for

end for

B PRELIMINARIES AND ASSUMPTIONS FOR THE PROOFS

B.1 PRELIMINARIES

Before going to the mathematical detail of the proof, we recall a few definition and useful facts.

Hermite expansion — Given the Gaussian measure γ_m on \mathbb{R}^m , we can build a scalar product on $\ell^2(\mathbb{R}^m, \gamma_m)$ as

$$\langle f, g \rangle_\gamma = \int_{\mathbb{R}^m} fg \, d\gamma_m = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}, I_m)}[f(\mathbf{z})g(\mathbf{z})]. \quad (22)$$

It turns out that there is a specific orthonormal basis of interest for this scalar product, that we present in tensor form:

Definition 4 (Hermite decomposition). *Let $f : \mathbb{R}^m \rightarrow \mathbb{R}$ be a function that is square integrable w.r.t the Gaussian measure. There exists a family of tensors $(C_j(f))_{j \in \mathbb{N}}$ such that $C_j(f)$ is of order j and for all $\mathbf{x} \in \mathbb{R}^m$,*

$$f(\mathbf{x}) = \sum_{j \in \mathbb{N}} \langle C_j(f), \mathcal{H}_j(\mathbf{x}) \rangle \quad (23)$$

where $\mathcal{H}_j(\mathbf{x})$ is the j -th order Hermite tensor (Grad, 1949).

Higher-order singular value decomposition — The higher-order singular value decomposition (HOSVD) of a tensor is defined as follows:

Definition 5 (Higher-order SVD). *Let $C \in \mathbb{R}^{m^k}$ be a symmetric tensor of order k . A higher-order SVD of C is an orthonormal set $(\mathbf{u}_1, \dots, \mathbf{u}_r)$ of $r \leq k$ vectors, as well as a tensor $S \in \mathbb{R}^{r^k}$ such that*

$$C = \sum_{j_1, \dots, j_k=1}^r S_{j_1, \dots, j_k} \mathbf{u}_{j_1} \otimes \dots \otimes \mathbf{u}_{j_k} \quad (24)$$

The singular values tensor S , as well as the rank r , are unique, but just as the regular SVD, the vectors $(\mathbf{u}_1, \dots, \mathbf{u}_r)$ are only unique up to rotations.

B.2 SETTING AND ASSUMPTIONS

We discuss the setting and main assumptions required. The first concerns the class of target functions we consider.

Assumption 2 (Data model). *The training inputs $\mathbf{z}^\nu \in \mathbb{R}^d$ are independently drawn from the Gaussian distribution $\mathcal{N}(0, I_d)$. Further, we assume that the target function $y^\nu = f^*(\mathbf{z})$ depends only on a few relevant directions. In other words, there exists a low-dimensional subspace $V^* \subset \mathbb{R}^d$ of fixed dimension r and a function $g^* : V^* \rightarrow \mathbb{R}$ such that:*

$$y = f^*(\mathbf{z}) := g^*(\Pi^* \mathbf{z}), \quad (25)$$

where Π^* is the orthogonal projection on V^* .

As we will show later, learning with GD can be seen as a hierarchical process, where depending on the batch size different directions of the target are progressively learned.

Given a batch of training data $(\mathbf{z}^\nu, y^\nu)_{\nu=1}^n \in \mathbb{R}^{d+1}$ drawn from the model (2) defined above, we now define how the network weights (W, a) are initialized and updated.

Assumption 3 (Training procedure). Consider the following random initialization for the weights:

$$\sqrt{p} \cdot a_i^0 \stackrel{i.i.d}{\sim} \text{Unif}([-1, 1]) \quad \text{and} \quad \mathbf{w}_i^0 \stackrel{i.i.d}{\sim} \text{Unif}(\mathbb{S}^{d-1}). \quad (26)$$

The distribution of the a_i can be replaced by any other continuous distribution with positive variance. Note that for $p = \mathcal{O}(1)$, we have $\hat{f}(\mathbf{z}; W^0, \mathbf{a}^0) \neq 0$. To further simplify the analysis, we assume p to be even and further impose the following symmetrization at initialization:

$$a_i^0 = -a_{p-i+1}^0 \quad \text{and} \quad \mathbf{w}_i^0 = \mathbf{w}_{p-i+1}^0 \quad \text{for all } i \in [p/2], \quad (27)$$

which ensures $\hat{f}(\mathbf{z}; W^0, \mathbf{a}^0) = 0$. Note that this simplification is common in the related literature, e.g. [Chizat et al. \(2019\)](#); [Damian et al. \(2022\)](#), and is mainly necessary when p is small. Given the initial conditions, the weights are trained with the following two-step full-batch gradient descent:

(i) First layer training: for every gradient step $t \leq T$, a fresh batch of training data $\{(\mathbf{z}^\nu, y^\nu)\}_{\nu=1}^n$ is drawn from the model in [Assumption 2](#), and the first layer weights are updated according to:

$$\mathbf{w}_i^{t+1} = \mathbf{w}_i^t - \frac{\eta}{2n} \sum_{\nu=1}^n \nabla_{\mathbf{w}_i} \left(y^\nu - \hat{f}(\mathbf{z}^\nu; W^t, \mathbf{a}^0) \right)^2, \quad (28)$$

Hence, the total sample complexity for this step is Tn .

(ii) Second layer training: once the first layer is trained for T steps, the second layer weights \mathbf{a} are trained to optimality by performing ridge regression with the features learned in the first step:

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a} \in \mathbb{R}^p} \frac{1}{2n} \sum_{\nu=1}^n \left(y^\nu - \hat{f}(\mathbf{z}^\nu; W^T, \mathbf{a}) \right)^2 + \lambda \|\mathbf{a}\|^2. \quad (29)$$

Such a separation of the training between the first and second layer is a common setup for the theoretical study of training ([Damian et al., 2022](#); [Abbe et al., 2023](#); [Berthier et al., 2023](#)), and allows for a more tractable study of convergence.

C GRADIENT DESCENT ON THE FIRST LAYER

C.1 TECHNICAL ASSUMPTIONS

We shall show our results under the following assumptions. First, since we assume that the leap index of f^* is at least one, and we setup the network to zero output the following assumption is unrestrictive:

Assumption 4. *The teacher function f^* and the student activation σ both have their zero-th Hermite coefficient equal to 0.*

We shall also need a smoothness assumption:

Assumption 5. *Both the student activation σ and g^* are continuous, and differentiable except possibly on a finite set of points. Further, the first two derivatives of g^* and the first three derivatives of σ are uniformly bounded in \mathbb{R} .*

C.2 PRELIMINARIES

More on Hermite expansion We recall a few properties of the Hermite tensors of Definition 4. Up to symmetry, the tensors $\mathcal{H}_k(\mathbf{x})$ are an orthonormal basis of $\ell^2(\mathbb{R}^m, \gamma)$, in the sense that for any $\mathbf{i}, \mathbf{j} \in \mathbb{R}^k$,

$$\langle \mathcal{H}_{k,\mathbf{i}}(\mathbf{x}), \mathcal{H}_{k,\mathbf{j}}(\mathbf{x}) \rangle_\gamma = \frac{1}{|\mathfrak{o}(\mathbf{i})|} \mathbf{1}_{\mathbf{i} \text{ is a permutation of } \mathbf{j}} \quad (30)$$

where $|\mathfrak{o}(\mathbf{i})|$ is the number of distinct permutations of \mathbf{i} . It can be checked from the definition in Grad (1949) that the \mathcal{H}_k , and hence the $C_k(f)$, are basis-invariant, and hence represent an actual k -linear form on \mathbb{R}^m . Further, the property (30) yields an immediate expression for the scalar product in $\ell^2(\mathbb{R}^m, \gamma_m)$:

$$\langle f, g \rangle_\gamma = \sum_{k \in \mathbb{N}} \langle C_k(f), C_k(g) \rangle. \quad (31)$$

Further, the Hermite coefficients of low-rank functions are straightforward to compute:

Lemma 1. *Let $g : \mathbb{R}^r \rightarrow \mathbb{R}$, and a linear map $A \in \mathbb{R}^{r \times d}$. Then the Hermite coefficients of $f(\mathbf{x}) = g(A\mathbf{x})$ are*

$$C_k(f) = C_k(g) \cdot (A, \dots, A), \quad (32)$$

where \cdot is the multilinear multiplication operator (Greub, 2012).

In particular, this implies that the singular vectors of C_k^* all belong to V^* .

Concentration in Orlicz spaces We recall the classical definition of Orlicz spaces:

Definition 6. *For any $\alpha \in \mathbb{R}$, let $\psi_\alpha(x) = e^{x^\alpha} - 1$. Let X be a real random variable; the Orlicz norm $\|X\|_{\psi_\alpha}$ is defined as*

$$\|X\|_{\psi_\alpha} = \inf \left\{ t > 0 : \mathbb{E} \left[\psi_\alpha \left(\frac{|X|}{t} \right) \right] \leq 1 \right\} \quad (33)$$

We refer to the monographs Ledoux and Talagrand (1991); van der Vaart and Wellner (1996) for more information. We say that a random variable is sub-gaussian (resp. sub-exponential) if its ψ_2 (resp. ψ_1) norm is finite. The main use of this definition is the following concentration inequality: for a variable X with finite Orlicz norm,

$$\mathbb{P}(|X - \mathbb{E}X| > t \|X\|_{\psi_\alpha}) \leq 2e^{-t^\alpha}. \quad (34)$$

The Orlicz norms are sub-multiplicative, in the following sense:

Lemma 2. Let X and Y be two random variables. Then, for any $\alpha > 0$, there exists a constant K_α such that

$$\|XY\|_{\psi_{\alpha/2}} \leq K_\alpha \|X\|_{\psi_\alpha} \|Y\|_{\psi_\alpha} \quad (35)$$

Finally, we shall use the following theorem:

Theorem 5 (Theorem 6.2.3 in [Ledoux and Talagrand \(1991\)](#) and Lemma 2.2.2 in [van der Vaart and Wellner \(1996\)](#)). Let X_1, \dots, X_n be n independent random variables with zero mean and second moment $\mathbb{E}X_i^2 = \sigma_i^2$. Then,

$$\left\| \sum_{i=1}^n X_i \right\|_{\psi_\alpha} \leq K_\alpha \log(n)^{1/\alpha} \left(\sqrt{\sum_{i=1}^n \sigma_i^2} + \max_i \|X_i\|_{\psi_\alpha} \right) \quad (36)$$

Preliminary computations We begin with a few useful preliminary computations. First, since $\mathbf{w}_i^0 \sim \text{Unif}(\mathbb{S}^{d-1})$, the following lemma holds:

Lemma 3. With probability at least $1 - cpe^{-c \log(d)^2}$, we have for any $i \in [p]$ and $k \in [r]$:

$$\|\boldsymbol{\pi}_i^0\| \leq \frac{\sqrt{r} \log(d)}{\sqrt{d}} \quad (37)$$

Let \mathbf{g}_i be the negative gradient for the i -th neuron at initialization:

$$\mathbf{g}_i = -\nabla_{\mathbf{w}_j} \mathcal{L} \left(\hat{f}(\mathbf{z}^\nu; W^0, \mathbf{a}), f^*(\mathbf{z}^\nu) \right). \quad (38)$$

Since at initialization the output of the network is exactly zero, we have

$$\mathbf{g}_i = \frac{a_i}{\sqrt{p}} \cdot \frac{1}{n} \sum_{\nu=1}^n \mathbf{z}^\nu \sigma'(\langle \mathbf{w}_i^0, \mathbf{z}^\nu \rangle) f^*(\mathbf{z}^\nu) \quad (39)$$

Finally, the update equation for $\|\mathbf{w}_i\|$ reads

$$\|\mathbf{w}_i^1\|^2 = 1 + 2\eta \langle \mathbf{w}_i^0, \mathbf{g}_i \rangle + \eta^2 \|\mathbf{g}_i\|^2 \quad (40)$$

C.3 COMPUTING EXPECTATIONS

We begin by a simple computation of the expectation of \mathbf{g}_i :

Lemma 4. For any $i \in [p]$, we have

$$\mathbb{E}[\mathbf{g}_i] = \frac{a_i}{\sqrt{p}} \left(\sum_{k=0}^{\infty} c_{k+2} \langle (\mathbf{w}_i^0)^{\otimes k}, C_k^* \rangle \mathbf{w}_i + \sum_{k=0}^{\infty} c_{k+1} C_{k+1}^* \times_{1\dots k} (\mathbf{w}_i^0)^{\otimes k} \right) \quad (41)$$

where the last multiplication is a product over the first k axes of C_{k+1}^* (and thus results in a vector).

Proof. By Stein's lemma, for any \mathbf{w} , we have

$$\begin{aligned} \mathbb{E}[\mathbf{z} \sigma'(\langle \mathbf{w}, \mathbf{z} \rangle) f^*(\mathbf{z})] &= \mathbb{E}[\nabla_{\mathbf{z}} \sigma'(\langle \mathbf{w}, \mathbf{z} \rangle) f^*(\mathbf{z})] + \mathbb{E}[\sigma'(\langle \mathbf{w}, \mathbf{z} \rangle) \nabla_{\mathbf{z}} f^*(\mathbf{z})] \\ &= \mathbf{w} \mathbb{E}[\sigma''(\langle \mathbf{w}, \mathbf{z} \rangle) f^*(\mathbf{z})] + \mathbb{E}[\sigma'(\langle \mathbf{w}, \mathbf{z} \rangle) \nabla_{\mathbf{z}} f^*(\mathbf{z})] \end{aligned}$$

From Lemma 1, the k -th Hermite coefficient of $\mathbf{z} \mapsto \sigma''(\langle \mathbf{w}, \mathbf{z} \rangle)$ is $c_{k+2} \mathbf{w}^{\otimes k}$, where the $(c_k)_{k \geq 0}$ are the Hermite coefficients of σ . By two applications of the scalar product formula (31), we find

$$\mathbb{E}[\mathbf{z} \sigma'(\langle \mathbf{w}, \mathbf{z} \rangle) f^*(\mathbf{z})] = \sum_{k=0}^{\infty} c_{k+2} \langle \mathbf{w}^{\otimes k}, C_k^* \rangle \mathbf{w} + \sum_{k=0}^{\infty} c_{k+1} C_{k+1}^* \times_{1\dots k} \mathbf{w}^{\otimes k}. \quad (42)$$

□

Truncating the expansions Now, we show that the expectations in Lemma 4 can be truncated at the leap index term.

Lemma 5. *With probability at least $1 - cpe^{-c \log(d)^2}$, for every $k \geq 0$ and $i \in [p]$, we have*

$$|\langle C_k^*, (\mathbf{w}_i^0)^{\otimes k} \rangle| \leq c \left(\frac{\sqrt{r} \log(d)}{\sqrt{d}} \right)^k \quad \text{and} \quad \|C_{k+1}^* \times_{1\dots k} (\mathbf{w}_i^0)^{\otimes k}\| \leq c \left(\frac{\sqrt{r} \log(d)}{\sqrt{d}} \right)^k \quad (43)$$

As a result, if ℓ is the leap index of f^* ,

$$\left\| \mathbb{E}[\mathbf{g}_i] - C_\ell^* \times_{1\dots(\ell-1)} (\mathbf{w}_i^0)^{\otimes(\ell-1)} \right\| = \mathcal{O} \left(\frac{r^{\ell/2} \text{polylog}(d)}{d^{\ell/2}} \right) \quad (44)$$

Proof. First, we have by Lemma 1,

$$|\langle C_k^*, (\mathbf{w}_i^0)^{\otimes k} \rangle| = |\langle C_k(g^*), (W^* \mathbf{w}_i^0)^{\otimes k} \rangle| \leq \|C_k(g^*)\|_2 \cdot \|\boldsymbol{\pi}_i^0\|^k,$$

where $\|C_k(g^*)\|_2$ is the operator norm of $C_k(g^*)$. Since

$$\|C_k(g^*)\|_2 \leq \|C_k(g^*)\|_F \leq \|g^*\|_\gamma,$$

the first inequality ensues by Lemma 3. Now, let A_{k+1} be the $(k+1)$ -th mode unfolding of $C_{k+1}(g^*)$; then

$$\|C_{k+1}^* \times_{1\dots k} (\mathbf{w}_i^0)^{\otimes k}\| = \|A_{k+1}(W^* \mathbf{w}_i^0)^{\otimes k}\| \leq \|A_{k+1}\|_2 \|\boldsymbol{\pi}_i^0\|^k$$

The norm of A_{k+1} is then bounded by the same argument as above. The final equality is obtained by using the above bounds on every term above $k = \ell$ in the first sum, and above $k = \ell - 1$ in the second. \square

Student norms We now move on to controlling (40), in expectation. We begin with the cross-term:

Lemma 6. *With probability at least $1 - cpe^{-c \log(d)^2}$, we have for any $i \in [p]$,*

$$\mathbb{E}[\langle \mathbf{w}_i^0, \mathbf{g}_i \rangle] = \mathcal{O} \left(\frac{r^{\ell/2} \text{polylog}(d)}{pd^{\ell/2}} \right) \quad (45)$$

Proof. From eq. (44), we have

$$\mathbb{E}[\langle \mathbf{w}_i^0, \mathbf{g}_i \rangle] = \frac{a_i}{\sqrt{p}} \left(\langle C_\ell^*, (\mathbf{w}_i^0)^{\otimes \ell} \rangle + \mathcal{O} \left(\frac{r^{\ell/2} \text{polylog}(d)}{d^{\ell/2}} \right) \right)$$

The first part of Lemma 5 gives

$$\langle C_\ell^*, (\mathbf{w}_i^0)^{\otimes \ell} \rangle = \mathcal{O} \left(\frac{r^{\ell/2} \text{polylog}(d)}{pd^{\ell/2}} \right),$$

and the lemma follows since $|a_i| \leq 1/\sqrt{p}$. \square

The main object of study is therefore $\|\mathbf{g}_i\|^2$. We can write it as

$$\begin{aligned} \|\mathbf{g}_i\|^2 &= \frac{a_i^2}{n^2 p^2} \sum_{\nu, \nu'=1}^n \langle \mathbf{z}^\nu, \mathbf{z}^{\nu'} \rangle \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle) f^*(\mathbf{z}^\nu) \sigma'(\langle \mathbf{w}_i, \mathbf{z}^{\nu'} \rangle) f^*(\mathbf{z}^{\nu'}) \\ &= \frac{a_i^2}{n^2 p^2} \left(\sum_{\nu \neq \nu'} \langle \mathbf{z}^\nu, \mathbf{z}^{\nu'} \rangle \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle) f^*(\mathbf{z}^\nu) \sigma'(\langle \mathbf{w}_i, \mathbf{z}^{\nu'} \rangle) f^*(\mathbf{z}^{\nu'}) \right. \\ &\quad \left. + \sum_{\nu=1}^n \|\mathbf{z}^\nu\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2 \right) \end{aligned} \quad (46)$$

Since $\mathbf{z}^\nu, \mathbf{z}^{\nu'}$ are independent for $\nu \neq \nu'$, this leaves

$$\mathbb{E} [\|\mathbf{g}_i\|^2] = \frac{n(n-1)}{n^2} \|\mathbb{E}[\mathbf{g}_i]\|^2 + \frac{a_i^2}{np^2} \mathbb{E} [\|\mathbf{z}^\nu\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] \quad (47)$$

We shall only need orders of magnitude for those terms. These are taken care of in the following lemma:

Lemma 7. *There exists a bounded random variable X independent from d such that, with probability at least $1 - cpe^{-\log(d)^2}$,*

$$\|\mathbb{E}[\mathbf{g}_i]\|^2 = a_i^2 X_i \cdot \frac{\|\boldsymbol{\pi}_i^0\|^{2(\ell-1)}}{p^2} + \mathcal{O}\left(\frac{r^{\ell/2} \text{polylog}(d)}{d^{\ell/2}}\right) \quad (48)$$

where $(X_i)_{i \in [p]}$ are i.i.d copies of X . Additionally, there exist two constants c, C such that

$$c \cdot d \leq \mathbb{E} [\|\mathbf{z}\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z} \rangle)^2 f^*(\mathbf{z})^2] \leq C \cdot d \quad (49)$$

Proof. We begin with (48). Define the unit norm vectors

$$\mathbf{r}_i = \frac{W^* \mathbf{w}_i^0}{\|\boldsymbol{\pi}_i^0\|},$$

since the \mathbf{w}_i are isotropic, the \mathbf{r}_i are uniform on \mathbb{S}^{r-1} . Then,

$$\left\| C_\ell^* \times_{1 \dots (\ell-1)} (\mathbf{w}_i^0)^{\otimes (\ell-1)} \right\|^2 = \underbrace{\left\| C_\ell(g^*) \times_{1 \dots (\ell-1)} \mathbf{r}_i^{\otimes (\ell-1)} \right\|^2}_{=: X_i} \cdot \|\boldsymbol{\pi}_i^0\|^{2(\ell-1)}.$$

The random variables X_i thus defined are i.i.d, independent from d , and have positive expectation since $C_\ell(g^*)$ is nonzero. Equation (48) then results from the expansion in (44).

We now move on to the second part; first, by Hölder's inequality,

$$\mathbb{E} [\|\mathbf{z}\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z} \rangle)^2 f^*(\mathbf{z})^2] \leq \sqrt{\mathbb{E} [\|\mathbf{z}\|^4]} \sqrt[4]{\mathbb{E} [\sigma'(\langle \mathbf{w}_i, \mathbf{z} \rangle)^8]} \sqrt[4]{\mathbb{E} [f^*(\mathbf{z})^8]} \leq C \cdot d, \quad (50)$$

since the last two expectations are independent from d . On the other hand, using the same inequality with $\|\mathbf{z}\|^2 - d$, we can write

$$\mathbb{E} [\|\mathbf{z}\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z} \rangle)^2 f^*(\mathbf{z})^2] = d \mathbb{E} [\sigma'(\langle \mathbf{w}_i, \mathbf{z} \rangle)^2 f^*(\mathbf{z})^2] + \mathcal{O}(\sqrt{d}).$$

Since $\mu_\ell \neq 0$ and f^* has leap index ℓ , there exists $\varepsilon > 0$ two subsets $\mathcal{A} \subseteq \mathbb{R}, \mathcal{B} \subseteq V^*$ of positive measure such that $\sigma'(x)^2 \geq \varepsilon$ if $x \in \mathcal{A}$ and $f^*(\mathbf{z}^*) > \varepsilon$ if $\mathbf{z}^* \in \mathcal{B}$. From the fact that $\pi_i \leq 1/2$ with high probability, we conclude that the set

$$\mathcal{C} := \{\mathbf{z} \in \mathbb{R}^p : \langle \mathbf{w}_i, \mathbf{z} \rangle \in \mathcal{A}, P_{V^*} \mathbf{z} \in \mathcal{B}\}$$

has positive (Gaussian) measure. It follows that

$$\mathbb{E} [\sigma'(\langle \mathbf{w}_i, \mathbf{z} \rangle)^2 f^*(\mathbf{z})^2] \geq \gamma(\mathcal{C}) \varepsilon^2, \quad (51)$$

which concludes the proof of Eq. (49). \square

C.4 CONCENTRATION

We now move on to concentrating the quantities of interest of the previous section. Our aim will be to show the following proposition:

Proposition 2. *With probability at least $1 - Cpe^{-c\log(n)^2} - Cpe^{-c\log(d)^2}$, for any $i \in [p], k \in [r]$,*

$$\|\boldsymbol{\pi}_i^1 - \mathbb{E}[\boldsymbol{\pi}_i^1]\| = \mathcal{O}\left(\frac{\eta\sqrt{r}\log(n)}{p\sqrt{n}}\right) \quad (52)$$

$$\left|\|\mathbf{w}_i^1\|^2 - \mathbb{E}[\|\mathbf{w}_i^1\|^2]\right| = \mathcal{O}\left(\frac{\eta\log(n)}{p\sqrt{n}} + \frac{\eta^2 d \log(n)^6}{p^2 n \sqrt{n}} + \frac{\eta^2 \log(d)}{p^2 n \sqrt{d}} + \frac{\eta^2 r \log(n)^2}{p^2 n} + \frac{\eta^2 \log(n)^\ell r^{(\ell-1)/2}}{p^2 d^{(\ell-1)/2} \sqrt{n}}\right) \quad (53)$$

Importantly, we do not claim that the whole vector \mathbf{w}_i^1 concentrates; only its norm and its projection on a low-dimensional subspace do. Throughout this section, we define the random vectors

$$\mathbf{X}^\nu = \mathbf{z}^\nu \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle) f^*(\mathbf{z}^\nu). \quad (54)$$

These vectors are i.i.d, with the same distribution as a random vector that we will call \mathbf{X} .

Concentration of linear functionals We begin with a simple bound, that implies both Eq. (52) and the first term of Eq. (53).

Lemma 8. *Let \mathbf{w} be a unit vector in \mathbb{R}^d . There exists a universal constant c such that with probability $1 - 2pe^{-c\log(n)^2}$, for any $i \in [p]$ and $k \in [r]$,*

$$|\langle \mathbf{w}, \mathbf{g}_i \rangle - \mathbb{E}[\langle \mathbf{w}, \mathbf{g}_i \rangle]| \leq \frac{\log(n)}{p\sqrt{n}} \quad (55)$$

Proof. By Assumption 5, the function f^* is Lipschitz, so $f^*(\mathbf{z})$ is a sub-gaussian random variable. The same is obviously true for $\langle \mathbf{w}, \mathbf{z} \rangle$, and since σ' is bounded the random variable $\langle \mathbf{w}, \mathbf{X} \rangle = \langle \mathbf{w}, \mathbf{z} \rangle \sigma'(\langle \mathbf{w}_i^0, \mathbf{z} \rangle) f^*(\mathbf{z})$ is sub-exponential with bounded sub-exponential norm. We can thus apply Bernstein's inequality (Vershynin, 2018, Corollary 2.8.3) with $t = \log(n)/\sqrt{n}$ to get

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_{\nu=1}^n \langle \mathbf{w}, \mathbf{X}^\nu \rangle - \mathbb{E}[\langle \mathbf{w}, \mathbf{X} \rangle]\right| \geq \frac{\log(n)}{\sqrt{n}}\right) \leq 2e^{-c\log(n)^2}. \quad (56)$$

The result ensues upon noticing that $\frac{1}{n} \sum \mathbf{X}^\nu$ differs from \mathbf{g}_i by a factor of at most $1/p$. \square

Decomposing the gradient norm We now move on to the concentration of the q_i . This allows us to write

$$\|\mathbf{g}_i\|^2 - \mathbb{E}[\|\mathbf{g}_i\|^2] = \frac{1}{n^2 p^2} \left(\underbrace{\sum_{\nu=1}^n \|\mathbf{X}^\nu\|^2 - n\mathbb{E}[\|\mathbf{X}\|^2]}_{S_1} + \underbrace{\sum_{\nu \neq \nu'} \langle \mathbf{X}^\nu, \mathbf{X}^{\nu'} \rangle - n(n-1)\mathbb{E}[\langle \mathbf{X}, \mathbf{X} \rangle]}_{S_2} \right) \quad (57)$$

We shall show the concentration of those two terms sequentially.

Concentrating the norms We first focus on S_1 :

Lemma 9. *Let $i \in [p]$. There exists a constant $c > 0$ such that with probability $1 - e^{-c \log(n)^2}$,*

$$\mathbb{P}(|S_1| \geq \log(n)^6 d \sqrt{n}) \leq e^{-c \log(n)^2}. \quad (58)$$

Proof. The random variable $\|z\|/\sqrt{d}$ is sub-gaussian, and so is $f^*(z^\nu)$. By Lemma 2 and the Hölder inequality, the random variable $\|\mathbf{X}^\nu\|^2$ satisfies

$$\|\|\mathbf{X}^\nu\|^2\|_{\psi_{1/2}} \leq C \cdot d \quad \text{and} \quad \text{Var}(\|\mathbf{X}^\nu\|^2) \leq C \cdot d^2$$

As a result, we can apply Theorem 5 to the random variables $\|\mathbf{X}^\nu\|^2 - \mathbb{E}[\|\mathbf{X}\|^2]$, which yields

$$\left\| \sum_{\nu=1}^n \|\mathbf{X}^\nu\|^2 - n \mathbb{E}[\|\mathbf{X}\|^2] \right\|_{\psi_{1/2}} \leq c \log(n)^2 d \sqrt{n}. \quad (59)$$

Equation (58) is then a consequence of the Orlicz concentration bound (34). \square

Decomposing the cross-term We now move on to S_2 . To handle this sum, we use the following decoupling result from [Pena and Montgomery-Smith \(1995\)](#):

Theorem 6. *Let $(f_{ij})_{i,j \in [n]}$ be a set of measurable functions from S^2 to a Banach space $(B, \|\cdot\|)$, and $(X_1, \dots, X_n), (Y_1, \dots, Y_n)$ two sets of independent random variables such that the laws of X_i and Y_i are the same. Then there exists a constant $C > 0$ such that*

$$\mathbb{P}\left(\left\| \sum_{i \neq j} f_{ij}(X_i, X_j) \right\| \geq t\right) \leq C \mathbb{P}\left(\left\| \sum_{i \neq j} f_{ij}(X_i, Y_j) \right\| \geq \frac{t}{C}\right) \quad (60)$$

We apply this theorem to the functions $f_{\nu, \nu'}(\mathbf{X}^\nu, \mathbf{X}^{\nu'}) = \langle \mathbf{X}^\nu, \mathbf{X}^{\nu'} \rangle - \|\mathbb{E}\mathbf{X}\|^2$. Let \mathbf{Y}^ν be an independent copy of the \mathbf{X}^ν for $\nu \in [n]$, we then have to estimate

$$\mathbb{P}\left(\left| \sum_{\nu \neq \nu'} \langle \mathbf{X}^\nu, \mathbf{Y}^{\nu'} \rangle - n(n-1) \|\mathbb{E}\mathbf{X}\|^2 \right| \geq t\right).$$

For convenience, let $\bar{x} = \|\mathbb{E}\mathbf{X}\|^2$. Since the \mathbf{X}^ν are sub-exponential vectors, the scalar product $\langle \mathbf{X}^\nu, \mathbf{Y}^\nu \rangle$ has finite $\psi_{1/2}$ -norm. The same bound as Lemma 9 then gives that

$$\mathbb{P}\left(\left| \sum_{\nu=1}^n \langle \mathbf{X}^\nu, \mathbf{Y}^\nu \rangle - n \bar{x} \right| \geq \sqrt{nd} \log(n)^6\right) \leq e^{-c \log(n)^2} \quad (61)$$

Hence, to show Proposition 2, we only need to study the overall sum

$$\tilde{S}_2 := \sum_{\nu, \nu'=1}^n \langle \mathbf{X}^\nu, \mathbf{Y}^{\nu'} \rangle - n^2 \bar{x}^2 \quad (62)$$

Recall that, as in the proof of Lemma 7, the vector $\mathbb{E}\mathbf{X}$ belongs to the space $V_i = V^* + \text{span}(\mathbf{w}_i^0)$. We thus make the decomposition

$$\mathbf{X}^\nu = \mathbf{X}_i^\nu + \mathbf{X}_\perp^\nu \quad \text{and} \quad \mathbf{Y}^\nu = \mathbf{Y}_i^\nu + \mathbf{Y}_\perp^\nu \quad (63)$$

where $\mathbf{X}_i^\nu, \mathbf{Y}_i^\nu \in V_i$. Hence,

$$\sum_{\nu, \nu'=1}^n \langle \mathbf{X}^\nu, \mathbf{Y}^{\nu'} \rangle - n^2 \bar{x}^2 = \underbrace{\left\langle \sum_{\nu=1}^n \mathbf{X}_i^\nu, \sum_{\nu=1}^n \mathbf{Y}_i^\nu \right\rangle - n^2 \bar{x}^2}_{S'_2} + \underbrace{\left\langle \sum_{\nu=1}^n \mathbf{X}_\perp^\nu, \sum_{\nu=1}^n \mathbf{Y}_\perp^\nu \right\rangle}_{S'_2} \quad (64)$$

Bounding the last two terms The main step in bounding S'_2 is the following lemma:

Lemma 10. *With probability at least $1 - Ce^{-c \log(n)^2}$,*

$$\left\| \sum_{\nu=1}^n \mathbf{X}_i^\nu - n\mathbb{E}\mathbf{X} \right\| \leq C\sqrt{r} \log(n)\sqrt{n} \quad (65)$$

Proof. Let $(\mathbf{u}_1, \dots, \mathbf{u}_{r+1})$ be an orthonormal basis of V_i . Since we have for any vector $\mathbf{x} \in V_i$

$$\|\mathbf{x}\|^2 = \sum_{k=1}^{r+1} \langle \mathbf{x}, \mathbf{u}_k \rangle^2,$$

it suffices to bound such a scalar product with high probability. Each term of the form $\langle \mathbf{X}_i^\nu - \mathbb{E}\mathbf{X}, \mathbf{u}_k \rangle$ is a sub-exponential random variable with zero mean and bounded variance, and hence by another application of Bernstein's inequality

$$\mathbb{P}\left(\left|\sum_{\nu=1}^n \langle \mathbf{X}_i^\nu - \mathbb{E}\mathbf{X}, \mathbf{u}_k \rangle\right| \geq \log(n)\sqrt{n}\right) \leq e^{-c \log(n)^2} \quad (66)$$

The result ensues from a union bound, and the equivalence of norms in finite-dimensional spaces. \square

As an easy corollary of this lemma, we get the following bound on S'_2 :

Corollary 2. *With probability at least $1 - Ce^{-c \log(n)^2}$,*

$$S'_2 = \mathcal{O}\left(rn \log(n)^2 + \frac{\log(n)^{\ell} r^{(\ell-1)/2} n \sqrt{n}}{d^{(\ell-1)/2}}\right) \quad (67)$$

Proof. We use the following decomposition:

$$\begin{aligned} S'_2 &= n \langle \sum_{\nu=1}^n \mathbf{X}_i^\nu - n\mathbb{E}\mathbf{X}, \mathbb{E}\mathbf{X} \rangle + n \langle \mathbb{E}\mathbf{X}, \sum_{\nu=1}^n \mathbf{Y}_i^\nu - n\mathbb{E}\mathbf{X} \rangle + \langle \sum_{\nu=1}^n \mathbf{X}_i^\nu - n\mathbb{E}\mathbf{X}, \sum_{\nu=1}^n \mathbf{Y}_i^\nu - n\mathbb{E}\mathbf{X} \rangle \\ &\leq n \|\mathbb{E}\mathbf{X}\| \left(\left\| \sum_{\nu=1}^n \mathbf{X}_i^\nu - n\mathbb{E}\mathbf{X} \right\| + \left\| \sum_{\nu=1}^n \mathbf{Y}_i^\nu - n\mathbb{E}\mathbf{X} \right\| \right) + \left\| \sum_{\nu=1}^n \mathbf{X}_i^\nu - n\mathbb{E}\mathbf{X} \right\| \cdot \left\| \sum_{\nu=1}^n \mathbf{Y}_i^\nu - n\mathbb{E}\mathbf{X} \right\| \end{aligned}$$

by the Cauchy-Schwarz inequality. The result ensues from the high probability bounds of Lemma 10, as well as the bound on $\|\mathbb{E}\mathbf{X}\|$ from Lemma 7. \square

We finally bound the last term, which closes the proof of Proposition 2.

Lemma 11. *Let $i \in [p]$. With probability at least $1 - 2e^{-c \log(n)^2} - e^{-c \log(d)^2}$, we have*

$$|S''_2| \leq 2 \log(d) n \sqrt{d} \quad (68)$$

Proof. Define $\alpha^\nu = \sigma'(\langle \mathbf{w}_i^0, \mathbf{z}^\nu \rangle) f^*(\mathbf{z}^\nu)$, and β^ν its equivalent for \mathbf{Y}^ν . Since α^ν only depends on \mathbf{z}_i , the distribution of $\sum \mathbf{X}_\perp^\nu$ is the same as $\|\alpha\| \mathbf{X}_\perp$, where \mathbf{X}_\perp is a normal random vector in V_i^\perp . Therefore, we have

$$S''_2 \stackrel{d}{=} \|\alpha\| \cdot \|\beta\| \cdot \langle \mathbf{X}_\perp, \mathbf{Y}_\perp \rangle$$

for two independent Gaussian vectors $\mathbf{X}_\perp, \mathbf{Y}_\perp$. Now, both $\|\alpha\|^2$ and $\|\beta\|^2$ are the sum of n sub-exponential random variables with bounded variance, and $\langle \mathbf{X}_\perp, \mathbf{Y}_\perp \rangle$ is the sum of d such variables. Hence, by Bernstein's inequality, with probability $1 - 2e^{-c \log(n)^2}$,

$$\|\alpha\|^2 \leq n + \log(n)\sqrt{n} \leq 2n \quad \text{and} \quad \|\beta\|^2 \leq 2n,$$

and with probability at least $1 - e^{-c \log(d)^2}$

$$\langle \mathbf{X}_\perp, \mathbf{Y}_\perp \rangle \leq \log(d)\sqrt{d},$$

which ends the proof. \square

C.5 PROOF OF THEOREMS 1 AND 2

We begin with a proposition that summarizes everything from the two previous sections.

Proposition 3. *Let ℓ be the leap index of f^* , and assume that $n = \Omega(d^{\ell-\delta})$ for some $\delta > 0$. There is an event with probability at least $1 - cpe^{-\log(d)^2}$ such that for $i \in [p]$:*

$$\left\| \boldsymbol{\pi}_i^1 - \left(\boldsymbol{\pi}_i^0 + \frac{\eta a_i}{\sqrt{p}} C_\ell^* \times_{1 \dots (\ell-1)} (\mathbf{w}_i^0)^{\otimes (\ell-1)} \right) \right\| = \mathcal{O} \left(\frac{r^{\ell/2} \text{polylog}(d)}{d^{\ell/2}} + \frac{\sqrt{r} \eta \log(d)}{p \sqrt{n}} \right) \quad (69)$$

$$\|\mathbf{w}_i^1\|^2 = \Theta \left(1 + \frac{\eta^2 X_i \|\boldsymbol{\pi}_i^0\|^2}{p^2} + \frac{\eta^2 d}{np^2} \right) \quad (70)$$

where the X_i are i.i.d random variables as in Lemma 7.

Proof. The proof amounts to checking that all the bounds proven so far are of the right order. The first equality is simply a combination of Lemma 5 and Proposition 2. For the second part, notice that Lemma 7 implies that

$$\mathbb{E} \left[\|\mathbf{w}_i^1\|^2 \right] = \Theta \left(1 + \frac{\eta^2 X_i \|\boldsymbol{\pi}_i^0\|^2}{p^2} + \frac{\eta^2 d}{np^2} \right),$$

and it is straightforward (albeit tedious) to check that all bounds in Proposition 2 are negligible with respect to the above expectation. \square

Proof of Theorem 1 We first consider the case where $n = \Theta(d^{\ell-\delta})$. A simple triangular inequality yields

$$\|\boldsymbol{\pi}_i^1\| = \mathcal{O} \left(\|\boldsymbol{\pi}_i^0\| + \frac{\eta \|\boldsymbol{\pi}_i^0\|^{\ell-1}}{p} \right)$$

where the second part is due to Lemma 7. On the other hand, the middle term in (70) becomes negligible w.r.t the rightmost one, so we get

$$\|\mathbf{w}_i^1\| = \Omega \left(1 + \frac{\eta d^{\delta/2}}{p} \right)$$

This implies

$$\frac{\|\boldsymbol{\pi}_i^1\|}{\|\mathbf{w}_i^1\|} = \mathcal{O} \left(\max \left(\|\boldsymbol{\pi}_i^0\|, \frac{\|\boldsymbol{\pi}_i^0\|^{\ell-1}}{d^{\delta/2}} \right) \right) = \mathcal{O} \left(\frac{\text{polylog}(d)}{d^{(1 \wedge \delta)/2}} \right) \quad (71)$$

where the last inequality is due to Lemma 3.

Proof of Theorem 2 Now, we take $n = \Omega(d^\ell)$, and $\eta = d^{(\ell-1)/2}$. Then, the bounds of Proposition 3 become

$$\left\| \boldsymbol{\pi}_i^1 - a_i d^{(\ell-1)/2} C_\ell^* \times_{1\dots(\ell-1)} (\mathbf{w}_i^0)^{\otimes(\ell-1)} \right\| = \mathcal{O} \left(\frac{\sqrt{r} \text{polylog}(d)}{\sqrt{d}} \right) \quad \text{and} \quad \|\mathbf{w}_i^1\|^2 = \mathcal{O}(1)$$

Hence, the first part of Theorem 2 is straightforward: from Lemma 7,

$$\frac{\|\boldsymbol{\pi}_i^1\|}{\|\mathbf{w}_i^1\|} = \Omega \left(a_i^2 X_i \cdot (\sqrt{d} \|\boldsymbol{\pi}_i^0\|)^{\ell-1} \right), \quad (72)$$

which is a random variable with positive expectation. The latter part is not independent from d , but it dominates e.g. a variable of the form $\|\mathbf{z}_r\|$ where $\mathbf{z}_r \sim \mathcal{N}(0, I_r/2)$ with probability $1 - ce^{-\log(d)^2}$.

For the second part, we write using the higher-order SVD of C_ℓ^*

$$C_\ell^* \times_{1\dots(\ell-1)} (\mathbf{w}_i^0)^{\otimes(\ell-1)} = \sum_{j_1, \dots, j_\ell=1}^{r_\ell} S_{j_1, \dots, j_\ell} \langle \mathbf{w}_i^0, \mathbf{u}_{j_1}^* \rangle \dots \langle \mathbf{w}_i^0, \mathbf{u}_{j_{\ell-1}}^* \rangle \mathbf{u}_{j_\ell}^*$$

which belongs to V_ℓ^* . Finally, since S is full-rank, each vector $C_\ell^* \times_{1\dots(\ell-1)} (\mathbf{w}_i^0)^{\otimes(\ell-1)}$ is an i.i.d random variable which is absolutely continuous w.r.t the Lebesgue measure in V_ℓ^* . This implies that the collection of such vectors is full-rank with probability one, and ends the proof of Theorem 2.

C.6 SPIKE+BULK DECOMPOSITION

Having proven Theorems 1 and 2, we move to investigate the behavior after multiple gradient steps. First, we relate the discussion above to a ‘‘spike+noise’’ decomposition of the gradient. We start from Equation (39):

$$\mathbf{g}_i = \frac{a_i}{\sqrt{p}} \cdot \frac{1}{n} \sum_{\nu=1}^n \mathbf{z}^\nu \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle) f^*(\mathbf{z}^\nu) \quad (73)$$

Define $\sigma'_{>1}(u) : \mathbb{R} \rightarrow \mathbb{R}$ as the following function:

$$\sigma'_{>1}(u) = \sigma'(u) - \mu_1, \quad (74)$$

so that $\mathbb{E}[\sigma'_{>1}(u)] = 0$. We have the following decomposition of the gradient:

$$\mathbf{g}_i = \frac{a_j}{\sqrt{p}} \frac{1}{n} \mu_1 \sum_{i=1}^n y_i \mathbf{x}_i + \underbrace{\frac{1}{n} \frac{a_j}{\sqrt{p}} \mu_1 \sum_{i=1}^n \sigma'_{>1}(\mathbf{x}_i^\top \mathbf{w}^0) \mathbf{x}_i y_i}_{\Delta_j}, \quad (75)$$

or in matrix form:

$$\mathbf{G} = \mathbf{u}\mathbf{v}^\top + \Delta, \quad (76)$$

where $\mathbf{u} = \frac{\mu_1}{\sqrt{p}} \mathbf{a}$, $\mathbf{v} = \frac{1}{n} \sum_{i=1}^n y_i \mathbf{x}_i$. A similar decomposition was utilized in Ba et al. (2022) to provide an asymptotic characterization of the training and generalization errors in the regime $n = \Theta(d)$ and step-size $\eta = \mathcal{O}(\sqrt{p})$. In particular, they show that the presence of this spike for $\eta = \mathcal{O}(\sqrt{p})$ is not enough to go beyond the linear kernel regime.

However, as we see below, it is possible to obtain a precise characterization in the feature learning regime $\eta = \Theta(p)$ and generalizing to multiple steps, with stronger concentration over the structure of Δ . In particular, we prove that Δ effectively acts as uniform noise that can be incorporated into the initialization $\mathbf{W}^{(0)}$.

This is expressed through the following Lemma:

Lemma 12. *With high probability over the initialization W^0 , as $n, d \rightarrow \infty$ with $n = \Omega(\max(p, d))$, the matrix Δ satisfies the following:*

- (i) For any $\mathbf{v} \in V^*$, with $\|\mathbf{v}\| = 1$, $\langle \Delta_j, \mathbf{v} \rangle = \mathcal{O}\left(\frac{\text{polylog}(d)}{p\sqrt{d}}\right)$.
- (ii) $\|\Delta\| = \mathcal{O}(\text{polylog } d/\sqrt{d})$.
- (iii) For any $i \neq j, i, j \in [p/2]$, $\Delta_j^\top \Delta_i = \mathcal{O}\left(\frac{\text{polylog}(d)}{p^2\sqrt{d}}\right)$,

where we only consider the first half neurons due to the choice of the symmetric initialization in Equation (26).

Proof. Without loss of generality, we assume that $\mu_1 = 0$ and hence that $\Delta_i = \mathbf{g}_i$. By Lemma 4, since $\mu_1 = 0$; we have $\mathbb{E}[\Delta_j^\top \mathbf{v}] = \mathcal{O}\left(\frac{\text{polylog}(d)}{p\sqrt{d}}\right)$. Furthermore, from Lemma 8, we obtain that, with high probability:

$$|\Delta_j^\top \mathbf{v} - \mathbb{E}[\Delta_j^\top \mathbf{v}]| = \mathcal{O}\left(\frac{\log(n)}{p\sqrt{d}}\right). \quad (77)$$

This proves Part (i). Part (ii) follows from Lemma 14 in Ba et al. (2022).

It remains to show Part (iii). The same proof as in Proposition 2 (Eq. (53)) implies that, with high probability,

$$\langle \mathbf{g}_i, \mathbf{g}_j \rangle = \mathbb{E}[\langle \mathbf{g}_i, \mathbf{g}_j \rangle] + \mathcal{O}\left(\frac{\text{polylog}(d)}{p^2\sqrt{d}}\right), \quad (78)$$

and hence we only need to bound the expectation $\mathbb{E}[\langle \mathbf{g}_i, \mathbf{g}_j \rangle]$. In turn, the decomposition of Equation (47) still holds, and we get

$$\mathbb{E}[\langle \mathbf{g}_i, \mathbf{g}_j \rangle] \leq \|\mathbb{E}[\mathbf{g}_i]\| \|\mathbb{E}[\mathbf{g}_j]\| + \frac{1}{np^2} \mathbb{E}[\|\mathbf{z}\|^2 \sigma'(\langle \mathbf{w}_i^0, \mathbf{z} \rangle) \sigma'(\langle \mathbf{w}_j^0, \mathbf{z} \rangle) f^*(\mathbf{z})^2] \quad (79)$$

Since $\mu_1 = 0$, the bound of Lemma 7 becomes

$$\|\mathbb{E}[\mathbf{g}_i]\| \leq \frac{\pi_i}{p} = \mathcal{O}\left(\frac{\log(d)}{p\sqrt{d}}\right),$$

and it remains to bound the cross term. The main argument is the following lemma, which is the generalization (with an identical proof) of Lemma D.4 in Arnaboldi et al. (2023):

Lemma 13. *Let $N \geq 0$ be fixed, and f_1, \dots, f_N be a sequence of functions with bounded first and second derivatives. Consider the function on $N \times N$ matrices*

$$F(\Sigma) = \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, \Sigma)}[f_1(z_1) \dots f_N(z_N)] \quad (80)$$

Then, for Σ, Σ' two semidefinite positive matrices with unit diagonal, we have

$$|F(\Sigma) - F(\Sigma')| \leq C \|\Sigma - \Sigma'\|_\infty. \quad (81)$$

Now, we first have by the same arguments as in Lemma 7

$$\mathbb{E}[\|\mathbf{z}\|^2 \sigma'(\langle \mathbf{w}_i^0, \mathbf{z} \rangle) \sigma'(\langle \mathbf{w}_j^0, \mathbf{z} \rangle) f^*(\mathbf{z})^2] = d \mathbb{E}[\sigma'(\langle \mathbf{w}_i^0, \mathbf{z} \rangle) \sigma'(\langle \mathbf{w}_j^0, \mathbf{z} \rangle) f^*(\mathbf{z})^2] + \mathcal{O}(\sqrt{d}),$$

so we only to bound the first term of the RHS. Expanding the definition of f^* , the latter is a sum of k^2 terms of the form

$$\mathbb{E}[\sigma'(\langle \mathbf{w}_i^0, \mathbf{z} \rangle) \sigma'(\langle \mathbf{w}_j^0, \mathbf{z} \rangle) \sigma_k^*(\langle \mathbf{w}_k^*, \mathbf{z} \rangle) \sigma_{k'}^*(\langle \mathbf{w}_{k'}^*, \mathbf{z} \rangle)],$$

which falls under the framework Lemma 13 for $N = 4$. In particular, since $\mu_1 = 0$, $F(\Sigma) = 0$ whenever we have $\Sigma_{1i} = \Sigma_{2j} = 0$ for $i \neq 1, j \neq 2$. Hence, by an application of Lemma 13, we have

$$\mathbb{E}[\sigma'(\langle \mathbf{w}_i^0, \mathbf{z} \rangle) \sigma'(\langle \mathbf{w}_j^0, \mathbf{z} \rangle) \sigma_k^*(\langle \mathbf{w}_k^*, \mathbf{z} \rangle) \sigma_{k'}^*(\langle \mathbf{w}_k^*, \mathbf{z} \rangle)] \leq C \max(\langle \mathbf{w}_i^0, \mathbf{w}_j^0 \rangle, \pi_i, \pi_j) \leq C \frac{\log(d)}{\sqrt{d}}$$

with high probability, which ends the proof. \square

We next prove that the norm of \mathbf{w}_i^1 after the first gradient step possesses a simplified dimension-independent limit:

Lemma 14. *Suppose $n = \Theta(d)$. Then, there exists a constant C , such that for any neuron i , with high-probability as $d \rightarrow \infty$, with step-size η :*

$$\|\mathbf{w}_i^1\|^2 = 1 + \eta C a_i^2 + \mathcal{O}\left(\frac{\text{polylog } d}{\sqrt{d}}\right) \quad (82)$$

Proof. Recall Equation 47:

$$\mathbb{E}[\|\mathbf{g}_i\|^2] = \frac{n(n-1)}{n^2} \|\mathbb{E}[\mathbf{g}_i]\|^2 + \frac{a_i^2}{np^2} \mathbb{E}[\|\mathbf{z}^\nu\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] \quad (83)$$

Lemma 7 implies that $\|\mathbb{E}[\mathbf{g}_i]\|^2$ is approximately $a_i^2 X_i \cdot \frac{\|\boldsymbol{\pi}_i^0\|^{2(\ell-1)}}{p^2}$ for a random variable X_i . When $\ell = 1$, X_i simply reduces to a constant depending only on g^* . The second term can be decomposed as:

$$\frac{a_i^2}{np^2} \mathbb{E}[\|\mathbf{z}^\nu\|^2 \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] = \frac{a_i^2}{np^2} \mathbb{E}[d \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] + \mathbb{E}[(d - \|\mathbf{z}^\nu\|^2) \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] + \quad (84)$$

Let $m_0^i \in R^r$ denote the vector of overlaps $\langle \mathbf{w}_i^0, \mathbf{w}_1^* \rangle, \dots, \langle \mathbf{w}_i^0, \mathbf{w}_k^* \rangle$. By Holder's inequality, the second term is of order $\mathcal{O}(\frac{1}{\sqrt{d}})$ while through a change of variables, the first term can be expressed as a function of the overlaps $\langle \mathbf{w}_i^0, \mathbf{w}_j^* \rangle$ for $j \in [r]$:

$$\begin{aligned} \frac{a_i^2}{np^2} \mathbb{E}[d \sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] &= \frac{a_i^2 d}{np^2} \mathbb{E}[\sigma'(\langle \mathbf{w}_i, \mathbf{z}^\nu \rangle)^2 f^*(\mathbf{z}^\nu)^2] = \frac{a_i^2 d}{np^2} F_{\sigma, g^*}(M_0) \\ &= \frac{a_i^2 d}{np^2} F_{\sigma, g^*}(0) + \frac{1}{\sqrt{d}} \end{aligned}$$

\square

C.7 SECOND STEP: PROOF SKETCH FOR THEOREM 3

Before providing detailed proof of Theorem 3 for general polynomial activation functions, and a general number of steps, we illustrate the essential idea by analyzing the second gradient step. Let \mathbf{Z}^0 denote the batch of inputs used for the first gradient step. We condition on \mathbf{Z}^0 and assume that the high-probability events in Lemma 12 hold. We independently sample another batch of n training inputs \mathbf{Z} and perform the gradient update:

$$\mathbf{g}_j^1 = -\nabla_{\mathbf{w}_j} \mathcal{L} \left(\hat{f}(\mathbf{z}^\nu; W^1, \mathbf{a}), f^*(\mathbf{z}^\nu) \right) \quad (85)$$

However, unlike the first gradient step, the weights \mathbf{w}^1 are no-longer approximately orthonormal across neurons and contain significant correlation along the teacher subspace.

We have:

$$\mathbf{w}_j^1 = \eta \frac{a_j \mu_1}{\sqrt{p}} \mathbf{v} + \mathbf{w}_j^0 + \Delta_j, \quad (86)$$

where $\mathbf{v} = \frac{1}{n} \sum_{i=1}^n y_i \mathbf{z}_i$. By theorem 2, we have that the projection of \mathbf{v} along the target subspace V^* converges in probability to $C_1(f)$. Let $\mathbf{v}^* = C_1(f)$. We show that the alignment of \mathbf{v} along \mathbf{v}^* affects the components of the second gradient step along the teacher subspace, allowing the gradient to be sensitive to directions linearly coupled with \mathbf{v}^* in the target function.

We proceed by analyzing the projection of the above update along a direction in the teacher subspace. Let $\mathbf{v}_j = P_{V^*}(\mathbf{w}_j^1)$ and consider the decomposition $\mathbf{w}_j^1 = \mathbf{v}_j + P_{V^*}^\perp(\mathbf{w}_j^1)$. We further have from Lemma 14 that $\|P_{V^*}^\perp(\mathbf{w}_j^1)\|_2^2$ concentrates to a positive bounded value c_j depending only on a_j . For each sample, \mathbf{z}_i let $\kappa_i = \langle \mathbf{v}_j, \mathbf{z}_i \rangle$ denote the projection along the ‘‘signal’’ \mathbf{v}_j .

We now introduce the following function for $z \in \mathbb{R}$:

$$\sigma_{\kappa,j}(z) = \sigma(c_j z + \eta \kappa). \quad (87)$$

Define $\mu_{1,\kappa,j} = \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\sigma_{\kappa,j}(z) z]$. Further, let:

$$\sigma'_{>1,\kappa}(u) = \sigma'_{\kappa,j}(u) - \mu_{1,\kappa,j}. \quad (88)$$

From Lemma 5, we have that $P_{V^*}(\mathbf{v}) \xrightarrow{\mathbb{P}} \mathbf{v}^*$. Now, let $\mathbf{u} \in V^*$ be a direction in the teacher subspace orthogonal to \mathbf{v}^* . Using, equation (42), we have:

$$\begin{aligned} \mathbb{E} [\langle \mathbf{u}, \mathbf{g}_j^1 \rangle] &= \mathbb{E} \left[(f^*(z) - \hat{f}(z, \mathbf{W}^1, \mathbf{a})) \sigma'(\langle z, \mathbf{w}_j^1 \rangle)_j \langle z, \mathbf{u} \rangle \right] \\ &= \mathbb{E} \left[(f^*(z) \mu_{1,\langle z, \mathbf{v}_j \rangle}(\langle z_i, \mathbf{u} \rangle)) \right] - \mathbb{E} \left[\hat{f}(z, \mathbf{W}^1, \mathbf{a}) \sigma'(\langle z_i, \mathbf{w}^1 \rangle)_j \langle z_i, \mathbf{u} \rangle \right] \end{aligned} \quad (89)$$

Where in the first term we took the expectation over $P_{V^*}^\perp(\mathbf{w}^1)$ since it is orthogonal to the teacher-subspace.

The second term can be expressed as:

$$\langle \mathbb{E} \left[\hat{f}(z, \mathbf{W}^1, \mathbf{a}) \sigma'(\langle z, \mathbf{w}^1 \rangle)_j \mathbf{z}_i \right], \mathbf{u} \rangle \quad (90)$$

We have that $\hat{f}(z, \mathbf{W}^1, \mathbf{a})$ depends only on the directions $\mathbf{w}_1^1, \dots, \mathbf{w}_p^1$. By Lemma 12, each of the directions, satisfies $\langle \mathbf{w}_i, \mathbf{u} \rangle = \mathcal{O}(\frac{\text{polylog}(d)}{p\sqrt{d}})$. Furthermore, one can show that $(\mathbb{E} \left[\hat{f}(z, \mathbf{W}^1, \mathbf{a}) \sigma'(\langle z, \mathbf{w}^1 \rangle)_j \mathbf{z}_i \right])$ lies in the span of $\mathbf{w}_1^1, \dots, \mathbf{w}_p^1$. Therefore, $\mathbb{E} \left[\hat{f}(z, \mathbf{W}^1, \mathbf{a}) \sigma'(\langle z_i, \mathbf{w}^1 \rangle)_j \mathbf{z}_i \right]^\top \mathbf{u} \xrightarrow{d \rightarrow \infty} 0$.

Now, consider the first term i.e $\mathbb{E} [f^*(z) \mu_{1,\langle z, \mathbf{v}_j \rangle}(\langle z, \mathbf{u} \rangle)]$. Let $\mathbf{v}^*, \mathbf{u}, \mathbf{u}'_1, \dots, \mathbf{u}'_{d-2}$ be an orthonormal basis of \mathbb{R}^d . Without loss of generality, assume that $\mathbf{v}^*, \mathbf{u}, \dots, \mathbf{u}'_{r-2}$. span the teacher subspace V^* . We express y using the product Hermite decomposition under the above basis:

$$y = f^*(z) = \sum_{j_1, \dots, j_r=1}^{\infty} \frac{c_{j_1, \dots, j_r}^*}{j_1! j_2! \dots j_r!} \text{He}_{j_1}(\langle \mathbf{v}^*, z \rangle) \text{He}_{j_2}(\langle \mathbf{u}, z \rangle) \dots \text{He}_{j_r}(\langle \mathbf{u}_{r-2}, z \rangle). \quad (91)$$

Since $\mathbf{v}_j \xrightarrow{\mathbb{P}} \mathbf{v}^*$ and $\mathbf{u} \perp \mathbf{u}'_1, \dots, \mathbf{u}'_{r-2}$, only the terms of the form $\text{He}_{j_1}(\langle \mathbf{v}^*, z \rangle) \text{He}_{j_2}(\langle \mathbf{u}, z \rangle)$ contribute to the expectation $\mathbb{E} [y \mu_{1,\langle z, \mathbf{v}_j \rangle}(\langle z, \mathbf{u} \rangle)]$ in the limit $d \rightarrow \infty$. Consider the contribution of one such term:

$$\mathbb{E} [\text{He}_{j_1}(\langle \mathbf{v}^*, z \rangle) \text{He}_{j_2}(\langle \mathbf{u}, z \rangle) \mu_{1,\langle z, \mathbf{v}_j \rangle, j}(\langle z, \mathbf{u} \rangle)] \rightarrow \mathbb{E} [\text{He}_{j_1}(\langle \mathbf{v}^*, z \rangle) \text{He}_{j_2}(\langle \mathbf{u}, z \rangle) \mu_{1,\langle z, \mathbf{v}^* \rangle, j}(\langle z, \mathbf{u} \rangle)] \quad (92)$$

Suppose $j_2 \neq 1$, then $\mathbb{E}[\text{He}_{j_2}(\langle \mathbf{u}, \mathbf{z} \rangle) \langle \mathbf{z}_i, \mathbf{u} \rangle] = 0$. Therefore, the non-zero contributions arise from terms of the form $\text{He}_{j_1}(\langle \mathbf{v}^*, \mathbf{z} \rangle) \langle \mathbf{u}, \mathbf{z} \rangle$. It can be checked that directions \mathbf{u} having non-zero terms of this form span U_2^* as defined in Theorem 3. However, in general, the RHS of equation 92 might be 0 for some choices of σ and a_j . Moreover, such non-zero contributions might cancel each other for a chosen direction in U_2^* . Furthermore, to obtain high-probability result on the alignment along U_t^* for a general number of t steps, one needs to quantitatively propagate the expectations and concentration bounds on the projections and norms of \mathbf{W} , and show that the magnitude of the projections can be bounded independent of the dimension. We tackle these issues in the next section and provide a full proof of Theorem 3.

C.8 PROOF OF THEOREM 3

The proof proceeds by induction on the number of time-steps t . To avoid certain degeneracy conditions in the proof, we restrict ourselves to polynomial activations. Let U_t^* be the learned subspace at time-step t according to the definition 2.

Let $Q_t \in \mathbb{R}^{p \times p}$ denote the overlap matrix for weights of the first-layer neurons at time t , i.e. $Q_{i,j}^t = \langle \mathbf{w}_i^t, \mathbf{w}_j^t \rangle \forall i, j \in [p]$. Let $M_t \in \mathbb{R}^{r \times p}$ denote the target-network overlap matrix i.e. $M_{i,j}^t = \langle \mathbf{w}_i^t, \mathbf{w}_j^t \rangle \forall i \in [p], j \in [r]$. Let $\mathbf{W}^* \in \mathbb{R}^{r \times d}$ denote the matrix with rows $\mathbf{w}_1^*, \dots, \mathbf{w}_r^*$.

We denote by \mathbf{Z}_t , the batch of input sampled at time $t \in [T]$. By assumption $\mathbf{Z}_1, \dots, \mathbf{Z}_T$ are independent. Let \mathcal{F}_t denote the natural filtration associated to $\mathbf{Z}_1, \dots, \mathbf{Z}_T$, i.e. \mathcal{F}_t is the σ -algebra generated by $\mathbf{Z}_1, \dots, \mathbf{Z}_t$, and let μ_t denote the corresponding joint-measure of $\mathbf{Z}_1, \dots, \mathbf{Z}_t$. We let \mathbf{g}_i^t denote the gradient for the i th neuron at time t obtained using the batch \mathbf{Z}^{t+1} .

For any time t , let r_t denote the dimension of U_t^* and let $\mathbf{W}_t^* \in \mathbb{R}^{r_t \times d}$ denote a matrix with rows forming a basis of U_t^* , such that $(\mathbf{W}_t^*)^\top \mathbf{W}_t^*$ is independent of d, n . Thus, \mathbf{W}_t^* represents a dimension independent basis of U_t^* . Let $\mathbf{v}_{j,a} \in \mathbb{R}^{r_t}$ denote the projections of \mathbf{w}_j^t along \mathbf{W}_t^* i.e. $\mathbf{v}_{j,a} = \mathbf{W}_t^* \mathbf{w}_j^t$. Similarly, for an input $\mathbf{z} \in \mathbb{R}^d$, we denote the projection of \mathbf{z} along \mathbf{W}_t^* by $\kappa = \mathbf{W}_t^* \mathbf{z}$. In what follows, we shall say that a sequence of events \mathcal{E}_n occurs with high-probability as $n, d \rightarrow \infty$ if there exist constants $c, C > 0$ such that $\mathbb{P}(\mathcal{E}_n) \geq 1 - Cpe^{-c \log(n)^2} + Cpe^{-c \log(d)^2}$

At any timestep $t \geq 1$, we prove that the following statements hold with high probability w.r.t μ_t :

- (i) $Q^t = \tilde{Q}_a^t + \mathcal{O}(\frac{\text{polylog}d}{\sqrt{d}})$, $M^t = \tilde{M}_a^t + \mathcal{O}(\frac{\text{polylog}d}{\sqrt{d}})$, where $\tilde{Q}_a^t, \tilde{M}_a^t$ denote dimension-independent matrices with each entry being a polynomial dependent only on \mathbf{a}, t of \mathbf{w}_i^t , dependent on the second layer i, \mathbf{a} .
- (ii) Let $\mathbf{v} \in U_t^*$, with $\|\mathbf{v}\| = 1$ be arbitrary. Denote by $\mathbf{v}^m \in \mathbb{R}^k$, the components of \mathbf{v} along $\mathbf{w}_1^*, \dots, \mathbf{w}_r^*$ i.e. $\mathbf{v}^m = \mathbf{W}^* \mathbf{v}$. Then there exists an almost surely positive random variable $q_{t,\mathbf{v},\mathbf{a}}$, independent of d, n such that $\langle \mathbf{w}_i, \mathbf{v} \rangle = q_{t,\mathbf{v},\mathbf{a}} + \mathcal{O}(\frac{\text{polylog}}{\sqrt{d}})$. Furthermore, $q_{t,\mathbf{v},\mathbf{a}}$ is a non-constant polynomial in $\mathbf{a}, v_1, \dots, v_k$.
- (iii) For any $\mathbf{v} \in U_t^{\perp*} \cap V^*$, $|\langle \mathbf{w}_i, \mathbf{v} \rangle| = \mathcal{O}(\frac{\text{polylog}}{\sqrt{d}})$, with high probability, for all $i \in [p]$.

Proof. We proceed by induction over t . Suppose that the statements hold at some timestep t . We start by proving that (i) holds at time $t + 1$ in expectation:

Lemma 15. $\mathbb{E}[Q_t] = \tilde{Q}_a^t + \mathcal{O}(\frac{\text{polylog}d}{\sqrt{d}})$ and $\mathbb{E}[M_t] = \tilde{M}_a^t + \mathcal{O}(\frac{\text{polylog}d}{\sqrt{d}})$ where each entry of $\tilde{Q}_a^t, \tilde{M}_a^t$ is a polynomial of \mathbf{a} with degree independent of d, n .

Proof. Recall that:

$$\begin{aligned}
Q_{i,j}^{t+1} &= \langle \mathbf{w}_i^{t+1}, \mathbf{w}_j^{t+1} \rangle \\
&= Q_{i,j}^t + \eta \langle \mathbf{g}_i^t, \mathbf{w}_j^t \rangle + \eta \langle \mathbf{w}_i^t, \mathbf{g}_j^t \rangle + \eta^2 \langle \mathbf{g}_i^t, \mathbf{g}_j^t \rangle. \\
M_{i,j}^{t+1} &= \langle \mathbf{w}_i^*, \mathbf{w}_j^{t+1} \rangle \\
&= M_{i,j}^t + \eta \langle \mathbf{w}_i^*, \mathbf{g}_j^t \rangle
\end{aligned} \tag{93}$$

By the induction hypothesis, the entries of $\mathbb{E}[Q^t]$, $\mathbb{E}[M^t]$ converge with high-probability to polynomial limits with error $\mathcal{O}\left(\frac{\text{polylog } d}{d}\right)$. Therefore, it suffices to show that $\mathbb{E}[\langle \mathbf{g}_i^t, \mathbf{w}_j^t \rangle]$, $\mathbb{E}[\langle \mathbf{w}_i^t, \mathbf{w}_j^t \rangle]$, $\mathbb{E}[\langle \mathbf{g}_i^t, \mathbf{g}_j^t \rangle]$ converge to dimension-independent polynomial limits. First, consider the case $i = j$. We have, analogous to Equation 47

$$\mathbb{E}[\|\mathbf{g}_i^t\|^2] = \underbrace{\frac{1}{np^2} \mathbb{E}[\|\mathbf{z}^\nu\|^2 \sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2]}_{T_1} + \underbrace{\frac{n(n-1)}{n^2} \|\mathbb{E}[\mathbf{g}_i]\|^2}_{T_2} \tag{94}$$

The first term T_1 can be decomposed as follows:

$$\begin{aligned}
\frac{1}{np^2} \mathbb{E}[\|\mathbf{z}^\nu\|^2 \sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2] &= \frac{d}{np^2} \mathbb{E}[\sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2] \\
&\quad + \frac{1}{np^2} \mathbb{E}[(d - \|\mathbf{z}^\nu\|^2) \sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2].
\end{aligned}$$

Similar to equation 50, Holder's inequality implies that conditioned on the event in \mathcal{F}_t of Q_t, M_t being bounded independent of d, n , the second term is of order $\mathcal{O}(\frac{\sqrt{d}}{n}) = \mathcal{O}(\frac{1}{\sqrt{d}})$. Consider the first term, conditioned on \mathcal{F}_t .

$$\frac{d}{np^2} \mathbb{E}[\sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2 | \mathcal{F}_t] \tag{95}$$

By assumption, $\frac{d}{np^2} = \frac{\alpha}{p^2}$ for some constant α . Therefore, by definition of $f^*(\mathbf{z}^\nu)$ and $\hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a})$, the term inside the expectation only depends on the overlaps of \mathbf{z}^ν with the neurons and teacher subspace i.e $\langle \mathbf{w}_1^t, \mathbf{z}^\nu \rangle, \dots, \langle \mathbf{w}_p^t, \mathbf{z}^\nu \rangle, \langle \mathbf{w}_1^*, \mathbf{z}^\nu \rangle, \dots, \langle \mathbf{w}_k^*, \mathbf{z}^\nu \rangle$. By a change of variables the above term can therefore be expressed as an expectation w.r.t the $k + j$ correlated variables corresponding to the above overlaps.

Concretely, we have:

$$\frac{d}{np^2} \mathbb{E}[\sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2 | \mathcal{F}_t] = F_g(Q_t, M_t), \tag{96}$$

for some function $F : \mathbb{R} \rightarrow \mathbb{R}$

Lemma 16. F_g is a polynomial in Q_t, M_t independent of d, n .

Proof. By assumption, σ' and f^* are polynomials in $\langle \mathbf{w}_1^t, \mathbf{z}^\nu \rangle, \dots, \langle \mathbf{w}_p^t, \mathbf{z}^\nu \rangle$ and $\langle \mathbf{w}_1^*, \mathbf{z}^\nu \rangle, \dots, \langle \mathbf{w}_k^*, \mathbf{z}^\nu \rangle$ respectively. Therefore, $\sigma'(\langle \mathbf{w}_i^t, \mathbf{z}^\nu \rangle)^2 (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}))^2$ is a polynomial in the zero mean correlated Gaussian variables $\langle \mathbf{w}_1^t, \mathbf{z}^\nu \rangle, \dots, \langle \mathbf{w}_p^t, \mathbf{z}^\nu \rangle, \langle \mathbf{w}_1^*, \mathbf{z}^\nu \rangle, \dots, \langle \mathbf{w}_k^*, \mathbf{z}^\nu \rangle$. Therefore, by Wick's theorem, F_g is a polynomial in Q_t, M_t . \square

By the induction hypothesis, with high-probability, $Q_t = \tilde{Q}_{t,\mathbf{a}} + \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$ and $M_t = \tilde{M}_{t,\mathbf{a}} + \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$, where $\tilde{Q}_{t,\mathbf{a}}, \tilde{M}_{t,\mathbf{a}}$ denote deterministic matrices with entries being polynomial functions of \mathbf{a} . By propagating the errors through the polynomial F_g , we obtain that $F_g(Q_t, M_t) = F_g(\tilde{Q}_{t,\mathbf{a}}, \tilde{M}_{t,\mathbf{a}}) + \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$.

Next, consider the term T_2 in Equation 94. By repeatedly applying Stein's Lemma w.r.t terms $\langle \mathbf{w}_i^t, \mathbf{z} \rangle$ for $i \in [p]$ and $\langle \mathbf{w}^*, \mathbf{z} \rangle$ for $j \in [r]$, analogous to Lemma 4, $\mathbb{E}[\mathbf{g}_i]$, can be expressed as a linear combination of $\mathbf{w}_1^t, \dots, \mathbf{w}_p^t$ and $\mathbf{w}_1^*, \dots, \mathbf{w}_r^*$. Furthermore, by Wick's theorem, the coefficients are independent of d and polynomial in Q_t, M_t . Therefore, propagating errors from time t , T_2 can be approximated by polynomials in Q_t, M_t with error $\mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$

Similarly, the terms $\mathbb{E}[\langle \mathbf{g}_i^t, \mathbf{g}_j^t \rangle], \mathbb{E}[\langle \mathbf{w}_i^*, \mathbf{g}_j^t \rangle]$ converge with high-probability to dimension-independent polynomials in Q_t, M_t \square

Next, we prove that (ii) and (iii) holds in expectation:

Lemma 17. *Let $\mathbf{v} \in V^*$, with $\|\mathbf{v}\| = 1$ be arbitrary with components $\mathbf{v}^m \in \mathbb{R}^r$ along $\mathbf{w}_1^*, \dots, \mathbf{w}_r^*$, then $\mathbb{E}[\langle \mathbf{v}, \mathbf{g}_j^t \rangle] = h(\mathbf{v}^m, \mathbf{a}, Q_t, M_t) + \mathcal{O}(\frac{\text{polylog } d}{p\sqrt{d}})$, where $h(\mathbf{v}^m, \mathbf{a}, Q_t, M_t)$ satisfies:*

- (i) $h(\mathbf{v}^m, \mathbf{a}, Q_t, M_t)$ is non-zero, almost surely over \mathbf{a} if $\mathbf{v} \in U_{t+1}^*$.
- (ii) $h(\mathbf{v}^m, \mathbf{a}, Q_t, M_t) = 0$ otherwise.

Consider the gradient w.r.t the j th neuron's parameters:

$$\mathbf{g}_j^t = -\nabla_{\mathbf{w}_j} \mathcal{L} \left(\hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a}), f^*(\mathbf{z}^\nu) \right) = \frac{1}{n} a_j \sum_{\nu=1}^n \mathbf{z}^\nu (f^*(\mathbf{z}^\nu) - \hat{f}(\mathbf{z}^\nu; \mathbf{W}^t, \mathbf{a})) \sigma'(\langle \mathbf{z}^\nu, \mathbf{w}_j^t \rangle) \quad (97)$$

\square

Suppose that $\mathbf{v} \in U_{t+1}^* \cap (U_t^*)^\perp$ i.e when \mathbf{v} is a new direction not yet learned upto time t . Using, equation (97), the expectation $\mathbb{E}[\langle \mathbf{v}, \mathbf{g}_j^t \rangle]$ can be expressed as:

$$\mathbb{E}[\langle \mathbf{v}, \mathbf{g}_j^t \rangle] = \mathbb{E} \left[a_j (f^*(\mathbf{z}) - \hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a})) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]. \quad (98)$$

We first consider the term $\mathbb{E} \left[\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]$. Through a change of variables, and Wick's theorem, one obtains that $\mathbb{E} \left[\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]$ is a polynomial in Q and the overlaps $\langle \mathbf{w}^i, \mathbf{v} \rangle$ for $i \in [p]$ having value 0 when $\langle \mathbf{w}^i, \mathbf{v} \rangle = 0$ for all $i \in [p]$. By the induction hypothesis, $\langle \mathbf{w}^i, \mathbf{v} \rangle = \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$ with high probability. Therefore $\mathbb{E} \left[\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle | \mathcal{F}_t \right] = \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$ with high probability. Similarly, $\mathbb{E} \left[\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle | \mathcal{F}_t \right] = \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$ holds when $\mathbf{v} \notin U_{t+1}^*$.

Now, consider the term $\mathbb{E} \left[a_j f^*(\mathbf{z}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]$. First, using Fubini's theorem, we take the expectation w.r.t the component \mathbf{z}^\perp of \mathbf{z} in $V^{\star\perp}$.

Recall that $\mathbf{v}_{j,\mathbf{a}} = \mathbf{W}_t^* \mathbf{w}_j^t$ and $\kappa = \mathbf{W}_t^* \mathbf{z}$. The resulting expectation converges in probability to a function of κ :

$$\mathbb{E}_{\mathbf{z}^\perp} \left[a_j f^*(\mathbf{z}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right] = \mathbb{E}_\kappa \left[a_j f^*(\mathbf{z}) f_1(\mathbf{a}, \kappa) \langle \mathbf{z}, \mathbf{v} \rangle \right], \quad (99)$$

where $f_1(\mathbf{a}, \kappa)$ is defined as follows:

$$\begin{aligned} f_1(\mathbf{a}, \kappa) &= \mathbb{E}_{\mathbf{z}^\perp} [\sigma'(\mathbf{z}^\top \mathbf{w}_j^t)] \\ &= \mathbb{E}_{u \sim \mathcal{N}(0,1)} [\sigma'(c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)] \end{aligned}$$

, where $c_{j,\mathbf{a}}$ denotes the norm of \mathbf{w}_j^t along the orthogonal complement of V^* . $f_1(\mathbf{a}, \kappa)$ generalizes the ‘‘shifted-hermite’’ $\mu_{1,\kappa,j}$ defined in the section C.7. By assumption on σ , $\sigma'(c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)$ is a polynomial in $c_{j,\mathbf{a}}u, \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle$. Furthermore, only the odd terms in $c_{j,\mathbf{a}}u$ are zero in expectation $u \sim \mathcal{N}(0, 1)$. Therefore, $f_1(\mathbf{a}, \kappa)$ is a polynomial in $\langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle$ and $c_{j,\mathbf{a}}$ with only even degree terms in $c_{j,\mathbf{a}}$. By the induction hypothesis, $c_{j,\mathbf{a}}^2$ converges in probability to a polynomial in \mathbf{a} . Therefore $f_1(\mathbf{a}, \kappa)$ is a polynomial in \mathbf{a}, κ .

Subsequently, we consider the expectation w.r.t $\langle \mathbf{z}, \mathbf{v} \rangle$, at a fixed value of κ . Define the following function of κ :

$$f_2(\kappa) = \mathbb{E}_{\langle \mathbf{z}, \mathbf{v} \rangle} [y \langle \mathbf{z}, \mathbf{v} \rangle | \kappa]. \quad (100)$$

Using the tower law of expectation, we obtain:

$$\mathbb{E} [a_j f^*(\mathbf{z}) \sigma'(\mathbf{z}^\top \mathbf{w}_j^t) \langle \mathbf{z}, \mathbf{v} \rangle] = \mathbb{E}_\kappa [f_1(a_j, \kappa) f_2(\kappa)], \quad (101)$$

When $\mathbf{v} \notin U_{t+1}^*$, $f_2(\kappa)$ is identically 0 and the above expectation vanishes.

We aim to show that the above expectation does not vanish except for a_i belonging to a zero-measure set. By the definition of subspace conditioning (definition 2), $\exists \kappa > 0$ such that $\mathbb{E}_{\langle \mathbf{z}, \mathbf{v} \rangle} [f^*(\mathbf{z}) \mathbf{z} | \kappa]$ has non-zero overlap with \mathbf{v} .

Therefore, $f_2(\kappa)$ is not identically zero. Furthermore, since f^* is a polynomial by assumption, and $\mathbf{v} \perp V^*$, a rotation of basis implies that f_2 is a polynomial in κ . Let $\mathcal{S}_{y,t}$ be the set of degrees $s \in \mathbb{N}_0$ such that $\mathbb{E}_{f_2(\kappa)\kappa^s} [\kappa] \neq 0$. Since f_2 is not identically 0, we have that $\mathcal{S}_{y,t} \neq \emptyset$.

Now, recall that:

$$\begin{aligned} f_1(\mathbf{a}, \kappa) &= \mathbb{E}_{u \sim \mathcal{N}(0,1)} [\sigma'(c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)] \\ &= \mathbb{E}_{u \sim \mathcal{N}(0,1)} \left[\sum_{r=0}^{\deg(\sigma)-1} (r+1) b_{r+1} (c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)^r \right] \\ &= \sum_{r=0}^{\deg(\sigma)-1} (r+1) b_{r+1} \mathbb{E}_{u \sim \mathcal{N}(0,1)} [(c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)^r]. \end{aligned}$$

Now, let $s \in \mathcal{S}_{y,t}$ be arbitrary. By assumption, $\deg(\sigma) - 1 \geq s$. Let $p_s(\mathbf{a})$ denote the coefficient of κ^s in $f_1(\mathbf{a}, \kappa)$. Since $c_{j,\mathbf{a}}, \mathbf{v}_{j,\mathbf{a}}$ are non-constant polynomials in \mathbf{a} , the coefficient of κ^s in $(c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)^r$ is a non-constant polynomial in \mathbf{a} for any r such that $r - s$ is even. Furthermore, the degree of the coefficient of κ^s in $(c_{j,\mathbf{a}}u + \langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)^r$ is strictly increasing in r . Therefore, for any $s \in \mathcal{S}_{y,t}$, $p_s(\mathbf{a})$ is a non-constant polynomial in \mathbf{a} . Now, consider the term in $p_s(\mathbf{a})$ with the least degree in a_j . From the definition of $\mathbf{v}_{j,\mathbf{a}}$, we have that $\mathbf{v}_{j,\mathbf{a}} = 0$ whenever $a_j = 0$. Let d_j denote the least $s \in \mathbb{N}_0$ such that the coefficient of a_j^s in $\langle \mathbf{v}_{j,\mathbf{a}}, \kappa \rangle$ is non-zero. We have that $d_j > 0$. Consequently, the minimum degree of a_j in $(c_{j,\mathbf{a}})^q (\langle \kappa, \mathbf{v}_{j,\mathbf{a}} \rangle)^s$, is $(d_j)^s$ for any q . Therefore, the minimum degree of $p_s(\mathbf{a})$ is strictly increasing in s . This implies that $p_s(\mathbf{a})$ are linearly independent for $s = 1, \dots, \deg(\sigma) - 1$.

Now, consider the function defined above in Equation 101:

$$h(t, \mathbf{a}) = \mathbb{E}_\kappa [f_1(\mathbf{a}, \kappa) f_2(\kappa)]. \quad (102)$$

By expanding f_1, f_2 along κ , the coefficient of κ^s for each $s \in \mathcal{S}_{y,t}$ results in a non-constant polynomial in \mathbf{a} . We obtain:

$$h(t, \mathbf{a}) = \sum_{s \in \mathcal{S}_{y,t}} c_s p_s(\mathbf{a}), \quad (103)$$

where c_s denote constants independent of d, n . Therefore, we have that $h(\mathbf{a})$ is a non-constant polynomial in \mathbf{a} . Using Fubini's theorem, we have that the set of zeros of non-zero multivariate polynomials has 0 measure w.r.t the Lebesgue measure (for a generalization, see (Mityagin, 2020)), we obtain that $q(\mathbf{a}) \neq 0$ almost surely.

Now, suppose that $\mathbf{v} \in (U_t^*)$, i.e when \mathbf{v} is an already learned direction. By the induction hypothesis, $\langle \mathbf{w}_i^t, \mathbf{v} \rangle$ converges to a non-constant polynomial in \mathbf{a} . Consider the term $\frac{a_j}{\sqrt{p}} \mathbb{E} \left[\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]$ in $\langle \mathbf{g}_i^t, \mathbf{v} \rangle$. By expanding $\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a})$ we obtain:

$$\frac{a_j}{\sqrt{p}} \mathbb{E} \left[\hat{f}(\mathbf{z}; \mathbf{W}^t, \mathbf{a}) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right] = \frac{a_j}{p} \sum_{i=1}^p a_i \mathbb{E} \left[\sigma(\langle \mathbf{w}_i^t, \mathbf{z} \rangle) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right] \quad (104)$$

The term correspondign to the j_{th} neuron has the form:

$$\frac{a_j^2}{p} \mathbb{E} \left[\sigma(\langle \mathbf{w}_j^t, \mathbf{z} \rangle) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right] \quad (105)$$

By Wick's theorem and assumption on σ , $\mathbb{E} \left[\sigma(\langle \mathbf{w}_j^t, \mathbf{z} \rangle) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]$ is a non-zero polynomial in $\langle \mathbf{w}_j^t, \mathbf{v} \rangle$. Let d_j be the degree of a_j in $\langle \mathbf{w}_j^t, \mathbf{v} \rangle$. Then, the degree of a_j in $\frac{a_j^2}{p} \mathbb{E} \left[\sigma(\langle \mathbf{w}_j^t, \mathbf{z} \rangle) \sigma'(\langle \mathbf{z}, \mathbf{w}_j^t \rangle) \langle \mathbf{z}, \mathbf{v} \rangle \right]$ is at-least $d_j + 2$. Proceeding similarly for the other terms, one can show that the degree of a_j in $\langle \mathbf{g}_i^t, \mathbf{v} \rangle$ is strictly larger than in $\langle \mathbf{w}_i^t, \mathbf{u} \rangle$. This ensures that $\langle \mathbf{w}_i^{t+1}, \mathbf{v} \rangle = \langle \mathbf{g}_i^t, \mathbf{v} \rangle + \eta \langle \mathbf{g}_i^t, \mathbf{v} \rangle$ remains a non-constant polynomial upto error $\mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$. Therefore, almost surely over \mathbf{a} , a direction is not "un-learned". Finally, by decomposing along a general $\mathbf{v} \in U_{t+1}^*$, along U_t^* and $U_{t+1}^* \cap (U_t^*)^\perp$, one obtains that points (ii) and (iii) of the induction statements hold in expectation.

Next, we prove that the events (i), (ii), (iii) hold with high probability. By the induction hypothesis, we have that and the above analysis, we have that:

Lemma 18. *Suppose that the induction hypothesis holds at time t . Then, the following events occur with high-probability for all $i, j \in [p]$*

$$(i) \quad \left| \|\mathbf{g}_i^{t+1}\|^2 - \mathbb{E}[\|\mathbf{g}_i^{t+1}\|^2] \right| = \mathcal{O}\left(\frac{\text{polylog } d}{\sqrt{d}}\right)$$

$$(ii) \quad \left| \langle \mathbf{g}_i, \mathbf{g}_j \rangle - \mathbb{E}[\langle \mathbf{g}_i, \mathbf{g}_j \rangle] \right| = \mathcal{O}\left(\frac{\text{polylog } d}{\sqrt{d}}\right)$$

(iii) For any $k \in [r]$, and any unit vector \mathbf{w}

$$\left| \langle \mathbf{w}, \mathbf{g}_i \rangle - \mathbb{E}[\langle \mathbf{w}, \mathbf{g}_i \rangle] \right| = \mathcal{O}\left(\frac{\text{polylog } d}{\sqrt{d}}\right) \quad (106)$$

Proof. We condition on the event in \mathcal{F}_t that Q_t, M_t are bounded by some constants independent of d, n . Subsequently, the proof proceeds similar to Proposition 2, with the additional incorporation of the term due to $\hat{f}(\mathbf{z}'; \mathbf{W}^t, \mathbf{a})$ in the gradient.

We have:

$$\mathbf{g}_j^t = \frac{1}{n} a_j \sum_{\nu=1}^n z_i (f^*(z^\nu) - \hat{f}(z^\nu; \mathbf{W}^t, \mathbf{a})) \sigma'(z_i^\top \mathbf{w}_j^t) \quad (107)$$

Define:

$$\mathbf{X}_i^\nu = z^\nu \sigma'(\langle \mathbf{w}_i^t, z^\nu \rangle) (f^*(z^\nu) - \hat{f}(z^\nu; \mathbf{W}^t, \mathbf{a})). \quad (108)$$

Analogous to the proof of Proposition 2, we have:

$$\|\mathbf{g}_i\|^2 - \mathbb{E}[\|\mathbf{g}_i\|^2] = \frac{1}{n^2 p^2} \left(\underbrace{\sum_{\nu=1}^n \|\mathbf{X}_i^\nu\|^2 - n \mathbb{E}[\|\mathbf{X}^\nu\|^2]}_{S_1} + \underbrace{\sum_{\nu \neq \nu'} \langle \mathbf{X}_i^\nu, \mathbf{X}_i^{\nu'} \rangle - n(n-1) \mathbb{E}[\langle \mathbf{X}_i^\nu, \mathbf{X}_i^{\nu'} \rangle]}_{S_2} \right) \quad (109)$$

Similarly, we have:

$$\langle \mathbf{g}_i, \mathbf{g}_j \rangle - \mathbb{E}[\langle \mathbf{g}_i, \mathbf{g}_j \rangle] = \frac{1}{n^2 p^2} \left(\underbrace{\sum_{\nu=1}^n \langle \mathbf{X}_i^\nu, \mathbf{X}_j^\nu \rangle - n \mathbb{E}[\langle \mathbf{X}_i^\nu, \mathbf{X}_j^\nu \rangle]}_{S'_1} + \underbrace{\sum_{\nu \neq \nu'} \langle \mathbf{X}_i^\nu, \mathbf{X}_j^{\nu'} \rangle - n(n-1) \mathbb{E}[\langle \mathbf{X}_i^\nu, \mathbf{X}_j^{\nu'} \rangle]}_{S'_2} \right)^2 \quad (110)$$

Note that $f^*(z^\nu)$ and $\hat{f}(z^\nu; \mathbf{W}^t, \mathbf{a})$ are polynomials in finite-number of correlated Gaussians $\langle \mathbf{w}_1^t, z^\nu \rangle, \dots, \langle \mathbf{w}_p^t, z^\nu \rangle, \langle \mathbf{w}_1^*, z^\nu \rangle, \dots, \langle \mathbf{w}_r^*, z^\nu \rangle$. Therefore, by repeatedly applying Lemma 2 and Theorem 5, we obtain that $\sigma'(\langle \mathbf{w}_i^t, z \rangle)$, $f^*(z^\nu)$ and $\hat{f}(z^\nu; \mathbf{W}^t, \mathbf{a})$ have bounded Orlicz norms of some finite order α_t .

Subsequently, similar to Lemma 9, through Holder's inequality, Lemma 2 and Theorem 5, we obtain that $\|\mathbf{X}_i^\nu\|^2, \langle \mathbf{X}_i^\nu, \mathbf{X}_j^{\nu'} \rangle, \langle \mathbf{X}_i^\nu, \mathbf{X}_j^\nu \rangle, \langle \mathbf{X}_i^\nu, \mathbf{X}_j^\nu \rangle$, have Orlicz norms of order $\mathcal{O}(d)$ with $\alpha = \alpha_t$ for some α_t independent of d .

The remaining proof follows by repeating the arguments in Lemmas 9, 10 for Orlicz norms of general order.

Similarly (iii) is obtained by replacing the application of Bernstein's inequality in Lemma 8 by Theorem 5. \square

Lemmas 15 and 18 together with C.8 and the induction hypothesis imply statement (iii) at time $t + 1$.

It remains to prove the base case i.e $t = 1$. If the leap index $\ell > 1$, $U_t^* = 0$ for all $t \geq 1$. Applying the above arguments then implies that (i) and (iii) hold for all timesteps t .

Therefore, we consider the case $\ell = 1$ At $t = 1$, U_1^* is simply the subspace along $(C_1(f^*))$. Let $\mathbf{v} = \pm \frac{1}{\|C_1(f^*)\|} C_1(f^*)$ be a vector as per (ii) let $i \in [p]$ be an arbitrary neuron. We have:

$$\begin{aligned} \langle \mathbf{v}, \mathbf{w}_i^1 \rangle &= \langle \mathbf{v}, \mathbf{w}_i^0 \rangle + \eta \langle \mathbf{v}, \mathbf{g}^i \rangle \\ &= \eta \langle \mathbf{v}, \mathbf{g}^i \rangle + \mathcal{O}\left(\frac{1}{\sqrt{d}}\right) \end{aligned}$$

It is straightforward to check that 4 hold when σ, g^* are polynomials, while Lemma 14 holds in expectation. Applying the concentration results for Orlicz norms of general order as in Lemma 18 imply that Lemma 14

also holds in probability for polynomial σ, g^* . To establish (i) at time $t = 1$, we note that Lemma 14 implies that $\mathbb{E}[\mathbf{w}_i]^2$ converges to $1 + ca_i^2$ where c is independent of d, n . By Lemma 4 the first term equals with high-probability, $\pm \frac{a_i \mu_1}{p} \|C_1(f^*)\| + \mathcal{O}(\frac{1}{\sqrt{d}})$. Since, $\frac{a_i \mu_1}{p} \|C_1(f^*)\|$ is a non-constant (linear) polynomial in a_i , this proves (ii) for $t = 1$. Lemmas 4 and part (iii) in Lemma 18 directly imply (iii) of the induction statements.

Points (ii), (iii) of the induction statements directly imply Theorem 3.

C.9 PREDICTION OF THE ALIGNMENT AT THE SECOND STEP

We now utilize the analysis in the previous section to obtain a theoretical prediction for the gradient orientation after two steps for the above target function. We follow the notation defined in the proof sketch in Section C.7

We have, using Equation (89):

$$\mathbb{E}[\langle \mathbf{u}, \mathbf{g}_j^1 \rangle] \rightarrow \mathbb{E}[y \mu_{1, \mathbf{x}^\top \mathbf{v}^*}(\mathbf{x}_i^\top \mathbf{u})] - \mathbb{E}[\hat{y}_i^1 \sigma'(\mathbf{x}_i^\top \mathbf{w}^1)_j(\mathbf{x}_i^\top \mathbf{u})]. \quad (111)$$

As explained in the previous section, the second term does not contribute to an alignment towards \mathbf{v}^\perp . Therefore we consider the ratio of the first term when $\mathbf{u} = \mathbf{v}^*$ or $\mathbf{u} = \mathbf{v}^\perp$. We obtain:

$$\langle \mathbf{g}^1, \mathbf{v}^* \rangle \approx \mathbb{E}[y \mu_{1, (\mu_1 \mathbf{x}^\top \mathbf{v}^*)/\sqrt{2p}}(\mathbf{x}_i^\top \mathbf{v}^*)], \quad (112)$$

where μ_1 denotes the first Hermite coefficient of the student activation σ , given by 0.5 for Relu.

$\mu_{1, \mathbf{x}^\top \mathbf{v}^*}$ corresponds to the first Hermite coefficient of a translated Relu function and is given by:

$$\mu_{1, \kappa, j} = (1 - \Phi(-a_j \eta \kappa)) = \frac{1}{2}(1 \pm \text{erf}(\eta \kappa / \sqrt{2})). \quad (113)$$

Therefore, when $\eta = 4\sqrt{p}$, we obtain:

$$\mu_{1, (\mu_1 \mathbf{x}^\top \mathbf{v}^*)/\sqrt{2p}} = \frac{1}{2}(1 \pm \text{erf}(\mathbf{x}^\top \mathbf{v}^*)). \quad (114)$$

Therefore, using the Hermite decomposition of erf, we obtain the following predicted orientations in the setting considered in the right panel of Fig. 6:

$$\mathbf{v}_1^{(t=2)} = (1 - \frac{2}{\sqrt{3\pi}})\mathbf{w}_1^* + (1 + \frac{2}{\sqrt{3\pi}})\mathbf{w}_2^* \quad \mathbf{v}_2^{(t=2)} = (1 + \frac{2}{\sqrt{3\pi}})\mathbf{w}_1^* + (1 - \frac{2}{\sqrt{3\pi}})\mathbf{w}_2^* \quad (115)$$

C.10 LIMITATIONS OF THE STAIRCASE STRUCTURE

We show that a natural class of teacher functions, containing neurons with identical activation functions and uniform second-layer weights does not contain a staircase structure:

Proposition 4. *Let $y = f^*(\mathbf{z}) = \sum_{k=1}^r \sigma^*(\langle \mathbf{w}_k^*, \mathbf{z} \rangle)$ for some σ^* having leap index 1, then $U_i^* = U_1^*$ for all $i \geq 1$.*

Proof. For any such target function, the direction \mathbf{v}^* is given by $\mathbf{v}^* = \frac{1}{\sqrt{r}}(\sum_{k=1}^r \mathbf{w}_k^*)$. Without loss of generality, assume that $\mathbf{w}_k^* = \mathbf{e}_k$, where \mathbf{e}_k denotes the unit vector corresponding to the k_{th} coordinate.

Now, consider any direction $\mathbf{u} \perp \mathbf{v}^*$ in the teacher subspace. Such a vector satisfies $\sum_{k=1}^r u_k = 0$. Therefore, for any $k \geq 0$, we have:

$$\begin{aligned} \mathbb{E} [f^*(\mathbf{z}) H_k((\mathbf{v}^*)^\top \mathbf{z}) ((\mathbf{u}_i)^\top \mathbf{z})] &= \left(\sum_{i=1}^p (u_i) \right) (\mathbb{E} [f^*(\mathbf{z}) H_k((\mathbf{v}^*)^\top \mathbf{z}) z_1]) \\ &= 0. \end{aligned} \tag{116}$$

Where we used the symmetry of $f^*(\mathbf{z})$ w.r.t permutations of the first r coordinates. Therefore, the Hermite decomposition of $f^*(\mathbf{z})$ does not contain any term that linearly couples \mathbf{u} to \mathbf{v}^* . \square

Therefore, the presence of a staircase structure requires asymmetry between the the dependence of the target function on different directions in the teacher subspace.

D LEARNING THE SECOND LAYER

D.1 PROOF OF PROPOSITION 1

We first prove the finite p case of Proposition 1. Let \mathbf{a} be a second layer vector with $a_i \leq c/\sqrt{p}$, and assume that W only learns a subspace $U \subseteq V^*$. We write $\mathbb{R}^d = U \oplus U^\perp \oplus V^{*\perp}$, where U^\perp is the orthogonal subspace of U in V^* . By assumption, we have $\|P_{U^\perp} \mathbf{w}_i\| \leq \varepsilon_d$ for every i ; where ε_d is going to zero as d grows.

For any $\mathbf{z} \in \mathbb{R}^d$, we have

$$\begin{aligned} \hat{f}(\mathbf{z}; W, \mathbf{a}) &= \sum_{i=1}^p \frac{a_i}{\sqrt{p}} \sigma(\langle \mathbf{w}_i, P_U \mathbf{z} \rangle + \langle \mathbf{w}_i, P_{U^\perp} \mathbf{z} \rangle + \langle \mathbf{w}_i, P_{V^{*\perp}} \mathbf{z} \rangle) \\ &= \sum_{i=1}^p \frac{a_i}{\sqrt{p}} \sigma(\langle \mathbf{w}_i, P_U \mathbf{z} \rangle + \langle \mathbf{w}_i, P_{V^{*\perp}} \mathbf{z} \rangle) + \frac{a_i}{\sqrt{p}} \varepsilon_d \tilde{\sigma}(\langle \mathbf{w}_i, P_{U^\perp} \mathbf{z} \rangle) \end{aligned}$$

where $\tilde{\sigma}$ is a Lipschitz function. We call the first term of the above expression $\tilde{f}(P_U \mathbf{z}, P_{V^{*\perp}} \mathbf{z})$, forgetting the structure of the function \hat{f} . Then, we can write the risk as

$$\mathcal{R}(W, \mathbf{a}) = \mathbb{E}_{\mathbf{z}} \left[\left(f^*(\mathbf{z}) - \tilde{f}(P_U \mathbf{z}, P_{V^{*\perp}} \mathbf{z}) \right)^2 \right] + O(\varepsilon_d), \quad (117)$$

having used the Cauchy-Schwarz inequality to bound the contribution of $\tilde{\sigma}$. Then, by successive expectations,

$$\begin{aligned} \mathcal{R}(W, \mathbf{a}) &= \mathbb{E}_{P_{V^{*\perp}} \mathbf{z}, P_U \mathbf{z}} \left[\mathbb{E}_{P_{U^\perp} \mathbf{z}} \left[\left(f^*(\mathbf{z}) - \tilde{f}(P_U \mathbf{z}, P_{V^{*\perp}} \mathbf{z}) \right)^2 \middle| P_{V^{*\perp}} \mathbf{z}, P_U \mathbf{z} \right] \right] + O(\varepsilon_d), \\ &\geq \mathbb{E}_{P_{V^{*\perp}} \mathbf{z}, P_U \mathbf{z}} \left[\inf_f \mathbb{E}_{P_{U^\perp} \mathbf{z}} \left[\left(f^*(\mathbf{z}) - f(P_U \mathbf{z}, P_{V^{*\perp}} \mathbf{z}) \right)^2 \middle| P_{V^{*\perp}} \mathbf{z}, P_U \mathbf{z} \right] \right] + O(\varepsilon_d) \end{aligned}$$

where the infimum is taken over all measurable functions $f : U \times V^{*\perp} \rightarrow \mathbb{R}$. But this infimum exactly corresponds to the definition of conditional expectation/conditional variance, which is independent from $P_{V^{*\perp}} \mathbf{z}$ (since f^* is). As a result,

$$\mathcal{R}(W, \mathbf{a}) \geq \mathbb{E}_{P_U \mathbf{z}} [\text{Var}(f^*(\mathbf{z}) | P_U \mathbf{z})] + O(\varepsilon_d), \quad (118)$$

which implies the statement of Proposition 1.

D.2 FULL STATEMENT OF THEOREM 4

We now provide the full statement of Theorem 4. It establishes the asymptotic equivalence of the training and generalization errors of the original features and the conditional Gaussian features defined by equation (15).

Consider the sequence of vectors $\mathbf{v}_n \in \mathbb{R}^d$ defined as in Equation (76) by $\mathbf{v}_n = \frac{1}{n} \sum_{i=1}^n y_i \mathbf{x}_i$. For simplicity, we omit the dependence of \mathbf{v}_n on n and denote each entry by \mathbf{v} . For any vector $\mathbf{z} \in \mathbb{R}^d$, define the decomposition $\mathbf{z} = z_v \mathbf{v} + \mathbf{z}^\perp$ and feature maps:

$$\phi_{\text{CK}}(\mathbf{z}) = \sigma(W^{(1)} \mathbf{z}), \quad (119)$$

where $W^{(1)}$ denotes the weight matrix obtained through the application of a single gradient step.

Then the random variable $\phi_{\text{CK}}(\mathbf{z})$ admits a regular conditional distribution conditioned on the values of $z_{\mathbf{v}}$ (Theorem 8.37 in [Klenke \(2013\)](#)). Therefore, the following mean, correlation, and covariance matrix are well-defined:

$$\begin{aligned}\mu(z_{\mathbf{v}}) &= \mathbb{E}[\phi_{\text{CK}}(\mathbf{z}) \mid z_{\mathbf{v}}], \quad \Psi(z_{\mathbf{v}}) = \mathbb{E}[\phi_{\text{CK}}(\mathbf{z})(z_{\perp}^{\perp})^{\top} \mid z_{\mathbf{v}}], \\ \Phi(z_{\mathbf{v}}) &= \text{Cov}[\phi_{\text{CK}}(\mathbf{z}) \mid z_{\mathbf{v}}] - \Psi(z_{\mathbf{v}})\Psi(z_{\mathbf{v}})^{\top}\end{aligned}\quad (120)$$

Now, for each value of $z_{\mathbf{v}}$, define the following random variable:

$$\phi_{\text{CL}}(\mathbf{z}; \mathbf{v}) = \mu(z_{\mathbf{v}}) + \Psi(z_{\mathbf{v}})\mathbf{z}^{\perp} + \Phi(z_{\mathbf{v}})\boldsymbol{\xi}. \quad (121)$$

Then $\phi_{\text{CL}}(\mathbf{z}; \mathbf{v})$ satisfies:

$$\mathbb{E}[\phi_{\text{CL}}(\mathbf{z}) \mid z_{\mathbf{v}}] = \mu(z_{\mathbf{v}}), \mathbb{E}[\phi_{\text{CL}}(\mathbf{z})(z_{\perp}^{\perp})^{\top} \mid z_{\mathbf{v}}] = \Psi(z_{\mathbf{v}}), \text{Cov}[\phi_{\text{CL}}(\mathbf{z}) \mid z_{\mathbf{v}}] = \text{Cov}[\phi_{\text{CK}}(\mathbf{z}) \mid z_{\mathbf{v}}]. \quad (122)$$

Therefore, $\phi_{\text{CL}}(\mathbf{z}; \mathbf{v})$ is a Gaussian variable having the same conditional mean, covariance as $\phi_{\text{CK}}(\mathbf{z}; \mathbf{v})$ and the same correlation with \mathbf{z}_{\perp} as $\phi_{\text{CK}}(\mathbf{z}; \mathbf{v})$. Since \mathbf{z}_{\perp} is Gaussian and independent of $z_{\mathbf{v}}$, this uniquely characterizes the conditional measure of $\phi_{\text{CL}}(\mathbf{z}; \mathbf{v})$.

Now, consider a set of n training inputs $\mathbf{z}_1, \dots, \mathbf{z}_n$. For each $i \in n$, generate an equivalent feature map Φ_{CL} through equation (1), with $\boldsymbol{\xi}$ being independently sampled for each example. Let Φ_{CK} and Φ_{CL} denote matrices in $\mathbb{R}^{n \times p}$ with rows $\phi_{\text{CK}}(\mathbf{z}_i)$ and $\phi_{\text{CL}}(\mathbf{z}_i)$ respectively,

Consider the following minimization problem:

$$\min_{\mathbf{a} \in \mathbb{R}^p} \frac{1}{n} \sum_{\nu=1}^n (\langle \mathbf{a}, \phi_{\text{CK}}(\mathbf{z}^{\nu}) \rangle - f^*(\mathbf{z}^{\nu}))^2 + \lambda \|\mathbf{a}\|^2 \quad (123)$$

Define the following constraint set:

$$\mathcal{S}_p = \{\boldsymbol{\theta} \in \mathbb{R}^d \mid \|\boldsymbol{\theta}\|_2 \leq R, \quad \|\boldsymbol{\theta}\|_{\infty} \leq Cp^{-\eta}\}. \quad (124)$$

We make the following assumption:

Assumption 6. *There exist constants R, C, η such that the minimizer $\hat{\mathbf{a}}_{\text{CK}}$ of the optimization problem defined by equation (123) lies in \mathcal{S}_p with high probability as $n, d \rightarrow \infty$.*

The above assumption can be enforced by utilizing constrained minimization for the second layer. Alternatively, for overparameterized models i.e $p/n > 1$, one could utilize the arguments in Theorem 5 of [Montanari and Saeed \(2022\)](#). Let $\hat{\mathcal{R}}_n^*(\Phi, \mathbf{y}(\mathbf{Z}))$, $\mathcal{R}_g^*(\Phi, \mathbf{y}(\mathbf{Z}))$ denote the training and generalization errors respectively with features Φ and labels $\mathbf{y}(\mathbf{Z})$.

Theorem 4. *Assume that $n, p = \Theta(d)$, and that the vector $V_1^* = \mathbf{v}^*$ defined in Theorem 2 is nonzero. Then, the sequence of vectors $\mathbf{v}_n = \frac{1}{n} \sum_{i=1}^n y_i \mathbf{x}_i \in \mathbb{R}^d$ satisfy:*

- (i) *As $n, d \rightarrow \infty$, $P_{V^*} \mathbf{v} \xrightarrow{\mathbb{P}} \frac{\mu}{\sqrt{p}} \mathbf{v}^*$.*
- (ii) *Under Assumption 6, the training and generalization errors obtained through the minimization of the objective (15) for training distribution defined by feature maps $\phi_{\text{CK}}(\mathbf{z})$ converge in distribution to the corresponding training and generalization errors for features $\phi_{\text{CL}}(\mathbf{z}; \mathbf{v})$.*

Concretely, we have that for any bounded Lipschitz function $\Psi : \mathbb{R} \rightarrow \mathbb{R}$:

$$\lim_{n, p \rightarrow \infty} \left| \mathbb{E} \left[\Psi \left(\hat{\mathcal{R}}_n^*(\Phi_{\text{CK}}, \mathbf{y}(\mathbf{Z})) \right) \right] - \mathbb{E} \left[\Psi \left(\hat{\mathcal{R}}_n^*(\Phi_{\text{CL}}, \mathbf{y}(\mathbf{Z})) \right) \right] \right| = 0$$

$$\lim_{n,p \rightarrow \infty} |\mathbb{E} [\Psi(\mathcal{R}_g(\Phi_{CK}, \mathbf{y}(\mathbf{Z})))] - \mathbb{E} [\Psi(\mathcal{R}_g(\Phi_{CL}, \mathbf{y}(\mathbf{Z})))]| = 0$$

In particular, for any $\mathcal{E} \in \mathbb{R}$, and denoting $\xrightarrow{\mathbb{P}}$ the convergence in probability:

$$\begin{aligned} \widehat{\mathcal{R}}_n^*(\Phi_{CK}, \mathbf{y}(\mathbf{Z})) \xrightarrow{\mathbb{P}} \mathcal{E} & \text{ if and only if } \widehat{\mathcal{R}}_n^*(\Phi_{CL}, \mathbf{y}(\mathbf{Z})) \xrightarrow{\mathbb{P}} \mathcal{E} \\ \mathcal{R}_g^*(\Phi_{CK}, \mathbf{y}(\mathbf{Z})) \xrightarrow{\mathbb{P}} \mathcal{E} & \text{ if and only if } \mathcal{R}_g(\Phi_{CL}, \mathbf{y}(\mathbf{Z})) \xrightarrow{\mathbb{P}} \mathcal{E}, \end{aligned} \quad (125)$$

Part (i) follows directly from Lemma 5. To prove the equivalence of training and generalization errors for the given direction, we rely on the framework of one-dimensional CLT (Central Limit Theorem), discussed in Goldt et al. (2022). One-dimensional CLT was recently shown to imply the universality of training and generalization errors for Random feature models in Hu and Lu (2022). However, in our setting where we train the model, and as verified empirically in Ba et al. (2022), a naive one-dimensional CLT with equivalent Gaussian features no longer holds.

Instead, we introduce a generalization termed ‘‘conditional one-dimensional CLT’’, given by the following Lemma:

Lemma 19. For any Lipschitz function $\varphi : \mathbb{R}^2 \rightarrow \mathbb{R}$,

$$\lim_{n,p \rightarrow \infty} \sup_{\theta_1 \in \mathcal{S}_p, \theta_2 \in \mathcal{S}^{d-1}} |\mathbb{E} [\varphi(\theta_1^\top \phi_{CK}(\mathbf{z}), \theta_2^\top \mathbf{z}) \mid z_v = k] - \mathbb{E} [\varphi(\theta^\top \phi_{CL}(\mathbf{z}), \theta^\top \mathbf{z}) \mid z_v = k]| = 0, \quad \forall k \in \mathbb{R}, \quad (126)$$

where \mathcal{S}^{d-1} denotes the unit sphere in \mathbb{R}^d

Proof. For an input $\mathbf{z} \sim \mathcal{N}(0, I_d)$, we consider the decomposition $\mathbf{z} = z_v \mathbf{v} + \mathbf{z}^\perp$. We note that conditioned on $z_v = k$, $\phi_{CL}(\mathbf{z})$, \mathbf{z}^\perp is a Gaussian random variable. Next, consider $\phi_{CK}(\mathbf{z})$. Our proof relies on the observation that while the features $\phi_{CK}(\mathbf{z})$ have complex non-linear dependence on z_v , for a fixed value of z_v , they are equivalent to a random-features mapping applied to \mathbf{z}^\perp . Concretely, we have by Lemma 12 that the weight matrix $\mathbf{W}^{(1)}$ has the following spike+bulk decomposition (equation (76)):

$$\mathbf{W}^{(1)} = \eta \mathbf{u} \mathbf{v}^\top + \mathbf{W}^{(0)} + \eta \Delta, \quad (127)$$

where $\mathbf{u} = \frac{\mu_1}{p} \mathbf{a}$

Let \mathbf{W}^\perp denote the combined matrix $\mathbf{W}^{(0)} + \eta \Delta$ with rows \mathbf{w}_i^\perp for $i \in [p]$.

Lemma 14 implies that there exist constants c_i for $i \in [p]$ depending only on a_i such that $\|\mathbf{w}_i^\perp\|^2 = c_i + \mathcal{O}(\frac{\text{polylog } d}{\sqrt{d}})$ with high-probability. Define the following neuron-wise activation functions:

$$\sigma_{i,z_v}(u) = \sigma(c_j u + \eta v_z) - \mathbb{E}_u [\sigma(c_j u + \eta u_i z_v)], \quad (128)$$

where $i \in [p]$ denotes the index of the neuron and the expectation is w.r.t $z \sim \mathcal{N}(0, 1)$. Under the choice of symmetric initialization in Equation (26), it suffices to restrict ourselves to the first half $p/2$ neurons.

For a fixed value of z_v , the feature map $\phi_{CK}(\mathbf{z}) = \sigma(\mathbf{W}^1 \mathbf{z})$ is equivalent to a random features mapping with neurons σ_{i,z_v} applies to inputs $\mathbf{z}^\perp \in \mathbb{R}^d$ with approximately orthogonal weights \mathbf{W}^\perp . Consider the following events for some positive constants C_1, C_2, C_3 :

$$\mathcal{A}_1 = \left\{ \sup_{i,j \in [p/2]} |\langle \mathbf{w}_i^\perp, \mathbf{w}_j^\perp \rangle - c_i \delta_{ij}| \leq C_1 \left(\frac{\text{polylog } d}{d} \right)^{1/2} \right\} \quad \mathcal{A}_2 = \left\{ \|\mathbf{W}^\perp\|_{\text{op}} \leq C_3 (\text{polylog } d) \right\}$$

We have, using Lemma 12 and a union bound, that for $p, d = \Theta(n)$, $\Pr[\mathcal{A}_1] \xrightarrow{n,d \rightarrow \infty} 1$. Furthermore, part (ii) of Lemma 14 in Ba et al. (2022) implies that $\Pr[\mathcal{A}_2] \rightarrow 1$. Next, we utilize Corollary 2 and Lemma 3

in [Hu and Lu \(2022\)](#). Note that the neuron wise activation functions (128) for a fixed value of z_v satisfy $\mathbb{E}_u [\sigma_{i,v_z}(u)] = 0$. We relax the requirement of odd-activation in [Hu and Lu \(2022\)](#) by noting that ϕ_{CK}, ϕ_{CL} have exactly equivalent means and covariances as in Theorem 6 of [Dandi et al. \(2023\)](#). \square

The above Lemma states that the one-dimensional projections of $\phi_{CK}(\mathbf{z})$ are asymptotically distributed as jointly Gaussian variables with \mathbf{z}_\perp .

D.3 CONDITIONAL GET

We now prove part (ii) of Theorem 4 using Lemma 19. This relies on the universality of training and generalization errors between the given distribution and the “conditional equivalent” distribution. The central idea of the proof again relies on the isolation of the effects of the “spikes” and the “noise” in the features.

The technique presented here is also of independent interest for proving the universality of training, generalization errors in related setups such as with spiked-covariance inputs.

We utilize the following properties of the features :

Lemma 20. *For any fixed z_v , the random variable $\phi_{CK} - \mu(z_v)$ is sub-Gaussian with sub-Gaussian norm independent of z_v and n .*

Proof. The result follows from the assumption of uniform boundedness of the derivative of σ^* and the Lipschitz concentration of Gaussian variables. \square

Lemma 21. *There exists a constant C such that the matrix $\bar{\Phi}_{CK}$ with rows $\phi_{CK} - \mu(z_v)$ satisfies:*

$$\Pr[\|\bar{\Phi}_{CK}\| \geq K\sqrt{p}] \leq 2\exp(-Cn) \quad (129)$$

Proof. By Lemma 20, each row of $\bar{\Phi}_{CK}$ is sub-Gaussian. Therefore, the result follows from the concentration of spectral norm of matrices with independent sub-Gaussian rows (Theorem 5.39 in [Vershynin \(2010\)](#)). \square

We start by proving certain properties of the optimal parameters a_i upon the training of the second layer:

Lemma 22. *Let $\hat{a}_{CK}(\lambda)$ be the parameters obtained through ridge regression on features $\phi_{CK}(\mathbf{z}_i)_{\{i=1,\dots,n\}}$ with regularization strength λ . Then, there exists a constants C such that with high probability as $n, d \rightarrow \infty$:*

$$\frac{1}{n} \sum_{i=1}^n (\hat{a}_{CK}^\top \mu(z_{i,v}))^2 \leq C \quad (130)$$

Proof. By assumption, $y_i(\mathbf{z}) = \frac{1}{\sqrt{p}} a^\top \sigma(W\mathbf{z})$ with σ' uniformly bounded. Therefore, from the concentration of Lipschitz functions of gaussian variables, $y_i(\mathbf{z})$ is sub-Gaussian. Thus $y_i^2(\mathbf{z})$ are sub-exponential variables. Using Bernstein’s inequality [Vershynin \(2018\)](#), we obtain:

$$\Pr\left[\frac{1}{n} \sum_{i=1}^n (y_i)^2 - \mathbb{E}[(y_i)^2] > K\right] \leq 2\exp(-\min(c_1 K, c_2 K^2)n). \quad (131)$$

For constants c_1, c_2 . \square

Let $\mathcal{A}_{\mathbf{y}}$ denote the following event:

$$\mathcal{A}_{\mathbf{y}} = \left\{ \frac{1}{n} \sum_{i=1}^n (y_i)^2 < C_1 \right\}. \quad (132)$$

By Equation (131), we have $\Pr[\mathcal{A}_{\mathbf{y}}] \rightarrow 1$ as $n, d \rightarrow \infty$.

Let $\hat{\mathcal{R}}(W, \mathbf{a})$ denote the empirical risk at given values of \mathbf{a}, W . We have:

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a}} \hat{\mathcal{R}}(W, \mathbf{a}) = \arg \min_{\mathbf{a}} \frac{1}{2n} \sum_{i=1}^n (y_i - \mathbf{a}^\top \phi_k(\mathbf{z}_i))^2. \quad (133)$$

We note that when $\mathbf{a} = \mathbf{0}$, we have:

$$\hat{\mathcal{R}}(W, \mathbf{0}) = \frac{1}{n} \sum_{i=1}^n (y_i)^2. \quad (134)$$

Since $\hat{\mathbf{a}}$ minimizes $\hat{\mathcal{R}}(W, \mathbf{a})$, we must have:

$$\hat{\mathcal{R}}(W, \hat{\mathbf{a}}) \leq \hat{\mathcal{R}}(W, \mathbf{0}). \quad (135)$$

We obtain:

$$\begin{aligned} & \frac{1}{2n} \sum_{i=1}^n (y_i - \mathbf{a}^\top \phi_k(\mathbf{z}_i))^2 \leq \frac{1}{n} \sum_{i=1}^n (y_i)^2 \\ \implies & \frac{1}{2n} \sum_{i=1}^n (\mathbf{a}^\top \phi_k(\mathbf{z}_i))^2 \leq \frac{1}{n} \sum_{i=1}^n \mathbf{a}^\top \phi_k(\mathbf{z}_i) y_i \\ \implies & \frac{1}{2n} \sum_{i=1}^n (\hat{\mathbf{a}}^\top \phi_k(\mathbf{z}_i))^2 \leq \sqrt{\frac{1}{n} \sum_{i=1}^n (\mathbf{a}^\top \phi_k(\mathbf{z}_i))^2} \sqrt{\frac{1}{n} \sum_{i=1}^n y_i^2}, \end{aligned}$$

where the last inequality follows from Cauchy-Schwarz. Therefore:

$$\begin{aligned} & \sqrt{\frac{1}{n} \sum_{i=1}^n (\mathbf{a}^\top \phi_k(\mathbf{z}_i))^2} \leq 2 \sqrt{\frac{1}{n} \sum_{i=1}^n y_i^2} \\ & \frac{1}{n} \sum_{i=1}^n (\mathbf{a}^\top \mu(\mathbf{z}_i))^2 + \frac{1}{n} \|\bar{\Phi}_{\text{CK}}^\top \mathbf{a}\|_2^2 \leq 4 \left(\frac{1}{n} \sum_{i=1}^n y_i^2 \right). \end{aligned}$$

Applying Lemma 21 and $\Pr[\mathcal{A}_{\mathbf{y}}] \xrightarrow{n, d \rightarrow \infty} 1$ then completes the proof.

Next, we prove the universality of the training, generalization error, conditioned on the values of the projections $z_{\mathbf{v}}$. This can be achieved through a number of techniques such as the Lindeberg's method in [Hu and Lu \(2022\)](#). We utilize the result of [Montanari and Saeed \(2022\)](#), who apply the interpolation technique to continuously transform the inputs \mathbf{x}_i to equivalent Gaussian vectors \mathbf{g}_i .

Instead, we interpolate between the features $\phi_{\text{CK}}(\mathbf{z})$ and $\phi_{\text{CL}}(\mathbf{z})$. Define:

$$\mathbf{u}_{t,i} = \boldsymbol{\mu}(z_{i,\mathbf{v}}) + \cos(t)(\phi_{\text{CK}}(\mathbf{z}) - \boldsymbol{\mu}(z_{i,\mathbf{v}})) + \sin(t)(\phi_{\text{CL}}(\mathbf{z}) - \boldsymbol{\mu}(z_{i,\mathbf{v}})),$$

Let \mathcal{A}_1 denote the event:

$$\mathcal{A}_1 = \left\{ \frac{1}{n} \sum_{i=1}^n (\hat{a}_{\text{CK}}^\top \mu(z_{i,v}))^2 \leq C_1 \right\} \quad (136)$$

Under the above interpolation path, we generalize Theorem 1 in [Montanari and Saeed \(2022\)](#) to obtain that for any bounded Lipschitz function $\Phi : \mathbb{R} \rightarrow \mathbb{R}$:

$$\lim_{n,p \rightarrow \infty} \sup_{v_{z_1}, \dots, v_{z_n}} \left| \mathbb{E} \left[\mathbf{1}_{\mathcal{A}_1} \Phi \left(\hat{\mathcal{R}}_n^* (\Phi_{\text{CK}}, \mathbf{y}(\mathbf{Z})) \right) \mid v_{z_1}, \dots, v_{z_n} \right] - \mathbb{E} \left[\mathbf{1}_{\mathcal{A}_1} \Phi \left(\hat{\mathcal{R}}_n^* (\Phi_{\text{CL}}, \mathbf{y}(\mathbf{Z})) \right) \mid v_{z_1}, \dots, v_{z_n} \right] \right| = 0. \quad (137)$$

Below, we explain the modifications to Theorem 1 in [Montanari and Saeed \(2022\)](#) that allow its applicability to our setting:

- (i) We replace equation (12) in Assumption 5 of [Montanari and Saeed \(2022\)](#) by the conditional 1d-CLT (Lemma 19). This is similar to the conditioning utilized in [Dandi et al. \(2023\)](#) for proving the universality in mixture models.
- (ii) Our target function $y = f^*(z)$ depends on the projection along the spike v_z as well as the orthogonal component z^\perp . Since we condition on the values of v_z , their dependence can be absorbed into the loss function for each input z_i^\perp .
- (iii) While Theorem 1 in [Montanari and Saeed \(2022\)](#) does not allow a dependence of the labels on the latent variables z , such a target function can be incorporated by considering the inputs to be the joint variables in $(\Phi_{\text{CK}}(z), z) \in \mathbb{R}^{p+d}$ and constraining the parameters to have 0 components along the last d directions.
- (iv) The event \mathcal{A}_1 and Lemma 21 ensure that Lemmas 5 and 6 in [Montanari and Saeed \(2022\)](#) hold under the presence of variable and unbounded means across samples $\mu(z_{i,v})$.

Next, using the Law of total expectation and Equation 137, we obtain:

$$\lim_{n,p \rightarrow \infty} \left| \mathbb{E} \left[\mathbf{1}_{\mathcal{A}_1} \Phi \left(\hat{\mathcal{R}}_n^* (\Phi_{\text{CK}}, \mathbf{y}(\mathbf{Z})) \right) \right] - \mathbb{E} \left[\mathbf{1}_{\mathcal{A}_1} \Phi \left(\hat{\mathcal{R}}_n^* (\phi_{\text{CL}}, \mathbf{y}(\mathbf{Z})) \right) \right] \right| = 0.$$

Finally, we note Lemma 22 implies that $\Pr[\mathcal{A}_1^c] \rightarrow 0$. Since Φ is bounded, we have that

$$\lim_{n,p \rightarrow \infty} \left| \mathbb{E} \left[\mathbf{1}_{\mathcal{A}_1^c} \Phi \left(\hat{\mathcal{R}}_n^* (\Phi_{\text{CK}}, \mathbf{y}(\mathbf{Z})) \right) \right] - \mathbb{E} \left[\mathbf{1}_{\mathcal{A}_1^c} \Phi \left(\hat{\mathcal{R}}_n^* (\phi_{\text{CL}}, \mathbf{y}(\mathbf{Z})) \right) \right] \right| = 0.$$

This completes the proof of Theorem 4.

D.4 GENERALIZATION ERROR LOWER BOUNDS: PROOF OF COROLLARY 1

From Theorem 4, it is sufficient to prove the lower bound for the generalization error corresponding to the equivalent features $\phi_{\text{CL}}(z)$. Let \mathbf{Z} denote the input design matrix with rows z_i . Similarly, let Ξ denote the matrix with rows containing n independent Gaussian vectors, denoting the uncorrelated noise in the equivalent conditional Gaussian features defined by equation (121). We have that $\hat{a}_{\text{CL}}(\lambda, \mathbf{Z}, \Xi) = (\Phi_{\text{CL}}^\top \Phi_{\text{CL}} + \frac{\lambda n}{N} \mathbf{I})^{-1} \Phi_{\text{CL}}^\top \mathbf{y}$. The generalization error can then be expressed as:

$$\begin{aligned} \mathcal{R}(W, \hat{a}_{\text{CL}}) &= \mathbb{E}_{z, \xi} \left[(f^*(z) - \hat{a}_{\text{CL}}(\lambda, \mathbf{Z}, \Xi)^\top \phi_{\text{CL}}(z))^2 \right] \\ &= \mathbb{E}_\xi \left[\mathbb{E}_z \left[(f^*(z) - \hat{a}_{\text{CL}}(\lambda, \mathbf{Z}, \Xi)^\top \phi_{\text{CL}}(z))^2 \right] \right], \end{aligned}$$

where the last line follows from Fubini's theorem.

We note that the predictor $\hat{f}(\mathbf{z}) = \frac{1}{\sqrt{p}} \hat{a}_{\text{CL}}^\top \phi_{\text{CL}}(\mathbf{z})$ is a linear function of \mathbf{z}^\perp with coefficients dependent on \mathbf{z}_v . Therefore, $\hat{f}(\mathbf{z}) \in \mathcal{P}_{v,1}$.

For a fixed value of ξ , we obtain the following expression for the generalization error:

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} \left[(f^*(\mathbf{z}) - \hat{f}(\mathbf{z}))^2 \right] &= \|f^* - \hat{f}\|^2 \\ &= \|P_{v,1}(f^* - \hat{f})\|^2 + \|P_{v,>1}(f^* - \hat{f})\|^2 \\ &\geq \|P_{v,>1}(f^*)\|^2, \end{aligned}$$

where we used that $P_{v,>1}(f^* - \hat{f}) = P_{v,>1}(f^*)$. Since the projection of f^* on the orthogonal complement of the teacher subspace is 0, Corollary 1 then follows using $P_{V^*} \mathbf{v} \xrightarrow{\mathbb{P}} \frac{\mu}{\sqrt{p}} \mathbf{v}^*$ and the dominated convergence theorem for the RHS.