

ON SINGLE INDEX MODELS BEYOND GAUSSIAN DATA

JOAN BRUNA^{*,*,†}, LOUCAS PILLAUD-VIVIEN^{*,†} AND AARON ZWEIG^{*}

ABSTRACT. Sparse high-dimensional functions have arisen as a rich framework to study the behavior of gradient-descent methods using shallow neural networks, showcasing their ability to perform feature learning beyond linear models. Amongst those functions, the simplest are single-index models $f(x) = \phi(x \cdot \theta^*)$, where the labels are generated by an arbitrary non-linear scalar link function ϕ applied to an unknown one-dimensional projection θ^* of the input data. By focusing on Gaussian data, several recent works have built a remarkable picture, where the so-called information exponent (related to the regularity of the link function) controls the required sample complexity. In essence, these tools exploit the stability and spherical symmetry of Gaussian distributions. In this work, building from the framework of [Ben Arous et al., 2021], we explore extensions of this picture beyond the Gaussian setting, where both stability or symmetry might be violated. Focusing on the planted setting where ϕ is known, our main results establish that Stochastic Gradient Descent can efficiently recover the unknown direction θ^* in the high-dimensional regime, under assumptions that extend previous works [Yehudai and Shamir, 2020, Wu, 2022].

CONTENTS

1. Introduction	2
2. Preliminaries and Problem Setup	3
3. Stochastic Gradient Descent under LPG	5
4. Typical cases of loss landscape with LPG property	7
5. Experiments	14
6. Conclusion and Perspectives	16
References	17
Appendix A. Proof of Lemma 2.1	20
Appendix B. Proofs on the SGD dynamics: Section 3	21
Appendix C. The LPG property in the symmetric case: proofs of Section 4.1	33
Appendix D. The LPG property in the non-symmetric case: proofs of Section 4.2	44

^{*} COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY

^{*} CENTER FOR DATA SCIENCE, NEW YORK UNIVERSITY

[†] CENTER FOR COMPUTATIONAL MATHEMATICS, FLATIRON INSTITUTE

1. INTRODUCTION

Over the past years, there has been sustained effort to enlarge our mathematical understanding of high-dimensional learning, particularly when using neural networks trained with gradient-descent methods — highlighting the interplay between algorithmic, statistical and approximation questions. An essential, distinctive aspect of such models is their ability to perform *feature learning*, or to extract useful low-dimensional features out of high-dimensional observations.

An appealing framework to rigorously analyze this behavior are sparse functions of the form $f(x) = \phi(\Theta_*^\top x)$, where the labels are generated by a generic non-linear, low-dimensional function $\phi : \mathbb{R}^k \rightarrow \mathbb{R}$ of linear features $\Theta_*^\top x$, with $\Theta_* \in \mathbb{R}^{d \times k}$ with $k \ll d$. While the statistical and approximation aspects of such function classes are by now well-understood [Barron, 1993, Bach, 2017], the outstanding challenge remains computational, in particular in understanding the ability of gradient-descent methods to succeed. Even in the simplest setting of *single-index models* ($k = 1$), and assuming that ϕ is known, the success of gradient-based learning depends on an intricate interaction between the data distribution $x \sim \nu$ and the ‘link’ function ϕ ; and in fact computational lower bounds are known for certain such choices [Yehudai and Shamir, 2020, Song et al., 2021, Goel et al., 2020, Diakonikolas et al., 2017, Shamir, 2018].

Positive results thus require to make specific assumptions, either about the data, or about the link function, or both. On one end, there is a long literature, starting at least with [Kalai and Sastry, 2009, Shalev-Shwartz et al., 2010, Kakade et al., 2011], that exploits certain properties of ϕ , such as invertibility or monotonicity, under generic data distributions satisfying mild anti-concentration properties [Soltanolkotabi, 2017, Frei et al., 2020, Yehudai and Shamir, 2020, Wu, 2022]. On the other end, by focusing on canonical high-dimensional measures such as the Gaussian distribution, the seminal works [Ben Arous et al., 2021, Dudeja and Hsu, 2018] built a harmonic analysis framework of SGD, resulting in a fairly complete picture of the sample complexity required to learn generic link functions ϕ , and revealing a rich asymptotic landscape beyond the proportional regime $n \asymp d$, characterized by the number of vanishing moments, or *information exponent* s of ϕ , whereby $n \asymp d^{s-1}$ samples are needed for recovery. Since then, several authors have built and enriched this setting to multi-index models [Abbe et al., 2022, Abbe et al., 2023, Damian et al., 2022, Arnaboldi et al., 2023], addressing the semi-parametric learning of the link function [Bietti et al., 2022], as well as exploring SGD-variants [Ben Arous et al., 2022, Barak et al., 2022, Berthier et al., 2023, Chen et al., 2023]. This harmonic analysis framework relies on two key properties of the Gaussian measure and their interplay with SGD: its spherical symmetry and its stability by linear projection. Together, they provide an optimization landscape that is well-behaved in the limit of infinite data, and enable SGD to escape the ‘mediocrity’ of initialisation, where the initial direction θ_0 , in the high-dimensional setting, has vanishingly small correlation $|\theta_0 \cdot \theta^*| \simeq 1/\sqrt{d}$ with the planted direction θ^* .

In this work, we study to what extent the ‘Gaussian picture’ is robust to perturbations, focusing on the planted setting where ϕ is known. Our motivation comes from the fact that real data is rarely Gaussian, yet amenable to being approximately Gaussian via CLT-type arguments. We establish novel positive results along two main directions: (i) when spherical symmetry

is preserved but stability is lost, and (ii) when spherical symmetry is lost altogether. In the former, we show that spherical harmonics can be leveraged to provide a benign optimization landscape for SGD under mild regularity assumptions, for initialisations that can be reached with constant probability with the same sample complexity as in the Gaussian case. In the latter, we quantify the lack of symmetry with robust projected Wasserstein distances, and show that for ‘quasi-symmetric’ measures with small distance to the Gaussian reference, SGD efficiently succeeds for link functions with information exponent $s \leq 2$. Finally, using Stein’s method, we address substantially ‘non-symmetric’ distributions, demonstrating the strength and versatility of the harmonic analysis framework.

2. PRELIMINARIES AND PROBLEM SETUP

The focus of this work is to understand regression problems with input/output data $(x, y) \in \mathbb{R}^d \times \mathbb{R}$ generated by *single-index* models. This is a class of problems where the data labels are produced by a non-linear map of a one-dimensional projection of the input, that is

$$(1) \quad y = \phi(x \cdot \theta^*),$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is also known as the *link function*, and $\theta^* \in \mathcal{S}_{d-1}$, the sphere of $\mathbb{R}^d$, is the *hidden direction* that the models wants to learn. Quite naturally, the learning is made through the family of generalized linear predictors $\mathcal{H} = \{\phi_\theta : x \rightarrow \phi(x \cdot \theta), \text{ for } \theta \in \mathcal{S}_{d-1}\}$, built upon the link function (which is assumed known) and parametrized by the sphere.

Loss function. We assume that the input data is distributed according to a probability $\nu \in \mathcal{P}(\mathbb{R}^d)$. Equation (1) implies that the target function that produces the labels, ϕ_{θ^*} , lies in this parametric class. The overall loss classically corresponds to the average over all the data of the square penalisation $l(\theta, x) := (\phi_\theta(x) - \phi_{\theta^*}(x))^2$ so that the population loss writes

$$(2) \quad L(\theta) := \mathbb{E}_\nu \left[(\phi(x \cdot \theta) - \phi(x \cdot \theta^*))^2 \right] = \|\phi_\theta - \phi_{\theta^*}\|_{L_\nu^2}^2,$$

where we used the notation $\|f\|_{L_\nu^p}^p = \mathbb{E}_\nu[|f|^p]$, valid for all $p \in \mathbb{N}^*$. Let us put emphasis on the fact that the loss L is a non-convex function of the parameter θ , hence it is not *a priori* guaranteed that gradient-based method are able to retrieve the ground-truth θ^* . This often requires a precise analysis of the *loss landscape*, and where the high-dimensionality can play a role of paramount importance: we place ourselves in this high-dimensional setting for which the dimension is fixed but considered very large $d \gg 1$. Finally, we assume throughout the article that ϕ_θ belongs to the weighted Sobolev space $W_\nu^{1,4} := \{\phi, \text{ such that } \sup_{\theta \in \mathcal{S}_{d-1}} [\|\phi_\theta\|_{L_\nu^4} + \|\phi'_\theta\|_{L_\nu^4}] < \infty\}$.

Stochastic gradient descent. To recover the signal given by $\theta^* \in \mathcal{S}_{d-1}$, we run *online stochastic gradient descent* (SGD) on the sphere $\mathcal{S}_{d-1}$. This corresponds to having at each iteration $t \in \mathbb{N}^*$ a *fresh sample* x_t drawn from ν and independent of the filtration $\mathcal{F}_t = \sigma(x_1, \dots, x_{t-1})$ and performing a spherical gradient step, with step-size $\delta > 0$, with respect to $\theta \rightarrow l(\theta, x_t)$:

$$(3) \quad \theta_{t+1} = \frac{\theta_t - \delta \nabla_\theta^S l(\theta_t, x_t)}{|\theta_t - \delta \nabla_\theta^S l(\theta_t, x_t)|}, \quad \text{with initialization } \theta_0 \sim \text{Unif}(\mathcal{S}_{d-1}),$$

An important scalar function that enables to track the progress of the SGD iterates is the correlation with the signal $m_\theta := \theta \cdot \theta^* \in [-1, 1]$. We will drop the subscript in case there is no ambiguity. Note that, due to the high-dimensionality of the setup, we have the following lemma:

Lemma 2.1. *For all $a > 0$, we have $\mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) \leq a^{-1}e^{-a^2/4}$. Additionally, for any $\delta > 0$ such that $\max\{a, \delta\} \leq \sqrt{d}/4$, we have the lower bound: $\mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) \geq \frac{\delta}{4}e^{-(a+\delta)^2}$.*

This fact implies that, when running the algorithm in practice, it is initialized with high probability near the equator of $\mathcal{S}_{d-1}$, or at least in a band of typical size $1/\sqrt{d}$ (see Figure 1 for a schematic illustration of this fact). Finally, we use the notation ∇_θ^S to denote the spherical gradient, that is $\nabla_\theta^S l(\theta, x) = \nabla_\theta l(\theta, x) - (\nabla_\theta l(\theta, x) \cdot \theta)\theta$. As $\nabla_\theta^S l(\cdot, x_t)$ is an unbiased estimate of $\nabla_\theta^S L$, it is expected that the latter gradient field rules how the SGD iterates travel across the loss landscape.

Loss landscape in the Gaussian case. As stressed in the introduction, this set-up has been studied by [Dudeja and Hsu, 2018, Ben Arous et al., 2021] in the case where ν is the standard Gaussian, noted as γ here to avoid any confusion for later. Let us comment a bit this case to understand what can be the typical landscape of this single-index problem. Thanks to the spherical symmetry, the loss admits a scalar summary statistic, given precisely by the correlation m_θ . Moreover, the loss admits an explicit representation in terms of the Hermite decomposition of the link function ϕ : if $\{h_j\}_j$ denotes the orthonormal basis of Hermite polynomials of L_γ^2 , then $L(\theta) = 2 \sum_j |\langle \phi, h_j \rangle|^2 (1 - m^j) := \bar{\ell}(m)$. As a result, the gradient field projected along the signal is a (locally simple) *positive function* of the correlation that behaves similarly to

$$(4) \quad -\nabla_\theta^S L(\theta) \cdot \theta^* \simeq C m^{s-1} (1 - m),$$

where $s \in \mathbb{N}^*$ is the index of the first non-zero of the Hermite coefficients $\{\langle \phi, h_j \rangle\}_j$. This has at least three important consequences for the gradient flow: (i) if initialized positively, the correlation is an increasing function along the dynamics and there is no bad local minima in the loss landscape, (ii) the parameter $s \in \mathbb{N}^*$ controls the flatness of the loss landscape near the origin and therefore controls the optimization speed of SGD in this region (iii) as soon as m is large enough, the contractive term $1 - m$ makes the dynamics converge exponentially fast.

Loss landscape in general cases. Obviously for general distributions ν , the calculation presented in Eq.(4) is no-longer valid. However, the crux of the present paper is that properties (i)-(ii)-(iii) are robust to the change of distribution and can be shown to be preserved under small adaptations. More precisely, we have the following definition.

Definition 2.2 (Local Polynomial Growth). *We say that L has the local polynomial growth of order $k \in \mathbb{N}^*$ and scale $b \geq 0$, if there exists $C > 0$ such that for all $m_\theta \geq b$,*

$$(5) \quad -\nabla_\theta^S L(\theta) \cdot \theta^* \geq C(1 - m_\theta) (m_\theta - b)^{k-1}.$$

In such a case we say that L satisfies LPG(k, b).

In this definition, and as showed later in specific examples given in Section 4, the scale parameter b should be thought as a small parameter proportional to $1/\sqrt{d}$. If ν is Gaussian, we can rewrite Eq.(4) and show that L verifies LPG($s, 0$) for $s \in \mathbb{N}^*$, referred to as the *information exponent*

of the problem in [Ben Arous et al., 2021]. An important consequence of satisfying $\text{LPG}(k, b)$ is that the population landscape is free of bad local minima outside the equatorial band $\Sigma_b := \{\theta \in \mathcal{S}_{d-1}, m_\theta \leq b\}$. Therefore, when b is of scale $1/\sqrt{d}$, Lemma 2.1 indicates that one can efficiently produce initializations that avoid it. Hence, this property is the fundamental ingredient that enables the description of the path taken by SGD that we derive it in the next Section. Section 4 is devoted to showcasing generic examples when this property is satisfied.

3. STOCHASTIC GRADIENT DESCENT UNDER LPG

In this section, we derive the main results on the trajectory of the stochastic gradient descent. They state that the property $\text{LPG}(s, b/\sqrt{d})$ is in fact *sufficient* to recover the same quantitative guarantees as the one depicted in [Ben Arous et al., 2021], despite the lack of Gaussianity of the distribution ν . Recall that the recursion satisfied by the SGD iterates is given by Eq.(3). To describe their movement, let us introduce the following notations: for all $t \in \mathbb{N}^*$, we denote the normalization by $r_t := |\theta_t - \delta \nabla_\theta^S \ell(\theta_t, x_t)|$ and the martingale induced by the stochastic gradient descent as $M_t := l(\theta_t, x_t) - \mathbb{E}_\nu[l(\theta_t, x)]$.

Moment growth assumptions. To be able to analyse the SGD dynamics, we make the following assumptions on the moments of the martingale increments induced by the random sampling. To shorten notations, let us denote for all $\theta \in \mathbb{R}^d$, $x_\theta = x \cdot \theta \in \mathbb{R}$ and $C(u, v) = \phi'(u)\phi(v)$, for $u, v \in \mathbb{R}$.

Assumption 3.1 (Moment Growth Assumption). *There exists a constant $K > 0$, independent of the dimension d , such that*

$$(6) \quad \sup_{\theta \in \mathcal{S}_{d-1}} \mathbb{E}_x [x_{\theta_*}^2 C^2(x_\theta, x_{\theta_*})] \vee \mathbb{E}_x [x_\theta^2 C^2(x_\theta, x_{\theta_*})] \leq K, \quad \text{and ,}$$

$$(7) \quad \sup_{\theta \in \mathcal{S}_{d-1}} \mathbb{E}_x [|x|^{2k} C^2(x_\theta, x_{\theta_*})] \leq K d^k, \quad \text{for } k = 1, 2.$$

A precise care is given to the dependency in the dimension in the upper bound to match the practical cases that we later discuss in Section 4. Note that these assumptions are typically true for sub-gaussian random variables if ϕ belongs to a Sobolev regularity class. These assumptions are similar to the one given in Eqs. (1.3)-(1.4) in [Arous et al., 2021] for the Gaussian case. In all the remainder of the section we assume that Assumption 3.1 is satisfied.

Tracking the correlation. Recall that the relevant signature of the dynamics is the one-dimensional correlation $m_t = \theta_t \cdot \theta^*$. For infinitesimal step-sizes $\delta \rightarrow 0$, it is expected that θ_t follow the spherical gradient flow $\dot{\theta}_t = -\nabla_\theta^S L(\theta_t)$, that translates naturally on the summary statistics m_t as the following time evolution

$$(8) \quad \dot{m}_t = -\nabla_\theta^S L(\theta_t) \cdot \theta^*.$$

The main idea behind the result of this section is to show that, even if the energy landscape near $m = m_0$ is rough at scale $1/\sqrt{d}$, the noise induced by SGD does not prevent m to grow as the

idealized dynamics described by the ODE (8). Let us write the iterative recursion followed by $(m_t)_{t \geq 0}$: with the notation recalled above, for $t \in \mathbb{N}^*$, we have

$$(9) \quad m_{t+1} = \frac{1}{r_t} \left(m_t - \delta \nabla_{\theta}^S L(\theta_t) \cdot \theta^* - \delta \nabla_{\theta}^S M_t \cdot \theta^* \right).$$

With this dynamics at hand, the proof consists in controlling both the discretization error through r_t and the directional martingale induced by the term $\delta \nabla_{\theta}^S M_t \cdot \theta^*$.

Weak recovery. As it is the case for the gradient flow, most of the time spent by the SGD dynamics is near the equator, or more precisely in a band of the type $\Sigma_{b,c} = \{\theta \in \mathcal{S}_{d-1}, b/\sqrt{d} \leq m_{\theta} \leq c/\sqrt{d}\}$, where $b < c$ are constants independent of the dimension. Hence, the real first step of the dynamics is to go out any of these bands. This is the reason why it is natural to define $S_a := \{\theta \in \mathcal{S}_{d-1}, m_{\theta} \geq a\}$, the spherical cap of level $a \in (0, 1)$ as well as the hitting time

$$(10) \quad \tau_a^+ := \inf\{t \geq 0, m_{\theta_t} \geq a\},$$

which corresponds to the first time $(\theta_t)_{t \geq 0}$ enters in S_a . We arbitrarily choose a numerical constant independent of dimension, say $a = 1/2$, and refer to the related hitting time $\tau_{1/2}^+$ as the *weak recovery time* of the algorithm.

Theorem 3.2 (Weak Recovery). *Let $(\theta_t)_{t \geq 0}$ follow the SGD dynamics of Eq.(3) and let L satisfy $\text{LPG}(s, b/\sqrt{d})$, with $b > 0$ and $s \in \mathbb{N}^*$, then, conditionally on the fact that $m_0 \geq 5b/\sqrt{d}$, for any $0 < \varepsilon \leq \varepsilon_*$, we have*

$$(11) \quad \tau_{1/2}^+ \leq \begin{cases} d \cdot K/\varepsilon & \text{when } s = 1, \text{ and with the choice } \delta = \varepsilon/d \\ d \log^2(d) \cdot K/\varepsilon & \text{when } s = 2, \text{ and with the choice } \delta = \varepsilon/(d \log d) \\ d^{s-1} \cdot K/\varepsilon & \text{when } s \geq 3, \text{ and with the choice } \delta = \varepsilon d^{-s/2} \end{cases}$$

with probability at least $1 - K\varepsilon$, for generic constants $K, \varepsilon_* > 0$ that depend solely on the link function ϕ and the distribution ν .

Let us comment on this result. It says that the integer s coming from the growth condition controls the hardness of exiting the equator of the sphere. Indeed, as can be seen in $\text{LPG}(s, b/\sqrt{d})$, the larger the s , the smaller the gradient projection is and hence the less information the SGD dynamics has to move from the initialization. This result can be seen as an extension of [Ben Arous et al., 2021, Theorem 1.3] valid only in the Gaussian case ($b = 0$). Furthermore, the Gaussian case shows that Theorem 3.2 is tight up to $\log(d)$ factors. Finally, note that the result is conditional to the fact that the initialization is larger than some constant factor of $1/\sqrt{d}$, which has at least constant probability to happen in virtue of Lemma 2.1. This probability can be lowered by any constant factor by sampling offline a constant factor (independent of d) of i.i.d. initializations and keeping the one that maximizes its correlation with θ^* . Finally, note that in all the cases covered by the analysis, it is possible to keep track of the constant C , and show that overall it depends only (i) on the property of ν w.r.t. the Gaussian on the one hand and (ii) on the Sobolev norm of the link function $\|\phi\|_{W_{\nu}^{1,4}}$ on the other hand.

Strong recovery. We place ourselves *after* the weak recovery time described in the previous question and want to understand if the $(\theta_t)_{t \geq 0}$ dynamics goes to θ^* and if yes, how fast it does so. This is what we refer to as *the strong recovery* question, captured by the fact that the one-dimensional summary statistics m go towards 1. Thanks to the Markovian property of the SGD dynamics, we have the equality between all time $s > 0$ marginal laws of

$$\left(\theta_{\tau_{1/2}^+ + s} \mid \tau_{1/2}^+, \theta_{\tau_{1/2}^+} \right) \stackrel{\text{Law}}{=} \left(\theta_s \mid \theta_s = \theta_{\tau_{1/2}^+} \right),$$

and hence the strong recovery question is equivalent to study the dynamics with initialization that has already weakly recovered the signal, i.e. such that $m_\theta = 1/2$. We show that this part of the loss landscape is very different that the equator band in which the dynamics spends most of its times: in all the cases, we can choose stepsizes independent of the dimension and show that the time to reach the vicinity of θ^* will be independent of d .

Theorem 3.3 (Strong Recovery). *Let $(\theta_t)_{t \geq 0}$ follow the SGD dynamics of Eq.(3) and let L satisfy $\text{LPG}(s, b/\sqrt{d})$, with $b > 0$ and $s \in \mathbb{N}^*$, then, for any $\varepsilon > 0$, taking $\delta = \varepsilon/d$, we have that there exists a time $T > 0$, such that*

$$(12) \quad |1 - m_T| \leq \varepsilon, \text{ and } |T - \tau_{1/2}^+| \leq Kd \log(1/\varepsilon) \varepsilon^{-1}$$

with probability at least $1 - K\varepsilon$, for some generic $K > 0$ that depends solely of the link function ϕ .

As introduced above, the important messages conveyed by this theorem are that (i) there is no difference between the different parameters setups captured by the information exponent s , and (ii) the time it takes to reach an ε -vicinity of θ^* is always strictly smaller than the one needed to exit the *weak recovery phase* (e.g. d compared to d^{s-1} when $s \geq 3$). This means that *the dynamics spends most of its time escaping the mediocrity*. Remark that we decided to present Theorem 3.3 resetting the step-size δ to put emphasis on the intrinsic difference between the two phases. Yet, we could have kept the same stepsize as in the weak recovery case: this obviously would slow down unnecessarily the second phase.

4. TYPICAL CASES OF LOSS LANDSCAPE WITH LPG PROPERTY

In this section, we showcase two prototypical cases where the LPG holds true: the section 4.1 deals with the spherically symmetric setting, whereas the section 4.2 describes a perturbative regime where the distribution is *approximately* Gaussian in a quantitative sense.

4.1. The symmetric case. We start our analysis with the spherically symmetric setting. We show that a spherical harmonic decomposition provides a valid extension of the Hermite decomposition in the Gaussian case, leading to essentially the same quantitative performance up to constant (in dimension) factors.




FIGURE 1. Sketch of the SGD dynamics. After a long time spent in a band of typical size $1/\sqrt{d}$, the dynamics escapes at *weak recovery* point θ_w and then goes rapidly to the ground-truth θ^* .

Spherical Harmonic Representation of the Population Loss. We express the data distribution ν as a mixture of uniform measures $\nu = \int_0^\infty \tau_{r,d} \rho(dr)$, where $\tau_{r,d} = \text{Unif}(r\mathcal{S}_{d-1})$. Let $\tau_d = \tau_{1,d}$ and $u_d \in \mathcal{P}([-1, 1])$ be the projection of τ_d onto one direction, with density given in close form by $u_d(dt) = Z^{-1}(1 - t^2)^{(d-3)/2} \mathbf{1}(|t| \leq 1) dt$, where Z is a normalizing factor. Let $\{P_{j,d}\}_{j \in \mathbb{N}}$ be the orthogonal basis of Gegenbauer polynomials of $L^2_{u_d}([-1, 1])$, normalized such that $P_{j,d}(1) = 1$ for all j, d . For each $r > 0$, consider

$$l_r(\theta) := \langle \phi_\theta, \phi_{\theta^*} \rangle_{\tau_{r,d}} = \langle \phi_\theta^{(r)}, \phi_{\theta^*}^{(r)} \rangle_{\tau_d},$$

where we define $\phi^{(r)} : [-1, 1] \rightarrow \mathbb{R}$ such that $\phi^{(r)}(t) := \phi(rt)$. We write its decomposition in $L^2_{u_d}([-1, 1])$ as $\phi^{(r)} = \sum_j \alpha_{j,r,d} P_{j,d}$, with $\alpha_{j,r,d} = \frac{\langle \phi^{(r)}, P_{j,d} \rangle}{\|P_{j,d}\|^2}$. Let Ω_d be the Lebesgue measure of $\mathcal{S}_{d-1}$, and $N(j, d) = \frac{2d+j-2}{d} \binom{d+j-3}{d-1}$ the so-called *dimension* on the spherical harmonics of degree j in dimension d . From the Hecke-Funk representation formula [Frye and Efthimiou, 2012, Lemma 4.23] and the chosen normalization $P_{j,d}(1) = 1$, we have $\|P_{j,d}\|_2 = \left(\frac{\Omega_{d-1}}{\Omega_{d-2} N(j, d)} \right)^{1/2}$ [Frye and Efthimiou, 2012, Proposition 4.15] and obtain finally

$$(13) \quad l_r(\theta) = \sum_j \bar{\alpha}_{j,r,d}^2 P_{j,d}(\theta \cdot \theta^*),$$

where we defined for convenience $\bar{\alpha}_{j,r,d} = \alpha_{j,r,d} / \sqrt{N(j, d)}$. As a result, it follows that the overall loss writes as solely the correlaton $m_\theta = \theta \cdot \theta^*$ as

$$(14) \quad L(\theta) = 2\|\phi\|_{L^2_\nu}^2 - 2 \sum_j \beta_{j,d} P_{j,d}(m_\theta) := \ell(m),$$

where $\beta_{j,d} = \int_0^\infty \bar{\alpha}_{j,r,d}^2 \rho(dr) \geq 0$. Unsurprisingly, we observe that, thanks to the spherical symmetry and analogous to the Gaussian case, the loss still admits $m = \theta \cdot \theta^*$ as a summary statistics. Yet, it is represented in terms of Gegenbauer polynomials, rather than monomials

as in the Gaussian case. The monomial representation is a consequence of the stability of the Gaussian measure, as seen by the fact that $\langle \phi_\theta, \phi_{\bar{\theta}} \rangle_{\gamma_d} = \langle \phi, A_{\theta, \bar{\theta}} \phi \rangle_\gamma$, where $(A_m f)(t) = \mathbb{E}_{z \sim \gamma}[f(mt + \sqrt{1 - m^2}z)]$ has a *semi-group* structure (it is even known in fact as the Ornstein-Uhlenbeck semi-group).

Let $\eta \in \mathcal{P}(\mathbb{R})$ be the marginal of ν along any direction. The following proposition gives a closed form formula of the coefficients $\beta_{j,d}$, represented as integrals over the radial distribution ρ and projections of the link function ϕ :

Proposition 4.1 (Loss representation). *The $\beta_{j,d}$ defined in (14) have the integral representation*

$$(15) \quad \beta_{j,d} = \langle \phi, \mathcal{K}_j \phi \rangle_{L_\eta^2},$$

where $\mathcal{K}_j$ is a positive semi-definite integral operator of L_η^2 that depends solely on ρ and ϕ .

Note that a closed form expression of $\mathcal{K}_j$ can be found in Appendix C.2. The above proposition is in fact the stepping stone to calculate properly the *information exponent* that plays a crucial role in the property LPG. This is given through the link between the spectrum $\beta_{j,d}$ and the decomposition of ϕ in the L_η^2 orthogonal basis of polynomials, that we denote by $\{q_j\}_j$.

Proposition 4.2. *Let $s = \inf\{j; \beta_{j,d} > 0\}$ and $\tilde{s} = \inf\{j; \langle \phi, q_j \rangle_\eta \neq 0\}$. Then $s \leq \tilde{s}$.*

Thus, the number of vanishing moments of ϕ with respect to the data marginal η provides an upper bound on the ‘effective’ information exponent of the problem s , as we will see next.

Local Polynomial Growth. From (14), and as $\nabla_\theta^S L(\theta) = \ell'(m)(\theta^* - m\theta)$, we directly obtain

$$(16) \quad -\nabla_\theta^S L(\theta) \cdot \theta^* = -(1 - m^2)\ell'(m) = 2(1 - m^2) \sum_j \beta_{j,d} P'_{j,d}(m),$$

which is the quantity we want to understand to exhibit the property LPG in this case. Hence, we now turn into the question of obtaining sufficient guarantees on the coefficients $(\beta_{j,d})_j$ that ensure local polynomial growth. Since the typical scale of initialization for m is $\Theta(1/\sqrt{d})$, our goal is to characterize sufficient conditions of local polynomial growth with $b = O(1/\sqrt{d})$.

For that purpose, let us define two key quantities of Gegenbauer polynomials:

$$(17) \quad v_{j,d} := - \min_{t \in (0,1)} P_{j,d}(t), \quad (\text{smallest value})$$

$$(18) \quad z_{j,d} := \arg \max \{t \in (0,1); P_{j,d}(t) = 0\}. \quad (\text{largest root})$$

We have the following sufficient condition based on the spectrum $(\beta_{j,d})_j$:

Proposition 4.3 (Spectral characterization of LPG). *Suppose there exist constants $K, C > 0$ and $s \in \mathbb{N}$ such that we both have $\beta_{s,d} \geq C$ and $\sum_{j>s} \beta_{j,d} j(j+d-2)v_{j-1,d+2} \leq Kd^{(3-s)/2}$. Then, taking s^* as the infimum of such s , L has the property $\text{LPG}(s^* - 1, z_{s^*,d})$. In particular, whenever $s^* \ll d$, we have $z_{s^*,d} \leq 2\sqrt{s^*/d}$.*

This proposition thus establishes that, modulo a mild regularity assumption expressed through the decay of the coefficients $\beta_{j,d}$, the spherically symmetric non-Gaussian setting has the same

geometry as the Gaussian setting, for correlations *slightly* above the equator. Crucially, the required amount of correlation to ‘feel’ the local polynomial growth is a $O(\sqrt{s})$ factor from the typical initialization, and can be thus obtained with probability $\simeq e^{-s}$ over the initialization, according to Lemma 2.1, a lower bound which is *independent of d* .

The sufficient condition for $\beta_{j,d}$ appearing in Proposition 4.3 involves the minimum values $v_{j,d}$ of Gegenbauer polynomials $P_{j,d}$, as well as sums of the form $\sum_j j^2 \beta_{j,d}$. In order to obtain a more user-friendly condition, we now provide an explicit control of $v_{j,d}$, and leverage mild regularity of ϕ to control $\sum_j j^2 \beta_{j,d}$. This motivates the following assumption on ϕ and ν :

Assumption 4.4. *The link function ϕ satisfies $\phi \in L_\eta^2$ and $\phi' \in L_\eta^4$, and the radial distribution ρ has finite fourth moment $\mathbb{E}_\rho[r^4] < \infty$.*

Theorem 4.5 (LPG for symmetric distributions). *Assume that ϕ and ν satisfy Assumption 4.4, and let $s^* = \inf\{j; \langle \phi, q_j \rangle_\eta \neq 0\}$. Then L has the property $\text{LPG}(s^* - 1, 2\sqrt{s^*/d})$.*

The proof is provided in Appendix C.5. At a technical level, the main challenge in proving Theorem 4.5 is to achieve a uniform control of $v_{j,d}$ in j , a result which may be of independent interest. We address it by combining state-of-the-art bounds on the roots of the Gegenbauer polynomials, allowing us to cover the regime where j is small or comparable to d , together with integral representations via the Cauchy integral formula, providing control in the regime of large j . On the other hand, we relate the sum $\sum_j j^2 \beta_{j,d}$ to a norm of ϕ' using a Cauchy-Schwartz argument, where we leverage the fourth moments from Assumption 4.4.

Remark 4.6. *Since we are in a setting where ϕ is known, an alternative to the original recovery problem from Eq (2) is to consider a pure Gegenbauer ‘student’ link function of the form $\tilde{\phi} = P_{s,d}$, where s is the information exponent from Proposition 4.2. Indeed, the resulting population loss $\tilde{L}(\theta) = \mathbb{E}[(\tilde{\phi}(x \cdot \theta) - \tilde{\phi}(x \cdot \theta^*))^2]$ satisfies the LPG property, as easily shown in Fact C.5.*

For the sake of completeness, we describe more precisely two concrete case studies below.

Example 4.7 (Uniform Measure on the Sphere). *When $\nu = \text{Unif}(\sqrt{d}S_{d-1})$, we have $\rho = \delta_{\sqrt{d}}$, and therefore $\beta_{j,d} = \bar{\alpha}_{j,\sqrt{d},d}^2$. In that case, the orthogonal polynomial basis $\{q_j(t)\}_j$ of L_η^2 coincides with the rescaled Gegenbauer polynomials, $q_j(t) = P_{j,d}(t/\sqrt{d})$. Consider now a link function ϕ with $s - 1$ vanishing moments with respect to L_η^2 , i.e. such that $\bar{\alpha}_{j,d} = \langle \phi, q_j \rangle_\eta = 0$ for $j < s$ and $\bar{\alpha}_{s,d} = \langle \phi, q_s \rangle_\eta \neq 0$; and with sufficient decay in the higher harmonics as to satisfy the bound on the sum presented in Proposition 4.3 (for example, $\phi(t) = q_s(t)$ trivially satisfies this condition). Then Proposition 4.3 applies and we conclude that the resulting population landscape satisfies $\text{LPG}(s - 1, O(\sqrt{s/d}))$.*

In [Yehudai and Shamir, 2020, Wu, 2022] it is shown that monotonically increasing link functions¹ lead to a benign population landscape, provided the data distribution ν satisfies mild anti-concentration properties. We verify that in our framework.

¹or link functions where their monotonic behavior dominates; see [Wu, 2022].

Example 4.8 (Non-decreasing ϕ). *Indeed, Proposition 4.3 is verified with $s = 1$, provided $\phi' \in L_\eta^2$. Indeed, if $\phi \neq 0$ is monotonic, then we have $\beta_1 = \langle \phi, \mathcal{K}_1 \phi \rangle_\eta = C_d (\mathbb{E}_\eta[t\phi(t)])^2 \neq 0$, since we can assume without loss of generality that $\phi(t) \geq 0$ for $t \geq 0$ and $\phi(t) \leq 0$ for $t \leq 0$.*

We emphasize that the results of [Yehudai and Shamir, 2020, Wu, 2022] extend beyond the spherically symmetric setting, which is precisely the focus of next section.

4.2. Non-Spherically Symmetric Case. We now turn to the setting where ν is no longer assumed to have spherical symmetry. By making further regularity assumptions on ϕ , our main insight is that distributions that are *approximately* symmetric (defined in an appropriate sense) still benefit from a well-behaved optimization landscape.

Two-dimensional Wasserstein Distance. When ν is not spherically symmetric, the machinery of spherical harmonics does not apply, and we thus need to rely on another structural property. Consider a centered and isometric data distribution $\nu \in \mathcal{P}_2(\mathbb{R}^d)$, i.e. such that $\mathbb{E}_\nu x = 0$ and $\Sigma_\nu = \mathbb{E}_\nu[xx^\top] = I_d$. We consider the two-dimensional 1-Wasserstein distance [Niles-Weed and Rigollet, 2022, Definition 1] –see also [Paty and Cuturi, 2019]– between a pair of distributions $\nu_a, \nu_b \in \mathcal{P}(\mathbb{R}^d)$, defined as

$$(19) \quad \widetilde{W}_{1,2}(\nu_a, \nu_b) := \sup_{P \in \text{Gr}(2,d)} W_1(P_\# \nu_a, P_\# \nu_b),$$

where the supremum runs for any two-dimensional subspace $P \in \text{Gr}(2, d)$, and $P_\# \nu \in \mathcal{P}(\mathbb{R}^2)$ is the projection (or marginal) of ν onto the span of P . $\widetilde{W}_{1,2}$ is a distance ([Paty and Cuturi, 2019, Proposition 1]) and measures the largest 1-Wasserstein distance between any two-dimensional marginals.

We are in particular interested in the setting where $\nu_a = \nu$ is our data distribution, and ν_b is a reference symmetric measure – for instance the standard Gaussian measure γ_d . Consider the fluctuations

$$(20) \quad \Delta_L(\theta) := |L(\theta) - \bar{\ell}(m_\theta)| \text{ and } ,$$

$$(21) \quad \Delta_{\nabla L}(\theta) := |\nabla_\theta^S L(\theta) \cdot \theta^* - \bar{\ell}'(m_\theta)(1 - m_\theta^2)| ,$$

where $\bar{\ell}(m)$ is the Gaussian loss defined in Section 2 and $L(\theta) = \mathbb{E}_\nu[|\phi(x \cdot \theta) - \phi(x \cdot \theta^*)|^2]$. Δ_L and $\Delta_{\nabla L}$ thus measure respectively the fluctuations of the population loss and the relevant (spherical) gradient direction. By making additional mild regularity assumptions on the link function ϕ , we can obtain a uniform control of the population loss geometry using the dual representation of the 1-Wasserstein distance.

Assumption 4.9 (Regularity of link function). *We assume that ϕ, ϕ' are both B -Lipschitz, and that $\phi''(t) = O(1/t)$.*

Assumption 4.10 (Subgaussianity). *The data distribution ν is M -subgaussian: for any $v \in \mathcal{S}_{d-1}$, we have $\|x \cdot v\|_{\psi_2} \leq M$, where $\|z\|_{\psi_2} := \inf\{t > 0; \mathbb{E}[\exp(z^2/t^2)] \leq 2\}$ is the Orlicz-2 norm.*

Proposition 4.11 (Uniform gradient approximation). *Under Assumptions 4.9 and 4.10, for all $\theta \in \mathcal{S}_{d-1}$,*

$$(22) \quad \Delta_{\nabla L}(\theta) = (1 - m^2)O\left(\widetilde{W}_{1,2}(\nu, \gamma) \log(\widetilde{W}_{1,2}(\nu, \gamma)^{-1})\right)$$

where the $O(\cdot)$ notation only hides constants appearing in Assumptions 4.9 and 4.10.

In words, the population gradient under ν is viewed as a perturbation of the population gradient under γ , which has the well-behaved geometry already described in Section 2. These perturbations can be uniformly controlled by the projected 1-Wasserstein distance, thanks to the subgaussian tails of ν .

Our focus will be in situations where $\widetilde{W}_{1,2}(\nu, \gamma) = O(1/\sqrt{d})$. This happens to be the ‘natural’ optimistic scale for this metric in the class of isotropic distributions $\Sigma_\nu = I_d$, as can be seen for instance when $\nu = \text{Unif}(\sqrt{d}\mathcal{S}_{d-1})$. Under such conditions, it turns out that link functions with information exponent $s \leq 2$ can be recovered with simple gradient-based methods, by paying an additional polynomial (in d) cost in time complexity.

Assumption 4.12. *The Gaussian information exponent of ϕ , $s := \arg \min\{j; \langle \phi, H_j \rangle \neq 0\}$ satisfies $s \leq 2$.*

Assumption 4.13. *The projected Wasserstein distance satisfies $\widetilde{W}_{1,2}(\nu, \gamma) \leq M'/\sqrt{d}$.*

Proposition 4.14 (LPG, non-symmetric setting). *Under Assumptions 4.9, 4.10, 4.12 and 4.13, L verifies $\text{LPG}\left(1, O\left(\sqrt{\frac{(\log d^\kappa)}{d}}\right)\right)$, where κ depends only on B, M, M' .*

This proposition illustrates the cost of breaking spherical symmetry in two aspects: (i) it requires additional regularity on ϕ , and notably restricts its (Gaussian) information exponent to $s = 2$, and (ii) the scale to reach LPG is now no longer dimension-free, but has a polynomial dependency on dimension, since from Lemma 2.1, picking δ any positive constant we have

$$\mathbb{P}_{\theta_0}\left(m_{\theta_0} \geq \sqrt{\log d^\kappa}/\sqrt{d}\right) \geq \frac{\delta}{4}e^{-(\sqrt{\log d^\kappa} + \delta)^2} \geq \Omega\left(d^{-(1+o(1))\kappa}\right).$$

At present, we are not able to rule this out as a limitation of our proof; establishing whether this polynomial dependency on dimension is an inherent cost of the symmetry breaking is an interesting question for future work.

While assumptions 4.9, 4.10 and 4.12 are transparent and impose only mild conditions on the link function and tails of ν , the ‘real’ assumption of Proposition 4.14 is the concentration of $\widetilde{W}_{1,2}(\nu, \gamma)$ (Assumption 4.13). The ball $\{\nu; \widetilde{W}_{1,2}(\nu, \gamma) = O(1/\sqrt{d})\}$ contains many non-symmetric measures, for instance empirical measures sampled from γ with $n = \omega(d^2)$ [Niles-Weed and Rigollet, 2022, Proposition 8], and we suspect it contains many other examples, such as convolutions of the form $\nu * \gamma_\sigma$ arising for instance in diffusion models. That said, one should *not* expect the distance $\widetilde{W}_{1,2}(\nu, \gamma)$ to be of order $1/\sqrt{d}$ for generic ‘nice’ distributions

ν ; for instance, log-concave distributions are expected to satisfy $W_1(P_{\#}\nu, \gamma) \simeq 1/\sqrt{d}$ for *most* subspaces P , as captured in Klartag's CLT for convex bodies [Klartag, 2007]. In summary, many situations of interest fall outside this regime, which motivates us to relax the uniform Wasserstein criterion.

Localized Growth via Stein's method. To illustrate the mileage of the previous techniques beyond this 'quasi-symmetric' case, we consider now an idealised setting where the data is drawn from a product measure $\nu = \eta^{\otimes d}$, with $\eta \in \mathcal{P}(\mathbb{R})$ and $W_1(\eta, \gamma_1) = \Theta(1)$. In other words, $x = (x_1, \dots, x_d) \sim \nu$ if $x_i \sim \eta$ are i.i.d. In this setting, the distances $W_1(P_{\#}\nu, \gamma)$ reflect a CLT phenomena, which requires the subspace P to 'mix' across independent variables. Consequently, one may expect the expression of the hidden direction θ^* in the canonical basis to play a certain role. For that purpose, we make the following additional regularity assumption on the tails of ϕ to simplify the quantitative bounds:

Assumption 4.15 (Additional Regularity in third derivatives). *ϕ admits four derivatives bounded by L , with $|\phi^{(3)}(t)| = O(1/t)$ and $|\phi^{(4)}(t)| = O(1/t^2)$. Moreover, the third moment of the data distribution is finite: $\tau_3 = \mathbb{E}_{t \sim \eta}[t^3] < \infty$.*

Stein's method provides a powerful control on $\Delta_L(\theta)$ and $\Delta_{\nabla L}(\theta)$, as shown by the following result:

Proposition 4.16 (Stein's method for product measure). *Let $\chi(\theta, \theta^*) := \|\theta\|_4^2 + \|\theta^*\|_4^2$. Under Assumptions 4.9, 4.10 and 4.15, there exists a universal constant $C = C(M, B, \tau_3)$ such that*

$$(23) \quad \Delta_L(\theta) \leq C\chi(\theta, \theta^*), \text{ and } \Delta_{\nabla L}(\theta) \leq C\sqrt{1 - m^2}\chi(\theta, \theta^*).$$

The proof is based on the Stein's method for multivariate variables [Röllin, 2013, Theorem 3.1] with independent entries, which provides a quantitative CLT bound. Contrary to the quasi-symmetric case, here the concentration is not uniform over the sphere, but crucially depends on the sparsity of *both* θ and θ^* , measured via the ℓ_4 norms $\|\theta\|_4, \|\theta^*\|_4$: for incoherent, non-sparse directions, we have $\|\theta\|_4^2 \simeq 1/\sqrt{d}$, recovering the concentration rate that led to Proposition 4.14, while for sparse directions we have $\|\theta\|_4^2 = \Theta(1)$, indicating an absence of concentration to the Gaussian landscape.

Therefore, the natural conclusion is to assume a planted model where θ^* is *incoherent* with the data distribution, i.e. $\|\theta^*\|_4 = O(d^{1/4})$. While the LPG property does not directly apply in this setting, we outline an argument that suggests that the single-index model can still be efficiently solved using gradient-based methods. For that purpose, we assume that θ^* is drawn uniformly in $\mathcal{S}_{d-1}$, which implies that its squared- L^4 norm $\|\theta^*\|_4^2$ is of order $d^{-1/2}$ with high probability:

Fact 4.17. *Assume $\theta^* \sim \text{Unif}(\mathcal{S}_{d-1})$. Then $\mathbb{P}(\|\theta^*\|_4^2 \leq C/\sqrt{d}) \geq 1 - C' \exp(-C)$.*

Because θ_0 is also drawn uniformly on the sphere, the typical value of $\chi(\theta, \theta^*)$ is of order $d^{-1/2}$. For Gaussian information exponent $s = 2$, the population gradient under Gaussian data satisfies $-\nabla_{\theta}^{S_{d-1}} L_{\gamma}(\theta) \cdot \theta^* \geq Cm_{\theta}$. As a consequence, by Proposition 4.16, whenever $|\theta_0 \cdot \theta^*| > c\sqrt{d}$

(which happens with constant probability lower bounded by e^{-c}), we enter a ‘local’ LPG region where $-\nabla_{\theta}^{\mathcal{S}_{d-1}} L(\theta) \cdot \theta^* \geq C(m - c/\sqrt{d}) > 0$. While this condition is sufficient in the quasi-symmetric setting to start accumulating correlation (Theorem 3.2), now this event is conditional on θ being dense, ie so that $\chi(\theta, \theta^*) = O(1/\sqrt{d})$.

Since the typical value of $\chi(\theta, \theta^*)$ is of scale $1/\sqrt{d}$, one would expect that SGD will rarely visit sparse points where $\chi(\theta, \theta^*) \gg O(1/\sqrt{d})$, and thus that the local LPG property will be valid for *most* times during the entropic phase of weak recovery — and therefore that the correlation m_{θ} will pile-up as in the quasi-symmetric setting.

We summarise this property in the following conjecture:

Conjecture 4.18 (SGD avoids sparse points). *Assume θ^*, θ_0 are drawn from the uniform measure, and let θ_t be the t -th iterate of SGD with $\delta \simeq 1/(d \log d)$. There exists a universal constant C such that for any $\xi > 0$, we have*

$$(24) \quad \mathbb{P} \left(\sup_{t \leq T} \|\theta_t\|_4^2 \geq \sqrt{\frac{\xi \log T}{d}} \right) \leq C \exp(-\xi^2 d).$$

Since the time to escape mediocrity in the case $s = 2$ is $T \simeq d \log(d)^2$, this conjecture would imply that SGD does not effectively ‘see’ any sparse points, and thus escapes mediocrity. If one assumed that in this phase the dynamics is purely noisy, now pretending that θ_i were drawn independently from the uniform measure, and that $\|\theta\|_4^4$ is approximately Gaussian with mean d^{-1} and variance d^{-3} , the result follows by simple concentration. The challenging aspect of Conjecture 4.18 is precisely to handle the dependencies across iterates, as well as the spherical projection steps.

5. EXPERIMENTS

In order to validate our theory, and inspect the degree to which our bounds may be pessimistic, we consider empirical evaluation of the training process in our two primary settings. Specifically, we consider random initialization on the half-sphere (with the sign chosen to induce positive correlation as in [Arous et al., 2021]), and investigate how often strong recovery occurs relative to the information exponent of the link function.

Symmetric Case. For the spherically symmetric setting, we experiment with the input distribution that is uniform on the sphere. We are primarily interested in verifying that, unlike the Gaussian case, strong recovery depends on whether the initial correlation is sufficiently high to avoid local minima and benefit from the LPG guarantee. This is not evident in the 2nd degree Gegenbauer case, which is monotonic and quickly reaches strong recovery, but it is clear from the 4th degree Gegenbauer link function.

In the infinite sample setting, Figure 3 exactly characterizes the loss landscape when learning the 4th degree Gegenbauer under inputs uniform on $\mathcal{S}_{d-1}$ for different values of d . Note that the largest zero for $d = 50$ occurs at $\approx \pm 0.31$, and the loss is monotonic for m values initialized




FIGURE 2. Correlation with the true signal throughout training, under different choices of link function and input distribution.

outside that region. This phenomenon persists for higher dimensions, and one may observe that d increases, the critical points become smaller in magnitude, according to the scaling $\simeq \sqrt{1/d}$. The bottom right subplot in Figure 2 indicates training runs in this setting, where red lines are initialized uniformly on the sphere, and blue lines are initialized uniformly conditioned on $m = 0.4$, which is slightly past the last zero of the polynomial. We observe that random




FIGURE 3. The 4th degree Gegenbauer polynomial with dimension 50, 75, 100. Each is equivalent (up to rescaling) to the loss landscape of learning the 4th degree Gegenbauer in the appropriate dimension. Points indicate largest zeros.

initialization infrequently exceeds the threshold necessary for strong recovery, but planting the initialization above this threshold gives a high probability of recovery.

Non-Symmetric Case. For the non-spherically symmetric setting, we compare the performance of Gaussian inputs with inputs that are approximately Gaussian under a two-dimensional projection. For simplicity, we loosen our assumptions slightly, and consider the input distribution as the d dimensional product distribution of uniform random variables (rescaled to have unit variance), and allow for a non-Lipschitz link function. Here, we are primarily interested in whether Assumption 4.12 is tight and $s \leq 2$ is necessary for recovery, as well as whether Conjecture 4.18 holds in practice.

To evaluate, we compare strong recovery rates when training on a "tricky" function with $s = 2$ (chosen to be $\frac{1}{2}(h_2 - h_3 - h_4 + h_5)$) versus a function with $s = 3$ (simply the degree three hermite polynomial h_3). We make this choice for the $s = 2$ function in order to produce a function which is not monotonic, for which learning is easy under many distributions, as discussed in [Yehudai and Shamir, 2020].

In Figure 2 we observe that strong recovery reliably occurs for both the Gaussian and hypercube input distributions when $s = 2$. There is more variance in the Gaussian runs, likely because the magnitude of the gradients will be larger due to the inclusion of high degree terms. But for the $s = 3$ case, the Gaussian distribution converges quickly while the hypercube distribution frequently cannot escape the equator.

6. CONCLUSION AND PERSPECTIVES

In this work, we have asked whether the remarkable properties of high-dimensional Gaussian SGD regression of single-index models are preserved as one loses some key aspects that make Gaussian distributions so special (and so appealing for theorists). Our results are mostly positive, indicating a robustness of the Gaussian theory, especially within the class of spherically symmetric distributions, where a rich spherical harmonic structure is still available. As one loses

spherical symmetry, the situation becomes more dire, motivating a perturbative analysis that we have shown is effective via projected Wasserstein and Stein couplings.

That said, there are several open and relevant avenues that our work has barely touched upon, such as understanding whether the robustness can be transferred to other algorithms beyond SGD, or addressing the semi-parametric problem when the link function is unknown, along the lines of [Bietti et al., 2022, Abbe et al., 2023, Damian et al., 2022, Berthier et al., 2023]. A particularly interesting direction of future work is to extend the analysis of product measures to ‘weakly dependent’ distributions, motivated by natural images where locality in pixels captures most (but not all) of the statistical dependencies. Stein’s method appears to be a powerful framework that can accommodate such weak dependencies, and deserves future investigation.

Acknowledgements. This work was partially supported by NSF DMS 2134216, NSF CAREER CIF 1845360, NSF IIS 1901091 and the Alfred P Sloan Foundation.

REFERENCES

- [Abbe et al., 2022] Abbe, E., Boix-Adsera, E., and Misiakiewicz, T. (2022). The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. *arXiv preprint arXiv:2202.08658*. 2
- [Abbe et al., 2023] Abbe, E., Boix-Adsera, E., and Misiakiewicz, T. (2023). Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. *arXiv preprint arXiv:2302.11055*. 2, 17
- [Area et al., 2004] Area, I., Dimitrov, D. K., Godoy, E., and Ronveaux, A. (2004). Zeros of gegenbauer and hermite polynomials and connection coefficients. *Math. Comput.*, 73:1937–1951. 34
- [Arnaboldi et al., 2023] Arnaboldi, L., Stephan, L., Krzakala, F., and Loureiro, B. (2023). From high-dimensional & mean-field dynamics to dimensionless odes: A unifying approach to sgd in two-layers networks. *arXiv preprint arXiv:2302.05882*. 2
- [Arous et al., 2021] Arous, G. B., Gheissari, R., and Jagannath, A. (2021). Online stochastic gradient descent on non-convex losses from high-dimensional inference. *The Journal of Machine Learning Research*, 22(1):4788–4838. 5, 14, 22, 33
- [Bach, 2017] Bach, F. (2017). Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681. 2
- [Barak et al., 2022] Barak, B., Edelman, B., Goel, S., Kakade, S., Malach, E., and Zhang, C. (2022). Hidden progress in deep learning: Sgd learns parities near the computational limit. *Advances in Neural Information Processing Systems*, 35:21750–21764. 2
- [Barron, 1993] Barron, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945. 2
- [Ben Arous et al., 2021] Ben Arous, G., Gheissari, R., and Jagannath, A. (2021). Online stochastic gradient descent on non-convex losses from high-dimensional inference. *Journal of Machine Learning Research (JMLR)*, 22:106–1. 1, 2, 4, 5, 6, 23
- [Ben Arous et al., 2022] Ben Arous, G., Gheissari, R., and Jagannath, A. (2022). High-dimensional limit theorems for sgd: Effective dynamics and critical scaling. *arXiv preprint arXiv:2206.04030*. 2
- [Berthier et al., 2023] Berthier, R., Montanari, A., and Zhou, K. (2023). Learning time-scales in two-layers neural networks. *arXiv preprint arXiv:2303.00055*. 2, 17
- [Bietti et al., 2022] Bietti, A., Bruna, J., Sanford, C., and Song, M. J. (2022). Learning single-index models with shallow neural networks. In *Advances in Neural Information Processing Systems*. 2, 17
- [Chen et al., 2023] Chen, S., Dou, Z., Goel, S., Klivans, A. R., and Meka, R. (2023). Learning narrow one-hidden-layer relu networks. *arXiv preprint arXiv:2304.10524*. 2

- [Damian et al., 2022] Damian, A., Lee, J., and Soltanolkotabi, M. (2022). Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*. 2, 17
- [De Carli, 2008] De Carli, L. (2008). Local lp inequalities for gegenbauer polynomials. In *Topics in classical analysis and applications in honor of Daniel Waterman*, pages 73–87. World Scientific. 40
- [Diakonikolas et al., 2017] Diakonikolas, I., Kane, D. M., and Stewart, A. (2017). Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84. IEEE. 2
- [Dimitrov and Nikolov, 2010] Dimitrov, D. K. and Nikolov, G. P. (2010). Sharp bounds for the extreme zeros of classical orthogonal polynomials. *Journal of Approximation Theory*, 162(10):1793–1804. Special Issue dedicated to the memory of Borislav Bojanov. 41
- [DLMF, 2022] DLMF (2022). *NIST Digital Library of Mathematical Functions*. <https://dlmf.nist.gov/>, Release 1.1.10 of 2023-06-15. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds. 41
- [Driver and Jordaan, 2012] Driver, K. and Jordaan, K. (2012). Bounds for extreme zeros of some classical orthogonal polynomials. *Journal of Approximation Theory*, 164(9):1200–1204. 42
- [Dudeja and Hsu, 2018] Dudeja, R. and Hsu, D. (2018). Learning single-index models in gaussian space. In Bubeck, S., Perchet, V., and Rigollet, P., editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1887–1930. PMLR. 2, 4
- [Freedman, 1975] Freedman, D. A. (1975). On tail probabilities for martingales. *the Annals of Probability*, pages 100–118. 32
- [Frei et al., 2020] Frei, S., Cao, Y., and Gu, Q. (2020). Agnostic learning of a single neuron with gradient descent. *Advances in Neural Information Processing Systems*, 33:5417–5428. 2
- [Frye and Efthimiou, 2012] Frye, C. and Efthimiou, C. J. (2012). Spherical harmonics in p dimensions. *arXiv preprint arXiv:1205.3548*. 8, 33, 34
- [Gautschi, 1959] Gautschi, W. (1959). Some elementary inequalities relating to the gamma and incomplete gamma function. *J. Math. Phys*, 38(1):77–81. 20
- [Goel et al., 2020] Goel, S., Gollakota, A., Jin, Z., Karmalkar, S., and Klivans, A. (2020). Superpolynomial lower bounds for learning one-layer neural networks using gradient descent. In *International Conference on Machine Learning*, pages 3587–3596. PMLR. 2
- [Kakade et al., 2011] Kakade, S. M., Kanade, V., Shamir, O., and Kalai, A. (2011). Efficient learning of generalized linear and single index models with isotonic regression. *Advances in Neural Information Processing Systems*, 24. 2
- [Kalai and Sastry, 2009] Kalai, A. T. and Sastry, R. (2009). The isotron algorithm: High-dimensional isotonic regression. In *COLT*. 2
- [Klartag, 2007] Klartag, B. (2007). A central limit theorem for convex sets. *Inventiones mathematicae*, 168(1):91–131. 13
- [Laforgia and Natalini, 2013] Laforgia, A. and Natalini, P. (2013). On some inequalities for the gamma function. *Advances in Dynamical Systems and Applications*, 8(2):261–267. 20
- [Niles-Weed and Rigollet, 2022] Niles-Weed, J. and Rigollet, P. (2022). Estimation of wasserstein distances in the spiked transport model. *Bernoulli*, 28(4):2663–2688. 11, 12
- [Paty and Cuturi, 2019] Paty, F.-P. and Cuturi, M. (2019). Subspace robust wasserstein distances. In *International conference on machine learning*, pages 5072–5081. PMLR. 11
- [Röllin, 2013] Röllin, A. (2013). Stein’s method in high dimensions with applications. In *Annales de l’IHP Probabilités et statistiques*, volume 49, pages 529–549. 13, 47
- [Shalev-Shwartz et al., 2010] Shalev-Shwartz, S., Shamir, O., and Sridharan, K. (2010). Learning kernel-based halfspaces with the zero-one loss. *arXiv preprint arXiv:1005.3681*. 2
- [Shamir, 2018] Shamir, O. (2018). Distribution-specific hardness of learning neural networks. *The Journal of Machine Learning Research*, 19(1):1135–1163. 2

- [Soltanolkotabi, 2017] Soltanolkotabi, M. (2017). Learning relus via gradient descent. *Advances in neural information processing systems*, 30. 2
- [Song et al., 2021] Song, M. J., Zadik, I., and Bruna, J. (2021). On the cryptographic hardness of learning single periodic neurons. *Advances in Neural Processing Systems (NeurIPS)*. 2
- [Stam, 1982] Stam, A. J. (1982). Limit theorems for uniform distributions on spheres in high-dimensional euclidean spaces. *Journal of Applied probability*, 19(1):221–228. 20
- [Szego, 1939] Szego, G. (1939). *Orthogonal polynomials*, volume 23. American Mathematical Soc. 41
- [Ursell, 2007] Ursell, F. (2007). Integrals with nearly coincident branch points: Gegenbauer polynomials of large degree. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2079):697–710. 43
- [Watson, 1922] Watson, G. N. (1922). *A treatise on the theory of Bessel functions*, volume 3. The University Press. 43
- [Wu, 2022] Wu, L. (2022). Learning a single neuron for non-monotonic activation functions. In *International Conference on Artificial Intelligence and Statistics*, pages 4178–4197. PMLR. 1, 2, 10, 11
- [Yehudai and Shamir, 2020] Yehudai, G. and Shamir, O. (2020). Learning a single neuron with gradient methods. In *Conference on Learning Theory*, pages 3756–3786. PMLR. 1, 2, 10, 11, 16

APPENDIX

We gather in the appendix the proofs of the theorems, propositions and lemmas stated in the main text. In Section A, the reader will find a short proof of Lemma 2.1. In Section B, we prove Theorems 3.2 and 3.3 on the SGD dynamics. Sections C and D are respectively devoted to prove that the LPG property holds in some spherical symmetric case and under some perturbative regime.

APPENDIX A. PROOF OF LEMMA 2.1

Let us first recall the Lemma before writing a proof of it.

Lemma. *For all $a > 0$, we have $\mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) \leq a^{-1}e^{-a^2/4}$. Additionally, for any $\delta > 0$ such that $\max\{a, \delta\} \leq \sqrt{d}/4$, we have the lower bound: $\mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) \geq \frac{\delta}{4}e^{-(a+\delta)^2}$.*

Proof. By rotation invariance of the uniform distribution of the sphere, m_{θ_0} is distributed according to $\theta_0[1]$, the first coordinate of the vector $\theta_0 \in \mathcal{S}_{d-1}$. By a particular case of Stam's formula [Stam, 1982, relation (3)], we know that for $d \geq 3$, both are distributed according to the probability of density, $\forall t \in \mathbb{R}$,

$$\tau(t) := \frac{\Gamma(d/2)}{\sqrt{\pi}\Gamma((d-1)/2)} (1-t^2)^{(d-3)/2} \mathbb{1}_{[-1,1]}.$$

First, note that we can upper and lower bound the constant by the following:

$$\sqrt{\frac{d}{3}} \leq \frac{\Gamma(d/2)}{\Gamma((d-1)/2)} \leq \sqrt{\frac{d}{2}},$$

for $d \geq 6$ by [Laforgia and Natalini, 2013, equality 3.2], which was already proved in [Gautschi, 1959].

Hence, in terms of the upper bound, we have:

$$\begin{aligned} \mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) &\leq \sqrt{\frac{d}{2\pi}} \int_{a/\sqrt{d}}^1 (1-t^2)^{(d-3)/2} dt \\ &\leq \sqrt{\frac{d}{2\pi}} \int_{a/\sqrt{d}}^1 e^{-\frac{d-3}{2}t^2} dt \\ &\leq \frac{1}{\sqrt{2\pi}} \frac{d}{a} \int_{a/\sqrt{d}}^1 te^{-\frac{d-3}{2}t^2} dt \\ &\leq \frac{1}{2a} e^{-a^2/4}, \end{aligned}$$

which concludes the first part of the result.

Second, let $a \leq \sqrt{d}/4$ and take any $0 < \delta < \sqrt{d}/4$. We have,

$$\begin{aligned} \mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) &\geq \sqrt{\frac{d}{3\pi}} \int_{a/\sqrt{d}}^1 (1-t^2)^{(d-3)/2} dt \\ &\geq \sqrt{\frac{d}{3\pi}} \int_{a/\sqrt{d}}^{(a+\delta)/\sqrt{d}} (1-t^2)^{(d-3)/2} dt \\ &\geq \sqrt{\frac{d}{3\pi}} \frac{\delta}{\sqrt{d}} \left(1 - \frac{(a+\delta)^2}{d}\right)^{(d-3)/2}, \end{aligned}$$

where the last inequality simply comes from the fact that $t \rightarrow (1-t^2)^{(d-3)/2}$ is non-increasing. Going further, if we lower bound the term with the negative $-3/2$ power by 1, we have

$$\begin{aligned} \mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) &\geq \frac{\delta}{4} \exp\left(\frac{d}{2} \log\left(1 - \frac{(a+\delta)^2}{d}\right)\right) \\ &\geq \frac{\delta}{4} \exp\left(-\frac{(a+\delta)^2}{2(1 - (a+\delta)^2/d)}\right), \end{aligned}$$

where the last inequality come from the classical bound $\log(1+x) \geq x/(1+x)$, for $x > -1$. Furthermore, as, $(a+\delta)^2 \leq d/2$, we have finally

$$\mathbb{P}_{\theta_0}(m_{\theta_0} \geq a/\sqrt{d}) \geq \frac{\delta}{4} e^{-(a+\delta)^2},$$

which finalizes the proof of the Lemma. ■

APPENDIX B. PROOFS ON THE SGD DYNAMICS: SECTION 3

We first recall the notations useful to fully describe the dynamics. In Section B.3, we prove Theorem 3.2 about weak recovery. In Section B.4, we prove Theorem 3.3 about strong recovery. Finally,

B.1. Recalling the dynamics. For the sake of clarity, let us recall the notations and facts developed in the main text. The overall loss classically corresponds to the average over all the *data* of a square penalisation $l(\theta, x) = (\phi_\theta(x) - \phi_{\theta^*}(x))^2$ so that

$$L(\theta) = \mathbb{E}_\nu[(\phi_\theta(x) - \phi_{\theta^*}(x))^2].$$

To recover the signal given by θ^* , we run *online stochastic gradient descent* on the sphere $\mathcal{S}_{d-1}$. This corresponds to have at each iteration $t \in \mathbb{N}^*$ a *fresh sample* x_t independent of the filtration $\mathcal{F}_t = \sigma(x_1, \dots, x_{t-1})$ and perform a spherical gradient step, with step-size $\delta > 0$, with respect to $\theta \rightarrow l(\theta, x_t)$:

$$(25) \quad \theta_{t+1} = \frac{\theta_t - \delta \nabla_\theta^S l(\theta_t, x_t)}{|\theta_t - \delta \nabla_\theta^S l(\theta_t, x_t)|},$$

initialized at θ_0 uniformly on the sphere: $\theta_0 \sim \text{Unif}(\mathcal{S}_{d-1})$. Recall that we use the notation ∇_θ^S to denote the spherical gradient, that is

$$\nabla_\theta^S l(\theta, x) = \nabla_\theta l(\theta, x) - (\nabla_\theta l(\theta, x) \cdot \theta) \theta.$$

Let us introduce the following frequently used notations: for all $t \in \mathbb{N}^*$, we denote the normalization by $r_t := r(\theta_t, x_t) = |\theta_t - \delta \nabla_\theta^S l(\theta_t, x_t)|$ and the martingale induced by the stochastic gradient descent as $M_t = M(\theta_t, x_t) = l(\theta_t, x_t) - \mathbb{E}_\nu[l(\theta_t, x)]$.

B.2. Tracking the correlation. Recall that the relevant signature of the dynamics is the one-dimensional correlation: $m_t = \theta_t \cdot \theta^*$. Let us re-write the iterative recursion followed by $(m_t)_{t \geq 0}$, with the notation recalled above, for $t \in \mathbb{N}^*$,

$$(26) \quad m_{t+1} = \frac{1}{r_t} (m_t - \delta \nabla^S l(\theta_t, x_t) \cdot \theta^*) = \frac{1}{r_t} (m_t - \delta \nabla^S L(\theta_t) \cdot \theta^* - \delta \nabla^S M_t \cdot \theta^*).$$

We want to lower bound the right hand side of (26). We begin by a lower bound on $1/r_t$.

Lemma B.1 (Bound on r_t). *For all $t \in \mathbb{N}^*$, we have $1/r_t \geq 1 - \delta^2 |\nabla_\theta l(\theta_t, x_t)|^2$.*

Proof. For all $t \in \mathbb{N}^*$, we have, by orthogonality of θ_t and $\nabla_\theta^S l(\theta_t, x_t)$, that

$$r_t^2 = |\theta_t - \delta \nabla_\theta^S l(\theta_t, x_t)|^2 = 1 + \delta^2 |\nabla_\theta^S l(\theta_t, x_t)|^2 \leq 1 + \delta^2 |\nabla_\theta l(\theta_t, x_t)|^2.$$

Hence, from the inequality $(1 + u)^{-1/2} \geq 1 - u$ for all $u > 0$, we conclude the proof. $\blacksquare$

Thanks the fact that L satisfies $\text{LPG}(s, b/\sqrt{d})$, ie $-\nabla^S L(\theta) \cdot \theta^* \geq C(1 - m)(m - b/\sqrt{d})^{s-1}$, we have that the dynamics satisfies the following inequality between iterates:

$$(27) \quad m_{t+1} \geq m_t + C\delta(1 - m_t) \left(m_t - \frac{b}{\sqrt{d}}\right)^{s-1} - \delta \nabla^S M_t \cdot \theta^* - \delta^2 |m_t| |\nabla_\theta l(\theta_t, x_t)|^2 - \delta^3 \xi_t,$$

where $\xi_t = |\nabla_\theta l(\theta_t, x_t)|^2 |\nabla^S l(\theta_t, x_t) \cdot \theta^*|$. All the terms of the inequality have a natural origin: the second term is the ideal term coming from the gradient flow and the growth condition, the third term corresponds to the martingale increments coming from the noise induced by SGD and the two final terms are simply discretization errors coming from discrete nature of the procedure and the projection step.

However, to have a tight dependency with respect to the dimension, we need to be extra careful. This is why, following [Arous et al., 2021], we decompose this term introducing a threshold $M > 0$, to be fixed later, such that:

$$|m_t| |\nabla_\theta l(\theta_t, x_t)|^2 = |m_t| |\nabla_\theta l(\theta_t, x_t)|^2 \mathbb{1}_{\{|\nabla_\theta l(\theta_t, x_t)|^2 \leq M\}} + |m_t| |\nabla_\theta l(\theta_t, x_t)|^2 \mathbb{1}_{\{|\nabla_\theta l(\theta_t, x_t)|^2 > M\}}$$

With the same notations and summing all these terms until time $T \in \mathbb{N}^*$, we can write

$$\begin{aligned} m_T \geq m_0 &+ C\delta \sum_{t=0}^{T-1} (1 - m_t)(m_t - b/\sqrt{d})^{s-1} - \delta \sum_{t=0}^{T-1} \nabla^S M_t \cdot \theta^* - \delta^2 \sum_{t=0}^{T-1} |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}} \\ &- \delta^2 \sum_{t=0}^{T-1} |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} - \delta^3 \sum_{t=0}^{T-1} \xi_t, \end{aligned}$$

where we use, for the sake of compactness, the shortcut notation $l_t = l(x_t, \theta_t)$. The strategy of the proof is the following: the first term is the drift term that makes the correlation grow, the second term is simply a martingale term that we deal with via standard martingale inequality, and the forth and fifth term are discretization error that we will bound loosely. The difficulty comes from the third term: the proof is based on the fact that we use a “part” of the drift term (say half) to control it. This is why we decide to rewrite finally our inequality as,

$$\begin{aligned} (28) \quad m_T \geq m_0 &+ \delta \frac{C}{2} \sum_{t=0}^{T-1} (1 - m_t)(m_t - b/\sqrt{d})^{s-1} - \delta \sum_{t=0}^{T-1} \nabla^S M_t \cdot \theta^* - \delta \sum_{t=0}^{T-1} D_t \\ &- \delta^2 \sum_{t=0}^{T-1} |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} - \delta^3 \sum_{t=0}^{T-1} \xi_t, \end{aligned}$$

where we have defined $D_t := \frac{C}{2}(1 - m_t)(m_t - b/\sqrt{d})^{s-1} - \delta |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}}$. The following section show how to control these five terms in a quantitative way.

B.3. Weak recovery.

Good initialization. During all this section, we condition on the event $\{m_0 \geq 5b/\sqrt{d}\}$.

Before stating these lemmas, let us introduce some new notations. As already introduce, we recall that we denote $S_{\eta} := \{\theta \in \mathcal{S}_{d-1}, m_{\theta} \geq \eta\}$, the spherical cap of level $\eta \in (0, 1)$. Moreover for $\alpha \in (-1, 1)$, similarly to what is done in [Ben Arous et al., 2021], we define the following stopping times $\tau_{\alpha}^+ := \inf\{t \geq 0, m_{\theta_t} \geq \alpha\}$ and $\tau_{\alpha}^- := \inf\{t \geq 0, m_{\theta_t} \leq \alpha\}$ reciprocally as the first time when $(\theta_t)_{t \geq 0}$ enters in S_{α} or leaves S_{α} .

B.3.1. Proof of Theorem 3.2. Thanks to Lemmas B.2, B.3, B.4, B.5 and B.6, that serve bounding all the terms in the m_T inequality, there exists a constant K that depend solely on the model such that we have the following lower bound: for all $\lambda > 0$, conditionally to the event on the events $\{T \leq \tau_{1/2}^+ \wedge \tau_{2b/\sqrt{d}}^-\}$,

$$m_T \geq m_0 + \frac{C}{2^{s+1}} \delta \sum_{t=0}^{T-1} m_t^{s-1} - 4\lambda,$$

with probability larger than $1 - \left(\frac{KT\delta^2}{\lambda^2} + \exp\left(-\frac{\lambda^2}{2K^2\delta^2T + \lambda\delta(C + \delta M)}\right) + \frac{KTd^2\delta^2}{\lambda M} + \frac{KTd\delta^3}{\lambda} \right)$.

Now we choose $\lambda = b/\sqrt{d}$ and $M = d^{3/2}$ so that

$$m_T \geq \frac{b}{\sqrt{d}} + \frac{C}{2^{s+1}} \delta \sum_{t=0}^{T-1} m_t^{s-1},$$

with probability at least $1 - p_{\delta,M}(T)$, where we defined naturally

$$p_{\delta,M}(T) := \left(\frac{KTd\delta^2}{b^2} + \exp\left(-\frac{b^2}{2K^2d\delta^2T + b\sqrt{d}\delta(C + \delta M)}\right) + \frac{KTd^{5/2}\delta^2}{bM} + \frac{KTd^{3/2}\delta^3}{b} \right).$$

Let us upper bound the probability $p_{\delta,M}(T)$. Let us set $\varepsilon > 0$ a small constant. First, in the exponential term, the term $b\sqrt{d}\delta(C + \delta d^{3/2})$ is negligible in virtue of the fact that in any of the cases of Theorem 3.2, we have $\delta \leq \varepsilon/d$. Moreover, for the sake of clarity, we gather all constant K, C, b as one constant generic K , as these depend only on the data distribution and the link function. Hence, for d large enough,

$$p_{\delta,M}(T) \leq K \left(dT\delta^2 + \exp\left(-\frac{1}{dT\delta^2}\right) + dT\delta^2 + d^{3/2}T\delta^3 \right),$$

and as $d^{3/2}T\delta^3 \lesssim dT\delta^2$ for the range of δ we choose, we have $p_{\delta,M}(T) \leq K \left(dT\delta^2 + \exp\left(-\frac{1}{dT\delta^2}\right) \right)$, and considering that we will take in any case $dT\delta^2 \leq 1$, as we have the inequality $\exp\left(-\frac{1}{dT\delta^2}\right) \leq dT\delta^2$, so that finally

$$p_{\delta,M}(T) \leq KdT\delta^2$$

We divide the proof into the three cases $s = 1$, $s = 2$, $s \geq 3$.

Case $s = 1$, $\delta = \varepsilon/d$. In this case, we have that with probability $1 - p_{\delta,M}(T)$,

$$m_T \geq \frac{b}{\sqrt{d}} + \frac{C\delta}{2^s} T.$$

The right and side is larger than $1/2$ as soon as $\delta T \geq 2^s/C$. From this we have that with probability at least $1 - p_{\delta,M}(T)$, the hitting time is upper bounded by

$$\tau_{1/2}^+ \leq \frac{2^s}{C\delta}.$$

Now, taking $\delta = \varepsilon d^{-1}$, we can check that for ε small enough, $dT\delta^2 \leq 2^s\varepsilon/C = \varepsilon\mathcal{O}(1)$ so that we have that with probability at least $1 - K\varepsilon$, we have

$$\tau_{1/2}^+ \leq \frac{K}{\varepsilon} d.$$

Case $s = 2$, $\delta = \varepsilon/(d \log d)$. Now by a discrete version of Grönwall inequality, recalled in Lemma B.10, we have with probability at least $1 - p_{\delta,M}(T)$,

$$m_T - \frac{b}{\sqrt{d}} \geq \frac{b}{\sqrt{d}} \left(1 + \delta \frac{C}{2}\right)^T \geq \frac{b}{\sqrt{d}} e^{C\delta T},$$

for d large enough. And as the right hand side is larger than $1/2 + b/\sqrt{d}$ whenever,

$$\delta T \geq \frac{1}{C} \log(\sqrt{d}/4b),$$

for d large enough compared to b . Then taking such a T , with probability at least $1 - p_{\delta,M}(T)$, the hitting time is upper bounded by

$$\tau_{1/2}^+ \leq \frac{2}{C\delta} \log(d).$$

Now, taking $\delta = \varepsilon d^{-1}(\log d)^{-1}$, we can check that for ε small enough, $dT\delta^2 \leq \frac{\varepsilon}{C} = \varepsilon\mathcal{O}(1)$ so that we have that with probability at least $1 - K\varepsilon$, we have

$$\tau_{1/2}^+ \leq \frac{K}{\varepsilon} d \log(d)^2.$$

Case $s \geq 3$, $\delta = \varepsilon d^{-s/2}$. Now by the discrete version of Bihari-LaSalle inequality, recalled in Lemma B.10, we have with probability at least $1 - p_{\delta,M}(T)$,

$$m_T - \frac{b}{\sqrt{d}} \geq \frac{b}{\sqrt{d}} \left(1 - \delta \frac{C(s-2)}{2} \left(\frac{b}{\sqrt{d}}\right)^{s-2} T\right)^{-\frac{1}{s-2}}.$$

And as the right hand side is larger than $1/2 + b/\sqrt{d}$ whenever,

$$\delta T \geq \frac{d^{(s-2)/2}}{C(s-2)b^{s-2}},$$

for d large enough compare to b . Then taking such a T , with probability at least $1 - p_{\delta,M}(T)$, the hitting time is upper bounded by

$$\tau_{1/2}^+ \leq \frac{1}{Cb^{s-2}} \frac{d^{\frac{s-2}{2}}}{\delta}.$$

Now, taking $\delta = \varepsilon d^{-s/2}$, we can check that for ε small enough, $dT\delta^2 \leq \frac{\varepsilon}{C(s-2)b^{s-2}} = \varepsilon\mathcal{O}(1)$ so that we have that with probability at least $1 - K\varepsilon$, we have

$$\tau_{1/2}^+ \leq \frac{K}{\varepsilon} d^{s-1}.$$

B.3.2. Technical intermediate result to lower bound each term of Eq. (28).

Lemma B.2 (ODE term). *Conditioned to the event $\{T \leq \tau_{1/2}^+ \wedge \tau_{2b/\sqrt{d}}^-\}$, we have the inequality*

$$\sum_{t=0}^{T-1} (1 - m_t)(m_t - b/\sqrt{d})^{s-1} \geq \frac{1}{2^s} \sum_{t=0}^{T-1} m_t^{s-1}.$$

Proof. This simply results from the fact that for all $t \leq T - 1$, we have $\{t \leq \tau_{1/2}^+ \wedge \tau_{2b/\sqrt{d}}^-\} \subset \{T \leq \tau_{1/2}^+ \wedge \tau_{2b/\sqrt{d}}^-\}$, so that we can use the inequalities $1 - m \geq 1/2$ and $m - b/\sqrt{d} \geq m/2$. Summing these terms until $T - 1$ gives the proof of the lemma. $\blacksquare$

Lemma B.3 (First martingale term). *For all $\lambda > 0$, we have that*

$$(29) \quad \mathbb{P} \left(\sup_{t \leq T} \delta \left| \sum_{k=0}^{t-1} \nabla^S M_k \cdot \theta^* \right| \geq \lambda \right) \leq \frac{KT\delta^2}{\lambda^2},$$

where $K > 0$, that depends solely on the model through f, ν .

Proof. This is a consequence of Doob's maximal inequality for (sub)martingale. Indeed, for $t \leq T$, let $H_{t-1} = \sum_{k=0}^{t-1} \nabla^S M_k \cdot \theta^*$. We have that H_t is a $\mathcal{F}_t$ -adapted martingale and we have the following upper bounded on its variance:

$$\begin{aligned} \mathbb{E}[H_{t-1}^2] &= \mathbb{E} \left[\left(\sum_{k=0}^{t-1} \nabla^S M_k \cdot \theta^* \right)^2 \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{t-1} (\nabla^S M_k \cdot \theta^*)^2 \right] \\ &\leq t \sup_{\theta} \mathbb{E}_x \left[(\nabla^S M_k \cdot \theta^*)^2 \right] \\ &\leq Kt, \end{aligned}$$

where the last inequality comes from the Lemma B.8. Now, thanks to Doob's maximal inequality, we have for all $\lambda > 0$,

$$\mathbb{P} \left(\sup_{t \leq T} \delta |H_{t-1}| \geq \lambda \right) \leq \frac{\mathbb{E}[H_{T-1}^2] \delta^2}{\lambda^2} \leq \frac{KT\delta^2}{\lambda^2},$$

and this concludes the proof of the lemma. $\blacksquare$

Lemma B.4 (Submartingale term). *For all $\lambda > 0$, if for all $t \leq T$, $m_t \in [2b/\sqrt{d}, 1/2]$, and δ is such that $\delta \leq \varepsilon/d$, with a small enough constant $\varepsilon > 0$, we have that*

$$(30) \quad \mathbb{P} \left(\delta \sum_{t=0}^{T-1} D_t \leq -\lambda \right) \leq \exp \left(-\frac{\lambda^2}{2K^2\delta^2T + \lambda\delta(C + \delta M)} \right)$$

where $K > 0$, that depends solely on the model through f, ν .

Proof. First, recall that we have defined $D_t = \frac{C}{2}(1-m_t)(m_t-b/\sqrt{d})^{s-1} - \delta|m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}}$. Let us notice that if $m_t \in [2b/\sqrt{d}, 1/2]$, then $1 - m_t \geq 1/2$ and $(m_t - b/\sqrt{d})^{s-1} \geq m_t^{s-1}/2^{s-1}$. Hence, if m_t lies in such an interval,

$$\begin{aligned} D_t &\geq \frac{C}{2^{s+1}} m_t^{s-1} - \delta|m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}} \\ &\geq \frac{C}{2^{s+1}} m_t^{s-1} \left(1 - 2^{s+1} \delta \frac{|\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}}}{C m_t^{s-2}} \right). \end{aligned}$$

Now, for δ such that $\mathbb{E} \left[1 - 2^{s+1} \delta \frac{|\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}}}{C m_t^{s-2}} \mid \mathcal{F}_{t-1} \right] \geq 0$, $(\sum_{k=1}^t D_k)_{t \geq 0}$ is a submartingale, which is true as soon as

$$\delta \leq \frac{C m_t^{s-2}}{2^{s+1} \sup_{\theta} \mathbb{E} \left[|\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}} \mid \mathcal{F}_{t-1} \right]},$$

which is itself true if

$$\delta \leq \frac{C}{4^s \sup_{\theta} \mathbb{E} \left[|\nabla_{\theta} l_t|^2 \right]},$$

which is implied by the condition required in the lemma given the upper bound on $\mathbb{E}[|\nabla_{\theta} l_t|^2]$ provided in Lemma B.8. In order to apply Freedman tail inequality for this submartingale, let us provide upper bound on the increments as well as their variance. Indeed, we have, for all $t \geq 0$,

$$\begin{aligned} |D_t| &\leq \frac{C|1-m_t||m_t-b\sqrt{d}|^{s-1}}{2} + \delta|m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}} \\ &\leq \frac{C + \delta M}{2}, \end{aligned}$$

and in virtue of the inequality $(a+b)^2 \leq 2(a^2 + b^2)$, we have

$$\begin{aligned} \mathbb{E} [D_t^2 \mid \mathcal{F}_{t-1}] &\leq 2 \left(\frac{C^2|1-m_t|^2|m_t-b\sqrt{d}|^{2(s-1)}}{4} + \delta^2|m_t|^2 \mathbb{E} \left[|\nabla_{\theta} l_t|^4 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 \leq M\}} \right] \right) \\ &\leq \frac{C^2 + \delta^2 \mathbb{E} [|\nabla l|^4]}{2} \\ &\leq \frac{C^2 + K\delta^2 d^2}{2} \\ &\leq K^2. \end{aligned}$$

Hence, by the Freedman tail inequality recalled in Theorem B.9, for all $\lambda > 0$,

$$\mathbb{P} \left(\delta \sum_{t=0}^{T-1} D_t \leq -\lambda \right) \leq \exp \left(-\frac{\lambda^2}{2K^2\delta^2 T + \lambda\delta(C + \delta M)} \right),$$

which concludes the proof of the Lemma. ■

Lemma B.5 (First discretization term). *We have that, almost surely*

$$(31) \quad \mathbb{P} \left(\sup_{t \leq T} \delta^2 \sum_{t=0}^{T-1} |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} \geq \lambda \right) \leq \frac{KT\delta^2 d^2}{\lambda M},$$

where $K > 0$ depends solely on the model through f, ν .

Proof. This term is handled via a combination of Markov and Cauchy-Schwartz inequalities. First, notice that,

$$\sup_{t \leq T} \delta^2 \sum_{t=0}^{T-1} |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} \leq T\delta^2 \sup_{t \leq T} \left\{ |m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} \right\}.$$

Furthermore, for all $t \leq T$, all $\lambda > 0$, via Markov inequality, then Cauchy-Schwartz inequality,

$$\begin{aligned} \mathbb{P} \left(|m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} \geq \lambda \right) &\leq \frac{\mathbb{E} \left[|m_t| |\nabla_{\theta} l_t|^2 \mathbb{1}_{\{|\nabla_{\theta} l_t|^2 > M\}} \right]}{\lambda} \\ &\leq \frac{\sqrt{\mathbb{E} [|\nabla_{\theta} l_t|^4]} \sqrt{\mathbb{P} (|\nabla_{\theta} l_t|^2 > M)}}{\lambda} \\ &\leq \frac{\sqrt{\mathbb{E} [|\nabla_{\theta} l_t|^4]} \sqrt{\mathbb{E} [|\nabla_{\theta} l_t|^4]} / M^2}{\lambda} \\ &\leq \frac{\mathbb{E} [|\nabla_{\theta} l_t|^4]}{\lambda M} \\ &\leq \frac{Kd^2}{\lambda M}, \end{aligned}$$

where the last inequality is due to Lemma B.8. Multiplying this bound by $T\delta^2$ ends the proof the lemma. $\blacksquare$

Lemma B.6 (Second discretization term). *For all $\lambda > 0$, we have that*

$$(32) \quad \mathbb{P} \left(\sup_{t \leq T} \delta^3 \sum_{k=0}^{t-1} \xi_k \geq \lambda \right) \leq \frac{KTd\delta^3}{\lambda},$$

where $K > 0$ depends solely on the model through f, ν .

Proof. Recall that $\xi_k = |\nabla_{\theta} l(\theta_k, x_k)|^2 |\nabla^S l(\theta_k, x_k) \cdot \theta^*|$. The bound follows from an application of Markov's inequality. Indeed, since all the terms of the sum are positive, the supremum is

attained in $t = T - 1$, and we shall only consider this case. For $\lambda > 0$,

$$\begin{aligned}
\mathbb{P} \left(\delta^3 \sum_{t=0}^{T-1} \xi_t \geq \lambda \right) &\leq \frac{\delta^3}{\lambda} \mathbb{E} \left[\sum_{t=0}^{T-1} \xi_t \right] \\
&\leq \frac{T\delta^3}{\lambda} \sup_{\theta} \left\{ \mathbb{E}_x [|\nabla_{\theta} l(\theta, x)|^2] |\nabla^S l(\theta, x) \cdot \theta^*| \right\} \\
&\leq \frac{T\delta^3}{\lambda} \sup_{\theta} \left\{ \sqrt{\mathbb{E}_x [|\nabla_{\theta} l(\theta, x)|^4]} \sqrt{\mathbb{E}_x [|\nabla^S l(\theta, x) \cdot \theta^*|^2]} \right\} \\
&\leq \frac{T\delta^3}{\lambda} \sqrt{\sup_{\theta} \mathbb{E}_x [|\nabla_{\theta} l(\theta, x)|^4]} \sqrt{\sup_{\theta} \mathbb{E}_x [|\nabla^S l(\theta, x) \cdot \theta^*|^2]} \\
&\leq \frac{T\delta^3}{\lambda} \sqrt{\sup_{\theta} \mathbb{E}_x [|\nabla_{\theta} l(\theta, x)|^4]} \sqrt{\sup_{\theta} \mathbb{E}_x [|\nabla^S l(\theta, x) \cdot \theta^*|^2]} \\
&\leq \frac{T\delta^3}{\lambda} \sqrt{Kd^2} \sqrt{K} \\
&\leq \frac{KTd\delta^3}{\lambda},
\end{aligned}$$

where the penultimate inequality comes from Lemma B.8. ■

B.4. Strong recovery. The reasoning is almost identical to the one of the previous section, except from the fact that instead of tracking the growing movement on $(m_t)_{t \geq 0}$, we will track the decaying movement of $(1 - m_t)_{t \geq 0}$.

B.4.1. Upper bound on the residual. As said in the main text, we place ourselves *after* the weak recovery time. Thanks to the Markovian property of the SGD dynamics, we have the equality between all time $s > 0$ marginal laws of

$$\left(\theta_{\tau_{1/2}^+ + s} \mid \tau_{1/2}^+, \theta_{\tau_{1/2}^+} \right) \stackrel{\text{Law}}{=} \left(\theta_s \mid \theta_s = \theta_{\tau_{1/2}^+} \right),$$

and hence the strong recovery question is equivalent to study the dynamics with initialization such that $m_{\theta} = 1/2$. As demonstrated before we have that $\mathbb{P}(\tau_{1/2}^+ < \infty) \geq 1 - K\varepsilon$ so that up to ε terms, this conditioning does not hurt the probability of the later events. In fact this conditioning seems even artificial as it seems provable that $\tau_{1/2}^+$ is almost surely finite. Yet, we leave this more precise study for another time.

B.4.2. A (slightly) different decomposition. Let us define for all $t \in \mathbb{N}$, the residual $u_t = 1 - m_{t+\tau_{1/2}^+} > 0$, and thanks to the lower bound given by Eq. (27), we have

$$u_{t+1} \leq u_t - C\delta u_t (m_t - b/\sqrt{d})^{s-1} + \delta \nabla^S M_t \cdot \theta^* + \delta^2 |m_t| |\nabla l(x_t, \theta_t)|^2 + \delta^3 \xi_t,$$

From there, the proof is similar to the weak recovery case, except that the extra-care we used for the term $\delta^2 |m_t| |\nabla l(x_t, \theta_t)|^2$ is not necessary. We use simply the decomposition of this term in a second martingale term

$$N_t = |m_t| |\nabla l(x_t, \theta_t)|^2 - \mathbb{E} [|m_t| |\nabla l(x_t, \theta_t)|^2 | \mathcal{F}_{t-1}]$$

and the drift that we directly upper bound as $\mathbb{E} [|m_t| |\nabla l(x_t, \theta_t)|^2 | \mathcal{F}_{t-1}] \leq Kd$. Now similarly to Lemma B.3, we have the upper bound:

Lemma B.7 (New martingale term). *For all $\lambda > 0$, we have that*

$$(33) \quad \mathbb{P} \left(\sup_{t \leq T} \delta^2 \left| \sum_{k=0}^{t-1} N_k \right| \geq \lambda \right) \leq \frac{Kd^2 T \delta^4}{\lambda^2},$$

where $K > 0$, that depends solely on the model through f, ν .

Proof. This is a consequence of Doob's maximal inequality for the martingale. Indeed, for $t \leq T$, let $H_{t-1} = \sum_{k=0}^{t-1} N_k$. We have that N_t is a $\mathcal{F}_t$ -adapted martingale and we have the following upper bounded on its variance:

$$\begin{aligned} \mathbb{E}[N_{t-1}^2] &= \mathbb{E} \left[\left(\sum_{k=0}^{t-1} \nabla^S M_k \cdot \theta^* \right)^2 \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{t-1} N_k^2 \right] \\ &\leq t \sup_{\theta} \mathbb{E}_x (N_k)^2 \\ &\leq Kd^2 t, \end{aligned}$$

where the last inequality comes from the Lemma B.8. Now, thanks to Doob's maximal inequality, we have for all $\lambda > 0$,

$$\mathbb{P} \left(\sup_{t \leq T} \delta^2 |H_{t-1}| \geq \lambda \right) \leq \frac{\mathbb{E}[H_{T-1}^2] \delta^4}{\lambda^2} \leq \frac{Kd^2 T \delta^4}{\lambda^2},$$

and this concludes the proof of the lemma. ■

Now, everything is in order to prove the Theorem 3.3.

B.4.3. Proof of Theorem 3.3. Let us fix a small number $\varepsilon > 0$. As previously, thanks to Lemmas B.3, B.6, B.7, there exists $K > 0$ that depends solely on the model such that we have the following upper bound: for all λ , and $t \leq \tau_{1/3}^- \wedge \tau_{1-\varepsilon}^+$ summing between times 0 and t ,

$$u_t \leq u_0 - \frac{C\delta}{4^{s-1}} \sum_{k=0}^{t-1} u_k + K\delta^2 d + 3\lambda,$$

with probability larger than $1 - \left(\frac{Kt\delta^2}{\lambda^2} + \frac{Kd^2t\delta^4}{\lambda^2} + \frac{Ktd\delta^3}{\lambda} \right)$ and d large enough. Let us choose $\lambda = 1/16$ and δ small enough so that $K\delta^2d \leq \lambda$. Hence, realizing that $u_0 \leq 1/2$, we have

$$u_t \leq \frac{3}{4} - \frac{C\delta}{4^{s-1}} \sum_{k=0}^{t-1} u_k ,$$

with probability at least $1 - Kt\delta^2(1 + d^2\delta^2 + d\delta) \gtrsim 1 - Kt\delta^2$, as we choose in any case $\delta = \varepsilon\mathcal{O}(1)$. Note that we used the same convention as in the weak recovery case that K denotes *any* constant that simply depend on the model. We have by Grönwall inequality (Lemma B.10)

$$u_t \leq \frac{3}{4} \left(1 - \frac{C\delta}{4^{s-1}} \right)^t \leq \frac{3}{4} e^{-\frac{C\delta}{4^{s-1}}t}.$$

Hence, as the right end side is smaller than ε for the time

$$t\delta \geq \frac{4^{s-1}}{C} \log(1/\varepsilon),$$

we choose such a t , so that with probability at least $1 - K\delta \log(1/\varepsilon)$, the delayed hitting time $\bar{\tau}_{1-\varepsilon}^+ := \inf\{t \geq 0, u_t \leq \varepsilon\}$ satisfies

$$\bar{\tau}_{1-\varepsilon}^+ \leq \frac{4^{s-1}}{C\delta} \log(1/\varepsilon),$$

and taking $\delta = \varepsilon/d$ gives that with a probability at least $1 - K\varepsilon \log(1/\varepsilon)/d$, we have

$$\bar{\tau}_{1-\varepsilon}^+ \leq \frac{4^{s-1}}{C\varepsilon} d \log(1/\varepsilon).$$

Considering that d is large and ε is simply a constant we get that $1 - K\varepsilon \log(1/\varepsilon)/d \geq 1 - K\varepsilon$ and this concludes the proof of Theorem 3.3.

B.5. Some technical bounds. We end this section by providing (i) some necessary technical bound on the quantities appearing in the SGD controls (ii) some discrete versions of Grönwall-type lemmas.

B.5.1. Technical bounds on models expectations.

Lemma B.8 (Technical bounds). *We have that there exists a constant $K > 0$ solely depending on the function ϕ and the distribution ν such that:*

$$(34) \quad \sup_{\theta \in \mathcal{S}_{d-1}} \mathbb{E}_x [\langle \nabla_{\theta}^S M(x, \theta), \theta_* \rangle^2] \leq K, \quad \text{and} \quad \sup_{\theta} \mathbb{E}_x [|\nabla^S l(\theta, x) \cdot \theta^*|^2] \leq K$$

$$(35) \quad \sup_{\theta \in \mathcal{S}_{d-1}} \mathbb{E}_x [|\nabla_{\theta} l(\theta, x)|^2] \leq Kd,$$

$$(36) \quad \sup_{\theta \in \mathcal{S}_{d-1}} \mathbb{E}_x [|\nabla_{\theta} l(\theta, x)|^4] \leq Kd^2.$$

Proof. In all the following proof we consider any $\theta \in \mathcal{S}_{d-1}$. Notice that we have the following calculation that is common to all the bounds we cover

$$\nabla l(\theta, x) = x\phi'(x \cdot \theta)\phi(x \cdot \theta_*)$$

We treat the three bounds separately.

First terms. We have that for all $x \in \mathbb{R}^d$,

$$M(x, \theta) = l(x, \theta) - \mathbb{E}_\nu[l(x, \theta)],$$

hence

$$\begin{aligned} \nabla_\theta^\mathcal{S} M(x, \theta) &= \nabla_\theta^\mathcal{S} l(x, \theta) - \mathbb{E}_\nu[\nabla_\theta^\mathcal{S} l(x, \theta)] \\ &= \nabla_\theta l(x, \theta) - \mathbb{E}_\nu[\nabla_\theta l(x, \theta)] - (\theta \cdot \nabla_\theta l(x, \theta))\theta + \mathbb{E}_\nu[(\theta \cdot \nabla_\theta l(x, \theta))\theta], \end{aligned}$$

and finally,

$$\nabla_\theta^\mathcal{S} M(x, \theta) \cdot \theta_* = \nabla_\theta l(x, \theta) \cdot \theta_* - \mathbb{E}_\nu[\nabla_\theta l(x, \theta) \cdot \theta_*] - (\theta \cdot \nabla_\theta l(x, \theta))m + \mathbb{E}_\nu[(\theta \cdot \nabla_\theta l(x, \theta))m].$$

hence thanks to applying the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, this amounts to bound first

$$\mathbb{E}_x (\nabla_\theta l(x, \theta) \cdot \theta_*)^2 = \mathbb{E}_x [(x \cdot \theta_*)^2 \phi'^2(x \cdot \theta) \phi^2(x \cdot \theta_*)] \leq K,$$

and second

$$\mathbb{E}_x ((\nabla_\theta l(x, \theta) \cdot \theta)m)^2 \leq \mathbb{E}_x [(x \cdot \theta)^2 \phi'^2(x \cdot \theta) \phi^2(x \cdot \theta_*)] \leq K.$$

Second term. We have

$$\mathbb{E}_x |\nabla_\theta l(x, \theta)|^2 = \mathbb{E}_x [|x|^2 \phi'^2(x \cdot \theta) \phi^2(x \cdot \theta_*)] \leq Kd.$$

Third term. We have similarly

$$\mathbb{E}_x |\nabla_\theta l(x, \theta)|^4 = \mathbb{E}_x [|x|^4 \phi'^2(x \cdot \theta) \phi^2(x \cdot \theta_*)] \leq Kd^2.$$

■

B.5.2. Standard tail probabilities for submartingales. We recall here a theorem on submartingales from Freedman. This is an adaptation from Theorem 4.1 stated in [Freedman, 1975].

Theorem B.9 (Submartingale tail bound). *Suppose that $(X_t)_{t \in \mathbb{N}}$ is random sequence adapted to a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$. For $T \geq 1$, suppose there exist $a, b > 0$ such that $\mathbb{E}[X_t | \mathcal{F}_{t-1}] \geq 0$, the almost sure upper-bound $\sup_{t \leq T} |X_t| \leq a$ as well as $\sup_{t \leq T} \mathbb{E}[X_t^2 | \mathcal{F}_{t-1}] \leq b$, then for all $\lambda > 0$,*

$$(37) \quad \mathbb{P} \left(\sum_{k=1}^T X_k \leq -\lambda \right) \leq \exp \left(-\frac{\lambda^2}{2(Tb + \lambda a)} \right)$$

B.5.3. Discrete Grönwall and Bihari-Lasalle bounds. We now turn to stating a classical comparison lemma for recursive inequalities.

Lemma B.10 (Grönwall and Bihari-Lasalle). *We have the bounds for the recursive inequalities:*

Case $s = 2$. Suppose $(m_t)_{t \in \mathbb{N}}$ satisfies for $s \geq 3$, and positives numbers $a, b > 0$, and $b < a/2 \wedge 1$,

$$(38) \quad m_t \geq a + b \sum_{k=0}^{t-1} m_k, \quad \text{then, } m_t \geq a(1+b)^t$$

$$(39) \quad m_t \leq a - b \sum_{k=0}^{t-1} m_k, \quad \text{then, } m_t \leq a(1-b)^t.$$

Case $s \geq 3$. Suppose $(m_t)_{t \in \mathbb{N}}$ satisfies for $s \geq 3$, and positives numbers $a, b > 0$:

$$(40) \quad m_t \geq a + b \sum_{k=0}^{t-1} m_k^{s-1}, \quad \text{then, } m_t \geq a(1 - (s-2)ba^{s-2}t)^{-\frac{1}{s-2}}.$$

Proof. The case $s = 2$ is known to be the discrete version of the Grönwall lemma and is treated in all standard textbooks, the case $s \geq 3$ referred to as the Bihari-Lasalle inequality is for example proven in Appendix C of [Arous et al., 2021]. $\blacksquare$

APPENDIX C. THE LPG PROPERTY IN THE SYMMETRIC CASE: PROOFS OF SECTION 4.1

C.1. Useful Facts about Gegenbauer Polynomials. We recall known facts on Gegenbauer Polynomials.

Definitions. Recall that $P_{j,d}$ denotes the Gegenbauer polynomial of degree j and dimension d , normalized so that $P_{j,d}(1) = 1$ for all j, d . We denote also $\bar{P}_{j,\lambda}$ the Gegenbauer polynomials normalized so that $\|\bar{P}_{j,\lambda}\|_{L^2(\mathbb{R}, u_{2\lambda+2})}^2 = \pi 2^{1-2\lambda} \frac{\Gamma(j+2\lambda)}{(j+\lambda)\Gamma^2(\lambda)\Gamma(j+1)}$. Throughout the proof, we will use either d , and from time to time the mute symbol λ to denote the dimension variable of Gegenbauer polynomials. They satisfy the following recurrence:

$$(41) \quad (j+1)\bar{P}_{j+1,\lambda}(t) = 2(j+\lambda)t\bar{P}_{j,\lambda}(t) - (j+2\lambda-1)\bar{P}_{j-1,\lambda}(t),$$

with first terms: $\bar{P}_{0,\lambda}(t) = 1$ and $\bar{P}_{1,\lambda}(t) = 2\lambda t$.

Rodrigues Formula for Gegenbauer Polynomials. The Gegenbauer polynomials can be represented as repeated derivatives of a simple polynome.

Proposition C.1 ([Frye and Efthimiou, 2012, Proposition 4.19]). *We have the formula*

$$(42) \quad P_{j,d}(t) = \frac{(-1)^j}{2^j(j+(d-3)/2)_j} (1-t^2)^{(3-d)/2} \left(\frac{d}{dt}\right)^j (1-t^2)^{j+(d-3)/2},$$

where $(x)_j = \prod_{k=0}^{j-1} (x-k)$ is the falling factorial.

Hecke-Funk Formula. Recall that we use the notation τ_d to denote the uniform distribution on the sphere and u_d the distribution of, e.g., its first coordinate: $u_d \propto (1 - t^2)^{(d-3)/2} \mathbb{1}_{[-1,1]}$.

Theorem C.2 ([Frye and Efthimiou, 2012, Theorem 4.24]). For $\theta, \theta' \in \mathcal{S}_{d-1}$, $f \in L^2_{u_d}(\mathbb{R})$ and $j \in \mathbb{N}$,

$$\begin{aligned} \langle f_\theta, (P_{j,d})_{\theta'} \rangle_{\tau_d} &= \Omega_{d-2} P_{j,d}(\theta \cdot \theta') \langle f, P_{j,d} \rangle_{u_d} \\ &= \Omega_{d-2} P_{j,d}(\theta \cdot \theta') \int_{-1}^1 f(t) P_{j,d}(t) (1 - t^2)^{(d-3)/2} dt . \end{aligned} \quad (43)$$

Fact C.3. [Derivative Representation] We have the following derivation property for all j, d :

$$P'_{j,d} = \frac{j(j+d-2)}{(d-1)} P_{j-1,d+2} . \quad (44)$$

Proof. Recall the normalization relationships $\lambda = \frac{d}{2} - 1$, $\bar{P}_{j,\lambda}(1) = \frac{\Gamma(j+2\lambda)}{\Gamma(j+1)\Gamma(2\lambda)}$, $P_{j,d} = \frac{\Gamma(j+1)\Gamma(d-2)}{\Gamma(j+d-2)} \bar{P}_{j,\frac{d}{2}-1}$, as well as the identity $\bar{P}'_{j,\lambda} = 2\lambda \bar{P}_{j-1,\lambda+1}$. Thus,

$$\begin{aligned} P'_{j,d} &= \frac{\Gamma(j+1)\Gamma(d-2)}{\Gamma(j+d-2)} \bar{P}'_{j,\frac{d}{2}-1} \\ &= 2 \frac{\Gamma(j+1)\Gamma(d-2)}{\Gamma(j+d-2)} \left(\frac{d}{2} - 1\right) \bar{P}_{j-1,\frac{d}{2}} \\ &= (d-2) \frac{\Gamma(j+1)\Gamma(d-2)}{\Gamma(j+d-2)} \frac{\Gamma(j-1+d)}{\Gamma(j)\Gamma(d)} P_{j-1,d+2} \\ &= \frac{j(j+d-2)}{(d-1)} P_{j-1,d+2} \end{aligned} \quad (45)$$

■

We have the following bound of the location of the largest root $z_{j,d}$ of $P_{j,d}$:

Fact C.4 (Bound on the Largest Root, [Area et al., 2004, Corollary 2.3]).

$$z_{j,d} \leq \sqrt{\frac{(j-1)(j+d-4)}{(j+d/2-3)(j+d/2-2)}} \cos(\pi/(j+1)) . \quad (46)$$

And we have the following bound on the Taylor expansion of the Gegenbauer polynomials:

Fact C.5 (Taylor Upper bound beyond largest root).

$$P_{j,d}(t) \geq (t - z_{j,d})^j , \text{ for } t \geq z_{j,d} ,$$

Proof. Note that all families of orthogonal polynomials have exclusively real, simple roots. Therefore, by Rolle's theorem, the $j-1$ critical points of $P_{j,d}$ must be interlaced with the j zeroes. So all zeroes of $P'_{j,d}$ are upper bounded by $z_{j,d}$. Furthermore, by Fact C.3, $P'_{j,d}$ is itself

an orthogonal polynomial. So applying this argument recursively, we see the zeros of $P_{j,d}^{(k)}$ for $k \leq j$ are all upper bounded by $z_{j,d}$.

Note also by Fact C.3 that, because $P_{j,d}(1) = 1$ for any choice of j and d , it follows that $P_{j,d}^{(k)}(1) > 0$. This implies $P_{j,d}^{(k)}(z_{j,d}) > 0$, as in order to flip signs there would need to be a zero in the range $[z_{j,d}, 1]$ which we've confirmed above cannot exist.

Now, consider a Taylor expansion

$$(47) \quad P_{j,d}(t) = \sum_{i=0}^j c_i (t - z_{j,d})^i$$

Observe that $P_{j,d}^{(k)}(z_{j,d}) = k!c_k$, and therefore by the above argument we have $c_k > 0$. So it remains to show that $c_j \geq 1$.

Consider applying Fact C.3 repeatedly, then we have:

$$(48) \quad P_{j,d}^{(j)}(1) = \frac{j! \prod_{l=1}^j (j + d - 3 + l)}{\prod_{l=1}^j (d - 3 + 2l)}$$

$$(49) \quad = j! \prod_{l=1}^j \frac{j + d - 3 + l}{d - 3 + 2l}$$

$$(50) \quad \geq j!$$

And from the fact that $P_{j,d}^{(j)}(1) = j!c_j$, we conclude $c_j \geq 1$. ■

C.2. Proof of Proposition 4.1.

Proposition C.6 (Loss representation, restated). *The $\beta_{j,d}$ defined in (14) have the integral representation*

$$(51) \quad \beta_{j,d} = \langle \phi, \mathcal{K}_j \phi \rangle_{L^2(\mathbb{R}, \eta)},$$

where $\mathcal{K}_j$ is a positive semi-definite integral operator of L_η^2 that depend solely on ρ and ϕ , with kernel

$$(52) \quad \mathcal{K}_j(t, t') = \frac{\Omega_{d-2} N(j, d)}{\Omega_{d-1}} \int_0^\infty P_j(r^{-1}t) P_j(r^{-1}t') \bar{u}_d(r^{-1}t) \bar{u}_d(r^{-1}t') \rho(dr),$$

where we defined the conditional density

$$\bar{u}_d(r^{-1}t) = \frac{r^{-1}u_d(r^{-1}t)}{\int_0^\infty (r')^{-1}u_d((r')^{-1}t) \rho(dr')}.$$

Moreover, we have

$$(53) \quad \mathbb{E}_\eta[\phi^2] = \frac{\Omega_{d-2}}{\Omega_{d-1}} \sum_j \beta_{j,d} = \frac{\Gamma((d-2)/2)}{\sqrt{\pi}\Gamma((d-1)/2)} \sum_j \beta_{j,d}.$$

Proof. The marginal conditioned on $\|x\| = r$ is precisely given by $\eta(x_1 = t \mid \|x\| = r) = r^{-1}u_d(r^{-1}t)$, so

$$\eta(t) = \int_0^\infty r^{-1}u_d(r^{-1}t)\rho(dr).$$

We have

$$(54) \quad \begin{aligned} \alpha_{j,r} &= \|P_j\|^{-2} \int_{-1}^1 P_j(t)\phi^{(r)}(t)\tau_d(dt) = \|P_j\|^{-2} \int_{-1}^1 P_j(t)\phi(rt)u_d(t)dt \\ &= \frac{\Omega_{d-2}N(j,d)}{\Omega_{d-1}} r^{-1} \int_{-\infty}^\infty \phi(t)P_j(r^{-1}t)(1-r^{-2}t^2)_+^{(d-3)/2} dt, \end{aligned}$$

so

$$(55) \quad \begin{aligned} \beta_{j,d} &= \frac{\Omega_{d-2}N(j,d)}{\Omega_{d-1}} \int_0^\infty r^{-2} \iint_{-\infty}^\infty \phi(t)P_j(r^{-1}t)(1-r^{-2}t^2)_+^{(d-3)/2} \phi(t')P_j(r^{-1}t')(1-r^{-2}(t')^2)_+^{(d-3)/2} dt dt' \rho(dr) \\ &= \langle \phi, \mathcal{K}_j \phi \rangle_{L^2(\mathbb{R}, \eta)}, \end{aligned}$$

with the $L^2(\mathbb{R}, \eta)$ positive semi-definite integral kernel operator

$$(56) \quad \begin{aligned} \mathcal{K}_j(t, t') &= \frac{\Omega_{d-2}N(j,d)}{\Omega_{d-1}} \eta(t)^{-1} \eta(t')^{-1} \int_0^\infty r^{-2} P_j(r^{-1}t)(1-r^{-2}t^2)_+^{(d-3)/2} P_j(r^{-1}t')(1-r^{-2}(t')^2)_+^{(d-3)/2} \rho(dr) \\ &= \frac{\Omega_{d-2}N(j,d)}{\Omega_{d-1}} \int_0^\infty P_j(r^{-1}t)P_j(r^{-1}t')\bar{u}_d(r^{-1}t)\bar{u}_d(r^{-1}t')\rho(dr), \end{aligned}$$

where we defined the conditional density

$$\bar{u}_d(r^{-1}t) = \frac{r^{-1}u_d(r^{-1}t)}{\int_0^\infty (r')^{-1}u_d((r')^{-1}t)\rho(dr')}.$$

Finally, let us establish (57). We have

$$\begin{aligned}
 \mathbb{E}_\eta \phi^2 &= \mathbb{E}_\rho [\mathbb{E}_{x_1 \| \|x\|=r} \mathbb{E} \phi(x_1)^2] \\
 &= \mathbb{E}_\rho [\mathbb{E}_{u_d} (\phi^{(r)})^2] \\
 &= \mathbb{E}_\rho \sum_j \alpha_{j,d,r}^2 \|P_j\|^2 \\
 &= \mathbb{E}_\rho \sum_j \alpha_{j,d,r}^2 \frac{\Omega_{d-1}}{\Omega_{d-2} N(j, d)} \\
 &= \frac{\Omega_{d-1}}{\Omega_{d-2}} \mathbb{E}_\rho \sum_j \bar{\alpha}_{j,r,d}^2 = \frac{\Omega_{d-1}}{\Omega_{d-2}} \sum_j \beta_{j,d}.
 \end{aligned}
 \tag{57}$$

■

C.3. Proof of Proposition 4.2.

Proof. If $\beta_{j,d} = 0$ for $j < s$, then $\alpha_{j,r,d} = 0$ for $j < s$ and ρ -ae r . We want to show that for any polynomial Q of degree $j' < s$, we must have $\langle \phi, Q \rangle_\eta = 0$.

For each r , consider $Q^{(r)}(t) = Q(rt)$, which is also a polynomial of degree $j' < s$, and its decomposition as $Q^{(r)} = \sum_{j=0}^{j'} b_{j,j',r} P_{j,d}$, which only involves terms of degree $j' < s$ since Gegenbauer polynomials of degree up to r span all polynomials of degree up to r . We have

$$\begin{aligned}
 \langle \phi, Q \rangle_\eta &= \mathbb{E}_\eta [\phi(x) Q(x)] \\
 &= \mathbb{E}_\rho \mathbb{E}_{x_1 \| \|x\|=r} [\phi(x) Q(x)] \\
 &= \mathbb{E}_\rho \mathbb{E}_{u_d} [\phi^{(r)}(x) Q^{(r)}(x)] \\
 &= \mathbb{E}_\rho \left[\sum_{j \leq j'} b_{j,j',r} \alpha_{j,r,d} \right] = 0.
 \end{aligned}
 \tag{58}$$

■

C.4. Proof of Proposition 4.3.

Proposition 4.3 (Spectral characterization of LPG). *Suppose there exist constants $K, C > 0$ and $s \in \mathbb{N}$ such that we both have $\beta_{s,d} \geq C$ and $\sum_{j>s} \beta_{j,d} j(j+d-2) v_{j-1,d+2} \leq K d^{(3-s)/2}$. Then, taking s^* as the infimum of such s , L has the property $\text{LPG}(s^* - 1, z_{s^*,d})$. In particular, whenever $s^* \ll d$, we have $z_{s^*,d} \leq 2\sqrt{s^*/d}$.*

Proof. Assume first that there are $\bar{C}, \bar{\zeta}$ such that

$$P'_{s,d}(t) \geq \bar{C}(t - \bar{\zeta})^{s-1}, \quad \text{for } t \geq \bar{\zeta}.
 \tag{59}$$

Now, let

$$(60) \quad B = \frac{1}{d-1} \sum_{j \geq s} \beta_{j,d} j(j+d-2) v_{j-1,d+2} < 1 ,$$

and define

$$(61) \quad \zeta^* := \left(\frac{B}{\beta_{s,d} \bar{C}} \right)^{1/(s-1)} + \bar{\zeta} .$$

From (59), (60) and (61) we verify that $\ell'(m) = \sum_j \beta_{j,d} P'_{j,d}(m)$ satisfies, for $m \geq \zeta^*$,

$$\ell'(m) \geq \beta_{s,d} \bar{C} \left((m - \bar{\zeta})^{s-1} - (m - \zeta^*)^{s-1} \right) \geq \beta_{s,d} \bar{C} \left[\left(\frac{1 - \bar{\zeta}}{1 - \zeta^*} \right)^{s-1} - 1 \right] (m - \zeta^*)^{s-1} .$$

Finally, we have that for any j, d , the largest root $z_{j,d}$ satisfies $z_{j,d} \leq \sqrt{\frac{(j-1)(j+2d-2)}{(j+d-2)(j+d-1)}} \simeq j/\sqrt{d}$ and

$$P_{j,d}(t) \geq \frac{1}{2} (t - z_{j,d})^s , \text{ for } t \geq z_{j,d} ,$$

which implies that

$$(62) \quad P'_{s,d}(t) \geq \frac{s(s+d-2)}{2(d-1)} (t - z_{s-1,d+2})^{s-1} , \text{ for } t \geq z_{s-1,d+2} .$$

We thus have $\bar{C} = \frac{s(s+d-2)}{2(d-1)}$ with $\bar{\zeta} = z_{s-1,d+2}$.

Finally, we verify that

$$(63) \quad \frac{d^{(s-1)/2}}{s(s+d-2)} \sum_{j > s} \beta_{j,d} j(j+d-2) v_{j-1,d+2} \leq K$$

ensures a local polynomial growth of order $s-1$ at scale $O(1/\sqrt{d})$. Indeed, plugging (63) into (60), together with $\beta_{s,d} \geq C$ yields

$$(64) \quad \left(\frac{B}{\beta_{s,d} \bar{C}} \right)^{1/(s-1)} \leq (CK)^{1/(s-1)} d^{-1/2} ,$$

which shows that $\zeta^* = O(1/\sqrt{d})$. Finally, we observe that $\bar{C} \geq s = \Theta(1)$ if $s < d$.

■

C.5. Proof of Theorem 4.5.

Proof. To prove the theorem, we will establish the sufficient conditions of Proposition 4.3 under our mild assumptions. The key technical results we need are explicit bounds for $v_{j,d}$ and for the sum $\sum_j j^2 \beta_{j,d}$, established in the following two lemmas. Since the parameter $\lambda = d/2 - 1$ is more convenient to express many relationships in Gegenbauer polynomials, we will adopt it in this proof instead of d , without loss of generality.

Lemma C.7 (Control of $v_{j,\lambda}$). *We have*

$$(65) \quad v_{j,\lambda} \lesssim \begin{cases} \left[1 - \left(\frac{\lambda}{j+\lambda}\right)^2\right]^{j/2} & \text{if } j = \Omega(1), \\ \lambda^{\frac{(\alpha-1)\lambda^\alpha}{2}} & \text{if } j = \Theta(\lambda^\alpha), \text{ with } 0 < \alpha < 1, \\ e^{-\frac{1}{2}\lambda^{2-\alpha}} & \text{if } j = \Theta(\lambda^\alpha) \text{ with } 1 \leq \alpha < 3/2, \\ e^{-\lambda} & \text{if } j = \Omega(\lambda^{3/2}). \end{cases}$$

Lemma C.8 (Decomposition of derivative). *If $\phi \in L^2(\mathbb{R}, \mu)$ is such that $\phi' \in L^4(\mathbb{R}, \eta)$ and $\mathbb{E}_\rho[r^4] < \infty$, then $\beta_j = \langle \phi, \mathcal{K}_j \phi \rangle$ satisfies*

$$(66) \quad \sum_j j^2 \beta_{j,d} \leq \frac{\Omega_{d-2}}{\Omega_{d-1}} \mathbb{E}_\rho[r^4]^{1/2} \|\phi'\|_{L^4(\eta)}^2 = O(1/d).$$

Let $s = \inf\{j; \beta_{j,d} \neq 0\}$. We need to verify that there exists a constant $K > 0$ such that

$$(67) \quad \sum_{j>s} \beta_{j,d} j(j+d-2) v_{j-1,d+2} \leq K d^{(3-s)/2}.$$

We will control the LHS by splitting it into appropriate regions, determined by $J_i, i \in \{1, 2, 3\}$.

Let $\alpha = \frac{4}{1+s}$ and $J_1 = \frac{\lambda^\alpha}{2}$. From Lemma C.7, part (i) we have that $v_{j,\lambda} \leq C \left(\frac{j(j+2\lambda)}{(j+\lambda)^2}\right)^{j/2}$, and in particular $v_{j,\lambda} \leq C \lambda^{(\alpha-1)j/2}$ for $j \leq J_1$. As a result, using Lemma C.8,

$$(68) \quad \begin{aligned} \sum_{j=s+1}^{J_1} \beta_j j(j+\lambda) v_{j-1,\lambda+1} &\leq \lambda^{(\alpha-1)(s+1)/2} \sum_{j=s+1}^{J_1} \beta_j j(j+\lambda) \\ &\leq \lambda^{(\alpha-1)(s+1)/2} (C_1 \lambda^{-1} + \lambda \sum_{j=s+1}^{J_1} \beta_j j) \\ &\leq \lambda^{(\alpha-1)(s+1)/2} (C_1 \lambda^{-1} + \lambda \sum_{j=s+1}^{J_1} \beta_j j^2) \\ &\leq \lambda^{(\alpha-1)(s+1)/2} C_2 \\ &\leq C_2 \lambda^{(3-s)/2}. \end{aligned}$$

Let $J_2 = \lambda$. We have

$$(69) \quad \begin{aligned} \sum_{j=J_1+1}^{J_2} \beta_j j(j+\lambda) v_{j-1,\lambda+1} &\leq \lambda^{\frac{(\alpha-1)\lambda^\alpha}{2}} C_3 \\ &\leq C_3 \lambda^{(3-s)/2}. \end{aligned}$$

Let $J_3 = \lambda^{3/2}$. We have

$$(70) \quad \sum_{j=J_2+1}^{J_3} \beta_j j(j+\lambda) v_{j-1, \lambda+1} \leq e^{-\frac{1}{2}\sqrt{\lambda}} C_4 \leq C_4 \lambda^{(3-s)/2}.$$

Finally, the remainder satisfies

$$(71) \quad \sum_{j>J_3} \beta_j j(j+\lambda) v_{j-1, \lambda+1} \leq C_5 (e/2)^{-\lambda} \leq C_5 \lambda^{(3-s)/2},$$

which proves (67). ■

Proof of Lemma C.7. We prove this result by analysing different regimes for j and λ . Concretely, we claim the following:

Claim C.9. *We have the following regimes:*

(1) *For $j = \Omega(1)$, we have*

$$(72) \quad v_{j, \lambda} \lesssim \left[1 - \left(\frac{\lambda}{j + \lambda} \right)^2 \right]^{j/2}.$$

(2) *For $j = \Theta(\lambda^\alpha)$, with $0 < \alpha < 1$, we have*

$$(73) \quad v_{j, \lambda} \lesssim \lambda^{\frac{(\alpha-1)\lambda^\alpha}{2}}.$$

(3) *For $j = \Theta(\lambda^\alpha)$, with $1 \leq \alpha < 2$, we have*

$$(74) \quad v_{j, \lambda} \lesssim e^{-\frac{1}{2}\lambda^{2-\alpha}}.$$

(4) *For $j = \Omega(\lambda^\alpha)$, with $\alpha > 3/2$, we have*

$$(75) \quad v_{j, \lambda} \lesssim e^{-\lambda}.$$

To prove the first three regimes of Claim C.9, we control $v_{j, \lambda}$ based on the distribution of the roots of $P_{j, \lambda}$. We recall that $(z_{k, j, \lambda})_{k \leq j}$ denotes the roots of $P_{j, \lambda}$ in increasing order, and $z_{j, \lambda} = z_{j, j, \lambda}$ its largest root.

Lemma C.10 (Representation of $P_{j, \lambda}$ in terms of its roots, [De Carli, 2008, Lemma 2.1]). *We have*

$$(76) \quad P_{j, \lambda}(t) = \begin{cases} \prod_{k=j/2}^j \frac{t^2 - z_{k, j, \lambda}^2}{1 - z_{k, j, \lambda}^2} & \text{if } j \text{ even,} \\ t \prod_{k=(j+1)/2}^j \frac{t^2 - z_{k, j, \lambda}^2}{1 - z_{k, j, \lambda}^2} & \text{if } j \text{ odd.} \end{cases}$$

From this representation, we deduce that $v_{j,\lambda}$ can be calculated explicitly. Indeed, as the local maxima of $|P_{j,\lambda}(t)|$ are increasing [Szego, 1939], [DLMF, 2022, Eq (18.14.15)], we have the following equation:

$$(77) \quad v_{j,\lambda} = -P_{j,\lambda}(z_{j-1,\lambda+1}) = - \begin{cases} \prod_{k=j/2}^j \frac{z_{j-1,\lambda+1}^2 - z_{k,j,\lambda}^2}{1 - z_{k,j,\lambda}^2} & \text{if } j \text{ even,} \\ z_{j-1,\lambda+1} \prod_{k=(j+1)/2}^j \frac{z_{j-1,\lambda+1}^2 - z_{k,j,\lambda}^2}{1 - z_{k,j,\lambda}^2} & \text{if } j \text{ odd.} \end{cases}$$

Let us focus first on the case j even, for simplicity. We can rewrite (77) more conveniently as

$$v_{j,\lambda} = \frac{z_{j,\lambda}^2 - z_{j-1,\lambda+1}^2}{1 - z_{j,\lambda}^2} \prod_{k=j/2}^{j-1} \frac{z_{j-1,\lambda+1}^2 - z_{k,j,\lambda}^2}{1 - z_{k,j,\lambda}^2}.$$

For $\delta \in (0, z_{j-1,\lambda+1})$ let

$$m(\delta, j, \lambda) := |\{k \in \{j/2, j\}; z_{k,j,\lambda} \geq \delta\}|$$

denote the number of zeros of $P_{j,\lambda}$ in the interval $(\delta, 1)$. Since the function $t \mapsto \frac{a^2 - t^2}{1 - t^2}$ is decreasing in $t \in (0, a)$, we have

Fact C.11. *We have the upper bound:*

$$(78) \quad v_{j,\lambda} \leq \frac{z_{j,\lambda}^2 - z_{j-1,\lambda+1}^2}{1 - z_{j,\lambda}^2} \inf_{\delta} \left(\frac{z_{j-1,\lambda+1}^2 - \delta^2}{1 - \delta^2} \right)^{m(\delta, j, \lambda)}.$$

Letting $\delta = z_{j/2,j,\lambda}$ the smallest positive root of $P_{j,\lambda}$ we have

$$(79) \quad v_{j,\lambda} \leq \frac{z_{j,\lambda}^2 - z_{j-1,\lambda+1}^2}{1 - z_{j,\lambda}^2} \left(\frac{z_{j-1,\lambda+1}^2 - z_{j/2,j,\lambda}^2}{1 - z_{j/2,j,\lambda}^2} \right)^{j/2}.$$

We can thus obtain an explicit control on $v_{j,\lambda}$ from bounds on the zeros of the Gegenbauer polynomials. We complement the upper bound on the largest root (Fact C.4) with lower bounds for all positive roots, as well as a sharp lower bound for its largest root [Dimitrov and Nikolov, 2010]:

Theorem C.12 (Upper and Lower bounds for Gegenbauer roots, [Dimitrov and Nikolov, 2010, Theorem 2]). *Let*

$$b_{j,\lambda} = j^3 + 2(\lambda - 1)j^2 - (3\lambda - 5)j + 4(\lambda - 1),$$

$$a_{j,\lambda} = 2(j + \lambda - 1)(j^2 + j(\lambda - 1) + 4(\lambda + 1)) \text{ and}$$

$$c_{j,\lambda} = j^2(j + 2\lambda)^2 + (2\lambda + 1)(j^2 + 2(\lambda + 3)j + 8(\lambda - 1)).$$

Then for every k, j, λ we have

$$(80) \quad \frac{b_{j,\lambda} - (j - 2)\sqrt{c_{j,\lambda}}}{a_{j,\lambda}} \leq z_{k,j,\lambda}^2 \leq \frac{b_{j,\lambda} + (j - 2)\sqrt{c_{j,\lambda}}}{a_{j,\lambda}}.$$

Theorem C.13 (Lower bound for largest root, [Driver and Jordaan, 2012, Section 2.3]).

$$(81) \quad z_{j,\lambda}^2 > 1 - \frac{(2\lambda + 1)(2\lambda + 3)}{(j - 1)(j + 2\lambda + 1) + (2\lambda + 1)(2\lambda + 3)} := 1 - \frac{g_{j,\lambda}}{h_{j,\lambda}}.$$

Rewriting Fact C.4 as $z_{j,\lambda}^2 \leq \frac{e_{j,\lambda}}{f_{j,\lambda}}$, with

$$e_{j,\lambda} = (j - 1)(j + 2\lambda - 2), \quad f_{j,\lambda} = (j + \lambda - 2)(j + \lambda - 1),$$

and using again the monotonicity of $t \mapsto \frac{t-p}{1-t}$ we can bound the first term in the RHS of (79) as

$$(82) \quad \frac{z_{j,\lambda}^2 - z_{j-1,\lambda+1}^2}{1 - z_{j,\lambda}^2} \leq \frac{e_{j,\lambda}/f_{j,\lambda} + g_{j-1,\lambda+1}/h_{j-1,\lambda+1} - 1}{1 - e_{j,\lambda}/f_{j,\lambda}}.$$

For $j, \lambda = \omega(1)$, we have

$$\begin{aligned} a_{j,\lambda} &\simeq 2j(j + \lambda)^2, & b_{j,\lambda} &\simeq j^2(j + 2\lambda), \quad \sqrt{c_{j,\lambda}} \simeq j(j + 2\lambda), \\ e_{j,\lambda} &\simeq j(j + 2\lambda), & f_{j,\lambda} &\simeq (j + \lambda)^2, \\ g_{j,\lambda} &\simeq 4\lambda^2, & h_{j,\lambda} &\simeq j(j + 2\lambda) + 4\lambda^2, \end{aligned}$$

and thus

$$(83) \quad \frac{z_{j,\lambda}^2 - z_{j-1,\lambda+1}^2}{1 - z_{j,\lambda}^2} \lesssim \frac{3j(j + 2\lambda)}{j(j + 2\lambda) + 4\lambda^2} \leq 3.$$

Therefore,

$$\begin{aligned} v_{j,\lambda} &\leq 3 \left(\frac{\frac{e_{j-1,\lambda+1}}{f_{j-1,\lambda+1}} - \frac{b_{j,\lambda} - (j-2)\sqrt{c_{j,\lambda}}}{a_{j,\lambda}}}{1 - \frac{b_{j,\lambda} - (j-2)\sqrt{c_{j,\lambda}}}{a_{j,\lambda}}} \right)^{j/2} \\ &\leq 3 \left(\frac{a_{j,\lambda}e_{j-1,\lambda+1} - f_{j-1,\lambda+1}(b_{j,\lambda} - (j-2)\sqrt{c_{j,\lambda}})}{f_{j-1,\lambda+1}(a_{j,\lambda} - b_{j,\lambda} + (j-2)\sqrt{c_{j,\lambda}})} \right)^{j/2} \\ &= 3 \left(\frac{2j(j + \lambda)^2j(j + 2\lambda) - (j + \lambda)^2(j^2(j + 2\lambda) - j^2(j + 2\lambda))}{(j + \lambda)^2(2j(j + \lambda)^2 - j^2(j + 2\lambda) + j^2(j + 2\lambda))} \cdot (1 + o_{j,\lambda}(1)) \right)^{j/2} \\ &\lesssim \left(\frac{j(j + 2\lambda)}{(j + \lambda)^2} \right)^{j/2} \\ (84) \quad &= \left[1 - \left(\frac{\lambda}{j + \lambda} \right)^2 \right]^{j/2}. \end{aligned}$$

As a direct consequence of (84), we immediately obtain Eqs (72), (73) and (74). The case where j is odd is treated analogously.

Let us now study the regime $j = \omega(\lambda^{3/2})$. Given $z \in \mathbb{C}$ with $|z| < 1$, Gegenbauer polynomials admit the following generating function [Watson, 1922, Section 3.32]:

$$(85) \quad \frac{1}{(1 - 2z \cos \theta + z^2)^\lambda} = \sum_{j \geq 0} \bar{P}_{j,\lambda}(\cos \theta) z^j .$$

From this generating function, the Cauchy integral formula leads to the following integral representation:

Fact C.14 ([Ursell, 2007, Eq (1.2)]). *For any $0 < \rho < 1$, we have*

$$(86) \quad \bar{P}_{j,\lambda}(\cos \theta) = \frac{1}{2\pi i} \oint_{|z|=\rho} \frac{dz}{(1 - 2z \cos \theta + z^2)^\lambda z^{j+1}} .$$

Assume $j = \Theta(\lambda^\alpha)$, with $\alpha > 3/2$. We are interested in the above representation for $\bar{\theta} = \arccos(z_{j-1,\lambda+1})$. From Theorem C.13, we have $z_{j-1,\lambda+1}^2 \geq 1 - d_{j,\lambda}/(2c_{j,\lambda})$, and thus

$$\bar{\theta}^2 \lesssim \frac{d_{j,\lambda}}{2c_{j,\lambda}} \simeq \frac{32\lambda^2 j^4}{16j^6} = \frac{2\lambda^2}{j^2} ,$$

so $\bar{\theta} = O(\lambda/j)$. Combining this upper bound with the lower bound obtained from Fact C.4 we have $\bar{\theta} = \Theta(\lambda/j)$.

Using $1 - \cos \theta \simeq \theta^2/2 \simeq \lambda^2/j^2$ and

$$(87) \quad \begin{aligned} |1 - 2z \cos \theta + z^2| &= |(1 - z)^2 + 2z(1 - \cos \theta)| \\ &\geq |1 - z|^2 - 2|z|(1 - \cos \theta) \\ &\geq 1 - \rho \left(2 + \Theta \left(\frac{\lambda^2}{j^2} \right) \right) + \rho^2 , \end{aligned}$$

we have

$$(88) \quad |\bar{P}_{j,\lambda}(\cos \theta)| \leq \inf_{0 < \rho < 1} |\rho|^{-(j+1)} (1 - \rho(2 + c\lambda^2/j^2) + \rho^2)^{-\lambda} := g(\rho) .$$

Optimizing the RHS over ρ we obtain $\rho^* = \frac{j - (\sqrt{2}-1)\lambda}{j+2\lambda}$; substituting, we obtain

$$(89) \quad g(\rho^*) \simeq e^{-(1+\sqrt{2})\lambda} \left(\frac{j + 2\lambda}{\lambda(1 + \sqrt{2})} \right)^{2\lambda} .$$

As a result, it follows that

$$(90) \quad P_{j,\lambda}(\cos \theta) = \bar{P}_{j,\lambda}(\cos \theta) \frac{j!(2\lambda - 1)!}{(2\lambda + j - 1)!}$$

satisfies, for $\theta = \Theta(\lambda/j)$ and $j = \omega(\lambda^{3/2})$,

$$\begin{aligned} \log |P_{j,\lambda}(\cos \theta)| &\simeq j \log j - j + 2\lambda \log(2\lambda) - 2\lambda - (j + 2\lambda) \log(2\lambda + j) + 2\lambda + j \\ &\quad - (1 + \sqrt{2})\lambda + 2\lambda \log(j + 2\lambda) - 2\lambda \log(\lambda(1 + \sqrt{2})) \\ (91) \quad &\simeq - (1 + \sqrt{2})\lambda, \end{aligned}$$

where we have used Stirling's approximation. This proves Eq (75) and completes the proof of Lemma C.7. $\blacksquare$

Proof of Lemma C.8. We have, using Fact C.3, that

$$\begin{aligned} \mathbb{E}_\rho[\mathbb{E}_{u_d}(\phi^{(r)'})^2] &= \frac{\Omega_{d-1}}{(d-1)^2 \Omega_{d-2}} \sum_j \mathbb{E}_\rho [\bar{\alpha}_{j,d,r}^2 (j(j+d-2))^2] \\ (92) \quad &\geq \frac{\Omega_{d-1}}{\Omega_{d-2}} \sum_j j^2 \mathbb{E}_\rho [\bar{\alpha}_{j,d,r}^2]. \end{aligned}$$

And we can upper bound via

$$\begin{aligned} \mathbb{E}_\rho[\mathbb{E}_{u_d}(\phi^{(r)'})^2] &= \mathbb{E}_\rho[r^2 \mathbb{E}_{u_d}((\phi')^{(r)})^2] \\ &= \mathbb{E}_\rho \mathbb{E}_{x_1 \|x\|=r} [r^2 (\phi(x_1)')^2] \\ (93) \quad &\leq \sqrt{\mathbb{E}_\rho[r^4] \mathbb{E}_\eta(\phi')^4}, \end{aligned}$$

where this last line is finite by our assumptions on ϕ and ρ ,

so from (92) we conclude that

$$(94) \quad \sum_j j^2 \beta_{j,d} \leq \frac{\Omega_{d-2}}{\Omega_{d-1}} \sqrt{\mathbb{E}_\rho[r^4] \mathbb{E}_\eta(\phi')^4}.$$

$\blacksquare$

APPENDIX D. THE LPG PROPERTY IN THE NON-SYMMETRIC CASE: PROOFS OF SECTION 4.2

D.1. Proof of Proposition 4.11.

Assumption 4.9 (Regularity of link function). *We assume that ϕ, ϕ' are both B -Lipschitz, and that $\phi''(t) = O(1/t)$.*

Assumption 4.10 (Subgaussianity). *The data distribution ν is M -subgaussian: for any $v \in \mathcal{S}_{d-1}$, we have $\|x \cdot v\|_{\psi_2} \leq M$, where $\|z\|_{\psi_2} := \inf\{t > 0; \mathbb{E}[\exp(z^2/t^2)] \leq 2\}$ is the Orlicz-2 norm.*

Proposition 4.11 (Uniform gradient approximation). *Under Assumptions 4.9 and 4.10, for all $\theta \in \mathcal{S}_{d-1}$,*

$$(22) \quad \Delta_{\nabla L}(\theta) = (1 - m^2)O\left(\widetilde{W}_{1,2}(\nu, \gamma) \log(\widetilde{W}_{1,2}(\nu, \gamma)^{-1})\right)$$

where the $O(\cdot)$ notation only hides constants appearing in Assumptions 4.9 and 4.10.

Proof. Recall the notation $\phi_\theta(x) = \phi(\langle x, \theta \rangle)$. Let $v = \theta^* - m\theta$. From the definition, we have that

$$\begin{aligned} \langle \nabla_\theta^\mathbb{S} L(\theta), \theta^* \rangle &= 2\mathbb{E}_\nu [\phi'_\theta(\phi_\theta - \phi_{\theta^*})(x \cdot v)] \\ &:= \mathbb{E}_\nu [g_{\theta, \theta^*}]. \end{aligned} \quad (95)$$

Since $\mathbb{E}_\gamma [g_{\theta, \theta^*}]$ is precisely $\bar{\ell}'(m)(1 - m^2)$, we need to establish that

$$\sup_\theta |\mathbb{E}_\nu g_{\theta, \theta^*} - \mathbb{E}_\gamma g_{\theta, \theta^*}| \leq C\sqrt{1 - m^2} \widetilde{W}_{1,2}(\nu, \gamma) \left(\log \widetilde{W}_{1,2}(\nu, \gamma) \right)^2. \quad (96)$$

Fix θ and let P_{θ, θ^*} be the orthogonal projection onto the subspace spanned by θ, θ^* . For $R > 0$ we consider $A_R = \{x \in \mathbb{R}^d; \|P_{\theta, \theta^*} x\| \leq R\}$.

$$\begin{aligned} |\mathbb{E}_\nu g_{\theta, \theta^*} - \mathbb{E}_\gamma g_{\theta, \theta^*}| &= \left| \int g_{\theta, \theta^*}(x)(\nu(dx) - \gamma(dx)) \right| \\ &\leq \underbrace{\left| \int_{x \in A_R} g_{\theta, \theta^*}(x)(\nu(dx) - \gamma(dx)) \right|}_{T_a} + \underbrace{\left| \int_{x \notin A_R} g_{\theta, \theta^*}(x)(\nu(dx) - \gamma(dx)) \right|}_{T_b}. \end{aligned} \quad (97)$$

Let us first bound T_a . Denote by $v = \theta^* - m\theta$, with $\|v\|^2 = 1 - m^2$. Since ϕ and ϕ' are Lipschitz and $|\phi''| \leq O((1 + t)^{-1})$ by Assumption 4.9, we have that

$$\nabla_x g_{\theta, \theta^*}(x) = \phi''_\theta(\phi_\theta - \phi_{\theta^*})x^\top v\theta + \phi'_\theta(\phi'_\theta\theta - \phi'_{\theta^*}\theta^*)x^\top v + \phi'_\theta(\phi_\theta - \phi_{\theta^*})v \quad (98)$$

satisfies

$$\begin{aligned} \|\nabla_x g_{\theta, \theta^*}(x)\| &\leq 2\|v\|C\text{Lip}(\phi)R + 4\|v\|\text{Lip}(\phi)^2R \\ &\leq C\|v\|R, \end{aligned} \quad (99)$$

and as a result we have that g_{θ, θ^*} is $C\|v\|R$ -Lipschitz when restricted to A_R , and thus

$$T_a \leq CR\|v\|\widetilde{W}_{1,2}(\nu, \mu). \quad (100)$$

Let us now control the tail T_b . Since $x^\top v$ is $\sqrt{2}M\|v\|$ -subgaussian and ϕ is Lipschitz, we have that $z = |g_{\theta, \theta^*}(x)|$ is $\tilde{M}\|v\|$ -subexponential where $\tilde{M}$ only depends on M and L . It follows that

$$\begin{aligned} T_b &\leq R(\mathbb{P}_\nu(z \geq R) + \mathbb{P}_\gamma(z \geq R)) \\ &\leq R \exp\left(-\frac{\beta}{\|v\|}R\right), \end{aligned} \quad (101)$$

where β is a constant that depends only on $\tilde{M}$. As a result, we have

$$|\mathbb{E}_\nu g_{\theta, \theta^*} - \mathbb{E}_\gamma g_{\theta, \theta^*}| \leq \inf_{R>0} \left(CR\|v\|\widetilde{W}_{1,2}(\nu, \mu) + R \exp\left(-\frac{\beta}{\|v\|}R\right) \right). \quad (102)$$

Setting

$$R = -\|v\|\beta^{-1} \log((C\|v\|\widetilde{W}_{1,2}(\nu, \gamma)))$$

we obtain

$$\begin{aligned} |\mathbb{E}_\nu g_{\theta, \theta^*} - \mathbb{E}_\gamma g_{\theta, \theta^*}| &\leq 2C(1 - m^2)\beta^{-1} \left| \log(C\|v\|\widetilde{W}_{1,2}(\nu, \gamma)) \right| \widetilde{W}_{1,2}(\nu, \gamma) \\ (103) \quad &\leq (1 - m^2)O\left(\widetilde{W}_{1,2}(\nu, \gamma) \log(\widetilde{W}_{1,2}(\nu, \gamma)^{-1})\right), \end{aligned}$$

as claimed. ■

D.2. Proof of Proposition 4.14. We leverage Proposition 4.11 and the fact that if ϕ has information exponent $s = 2$, then $\bar{\ell}'(m) \simeq m$ for small m .

We need to show that for $b = \Theta(\log d)$ we have

$$(104) \quad \langle \nabla_\theta L(\theta), \theta^* \rangle \geq C \left(m - \frac{b}{\sqrt{d}} \right), \text{ for } \frac{b}{\sqrt{d}} \leq m \leq \frac{1}{2},$$

as well as

$$(105) \quad \langle \nabla_\theta L(\theta), \theta^* \rangle \geq C'(1 - m^2)$$

for $m \geq \frac{1}{2}$.

From (103) and $\widetilde{W}_{1,2}(\nu, \gamma) \leq C/\sqrt{d}$, we obtain

$$\begin{aligned} \langle \nabla_\theta L(\theta), \theta^* \rangle &= \bar{\ell}'(m)(1 - m^2) + \langle \nabla_\theta L(\theta), \theta^* \rangle - \bar{\ell}'(m)(1 - m^2) \\ &\geq 2\alpha_2^2 m(1 - m^2) - (1 - m^2)\tilde{C}\widetilde{W}_{1,2}(\nu, \gamma) \log(\widetilde{W}_{1,2}(\nu, \gamma)^{-1}) \\ &\geq \left(\alpha_2^2 m - \tilde{C} \frac{C}{\sqrt{d}} \log(\sqrt{d}/C) \right) (1 - m^2) \\ &\geq \alpha_2^2 \left(m - \tilde{C} \frac{C}{\alpha_2^2 \sqrt{d}} \log(\sqrt{d}/C) \right) (1 - m^2) \\ (106) \quad &\geq \alpha_2^2 \left(m - \frac{\log d^{C'}}{\sqrt{d}} \right) (1 - m^2), \end{aligned}$$

which proves (104) and (105).

D.3. Proof of Proposition 4.16.

Assumption 4.15 (Additional Regularity in third derivatives). ϕ admits four derivatives bounded by L , with $|\phi^{(3)}(t)| = O(1/t)$ and $|\phi^{(4)}(t)| = O(1/t^2)$. Moreover, the third moment of the data distribution is finite: $\tau_3 = \mathbb{E}_{t \sim \eta}[t^3] < \infty$.

Proposition 4.16 (Stein's method for product measure). *Let $\chi(\theta, \theta^*) := \|\theta\|_4^2 + \|\theta^*\|_4^2$. Under Assumptions 4.9, 4.10 and 4.15, there exists a universal constant $C = C(M, B, \tau_3)$ such that*

$$(23) \quad \Delta_L(\theta) \leq C\chi(\theta, \theta^*) , \text{ and } \Delta_{\nabla L}(\theta) \leq C\sqrt{1 - m^2}\chi(\theta, \theta^*) .$$

Proof. Recall the notation $\phi_\theta(x) = \phi(\langle x, \theta \rangle)$, and, using $v = \theta^* - m\theta$,

$$(107) \quad h_{\theta, \theta^*}(x) := \phi_\theta^2 - 2\phi_\theta\phi_{\theta^*} ,$$

$$(108) \quad g_{\theta, \theta^*}(x) := 2\phi'_\theta(\phi_\theta - \phi_{\theta^*})(x \cdot v) ,$$

so that

$$(109) \quad \Delta_L(\theta) = \mathbb{E}_\nu[h_{\theta, \theta^*}(x)] - \mathbb{E}_\gamma[h_{\theta, \theta^*}(x)] ,$$

$$(110) \quad \Delta_{\nabla L}(\theta) = \mathbb{E}_\nu[g_{\theta, \theta^*}(x)] - \mathbb{E}_\gamma[g_{\theta, \theta^*}(x)] .$$

The result is obtained via the following Stein coupling method for product measures:

Theorem D.1 (Stein Coupling, [Röllin, 2013, Theorem 3.1]). *Let X be a d -dimensional random vector of independent coordinates, such that $\mathbb{E}X = 0$, $\mathbb{E}[XX^\top] = I_d$ and $\mathbb{E}|X_i|^3 = \tau_i^3 < \infty$. If Z is a standard Gaussian random vector, and $h : \mathbb{R}^d \rightarrow \mathbb{R}$ is three-times differentiable, then*

$$(111) \quad |\mathbb{E}h(X) - \mathbb{E}h(Z)| \leq \frac{5}{6} \sum_{i=1}^d \tau_i^3 \|\partial_{x_i}^3 h\|_\infty .$$

We verify that, thanks to the decay assumptions in Assumption 4.15, we have

$$(112) \quad \partial_{x_i}^3 g_{\theta, \theta^*}(x) = \lambda_1(x)\theta_i^3 + \lambda_2(x)\theta_i^2\theta_i^* + \lambda_3(x)\theta_i(\theta_i^*)^2 + \lambda_4(x)(\theta_i^*)^3 ,$$

$$(113) \quad \partial_{x_i}^3 h_{\theta, \theta^*}(x) = \lambda_5(x)\theta_i^3 + \lambda_6(x)\theta_i^2\theta_i^* + \lambda_7(x)\theta_i(\theta_i^*)^2 + \lambda_8(x)(\theta_i^*)^3 ,$$

where

$$(114) \quad \sup_{k \in \{1, 2, 3, 4\}} |\lambda_k(x)| \leq C\|v\| , \quad \sup_{k \in \{5, 6, 7, 8\}} |\lambda_k(x)| \leq \tilde{C} .$$

Observing by Cauchy-Schwartz that

$$\begin{aligned} \max \left\{ \sum_i |\theta_i|^2 |\theta_i^*|, \sum_i |\theta_i|^3 \right\} &\leq \|\theta\|_4^2 , \\ \max \left\{ \sum_i |\theta_i^*|^2 |\theta_i|, \sum_i |\theta_i^*|^3 \right\} &\leq \|\theta^*\|_4^2 , \end{aligned}$$

we obtain from Theorem D.1 that

$$\Delta_L(\theta) \leq C\chi(\theta, \theta^*) , \quad \Delta_{\nabla L}(\theta) \leq C'\|v\|\chi(\theta, \theta^*) ,$$

as claimed. ■