Fundamental Limits of Learning Single-Index Models under Structured Data

author names withheld

Under Review for the Workshop on High-dimensional Learning Dynamics, 2025

Abstract

Recent works have developed a comprehensive picture of gradient-based learning of isotropic Gaussian single-index models by developing computational lower bounds along with optimal algorithms. In this work, we demonstrate that the picture can change significantly when the data covariance is structured and contains some degree of information about the target. Through studying a spiked covariance model, we show that for the class of Correlational Statistical Query (CSQ) learners, a simple preconditioning of online SGD already achieves an almost optimal sample complexity. Unlike the isotropic case, further smoothing the landscape does not improve this complexity. We prove similar lower bounds in the Statistical Query (SQ) class, where we demonstrate a gap between the SQ lower bound and the performance of the algorithms that are known to be optimal in the isotropic setting. Finally, we show a stark contrast in the information-theoretic limit, where the tight lower bound goes through a sudden phase transition from d to 1 depending on covariance structure, where d is the dimension of the input. Overall, our analysis provides a clear characterization of when and how the spike simplifies learning by improving over isotropic covariance.

1. Introduction

The single-index model with isotropic input data has been thoroughly studied in recent years [1, 2, 5], but it has remained relatively underexplored in the context of structured input data. In practice, data contain structure that simplifies learning [13, 15]. In this work, we consider the problem of learning a single-index model in high dimensions under structured input data:

$$y = f^{\star} \left(\frac{\langle \boldsymbol{u}^{\star}, \boldsymbol{x} \rangle}{\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{u}^{\star}\|} \right) + \xi, \quad \boldsymbol{x} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}), \quad \xi \sim \mathcal{N}(\boldsymbol{0}, \sigma^2).$$
(1.1)

where $f^* : \mathbb{R} \to \mathbb{R}$ is the link function, $x \in \mathbb{R}^d$ are the inputs with covariance Σ and $\xi \in \mathbb{R}$ corresponds to label noise with finite variance $\sigma^2 > 0$. Learning the model consists in estimating the target to achieve small population loss. In particular, it can be achieved by approximately recovering the unknown link function and the hidden dimension u^* . To avoid identifiability (scaling) issues, we standardize the projected inputs by $\|\Sigma^{1/2}u^*\|$. To see concretely how covariance structure affects sample complexity, we focus on the *spiked single-index model* that was studied in [14]:

$$\boldsymbol{\Sigma} = \frac{1}{1+\kappa} (\mathbf{I}_d + \kappa \boldsymbol{\theta} \boldsymbol{\theta}^\top), \quad \kappa \asymp d^{r_2}, \quad \langle \boldsymbol{u}^\star, \boldsymbol{\theta} \rangle \asymp d^{-r_1}, \quad r_1 \in [0, 1/2], \quad r_2 \in [0, 1].$$
(1.2)

Here, $\kappa \simeq d^{r_2}$ captures the magnitude of the spike, and $\langle u^*, \theta \rangle \simeq d^{-r_1}$ the alignment of the signal direction with the spike. Our main theorems show that it is $r \coloneqq \max\{r_1 - r_2/2, 0\}$ that accurately captures their interplay.

1.1. Contributions

We prove fundamental limits for learning the spiked single-index model with different classes of algorithms. More precisely:

• Computational lower bounds. We prove two Statistical Query (SQ) lower bounds. Notably, the sample complexity of online SGD with squared loss is captured by the Correlational Statistical Query (CSQ) framework. Additionally, online SGD coupled with "label transformations" falls within the general SQ framework. Our lower bounds provide an extension to the recent treatment of the isotropic single-index model in [5]. More precisely, for l^* , k^* —the information exponent, respectively the generative exponent-of the single-index model, the lower bounds depend Figure 1: Sample complexity to learn the on r as soon as r < 1/4: we get the informal spiked SIM. Smaller r_1 denotes a better sample complexity from the CSQ lower bound spike-target alignment, while larger r_2 de $n = \Omega(d^{2rl^{\star}})$ and the SQ lower bound based on notes a larger spike magnitude. VSTAT $n = \Omega(d^{2rk^*})$. When r > 1/4, our lower



bounds coincide with the isotropic ones [4][5]: for CSQ, $n = \Omega(d^{l^*/2})$ and for SQ, n = $\Omega(d^{k^{\star}/2})$. This suggests that the spike is helpful if and only if r < 1/4.

- Information-theoretic phase transition. We also prove an information-theoretic lower bound, which is the same for the isotropic and spiked settings as long as r > 0. Namely, $n = \Omega(d)$. Although this bound was already informally known in the isotropic setting, we believe that it has not been proved before for the single-index model. Note that it matches the informationtheoretic upper bound in [5]. Moreover, for r = 0, and $r_1 < r_2/2$, the sample complexity undergoes a stark phase transition as $n = \Omega(1)$ samples are sufficient in this regime.
- Near-matching upper bounds. Finally, we provide upper bounds that nearly match our statistical query lower bounds by showing how initializing existing algorithms such as online SGD ([2]) and tensor power iteration ([5]) at the covariance spike reduces their sample complexity. These algorithms require $n = O(d^{2rl^{\star}+1-4r})$ samples. In addition, by coupling them with a SQ algorithm that transform labels to reduce the information to the generative exponent as in [5], we get the SQ upper bound $n = \mathcal{O}(d^{2rk^*+1-4r})$. Even though they are not strictly speaking SQ algorithms, their sample complexity captures the same hardness as the lower bounds.

2. Preliminaries

2.1. Statistical queries

Given a regression problem where we have a joint distribution $\mathbb{P}_{X,Y}$ on $(X,Y) \in \mathbb{R}^d \times \mathbb{R}$, the goal is to predict Y given X by choosing a predictor $h : \mathbb{R}^d \to \mathbb{R}$ among a class of predictors \mathcal{H} . The statistical query framework [12] contains the following class of algorithms. At each iteration, the algorithm can (possibly adaptively) choose a query $q : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$ from a class of queries \mathcal{Q} , and the oracle returns a response $\hat{\mathbb{E}}[q]$ with the following guarantee

$$\left| \hat{\mathbb{E}}[q] - \mathbb{E}_{\boldsymbol{X},Y}[q(\boldsymbol{X},Y)] \right| \le \tau,$$

where τ is the noise tolerance of the oracle. In particular, note that the tolerance can be chosen adversarially by the oracle. One widely considered restriction of the class of queries is known as Correlational Statistical Queries (CSQ) [3], where given some class of real-valued functions on \mathbb{R}^d denoted by \tilde{Q} , the class Q is given by

$$\mathcal{Q} = \{q(\boldsymbol{x}, y) = \tilde{q}(\boldsymbol{x})y : \tilde{q} \in \mathcal{Q}\}.$$

Example 1. The online SGD update with respect to the squared error loss fits the CSQ framework, apart from the noise not being adversarial. Indeed, for an estimator $f_w, w \in \mathcal{W}$ (for example a neural network), the gradient of the population risk admits the decomposition

$$\nabla \mathcal{R}(\boldsymbol{w}) = \frac{1}{2} \mathbb{E}[\nabla (Y - f_{\boldsymbol{w}}(\boldsymbol{X})^2] = \mathbb{E}[Y \nabla f_{\boldsymbol{w}}(\boldsymbol{X})] - \mathbb{E}[f_{\boldsymbol{w}}(\boldsymbol{X}) \nabla f_{\boldsymbol{w}}(\boldsymbol{X})],$$

from which we see that the target and the inputs only interact through a correlational query (the first term on the right-hand side).

2.2. Reformulation as an isotropic SIM

Lemma 2 (Single-index model translation). *The spiked single-index model defined in Equations* (1.1), (1.2) *is equivalent to the isotropic single-index model*

$$y = f^{\star}(\langle \overline{\boldsymbol{u}}^{\star}, \boldsymbol{z} \rangle) + \xi, \quad \boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad \boldsymbol{u}^{\star} \in \mathbb{S}^{d-1} \quad subject \ to \quad |\langle \boldsymbol{\theta}, \overline{\boldsymbol{u}}^{\star} \rangle| = cd^{-r},$$

for some absolute constant c > 0. Furthermore, this imposes a restriction on the unknown direction:

$$\overline{\boldsymbol{u}}^{\star} \in S_{\boldsymbol{\theta},r} \coloneqq \{ \overline{\boldsymbol{u}} \in \mathbb{S}^{d-1} : |\langle \overline{\boldsymbol{u}}, \boldsymbol{\theta} \rangle| = cd^{-r} \}.$$

Proof. The spiked single-index model can be reformulated as an isotropic single-index model by whitening the data. For this purpose, let

$$\overline{oldsymbol{u}}^\star\coloneqq rac{oldsymbol{\Sigma}^{1/2}oldsymbol{u}^\star}{\|oldsymbol{\Sigma}^{1/2}oldsymbol{u}^\star\|}, \hspace{1em} extsf{which yields} \hspace{1em} rac{\langleoldsymbol{u}^\star,oldsymbol{x}
angle}{\|oldsymbol{\Sigma}^{1/2}oldsymbol{u}^\star\|}=\langle\overline{oldsymbol{u}}^\star,oldsymbol{z}
angle,$$

where z is a standard Gaussian random vector in \mathbb{R}^d . Then, we can relate the constraints on the alignment between \overline{u}^* and θ and the alignment between \overline{u}^* and θ as follows. For any absolute constant $c_1 > 0$, there exist absolute constants $c_2, c > 0$ such that

$$|\langle \boldsymbol{\theta}, \boldsymbol{u}^{\star} \rangle| = c_1 d^{-r_1} \Leftrightarrow |\langle \boldsymbol{\theta}, \overline{\boldsymbol{u}}^{\star} \rangle| = c_2 \frac{d^{-r_1} \sqrt{1 + d^{r_2}}}{\sqrt{d^{r_2} (1 + d^{-2r_1})}} = c d^{-r}.$$
(2.1)

Remark. [Invariance of SQ under whitening] For any bounded query $q : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$, setting

$$q_{\boldsymbol{x}}(\boldsymbol{x},y) = q\left(\boldsymbol{\Sigma}^{-1/2}\boldsymbol{x},y\right)$$
 and $q_{\boldsymbol{z}}(\boldsymbol{z},y) = q\left(\boldsymbol{\Sigma}^{1/2}\boldsymbol{z},y\right)$

gives a one-to-one correspondence between queries with respect to x and queries with respect to z. In particular, any SQ algorithm for one model can be simulated with identical query complexity and tolerance for the other.

3. Main Results: Computational and Statistical Bounds for Learning the Spiked SIM

We prove computational lower bounds for learning the spiked single-index model on algorithms having access to statistical queries, which informally include online SGD. We also prove an informationtheoretic lower bound. Additionally, we prove complementary upper bounds, by characterizing the sample complexity required for online SGD and tensor power iteration to learn the spiked SIM. The general strategy for constructing these types of lower bounds is to exhibit a class of almost orthogonal hypotheses and to determine how hard they are to distinguish under our statistical and computational constraints. In their basic form, statistical query lower bounds are solving an easier question, which is to quantify the hardness of distinguishing between a hypothesis class with a planted structure versus a unplanted hypothesis. Instead of describing an estimation problem, they describe a testing/detection problem. Even though this problem seems much easier and hence unable to provide tight lower bounds, it turns out that these lower bounds successfully capture most of the hardness in notorious estimation problems such as Tensor PCA [7] and planted k-SAT [9]. However, there is a well-documented estimation/testing gap in some of these lower bounds, resulting in a small gap between the upper bounds and lower bounds. For a more thorough exposition, we refer the reader to [8]. Our results are subject to such a gap and we hypothesize that it stems from this estimation/testing discrepancy.

Remark. The bounds we present depend on well-known complexity measures of the single-index model, namely, the information and generative exponents. We delay their definition to the appendix.

3.1. Lower bounds

Theorem 3 (CSQ lower bound). Consider the class of bounded correlational queries $\tilde{q} \in \mathcal{Q}$, $\mathbb{E}[\tilde{q}(\boldsymbol{z})^2] = 1$. For $l^* \geq 3$, to learn the class of functions

$$\mathcal{F}_{l^{\star},\boldsymbol{\theta}.r} \coloneqq \Big\{ \boldsymbol{z} \mapsto g(\langle \overline{\boldsymbol{u}}, \boldsymbol{z} \rangle) : \overline{\boldsymbol{u}} \in S_{\boldsymbol{\theta},r}, \quad \|g\|_{L^{2}(\gamma)} = 1, \quad \mathrm{IE}(g) = l^{\star} \Big\},$$

with squared L^2 error at most 1, any correlational statistical query learner with q = poly(d) queries requires a tolerance of at most

$$\tau \lesssim \begin{cases} d^{-rl^{\star}} & \text{if } r < \frac{1}{4}, \\ \left(\frac{\log(qd)}{d}\right)^{l^{\star}/4} & \text{if } r \ge \frac{1}{4}. \end{cases}$$

Remark. Using the informal translation to sample complexity lower bounds via $\tau \simeq 1/\sqrt{n}$, we obtain the sample complexity lower bound $n = \Omega(d^{2rl^*})$ when r < 1/4 and $n = \tilde{\Omega}(d^{l^*/2})$ when $r \ge 1/4$.

Theorem 4 (SQ lower bound). Any SQ algorithm using poly(d) queries to VSTAT(n) to learn the spiked single-index model with generative exponent $k^* \ge 3$ requires at least the following number of samples:

$$n = \begin{cases} \tilde{\Omega}(d^{k^{\star}/2}) & \text{if } r \ge \frac{1}{4}, \\ \Omega(d^{2rk^{\star}}) & \text{if } r < \frac{1}{4}. \end{cases}$$

Remark. Both the CSQ and SQ lower bounds suggest that the spike only helps when r < 1/4, since otherwise they coincide with the isotropic setting [4, 6].

Theorem 5 (Information-theoretic lower bound). *Any algorithm using (possibly) infinitely many queries to learn the spiked single-index model requires at least the following number of samples:*

$$n = \begin{cases} \Omega(d) & \text{if } r_1 \ge r_2/2, \\ \Omega(1) & \text{if } r_1 < r_2/2. \end{cases}$$

3.2. Upper bounds

Theorem 6 (CSQ upper bound). A version of online SGD and tensor power iteration learn the spiked single-index model with $n = \tilde{\mathcal{O}}(d^{2r(l^*-2)+1})$ samples, where $l^* = IE(f^*) \ge 3$ is the information exponent of f^* .

More precisely this sample complexity is achieved by tensor power iteration initialized at the spike θ , which directly extends the results of [5] to the spiked setting. However, note that it is not a SQ algorithm. Interestingly, online spherical SGD initialized at the spike also achieves this sample complexity, extending the results of [2]. Although not completely formal due to the adversarial noise in statistical queries, the sample complexity of online SGD on squared loss is widely assumed to be captured by the statistical query framework. Fundamentally, the added structure in the inputs helps us boost the initialization from random alignment of order $d^{-1/2}$ to alignment of order d^{-r} .

Remark. Note that this upper bound coincide with the isotropic upper bound (and lower bound) for r = 1/4. The fact that both our versions of tensor power iteration and online SGD coincide in their sample complexity is of particular interest. Indeed, [2] showed that online SGD can only learn the isotropic SIM with sample complexity $\Omega(d^{l^*-1})$, whereas tensor power iteration succeeds with $\mathcal{O}(d^{l^*/2})$ ([5]), which suggests that online SGD is not an optimal learning algorithm for the task. This might help explain the success of gradient-based approaches in practice, as making the modeling assumptions closer to reality by adding structure in the inputs enables online SGD to get closer to a theoretically optimal algorithm.

Corollary 7 (SQ upper bound). A version of tensor power iteration learns the spiked single-index model with $n = \tilde{\mathcal{O}}(d^{2r(k^*-2)+1})$ samples, where $k^* \geq 3$ is the generative exponent.

4. Conclusion

In this paper, we have provided a precise description of the fundamental limits of learning singleindex models when the input covariance has a spiked structure. Our work illustrates two broader insights:

- Preconditioning suffices. A simple whitening and a spike-based initialization of online SGD achieves nearly optimal CSQ performance, demonstrating how covariance alignment enables algorithms to learn more efficiently.
- 2. *Computational-statistical gap behavior*. As the covariance structure becomes more informative about the target, the computational-statistical gap tightens and the sample complexity can reach the information-theoretic limit irrespective of the information/generative exponent.

Future directions. Two natural extensions of this work are to consider general covariance structure (e.g. power-law spectrum) and multi-index models. There also remains to close the gap between the lower bounds and the upper bounds by finding algorithms that formally coincide with the queries.

References

- Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2552–2623. PMLR, 2023.
- [2] Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Online stochastic gradient descent on non-convex losses from high-dimensional inference. J. Mach. Learn. Res., 22:106–1, 2021.
- [3] Nader H Bshouty and Vitaly Feldman. On using extended statistical queries to avoid membership queries. *Journal of Machine Learning Research*, 2(Feb):359–395, 2002.
- [4] Alex Damian, Eshaan Nichani, Rong Ge, and Jason D Lee. Smoothing the landscape boosts the signal for sgd: Optimal sample complexity for learning single index models. Advances in Neural Information Processing Systems, 36, 2023.
- [5] Alex Damian, Loucas Pillaud-Vivien, Jason Lee, and Joan Bruna. Computational-statistical gaps in gaussian single-index models. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 1262–1262. PMLR, 2024.
- [6] Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pages 5413–5452. PMLR, 2022.
- [7] Rishabh Dudeja and Daniel Hsu. Statistical query lower bounds for tensor pca. *J. Mach. Learn. Res.*, 22(1), January 2021. ISSN 1532-4435.
- [8] Vitaly Feldman. A general characterization of the statistical query complexity. volume 65 of Proceedings of Machine Learning Research, pages 785–830, Amsterdam, Netherlands, 07– 10 Jul 2017. PMLR. URL http://proceedings.mlr.press/v65/feldman17c. html.
- [9] Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, page 77–86, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450335362. doi: 10.1145/2746539.2746577. URL https://doi.org/10.1145/2746539.2746577.
- [10] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S. Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. J. ACM, 64(2), April 2017. ISSN 0004-5411. doi: 10.1145/3046674. URL https://doi.org/10.1145/3046674.
- [11] Taj Jones-McCormick, Aukosh Jagannath, and Subhabrata Sen. Provable benefits of unsupervised pre-training and transfer learning via single-index models, 2025. URL https: //arxiv.org/abs/2502.16849.
- [12] Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM* (*JACM*), 45(6):983–1006, 1998.

- [13] Noam Levi and Yaron Oz. The underlying scaling laws and universal statistical structure of complex datasets, 2024. URL https://arxiv.org/abs/2306.14975.
- [14] Alireza Mousavi-Hosseini, Denny Wu, Taiji Suzuki, and Murat A Erdogdu. Gradient-based feature learning under structured data. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [15] D. L. Ruderman and W. Bialek. Statistics of natural images, 1994.
- [16] Balázs Szörényi. Characterizing statistical query learning: simplified notions and proofs. In *International Conference on Algorithmic Learning Theory*, pages 186–200. Springer, 2009.
- [17] Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.

Contents

1	Introduction	1						
	1.1 Contributions	2						
2	Preliminaries							
	2.1 Statistical queries	2						
	2.2 Reformulation as an isotropic SIM	3						
3	Main Results: Computational and Statistical Bounds for Learning the Spiked SIM							
	3.1 Lower bounds	4						
	3.2 Upper bounds							
4	onclusion							
A	Proof of Computational Lower Bounds							
	A.1 CSQ lower bound							
	A.2 SQ lower bound							
	A.2.1 SQ framework							
	A.2.2 Instantiation for single-index model	10						
	A.2.3 Proof of SQ lower bound	11						
B	Proof of Information-Theoretic Lower Bound	12						
С	Proofs of Upper Bounds							
	C.1 Tensor Power Iteration	15						
	C.2 Online SGD	16						
	C.2.1 Proofs of Propositions 35 and 36	18						

Notations. For vectors, we use $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ to denote Euclidean inner product and norm respectively. For matrices, $\|\cdot\|$ denotes the operator norm. For asymptotic orders with respect to d as $d \to \infty$, $\mathcal{O}(\cdot), \Theta(\cdot), \Omega(\cdot)$ stand for the standard Big-O notations, and $\tilde{\mathcal{O}}(\cdot), \tilde{\Theta}(\cdot), \tilde{\Omega}(\cdot)$ hide (poly-)logarithmic factors. We write $f(d) \asymp g(d)$ and $f(d) = \Theta(g(d))$ interchangeably. $o(\cdot)$ and $\omega(\cdot)$ correspond to the standard little-o, respectively little-omega notations. $\nabla_{\boldsymbol{w}}$ denotes the Euclidean gradient with respect to \boldsymbol{w} , and we omit the subscript when it is clear from context. We denote the expectation of a random variable $X \sim P$ by \mathbb{E}_X or \mathbb{E}_P and we omit the subscript when clear from context. The s-dimensional Gaussian measure on \mathbb{R}^s is denoted by γ_s , and $\gamma \coloneqq \gamma_1$.

Appendix A. Proof of Computational Lower Bounds

Lemma 8 (Number of almost orthogonal vectors in $S_{\theta,r}$). For any absolute constant c > 0, there exists an absolute constant C > 0 such that for any $M = \mathcal{O}(\text{poly}(d))$, there exist M elements in $S_{\theta,r}$ that satisfy, for all $i, j \in [M], i \neq j$,

$$\langle \overline{\boldsymbol{u}}_i, \overline{\boldsymbol{u}}_j \rangle \le \epsilon + cd^{-r},$$
 (A.1)

for $\epsilon = \sqrt{\frac{C \log M}{d}}$.

Proof. Given M > 0, let us construct a subset $S_{\epsilon} \subset S_{\theta,r}$ of cardinality $|S_{\epsilon}| = M$ that satisfies, for all $i, j \in [M], i \neq j$, $|\langle \overline{u}_i, \overline{u}_j \rangle| \leq \epsilon + c^2 d^{-2r}$. Afterwards, we find an upper bound on Mwith the probabilistic method. To construct S_{ϵ} , we use the following procedure. Let $\{t_i\}_{i=1}^M \sim$ $\operatorname{Unif}\left(\frac{1}{\sqrt{1-c^2d^{-2r}}}\mathbb{S}^{d-2}\right)$. Suppose $\{\theta, \theta_2, \ldots, \theta_d\}$ form an orthonormal basis of \mathbb{R}^d , and construct the matrix $\mathbf{A} = (\theta_2, \ldots, \theta_d) \in \mathbb{R}^{d \times (d-1)}$. Note that $\langle \mathbf{A}t, \theta \rangle = 0$ for all $t \in \mathbb{R}^{d-1}$, and that $\mathbf{A}^{\top}\mathbf{A} = \mathbf{I}_{d-1}$. Then, we let $\overline{u}_i = \mathbf{A}t_i + cd^{-r}\theta$. As a result, we have

$$\langle \overline{u}_i, \overline{u}_j \rangle = \langle t_i, t_j \rangle + c^2 d^{-2r}, \quad and \quad \langle \theta, \overline{u}_i \rangle = c d^{-r}, \quad for \ all \quad i \neq j, \quad i, j \in [M].$$
 (A.2)

Note that for fixed $i \neq j$, Hoeffding's inequality yields $\mathbb{P}(|\langle \mathbf{t}_i, \mathbf{t}_j \rangle| > \epsilon) \leq 2e^{-c_1 d\epsilon^2}$, for some absolute constant $c_1 > 0$. Thus by a union bound, the probability that $|\langle \mathbf{t}_i, \mathbf{t}_j \rangle| \leq \epsilon$ for all $i \neq j$ is at least $1 - M^2 e^{-c_1 d\epsilon^2}$. Therefore, if we take $\epsilon = \sqrt{\frac{C \log M}{d}}$, for some absolute constant C > 0, it is guaranteed that there exist $\{\mathbf{t}_i\}_{i=1}^M$ such that $|\langle \mathbf{t}_i, \mathbf{t}_j \rangle| \leq \epsilon$ for $i \neq j$. With this construction, we have

$$|\langle \overline{u}_i, \overline{u}_j \rangle| \le \epsilon + c^2 d^{-2r}, \quad \text{for all} \quad i \ne j, \, i, j \in [M]$$

Therefore, $S_{\theta,r}$ contains at least M vectors that are well-separated, and to ensure $\epsilon = o(1)$, we pick M to be at most poly(d).

A.1. CSQ lower bound

The Hermite expansion is a central tool enabling the study of the Gaussian single-index model. The (normalized) Hermite polynomials $\{h_i\}_{i>0}$ are defined as

$$h_j(w) \coloneqq \frac{(-1)^j e^{w^2/2}}{\sqrt{j!}} \frac{d^j}{dw^j} e^{-w^2/2},$$

and form an orthonormal basis for $L^2(\gamma) \coloneqq \{g : \mathbb{R} \to \mathbb{R} : \int_{\mathbb{R}} g(w)^2 d\gamma(w) < \infty\}$ leading to a measure of complexity of the link function f^* .

Definition 9 (Information exponent). Given $f \in L^2(\gamma)$, with Hermite expansion $f = \sum_{j\geq 0} \alpha_j h_j$, the information exponent of f is $\text{IE}(f) := \inf\{j > 0 : \alpha_j \neq 0\}$.

The proof of the CSQ lower bound relies on constructing a subclass $\mathcal{F} \subseteq \mathcal{F}_{l,\theta,r}$ of highly uncorrelated functions. In particular, we use the following lemma, adapted from Szörényi [16, Theorem 2].

Lemma 10 (Damian et al. [6, Lemma 2]). Let $\epsilon > 0$, and let \mathcal{F} be a class of bounded real-valued functions such that for every pair $f, g \in \mathcal{F}$ where $f \neq g$ we have $|\mathbb{E}[f(\mathbf{X})g(\mathbf{X})]| \leq \epsilon$ and $\mathbb{E}[f(\mathbf{X})^2] = 1$ for all $f \in \mathcal{F}$. Then, any correlational statistical query learner whose queries are bounded in L^2 norm, i.e. $\mathbb{E}[\tilde{q}(\mathbf{X})^2] = 1$, requires at least $\frac{|\mathcal{F}|(\tau^2 - \epsilon)}{2}$ queries to learn \mathcal{F} up to $2(1 - \epsilon)$ L^2 error.

Proof. [Proof of Theorem 3] Let $S_{\epsilon} \subset S_{\theta,r}$ with cardinality M be defined as in the proof of Lemma 8, and let $\mathcal{F}_{\epsilon} \coloneqq \{ \boldsymbol{z} \mapsto h_l(\langle \boldsymbol{u}, \boldsymbol{z} \rangle) : \overline{\boldsymbol{u}} \in S_{\epsilon} \} \subset \mathcal{F}_{l,\theta,r}$. Recall from the properties of Hermite polynomials that for every $\boldsymbol{u}, \boldsymbol{u}' \in S$, we have

$$\mathbb{E}_{\boldsymbol{Z}\sim\mathcal{N}(0,\mathbf{I}_d)}\left[h_l(\langle \overline{\boldsymbol{u}},\boldsymbol{Z}\rangle)h_l(\langle \overline{\boldsymbol{u}}',\boldsymbol{Z}\rangle)\right] = \langle \overline{\boldsymbol{u}},\overline{\boldsymbol{u}}'\rangle^l.$$

As a result, for $f \neq g$, $f, g \in \mathcal{F}_{l,\theta,r}$, we have $|\mathbb{E}[f(\mathbf{Z})g(\mathbf{Z})]| \leq C_l(\epsilon^l + c^{2l}d^{-2rl})$, for a constant C_l depending only on l. By Lemma 10, any CSQ learner requires a tolerance of at most $\tau^2 \leq 2q/M + C_l(\epsilon^l + c^{2l}d^{-2rl})$, where q denotes the number of queries. Taking M = 2qd, we obtain $\tau^2 \leq 1/d + C_l\left(\left(\frac{\log(2qd)}{d}\right)^{l/2} + c^{2l}d^{-2rl}\right)$, which completes the proof. \Box

A.2. SQ lower bound

A.2.1. SQ FRAMEWORK

Following [5], we introduce the SQ framework of [10] that enables to capture the complexity of a great variety of testing and estimation problems with a single number called the statistical dimension, for algorithms having access to a certain oracle. This framework generalizes the approach used for CSQ by allowing joint queries over the target and inputs, and we instantiate it in the setting of the single-index model.

Definition 11 (Search problem over distributions). Let \mathcal{X} be a domain, \mathcal{D} be a set of distributions over \mathcal{X}, \mathcal{F} be a set of solutions, and $\mathcal{Z} : \mathcal{D} \to 2^{\mathcal{F}}$ be a map to the set of valid solutions. The distributional search problem consists in finding a valid solution $f \in \mathcal{Z}(D)$ given oracle access to samples from an unknown distribution $D \in \mathcal{D}$. We will also use \mathcal{Z}_f to denote the set of distributions \mathcal{D} for which f is a valid solution.

Definition 12 (Relative pairwise correlation). *Given two distributions* D_1 , D_2 and a reference distribution D,

$$\chi_D(D_1, D_2) \coloneqq \int \frac{D_1(x)D_2(x)}{D(x)} dx - 1.$$
 (A.3)

Definition 13 $((\gamma, \beta)$ -correlation). We say that a set of m distributions $\mathcal{D} = \{D_1, \ldots, D_m\}$ is (γ, β) -correlated relative to a distribution D_0 over \mathcal{X} if $|\chi_{D_0}(D_i, D_j)| \leq \gamma$ for $i \neq j$ and $|\chi_{D_0}(D_i, D_i)| \leq \beta$ for all $i \in [m]$. **Definition 14** (SQ dimension). *Given a search problem* Z *and parameters* $\gamma, \beta > 0$, we define the statistical query dimension $SD(Z, \gamma, \beta)$ to be the largest integer m such that there exists a distribution D_0 over X and a finite set of distributions $D_D \subset D$ with $|D_D| \ge m$ such that for any $f \in \mathcal{F}, D_f := D_D \setminus Z_f$ is (γ, β) -correlated relative to D_0 .

Definition 15 (VSTAT Oracle). Let *D* be the input distribution over the domain \mathcal{X} . Given *n* samples from *D*, for any query function $h : \mathcal{X} \to [0,1]$, VSTAT(*n*) oracle returns a value $v \in [\mathbb{E}_{x \sim D}[h(x)] - \tau, \mathbb{E}_{x \sim D}[h(x)] + \tau]$, where $\tau = \max\left\{\frac{1}{n}, \sqrt{\frac{\mathbb{E}_{x \sim D}[h(x)](1 - \mathbb{E}_{x \sim D}[h(x)])}{n}}\right\}$.

The VSTAT oracle corresponds to the response to a query having access to n i.i.d. samples with the noise tolerance τ corresponding to at most the concentration bound.

Lemma 16 (General SQ lower bound, Corollary 3.12 in [10]). For any $\gamma' > 0$, an SQ algorithm requires at least $SD(\mathcal{Z}, \gamma, \beta) \cdot \frac{\gamma'}{\beta - \gamma}$ queries to $VSTAT(\frac{1}{3(\gamma + \gamma')})$ to solve \mathcal{Z} .

Corollary 17. For any $\beta, \gamma, \tau \ge 0$, any algorithm requires at least $SD(\mathcal{Z}, \gamma, \beta) \cdot \frac{\frac{n}{\beta} - \gamma}{\beta - \gamma}$ queries to VSTAT(n) to solve \mathcal{Z} .

A.2.2. INSTANTIATION FOR SINGLE-INDEX MODEL

We follow the exposition in [5] and apply it to our specific setting to get the SQ lower bound for the spiked single-index model.

Definition 18 (Isotropic Gaussian single-index model). We say that a joint distribution $\mathbb{P} \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R})$ follows an isotropic Gaussian single index model if there exists a probability measure $P \in \mathcal{G} \subset \mathcal{P}(\mathbb{R}^2)$ and $v \in \mathbb{S}^{d-1}$ such that $\mathbb{P} = [R_v \otimes I_2]_{\sharp}[\gamma_{d-1} \otimes P]$, where $R_v \in \mathcal{O}_d$ is any orthogonal matrix whose last column is v, i.e. of the form $R_v = [R_1 v]$, and

$$\mathcal{G} \coloneqq \{\nu_{(w,y)} \in \mathcal{P}(\mathbb{R} \times \mathbb{R}); \ \nu_w = \gamma_1; \ \mathbb{E}_{\nu}[Y^2] < \infty; \ \chi^2(\nu_{(w,y)} || \gamma_1 \otimes \nu_y) > 0\}.$$
(A.4)

Having restated the problem as an isotropic single-index model in Subsection 2.2 enables us to use the lower bound from [5] and apply it to our more general setting. Let $W := \langle \overline{u}, Z \rangle$, for $\overline{u} \in S_{\theta,r}$ and $Z \sim \mathcal{N}(0, \mathbf{I}_d)$ so that $W \sim \mathcal{N}(0, 1)$. Then the search problem Z is of the following form:

- **Domain:** $\mathcal{X} = \mathbb{R}^d \times \mathbb{R}$ (it represents the (\mathbf{Z}, Y) pair),
- Distributions: $\mathcal{D} = \{\mathbb{P}_{\overline{u}} : \overline{u} \in S_{\theta,r}\},\$
- Solution Set: $\mathcal{F} = S_{\theta,r}$,
- Valid Solutions: $\mathcal{Z}(\mathbb{P}_{\overline{u}^{\star}}) = \{\overline{u} \in \mathcal{F} : |\langle \overline{u}, \overline{u}^{\star} \rangle| = \Theta(1)\},\$
- Reference Distribution: $D_0 = \gamma_d \otimes P_y$.

Then, we can use Corollary 17 to obtain a SQ lower bound in our setting. In order to do this, we must find a lower bound on $SD(\mathcal{Z}, \gamma, \beta)$ for some suitable γ, β .

A.2.3. PROOF OF SQ LOWER BOUND

Definition 19. For a distribution $P \in \mathcal{G}$, let

$$\zeta_k(Y) \coloneqq \mathbb{E}_P[h_k(W)|Y]. \quad and \quad \lambda_k \coloneqq \|\zeta_k\|_{P_y} = \left(\mathbb{E}[\zeta_k(Y)^2]\right)^{1/2}. \tag{A.5}$$

Definition 20. The generative exponent of P in G is defined as $GE(P) := \inf\{k > 0 : \lambda_k \neq 0\}$.

Remark. The generative exponent admits a variational form that makes its relationship with the information clearer. By reformulating the information exponent as

$$IE(P) = \inf\{l > 0 : \beta_l \neq 0\} \text{ where } \beta_l \coloneqq \mathbb{E}_P[Yh_l(W)],$$

we have that the variational formula for the generative exponent:

$$\operatorname{GE}(P) \coloneqq \inf_{T \in L^2(P_y)} \operatorname{IE}((\mathbf{I}_d \otimes T)_{\sharp} P).$$

We introduce [5, Lemma 3.1] without proof.

Lemma 21 (Relative pairwise correlation expansion). Let $\overline{u}, \overline{u}' \in S_{\theta,r}$, and let $m = \langle \overline{u}, \overline{u}' \rangle$. Then we have

$$\chi_0(\mathbb{P}_{\overline{\boldsymbol{u}}}, \mathbb{P}_{\overline{\boldsymbol{u}}'}) = \mathbb{E}_{\mathbb{P}_0}\left[\frac{d\mathbb{P}_{\overline{\boldsymbol{u}}}}{d\mathbb{P}_0} \cdot \frac{d\mathbb{P}_{\overline{\boldsymbol{u}}'}}{d\mathbb{P}_0}\right] - 1 = \sum_{k \le k^*} \lambda_k^2 m^k.$$
(A.6)

Lemma 22 (SQ dimension lower bound). Given a positive integer M = O(poly(d)), and $\epsilon = \sqrt{\frac{C \log(M)}{d}}$ for some absolute constant C > 0, the statistical dimension of the search problem \mathcal{Z} satisfies the following:

$$\mathcal{SD}(\mathcal{Z}, 2C_{k^{\star}}(\lambda_{k^{\star}}^{2}\epsilon^{k^{\star}} + c^{2k^{\star}}d^{-2rk^{\star}}), 1) \ge M,$$
(A.7)

for some constant $C_{k^*} > 0$ that only depends on k^* .

Proof. Let $M = \mathcal{O}(\text{poly}(d))$. Let $\epsilon > 0$ and let S_{ϵ} be defined as in the proof of Lemma 8. By construction, the cardinality of S_{ϵ} is M. Let $\mathcal{D} = \{\mathbb{P}_{\overline{u}_i} : \overline{u}_i \in S_{\epsilon}\}$. For any $i \neq j$,

$$\chi_0(\mathbb{P}_{\overline{\boldsymbol{u}}_i}, \mathbb{P}_{\overline{\boldsymbol{u}}_j}) = \sum_{k \ge 1} \lambda_k^2 \langle \overline{\boldsymbol{u}}_i, \overline{\boldsymbol{u}}_j \rangle^k \le \lambda_{k^\star}^2 \langle \overline{\boldsymbol{u}}_i, \overline{\boldsymbol{u}}_j \rangle^{k^\star} + \frac{\langle \overline{\boldsymbol{u}}_i, \overline{\boldsymbol{u}}_j \rangle^{k^\star + 1}}{1 - \langle \overline{\boldsymbol{u}}_i, \overline{\boldsymbol{u}}_j \rangle}$$
$$\le \lambda_{k^\star}^2 (\epsilon + c^2 d^{-2r})^{k^\star} + 2(\epsilon + c^2 d^{-2r})^{k^\star + 1}$$
$$\le C_{k^\star} \left(\lambda_{k^\star}^2 \epsilon^{k^\star} + (c^2 d^{-2r})^{k^\star} + 2(\epsilon^{k^\star + 1} + (c^2 d^{-2r})^{k^\star + 1}) \right)$$

where $C_{k^*} > 0$ is a constant depending only on k^* . Therefore, for $\lambda_{k^*}^2 \ge 2 \max\{\epsilon, c^2 d^{-2r}\}$, we can upper bound the right-hand side by $2C_{k^*}(\lambda_{k^*}^2 \epsilon^{k^*} + c^{2k^*} d^{-2rk^*}))$. Therefore, $\mathcal{SD}(\mathcal{Z}, 2C_{k^*}(\lambda_{k^*}^2 \epsilon^{k^*} + c^{2k^*} d^{-2rk^*}))$. \square

Theorem 23 (SQ lower bound). Any SQ algorithm using poly(d) queries to VSTAT(n) to learn the anisotropic single-index model with generative exponent k^* requires the following number of samples:

$$n = \begin{cases} \tilde{\Omega}(d^{k^{\star}/2}) & \text{if } r \geq \frac{1}{4}, \\ \Omega(d^{2rk^{\star}}) & \text{if } r < \frac{1}{4}. \end{cases}$$

Proof. By Lemma 22 and Corollary 17, the number of queries q to VSTAT(n) satisfies the following:

$$q \ge M \cdot \frac{\frac{3}{n} - 2C_{k^{\star}}(\lambda_{k^{\star}}^{2}\epsilon^{k^{\star}} + d^{-2rk^{\star}})}{1 - 2C_{k^{\star}}(\lambda_{k^{\star}}^{2}\epsilon^{k^{\star}} + d^{-2rk^{\star}})} \ge \frac{M}{2} \cdot \left(\frac{3}{n} - 2C_{k^{\star}}(\lambda_{k^{\star}}^{2}\epsilon^{k^{\star}} + d^{-2rk^{\star}})\right),$$
(A.8)

which implies

$$\frac{3}{n} \le \frac{2q}{M} + 2C_{k^{\star}} (\lambda_{k^{\star}}^2 \epsilon^{k^{\star}} + c^{2k^{\star}} d^{-2rk^{\star}}).$$
(A.9)

Setting M = 2qn = poly(d) yields

$$n \ge \frac{1}{C_{k^\star}(\lambda_{k^\star}^2 \epsilon^{k^\star} + c^{2k^\star} d^{-2rk^\star})},$$

which implies that

$$n = \begin{cases} \tilde{\Omega}(d^{k^{\star}/2}) & \text{if } r \geq \frac{1}{4}, \\ \Omega(d^{2rk^{\star}}) & \text{if } r < \frac{1}{4}. \end{cases}$$

Appendix B. Proof of Information-Theoretic Lower Bound

In order to provide an information-theoretic lower bound, we introduce the framework and some results from Wainwright [17, Chapter 15]. By reyling on Fano's inequality and Yang-Barron's method, we can lower bound the minimax risk by using the packing number of the set of parameters $S \in \{\mathbb{S}^{d-1}, S_{\theta,r}\}$ (depending on whether we tackle the isotropic case or the spiked covariance case), as well as the covering number of the parameterized set of joint probabilities $\mathcal{P} := \{\mathbb{P}_u : v \in S\}$ associated with single-index models.

Setup Given the class of isotropic Gaussian single-index models \mathcal{K} , we will derive a sample complexity lower bound on the hardness of learning $\mathbb{P} \in \mathcal{K}$, by constructing a minimax lower bound for the subset of distributions $\mathcal{P} := {\mathbb{P}_{v} : v \in S} \subset \mathcal{K}$ consisting of the joint distributions corresponding to the single-index model

$$y = f(\langle \boldsymbol{x}, \boldsymbol{v} \rangle) + \xi, \quad \boldsymbol{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad \xi \sim \mathcal{N}(0, \boldsymbol{\Sigma}^2),$$

for a fixed $f \in L^2(\gamma)$ with $||f||_{L^2(\gamma)} = 1$ and $\operatorname{IE}(f) = l^*$. Now, given a family $\overline{S} \subset S$ of cardinality M, let $\overline{\mathcal{P}} := \{\mathbb{P}_{\boldsymbol{v}} : \boldsymbol{v} \in \overline{S}\}$ be a family of single-index hypotheses, also of cardinality M. Let $\boldsymbol{Z} \in \mathbb{R}^d \times \mathbb{R}$ be a sample generated by uniformly sampling an index $J \sim \operatorname{Unif}([M])$ and then generating data (\boldsymbol{X}, Y) according to $\mathbb{P}_{\boldsymbol{v}_J}$. In this way, the observation \boldsymbol{Z} follows the mixture distribution $\mathbb{Q}_{\boldsymbol{Z}} = \overline{\mathbb{Q}} \coloneqq \frac{1}{M} \sum_{j=1}^{M} \mathbb{P}_{\boldsymbol{v}_j}$, which is the average over all hypotheses. Then the mutual information between the random variables \boldsymbol{Z} and J is defined as

$$I(\mathbf{Z}, J) \coloneqq \mathrm{KL}(\mathbb{Q}_{\mathbf{Z}, J} || \mathbb{Q}_{\mathbf{Z}} \mathbb{Q}_{J}),$$

where $\mathbb{Q}_{Z,J}$ denotes the joint distribution of (Z, J) and \mathbb{Q}_J is the distribution of J.

Definition 24. The minimax risk in our setting is defined as

$$R^{\star} \coloneqq \inf_{\hat{\boldsymbol{u}}} \sup_{\boldsymbol{u} \in S} \mathbb{E}_{\mathbb{P}_{\boldsymbol{u}}}(\|\hat{\boldsymbol{u}} - \boldsymbol{u}\|^2),$$

where the infimum ranges over all possible estimators.

We present two lemmas that we will use for our minimax lower bounds, and whose proofs can be found in [17, Chapter 15].

Lemma 25 (Fano's inequality). Given $\delta > 0$, let $\{v_1, \ldots, v_M\} \subset S$ be a 2δ -packing of S in the Euclidean norm $\|\cdot\|$, and suppose that J is uniformly distributed over the index set [M], and $(\mathbf{Z}|J=j) \sim \mathbb{P}_{v_j}$. Then the minimax risk is lower bounded as

$$R^{\star} \ge \delta^2 \cdot \left(1 - \frac{I(\mathbf{Z}, J_{\delta}) + \log 2}{\log M}\right).$$

Lemma 26 (Yang-Barron method). Let $\mathcal{N}_{KL}(\eta, \mathcal{P})$ denote the η -covering number of \mathcal{P} with respect to the square-root KL divergence. Then the mutual information is upper bounded as

$$I(\mathbf{Z}, J_{\delta}) \leq \inf_{\eta > 0} \{\eta^2 + \log(\mathcal{N}_{\mathcal{KL}}(\eta, \mathcal{P}))\}$$

Lemma 27. Given $v, v' \in S$, the KL divergence between \mathbb{P}_v and $\mathbb{P}_{v'}$ satisfies

$$KL(\mathbb{P}_{\boldsymbol{v}} \| \mathbb{P}_{\boldsymbol{v}'}) \asymp \frac{1}{\boldsymbol{\Sigma}^2} (1 - \left\langle \boldsymbol{v}, \boldsymbol{v}' \right\rangle^{l^*}), \tag{B.1}$$

where k^* denotes the information exponent of the single index.

Proof. The conditional probability distribution of $Y|\mathbf{X}$ for $(Y, \mathbf{X}) \sim \mathbb{P}_{v}$ is given by the Gaussian distribution $\mathcal{N}(f(\langle \mathbf{X}, v \rangle), \Sigma^{2})$. Hence, by standard computations of the KL for the normal location model (see e.g. [17][Example 15.13]),

$$\begin{split} KL(\mathbb{P}_{\boldsymbol{v}} \| \mathbb{P}_{\boldsymbol{v}'}) &= \frac{1}{2\boldsymbol{\Sigma}^2} \mathbb{E}_{\boldsymbol{X}} [f(\langle \boldsymbol{X}, \boldsymbol{v} \rangle) - f(\langle \boldsymbol{X}, \boldsymbol{v}' \rangle)^2] \\ &= \frac{1}{\boldsymbol{\Sigma}^2} (1 - (\alpha_{k^\star}^2 \langle \boldsymbol{v}, \boldsymbol{v}' \rangle^{l^\star} + o(\langle \boldsymbol{v}, \boldsymbol{v}' \rangle^{l^\star + 1})) \asymp \frac{1}{\boldsymbol{\Sigma}^2} (1 - \langle \boldsymbol{v}, \boldsymbol{v}' \rangle^{l^\star}), \end{split}$$

where f admits the Hermite expansion $f = \sum_{l>l^*} \alpha_l h_l$.

Assuming that $C_{\epsilon}(S)$ is an ϵ -covering, that is, $\|\boldsymbol{v} - \boldsymbol{v}'\| \leq \epsilon$ for all distinct $\boldsymbol{v}, \boldsymbol{v}'$ in $C_{\epsilon}(S)$, then $1 - \langle \boldsymbol{v}, \boldsymbol{v}' \rangle \leq \frac{\epsilon^2}{2}$. Thus,

$$KL(\mathbb{P}_{\boldsymbol{v}} \| \mathbb{P}_{\boldsymbol{v}'}) \leq \frac{1}{\boldsymbol{\Sigma}^2} (1 - \left\langle \boldsymbol{v}, \boldsymbol{v}' \right\rangle) \leq \frac{1}{\boldsymbol{\Sigma}^2} \cdot \epsilon^2.$$
(B.2)

Consequently, for an i.i.d. sample of n data points, we get the bound

$$\sqrt{KL(\mathbb{P}_{\boldsymbol{v}}^{\otimes n} \| \mathbb{P}_{\boldsymbol{v}'}^{\otimes n})} \le \frac{\sqrt{n}}{\Sigma} \epsilon \eqqcolon \eta.$$
(B.3)

Therefore, $\log(\mathcal{N}_{\mathcal{KL}}(\eta, \mathcal{P})) \leq \log(\mathcal{C}_{\epsilon}(\mathcal{S})).$

By standard properties of coverings and packings, for $P_{2\delta}(S)$ a 2δ -packing of size M, that is, $\|\boldsymbol{v} - \boldsymbol{v}'\| \ge 2\delta$ for all distinct $\boldsymbol{v}, \boldsymbol{v}'$ in $P_{2\delta}(S)$,

$$\log(P_{2\delta}(S) \asymp \log(C_{\delta}(S)). \tag{B.4}$$

In order to get a bound on the minimax risk that only depends on δ , we must select a value of δ for which

$$\frac{I(\boldsymbol{Z}, J_{\delta}) + \log 2}{\log M} \ge 1/2.$$

By following the recipe in [17], we first prescribe $\eta > 0$ such that

$$\eta^2 \ge \log\left(\mathcal{C}_{\frac{\Sigma}{\sqrt{n}}\eta}(S)\right). \tag{B.5}$$

Then we pick the largest $\delta > 0$ that satisfies the lower bound

$$\log(M(\delta)) \ge 4\eta^2 + 2\log 2. \tag{B.6}$$

Fact 28. For $\epsilon > 0$, the covering number of the sphere of radius ρ in Euclidean norm is given by

$$\mathcal{N}(\epsilon,\rho\mathbb{S}^{d-1},\|\cdot\|) \asymp \left(\frac{\rho}{\epsilon}\right)^{d-1}$$

Let us focus on the case $S = \mathbb{S}^{d-1}$ first. In that case, Equation (B.5) yields the inequality

$$\eta^2 \gtrsim d \log\left(\frac{\sqrt{n}}{\eta}\right),$$
(B.7)

which is fulfilled by setting $\eta \approx \sqrt{d}$. We immediately see from Equations (B.4)-(B.6) that the pair (δ, ϵ) must satisfy $\delta \leq \epsilon$. Since we must pick the largest such δ , taking $\delta \approx \epsilon$ ensures that Equations (B.5)-(B.6) are satisfied. Hence, $R^* \gtrsim \delta^2 \approx \eta^2/n \approx d/n$, which yields the sample complexity lower bound

$$n = \Omega(d)$$

Proposition 29. For r = 0, there are two cases for the constraint $\langle \overline{u}, \theta \rangle \approx d^{-r}$. If $r_1 = r_2/2$, then $\langle \overline{u}, \theta \rangle = c$, for some absolute constant $c \in (0, 1]$. However, if $r_1 < r_2/2$, it further holds that $\langle \overline{u}, \theta \rangle = c \approx 1$.

Proof. For $r_1 < r_2/2$,

$$\langle \overline{\boldsymbol{u}}, \boldsymbol{ heta}
angle = rac{\left\langle \boldsymbol{\Sigma}^{1/2} \boldsymbol{ heta}, \boldsymbol{\Sigma}^{1/2} \boldsymbol{u}
ight
angle}{\left\| \boldsymbol{\Sigma}^{1/2} \boldsymbol{u}
ight\|} = rac{\sqrt{1+\kappa} \langle \boldsymbol{ heta}, \boldsymbol{u}
angle}{\sqrt{1+\kappa} \langle \boldsymbol{ heta}, \boldsymbol{u}
angle^2} pprox rac{\sqrt{\kappa} \langle \boldsymbol{ heta}, \boldsymbol{u}
angle}{\sqrt{\kappa} |\langle \boldsymbol{ heta}, \boldsymbol{u}
angle|} = 1$$

where the approximation comes from

$$\kappa \langle \boldsymbol{\theta}, \boldsymbol{u} \rangle^2 = c_2 d^{r_2 - 2r_1} = \omega(1), \text{ which implies } \sqrt{1 + \kappa \langle \boldsymbol{\theta}, \boldsymbol{u} \rangle^2} \approx \sqrt{\kappa} |\langle \boldsymbol{\theta}, \boldsymbol{u} \rangle|.$$

Remark. We will prove below that Proposition 29 implies that the information-theoretic lower bound undergoes a sharp transition from $n = \Omega(d)$ to n = 1 at $r_1 = r_2/2$. Moreover, when $r_1 < r_2/2$, taking θ as our estimator for \overline{u}^* is enough to learn the single-index model's unknown direction.

Lemma 30 (Covering number of $S_{\theta,r}$). For $\epsilon > 0$, the covering number of $S_{\theta,r} = \{\overline{u} \in \mathbb{S}^{d-1} : \langle \overline{u}, \theta \rangle = cd^{-r}\}$ satisfies

$$\log \mathcal{N}(\epsilon, S_{\boldsymbol{\theta}, r}, \|\cdot\|) \asymp \begin{cases} d \log(1/\epsilon) & \text{if } r_1 \ge \frac{r_2}{2} \\ 0 & \text{if } r_1 < \frac{r_2}{2} \end{cases}$$

Proof. For $r_1 < \frac{r_2}{2}$, by Proposition 29, $S_{\theta,r} = \{\overline{u} \in \mathbb{S}^{d-1} : \langle \overline{u}, \theta \rangle \approx 1\}$. For $r_1 \ge \frac{r_2}{2}$, given the orthonormal basis $\{\theta_1, \ldots, \theta_d\}$ where $\theta_d = \theta$,

$$\langle \overline{\boldsymbol{u}}, \boldsymbol{\theta} \rangle = \sum_{i=1}^{d-1} \langle \overline{\boldsymbol{u}}, \boldsymbol{\theta}_i \rangle + cd^{-r}\boldsymbol{\theta},$$

and thus the covering number of $S_{\theta,r}$ satisfies

$$\log C_{\epsilon}(S_{\theta,r}) \asymp \log C_{\epsilon} \left(\frac{1}{\sqrt{1 - c^2 d^{-2r}}} \mathbb{S}^{d-2} \right) \asymp (d-2) \log \left(\frac{1}{\sqrt{1 - c^2 d^{-2r} \epsilon}} \right)$$
$$\asymp d \log \left(\frac{1}{\epsilon} (1 + c^2 d^{-2r}) \asymp d \log(1/\epsilon). \right)$$

In the case of $S = S_{\theta,r}$ and $r_1 \ge r_2/2$, Equation (B.6) yields the inequality

$$\eta^2 \gtrsim d \log\left(\frac{\sqrt{n}}{\eta}\right).$$

Thus, we must pick $\eta \asymp \sqrt{d}$. Hence $R^{\star} \ge \delta^2 \asymp \eta^2/n$, which yields the sample complexity lower bound

$$n = \Omega(d).$$

However, when $r_1 < r_2/2$, we can pick $\eta^2 = 1$ to obtain $R^* \ge 1/n$, which yields the sample complexity lower bound $n = \Omega(1)$.

Appendix C. Proofs of Upper Bounds

We derive upper bounds on learning the spiked single-index model by initializing existing algorithms at the covariance spike. Our upper bounds only consider learning the unknown direction \overline{u}^* of the isotropic formulation of the spiked single-index model; they do not include the sample complexity to learn the link function f^* . However, once \overline{u}^* (or u^*) is approximately recovered, recovering f^* is a simple convex optimization problem. A more detailed exposition of how this is can be done with two-layer neural networks and gradient descent is provided in [6].

C.1. Tensor Power Iteration

We use the tensor power iteration algorithm from [5] initialized at the covariance spike. By simply plugging in the initial alignment $m_0 = \Theta(d^{-r})$ in [5], we obtain the upper bound on the sample complexity

$$n = \tilde{\mathcal{O}}(d^{2rl^{\star} + 1 - 4r}) \tag{C.1}$$

to learn the spiked single-index model.

Proof. By [5, Lemma F.4], for any initial alignment $m_0 = \Theta(d^{-r})$, for any $r \in [0, 1/4]$, one step of tensor power iteration with $\Theta(d^{1+2rl^*-4r})$ samples achieves an alignment $m_1 \ge 2m_0$. Therefore, taking $s = \Theta(\log d)$ steps of tensor power iteration yields $m_s = \Theta(1)$. Hence, $n = \tilde{\mathcal{O}}(d^{1+2rl^*-4r})$ are sufficient to learn the hidden dimension \overline{u}^* .

This upper bound coincides with the one we obtain below with Online SGD initialized at the covariance spike.

For the SQ upper bound, we can analogously use [5, Algorithm 2] with initialization at θ . At a high level, the algorithm finds a label transformation with a denoiser $T \in L^2(P_y)$ such that IE(T(P)) = GE(P). This yields the sample complexity upper bound

$$n = \tilde{\mathcal{O}}(d^{1+2rk^{\star}-4r}).$$

depending on the generative exponent k^{\star} .

C.2. Online SGD

By extending the work of [2] to consider the case where the initial alignment between the weight and the hidden direction is of order d^{-r} , for $r \in [0, 1/2]$, we can derive a sample complexity bound for online SGD initialized at the spike of the covariance matrix. This contrasts with the uninformative weight initialization which achieves an alignment of order $1/\sqrt{d}$ with high probability. In the analysis performed in [2], the alignment at timestep t is decomposed into four components: the initial alignment, a drift term, a martingale term corresponding to the sample-wise error, and higher-order terms in the learning rate, corresponding to the projection onto the sphere. Then, the game consists of finding the appropriate learning rate and number of samples, so that the initial alignment and the drift term dominate the dynamics and make online SGD achieve weak recovery of the hidden direction. Therefore, starting with a better initialization yields a better sample complexity as it allows the selection of a larger learning rate. There is no technical innovation in the proof and the improved sample complexity has been described in a concurrent work [11] in a different context and without proof, but we show here how the proof of [2] generalizes for completeness. We only adapt the weak recovery proof, as it is the phase where the information exponent appears. The descent phase/strong recovery always takes O(d) samples.

Setup We study the isotropic single-index model $y = f^*(\langle \boldsymbol{z}, \overline{\boldsymbol{u}}^* \rangle) + \xi$, $\boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, $\xi \sim \mathcal{N}(0, \Sigma^2)$. We denote the loss by $\mathcal{L} : \mathbb{S}^{d-1} \times \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$, and study the case where the population loss is of the form

$$\Phi(\boldsymbol{w})\coloneqq\mathbb{E}_{\boldsymbol{Z},Y}[\mathcal{L}]=\phi(m(\boldsymbol{w}))\quad\text{where}\quad m(\boldsymbol{w})\coloneqq\langle\boldsymbol{w},\overline{\boldsymbol{u}}^*\rangle,$$

for some $\phi : [-1, 1] \to \mathbb{R}$. Moreover, we define the sample-wise error

$$H^t(\boldsymbol{w}) \coloneqq \mathcal{L}(\boldsymbol{w}; \boldsymbol{z}_t, y_t) - \Phi(\boldsymbol{w})$$

In this subsection, we denote the information exponent by k, and we follow the more general definition of [2].

Algorithm 1: Spherical Online SGD Initialized at Covariance Spike

Data: $(\boldsymbol{z}_i, y_i)_{i=1}^n \sim \mathbb{P}_{\overline{\boldsymbol{u}}^*}, \quad \boldsymbol{\theta} \in \mathbb{S}^{d-1}$ with $|\langle \boldsymbol{\theta}, \overline{\boldsymbol{u}}^* \rangle| = cd^{-r},$ step size $\eta > 0$ **Result:** Final iterate \boldsymbol{w} on the sphere $\boldsymbol{w} \leftarrow \boldsymbol{\theta};$ for $t \in [n]$ do $| \boldsymbol{w} \leftarrow \boldsymbol{w} - \frac{\eta}{d} (\mathbf{I}_d - \boldsymbol{w} \boldsymbol{w}^\top) \nabla \mathcal{L}(\boldsymbol{w}; \boldsymbol{z}_t, y_t);$ $\boldsymbol{w} \leftarrow \boldsymbol{w} / || \boldsymbol{w} ||;$ end return $\boldsymbol{w};$

Definition 31 (Information exponent). The population loss Φ has information exponent k if $\phi \in C^{k+1}$) and there exist C, c > 0 such that

$$\begin{cases} \frac{d^l \phi}{dm^l}(0) = 0 \quad 1 \le l \le k \\ \frac{d^k \phi}{dm^k}(0) \le -c < 0 \\ \left\| \frac{d^{k+1} \phi}{dm^{k+1}}(m) \right\|_{\infty} \le C. \end{cases}$$

Definition 32. For a population loss of the form $\Phi(w) \coloneqq \mathbb{E}[\mathcal{L}] = \phi(m(w))$ where $m(w) \coloneqq \langle w, \overline{u}^* \rangle$, for some $\phi : [-1, 1] \to \mathbb{R}$, we say that Assumption A_{ϱ} for some ϱ in (0, 1] holds if

- *1.* ϕ *is differentiable, and*
- 2. ϕ' is strictly negative on the interval $(0, \varrho)$.

When $\rho = 1$, we say that Assumption A holds.

Definition 33. For a data distribution and loss pair $(\mathbb{P}, \mathcal{L})$, Assumption B holds if there exist $C_1, \iota > 0$ such that the following two moment bounds hold for all d:

1. We have that

$$\sup_{\boldsymbol{w}\in\mathbb{S}^{d-1}}\mathbb{E}\left[|\langle\nabla H(\boldsymbol{w}),\overline{\boldsymbol{u}}^{\star}\rangle|^{2}\right]\leq C_{1},$$

2. and that

$$\sup_{\boldsymbol{w}\in\mathbb{S}^{d-1}}\mathbb{E}\left[\|\nabla H(\boldsymbol{w})\|^{4+\iota}\right]\leq C_1d^{2+\iota/2}.$$

For every μ , define the hitting times

$$\tau_{\mu}^{+} \coloneqq \min\{t \ge 0 : m(\boldsymbol{W}_{t}) \ge \mu\}, \quad \text{and} \quad \tau_{\mu}^{-} \coloneqq \min\{t \ge 0 : m(\boldsymbol{W}_{t}) \le \mu\}.$$

Fix ι given in Assumption B and define

$$\bar{L} \coloneqq \sup_{\boldsymbol{w}} \mathbb{E}\left[\left\|\frac{1}{\sqrt{d}}\nabla H(\boldsymbol{w})\right\|^{4+\iota}\right] \vee \sup_{\boldsymbol{w}} \mathbb{E}\left[\left\|\frac{1}{\sqrt{d}}\nabla H(\boldsymbol{w})\right\|^{2}\right] \vee 1.$$
(C.2)

Additionally, let $a_k = ck$, $a_{k+1} = C(k+1)$, where C, c are as in the definition of the information exponent, and let m_t denote $m(W_t)$.

Theorem 34. Suppose there exists $\varrho > 0$ such that Assumptions A_{ϱ} and B hold and that the population loss has information exponent k. Let (α_d, η_d) be as in Proposition 35. Then there exists $\nu > 0$ such that if $\mathbf{W}_t = \mathbf{W}_t^{d,\eta}$ is the online SGD with step size η , we have for every $\gamma > 0$,

$$\lim_{d \to \infty} \inf_{\boldsymbol{w}_0: m(\boldsymbol{w}_0) \ge \gamma/d^r} \mathbb{P}_{\boldsymbol{w}_0} \Big(\tau_{\nu}^+ < \alpha d \Big) = 1.$$

where, we recall, τ_{ν}^+ is the stopping time $\inf\{t: m_t > \nu\}$.

Proposition 35. Suppose that Assumptions A_{ϱ} and B hold and that the population loss has information exponent $k \geq 3$. Let $D = D_d$ and $\epsilon = \mathcal{O}(d^{1+(k-3)r})$, and suppose $\alpha = \alpha_d$ is of at most polynomial growth in d, and $\eta = \eta_d$ is such that $\alpha \eta^2 \leq \epsilon$ and for some K > 0,

$$\eta \le \bar{\eta}_d(k) := \frac{a_k \gamma^{k-2}}{K \bar{L} d^{r(k-2)} \log d}.$$
(C.3)

Then for every $\gamma > 0$ and every $T \le M := \alpha d$ satisfying

$$T \le \frac{d^{2(1-r)}\gamma^2}{D^2\eta^2} \eqqcolon \bar{t}, \qquad (C.4)$$

,

online SGD with step-size η satisfies the following as $d \to \infty$ for some $\nu > 0$, uniformly over the choice of D, ϵ, K : there exists a constant $C(C_1, a_k, a_{k+1}) > 0$ such that

$$\inf_{\boldsymbol{w}_0 \in E_{\gamma/d^r}} \mathbb{P}_{\boldsymbol{w}_0} \Big(m_t \ge \frac{m_0}{2} + \frac{\eta a_k}{8d} \sum_{j=0}^{t-1} m_j^{k-1} \quad \forall t \le \tau_{\gamma/2d^r}^- \land \tau_{\nu}^+ \land T \Big) \ge 1 - \frac{C}{D^2} - o(1) \,. \quad (C.5)$$

Proposition 36 (Weak recovery). Under the assumptions of Proposition 35, for $\alpha = \Omega(d^{2r(k-2)-1})$ and $\eta \leq \bar{\eta}(k)$ satisfying $\alpha \eta^2 \leq \epsilon$ for $\epsilon = \mathcal{O}(d^{1+(k-3)r})$, there exists $\nu_0(\varrho, a_k, a_{k+1}) > 0$ such that for every $\nu < \nu_0$, for every $\gamma > 0$, we have

$$\lim_{d\to\infty} \inf_{\boldsymbol{w}_0\in E_{\gamma/d^r}} \mathbb{P}_{\boldsymbol{w}_0}\left(\tau_{\nu}^+ \leq \bar{t} \wedge M\right) = 1.$$

Corollary 37. We can weakly recover the single-index model with online SGD initialized at θ as soon as the sample complexity is

$$n = \alpha d = \tilde{\Omega}(d^{2rk+1-4r}).$$

C.2.1. PROOFS OF PROPOSITIONS 35 AND 36

We need to adapt the proofs of the bounds of all the different components appearing in the decomposition of the alignment m_t . Without loss of generality, we assume that $\overline{u}^* = e_1$, the first vector of the canonical basis. From [2], we have the inequality

$$m_{t} \geq m_{t-1} - \frac{\eta}{d} \langle \nabla \Phi(\mathbf{W}_{t-1}), \mathbf{e}_{1} \rangle - \frac{\eta}{d} \langle \nabla H^{t}(\mathbf{W}_{t-1}), \mathbf{e}_{1} \rangle - \frac{\eta^{2}}{d} L_{t} \mathbf{1}_{\{L_{t} < \hat{L}\}} |m_{t-1}|$$
(C.6)
$$- \eta^{2} \left(\frac{A}{d^{2}} + \frac{1}{d} L_{t} \mathbf{1}_{\{L_{t} \ge \hat{L}\}} \right) |m_{t-1}| - \eta^{3} \left(\frac{A}{d^{2}} + \frac{L_{t}}{d} \right) \left(\left| \frac{\langle \nabla \Phi(\mathbf{W}_{t-1}), \mathbf{e}_{1} \rangle}{d} \right| + \left| \frac{\langle \nabla H^{t}(\mathbf{W}_{t-1}), \mathbf{e}_{1} \rangle}{d} \right| \right)$$
(C.7)

where A is a constant depending only on a_k, a_{k+1} and $L_t := \left\| \frac{1}{\sqrt{d}} \nabla H^t(\mathbf{W}_{t-1}) \right\|^2$. Let us start by controlling the higher order corrections that arise from the projection on the sphere of the parameters at each iteration of the algorithm. The bounds do not change much, except carefully pick \hat{L} .

Observe that for every $\nu < 1/2$, for every $\boldsymbol{w} \in \left\{ \boldsymbol{w} : m(\boldsymbol{w}) \in [0, \nu] \right\}$,

$$\langle \nabla m(\boldsymbol{w}), \boldsymbol{e}_1 \rangle = (\boldsymbol{e}_1 - \langle \boldsymbol{w}, \boldsymbol{e}_1 \rangle \boldsymbol{w}) = 1 - m(\boldsymbol{w})^2 \ge 1 - \nu^2 \ge 1/2.$$

Then there exists $\nu_0(\varrho, a_k, a_{k+1}) > 0$ such that for all $\nu < \nu_0$, for all $x \in \{ \boldsymbol{w} : m(\boldsymbol{w}) \in [0, \nu] \}$,

$$\frac{1}{4}a_km(\boldsymbol{w})^{k-1} \leq -\langle \nabla \Phi(\boldsymbol{w}), \boldsymbol{e}_1 \rangle \leq \frac{3}{2}a_km(\boldsymbol{w})^{k-1}$$

Lemma 38. Let $\nu, \gamma > 0$ with $\nu < 1/2$. For all K > 0, for all $\eta < \overline{\eta}_d(k)$, for $k \ge 2$,

$$\forall \boldsymbol{w} \in \mathcal{A} \coloneqq \left\{ \boldsymbol{w} : m(\boldsymbol{w}) \in \left[\frac{\gamma}{2d^r}, \nu\right] \right\}, \quad \frac{\eta^2}{d} |m(\boldsymbol{w})| \le \frac{\eta}{d} 2^{k-2} a_k \frac{|m(\boldsymbol{w})|^{k-1}}{K\bar{L}\log d}.$$

Proof. Since $w \in A$ and given Equation (C.3), we have

$$\eta^{2}|m(\boldsymbol{w})| \leq \frac{a_{k}}{K\bar{L}\log d} \cdot \eta|m(\boldsymbol{w})| \cdot \left(\frac{\gamma}{d^{r}}\right)^{k-2} \leq \frac{a_{k}}{K\bar{L}\log d}\eta|m(\boldsymbol{w})|^{k-1},$$

where the last inequality follows from the fact that $|m(\boldsymbol{w})| > \frac{\gamma}{2d^r}$.

Lemma 39. Suppose that $\alpha \eta^2 \leq \epsilon$ for some $\epsilon = \mathcal{O}(d^{1+(k-3)r})$, and let \overline{L} be as in Equation (C.2). There exists $C = C(\overline{L}, A, C_1, C_2) > 0$ such that the following holds uniformly over $w_0 \in \mathbb{S}^{d-1}$:

$$\mathbb{P}_{\boldsymbol{w}_0}\left(\sup_{t\leq M}\eta^3\sum_{j=0}^{t-1}\left(\frac{A}{d^2}+\frac{L_{j+1}}{d}\right)\left(\left|\frac{\langle\nabla\Phi(\boldsymbol{W}_j),\boldsymbol{e}_1\rangle}{d}\right|+\left|\frac{\langle\nabla H^{j+1}(\boldsymbol{W}_j),\boldsymbol{e}_1\rangle}{d}\right|\right)>\frac{\gamma}{10d^r}\right)\leq\frac{C\epsilon\eta}{\gamma d^{1-r}},$$
(C.8)

$$\mathbb{P}_{\boldsymbol{w}_{0}}\left(\sup_{t\leq M}\eta^{2}\sum_{j=0}^{t-1}\left(\frac{A}{d^{2}}+\frac{L_{j+1}\mathbf{1}_{\{L_{j+1}<\hat{L}\}}}{d}\right)|m(\boldsymbol{W}_{j})| > \frac{\gamma}{10d^{r}}\right) \leq \frac{C\epsilon d^{r}}{\hat{L}^{1+\iota/4}\gamma}.$$
(C.9)

Proof. We start by proving the first bound, by using Markov's inequality and Cauchy-Schwarz inequality.

$$\begin{split} \mathbb{P}_{\boldsymbol{w}_{0}} \left(\sup_{t \leq M} \eta^{3} \sum_{j=0}^{t-1} \left(\frac{A}{d^{2}} + \frac{L_{j+1}}{d} \right) \left(\left| \frac{\langle \nabla \Phi(\boldsymbol{W}_{j}), \boldsymbol{e}_{1} \rangle}{d} \right| + \left| \frac{\langle \nabla H^{j+1}(\boldsymbol{W}_{j}), \boldsymbol{e}_{1} \rangle}{d} \right| \right) > \lambda \right) \\ &\leq \frac{M\eta^{3}}{\lambda d^{2}} \sup_{\boldsymbol{w}} \mathbb{E} \left[\frac{|\nabla H(\boldsymbol{w})|^{2}}{d} + \frac{A}{d} \right] \left(|\langle \nabla \Phi(\boldsymbol{w}), \boldsymbol{e}_{1} \rangle| + |\langle \nabla H(\boldsymbol{w}), \boldsymbol{e}_{1} \rangle| \right) \\ &\leq \frac{\alpha d\eta^{3}}{\lambda d^{2}} \sqrt{\sup_{\boldsymbol{w}} \mathbb{E} \left[\left| \frac{\nabla H(\boldsymbol{w})}{d} \right|^{4} \right]} \sqrt{\sup_{\boldsymbol{w}} |\langle \nabla \Phi(\boldsymbol{w}), \boldsymbol{e}_{1} \rangle|^{2} + \sup_{\boldsymbol{w}} |\langle \nabla H(\boldsymbol{w}), \boldsymbol{e}_{1} \rangle|^{2}} \\ &\leq \frac{\alpha d\eta^{3}}{\lambda d^{2}} \sqrt{\bar{L} + \frac{A^{2}}{d^{2}}} \cdot \sqrt{A + C_{1}}. \end{split}$$

Take $\lambda = \frac{\gamma}{10d^r}$ and use $\alpha \eta^2 \leq \epsilon$, which yields the upper bound

$$\frac{d^{r-1}\eta\epsilon}{\gamma}\sqrt{\left(\bar{L}+\frac{A^2}{d^2}\right)(A+C_1)}.$$

For the second bound, we also use Markov's and Cauchy-Schwarz inequalities. To ease notations, denote $L_{j+1}\mathbf{1}_{\{L_{j+1}<\hat{L}\}}$ by \tilde{L}_{j+1} .

$$\mathbb{P}_{\boldsymbol{w}_0}\left(\sup_{t\leq M}\sum_{j=0}^{t-1}\eta^2\left(\frac{A}{d^2}+\frac{\tilde{L}_{j+1}}{d}\right)|m_j|>\lambda\right)\leq \mathbb{P}_{\boldsymbol{w}_0}\left(\eta^2\sum_{j=0}^{M-1}\frac{\tilde{L}_{j+1}}{d}>\lambda-\frac{A\epsilon}{d}\right).$$

By Markov's inequality,

$$\mathbb{P}_{\boldsymbol{w}_{0}}\left(\sum_{j=0}^{M-1}\frac{\tilde{L}_{j+1}}{d} > \Lambda\right) \leq \frac{M}{\Lambda} \sup_{j \leq M} \mathbb{E}_{\boldsymbol{w}_{0}}\left[\tilde{L}_{j+1}\right]$$

$$\leq \frac{M}{\Lambda}\sqrt{\sup_{\boldsymbol{w}} \mathbb{E}\left(\left|\frac{1}{\sqrt{d}}\nabla H(\boldsymbol{w})\right|^{4}\right) \cdot \sup_{\boldsymbol{w}} \mathbb{P}\left(\left|\frac{1}{\sqrt{d}}\nabla H(\boldsymbol{w})\right|^{2} > \hat{L}\right)}$$

$$\leq \frac{\alpha d}{\Lambda}\sqrt{\bar{L}\sup_{\boldsymbol{w}} \frac{\mathbb{E}\left(\left|\frac{1}{\sqrt{d}}\nabla H(\boldsymbol{w})\right|^{4+\iota}\right)}{\hat{L}^{2+\iota/2}}} \leq \frac{\alpha d}{\Lambda}\frac{\bar{L}}{\hat{L}^{1+\iota/4}}.$$

By setting $\Lambda = \frac{d}{\eta^2} \left(\lambda - \frac{A\epsilon}{d} \right)$ with $\lambda = \frac{\gamma}{10d^r}$, the left-hand side in Equation (C.9) is bounded by $C\epsilon d^r$

$$\frac{\hat{L}^{1+\iota/4}\gamma}{\hat{L}^{1+\iota/4}\gamma}$$

for some $C(A, \overline{L}) > 0$.

By setting $\hat{L} = d^{1+(k-2)r}$, we can ensure that the term in Equation (C.9) vanishes as d goes to infinity. By summing over all times $t \ge 1$, and plugging Lemma 38, as well as Lemma 39 into the inequality given by Equations (C.6) and (C.7), we obtain that uniformly over $w_0 \in \mathbb{S}^{d-1}$, with probability $1 - o_d(1)$,

$$m_{t} \geq \frac{4}{5}m_{0} + \sum_{j=0}^{t-1} \frac{\eta a_{k} |m_{j}|^{k-1}}{4d} \left(1 - \frac{L_{j+1} \mathbf{1}_{\{L_{j+1} < \hat{L}\}}}{K\bar{L}\log d}\right) - \frac{\eta}{d} \sum_{j=0}^{t-1} \langle \nabla H^{j+1}(\boldsymbol{W}_{j}), \boldsymbol{e}_{1} \rangle, \quad \forall t \leq \tau_{\frac{\gamma}{2d^{r}}}^{-} \wedge \tau_{\nu}^{+}$$
(C.10)

Proposition 40. Let $\hat{L} = d^{1+(k-2)r}$. For $k \ge 2$, if $\alpha \eta^2 \le \epsilon$ for some $\epsilon > 0, \eta \le \overline{\eta}_d(k)$, and α is at most polynomial in d, then for every $\gamma > 0$,

$$\lim_{d \to \infty} \inf_{\boldsymbol{w}_0} \mathbb{P}_{\boldsymbol{w}_0} \left(\sum_{j=0}^{t-1} \frac{\eta a_k |m_j|^{k-1}}{4d} \left(1 - \frac{L_{j+1} \mathbf{1}_{\{L_{j+1} < \hat{L}\}}}{K\bar{L}\log d} \right) \ge -\frac{\gamma}{10d^r} + \sum_{j=0}^{t-1} \frac{\eta a_k |m_j|^{k-1}}{8d}, \quad \forall t \le M \right) = 1$$

Moreover, this limit holds uniformly over choices of $\epsilon > 0$.

г	-	-	
L			

Proof. The strategy of the proof is to define a submartingale, and control its conditional second moment, as well as its martingale increments, to use a martingale inequality due to Freedman. To ease notations, denote $L_{j+1}\mathbf{1}_{\{L_{j+1}<\hat{L}\}}$ by \tilde{L}_{j+1} . It suffices to prove the following: for $\hat{L} = d^{1+(k-2)r}$, and α at most polynomial in d with $\alpha \eta^2 \leq \epsilon$ for some $\epsilon > 0$, we have for every $\gamma > 0$,

$$\lim_{d \to \infty} \sup_{\boldsymbol{w}_0 \in \mathbb{S}^{d-1}} \mathbb{P}_{\boldsymbol{w}_0} \left(\inf_{t \le M} \sum_{j=0}^{t-1} \frac{\eta a_k |m_j|^{k-1}}{4d} \left(\frac{1}{2} - \frac{\tilde{L}_{j+1}}{K\bar{L}\log d} \right) < -\frac{\gamma}{10d^r} \right) = 0.$$

Fix any $w_0 \in \mathbb{S}^{d-1}$.

$$\mathbb{P}_{\boldsymbol{w}_{0}}\bigg(\inf_{t\leq M}\sum_{j=0}^{t-1}\frac{\eta a_{k}|m_{j}|^{k-1}}{4d}\bigg(\frac{1}{2}-\frac{\tilde{L}_{j+1}}{K\bar{L}\log d}\bigg)<-\frac{\gamma}{10d^{r}}\bigg)\\\leq M\sup_{t\leq M}\mathbb{P}_{\boldsymbol{w}_{0}}\bigg(\sum_{j=0}^{t-1}\frac{\eta a_{k}|m_{j}|^{k-1}}{4d}\bigg(\frac{1}{\log d}-\frac{\tilde{L}_{j+1}}{K\bar{L}\log d}\bigg)<-\frac{\gamma}{10d^{r}}\bigg).$$

Let $Z_t \coloneqq \sum_{j=0}^{t-1} \frac{\eta a_k |m_j|^{k-1}}{4d} \left(\frac{1}{\log d} - \frac{\tilde{L}_{j+1}}{K\bar{L}\log d} \right)$. Since for all j, m_j is \mathcal{F}_j -measurable and

$$\mathbb{E}[\tilde{L}_{j+1}|\mathcal{F}_j] \leq \sup_{\boldsymbol{w}} \mathbb{E}\left[\left|\frac{1}{\sqrt{d}}\nabla H(\boldsymbol{w})\right|^2\right] \leq \bar{L},$$

we have that for all $K \ge 1, Z_t$ is an \mathcal{F}_t -submartingale. The martingale increments are bounded as follows.

$$|Z_t - Z_{t-1}| = \frac{\eta a_k |m_{t-1}|^{k-1}}{4d} \left| \frac{1}{\log d} - \frac{\tilde{L}_t}{K\bar{L}\log d} \right| \le \frac{\eta a_k}{4d\log d} (1 \vee \frac{\hat{L}}{K\bar{L}}) \le \frac{\eta a_k}{4d\log d} \left(1 + \frac{\hat{L}}{\bar{L}}\right).$$

For the conditional variances, we have

$$\mathbb{E}[(Z_t - Z_{t-1})^2 | \mathcal{F}_{t-1}] = \left(\frac{\eta a_k}{4d \log d}\right)^2 \mathbb{E}\left[|m_{t-1}|^{2(k-1)} \left(1 - \frac{\tilde{L}_t}{K\bar{L}}\right)^2 \middle| \mathcal{F}_{t-1}\right]$$

$$\leq \left(\frac{\eta a_k}{4d \log d}\right)^2 \left(1 + \frac{1}{\bar{L}^2} \mathbb{E}\left[\tilde{L}_t^2 \middle| \mathcal{F}_{t-1}\right]\right) \leq \left(\frac{\eta a_k}{4d \log d}\right)^2 \left(1 + \frac{1}{\bar{L}^2} \mathbb{E}\left[L_t^2 \middle| \mathcal{F}_{t-1}\right]\right)$$

$$= \left(\frac{\eta a_k}{4d \log d}\right)^2 \left(1 + \frac{1}{\bar{L}^2} \left[\left|\frac{1}{\sqrt{d}} \nabla H^t(X_{t-1})\right|^4 \middle| \mathcal{F}_{t-1}\right]\right)$$

$$\leq \left(\frac{\eta a_k}{4d \log d}\right)^2 \left(1 + \frac{1}{\bar{L}}\right) \leq \left(\frac{\eta a_k}{4d \log d}\right)^2 \quad a.s.$$

Therefore, by Freedman's inequality:

$$\sup_{t \le M} \mathbb{P}_{\boldsymbol{w}_0} \left(\sum_{j=0}^{t-1} \frac{\eta a_k |m_j|^{k-1}}{4d} \left(\frac{1}{\log d} - \frac{\tilde{L}_{j+1}}{K\bar{L}\log d} \right) \ge -\frac{\gamma}{10d^r} \right) \le \exp\left(\frac{-\gamma/100d^{2r}}{M\left(\frac{\eta a_k}{4d\log d}\right)^2 + \frac{1}{3}\frac{\eta a_k}{4d\log d} \cdot \frac{\gamma}{10d^r} \left(1 + \frac{\hat{L}}{\bar{L}} \right)} \right)$$

Now there remains to choose \hat{L} appropriately, to ensure that the bound is $o(\frac{1}{M})$, so that the overall bound is o(1). This constrains us to take either $\epsilon = \Omega(d^{1-2r+\mu} \log d)$ or $\eta \hat{L} = \Omega(d^{1-r+\mu})$, for some $\mu > 0$. Assuming we take the largest $\eta = \Theta(d^{-r(k-2)}/\log d)$, taking $\hat{L} = d^{1+(k-2)r}$ ensures we have the right bound, irrespective of $\epsilon > 0$.

Proposition 41. If C_1 is as in Assumption B, for every $\lambda > 0$, we have

$$\sup_{T \le M} \sup_{\boldsymbol{w}_0 \in \mathbb{S}^{d-1}} \mathbb{P}_{\boldsymbol{w}_0}\left(\max_{t \le T} \frac{1}{\sqrt{T}} \Big| \sum_{j=0}^{t-1} \langle \nabla H^{j+1}(\boldsymbol{W}_j), \boldsymbol{e}_1 \rangle \Big| \ge \lambda \right) \le \frac{2C_1}{\lambda^2}, \quad (C.11)$$

Proof. Let $\tilde{M}_t = \frac{dM_t}{\eta}$. It is a martingale with variance

$$\sup_{\boldsymbol{w}_0 \in \mathbb{S}^{d-1}} \mathbb{E}_{\boldsymbol{w}_0} \left[\sum_{j=0}^{t-1} \left(\left\langle \nabla H^{j+1}(\boldsymbol{W}_j), \boldsymbol{e}_1 \right\rangle \right)^2 \right] \le t \sup_{\boldsymbol{w}} \mathbb{E} \left[\left(\left\langle \nabla H(\boldsymbol{w}), \boldsymbol{e}_1 \right\rangle \right)^2 \right] \le C_1 t.$$

By Doob's maximal inequality,

$$\sup_{\boldsymbol{w}_0} \mathbb{P}_{\boldsymbol{w}_0} \Big(\sup_{t \le T} \Big| \tilde{M}_t > \lambda \sqrt{T} \Big| \Big) \le \frac{2 \sup_{\boldsymbol{w}_0} \mathbb{E}_{\boldsymbol{w}_0} [\tilde{M}_t^2]}{\lambda^2 t} \le \frac{2C_1}{\lambda^2}.$$

Therefore, by Proposition 41, for all b > 0,

$$\sup_{T \le M} \sup_{\boldsymbol{w}_0 \in \mathbb{S}^{d-1}} \mathbb{P}_{\boldsymbol{w}_0}\left(\max_{t \le T} \frac{\eta}{d} \Big| \sum_{j=0}^{t-1} \langle \nabla H^{j+1}(\boldsymbol{W}_j), \boldsymbol{e}_1 \rangle \Big| \ge \frac{\eta b \sqrt{T}}{10d} \right) \le \frac{200C_1}{b^2 T}.$$
 (C.12)

We can now prove Proposition 35 by combining Equation (C.10), Proposition 40, and Proposition 41. For every $\gamma > 0$ and $\nu > \nu_0(\varrho, a_k, a_{k+1})$,

$$\lim_{d \to \infty} \inf_{\boldsymbol{w}_0 \in E_{\gamma/d^r}} \mathbb{P}_{\boldsymbol{w}_0} \Big(m_t \ge \frac{7}{10} m_0 + \sum_{j=0}^{t-1} \frac{\eta a_k |m_j|^{k-1}}{8d} - \frac{\eta}{d} \sum_{j=0}^{t-1} \langle \nabla H^{j+1}(\boldsymbol{W}_j), \boldsymbol{e}_1 \rangle, \quad \forall t \le \tau_{\frac{\gamma}{2d^r}}^- \wedge \tau_{\nu}^+ \wedge M \Big) = 1$$

Furthermore, if $D = D_d$, $\eta \leq \bar{\eta}_d(k)$ and \bar{t} are as in Proposition 35, for all $w_0 \in E_{\gamma/d^r}$, if $T \leq \bar{t}$, then

$$\frac{\eta D\sqrt{T}}{10d} \le \frac{\gamma}{10d^r} \le \frac{m_0}{10}.$$

Therefore, by applying the directional error martingale bound (Equation (C.12)) with b = D, we obtain the desired bound.

Proof. [Proof of Proposition 36] Observe the following discrete analogue of Bihari-LaSalle inequality: suppose that $(m_t)_{t\in\mathbb{N}}$ is a sequence satisfying, for some $k \ge 3$ and a, b > 0,

$$m_t \ge a + \sum_{j=0}^{t-1} b m_j^{k-1} \quad \text{then} \quad m_t \ge \frac{a}{(1-(k-2)ba^{k-2}t)^{1/(k-2)}}.$$

Therefore, for $a = \frac{m_0}{2}, b = \frac{\eta a_k}{8d}$, we obtain

$$m_t \ge \frac{m_0}{(1 - (k-2)\frac{\eta a_k}{8d}m_0^{k-2}t)^{1/(k-2)}} \eqqcolon g_k(t).$$

 $g_k(t) \ge \nu$ provided

$$\nu^{k-2} \left(1 - (k-2) \frac{\eta a_k}{8d} \left(\frac{\gamma}{d^r} \right)^{k-2} t \right) = o\left(\left(\frac{\gamma}{d^r} \right)^{k-2} \right).$$

Thus, for $t \ge t^* = \left\lceil \frac{8d}{\eta a_k (k-2)\gamma^{k-2}} d^{r(k-2)} \left(1 - \frac{K}{d^{r(k-2)}}\right) \right\rceil$, for K large enough, we get that $g_k(t) \ge \nu$. The only remaining thing is to ensure $t^* \le \bar{t} \land M$. $t^* \le \bar{t}$ always holds for $\eta \le \bar{\eta}(k)$. To ensure $t^* \le M = \alpha d$, we must have

$$\alpha = \Omega(d^{r(k-2)}/\eta).$$

By optimizing over the values for η , and taking into account the constraint that $\alpha \eta^2 \leq \epsilon = \mathcal{O}(d^{1+(k-3)r})$, we must pick $\eta = \Theta\left(\frac{d^{-r(k-2)}}{\log d}\right)$, and hence,

$$\alpha = \Omega(d^{2r(k-2)}\log d).$$

Corollary 42. Under Assumptions A and B, for $\eta = \overline{\eta}(k), \alpha \gtrsim d^{2r(k-2)-1} \log d$ there exists $\nu_0(a_k, a_{k+1})$, such that for all $\nu < \nu_0$, for every $\gamma > 0$, we have

$$\lim_{d \to \infty} \inf_{\boldsymbol{w}_0 \in E_{\gamma/d^r}} \mathbb{P}_{\boldsymbol{w}_0} \Big(\tau_{\nu}^+ \leq \bar{t} \wedge M \Big) = 1.$$

Hence the sample complexity $M = \Omega(d^{2r(k-2)+1} \log d)$ is sufficient for weak recovery of the single-index model.