

Spectral Regularization Allows Data-frugal Learning over Combinatorial Spaces

Anonymous authors

Paper under double-blind review

Abstract

Data-driven machine learning models are being increasingly employed in several important inference problems in biology, chemistry, and physics which require learning over combinatorial spaces. Recent empirical evidence (see, e.g., Tseng et al. (2020); Aghazadeh et al. (2021); Ha et al. (2021)) suggests that regularizing the spectral representation of such models improves their generalization power when labeled data is scarce. However, despite these empirical studies, the theoretical underpinning of when and how spectral regularization enables improved generalization is poorly understood. In this paper, we focus on learning pseudo-Boolean functions and demonstrate that regularizing the empirical mean squared error by the L_1 norm of the spectral transform of the learned function reshapes the loss landscape and allows for data-frugal learning, under a restricted secant condition on the learner’s empirical error measured against the ground truth function. Under a weaker quadratic growth condition, we show that stationary points which also approximately interpolate the training data points achieve statistically optimal generalization performance. Complementing our theory, we empirically demonstrate that running gradient descent on the regularized loss results in a better generalization performance compared to baseline algorithms in several data-scarce real-world problems.

1 Introduction

Machine learning (ML) models have become increasingly commonplace in learning pseudo-Boolean functions $f(\cdot)$, which map a d -dimensional binary vector $x \in \{\pm 1\}^d$ to a real number $f(x) \in \mathbb{R}$. In biology, ML models are being used to infer the functional properties of macro-molecules (e.g., proteins) from a small set of mutations Riesselman et al. (2018); Gelman et al. (2021). In physics, ML models are being used to infer the thermodynamic properties of combinatorial systems defined over a set of binary (Ising) state variables Carleo et al. (2019); Noé et al. (2019). These highly nonlinear and complex ML models are trained using modern techniques in continuous optimization, yet they inherit the elegant properties of pseudo-Boolean functions over the discrete binary hypercube $\{\pm 1\}^d$. In particular, it is known that the spectral representation of a pseudo-Boolean function is defined as the Walsh-Hadamard transform (WHT) of the resulting vector of combinatorial function evaluations, that is, $\mathbf{H}f(\mathbf{X})$, where \mathbf{H} is a $2^d \times 2^d$ Walsh matrix and the function evaluation vector $f(\mathbf{X}) = [f(x) : x \in \mathbf{X}]^T$ is sorted in the lexicographic ordering of all the length- d binary strings in \mathbf{X} . The WHT enables writing the pseudo-Boolean function $f(x)$ as a multi-linear polynomial $f(x) = \sum_{S \subseteq [d]} \alpha_S \prod_{i \in S} x_i$, where $[d] = \{1, 2, \dots, d\}$ and $\alpha_S \in \mathbb{R}$ is the WHT coefficient corresponding to the monomial $\prod_{i \in S} x_i$ Boros & Hammer (2002).¹

Recent studies in biology (e.g., protein function prediction) have measured the output of several of these real-world functions $f(x)$ to all the 2^d enumerations of the input x and analyzed their spectral representation. These costly high-throughput experiments on combinatorially complete datasets demonstrate a curious phenomenon: such pseudo-Boolean functions often have low dimensional structures that manifest in the form of an approximately-sparse polynomial representation (i.e., approximately-sparse WHT) with the top- k coeffi-

¹We use these terms interchangeably for pseudo-Boolean functions: spectral, Fourier, and Walsh-Hadamard.

cients corresponding to physically-meaningful interactions Poelwijk et al. (2019); Eble et al. (2020); Brookes et al. (2022) (see Appendix A for empirical evidences on protein functions).

The problem of learning pseudo-Boolean functions with sparse polynomial representations has a rich history from both theoretical and empirical viewpoints. In particular, since the pseudo-Boolean functions being learned in their polynomial representations are linear functions in their coefficients, one may think about this problem as a special instance of *sparse linear regression* in dimension 2^d . There are a number of algorithms for sparse linear regression such as LASSO Tibshirani (1996), OMP Tropp & Gilbert (2007), AMP Donoho et al. (2009), and FISTA Beck & Teboulle (2009). In particular, one may apply LASSO to get statistical guarantees for this problem which only scale linearly in the sparsity of the underlying ground truth polynomial k and logarithmically in the problem dimension $\log(2^d)$ Candes et al. (2006). Likewise, from another perspective, one may view the polynomial instead as a vector of 2^d outcomes $f(\mathbf{X})$; the objective of the learner is to approximate $f(\mathbf{X})$, when the learner can observe any chosen set of a few entries of this vector (corrupted by noise), under the assumption that $f(\mathbf{X})$ itself has a sparse WHT. In the literature, this problem has been referred to by a number of names, and we refer to it as the *sparse WHT problem*. There are yet again a number of computationally and statistically efficient algorithms for solving this problem, such as SPRIGHT Li & Ramchandran (2015) and Sparse-FFT Amrollahi et al. (2019), among others.

A common issue with these alternate views of the problem such as sparse linear regression or sparse WHT is that the resulting algorithms are not suited for *function approximation*. In particular, real-world physical and biological functions often have additional structure which is either unknown *a priori* or cannot be described succinctly, and which nonlinear function classes are often implicitly biased towards learning. Indeed, several deep neural networks (DNNs) have been shown to exhibit strong generalization performance in biological and physical problems Gelman et al. (2021); Ching et al. (2018); Sarkisyan et al. (2016). This leads to a fundamental disconnect between algorithms for which strong theoretical guarantees are known (e.g., LASSO) and practical ML models based on nonlinear and complex function classes which are well-suited for function approximation.

Another issue is specific to algorithms for solving the sparse WHT problem: some of these approaches require observing $f(\cdot)$ at any chosen input Li & Ramchandran (2015); Li et al. (2015). In many practical applications, where data is prohibitively expensive to collect, one is often forced to work with a handful of random samples. For example, in the case of proteins, a common approach to measure biological functions is through a procedure called *random mutagenesis* which allows only a random sampling of the combinatorial sequence space Sarkisyan et al. (2016). This constraint renders several algorithms for the sparse WHT problem ill-suited for learning, even though they admit near-optimal computational and sample complexity guarantees.

From a more practical point of view, several of the recent approaches for learning over combinatorial spaces (see, e.g., Tseng et al. (2020); Aghazadeh et al. (2021); Ha et al. (2021); Li et al. (2020)) follow similar ideas for improving empirical generalization. Instead of directly learning the 2^d -dimensional polynomial, they have converged on *explicitly* promoting sparsity in the spectral representation of the learned compactly-represented nonlinear function—also known as spectral regularization. These empirical studies motivate several important theoretical questions regarding the underlying mechanism of spectral regularization in improving generalization in practice. In this paper, we focus on the question: **When and how does spectral regularization improve generalization?** The answer to this question is important from both theoretical and practical perspectives. Theoretically, it helps us understand how the loss landscape is reshaped as the result of spectral regularization in favor of data-frugal learning. Practically, it tells us when and how to use spectral regularization in learning real-world combinatorial functions.

Contributions. In this paper, we theoretically analyze spectral regularization of the empirical risk from the perspective of learning pseudo-Boolean functions. We demonstrate that regularizing with the L_1 norm of the spectrum of the learned function improves the generalization performance of the learner. In particular, under a particular Restricted Secant Inequality (RSI) condition on the learner’s empirical error, measured against the ground truth, stationary points of the regularized loss admit optimal generalization performance. We relax this to a weaker quadratic growth (QG) condition and show that under an additional data-interpolation condition, stationary points of the regularized loss also admit statistical optimality. We empirically demon-

strate that these conditions are satisfied for a wide class of nonlinear functions. Our analysis provides a new generalization bounds when the underlying learned functions are sparse in the spectral domain. Empirical demonstrations on several real-world data-sets complements our theoretical findings.

2 Learning Pseudo-Boolean Functions

Problem statement. We consider the problem of learning structured pseudo-Boolean functions on the binary hypercube $\{\pm 1\}^d$. In the finite sample setting, we assume that the learner has access to a dataset D_n comprised of n labeled data points $(x^i, y^i)_{i=1}^n$ where $x^i \in \{\pm 1\}^d$ are drawn from an input distribution \mathcal{D} and the real-valued output $y_i \in \mathbb{R}$ is a noisy realization of an *unknown* pseudo-Boolean function of the input ². Namely we assume that there is a rich nonlinear function class $\mathcal{F} = \{f_{\theta} : \theta \in \mathbb{R}^m\}$ parameterized by θ , where $f_{\theta} : \{\pm 1\}^d \rightarrow \mathbb{R}$. The labels are generated as $y^i = f_{\theta^*}(x^i) + Z^i$ where Z^i is the noise in the measurement for input x^i , assumed to be independent and normally distributed $\sim \mathcal{N}(0, \sigma^2)$ ³. The objective of the learner is to learn θ^* .

Multi-linear polynomial representation. Any pseudo-Boolean function $f(\cdot)$ can be uniquely represented as a multi-linear polynomial, $f(x) = \sum_{z \in \{\pm 1\}^d} \alpha_z \prod_{i: z_i = +1} x_i$, where the scalar $\alpha_z \in \mathbb{R}$ is the coefficient corresponding to the monomial $\prod_{i: z_i = +1} x_i$, with order $\sum_{i=1}^d \mathbb{1}(z_i = +1)$ Boros & Hammer (2002). The problem of learning a pseudo-Boolean function $f(\cdot)$ is equivalent to finding the 2^d unknown coefficients α_z . In particular, the evaluations of f on all vertices of the binary hypercube $\{\pm 1\}^{2^d}$ results in a linear system of equations over the α_z variables,

$$[f(x) : x \in \{\pm 1\}^d]^T = \mathbf{H} [\alpha_z : z \in \{\pm 1\}^d]^T, \quad (1)$$

where the variables x and z enumerate the vertices of the binary hypercube $\{\pm 1\}^d$ in lexicographic order, and $\mathbf{H} \equiv \mathbf{H}_{2^d}$ denotes the $2^d \times 2^d$ Hadamard matrix defined recursively as,

$$\mathbf{H}_{2^{d+1}} = \begin{bmatrix} \mathbf{H}_{2^d} & \mathbf{H}_{2^d} \\ \mathbf{H}_{2^d} & -\mathbf{H}_{2^d} \end{bmatrix}, \quad (2)$$

where $\mathbf{H}_1 = [+1]$. Thus the vector of function evaluations can be obtained by taking the WHT of the vector of coefficients in the polynomial representation of f . Furthermore, by inverting the above linear system (note that \mathbf{H} is an orthonormal matrix), the vector of polynomial coefficients $[\alpha_z : z \in \{\pm 1\}^d]$ can, in turn, be obtained by taking the WHT of the vector $[f(x) : x \in \{\pm 1\}^d]^T$. For succinctness of notation, for a function f , we define $f(\mathbf{X})$ as $[f(x) : x \in \{\pm 1\}^d]^T$ where in the LHS the binary strings are enumerated in lexicographic order. The coefficients in the polynomial representation of $f(\cdot)$ can be collected in the vector $\mathbf{H}f(\mathbf{X})$. For functions, we use “sparsity in WHT” and “sparse polynomial representation” interchangeably.

Sparsity in real-world functions. Even in the noiseless setting, in order to perfectly learn a general pseudo-Boolean function $f(\cdot)$, its evaluations on all input binary vectors $x \in \{\pm 1\}^d$ are required, which may be very expensive. In the presence of additional structures, this significant statistical and computational cost can be mitigated. One such structure that has been observed in real-world biological and physical functions is sparsity in WHT Poelwijk et al. (2019); Eble et al. (2020); Brookes et al. (2022). In particular, in the theoretical analysis of this paper, we assume that the ground-truth function f_{θ^*} has a *sparse polynomial representation* composed of at most k monomials. Namely, $\|\mathbf{H}f_{\theta^*}(\mathbf{X})\|_0 \leq k$, where $k \ll 2^d$ is typically unknown.

Spectral regularization. In attempting to learn pseudo-Boolean functions with a sparse polynomial representation, a natural question to ask is: how can one encourage the learned functions to also have sparse polynomial representations. One solution is to regularize the training loss with an additional functional which

²We use the subscript x_i to refer to the i^{th} digit in the binary string $x = (x_1, x_2, \dots, x_d)$.

³In fact, our results only require the noise to be independent subgaussian random variables with variance parameter σ^2 , but for ease of exposition, we impose the Gaussian noise condition.

promotes sparsity in the polynomial representation of the learned function \hat{f} . Denoting the polynomial representation of \hat{f} by $\sum_{z \in \{\pm 1\}^d} \hat{\alpha}_z \prod_{i: z_i = +1} x_i$, a natural regularization functional would be $\|\hat{\alpha}\|_0 \equiv \|\mathbf{H}\hat{f}(\mathbf{X})\|_0$ which is also the sparsity of \hat{f} in WHT.

However, since $\|\cdot\|_0$ is not a continuous function, we proxy the $\|\cdot\|_0$ by the L_1 norm. The resulting regularization functional is, $\|\mathbf{H}\hat{f}(\mathbf{X})\|_1$. When learning over a parameterized function family, $\hat{f} = f_{\boldsymbol{\theta}}$, the resulting regularized ERM we consider in this paper is,

$$\min_{\boldsymbol{\theta}} \mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta}), \text{ where } R(\boldsymbol{\theta}) \triangleq \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1, \quad (\text{OBJ})$$

where $\mathcal{L}_n(\boldsymbol{\theta}) \triangleq \frac{1}{n} \sum_{i=1}^n \left[(f_{\boldsymbol{\theta}}(x^i) - y^i)^2 \right]$ is the empirical mean squared error (MSE). Regularizing by $R(\boldsymbol{\theta})$ of the above form is known as *spectral regularization (SP)*.

Remark 1. The scaling of the regularization parameter, $\frac{\lambda}{\sqrt{2^d}}$, is motivated from the linearly parameterized setting of $\mathcal{F} = \{\langle \boldsymbol{\theta}, x \rangle : \boldsymbol{\theta} \in \mathbb{R}^d\}$. An explicit computation gives, $\frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1 = \lambda \|\boldsymbol{\theta}\|_1$, which is the scaling of the regularization parameter as used in LASSO.

The regularization weight $\lambda > 0$ strikes a balance between the empirical MSE and the spectral regularization, and is set empirically using cross validation. Since the aggregate loss $\mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$ is subdifferentiable with respect to the model parameters $\boldsymbol{\theta}$, we can apply stochastic (sub-)gradient methods on the aggregate loss. Note however that, in general, as a function of the parameter $\boldsymbol{\theta}$, both the MSE, as well as the SP regularization are non-convex functions.

3 Related Works

A recent line of theoretical works on learning pseudo-Boolean functions demonstrate a staircase-like property of gradient descent in learning DNNs in that the WHT coefficients corresponding to higher order monomials (e.g., $x_1 x_2 x_3$) are reachable from lower order ones along increasing chains (i.e., $x_1 x_2$ and x_1), and are thus learnable in polynomial time and sample cost Abbe et al. (2021). Other works have shown that under certain distributions, low order parity functions are learnable by means of gradient decent on depth-2 networks, while they cannot be learned efficiently using linear methods Daniely & Malach (2020). These analyses are limited to certain DNN architectures or assume (linear) approximations to DNNs (e.g., neural tangent kernels) which entirely disallows the analysis of spectral regularization as they manifest only in nonlinear function classes.

Spectral bias of DNNs have been the subject of several other empirical and theoretical works Yang & Salman (2019). Approximations to the Fourier transform of two-layer Zhang et al. (2019) and multi-layer Rahaman et al. (2019) ReLU networks show that these networks have a learning bias towards low frequency functions Xu et al. (2019). To improve the limitations of DNNs in learning high frequency components, empirical works use Fourier features explicitly as part of the input Tancik et al. (2020). A parametrization of polynomial DNNs has also been shown to speed up the learning of higher frequency components in two-layer networks Choraria et al. (2022). Different notions of spectral priors have also been empirically investigated in graph neural networks Li et al. (2020) and elsewhere Yoshida & Miyato (2017). These results support the implicit bias of DNNs towards dense and low-frequency spectral representation and only motivate our analysis of sparsity as an explicit spectral prior.

4 Theoretical Analysis

To develop some intuition about the general problem, consider the special case of learning over the set of linear functions. Namely, $\mathcal{F} = \{\langle \boldsymbol{\theta}, \cdot \rangle : \boldsymbol{\theta} \in \mathbb{R}^d\}$. The polynomial representation of any linear function $f_{\boldsymbol{\theta}}(x) = \langle \boldsymbol{\theta}, x \rangle$ is simply $\sum_{i \in [d]} \theta_i x_i$ where only the order-1 coefficients are non-zero. Thus, the assumption that the polynomial representation of $f_{\boldsymbol{\theta}}$ is composed of at most k monomials is equivalent to saying that $\boldsymbol{\theta}$ is k -sparse. Thus in the linear setting, the problem boils down to sparse linear regression Candes et al. (2006). Furthermore, in the linear setting, spectral regularization has an explicit representation in terms of its parameter $\boldsymbol{\theta}$. In particular, here $R(\boldsymbol{\theta}) = \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1 = \lambda \|\boldsymbol{\theta}\|_1$. Thus, the proposed objective function,

(OBJ), when specialized to the linear setting boils down to the LASSO objective, $\frac{1}{n} \sum_{i=1}^n (\langle x^i, \theta \rangle - y^i)^2 + \lambda \|\theta\|_1$, which can be efficiently solved by gradient descent, and is known to be statistically optimal in the finite sample regime Raskutti et al. (2011).

4.1 Going beyond the linear setting - Restricted Secant Inequality

In the linear setting, a sufficient condition for the finite sample statistical optimality of LASSO is the restricted eigenvalue condition on the empirical covariance matrix $\Sigma_n = \frac{1}{n} \sum_{i=1}^n x^i x^{iT}$. The restricted eigenvalue condition is automatically satisfied if the smallest eigenvalue of Σ_n is bounded away from 0. The condition induces a particular sense of curvature around the true parameter θ^* of the loss $\text{Err}_n(\theta, \theta^*)$, that is the MSE of the learned function compared to the ground truth,

$$\text{Err}_n(\theta, \theta^*) = \frac{1}{n} \sum_{i=1}^n (f_\theta(x^i) - f_{\theta^*}(x^i))^2. \quad (3)$$

In contrast, when \mathcal{F} is a non-linear function family, we circumvent these sufficient conditions, and directly study assumptions which induce curvature in the loss $\text{Err}_n(\theta, \theta^*)$. These assumptions, however, do not concern the structure of the spectral regularizer, which unlike in the linear case, can no longer be represented as a closed-form function of θ . A major technical challenge in the analysis relates to circumventing the use of this explicit representation, which we expand upon later.

Definition 1 (Restricted secant inequality (RSI) Zhang & Yin (2013)). *The set of functions satisfying the restricted secant inequality with parameter C , denoted $\text{RSI}(C)$ is defined as follows. A function $g : \mathbb{R}^m \rightarrow \mathbb{R}$ belongs to $\text{RSI}(C)$ iff for some $z^* \in \arg \min_{z \in \mathbb{R}^m} g(z)$ and for all $z \in \mathbb{R}^m$,*

$$\langle \nabla g(z), z^* - z \rangle \geq C \|z^* - z\|^2. \quad (4)$$

Remark 2. *The RSI is known to generalize several extensions of convexity, such as quasar-star convexity Hinder et al. (2020) and strong star-convexity Lee & Valiant (2015). Refer to Karimi et al. (2016) for a comparison with other constraints such as the Polyak-Lojasiewicz (PL) inequality, and the quadratic growth condition which we study in Section 4.2. In general, the RSI is a significantly weaker condition than global strong convexity.*

The first main assumption we study is when the $\text{Err}_n(\theta, \theta^*)$ satisfies the RSI. This assumption captures a notion of curvature, in which the gradients of $\text{Err}_n(\theta, \theta^*)$ are informative about (i.e., positively correlated with) the error in the parameter space $\theta - \theta^*$.

From this behavior, it may be expected that all stationary points of functions which satisfy the RSI are global minima, which is indeed true Karimi et al. (2016). However, note that we impose the RSI condition on $\text{Err}_n(\theta, \theta^*)$ (which cannot be computed by the learner) and not the empirical risk $\mathcal{L}_n(\theta)$. Even under the RSI condition on $\text{Err}_n(\theta, \theta^*)$, it is not trivial to find global minima of the training objective, $\mathcal{L}_n(\theta) + R(\theta)$ which is composed of two non-convex (in θ) functions and is non-smooth. Thus, we restrict our attention to the analysis of *first-order stationary points* of this objective.

Assumption 1(a). *Assume the function $\text{Err}_n(\theta, \theta^*)$ satisfies the RSI with high probability over the dataset. Namely, there is a constant $C_{n,\delta}^* > 0$ such that with probability $\geq 1 - \delta$, for every $\theta \in \mathbb{R}^m$,*

$$\langle \theta - \theta^*, \nabla_\theta \text{Err}_n(\theta, \theta^*) \rangle \geq C_{n,\delta}^* \|\theta - \theta^*\|_2^2. \quad (5)$$

Remark 3. *The RSI is true for linear families $\mathcal{F} = \{\langle \theta, \cdot \rangle : \theta \in \mathbb{R}^m\}$ as long as with probability $\geq 1 - \delta$, the data covariance satisfies $\mathbb{E}_{x \sim \text{Unif}(D_n)}[xx^T] \succeq C_{n,\delta}^* I$, i.e., is well conditioned. The quadratic growth condition in Section 4.2 also holds under the same conditions with parameter $C_{n,\delta}^*$.*

In addition to the RSI, we make some mild regularity assumptions on the function classes we study. We emphasize that none of these assumptions imply convexity or smoothness of the training objective $\mathcal{L}_n(\theta) + R(\theta)$ or $\text{Err}_n(\theta, \theta^*)$. First, we assume that the function class \mathcal{F} is sufficiently smooth, having Lipschitz continuous gradients Cisse et al. (2017). Note that we study the case when the function class \mathcal{F} is smooth, and do not impose this condition on the empirical risk, $\mathcal{L}_n(\theta)$, as is often assumed Jin et al. (2018).

Assumption 2 (Lipschitz continuous gradients on \mathcal{F}). *For each $x \in \mathbf{X}$ and constant μ , assume that $f_{\boldsymbol{\theta}}(x)$ is twice differentiable in $\boldsymbol{\theta}$ and the Hessian of $f_{\boldsymbol{\theta}}(x)$ satisfies $-\mu I \leq \nabla^2 f_{\boldsymbol{\theta}}(x) \preceq \mu I$.*

The final assumption we impose assumes that the ground truth function $\boldsymbol{\theta}^*$ has well behaved gradients in that a certain covariance matrix induced by the gradients of the matrix is bounded.

Assumption 3 (Bounded average gradient norm at $\boldsymbol{\theta}^*$). *Under this assumption,*

$$\mathbb{E}_{x \sim \text{Unif}(\mathbf{X})} [\nabla f_{\boldsymbol{\theta}^*}(x)(\nabla f_{\boldsymbol{\theta}^*}(x))^T] \preceq L^2 I. \quad (6)$$

Assumption 3 can be interpreted as a one-point Lipschitzness condition, showing that on average across the inputs $x \in \mathbf{X}$, the gradients of $f_{\boldsymbol{\theta}}(x)$ are well behaved at the singular point $\boldsymbol{\theta} = \boldsymbol{\theta}^*$. Under the previous assumptions, we prove a bound on the error in the parameter space made by the learner. This theorem is a consequence of a more general result we prove in Appendix C for an arbitrary regularization scale $\lambda > 0$.

Theorem 1. *Suppose Assumptions 1(a), 2 and 3 are satisfied, λ is chosen as $4C_0\sigma\sqrt{\frac{d+\log(1/\delta)}{n}}$ for a sufficiently large constant $C_0 > 0$ and the size of the dataset is sufficiently large, namely $n > n_0$, defined as,*

$$n_0 \triangleq C_1 \frac{\sigma^2 (d^2 \mu^2 + L^2 k) (d + \log(1/\delta))}{(C_{n,\delta}^*)^2}, \quad (7)$$

for a sufficiently large constant $C_1 > 0$. Consider a learner which returns a first order stationary point, $\hat{\boldsymbol{\theta}}$, of the loss $\mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$. Then, with probability $\geq 1 - 2\delta$,

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \lesssim \frac{\sigma L}{C_{n,\delta}^*} \sqrt{\frac{k(d + \log(1/\delta))}{n}}. \quad (8)$$

Note that in the dependence on d and k , the sample size threshold n_0 scales asymptotically as $d^3 + dk$.

4.2 Going beyond the RSI - The Quadratic Growth condition

While RSI is a much weaker condition than convexity, at a high level, it assumes that the gradients of $\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$ are informative about the error in the parameter space $\boldsymbol{\theta} - \boldsymbol{\theta}^*$. We can further relax this assumption. In this section, we consider a weakening of this assumption which only supposes that $\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$ grows at least quadratically in $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2$. This is known as the *quadratic growth* (QG) condition Anitescu (2000). This condition no longer characterizes the behavior of the gradients of the function $\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$, much less those of the empirical loss \mathcal{L}_n . Stationary points of a function satisfying the QG condition are no longer global minima, in contrast with the behavior under the RSI (Definition 1).

Definition 2 (Quadratic Growth (QG) condition Anitescu (2000)). *The set of functions satisfying the QG condition with parameter C , denoted $\text{QG}(C)$ is defined by the inclusion, $g : \mathbb{R}^m \rightarrow \mathbb{R}$ belongs to $\text{QG}(C)$ iff for some $z^* \in \arg \min_{z \in \mathbb{R}^m} g(z)$ and for all $z \in \mathbb{R}^m$,*

$$g(z) - g(z^*) \geq C \|z^* - z\|_2^2. \quad (9)$$

For $C > 0$, functions in $\text{QG}(C)$ have a unique global minimizer.

Assumption 1(b) (Quadratic Growth (QG) condition Anitescu (2000)). *The function family $\mathcal{F} = \{f_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \mathbb{R}^m\}$ is assumed to satisfy the QG condition with parameter $C_{n,\delta}^*$, if with probability $\geq 1 - \delta$ over the dataset D_n ,*

$$\forall \boldsymbol{\theta} \in \mathbb{R}^m, \text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*) \geq C_{n,\delta}^* \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2. \quad (10)$$

Remark 4. (Anitescu, 2000, Theorem 2) *Restricting to functions f which have L -Lipschitz continuous gradients, if $f \in \text{RSI}(C)$, then it implies that $f \in \text{QG}(2C/L)$. This implies that up to the value of the parameter, the QG condition is weaker than RSI, under a smoothness constraint on the considered functions.*

Note that under the QG condition, gradients of $\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$ are no longer constrained to be positively correlated with the $\boldsymbol{\theta} - \boldsymbol{\theta}^*$. In fact, the absence of this feature proves to be a significant barrier for first-order methods to generalize well. To circumvent this issue, we motivate and introduce the notion of *approximate first-order stationary interpolators* in the next section, and characterize their statistical performance. This imposes a stronger requirement on the algorithm than just returning an arbitrary stationary point of the training objective $\mathcal{L}(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$.

Approximate First-order Interpolators. When \mathcal{F} is an expressive family of nonlinear models, it has been observed empirically that standard stochastic gradient methods can be run until the model begins to perfectly interpolate the training data, without hurting test-time performance. This phenomenon has been referred to in the literature as benign overfitting Bartlett et al. (2020), and has been seen to hold frequently in the training of DNNs in practice, when stochastic gradient descent is run for sufficiently many epochs. In fact, in all our experiments (see Figs. 5 and 6 in Appendix A), we observe that upon running stochastic gradient descent for sufficiently long on the aggregate loss $\mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$, the parameter $\hat{\boldsymbol{\theta}}$ eventually converges to a solution *which interpolates the labelled examples well*, i.e., the mean square error $\mathcal{L}_n(\hat{\boldsymbol{\theta}})$ is upper bounded by a sufficiently small $\Delta > 0$. Note that we *do not* assume that Δ is so small that the condition essentially imposes that $\hat{\boldsymbol{\theta}}$ is an approximate global minimizer of the training loss. In theory, our results also hold under the stronger condition that $\mathcal{L}_n(\hat{\boldsymbol{\theta}}) + R(\hat{\boldsymbol{\theta}})$ is upper bounded by Δ , which is the loss function being optimized by (stochastic) gradient descent. This phenomenon motivates the following definition of approximate first order stationary interpolators.

Definition 3 (Δ -approximate first order stationary interpolator). *A point $\hat{\boldsymbol{\theta}}$ is defined to be a Δ -approximate first-order stationary interpolator if,*

1. $\hat{\boldsymbol{\theta}}$ is a first order stationary point of the aggregate loss $\mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$. Namely,

$$0 \in (\nabla \mathcal{L}_n)(\hat{\boldsymbol{\theta}}) + (\nabla R)(\hat{\boldsymbol{\theta}}). \quad (11)$$

2. $f_{\hat{\boldsymbol{\theta}}}$ approximately interpolates the observed training data. Namely, the empirical mean squared error of $\hat{\boldsymbol{\theta}}$ evaluated on the training dataset satisfies,

$$\mathcal{L}_n(\hat{\boldsymbol{\theta}}) \leq \Delta. \quad (12)$$

Note that our results hold under the stronger condition $\mathcal{L}_n(\hat{\boldsymbol{\theta}}) + R(\hat{\boldsymbol{\theta}}) \leq \Delta$, the training loss function being minimized by the (stochastic) gradient methods in our experiments.

We are ready to establish a guarantee on the statistical performance of approximate first-order stationary interpolators.

Theorem 2. *Suppose λ is chosen $= 4C_0\sigma\sqrt{\frac{d+\log(1/\delta)}{n}}$ for a sufficiently large constant $C_0 > 0$. Assume that the size of the dataset is sufficiently large, namely, $n > n_0$ (as defined in Equation (7)). Consider a learner returns $\hat{\boldsymbol{\theta}}$ which is a Δ -approximate first order stationary interpolator, where $\Delta \leq C_1(C_{n,\delta}^*/\mu)^2$ for a sufficiently large constant C_1 . Then, with probability $\geq 1 - 2\delta$,*

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \lesssim \frac{L}{2C_{n,\delta}^*} \sqrt{\frac{kd}{n}}. \quad (13)$$

Remark 5. *For pseudo-Boolean functions, the log-covering number in any norm up to log factors is $\approx \log \binom{2^d}{k} \approx kd$. Therefore, an algorithm which returns a function which is an exact minimizer of the mean squared error among all functions with a k -sparse polynomial representation, $\boldsymbol{\theta}_k^{\text{LS}}$ admits the guarantee $\|\boldsymbol{\theta}_k^{\text{LS}} - \boldsymbol{\theta}^*\|_2 \lesssim \sqrt{kd/n}$ up to scaling constants, with high probability, up to log factors. However, when Δ is large, the learner considered in Theorem 2 is not constrained to return the exact minimizer of the squared error (subject to the sparsity constraint). Thus, under the additional condition of first-order stationarity, Theorem 2 imposes a much weaker condition, $\mathcal{L}_n(\hat{\boldsymbol{\theta}}) \lesssim \Delta$ rather than $\hat{\boldsymbol{\theta}} \in \arg \min \mathcal{L}_n(\boldsymbol{\theta})$.*




Figure 1: The plots demonstrate the generalization power of a depth-4 feed-forward neural network in learning 3 classes of sparse pseudo-Boolean functions $f(x) : \{\pm 1\}^{13} \rightarrow \mathbb{R}$ using the mean squared error (MSE) loss before (first row) and after (second row) adding the spectral (SP) regularizer, as a function of the number of training samples. MONOMIAL are 1-sparse functions in Walsh-Hadamard transform (WHT) with an increasing order of interactions (1 to 5). POWER LAW are 10-sparse functions in WHT with second order interactions and coefficients set using a power law function. STAIRCASE are 18-sparse functions in WHT with 3 first order, 6 second order, and 9 third order interactions, each with an equal coefficients and set using a power law function. The heat maps show the fraction of times (among 5 random independent trials) the models generalize on unseen data with the coefficient of determination larger than $R^2 \geq 0.45$ on the test data points.

Finally, in the context of the previous two results, we prove a lower bound showing statistical optimality under the imposed assumptions.

Theorem 3. Suppose $d \geq 3$ and $k \leq 2^d/4$. Then, there exists a parameter class Θ , and an associated function class $\mathcal{F} = \{f_\theta : \theta \in \Theta\}$ such that for any learner $\hat{\theta}$, there exists a ground truth function $f_{\theta^*} : \{\pm 1\}^d \rightarrow \mathbb{R}$ having a k -sparse polynomial representation, such that given a sufficiently large dataset of n samples,

1. Assumptions 1(a), 1(b), 2 and 3 are satisfied with constants $L = 1$, $\mu = 0$ and $C_{n,\delta^*} \in [\frac{1}{2}, \frac{3}{2}]$ with probability $\geq 1 - \delta$.
2. $\mathbb{E} [\|\hat{\theta} - \theta^*\|_2] \gtrsim \sigma \sqrt{\frac{kd}{n}}$, where σ^2 denotes the noise variance in the observed labels.

Extensions to cases with multiple local minima. While we focus on general parametric function classes in Theorems 1 and 2, for the specific case of DNNs, the issue of “permutation invariance” can appear. In particular, a permutations (i.e., relabeling) of the neurons in the same layer of a network can result in a different network with the same functional relationship between the input and the output. In this case, Assumptions 1(a) and 1(b) cannot hold globally for all $\theta \in \mathbb{R}^m$, as the loss $\text{Err}_n(\theta, \theta^*)$ and its gradient become 0 at any parameter θ_σ^* , where θ_σ^* denotes the weight matrices corresponding to a functionally invariant permutation σ of the neurons of the base network with parameter θ^* . However, even in this case, we can extract a guarantee from Theorem 1 and Theorem 2 by modifying the underlying assumptions slightly.

In particular, consider an arbitrary norm $\|\cdot\|$ and a nearest neighbor partitioning of the parameter space depending on the nearest permutation of θ^* . Namely, for a permutation σ , $\theta \in \mathcal{K}_\sigma \subset \mathbb{R}^m$ iff $\sigma = \arg \min_{\sigma'} \|\theta - \theta_{\sigma'}^*\|$. Suppose for each $\theta \in \mathbb{R}^m$, belonging to the partition \mathcal{K}_σ , the condition, $\langle \theta - \theta_\sigma^*, \nabla_\theta \text{Err}_n(\theta, \theta_\sigma^*) \rangle \geq C_{n,\delta}^* \|\theta - \theta_\sigma^*\|_2^2$ (modified Assumption 1(a)) and $\text{Err}_n(\theta, \theta_\sigma^*) \geq C_{n,\delta}^* \|\theta - \theta_\sigma^*\|_2^2$ (modified Assumption 1(b)) hold. In words, instead of globally, we impose the RSI and QG conditions in a local neighborhood (corresponding to the partition \mathcal{K}_σ) around each permutation θ_σ^* of the ground truth parameter θ^* . This avoids the issue alluded to before, since the modified versions of Assumptions 1(a) and 1(b) are now true trivially at $\theta = \theta_\sigma^*$, corresponding to any within-layer permutation of the neurons of θ^* .




Figure 2: We estimate the empirical lower bound $\hat{C}_{n,\delta}^*$ in the quadratic growth (QG) condition by finding the minimum change in the network’s output as a result of perturbations to the weights with noise drawn from the normal distribution $W \sim \mathcal{N}(0, \sigma^2)$. Experiments on 2 fully connected networks (FCNs) and a convolutional neural network (CNN) reveals empirical lower bounds.

Corresponding to the new modified assumptions, a learner which returns any $\hat{\theta}$ which is a stationary point under Assumption 1(a) (resp. approximate first order stationary interpolator under Assumption 1(b)), guarantees to approximate θ_σ^* , the nearest permutation of θ^* to it in the norm $\|\cdot\|$. The proofs of these results follow identically to Theorem 1 and Theorem 2. In particular, replacing θ^* by the parameter θ_σ^* where $\hat{\theta} \in \mathcal{K}_\sigma$ and following the same argument completes the proof of the result showing that θ_σ^* can be recovered approximately.

In summary, rather than assuming that the loss $\text{Err}_n(\theta, \theta^*)$ is globally bowl shaped as in Assumption 1(b), which cannot hold globally because of the permutation invariance of neural networks, it suffices to assume that the loss function is locally bowl shaped around each functionally invariant permutation θ_σ^* of the ground truth parameter θ^* . This line of reasoning can be extended to incorporate other symmetries in the function class which prevent exact parameter recovery, as in the case of permutation invariance of neural networks. The corresponding parameter recovery guarantee is also be modified to be up to this symmetry.

5 Empirical Studies

We design our experiments in a way to address these questions:

- Does SP improve generalization accuracy in learning sparse polynomials?
- When does QG hold for common DNN architectures? what is the empirical lower bound $\hat{C}_{n,\delta}^*$?
- Does SP improve generalization accuracy in real-world problems?
- How does L_1 -regularization compare to SP-regularization in practice in terms of generalization?

Sparse polynomials. We compare the generalization performance of a depth-4 fully connected network with and without the SP regularization on 3 classes of sparse polynomials in Fig. 1: 1) MONOMIAL are randomly-drawn 1-sparse functions in WHT with increasing order of interactions from 1 to 5. 2) POWER LAW are randomly-drawn 10-sparse functions in WHT all with order-2 interactions and coefficients set based on a power law function with decreasing exponent. 3) STAIRCASE are randomly-drawn 18-sparse functions in the WHT with 3 first order, 6 second order, and 9 third order interactions, each with equal coefficients and set based on a power law function. These synthetic sparse polynomials are inspired by physical models for real-world pseudo-Boolean functions (e.g., protein functions Brookes et al. (2022); Qian et al. (2001); Qin & Colwell (2018)). We observe clear transitions in generalization power: it is harder in terms of sample cost to learn nonlinear models with higher order interactions and lower sparsity exponent (i.e., denser functions in WHT). This analysis adds a new axis to the accuracy-vs-sample-cost phase transition curves in compressed sensing Donoho et al. (2011). SP regularization consistently improves the transition curves for generalization power as a function of order and sample cost.




Figure 3: The plot demonstrates the generalization error of depth-4 neural networks with the mean squared error (MSE) loss before and after adding the spectral (SP) regularizer as compared to Random Forest, XGBoost, LASSO, and SVM. SP drastically improves the generalization error of neural network (reduced mean and variance) and allows for a superior performance over LASSO with L_1 norm regularization over the polynomial representation directly.

QG condition. To empirically estimate the lower bound $\hat{C}_{n,\delta}^*$, we follow this procedure. We collect a set of training data points $(x^i, y^i)_{i=1}^{n=1000}$, initialize a DNN at θ^* (see below for more detail), and repeatedly perturb the weights with Gaussian noise to generate $\theta^{\text{pert},k} = \theta^* + W_k$, where $W \sim \mathcal{N}(0, \sigma^2 I)$ is independent and normally distributed, for $k = 1, \dots, K$. For each k , we compute the ratio $\sum_{i=1}^n (f_{\theta^{\text{pert},k}}(x^i) - f_{\theta^*}(x^i))^2 / \sigma^2$ and report the minimum value across $k \in [K]$ as a function of σ^2 in Fig. 2. Next, we train the DNN with SP regularization for certain number of epochs to arrive at a new θ^* and repeat the same procedure again to observe how the lower bound changes with training. The outputs y are generated from binary vectors $x \in \{\pm 1\}^{13}$ using the (randomly-drawn) sparse polynomial $f(x) = 3x_1 + 4x_2x_3 + 5x_4x_5 + x_{12}$. We choose sufficiently small architectures compared to the number of perturbations to ensure a reliable empirical estimate for $\hat{C}_{n,\delta}^*$. Fig. 2 shows the results for Xavier-initialized, depth-2 FCNs with 222 parameters, perturbed $K = 100$ (FCN-1) and $K = 500$ (FCN-2) times, with different ranges of σ^2 (learning rate $= 5 \times 10^{-3}$). Fig. 2 also shows the result for a depth-4 CNN with 1551 parameters, initialized by a trained network using MSE loss and perturbed $K = 6000$ times (learning rate $= 10^{-3}$). The plots demonstrate that the QG condition is satisfied for these instances with empirical lower bounds 6.02 and 19.5, respectively. The bound also improves drastically with training; moreover, initializing the networks with trained models consistently improves the lower bound. See Appendix A for more detailed empirical studies and discussions of these results.

Real world experiments. We compare the generalization error with and without the SP regularization to standard baseline algorithms for data-scarce learning in the real-world (Fig. 3). PROTEIN is a dataset which measures the fluorescence level of 2^{13} protein sequences that link two variants of the *Entacmaea quadricolor* proteins different at exactly 13 amino acids Poelwijk et al. (2019). T CELL is a dataset which measures the DNA repair outcome of T cells (average length of deletions) on 1521 sites on human genome after applying double-strand breaks (DSBs) using CRISPR Leenay et al. (2019). CANCER is a similar dataset on 287 sites on cancer genome Leenay et al. (2019). In the two latter datasets, a one-hot-encoded context sequence of size 20 around the DSB is used as the input to predict the DNA repair outcome. Following the low- n experimental setting in Biswas et al. (2021), we use a subset of 30 sequences drawn uniformly at random for training and validation and use the rest for testing. We repeat each experiments $10\times$ with independent random splits of the data and report the RMSE in predicting the phenotype. We initialize DNNs using Xavier (equal seeds). We use the default hyperparameters in scikit-learn for the baselines. Despite minimal hyperparameter tuning and no architecture search, SP allows for a competitive performance. In particular, the SP-regularized model outperforms LASSO which applies the L_1 norm penalty on the coefficients in the polynomial representation (Fig. 8 in Appendix A compares SP with different regularizers).

6 Summary, Discussion, and Future Vision

In the paper, we focused on learning sparse pseudo-Boolean functions with explicit regularization in the spectral domain and established the conditions on general nonlinear functions under which the stationary points which approximately interpolate the training data achieve statistically optimal generalization performance. We demonstrated how the assumptions can be extended to nonlinear functions with multiple local minimas. Real-world experiments with neural networks demonstrated the utility of these assumptions. Naturally, some of the bounds might result in larger constants for very deep neural network (due to higher Lipschitz constants). Future works involve analyzing the algorithms that can achieve such statistical performance. It would be tempting to ask under what conditions stochastic gradient descent achieves stationary points with restricted secant and quadratic growth conditions (for which we have clear empirical evidences). Further, the statistical analysis of the computationally-efficient optimization algorithms for spectral regularization is still poorly understood. On the algorithmic side, spectral algorithms for pseudo-Boolean functions can be extended to generalized Fourier transform to accommodate a larger class of data-frugal combinatorial problems. Overall, our work provides a concrete framework to connect combinatorial algorithms with strong theoretical guarantees and nonlinear machine learning models with strong generalization power.

Broader Impact Statement

We do not anticipate this work to have any potential negative impact on a user of this research.

References

- E. Abbe, E. Boix Adsera, M. Brennan, G. Bresler, and D. Nagaraj. The staircase property: How hierarchical structure can guide deep learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- A. Aghazadeh, H. Nisonoff, O. Ocal, D. Brookes, Y. Huang, O. Koyluoglu, J. Listgarten, and K. Ramchandran. Epistatic net allows the sparse spectral regularization of deep neural networks for inferring fitness functions. *Nature Communications*, 12(1):1–10, 2021.
- A. Amrollahi, A. Zandieh, M. Kapralov, and A. Krause. Efficiently learning Fourier sparse set functions. *Advances in Neural Information Processing Systems*, 32, 2019.
- M. Anitescu. Degenerate nonlinear programming with a quadratic growth condition. *SIAM Journal on Optimization*, 10(4):1116–1135, 2000.
- P. Bartlett, P. Long, G. Lugosi, and A. Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- S. Biswas, G. Khimulya, E. Alley, K. Esvelt, and G. Church. Low-N protein engineering with data-efficient deep learning. *Nature Methods*, 18(4):389–396, 2021.
- E. Boros and P. Hammer. Pseudo-Boolean optimization. *Discrete Applied Mathematics*, 123(1-3):155–225, 2002.
- D. Brookes, A. Aghazadeh, and J. Listgarten. On the sparsity of fitness functions and implications for learning. *Proceedings of the National Academy of Sciences*, 119(1), 2022.
- E. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.
- G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4):045002, 2019.

- T. Ching, D. Himmelstein, B. Beaulieu-Jones, A. Kalinin, B. Do, G. Way, E. Ferrero, P. Agapow, M. Zietz, M. Hoffman, et al. Opportunities and obstacles for deep learning in biology and medicine. *Journal of The Royal Society Interface*, 15(141):20170387, 2018.
- M. Choraria, L. Dadi, G. Chrysos, J. Mairal, and V. Cevher. The spectral bias of polynomial neural networks. *arXiv preprint arXiv:2202.13473*, 2022.
- M. Cisse, P. Bojanowski, E. Grave, Y. Dauphin, and N. Usunier. Parseval networks: Improving robustness to adversarial examples. In *International Conference on Machine Learning*, pp. 854–863. PMLR, 2017.
- A. Daniely and E. Malach. Learning parities with neural networks. *Advances in Neural Information Processing Systems*, 33:20356–20365, 2020.
- D. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- D. Donoho, A. Maleki, and A. Montanari. The noise-sensitivity phase transition in compressed sensing. *IEEE Transactions on Information Theory*, 57(10):6920–6941, 2011.
- H. Eble, M. Joswig, L. Lamberti, and W. Ludington. Higher-order interactions in fitness landscapes are sparse. *arXiv preprint arXiv:2009.12277*, 2020.
- Á. Figula. and V. Kvaratskhelia. Some numerical characteristics of Sylvester and Hadamard matrices, 2015.
- S. Gelman, S. Fahlberg, P. Heinzelman, P. Romero, and A. Gitter. Neural networks to learn protein sequence–function relationships from deep mutational scanning data. *Proceedings of the National Academy of Sciences*, 118(48), 2021. doi: 10.1073/pnas.2104878118.
- W. Ha, C. Singh, F. Lanusse, S. Upadhyayula, and B. Yu. Adaptive wavelet distillation from neural networks through interpretations. *Advances in Neural Information Processing Systems*, 34, 2021.
- O. Hinder, A. Sidford, and N. Sohoni. Near-optimal methods for minimizing star-convex functions and beyond. In *Conference on learning theory*, pp. 1894–1938. PMLR, 2020.
- C. Jin, L. Liu, R. Ge, and M. Jordan. On the local minima of the empirical risk. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 795–811. Springer, 2016.
- J. Lee and P. Valiant. Optimizing star-convex functions, 2015.
- R. Leenay, A. Aghazadeh, J. Hiatt, D. Tse, T. Roth, R. Apathy, E. Shifrut, J. Hultquist, N. Krogan, Z. Wu, et al. Large dataset enables prediction of repair after CRISPR–Cas9 editing in primary T cells. *Nature Biotechnology*, 37(9):1034–1037, 2019.
- M. Li, Z. Ma, Y. Wang, and X. Zhuang. Fast Haar transforms for graph neural networks. *Neural Networks*, 128:188–198, 2020.
- X. Li and K. Ramchandran. An active learning framework using sparse-graph codes for sparse polynomials and graph sketching. *Advances in Neural Information Processing Systems*, 28, 2015.
- X. Li, J. Bradley, S. Pawar, and K. Ramchandran. SPRIGHT: A fast and robust framework for sparse Walsh-Hadamard transform. *arXiv preprint arXiv:1508.06336*, 2015.
- P. Loh and M. Wainwright. Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima, 2015.

- F. Noé, S. Olsson, J. Köhler, and H. Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457), 2019.
- F. Poelwijk, M. Socolich, and R. Ranganathan. Learning the pattern of epistasis linking genotype and phenotype in a protein. *Nature Communications*, 10(1):1–11, 2019.
- J. Qian, N. Luscombe, and M. Gerstein. Protein family and fold occurrence in genomes: power-law behaviour and evolutionary model. *Journal of Molecular Biology*, 313(4):673–681, 2001.
- C. Qin and L. Colwell. Power law tails in phylogenetic systems. *Proceedings of the National Academy of Sciences*, 115(4):690–695, 2018.
- N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301–5310. PMLR, 2019.
- G. Raskutti, M. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994, 2011. doi: 10.1109/TIT.2011.2165799.
- A. Riesselman, J. Ingraham, and D. Marks. Deep generative models of genetic variation capture the effects of mutations. *Nature Methods*, 15(10):816–822, 2018.
- J. Rohn. Computing the norm $\|A\|_{\infty,1}$ is NP-hard. *Linear and Multilinear Algebra*, 47:195–204, 05 2000.
- K. Sarkisyan, D. Bolotin, M. Meer, D. Usmanova, A. Mishin, G. Sharonov, D. Ivankov, N. Bozhanova, M. Baranov, O. Soylemez, et al. Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603):397–401, 2016.
- M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- J. Tropp and A. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007.
- J. A Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- A. Tseng, A. Shrikumar, and A. Kundaje. Fourier-transform-based attribution priors improve the interpretability and stability of deep learning models for genomics. *Advances in Neural Information Processing Systems*, 33:1913–1923, 2020.
- Z. Xu, Y. Zhang, T. Luo, Y. Xiao, and Z. Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019.
- G. Yang and H. Salman. A fine-grained spectral perspective on neural networks. *arXiv preprint arXiv:1907.10599*, 2019.
- Y. Yoshida and T. Miyato. Spectral norm regularization for improving the generalizability of deep learning. *arXiv preprint arXiv:1705.10941*, 2017.
- H. Zhang and W. Yin. Gradient methods for convex minimization: better rates under weaker conditions. *arXiv preprint arXiv:1303.4645*, 2013.
- Y. Zhang, Z. Xu, T. Luo, and Z. Ma. Explicitizing an implicit bias of the frequency principle in two-layer neural networks. *arXiv preprint arXiv:1905.10264*, 2019.

A Additional Empirical Validations

A.1 Visualization of real-world functions in WHT

In this subsection, we visualize the combinatorial function evaluations and the WHT of the evaluation vector, that is, the left and right hand sides of the equation below:

$$[f(x) : x \in \{\pm 1\}^d]^T = \mathbf{H} [\alpha_z : z \in \{\pm 1\}^d]^T, \quad (14)$$

for the experimental data obtained from the fluorescence protein Poelwijk et al. (2019). In addition to these plots, we have attached videos as supplemental materials where we visualize the learning trajectory of a DNN initialized using Xavier initialization under data-scarce (red) and data-sufficient (blue) regimes. DNNs start from a local minima that does not have a well-structured WHT representation, and then gets sparser with training. However, if a sufficient amount of data is not available for training, the network does not converge to a good solution. Spectral regularisation enable DNN converge to the sparse solution in the data-scarce regime.

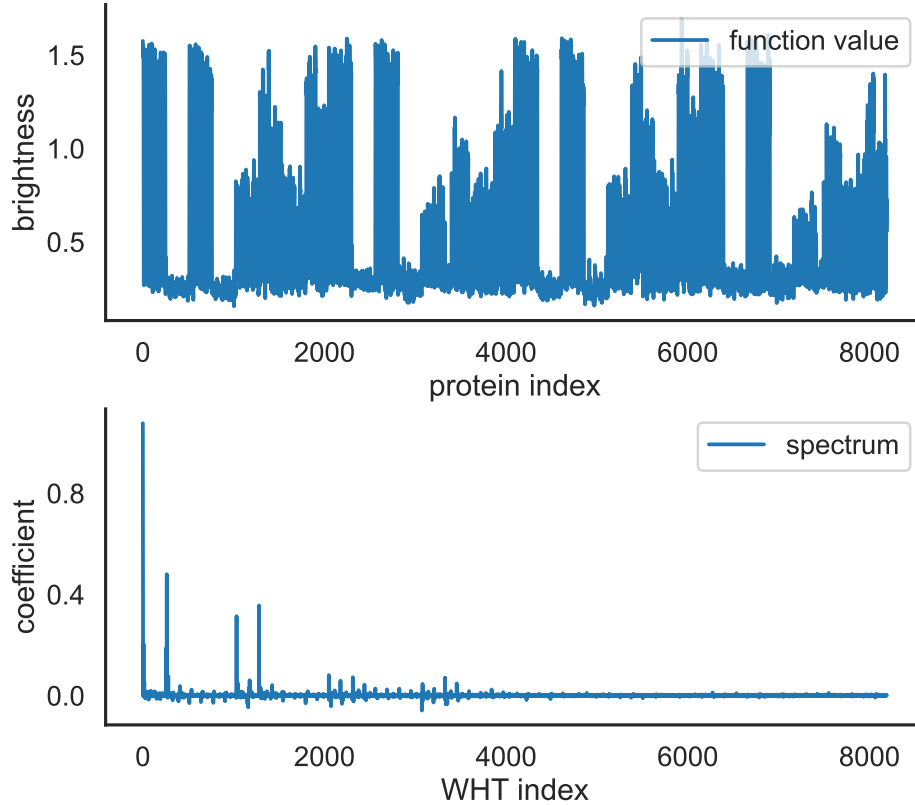


Figure 4: Combinatorial visualization of the brightness of 2^{13} proteins (top) and their Walsh-Hadamard transform (bottom) linking two variants of a fluorescence protein that are different in 13 locations on their amino acid chain Poelwijk et al. (2019). The brightness is a pseudo-Boolean function which maps from $f : \{\pm 1\}^{13} \rightarrow \mathbb{R}$. The sparse spectrum reveals low and high order interactions among amino acid sites on the protein.

A.2 Convergence: training and validation

In this subsection, we include additional empirical results which focus on the convergence properties of spectral regularization for training DNNs. The first objective we study is in validating whether stochastic gradient descent (SGD) indeed converges to approximate first order stationary interpolators (Definition 3).

The goal of these experiments are to see whether, upon running SGD, the MSE loss gets sufficiently small even when the training loss, against which stochastic gradients are computed, is augmented with spectral regularization. In Fig. 5, we plot the empirical training and validation loss of DNNs trained both with the MSE loss and the additional spectral regularization. We use a depth-4 fully connected network (learning rate= 1×10^{-1}) to train on the fluorescence protein Poelwijk et al. (2019) dataset (also used in the main paper).




Figure 5: We plot the training and validation error in terms of the empirical mean squared error (MSE) in predicting the brightness of a fluorescence protein Poelwijk et al. (2019) using DNNs trained with and without the spectral regularization. We use a depth-4 fully connected neural network with Xavier initialization in both cases. The plots demonstrate that 1) the training MSE eventually gets sufficiently small, even when the original loss function on which SGD is run is augmented with spectral regularization, even when network uses a random initialization and 2) spectral regularization consistently allows for a significantly better generalization gap all along the trajectory of training process.




Figure 6: We repeat the experiment above with the only difference that instead of a Xavier initialization, we carry out what is known as “weight initialization” (Definition 4). Here, we train the DNN with spectral regularization (blue curve) in comparison with a DNN without spectral regularization (orange curve), but with weight initialization. The plots demonstrate that even with weight initialization, spectral regularization does not result in an increase in training, and results in a decrease in validation MSE. Weight initialization has the advantage that the QG constant encountered along the training trajectory of SGD are higher (Figure 7) compared to with a Xavier initialization

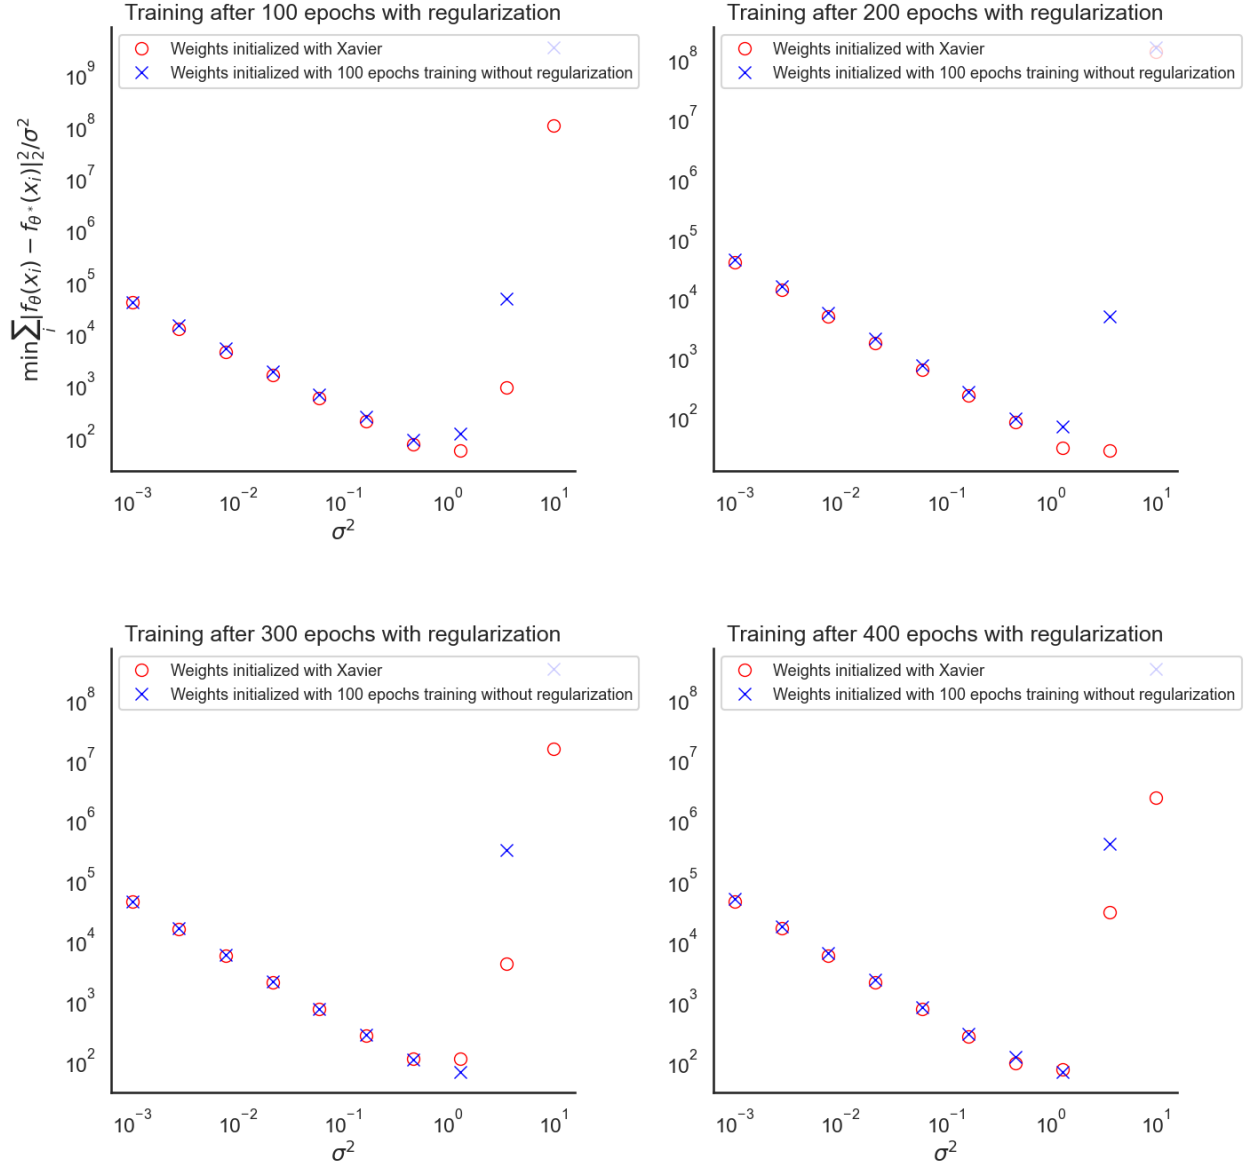


Figure 7: We plot the minimum ratio $\min_{k \in [K]} \sum_{i=1}^n (f_{\theta^{\text{pert},k}}(x^i) - f_{\theta^*}(x^i))^2 / \sigma^2$ against σ^2 for 4-layer CNN models trained initialized from (i) Xavier initialization, and (ii) weight initialization, with a learning rate of 10^{-3} . Similar to the experiments plotted in Fig. 2, we perturb θ^* , $K = 6000$ times to generate $\theta^{\text{pert},k}$.

A.3 Validating assumption 1(b): additional plots

In this subsection, we include additional plots on empirical experiments on validating **Assumption 1(b)**. The goal of this experiment is to also provide more details about the new weight initialization method discussed in the Empirical Studies section of the main paper. Note that Assumption 1(b) requires showing that for any θ ,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} \left[(f_{\theta}(x) - f_{\theta^*}(x))^2 \right] \geq C_{n,\delta}^* \|\theta - \theta^*\|_2^2. \quad (15)$$

However, since it is prohibitively expensive to check this for all choices of θ , and moreover since the models we consider exhibit some smoothness, we resort to checking this condition only for randomly sampled θ around




Figure 8: Comparison of the generalization error of different regularization schemes: spectral (SP), L1, and L2 norm regularization.

the reference θ^* . In particular, we collect a set of training data points $(x^i, y^i)_{i=1}^{n=1000}$ from an arbitrary sparse polynomial, $f(x) : \{\pm 1\}^{13} \rightarrow \mathbb{R}$, defined as $3x_1 + 4x_2x_3 + 5x_4x_5 + x_{12}$.

First we train a DNN using SGD on the unregularized MSE, and define this as θ^* . Then, we repeatedly perturb the weights K times independently with Gaussian noise to generate $\theta^{\text{pert},k} = \theta^* + W_k$, where $W \sim \mathcal{N}(0, \sigma^2 I)$ is independent and normally distributed, for $k = 1, \dots, K$. By concentration of measure, under the Gaussian perturbations, note that $\|\theta^{\text{pert},k} - \theta^*\|_2^2$ concentrates around $M\sigma^2$, where M is the number of parameters in the network. Therefore, instead of computing the ratio,

$$C_{n,\delta}^* \min_{\theta} \frac{\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\theta}(x) - f_{\theta^*}(x))^2]}{\|\theta - \theta^*\|_2^2} \quad (16)$$

we instead approximate it by the ratio (up to scaling by M),

$$\min_{k \in [K]} \frac{\sum_{i=1}^n [(f_{\theta^{\text{pert},k}}(x^i) - f_{\theta^*}(x^i))^2]}{\sigma^2} \quad (17)$$

as an approximate proxy for perturbations at the scale of σ around θ^* in L_{∞} distance, which is more accurate, as K grows larger. In Fig. 2, we plot the ratio in eq. (17) as a function of σ .

We repeat this experiment when θ^* is trained starting from weight initialization, and plot the estimated ratio in Fig. 7.

Definition 4 (Weight initialization). *Weight initialization, refers to initialization of the network by first running 100 epochs of SGD against the unregularized MSE. From this starting point, subsequently we “turn on” the regularization and run SGD against the MSE with spectral regularization.*

Each subplot shows the estimated QG ratio (eq. (17)) for models initialized with weight initialization, at the 100, 200, 300 and 400 epochs mark respectively. The plots demonstrate that our weight initialization method significantly improves the QG constant $C_{n,\delta}^*$ along sample trajectory encountered by the SGD, compared to models trained with spectral regularization starting from a random Xavier initialization.

Remark 6. *The guarantees hold for a specific path of regularization parameter λ that might depend on the level of noise. Therefore, in practice, the regularization parameter λ is chosen by cross-validation.*

A.4 Comparison to other regularization schemes

In this subsection, we compare the performance of neural network under SP regularization with other regularization schemes which directly penalize the ℓ_1 and ℓ_2 norm of the weights of the neural network. The goal

of these experiments is to show that promoting sparsity among the weights does not have the same effect as promoting sparsity in the spectral representation of neural networks. We test the generalization power of these networks using the datasets in Fig. 3 under the same experimental conditions. In Fig. 8 we demonstrate that while direct weight-regularization schemes such as ℓ_1 and ℓ_2 norm improve the generalization gap of neural networks, the generalization gap is significantly larger for SP regularization.

B Technical lemmas

In this section we outline and prove the lemmas used for the proof of the main theorems.

Lemma 1. $\|f_{\theta^*}(\mathbf{X}) - f_{\hat{\theta}}(\mathbf{X})\|_2 \leq 2\mu\sqrt{2^d}\|\hat{\theta} - \theta^*\|_2^2 + 2L\sqrt{2^d}\|\hat{\theta} - \theta^*\|_2.$

Proof. By the Lagrange form of the Taylor series expansion, for any x , there exists a θ_x such that,

$$g_{\hat{\theta}}(x) - g_{\theta^*}(x) = (\hat{\theta} - \theta^*)^T (\nabla^2 g_{\theta_x}(x)) (\hat{\theta} - \theta^*) + \langle \nabla g_{\theta^*}(x), \hat{\theta} - \theta^* \rangle \quad (18)$$

Then, for each $x \in \mathbf{X}$,

$$(g_{\hat{\theta}}(x) - g_{\theta^*}(x))^2 \leq 2\mu\|\hat{\theta} - \theta^*\|_2^4 + 2\langle \nabla g_{\theta^*}(x), \hat{\theta} - \theta^* \rangle^2 \quad (19)$$

Summing over $x \in \mathbf{X}$ results in,

$$\|\nu\|_2^2 = \|g_{\hat{\theta}}(\mathbf{X}) - g_{\theta^*}(\mathbf{X})\|_2^2 \leq 2\mu^2 2^d \|\hat{\theta} - \theta^*\|_2^4 + L^2 2^d \|\hat{\theta} - \theta^*\|_2. \quad (20)$$

where recall the assumption, $\frac{1}{2^d} \sum_{x \in \mathbf{X}} \nabla g_{\theta^*}(x) (\nabla g_{\theta^*}(x))^T \preceq L^2$.

□

Lemma 2. $\max_{v: \|v\|_\infty \leq 1} \|\mathbf{H}v\|_1 \leq (2d+1)\sqrt{2^d}.$

Proof. Note that $\max_{v: \|v\|_\infty \leq 1} \|\mathbf{H}v\|_1 = \max_{v \in \{-1, +1\}^{2^d}} \|\mathbf{H}v\|_1$, by (Rohn, 2000, Proposition 1). Therefore, it suffices to maximize over $v \in \{-1, +1\}^{2^d}$. It is known from Figula. & Kvaratskhelia (2015) that $\hat{\varrho}^{(d)} \leq d2^d$ where $\hat{\varrho}^{(d)} \triangleq \max_{m \in [2^d]} \sqrt{2^d} \|\Phi \mathbf{H} \mathbf{1}_m\|_1$ where Φ is an arbitrary set of vectors from the unit $\|\cdot\|_1$ ball in \mathbb{R}^{2^d} , $\mathbf{1}_m$ is the vector with the first m entries 1 and the remaining entries 0. Since Φ can be chosen as any permutation matrix, this result implies that $\max_{v \in \{0, 1\}^{2^d}} \|\mathbf{H}v\|_1 \leq d2^d$. To compare with the previous statement, note that the number of non-zeros in v here equals m in the optimization problem defining $\hat{\varrho}^{(d)}$. Then, we have that,

$$\max_{v \in \{-1, +1\}^{2^d}} \|\mathbf{H}v\|_1 = \max_{v \in \{0, 1\}^{2^d}} \|\mathbf{H}(2v - \mathbf{1})\|_1 \quad (21)$$

$$\stackrel{(i)}{\leq} 2\|\mathbf{H}v\|_1 + \|\mathbf{H}\mathbf{1}\|_1 \quad (22)$$

$$\leq 2d\sqrt{2^d} + \sqrt{2^d}, \quad (23)$$

where (i) follows by triangle inequality. □

Lemma 3. Under Assumptions 1(a) and 2, for any $\theta \in \mathbb{R}^d$ and any subgradient $G \in (\nabla R)(\theta)$, under the event \mathcal{E}_1 ,

$$\langle \theta - \theta^*, \nabla \mathcal{L}_n(\theta) + G \rangle - 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\theta^*}(x) - y(x)) \langle \theta - \theta^*, \nabla f_{\theta}(x) \rangle] \quad (24)$$

$$\geq \left(\frac{C_{n, \delta}^*}{2} - (2d+1)\lambda\mu \right) \|\theta - \theta^*\|_2^2 + R(\theta) - R(\theta^*) \quad (25)$$

Proof. Observe that,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) \rangle = 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (26)$$

Plugging this in below,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) \rangle - 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (27)$$

$$= 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla f_{\boldsymbol{\theta}}(x) \rangle - (f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (28)$$

$$= 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}}(x) - f_{\boldsymbol{\theta}^*}(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (29)$$

$$= \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}}(x) - f_{\boldsymbol{\theta}^*}(x))^2] \rangle \quad (30)$$

$$\geq C_{n,\delta}^* \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \quad (31)$$

where the last inequality follows by the restricted secant condition imposed on $\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$ in Assumption 1(a). Putting together eq. (31) with Lemma 5,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) + G \rangle - 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (32)$$

$$\geq C_{n,\delta}^* \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 + R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}^*) - (2d+1)\lambda\mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \quad (33)$$

This completes the proof. \square

In the remaining proofs, we will assume a lower bound on the regularization parameter to ensure that the regularization actually plays a role. In particular,

Condition 1. For a sufficiently large absolute constant $C_0 > 0$, assume that the regularization parameter λ satisfies,

$$\lambda \geq 4C_0\sigma \sqrt{\frac{d + \log(1/\delta)}{n}}. \quad (34)$$

Lemma 4. Suppose λ satisfies the lower bound in Condition 1. Define \mathcal{E}_2 as the event that, for all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta}^* - \boldsymbol{\theta}, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (35)$$

$$\leq \frac{\lambda}{4} \sqrt{\frac{1}{2^d}} \|\mathbf{H}(f_{\boldsymbol{\theta}^*}(\mathbf{X}) - f_{\boldsymbol{\theta}}(\mathbf{X}))\|_1 + \frac{\lambda\mu(2d+1)}{4} \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2. \quad (36)$$

Then, $\Pr(\mathcal{E}_2) \geq 1 - \delta$.

Proof. For each sample $x' \in D_n$, define the noise in the sample as $z(x') = f_{\boldsymbol{\theta}^*}(x') - y(x')$. Likewise, define the noise vector $\mathbf{z}_{D_n} = \{\mathbb{E}_{x \sim \text{Unif}(D_n)}[z(x')\mathbb{I}(x' = x)] : x \in [2^d]\}$. Note that the coordinates of \mathbf{z}_{D_n} corresponding to inputs unobserved in the dataset D_n are 0.

By Taylor series expanding, $f_{\boldsymbol{\theta}^*}(x) - f_{\boldsymbol{\theta}}(x) = \langle \boldsymbol{\theta}^* - \boldsymbol{\theta}, \nabla f_{\boldsymbol{\theta}}(x) \rangle + (\boldsymbol{\theta}^* - \boldsymbol{\theta})^T (\nabla^2 f_{\boldsymbol{\theta}_x}(x)) (\boldsymbol{\theta}^* - \boldsymbol{\theta})$ for some $\boldsymbol{\theta}_x \in \text{conv}(\{\boldsymbol{\theta}, \boldsymbol{\theta}^*\})$. Therefore,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta}^* - \boldsymbol{\theta}, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (37)$$

$$= \mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) (f_{\boldsymbol{\theta}^*}(x) - f_{\boldsymbol{\theta}}(x) - (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T (\nabla^2 f_{\boldsymbol{\theta}_x}(x)) (\boldsymbol{\theta} - \boldsymbol{\theta}^*))] \quad (38)$$

$$= \langle \mathbf{z}_{D_n}, f_{\boldsymbol{\theta}^*}(\mathbf{X}) - f_{\boldsymbol{\theta}}(\mathbf{X}) - \mathbf{A}(\boldsymbol{\theta}, \boldsymbol{\theta}^*) \rangle, \quad (39)$$

where $\mathbf{A}(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$ denotes the vector $\{(\boldsymbol{\theta}^* - \boldsymbol{\theta})^T \nabla^2 f_{\boldsymbol{\theta}_x}(x) (\boldsymbol{\theta}^* - \boldsymbol{\theta}) : x \in [2^d]\}$.

By an application of Holder's inequality and triangle inequality of the L_1 -norm, we have,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta}^* - \boldsymbol{\theta}, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (40)$$

$$\leq \|\mathbf{H}\mathbf{z}_{D_n}\|_{\infty} \|\mathbf{H}(f_{\boldsymbol{\theta}^*}(\mathbf{X}) - f_{\boldsymbol{\theta}}(\mathbf{X}))\|_1 + \|\mathbf{H}\mathbf{z}_{D_n}\|_{\infty} \|\mathbf{H}\mathbf{A}(\boldsymbol{\theta}, \boldsymbol{\theta}^*)\|_1 \quad (41)$$

Note that for each fixed row $i \in [2^d]$, $\langle \mathbf{H}_i, \mathbf{z}_{D_n} \rangle = \sum_{j \in [2^d]} \mathbf{H}_{ij} \mathbf{z}_{D_n}(j)$. Note that the coordinates of \mathbf{z}_{D_n} are independently distributed and subgaussian. Therefore, by Hoeffding's inequality, for some constant C_0 ,

$$\Pr \left(\langle \mathbf{H}_i, \mathbf{z}_{D_n} \rangle \geq C_0 \sqrt{\frac{\left(\sum_{j \in [2^d]} \text{Var}(\mathbf{z}_{D_n}(j)) \right) \log(1/\delta)}{2^d}} \right) \leq \delta \quad (42)$$

Note that the coordinate of \mathbf{z}_{D_n} labelled by $x \in \{\pm 1\}^d$, $\mathbb{E}_{x \sim \text{Unif}(D_n)}[z(x') \mathbb{I}(x' = x)]$, is the sum of $D_n(x)$ independent $\mathcal{N}(0, \sigma^2)$ Gaussians scaled by $1/n$, where $D_n(x)$, defined as the number of times x is sampled in D_n . Therefore, $\text{Var}(\mathbf{z}_{D_n}(i)) = \frac{\sigma^2 D_n(x)}{n^2}$. By union bounding over the 2^d rows of \mathbf{H} , with probability $\geq 1 - \delta$,

$$\Pr \left(\|\mathbf{H} \mathbf{z}_{D_n}\|_\infty \geq C_0 \sqrt{\frac{\left(\sum_{x \in [2^d]} \frac{\sigma^2 D_n(x)}{n^2} \right) \log(2^d/\delta)}{2^d}} \right) \leq \delta. \quad (43)$$

Note that $\sum_{x \in [2^d]} D_n(x) = n$, and therefore, with probability $\geq 1 - \delta$,

$$\|\mathbf{H} \mathbf{z}_{D_n}\|_\infty \leq C_0 \sigma \sqrt{\frac{1}{2^d}} \sqrt{\frac{d + \log(1/\delta)}{n}} \leq \frac{\lambda}{4\sqrt{2^d}}. \quad (44)$$

where the last inequality follows by the assumption on λ in Condition 1.

Note that $\|\mathbf{A}(\boldsymbol{\theta}, \boldsymbol{\theta}^*)\|_\infty \leq \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2$. Thus, the term $\|\mathbf{H} \mathbf{A}(\boldsymbol{\theta}, \boldsymbol{\theta}^*)\|_1$ can be upper bounded by $\sup_{v: \|v\|_\infty \leq \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2} \|\mathbf{H} v\|_1$. This itself can be further upper bounded using Lemma 2 by $\mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 (2d+1) \sqrt{2^d}$.

All in all, combining this argument and eq. (44) with eq. (41), under the events \mathcal{E}_1 and \mathcal{E}_2 , with probability $\geq 1 - \delta$,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta}^* - \boldsymbol{\theta}, \nabla f_{\boldsymbol{\theta}}(x) \rangle] \quad (45)$$

$$\leq \frac{\lambda}{4} \sqrt{\frac{1}{2^d}} \|\mathbf{H} (f_{\boldsymbol{\theta}^*}(\mathbf{X}) - f_{\boldsymbol{\theta}}(\mathbf{X}))\|_1 + \frac{\lambda \mu (2d+1)}{4} \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2. \quad (46)$$

□

Lemma 5. *Under the Lipschitz gradient condition, Assumption 2, consider any $\boldsymbol{\theta} \in \mathbb{R}^d$ and any subgradient $G \in (\nabla R)(\boldsymbol{\theta})$. Then,*

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, G \rangle \geq R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}^*) - (2d+1) \lambda \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2. \quad (47)$$

Proof. Recall that the regularization function $R(\boldsymbol{\theta})$ is defined as $\frac{\lambda}{\sqrt{2^d}} \|\mathbf{H} f_{\boldsymbol{\theta}}(\mathbf{X})\|_1$. Consider any subgradient $G \in (\nabla R)(\boldsymbol{\theta})$ of the regularization function $R(\boldsymbol{\theta})$. By Proposition 1, this is of the form $\frac{\lambda}{\sqrt{2^d}} (\nabla f_{\boldsymbol{\theta}}(\mathbf{X})) \mathbf{H} z$, where $z \in \text{sgn}(\mathbf{H} f_{\boldsymbol{\theta}}(\mathbf{X}))$. In particular, using this representation, we have the equation,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, G \rangle = \frac{\lambda}{\sqrt{2^d}} (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T (\nabla f_{\boldsymbol{\theta}}(\mathbf{X})) \mathbf{H} z. \quad (48)$$

Since $f_{\boldsymbol{\theta}}(x)$ is in general non-linear in $\boldsymbol{\theta}$, we can relate $(\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T \nabla f_{\boldsymbol{\theta}}(x)$ to $f_{\boldsymbol{\theta}}(x) - f_{\boldsymbol{\theta}^*}(x)$ by using a Taylor series expansion. In particular, for each $x \in [2^d]$ and each $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^*$, by the Lagrange form of the Taylor series expansion, there exists a $\boldsymbol{\theta}_x \in \text{conv}\{\boldsymbol{\theta}, \boldsymbol{\theta}^*\}$ such that $f_{\boldsymbol{\theta}^*}(x) - f_{\boldsymbol{\theta}}(x) = \langle \nabla f_{\boldsymbol{\theta}}(x), \boldsymbol{\theta}^* - \boldsymbol{\theta} \rangle + (\boldsymbol{\theta}^* - \boldsymbol{\theta})^T \nabla^2 f_{\boldsymbol{\theta}_x}(x) (\boldsymbol{\theta}^* - \boldsymbol{\theta})$. In addition, note from Assumption 2 that $-\mu I \preceq \nabla^2 f_{\boldsymbol{\theta}}(x) \preceq \mu I$. This results in the

following set of inequalities,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, G \rangle = \frac{\lambda}{\sqrt{2^d}} \left[\begin{array}{c} f_{\boldsymbol{\theta}}(x) - f_{\boldsymbol{\theta}^*}(x) + (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T \nabla^2 f_{\boldsymbol{\theta}_x}(x) (\boldsymbol{\theta} - \boldsymbol{\theta}^*) \\ \vdots \end{array} \right]_{x \in [2^d]} \mathbf{H}z \quad (49)$$

$$\geq \frac{\lambda}{\sqrt{2^d}} (f_{\boldsymbol{\theta}}(\mathbf{X}) - f_{\boldsymbol{\theta}^*}(\mathbf{X}))^T \mathbf{H}z - \frac{\lambda}{\sqrt{2^d}} \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \|\mathbf{H}z\|_1 \quad (50)$$

$$\geq \frac{\lambda}{\sqrt{2^d}} (\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X}) - \mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X}))^T z - (2d+1)\lambda\mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \quad (51)$$

where the last inequality follows from Lemma 2, where we show that for any vector $v : \|v\|_{\infty} \leq 1$ (a condition which is satisfied by the sign vector z), $\|\mathbf{H}v\|_1 \leq (2d+1)\sqrt{2^d}$.

A naive attempt to prove such a bound turns out to result in a loose bound. Indeed, $\|\mathbf{H}v\|_1 \leq \sqrt{2^d}\|\mathbf{H}v\|_2 = \sqrt{2^d}\|v\|_2 \leq \sqrt{2^d}\sqrt{2^d} = 2^d$. The improvement of one of the $\sqrt{2^d}$ factors to $(2d+1)$ turns to be quite a deep mathematical fact, and we invoke a result of Figula & Kvaratskhelia (2015) to prove this result in Lemma 2.

Finally, using the convexity of $\|\cdot\|_1$, for any $z \in \text{sgn}(\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X}))$,

$$(\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X}) - \mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X}))^T z = \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1 - (\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X}))^T z \quad (52)$$

$$\geq \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1 - \sup_{v: \|v\|_{\infty} \leq 1} \langle v, \mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X}) \rangle \quad (53)$$

$$\geq \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1 - \|\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 \quad (54)$$

Combining this with eq. (51) and using the definition of $R(\boldsymbol{\theta})$ completes the proof. \square

Lemma 6. Under Assumption 2, for any $\boldsymbol{\theta} \in \mathbb{R}^m$ and any subgradient $G \in (\nabla R)(\boldsymbol{\theta})$,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) + G \rangle + 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}^*}(x) - g_{\boldsymbol{\theta}}(x))] \quad (55)$$

$$\geq 2\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*) + R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}^*) - \left((2d+1)\lambda\mu + 2\mu\sqrt{\mathcal{L}_n(\boldsymbol{\theta})} \right) \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \quad (56)$$

Proof. The proof of this result builds on the analysis in Lemma 3. Observe that,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) \rangle = \mathbb{E}_{x \sim \text{Unif}(D_n)} [2(g_{\boldsymbol{\theta}}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla g_{\boldsymbol{\theta}}(x) \rangle] \quad (57)$$

For each $x \in [2^d]$ and each $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^*$, there exists a $\boldsymbol{\theta}_x \in \text{conv}(\{\boldsymbol{\theta}, \boldsymbol{\theta}^*\})$ such that $g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x) = \langle \nabla g_{\boldsymbol{\theta}}(x), \boldsymbol{\theta} - \boldsymbol{\theta}^* \rangle - (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T \nabla^2 g_{\boldsymbol{\theta}_x}(x) (\boldsymbol{\theta} - \boldsymbol{\theta}^*)$. Plugging this in, and using the assumption that $-\mu I \preceq \nabla^2 g_{\boldsymbol{\theta}}(x) \preceq \mu I$,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) \rangle - 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))] \quad (58)$$

$$= \mathbb{E}_{x \sim \text{Unif}(D_n)} [2(g_{\boldsymbol{\theta}}(x) - y(x)) \langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla g_{\boldsymbol{\theta}}(x) \rangle - 2(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))] \quad (59)$$

$$\geq \mathbb{E}_{x \sim \text{Unif}(D_n)} [2(g_{\boldsymbol{\theta}}(x) - y(x)) (g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x)) - 2(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))] \quad (60)$$

$$- 2\mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 (\mathbb{E}_{x \sim \text{Unif}(D_n)} [|g_{\boldsymbol{\theta}}(x) - y(x)|]) \quad (61)$$

$$\geq 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))^2] - 2\mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \left(\sqrt{\mathcal{L}_n(\boldsymbol{\theta})} \right) \quad (62)$$

where the last inequality follows by Jensen's inequality.

Putting together Lemma 5 and eq. (62), for any subgradient $G \in (\nabla R)(\boldsymbol{\theta})$, we have,

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \nabla \mathcal{L}_n(\boldsymbol{\theta}) + G \rangle - 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))] \quad (63)$$

$$\geq 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))^2] + \frac{\lambda}{\sqrt{2^d}} (\|\mathbf{H}g_{\boldsymbol{\theta}}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1) \quad (64)$$

$$- \left((2d+1)\lambda\mu + 2\mu\sqrt{\mathcal{L}_n(\boldsymbol{\theta})} \right) \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \quad (65)$$

This completes the proof. \square

Lemma 7. Suppose the regularization parameter satisfies the lower bound in Condition 1. Define \mathcal{E}_4 as the event that for all $\boldsymbol{\theta} \in \mathbb{R}^m$,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}^*}(x) - g_{\boldsymbol{\theta}}(x))] \leq \frac{\lambda}{2} \sqrt{\frac{1}{2^d}} \|\mathbf{H} (g_{\boldsymbol{\theta}^*}(\mathbf{X}) - g_{\boldsymbol{\theta}}(\mathbf{X}))\|_1 \quad (66)$$

Then, $\Pr(\mathcal{E}_4) \geq 1 - \delta$.

Proof. The proof of this result is similar to that of Lemma 4, and we include the details for completeness. For each sample $x' \in D_n$, define the noise in the sample as $z(x') = g_{\boldsymbol{\theta}^*}(x') - y(x')$. Likewise, define the noise vector $\mathbf{z}_{D_n} = \{\mathbb{E}_{x \sim \text{Unif}(D_n)}[z(x') \mathbb{I}(x' = x)] : x \in [2^d]\}$. Then, by an application of Holder's inequality, we have,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}^*}(x) - g_{\boldsymbol{\theta}}(x))] \leq \|\mathbf{H} \mathbf{z}_{D_n}\|_{\infty} \|\mathbf{H} (g_{\boldsymbol{\theta}^*}(\mathbf{X}) - g_{\boldsymbol{\theta}}(\mathbf{X}))\|_1. \quad (67)$$

Note that for each fixed row $i \in [2^d]$, $\langle \mathbf{H}_i, \mathbf{z}_{D_n} \rangle = \sum_{j \in [2^d]} \mathbf{H}_{ij} \mathbf{z}_{D_n}(j)$. Note that the coordinates of \mathbf{z}_{D_n} are independently distributed and subgaussian. Therefore, by Hoeffding's inequality, for some constant C_0 ,

$$\Pr \left(\langle \mathbf{H}_i, \mathbf{z}_{D_n} \rangle \geq C_0 \sqrt{\frac{\left(\sum_{j \in [2^d]} \text{Var}(\mathbf{z}_{D_n}(j)) \right) \log(1/\delta)}{2^d}} \right) \leq \delta \quad (68)$$

Note that the coordinate of \mathbf{z}_{D_n} labelled by $x \in \{\pm 1\}^d$, $\mathbb{E}_{x \sim \text{Unif}(D_n)}[z(x') \mathbb{I}(x' = x)]$, is the sum of $D_n(x)$ independent $\mathcal{N}(0, \sigma^2)$ Gaussians scaled by $1/n$, where $D_n(x)$, defined as the number of times x is sampled in D_n . Therefore, $\text{Var}(\mathbf{z}_{D_n}(i)) = \frac{\sigma^2 D_n(x)}{n^2}$. By union bounding over the 2^d rows of \mathbf{H} , with probability $\geq 1 - \delta$,

$$\Pr \left(\|\mathbf{H} \mathbf{z}_{D_n}\|_{\infty} \geq C_0 \sqrt{\frac{\left(\sum_{x \in [2^d]} \frac{\sigma^2 D_n(x)}{n^2} \right) \log(2^d/\delta)}{2^d}} \right) \leq \delta. \quad (69)$$

Note that $\sum_{x \in [2^d]} D_n(x) = n$, and therefore, with probability $\geq 1 - \delta$,

$$\|\mathbf{H} \mathbf{z}_{D_n}\|_{\infty} \leq C_0 \sigma \sqrt{\frac{1}{2^d}} \sqrt{\frac{d + \log(1/\delta)}{n}} \leq \frac{\lambda}{4\sqrt{2^d}}. \quad (70)$$

where the last inequality follows by the assumption on λ . This implies that with probability $\geq 1 - \delta$,

$$\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}^*}(x) - g_{\boldsymbol{\theta}}(x))] \leq \frac{\lambda}{2} \sqrt{\frac{1}{2^d}} \|\mathbf{H} (g_{\boldsymbol{\theta}^*}(\mathbf{X}) - g_{\boldsymbol{\theta}}(\mathbf{X}))\|_1 \quad (71)$$

□

Lemma 8. Assume that $d \geq 3$ and $k \leq 2^d/4$. Then, there exists a set of $2^{(d-1)\lfloor k/2 \rfloor}$ binary vectors of length 2^d , denoted \mathcal{C}_k , each having at most k ones, such that: the hamming distance between any pair of vectors is at least $\lfloor \frac{k}{2} \rfloor$.

Proof. The number of vectors within hamming distance $\kappa = \lfloor k/2 \rfloor$ of any vector is $\binom{2^d}{\kappa}$. Note that $k - \kappa \geq \kappa$.

We can greedily construct a packing of size at least $\binom{2^d}{k} / \binom{2^d}{\kappa} \geq \frac{(2^d - \kappa)!}{(2^d - k)!} \geq 2^{d(k-\kappa)} \left(\frac{3}{4}\right)^{k-\kappa} \geq 2^{d\kappa} \left(\frac{3}{4}\right)^{\kappa} \geq 2^{(d-1)\kappa}$ vectors before running out of binary vectors to choose. By construction, every pair of vectors has Hamming distance at least κ . □

C Proof of Theorem 1 - Statistical performance under RSI

In this section we discuss the proof of Theorem 1. We prove a slightly more general result which characterizes the performance of stationary points of the MSE with spectral regularization under the RSI, when the regularization parameter is chosen arbitrarily.

Theorem 4. *Suppose the regularization parameter λ satisfies Condition 1. Namely, for a sufficiently large absolute constant $C_0 > 0$,*

$$\lambda \geq 4C_0\sigma\sqrt{\frac{d + \log(1/\delta)}{n}} \quad (72)$$

Consider a learner which returns any first order stationary point of the loss $\mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$. Under Assumptions 1(a), 2 and 3, if $\frac{C_{n,\delta}^}{2} - \frac{5}{4}(2d+1)\lambda\mu - 3\lambda L\sqrt{k} > 0$, with probability $\geq 1 - 2\delta$,*

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{6\lambda L\sqrt{k}}{C_{n,\delta}^*} \quad (73)$$

The subgradient of the regularizer $R(\boldsymbol{\theta}) = \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1$ is characterized below.

Proposition 1. *The subgradient of the regularization $R(\boldsymbol{\theta}) = \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1$ is,*

$$(\nabla R)(\boldsymbol{\theta}) = \left\{ \frac{\lambda}{\sqrt{2^d}} (\nabla f_{\boldsymbol{\theta}}(\mathbf{X})) \mathbf{H}z : z \in \text{sgn}(\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})) \right\} \quad (74)$$

Here, for $z \in \mathbb{R}$, $\text{sgn}(z) = \begin{cases} \{+1\} & \text{if } z > 0 \\ \{-1\} & \text{if } z < 0 \\ [-1, +1] & \text{otherwise} \end{cases}$ and applied on a vector $z = (z_1, \dots, z_n)$ is the Cartesian product of the sets $\text{sgn}(z_1) \times \dots \times \text{sgn}(z_n)$.

We are now ready to state the proof.

Proof. Plugging in $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$ into Lemma 3, and noting that $0 \in \nabla \mathcal{L}_n(\hat{\boldsymbol{\theta}}) + (\nabla R)(\hat{\boldsymbol{\theta}})$, choosing G appropriately we obtain,

$$\begin{aligned} & 2\mathbb{E}_{x \sim \text{Unif}(D_n)} \left[(f_{\boldsymbol{\theta}^*}(x) - y(x)) \langle \boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}, \nabla f_{\hat{\boldsymbol{\theta}}}(x) \rangle \right] \\ & \geq R(\hat{\boldsymbol{\theta}}) - R(\boldsymbol{\theta}^*) + (C_{n,\delta}^* - (2d+1)\lambda\mu) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \end{aligned} \quad (75)$$

Now, we upper bound the LHS of Lemma 3 using Lemma 4. Plugging Lemma 4 into eq. (75) with the choice of $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, and rearranging both sides, under the events \mathcal{E}_1 and \mathcal{E}_2 which jointly occur with probability $\geq 1 - 2\delta$,

$$\begin{aligned} & \left(C_{n,\delta}^* - \frac{3}{2}(2d+1)\lambda\mu \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \\ & \stackrel{(i)}{\leq} \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}f_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 + \frac{\lambda}{2\sqrt{2^d}} \|\mathbf{H}(f_{\boldsymbol{\theta}^*}(\mathbf{X}) - f_{\hat{\boldsymbol{\theta}}}(\mathbf{X}))\|_1 \end{aligned} \quad (76)$$

$$\stackrel{(ii)}{\leq} \frac{\lambda}{\sqrt{2^d}} \left(\|\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}f_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 + \frac{1}{2} \|\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 + \frac{1}{2} \|\mathbf{H}f_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 \right) \quad (77)$$

$$= \frac{\lambda}{2\sqrt{2^d}} (3 \|\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}f_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1) \quad (78)$$

where (i) follows from the definition of the regularization term, $R(\boldsymbol{\theta}) = \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}f_{\boldsymbol{\theta}}(\mathbf{X})\|_1$. On the other hand, (ii) follows by triangle inequality of the norm $\|\cdot\|_1$. By the assumption $\frac{1}{2}C_{n,\delta}^* \geq \frac{3}{2}(2d+1)\lambda\mu + 3\lambda L\sqrt{k}$, the LHS is non-negative under the event \mathcal{E}_1 . Plugging this into eq. (78) results in the inequality,

$$\frac{\lambda}{2\sqrt{2^d}} (3 \|\mathbf{H}f_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}f_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1) \geq 0. \quad (79)$$

Conditioned on the event \mathcal{E}_1 .

Next, we apply (Loh & Wainwright, 2015, Lemma 5) to the function $\rho_\lambda(\cdot) = \|\cdot\|_1$, and note that by assumption $\mathbf{H}f_{\theta^*}(\mathbf{X})$ is k -sparse. Define $\nu = \mathbf{H}(f_{\theta^*}(\mathbf{X}) - f_{\hat{\theta}}(\mathbf{X}))$ and A as the set of k largest indices of ν in absolute value. Using the fact that $3\|\mathbf{H}f_{\theta^*}(\mathbf{X})\|_1 - \|\mathbf{H}f_{\hat{\theta}}(\mathbf{X})\|_1 \geq 0$,

$$\frac{\lambda}{2\sqrt{2d}} (3\|\mathbf{H}f_{\theta^*}(\mathbf{X})\|_1 - \|\mathbf{H}f_{\hat{\theta}}(\mathbf{X})\|_1) \stackrel{(i)}{\leq} \frac{\lambda}{\sqrt{2d}} (3\|\nu_A\|_1 - \|\nu_{A^c}\|_1) \quad (80)$$

$$\leq \frac{3\lambda\sqrt{k}}{2\sqrt{2d}} \|\nu_A\|_2 \quad (81)$$

$$\leq \frac{3\lambda\sqrt{k}}{2\sqrt{2d}} \|\nu\|_2. \quad (82)$$

Plugging this back into eq. (78) results in the following inequality, conditioned on the events \mathcal{E}_1 and \mathcal{E}_2 ,

$$\left(C_{n,\delta}^* - \frac{3}{2}(2d+1)\lambda\mu\right) \|\hat{\theta} - \theta^*\|_2^2 \leq \frac{3\lambda\sqrt{k}}{2\sqrt{2d}} \|f_{\theta^*}(\mathbf{X}) - f_{\hat{\theta}}(\mathbf{X})\|_2 \quad (83)$$

Finally, we relate $\|f_{\theta^*}(\mathbf{X}) - f_{\hat{\theta}}(\mathbf{X})\|_2$ to $\|\hat{\theta} - \theta^*\|_2$.

Plugging the bound in Lemma 1 into eq. (83) results in the following inequality, conditioned on \mathcal{E}_1 and \mathcal{E}_2 ,

$$\left(C_{n,\delta}^* - \frac{3}{2}(2d+1)\lambda\mu\right) \|\hat{\theta} - \theta^*\|_2^2 \leq 3\mu\lambda\sqrt{k} \|\hat{\theta} - \theta^*\|_2^2 + 3\lambda L\sqrt{k} \|\hat{\theta} - \theta^*\|_2. \quad (84)$$

Furthermore, recalling the assumption that $\frac{C_{n,\delta}^*}{2} - \frac{3}{2}(2d+1)\lambda\mu - 3\mu\lambda\sqrt{k} \geq 0$, this results in the overall bound,

$$\|\hat{\theta} - \theta^*\|_2 \leq \frac{3\lambda L\sqrt{k}}{C_{n,\delta}^* - \frac{3}{2}(2d+1)\lambda\mu - 3\mu\lambda\sqrt{k}} \leq \frac{6\lambda L\sqrt{k}}{C_{n,\delta}^*}, \quad (85)$$

under the events \mathcal{E}_1 and \mathcal{E}_2 which jointly occur with probability $\geq 1 - 2\delta$. □

Proof of Theorem 1. Theorem 1 follows from Theorem 4 by choosing λ as its lower bound in Condition 1, equal to $4C_0\sigma\sqrt{\frac{d+\log(1/\delta)}{n}}$. It is easily verified that when the size of the dataset, n , is larger than the quantity n_0 as defined in the statement of Theorem 1, the condition $\frac{C_{n,\delta}^*}{2} - \frac{5}{4}(2d+1)\lambda\mu - 3\lambda L\sqrt{k} > 0$ as required in Theorem 4 is satisfied.

Notes on the constants. We chose to leave the constants C_0 and C_1 implicit in the proof to keep the theorem statements and proofs simple. A crude bound shows that C_0 and C_1 can be chosen as 10, and with more careful analysis it can be brought down further.

D Proof of Theorem 2 - Statistical performance under QG condition

In this section we discuss the proof of Theorem 2. This result is a consequence of a more general result which characterizes the performance of stationary points of the MSE with spectral regularization when the regularization parameter, λ , is chosen arbitrarily.

Theorem 5. Define $\Delta = \frac{C_{n,\delta}^*}{2\mu} - \frac{(2d+1)\lambda}{2} - \frac{3}{2}\lambda\sqrt{k}$ and assume $\Delta > 0$. Suppose the regularization parameter λ satisfies Condition 1. Namely, for a sufficiently absolute large constant $C_0 > 0$,

$$\lambda \geq 4C_0\sigma\sqrt{\frac{d+\log(1/\delta)}{n}} \quad (86)$$

Consider a learner which returns a Δ^2 -approximate first order stationary interpolator (Definition 3) of the loss $\mathcal{L}_n(\boldsymbol{\theta}) + R(\boldsymbol{\theta})$. Then, under Assumptions 1(b), 2 and 3, with probability $\geq 1 - 2\delta$,

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{3\lambda L\sqrt{k}}{C_{n,\delta}^*}. \quad (87)$$

Proof. Plugging in $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$ into Lemma 6, and noting that the learner returns a first order stationary point of the regularized loss, $0 \in \nabla \mathcal{L}_n(\hat{\boldsymbol{\theta}}) + (\nabla R)(\hat{\boldsymbol{\theta}})$, choosing G appropriately we have,

$$2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}^*}(x) - g_{\hat{\boldsymbol{\theta}}}(x))] \quad (88)$$

$$\geq 2\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) + R(\hat{\boldsymbol{\theta}}) - R(\boldsymbol{\theta}^*) - \left((2d+1)\lambda\mu + 2\mu\sqrt{\mathcal{L}_n(\hat{\boldsymbol{\theta}})} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \quad (89)$$

Under assumption (A1), recall that we assume that $\hat{\boldsymbol{\theta}}$ is a sufficiently good interpolator in that, $\mathcal{L}_n(\hat{\boldsymbol{\theta}}) \leq \left(\frac{C_{n,\delta}^*}{2\mu} - \frac{(2d+1)\lambda}{2} - \frac{3}{2}\sqrt{k}\lambda \right)^2$. Simplifying eq. (89) under this assumption gives,

$$\begin{aligned} & 2\mathbb{E}_{x \sim \text{Unif}(D_n)} [(g_{\boldsymbol{\theta}^*}(x) - y(x)) (g_{\boldsymbol{\theta}^*}(x) - g_{\hat{\boldsymbol{\theta}}}(x))] \\ & \geq 2\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) + R(\hat{\boldsymbol{\theta}}) - R(\boldsymbol{\theta}^*) - \left(C_{n,\delta}^* - 3\mu\lambda\sqrt{k} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \end{aligned} \quad (90)$$

Next we bound the LHS of eq. (90) using Lemma 7.

Plugging Lemma 7 into (90) and rearranging both sides, under the event \mathcal{E}_4 ,

$$\begin{aligned} & 2\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) - \left(C_{n,\delta}^* - 3\lambda\mu\sqrt{k} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \\ & \stackrel{(i)}{\leq} \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 + \frac{\lambda}{2\sqrt{2^d}} \|\mathbf{H}(g_{\boldsymbol{\theta}^*}(\mathbf{X}) - g_{\hat{\boldsymbol{\theta}}}(\mathbf{X}))\|_1 \end{aligned} \quad (91)$$

$$\stackrel{(ii)}{\leq} \frac{\lambda}{\sqrt{2^d}} \left(\|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 + \frac{1}{2} \|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 + \frac{1}{2} \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 \right) \quad (92)$$

$$= \frac{\lambda}{2\sqrt{2^d}} (3 \|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1) \quad (93)$$

where (i) follows from the definition of the regularization term, $R(\boldsymbol{\theta}) = \frac{\lambda}{\sqrt{2^d}} \|\mathbf{H}g_{\boldsymbol{\theta}}(\mathbf{X})\|_1$. On the other hand, (ii) follows by triangle inequality of the norm $\|\cdot\|_1$. Next we focus on the LHS of the above expression and simplify it further. By the quadratic growth condition in assumption 1(b), under the event \mathcal{E}_3 , which is assumed to happen with probability $\geq 1 - \delta$, $\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) \geq C_{n,\delta}^* \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2$. Therefore, under \mathcal{E}_3 and \mathcal{E}_4 ,

$$2\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) - \left(C_{n,\delta}^* - 3\lambda\mu\sqrt{k} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \quad (94)$$

$$\geq \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \left(C_{n,\delta}^* + 3\lambda\mu\sqrt{k} \right) \quad (95)$$

$$\geq 0. \quad (96)$$

Plugging this into eq. (93) results in the inequality,

$$\frac{\lambda}{2\sqrt{2^d}} (3 \|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1) \geq 0. \quad (97)$$

Under the events \mathcal{E}_3 and \mathcal{E}_4 .

Next, we apply (Loh & Wainwright, 2015, Lemma 5) to the function $\rho_\lambda(\cdot) = \|\cdot\|_1$, and note that by assumption $\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})$ is k -sparse. By defining $\nu = \mathbf{H}(g_{\boldsymbol{\theta}^*}(\mathbf{X}) - g_{\hat{\boldsymbol{\theta}}}(\mathbf{X}))$ and A as the set of k largest indices

of ν in absolute value.

$$\frac{\lambda}{2\sqrt{2^d}} (3\|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1) \stackrel{(i)}{\leq} \frac{\lambda}{\sqrt{2^d}} (3\|\nu_A\|_1 - \|\nu_{A^c}\|_1) \quad (98)$$

$$\leq \frac{3\lambda\sqrt{k}}{2\sqrt{2^d}} \|\nu_A\|_2 \quad (99)$$

$$\leq \frac{3\lambda\sqrt{k}}{2\sqrt{2^d}} \|\nu\|_2. \quad (100)$$

here, (i) uses the fact that $3\|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1 \geq 0$ from eq. (97).

Finally, we plug the relation between $\|v\|_2 = \|g_{\boldsymbol{\theta}^*}(\mathbf{X}) - g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_2$ to $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2$ proved in Lemma 1 into eq. (100). This results in the inequality,

$$\frac{\lambda}{2\sqrt{2^d}} (3\|\mathbf{H}g_{\boldsymbol{\theta}^*}(\mathbf{X})\|_1 - \|\mathbf{H}g_{\hat{\boldsymbol{\theta}}}(\mathbf{X})\|_1) \leq 3\mu\lambda\sqrt{k}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 + 3\lambda L\sqrt{k}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \quad (101)$$

Plugging this back into eq. (93), under \mathcal{E}_3 and \mathcal{E}_4 ,

$$2\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) - \left(C_{n,\delta}^* - 3\lambda\mu\sqrt{k}\right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 \leq 3\mu\lambda\sqrt{k}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2 + 3\lambda L\sqrt{k}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2. \quad (102)$$

Resulting in the bound,

$$2\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) \leq 3\lambda L\sqrt{k}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 + C_{n,\delta}^*\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2. \quad (103)$$

Under the quadratic growth condition, by the event \mathcal{E}_3 in Assumption 1(b), $\text{Err}_n(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}^*) \geq C_{n,\delta}^*\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2^2$. Therefore, under the events \mathcal{E}_3 and \mathcal{E}_4 which jointly occur with probability $\geq 1 - 2\delta$,

$$\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{3\lambda L\sqrt{k}}{C_{n,\delta}^*}. \quad (104)$$

Proof of Theorem 2. Theorem 2 follows from Theorem 5 by choosing λ as its lower bound in Condition 1, equal to $4C_0\sigma\sqrt{\frac{d+\log(1/\delta)}{n}}$. It is easily verified that when $n > n_0$, as defined in the statement of Theorem 2, the condition $\Delta = \frac{C_{n,\delta}^*}{2} - \frac{5}{4}(2d+1)\lambda\mu - 3\lambda L\sqrt{k} > 0$ as required in Theorem 5 is satisfied. \square

E Proof of Theorem 3

In this section, we prove a lower bound on the statistical error of parameter estimation. Loosely speaking, the objective is to show that for every learner $\hat{\boldsymbol{\theta}}$,

$$\sup_{\boldsymbol{\theta}^*} \mathbb{E} \left[\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \right] \gtrsim \sigma \sqrt{\frac{kd}{n}}. \quad (105)$$

Proof. We first introduce an auxiliary result related to packing binary vectors with bounded Hamming weight in Lemma 8. Now define the function space $\mathcal{G} = \{g_{\boldsymbol{\theta}}(\cdot) : \boldsymbol{\theta} \in \mathbb{R}^{2^d}\}$, where $g_{\boldsymbol{\theta}}(x)$ is defined as the polynomial with coefficients specified by $\boldsymbol{\theta}$. Namely, $g_{\boldsymbol{\theta}}(x) = \sum_{S \subseteq [d]} \theta_S \prod_{i \in S} x_i$, where we index the 2^d coefficients of $\boldsymbol{\theta}$ by the 2^d subsets of $[d]$. In an alternate notation, we may represent,

$$g_{\boldsymbol{\theta}}(x) = \langle \boldsymbol{\theta}, 2^{[x]} \rangle \quad (106)$$

where $2^{[x]}$ denotes the 2^d length vector whose element indexed by some subset $S \subseteq [d]$ is $\prod_{i \in S} x_i$.

Furthermore, we assume that the data generating distribution independently samples n pairs (x_i, y_i) where $x_i \sim \text{Unif}(\{\pm 1\}^d)$ and $y_i = g_{\boldsymbol{\theta}^*}(x_i) + Z_i$ where $Z_i \sim \mathcal{N}(0, \sigma^2)$. Denote $D_n = \{x_1, \dots, x_n\}$.

By Lemma 8, the binary vectors belonging to \mathcal{C}_k can be used to construct a subset of the function space \mathcal{G} , defined as \mathcal{S}_k ,

$$\mathcal{S}_k = \{g_{\boldsymbol{\theta}}(\cdot) : \boldsymbol{\theta} \in \Delta\mathcal{C}_k\}. \quad (107)$$

where $\Delta > 0$ is a scaling factor and $\Delta\mathcal{C}_k = \{\Delta\boldsymbol{\theta} : \boldsymbol{\theta} \in \mathcal{C}_k\}$.

Henceforth, we will consider ourselves with learning functions (resp. parameters) in the class \mathcal{S}_k (resp. $\Delta\mathcal{C}_k$).

First we show the properties on $\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*)$ and \mathcal{G} in the statement of Theorem 3. Note that \mathcal{G} is a linear family by the representation in eq. (106), and therefore $\mu = 0$.

In addition, note that,

$$\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*) = \mathbb{E}_{x \sim D_n} [(g_{\boldsymbol{\theta}}(x) - g_{\boldsymbol{\theta}^*}(x))^2] \quad (108)$$

$$= \mathbb{E}_{x \sim D_n} \left[\left\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, 2^{[x]} \right\rangle^2 \right] \quad (109)$$

$$= (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T \mathbb{E}_{x \sim D_n} \left[2^{[x]} (2^{[x]})^T \right] (\boldsymbol{\theta} - \boldsymbol{\theta}^*) \quad (110)$$

Note that $A_x = 2^{[x]}(2^{[x]})^T$ is a matrix whose entries (indexed by pairs of subsets of $[d]$) can be described as $A_x(S, T) = \prod_{i \in S} x_i \prod_{j \in T} x_j$. Note that in expectation over $x \sim \text{Unif}(\mathbf{X})$, we have that,

$$\mathbb{E}_{x \sim \text{Unif}(\mathbf{X})} [A_x(S, T)] = \mathbb{I}(S = T) \quad (111)$$

This is because if $S \neq T$, there exists an element $i \in (S \setminus T) \cup (T \setminus S)$ (i.e. the symmetric difference of the two sets) and since $\mathbb{E}_{x \sim \text{Unif}(\mathbf{X})}[x_i] = 0$, we get the required statement. Therefore,

$$\mathbb{E}_{x \sim \text{Unif}(\mathbf{X})}[A_x] = I \quad (112)$$

Now, given n samples from the uniform distribution, $\mathbb{E}_{x \sim \text{Unif}(D_n)}[A_x]$ is expected to concentrate around its expectation $\mathbb{E}_{x \sim \text{Unif}(\mathbf{X})}[A_x]$. In particular, by invoking the matrix Bernstein inequality Tropp et al. (2015), we have that,

$$\Pr \left(\left\| \mathbb{E}_{x \sim \text{Unif}(D_n)}[A_x] - \mathbb{E}_{x \sim \text{Unif}(\mathbf{X})}[A_x] \right\|_{\text{op}} \geq t \right) \leq 2(2^d) \exp(-nt^2/2L^2) \quad (113)$$

where L is an almost sure upper bound on $\|A_x\|_{\text{op}}$. Note that $\|A_x\|_{\text{op}} = \|2^{[x]}\|_2 = \sqrt{2^d}$ and therefore we may choose $L = \sqrt{2^d}$. This results in the bound,

$$\Pr \left(\left\| \mathbb{E}_{x \sim \text{Unif}(D_n)}[A_x] - I \right\|_{\text{op}} \geq \frac{1}{2} \right) \leq 2(2^d) \exp(-n/2^{d+3}) \quad (114)$$

Therefore, if $n \gtrsim (d + \log(1/\delta))2^{d+3}$, with probability $\geq 1 - \delta$,

$$\frac{1}{2}I \preceq \mathbb{E}_{x \sim \text{Unif}(D_n)}[A_x] \preceq \frac{3}{2}I. \quad (115)$$

In eq. (110), this implies that, with probability $\geq 1 - \delta$,

$$\text{Err}_n(\boldsymbol{\theta}, \boldsymbol{\theta}^*) \geq \frac{1}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2^2 \quad (116)$$

And likewise, from the linear representation in eq. (106),

$$\mathbb{E}_{x \sim \text{Unif}(\mathbf{X})} [\nabla g_{\boldsymbol{\theta}^*}(x)(\nabla g_{\boldsymbol{\theta}^*}(x))^T] = \mathbb{E}_{x \sim \text{Unif}(\mathbf{X})} [2^{[x]}(2^{[x]})^T] = I \quad (117)$$

where the last equation follows from eq. (112).

To generate the lower bound instance, suppose the ground truth parameter θ^* is sampled uniformly from $\Delta\mathcal{C}_k$. Suppose the learner outputs a candidate parameter $\hat{\theta}$. The population level mean squared error of the learner is lower bounded by the testing error,

$$\sup_{\theta^* \in \mathbb{R}^{2d}} \mathbb{E} \left[\|\hat{\theta} - \theta^*\|_2 \right] \geq \mathbb{E}_{\theta^* \sim \text{Unif}(\Delta\mathcal{C}_k)} \left[\mathbb{E} \left[\|\hat{\theta} - \theta^*\|_2 \right] \right] \quad (118)$$

$$\geq \frac{1}{4} \sqrt{[k/2]} \Delta \inf_{\Psi} \mathbb{E}_{\theta^* \sim \text{Unif}(\mathcal{C}_k)} [\mathbb{I}(\Psi(D_n)) \neq \theta^*] \quad (119)$$

where the infimum is over all tests functions which return a function in \mathcal{C}_k . This uses the fact that any estimator $\hat{\theta}$ induces a testing function for θ^* by returning the $\theta \in \Delta\mathcal{C}_k$ such that $\|\theta - \hat{\theta}\|_2$ is smallest. If this test makes a mistake, then $\hat{\theta}$ must have predicted $> \frac{1}{2} \sqrt{[k/2]}$ entries of θ^* as $< \Delta/2$ instead of Δ or $> \Delta/2$ instead of 0 (making an error of at least $\Delta/2$ on these coordinates).

The optimal hypothesis testing error can be lower bounded by Fano's inequality as follows,

$$\inf_{\Psi} \mathbb{E}_{\theta^* \sim \text{Unif}(\Delta\mathcal{C}_k)} [\mathbb{I}(\Psi(D_n)) \neq \theta^*] \geq 1 - \frac{I(\theta^*; D_n) + \log(2)}{\log |\Delta\mathcal{C}_k|} \quad (120)$$

For each parameter $\theta \in \Delta\mathcal{C}_k$, define the distribution,

$$P_{\theta}(D_n) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} \exp \left(-\frac{1}{2\sigma^2} (y_i - g_{\theta}(x_i))^2 \right) \times \left(\frac{1}{2^d} \right). \quad (121)$$

which captures the data distribution when the underlying ground truth parameter is θ . Note that $D_n \sim P_{\theta^*}$; marginally x_i is uniformly distributed on the hypercube and conditionally y_i is normally distributed with mean $g_{\theta^*}(x_i)$ and variance σ^2 . Note that the mutual information can be bounded as,

$$I(\theta^*, D_n) = \inf_Q \mathbb{E}_{\theta^* \sim \text{Unif}(\Delta\mathcal{C}_k)} [\text{KL}(P_{\theta^*} \| Q)] \quad (122)$$

$$\leq \mathbb{E}_{\theta^* \sim \text{Unif}(\Delta\mathcal{C}_k)} [\text{KL}(P_{\theta^*} \| P_0)] \quad (123)$$

where P_0 is the distribution of D_n when the ground truth function g_{θ} in eq. (121) is chosen as 0 everywhere. Then,

$$\text{KL}(P_{\theta^*} \| P_0) = \frac{1}{2\sigma^2} \mathbb{E}_{D_n \sim P_{\theta^*}} \left[\sum_{i=1}^n 2y_i g_{\theta^*}(x_i) - (g_{\theta^*}(x_i))^2 \right] \quad (124)$$

$$= \frac{1}{2\sigma^2} \mathbb{E}_{D_n \sim P_{\theta^*}} \left[\sum_{i=1}^n (g_{\theta^*}(x_i))^2 \right] \quad (125)$$

$$= \frac{n}{2\sigma^2} \mathbb{E}_{x_i \sim \text{Unif}(\{\pm 1\}^d)} [(g_{\theta^*}(x_i))^2] \quad (126)$$

where the last equation just uses the fact that in P_{θ^*} , the marginal distribution of x_i is uniform. For each θ^* , by Parseval's theorem, $\mathbb{E}_{x_i \sim \text{Unif}(\{\pm 1\}^d)} [(g_{\theta^*}(x_i))^2] = k\Delta^2$. Overall, plugging into eq. (126) and subsequently into eq. (123),

$$I(\theta^*, D_n) \leq \mathbb{E}_{\theta^* \sim \text{Unif}(\Delta\mathcal{C}_k)} [\text{KL}(P_{\theta^*} \| P_0)] = \frac{nk\Delta^2}{2\sigma^2} \quad (127)$$

Putting these bounds into eq. (120) and subsequently into eq. (119) results in,

$$\mathbb{E}_{\theta^* \sim \text{Unif}(\Delta\mathcal{C}_k)} \left[\mathbb{E} \left[\|\hat{\theta} - \theta^*\|_2 \right] \right] \geq \frac{1}{4} \sqrt{[k/2]} \Delta \left(1 - \frac{nk\Delta^2/2\sigma^2 + \log(2)}{\log 2^{(d-1)[k/2]}} \right) \quad (128)$$

$$\geq \frac{1}{4} \sqrt{[k/2]} \Delta \left(1 - 10 \frac{(nk\Delta^2/\sigma^2 + 1)}{(d-1)[k/2]} \right) \quad (129)$$

Choosing $\Delta = 8\epsilon/\sqrt{\lfloor k/2 \rfloor}$, for a sufficiently large constant C , we get that for any learner, if $n \leq C \frac{\sigma^2 d}{\Delta^2} \asymp \frac{\sigma^2 k d}{\epsilon^2}$,

$$\mathbb{E}_{\boldsymbol{\theta}^* \sim \text{Unif}(\Delta \mathcal{C}_k)} \left[\mathbb{E} \left[\left\| \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^* \right\|_2 \right] \right] \geq \epsilon \asymp \sigma \sqrt{\frac{kd}{n}}. \quad (130)$$

□