The geometry of off-the-grid compressed sensing

Clarice Poon^{*}, Nicolas Keriven[†] Gabriel Peyré[‡]

February 12, 2020

Abstract

This paper presents a sharp geometric analysis of the recovery performance of sparse regularization. More specifically, we analyze the BLASSO method which estimates a sparse measure (sum of Dirac masses) from randomized sub-sampled measurements. This is a "continuous", often called off-the-grid, extension of the compressed sensing problem, where the ℓ^1 norm is replaced by the total variation of measures. This extension is appealing from a numerical perspective because it avoids to discretize the the space by some grid. But more importantly, it makes explicit the geometry of the problem since the positions of the Diracs can now freely move over the parameter space. On a methodological level, our contribution is to propose the Fisher geodesic distance on this parameter space as the canonical metric to analyze super-resolution performances in a way which is invariant to reparameterization of this space. While previous works express recovery conditions using a flat Euclidean distance, switching to the Fisher metric allows us to take into account measurement operators which are not translation invariant, which is crucial for applications such as Laplace inversion in imaging, Gaussian mixtures estimation and training of multilayer perceptrons with one hidden layer. On a theoretical level, our main contribution shows that if the Fisher distance between spikes is larger than a Rayleigh separation constant, then the BLASSO recovers in a stable way a stream of Diracs, provided that the number of measurements is proportional (up to log factors) to the number of Diracs. We measure the stability using an optimal transport distance constructed on top of the Fisher geodesic distance. Our result is (up to log factor) sharp and does not require any randomness assumption on the amplitudes of the underlying measure. Our proof technique relies on an infinite-dimensional extension of the so-called "golfing scheme" which operates over the space of measures and is of general interest.

1 Introduction

Sparse regularization, and in particular convex approaches based on ℓ^1 minimization, is one of the workhorses to ill-posed linear inverse models. It finds numerous applications ranging from signal processing [19] to machine learning [55]. In many situations, it makes sense to consider a "continuous" counterpart to ℓ^1 minimization, which avoids gridding the parameter space, thus enabling more efficient solvers and a sharper theoretical analysis. The most natural continuous extension encodes the positions and amplitudes of the sought after solution into a Radon measure, so that the ℓ^1 norm is replaced by the total variation (total mass) of the measure. A measure is then naturally said to be "sparse" when it is a *sum of Diracs* at the desired positions and amplitudes. The corresponding infinite dimensional optimization problem is called BLASSO in [27] following theoretical works on spectral extrapolation [7]. This setting of optimization on measures has also been considered in the inverse problem community [9]. Successful examples of applications of such "off-the-grid methods" include single-molecule fluorescent imaging [8], spikes sorting in neurosciences [35], mixture model estimation [39] and training shallow neural networks [5].

^{*}Department of Mathematical Sciences, University of Bath, cmhsp20@bath.ac.uk

[†]CNRS and GIPSA-lab, Univ. Grenoble Alpes, Grenoble INP nicolas.keriven@gipsa-lab.grenoble-inp.fr

[‡]CNRS and DMA, Ecole Normale Supérieure, 45 rue d'Ulm, F-75230 PARIS cedex 05, FRANCE, gabriel.peyre@ens.fr

1.1 Sparse spikes recovery using the BLASSO

Observation model. We consider the general problem of estimating a complex-valued unknown Radon measure $\mu_0 \in \mathcal{M}(\mathcal{X})$ defined over some metric space \mathcal{X} from a small number m of randomized linear observations $y \in \mathbb{C}^m$. In this paper, \mathcal{X} will either be a connected bounded open subset of \mathbb{R}^d or the d-dimensional torus \mathbb{T}^d , even though some of our results extend beyond this case. We define the product between a complex-valued continuous function $f \in \mathscr{C}(\mathcal{X})$ and complex-valued measure $\mu \in \mathcal{M}(\mathcal{X})$ as $\langle f, \mu \rangle_{\mathcal{M}} \stackrel{\text{def.}}{=} \int_{\mathcal{X}} \overline{f(x)} d\mu(x)$. The (forward) measurement operator $\Phi : \mathcal{M}(\mathcal{X}) \mapsto \mathbb{C}^m$ that we consider in this paper is of the form

$$\Phi\mu \stackrel{\text{def.}}{=} \frac{1}{\sqrt{m}} \left(\langle \varphi_{\omega_k}, \mu \rangle_{\mathcal{M}} \right)_{k=1}^m \tag{1}$$

where $(\omega_1, \ldots, \omega_m)$ are parameters identically and independently distributed according to a probability distribution $\Lambda(\omega)$ over some space Ω , and $\varphi_{\omega} : \mathcal{X} \to \mathbb{C}$ are smooth functions parameterized by ω . We further assume that φ_{ω} is normalized, that is $\mathbb{E}_{\omega \sim \Lambda}[|\varphi_{\omega}(x)|^2] = 1$ for all $x \in \mathcal{X}$. Our observations are of the form

$$y = \Phi(\mu_0 + \tilde{\mu}_0) + w, \qquad (2)$$

where $\mu_0 = \sum_{i=1}^s a_i \delta_{x_i}$ with $(x_i, a_i) \in \mathcal{X} \times \mathbb{C}$ is the *s*-sparse measure we are interested in, $\tilde{\mu}_0 \in \mathcal{M}(\mathcal{X})$ accounts for modelling error, and $w \in \mathbb{C}^m$ is measurement noise. In the rest of the paper, we naturally assume that the support of $\tilde{\mu}_0$ does not include any of the x_i .

BLASSO. An increasingly popular method to estimate such a sparse measure corresponds to solving a infinite-dimensional analogue of the Lasso regression problem

$$\min_{\mu \in \mathcal{M}(\mathcal{X})} \frac{1}{2} \left\| \Phi \mu - y \right\|_{2}^{2} + \lambda |\mu|(\mathcal{X}). \tag{$\mathcal{P}_{\lambda}(y)$}$$

Following [27], we call this method the BLASSO (for Beurling-Lasso). Here $|\mu|(\mathcal{X})$ is the so-called total variation (or total mass) of the measure μ , and is defined as

$$|\mu|(\mathcal{X}) \stackrel{\text{\tiny def.}}{=} \sup \left\{ \operatorname{Re}\left(\langle f, \mu \rangle_{\mathcal{M}} \right) \; ; \; f \in \mathscr{C}(\mathcal{X}), \left\| f \right\|_{\infty} \leqslant 1 \right\}$$

Note that on unbounded \mathcal{X} , one needs to impose that f vanishes at infinity. If $\mathcal{X} = \{x_i\}_i$ is a finite space, then this would correspond to the classical finite-dimensional Lasso problem [55], because $|\mu|(\mathcal{X}) = ||a||_1 \stackrel{\text{def.}}{=} \sum_i |a_i|$ where $a_i = \mu(\{x_i\})$. Similarly, when \mathcal{X} is possibly infinite but $\mu = \sum_i a_i \delta_{x_i}$, one also has that $|\mu|(\mathcal{X}) = ||a||_1$.

1.2 Previous works

From a theoretical perspective, understanding the performance of this approach corresponds to establishing a "Rayleigh criterion", which is the minimum allowable separation distance between two spikes $\min_{i \neq j} ||x_i - x_j||$ for the method to recover them from linear measurements. The first result in this direction is due to Candès and Fernandez-Granda [15], who prove that for Fourier measurements, this separation distance is (almost) equal to the inverse of the maximum sample frequency. This results has been extended to provide robustness to noise [14, 37, 4, 32] and to cope with more general measurement operators [6]. All these previous theoretical works however strongly rely on the translation invariance of the linear operator (Fourier measurements or convolutions) and the underlying domain (either Euclidean space or the periodic torus). Applying these results to spatially varying operator (such as for instance when imaging with non-stationary point spread functions) generally leads to overly pessimistic minimum separation condition.

In parallel, it has been shown by Tang et al [54] that the recovery guarantees of Candès and Fernandez-Granda [15] remain valid with high probability when only a small number of (Fourier) measurements are *randomly selected*, of the order (up to log factors) of the sparsity of the underlying measure. However, this result is only valid under a *random signs assumption* on the amplitudes of the sought-after Dirac masses, which is a well-known assumption in classical discrete compressed sensing [17, 31] but appears somewhat

unrealistic. While more detailed conditions can be derived when the amplitudes a_i are all assumed real and positive [29], in the general case $a_i \in \mathbb{C}$ it was not known until this paper whether the random signs assumption could be removed.

Although it is not the topic of this paper, let us note that lifting the minimum separation condition requires to impose positivity of the weights [27, 50] and the price to pay is an explosion of the instabilities as spikes cluster together [29].

Numerical solvers and alternative approaches. The focus of this paper is on the theoretical analysis of the performance BLASSO method, not on the development and analysis of efficient numerical solvers. Although the BLASSO problem is infinite dimensional, there are efficient numerical solvers that use the fact that the sought-after sparse solution is parameterized by a small number of parameters (positions and amplitudes of the spikes). This open the door to algorithms which do not scale with some grid size, and hence can scale beyond 1-D and 2-D problems. Let us mention in particular: (i) discretization on a grid [53, 33], (ii) semi-definite programming (SDP) relaxation using Lasserre Hierarchy [15, 28], (iii) Frank-Wolfe and its variants [9, 8], (iv) non-convex particle flows [20]. We also emphasize that the BLASSO is by no means the only method for estimating sparse measures in an off-the-grid setup. Let us, among other approaches, cite Prony-type spectral methods such as MUSIC and ESPRIT [51, 48, 42] and non-convex approaches for instance based on ℓ^0 or greedy minimization (see for instance [39, 52] for recent contributions in this direction). In practice, these methods often surpasses BLASSO in term of performance, in particular when the noise is small and the spikes tends to cluster so that the minimum separation distance condition does not hold. A rule of thumb is that ℓ^1 -regularization is however good baseline, which benefit from both efficient and stable numerical solvers and an in-depth theoretical analysis which leverage the convexity of the problem.

1.3 Contributions

The goal of this paper is twofold: remove the random signs assumption of Tang et al [54] while still keeping a sharp number of random measurements, and extend the framework to encompass non-translation invariant operators in a natural manner with improved separation conditions. The former is achieved by extending the so-called golfing scheme [41, 16] to the infinite-dimensional setting, while the latter is done through a particular Riemannian geometric framework, defined by the metric tensor associated to the covariance kernel of the measurement operator. We will show that, by imposing a minimal separation between Diracs with respect to the geodesic distance associated to this tensor, previous strategies can be naturally extended.

Informally, our main result reads as follows. Define the *limit covariance kernel* $K(x, x') \stackrel{\text{def}}{=} \mathbb{E}_{\omega} \overline{\varphi_{\omega}(x)} \varphi_{\omega}(x')$, which measures how much two Diracs at x and x' interact with each other in the large samples limit as $m \to \infty$, and assume that K is real-valued (primary examples include the Gaussian kernel, or the so-called Jackson kernel used in [15]). Define the metric tensor $\mathfrak{g}_x \stackrel{\text{def}}{=} \nabla_1 \nabla_2 K(x, x) \in \mathbb{R}^{d \times d}$, where ∇_i indicates the gradient with respect to the *i*th variable, and assume that for all $x \in \mathcal{X}$ it is a positive definite matrix. Finally, define the associated geodesic distance $\mathfrak{d}_{\mathfrak{g}}(x, x') = \inf_{\gamma} \int_0^1 \sqrt{\gamma'(t)^\top \mathfrak{g}_{\gamma(t)} \gamma'(t)} dt$, where the infimum is taken over all continuous path $\gamma : [0, 1] \to \mathcal{X}$ such that $\gamma(0) = x$ and $\gamma(1) = x'$ (more details about this geodesic distance are given in Section 3.1). Denote by $\mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x; r)$ the ball of radius r centered on x, for the metric $\mathfrak{d}_{\mathfrak{g}}$.

Theorem 1 (Main result, informal). Let $R_{\mathcal{X}} \stackrel{\text{def.}}{=} \sup_{x,x'} \mathfrak{d}_{\mathfrak{g}}(x,x')$. Under some assumptions on the kernel K (see Assumption 1 in Sec. 4) and features φ_{ω} (see Assumption 2 in Sec. 5), there are constants $r, \Delta > 0$, that only depends on K, and $C_1, C_2 > 0$ which depend on K and the regularity of φ_{ω_k} (up to 2nd order), such that the following holds. Suppose that y is of the form (2) with $\min_{i\neq j} \mathfrak{d}_{\mathfrak{g}}(x_i, x_j) \geq \Delta$ and

$$m \ge C_1 \cdot s \cdot \log(s) \log((C_2 R_{\mathcal{X}})^d / \rho). \tag{3}$$

Then with probability $1 - \rho$, when $||w|| \leq \delta$ and $\lambda \sim \frac{\delta}{\sqrt{s}}$, any solution $\hat{\mu}$ to $(\mathcal{P}_{\lambda}(y))$ satisfies

$$\mathcal{T}^{2}_{\mathfrak{d}_{\mathfrak{g}}}\left(\sum_{j=1}^{s}\hat{A}_{j}\delta_{x_{j}}, |\hat{\mu}|\right) \lesssim \sqrt{s}\delta + |\tilde{\mu}_{0}|\left(\mathcal{X}\right) \quad and \quad \max_{j=1}^{s}|\hat{a}_{j} - a_{j}| \lesssim \sqrt{s}\delta + |\tilde{\mu}_{0}|\left(\mathcal{X}\right), \tag{4}$$

where $\hat{A}_j \stackrel{\text{def.}}{=} |\hat{\mu}| (\mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_j; r)), \hat{a}_j \stackrel{\text{def.}}{=} \hat{\mu}(\mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_j; r)), \text{ and } \mathcal{T}_{\mathfrak{d}_{\mathfrak{g}}} \text{ is the partial optimal transport distance associated to } \mathfrak{d}_{\mathfrak{g}} \text{ (see Def. 1).}$

Let us comment a bit on this result. On a technical level, the most salient feature of Theorem 1 is that, up to log factors, the bound (3) is *linear* in the sparsity of the underlying measure. This improves over the best known result of Tang et al [54], since in our case we do not require the random signs assumption.

On a methodological level, the assumptions on the kernel K(x, x') mainly state that it must *decrease* sufficiently when x and x' are far apart, or, in other words, that the *coherence* between $\Phi \delta_x$ and $\Phi \delta_{x'}$ must be low. The main novelty of our approach is that we measure this separation in term of the geodesic metric $\mathfrak{d}_{\mathfrak{g}}$, which allows to account for non-translation invariant kernels in an intrinsic and natural manner. The assumptions on the features φ_{ω} , which are more technical in nature, relates to their regularity and the boundedness of their various derivatives.

Concerning the recovery bound (4), the first part states that the measure $\hat{\mu}$ concentrates around the true positions of the Diracs, while the second part guarantees that the complex amplitudes of $\hat{\mu}$ around the Diracs are close to their true values. The discrepancy in the first part is measured in terms of a *partial optimal transport* distance associated to $\mathfrak{d}_{\mathfrak{g}}$ (Def. 1 in Sec. 3).

Finally, the constants C_1, C_2 that appear in (3) can depend (generally polynomially) on the dimension d but not on the sparsity s. As we will see in Section 5 and the detailed version of Theorem 1 (Theorem 3), the bound (3) is actually valid when we suppose the features φ_{ω} and their derivatives to be *uniformly bounded* for all x and ω . When this is not the case, we will be able to relax this assumption, similar to the notion of stochastic incoherence [16] in compressed sensing. As a result, m can actually appear in C_1, C_2 , generally in a logarithmic form (see examples in Section 2), which only adds logarithmic terms in s and d in the final number of measurements.

Outline of the paper. The paper is organized as follows. In Section 2 we give example applications of Theorem 1, including non-translation invariant frameworks such as Laplace measurements used in microscopy [30]. In Section 3, we introduce our Riemannian geometry framework and prove intermediate recovery results based on the existence of a so-called *non-degenerate dual certificate*, which is known in the literature to be the key object in the analysis of the BLASSO model. In Section 4, we study in more detail the relationship between the minimal separation condition and the covariance kernel. We prove that, under some conditions on K, in the limit $m \to \infty$, one can indeed prove the existence of a non-degenerate dual certificate when minimal separation is imposed with respect to $\mathfrak{d}_{\mathfrak{g}}$. Finally, in Section 5, we state our main result with finite number of measurements m (Theorem 3, which is a detailed version of Theorem 1). Section 6 is dedicated to its proof using an infinite-dimensional extension of the celebrated golfing scheme [16], with technical computations in the appendix.

Relationship to our previous work [45] This article is a substantially extended version of the conference publication [45]. The results of Section 4 are in most part already published (under slightly more restrictive assumptions) in this conference paper. The remainder of the paper is however entirely novel. We remove the random signs assumption of [45] thanks to a new proof technique with the golfing scheme. Furthermore, the results in [45] are restricted to the small noise setting and focus on exact support stability, while we study here arbitrary noise levels and establish more general stability bounds in terms of optimal transport distances.

Notations. Given $n \in \mathbb{N}$, we denote by $[n] \stackrel{\text{def.}}{=} \{1, 2, \dots, n\}$ the first *n* integers. We write 1_n to denote the vector of length *n* whose entries are all 1's, and 0_n to denote the vector of length *n* whose entries are all 0's.

Given two matrices A and B, we write $A \prec B$ to mean that B - A is positive definite and $A \preceq B$ to mean that B - A is semi-positive definite. Given two positive numbers a, b, we write $a \leq b$ to mean that there exists some universal constant C > 0 so that $a \leq Cb$. Given $(\mathcal{X}, \mathfrak{d})$ a metric space, $x \in \mathcal{X}$ and r > 0, we define $\mathcal{B}_{\mathfrak{d}}(x; r) \stackrel{\text{def.}}{=} \{z \in \mathcal{X} ; \mathfrak{d}(x, z) < r\}$ the ball centered on x of radius r, or just $\mathcal{B}_{\|\cdot\|}(r) = \{z \in \mathcal{X} ; \|z\| < r\}$ the ball centered on 0 for a norm $\|\cdot\|$.

We write $\|\cdot\|_p$ to denote the ℓ_p norm, and $\|\cdot\|$ without any subscript denotes the spectral norm for matrices or ℓ_2 norm for vectors. For any norm $\|\cdot\|_X$ on vectors, the corresponding matrix norm is $\|A\|_{X \to Y} = \sup_{\|x\|_X = 1} \|Ax\|_Y$ and $\|A\|_X = \|A\|_{X \to X}$ for short. Given a vector $x \in \mathbb{C}^{sd}$ decomposed in blocks $x = [x_1^\top, \ldots, x_s^\top]^\top$ with $x_i \in \mathbb{C}^d$, where s and d will always be defined without ambiguity, we define the block norm $\|x\|_{\text{block}} \stackrel{\text{def.}}{=} \max_{1 \leq i \leq s} \|x_i\|$. Given a vector $x \in \mathbb{C}^{s(d+1)}$ decomposed as $x = [x_0^\top, X_1^\top, \ldots, X_s^\top]^\top$ where $x_0 \in \mathbb{C}^s$ and $X_j \in \mathbb{C}^d$, we define $\|x\|_{\text{Block}} \stackrel{\text{def.}}{=} \max(\|x_0\|_{\infty}, \max_{j=1}^s \|X_j\|_2)$. For a complex number a, its sign is denoted by $\operatorname{sign}(a) = \frac{a}{|a|}$. Given a complex-valued measure $\mu \in \mathcal{M}(\mathcal{X})$

For a complex number a, its sign is denoted by $\operatorname{sign}(a) = \frac{a}{|a|}$. Given a complex-valued measure $\mu \in \mathcal{M}(\mathcal{X})$ and complex-valued continuous function $f \in \mathscr{C}(\mathcal{X})$, we recall that $\langle f, \mu \rangle_{\mathcal{M}} \stackrel{\text{def.}}{=} \int_{\mathcal{X}} \overline{f(x)} d\mu(x)$. For two complex vectors v and w, $\langle v, w \rangle_2 \stackrel{\text{def.}}{=} v^* w$, where $v^* = \overline{v}^\top$ denotes conjugate transpose.

2 Examples

In this section, we illustrate Theorem 1 for some special cases of practical interest in imaging and machine learning. The following statements are obtained by bounding the constants in Theorem 3 in Section 5 (the detailed version of Theorem 1). These computations, which can be somewhat verbose, are delayed to Appendices C, D and E.

Off-the-grid Compressed Sensing. Off-the-grid Compressed sensing, initially introduced in the special case of 1-D Fourier measurements on $\mathcal{X} = \mathbb{T} = \mathbb{R}/\mathbb{Z}$ by [54], corresponds to Fourier measurements of the form (1). This is a "continuous" analogous of the celebrated compressed sensing line of works [17, 31]. We give a multi-dimensional version below.

Let $f_c \in \mathbb{N}$ with $f_c \ge 128$ (for simplicity) and $\mathcal{X} = \mathbb{T}^d$ the *d*-dimensional torus. Let $\varphi_{\omega}(x) \stackrel{\text{def.}}{=} e^{i2\pi\omega^\top x}$, $\Omega \stackrel{\text{def.}}{=} \{\omega \in \mathbb{Z}^d ; \|\omega\|_{\infty} \le f_c\}$, and $\Lambda(\omega) = \prod_{j=1}^d g(\omega_j)$ where $g(j) = \frac{1}{f_c} \sum_{k=\max(j-f_c,-f_c)}^{\min(j+f_c,f_c)} (1-|k/f_c|)(1-|(j-k)/f_c|)$. The Fisher metric is, up to a constant C, the Euclidean metric $\mathfrak{d}_g(x,x') = Cf_c \|x-x'\|$. Provided that $\min_{i\neq j} \|x_i - x_j\| \gtrsim \frac{d^{\frac{1}{2}s^{\frac{1}{4}}}{f_c}}{s}$, stable recovery is guaranteed with

$$m\gtrsim d^2s\left(\log(s)\log\left(\frac{s}{\rho}\right)+\log\left(\frac{(sf_cd)^d}{\rho}\right)\right).$$

Note that, compared to the uni-dimensional case, the minimal separation Δ depends on s: this could actually be replaced by a bound exponential in the dimension d, which we prefer not to do here. Indeed, during the proof, one must bound a quantity of the form $\sum_{i=2}^{s} ||x_1 - x_i||^{-4}$, for Δ -separated Diracs. Since in one dimension only 2 Diracs can be situated at distance $k\Delta$ from x_1 for each integer k > 0, this can be easily bounded by a global bound $\Delta^{-4} \sum_{k=1}^{\infty} k^{-4}$ that does not depend on s. In the multidimensional case however, an *exponential* number of Diracs can be packed around x_1 , and applying the same strategy would lead to a bound on Δ which is exponential in the dimension.

Continuous sampling Fourier transform A variant of the previous framework is *continuous* Fourier sampling, for instance with frequencies distributed according to a Gaussian distribution. Let $\mathcal{X} \subset \mathbb{R}^d$ be a bounded open subset of \mathbb{R}^d . The space of frequencies is $\Omega = \mathbb{R}^d$, $\varphi_{\omega}(x) \stackrel{\text{def.}}{=} e^{i\omega^\top x}$, and $\Lambda(\omega) = \mathcal{N}(0, \Sigma^{-1})$ for some known symmetric positive definite matrix Σ . Note that, for simplicity, the frequencies are drawn according to a Gaussian with *precision matrix* Σ (the inverse of the covariance matrix), such that the kernel K is the classical *Gaussian kernel* $K(x, x') = e^{-\frac{1}{2} \left\| \Sigma^{-\frac{1}{2}}(x-x') \right\|^2}$. The Fisher metric is $\mathfrak{d}_{\mathfrak{g}}(x, x') = \left\| \Sigma^{-\frac{1}{2}}(x-x') \right\|$.

In this case, provided that $\min_{i\neq j} \mathfrak{d}_{\mathfrak{g}}(x_i, x_j) \gtrsim \sqrt{\log(s)}$, stable recovery is guaranteed with

$$m \gtrsim s\left(L\log(s)\log\left(\frac{s}{\rho}\right) + L^2\log\left(\frac{(sL\mathcal{R}_{\mathcal{X}})^d}{\rho}\right)\right).$$

where $L = d + \log^2\left(\frac{dm}{\rho}\right)$. Note that $\log(m)$ appears in L in the r.h.s. of the expression above, which only incurs additional logarithmic terms in the bound on m, as mentioned in the introduction.

Learning of Gaussian mixtures with fixed covariances An original framework for continuous sparsity is sketched learning of mixture models [39], and in particular Gaussian mixture models (GMM), for which we can exploit the computations of the previous case of Fourier measurements sampled in accordance to a Gaussian distribution. Assume that we have data samples $z_1, \ldots, z_n \in \mathbb{R}^d$ drawn *i.i.d.* from a mixture of Gaussians $\xi \stackrel{\text{def.}}{=} \sum_{i=1}^s a_i \mathcal{N}(x_i, \Sigma)$ with known covariance Σ . The means $x_1, \ldots, x_s \in \mathcal{X} \subset \mathbb{R}^d$ and weights $a_1, \ldots, a_s > 0$ are the objects which we want to estimate. Sample frequencies $\omega_1, \ldots, \omega_m \in \mathbb{R}^d$ *i.i.d.* from a Gaussian $\Lambda = \mathcal{N}(0, \Sigma^{-1}/d)$, and construct the following *linear sketch* [24] of the data:

$$y = \frac{C}{n} \sum_{i=1}^{n} (e^{-i\omega_k^\top z_i})_{k=1}^m$$
(5)

where the constant $C = (1 + \frac{2}{d})^{\frac{d}{4}} \leq e^{\frac{1}{2}}$ is here for normalization purpose. Linear sketches are mainly used for computational gain: they are easy to compute in a streaming of distributed context, and are much smaller to store in memory than the whole database [24, 39]. It is easy to see that the sketch can be reformulated as (1), by writing

$$y \approx \mathbb{E}_z (C e^{-\mathrm{i}\omega_k^\top z})_{k=1}^m = \Phi \mu_0 \tag{6}$$

where $\mu_0 = \sum_i a_i \delta_{x_i}$, and Φ is defined using the feature functions

$$\varphi_{\omega}(x) = \mathbb{E}_{z \sim \mathcal{N}(x, \Sigma)} C e^{\mathrm{i}\omega^{\top} z} = C e^{\mathrm{i}\omega^{\top} x} e^{-\frac{1}{2}\omega^{\top} \Sigma \omega}.$$

The "noise" $w \stackrel{\text{def.}}{=} y - \mathbb{E}_z (Ce^{-i\omega_k^\top z})_{k=1}^m$ is simply the difference between empirical and true expectations, and using simple concentration inequalities that we skip here for simplicity, it is possible to show that with high probability, $\|w\| \leq \mathcal{O}\left(n^{-\frac{1}{2}}\right)$. Applying the previous computations we obtain the following result: provided that $\min_{i\neq j} \left\| \Sigma^{-\frac{1}{2}}(x_i - x_j) \right\|_2 \gtrsim \sqrt{d\log(s)}$, stable recovery of μ_0 is guaranteed when

$$m \gtrsim s\left(d\log(s)\log\left(\frac{s}{\rho}\right) + d^2\log\left(\frac{(sd\mathcal{R}_{\mathcal{X}})^d}{\rho}\right)\right)$$

and the concentration in the recovery bound (4) is given by $\delta = ||w|| = \mathcal{O}\left(n^{-\frac{1}{2}}\right)$.

Gaussian mixtures with varying covariances The case of simultaneously recovering both the means and covariance matrices is an interesting venue for future research. We simply describe here the associated metric and distance in the univariate case. The geodesic distance between univariate Gaussian distributions is well known [25]: Given x = (m, u) and x' = (n, v) with $m, n \in \mathbb{R}$ and $u, v \in \mathbb{R}_+$, let $\varphi(x) \stackrel{\text{def.}}{=} \frac{1}{4/\pi} \sqrt{u} e^{-(m-\cdot)/(2u^2)}$, then the covariance kernel is

$$K_0(x,x') = \langle \varphi(x), \, \varphi(x') \rangle_{L^2} = \frac{\sqrt{2uv}}{\sqrt{u^2 + v^2}} e^{-\frac{(m-n)^2}{2(u^2 + v^2)}}.$$

The associated metric at x = (m, u) is $\mathfrak{g}_x = \frac{1}{2u^2} \mathrm{Id}_2$ and the Fisher-Rao distance is the Poincaré half plane distance

$$\mathfrak{d}_0(x,x') = 2\operatorname{arsinh}\left(\frac{\|x-x'\|}{2\sqrt{uv}}\right), \quad \text{where} \quad \operatorname{arsinh}(x) \stackrel{\text{def.}}{=} \ln\left(x+\sqrt{x^2-1}\right). \tag{7}$$

Consider now the case of Gaussian mixture $\xi \stackrel{\text{def.}}{=} \sum_{i=1}^{s} a_i \mathcal{N}(x_i, v_i^2)$, where the unknowns are $a_i > 0$, $x_i \in \mathbb{R}$ and $v_i > 0$, and we are given data $\{z_i\}_{i=1}^n$ drawn iid from ξ and we construct the linear sketch (5) as before, where $\omega_k \in \mathbb{R}$ are iid from $\mathcal{N}(0, \sigma^2)$. This corresponds to the normalised random features

$$\varphi_{\omega}(m,u) = \left(2u^2\sigma^2 + 1\right)^{\frac{1}{4}}e^{-\mathrm{i}m\omega}e^{-\frac{1}{2}u^2\omega^2},$$

and

$$K((m,u),(n,v)) = \frac{\sqrt{2u_{\sigma}v_{\sigma}}}{\sqrt{u_{\sigma}^2 + v_{\sigma}^2}} e^{-\frac{(m-n)^2}{2(u_{\sigma}^2 + v_{\sigma}^2)}}$$
(8)

where $u_{\sigma}^2 = \frac{1}{2\sigma^2} + u^2$ and $v_{\sigma}^2 = \frac{1}{2\sigma^2} + v^2$. The metric at x = (m, u) is $\mathfrak{g}_x = \frac{1}{2u_{\sigma}^2} \operatorname{Id}_2$. Note that since (8) also corresponds the the kernel between Gaussian distributions with mean and standard deviation as $x_{\sigma} \stackrel{\text{def.}}{=} (m, u_{\sigma})$ and $x'_{\sigma} \stackrel{\text{def.}}{=} (n, v_{\sigma})$, the associated geodesic distance is therefore $\mathfrak{d}_0(x_{\sigma}, x'_{\sigma})$ where \mathfrak{d}_0 is the Poincaré half plane distance described in (7) (as mentioned in (16), geodesic distances on random features and parameter space are equivalent).

Sampling the Laplace transform. In some fluorescence microscopy applications (see [30] and the references therein), depth measurements are obtained from the Laplace transform of the signal. Contrary to Fourier measurements, this gives rise to a non-translation invariant kernel K, and was therefore not covered by existing theory. Using the proposed Riemannian geometry framework, we can cover this setting.

Let $\mathcal{X} = (0, 1)^d \subset \mathbb{R}^d_+$. Let $\Omega = \mathbb{R}^d_+$. Define for $x \in \mathcal{X}$ and $\omega \in \Omega$,

$$\varphi_{\omega}(x) \stackrel{\text{\tiny def.}}{=} \exp\left(-x^{\top}\omega\right) \prod_{i=1}^{d} \sqrt{\frac{x_i + \alpha_i}{\alpha_i}} \quad \text{and} \quad \Lambda(\omega) = \exp(-2\alpha^{\top}\omega) \prod_{i=1}^{d} (2\alpha_i).$$

where $\alpha_i \sim d$ are positive and distinct real numbers. The Fisher metric is

$$\mathfrak{d}_{\mathfrak{g}}(x,x') = \sqrt{\sum_{i=1}^{d} \left| \log \left(\frac{x_i + \alpha_i}{x'_i + \alpha_i} \right) \right|^2},$$

and provided that $\min_{i\neq j} \mathfrak{d}_{\mathfrak{g}}(x_i, x_j) \gtrsim d + \log(d^{3/2}s)$, stable recovery is guaranteed with

$$m \gtrsim s\left(C\log(s)\log\left(\frac{s}{\rho}\right) + C^2\log\left(\frac{C^d}{\rho}\right)\right)$$

where $C \stackrel{\text{def.}}{=} d^2 \left(d + \log^2(m) + \log^2\left(\frac{d}{\rho}\right) \right)$. Similar to the Gaussian example, $\log(m)$ appears in C.

Training a two-layers neural network Following [5], in the large number of neurons limit, regression using a neural network with a single hidden layer can be formulated using our framework. Given a set of *m* training samples $(\omega_k, y_k)_{k=1}^m$, one aims to predicts the values $y_k \in \mathbb{R}$ from the features $\omega_k \in \Omega$ using a continuous dictionary of functions $\omega \mapsto \varphi_{\omega}(x)$ (here $x \in \mathcal{X}$ parameterizes the dictionary), as $y_k \approx$ $\int_{\mathcal{X}} \varphi_{\omega_k}(x) d\mu_0(x) = \sum_{i=1}^s a_i \varphi_{\omega_k}(x_i)$. In the context of neural networks, one uses ridge functions of the form $\varphi_{\omega}(x) = \xi(\langle x, \omega \rangle)$, For instance, one can consider $\xi(u) \stackrel{\text{def.}}{=} \xi_n(u) \stackrel{\text{def.}}{=} \max(u, 0)^n$, where n = 1 is the ReLu non-linearity. Detailed treatment of this example will be left for future work, however, we simply mention that the associated kernel and metric was studied in [22, 23]. We recall their results here. Suppose that $\omega_k \sim \mathcal{N}(0, \mathrm{Id}_d)$, then the associated kernel is

$$k_n(x,y) = 2 \int \frac{e^{-\|\omega\|^2/2}}{(2\pi)^{d/2}} \xi_n(\langle\omega, x\rangle) \xi_n(\langle\omega, y\rangle) d\omega = \frac{1}{\pi} \|x\|^n \|y\|^n J_n(\theta),$$

where
$$\theta = \cos^{-1}\left(\frac{x^{\top}y}{\|x\| \|y\|}\right)$$
 and $J_n(\theta) = (-1)^n (\sin\theta)^{2n+1} \left(\frac{1}{\sin\theta} \frac{\partial}{\partial\theta}\right)^n \left(\frac{\pi-\theta}{\sin\theta}\right)$

Since $J_n(0) = \pi (2n-1)!!$, the normalised random features the normalised kernel are therefore

$$\varphi_{\omega}(x) = \frac{\pi}{\sqrt{J_n(0)}} \xi_n\left(\frac{\langle \omega, x \rangle}{\|x\|}\right) \text{ and } K_n(x,y) = \frac{J_n(\theta)}{J_n(0)},$$

and the associated metric is

$$\mathfrak{g}_{x} = \nabla_{1} \nabla_{2} K_{n}(x, x) = \frac{n^{2}}{\|x\|^{2} (2n-1)} \left(\mathrm{Id}_{d} + (4n-3) \frac{xx^{\top}}{\|x\|^{2}} \right).$$

Note that since $\varphi_{\omega}(x) = \varphi_{\omega}(x/||x||)$ for all x, the geodesic path between any x and y must lie on the unit sphere, moreover, given any $\gamma : [0,1] \to \mathcal{X}$ with $||\gamma(t)|| = 1$ for all $t \in [0,1]$, we have

$$\int_{0}^{1} \sqrt{\langle \mathfrak{g}_{\gamma(t)}\gamma'(t), \gamma'(t)\rangle} dt = \int_{0}^{1} \sqrt{\frac{n^{2}}{(2n-1)} \left(\|\gamma'(t)\|^{2} + (4n-3) \left|\langle \gamma'(t), \gamma(t) \rangle\right|^{2} \right)} dt = \sqrt{\frac{n^{2}}{2n-1}} \int_{0}^{1} \|\gamma'(t)\| dt,$$

therefore, $\mathfrak{d}_{\mathfrak{g}}(x,y) = \sqrt{\frac{n^2}{2n-1}} d_{\mathbb{S}}\left(\frac{x}{\|x\|}, \frac{y}{\|y\|}\right)$, where $d_{\mathbb{S}}$ is the geodesic distance on the sphere.

3 Stability and the Fisher information metric

In this section, we introduce the proposed Riemmanian geometry framework, and give intermediate recovery guarantees which constitute the first building block of our main result. Namely, we introduce so-called *dual certificates*, which are known to be key objects in the study of the BLASSO, and show how they lead to sparse recovery guarantees in our Riemannian framework.

3.1 Fisher and Optimal Transport Distances

Let us first introduce the proposed Riemannian geometry framework, and define objects related to it.

3.1.1 The covariance kernel and the Fubini-Study metric

A natural property to analyse in our problem is the way two Diracs *interact* with each other, which is linked to the well-known notion of *coherence* (or, rather, incoherence) between measurements in compressive sensing [38]. This is done through what we refer to as the *covariance kernel* $\hat{K} : \mathcal{X} \times \mathcal{X} \to \mathbb{C}$, defined as

$$\hat{K}(x,x') \stackrel{\text{\tiny def.}}{=} \langle \Phi \delta_x, \, \Phi \delta_{x'} \rangle_2 = \frac{1}{m} \sum_{j=1}^m \overline{\varphi_{\omega_k}(x)} \varphi_{\omega_k}(x'), \qquad \forall x, x' \in \mathcal{X} \,.$$
(9)

In the limit case $m \to \infty$, the law of large number states that \hat{K} converges almost surely to the *limit* covariance kernel:

$$K(x, x') \stackrel{\text{def.}}{=} \mathbb{E}_{\omega} \overline{\varphi_{\omega}(x)} \varphi_{\omega}(x') \tag{10}$$

where we recall that $\omega \sim \Lambda$. This object naturally governs the geometry of the space, and we use it to define our Riemmanian metric, which as we will see is linked to a notion of Fisher information metric. In the rest of the paper, we assume throughout that K is real-valued, even though \hat{K} may be complex-valued.

Given the normalisation $\mathbb{E}_{\omega} |\varphi_{\omega}(x)|^2 = 1$ for all $x \in \mathcal{X}$, $\varphi_{\omega}(x)$ can be interpreted as a complex-valued probability amplitude with respect to ω (parameterized by x), a classical notion in quantum mechanics (see [40]). When x varies, a natural metric between probability amplitudes is the so-called Fubini-Study metric, which is the complex equivalent of the well-known Fisher information metric. Writing $\varphi_{\omega}(x) =$

 $\sqrt{p(\omega, x)}e^{i\alpha(\omega, x)}$ where $p(\omega, x) \stackrel{\text{def.}}{=} |\varphi_{\omega}(x)|^2$ and $\alpha(\omega, x) \stackrel{\text{def.}}{=} \arg(\varphi_{\omega}(x))$, the Fubini-Study metric is defined by the following metric tensor in $\mathbb{C}^{d \times d}$ [36]:

$$\mathfrak{g}_{x} \stackrel{\text{def}}{=} \frac{1}{4} \mathbb{E}_{p} [\nabla_{x} \log(p) \nabla_{x} \log(p)^{\top}] + \mathbb{E}_{p} [\nabla_{x} \alpha \nabla_{x} \alpha^{\top}] - \mathbb{E}_{p} [\nabla_{x} \alpha] \mathbb{E}_{p} [\nabla_{x} \alpha]^{\top} - \frac{\mathrm{i}}{2} \mathbb{E}_{p} [\nabla_{x} \log(p) \nabla_{x} \alpha - \nabla_{x} \alpha \nabla_{x} \log(p)^{\top}].$$
(11)

where we use the notation $\mathbb{E}_p[f] = \int f(\omega)p(\omega, x)d\Lambda(\omega)$. If φ_{ω} is real-valued, then $\alpha = 0$ and this is indeed the Fisher metric up to a factor of $\frac{1}{4}$. The following simple Lemma shows the link between this metric and the derivatives of the covariance kernel K.

Lemma 1. For any kernel $K(x, x') \stackrel{\text{def.}}{=} \mathbb{E}_{\omega} \overline{\varphi_{\omega}(x)} \varphi_{\omega}(x')$, the Fubini-Study metric defined in (11) satisfies

$$\mathfrak{g}_x = \nabla_1 \nabla_2 K(x, x) - \mathbb{E}_p [\nabla_x \alpha] \mathbb{E}_p [\nabla_x \alpha]^\top$$
(12)

If furthermore K(x, x') is assumed real-valued, then $\mathbb{E}_p[\nabla_x \alpha] = 0$, and $\mathfrak{g}_x = \nabla_1 \nabla_2 K(x, x)$.

Proof. Using $p = |\varphi_{\omega}|^2$ and $\nabla \varphi_{\omega} = \left(\frac{\nabla p}{2p} + i\nabla \alpha\right)\varphi_{\omega}$, a direct computation shows that

$$\nabla_x \log(p) = \frac{2}{p} \operatorname{Re}\left(\overline{\varphi_\omega} \nabla \varphi_\omega\right) \quad \text{and} \quad \nabla_x \alpha = \frac{1}{p} \operatorname{Im}\left(\overline{\varphi_\omega} \nabla \varphi_\omega\right) \tag{13}$$

Therefore,

$$\frac{1}{4} \mathbb{E}_{p} [\nabla_{x} \log(p) \nabla_{x} \log(p)^{\top}] + \mathbb{E}_{p} [\nabla_{x} \alpha \nabla_{x} \alpha^{\top}] \\
= \int \frac{1}{p^{2}} \left(\operatorname{Re} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right) \operatorname{Re} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right)^{\top} + \operatorname{Im} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right) \operatorname{Im} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right)^{\top} \right) p d\Lambda \\
= \int \frac{1}{p} \operatorname{Re} \left(|\varphi_{\omega}|^{2} \overline{\nabla \varphi_{\omega}} \nabla \varphi_{\omega}^{\top} \right) d\Lambda = \int \operatorname{Re} \left(\overline{\nabla \varphi_{\omega}} \nabla \varphi_{\omega}^{\top} \right) d\Lambda = \operatorname{Re} \left(\nabla_{1} \nabla_{2} K(x, x) \right)$$

Similarly,

$$\begin{aligned} -\frac{\mathrm{i}}{2} \mathbb{E}_{p} [\nabla_{x} \log(p) \nabla_{x} \alpha - \nabla_{x} \alpha \nabla_{x} \log(p)^{\top}] \\ &= -\mathrm{i} \int \frac{1}{p^{2}} \left(\operatorname{Re} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right) \operatorname{Im} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right)^{\top} + \operatorname{Im} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right) \operatorname{Re} \left(\overline{\varphi_{\omega}} \nabla \varphi_{\omega} \right)^{\top} \right) p \mathrm{d}\Lambda \\ &= -\mathrm{i} \int \frac{1}{p} \operatorname{Im} \left(|\varphi_{\omega}|^{2} \overline{\nabla \varphi_{\omega}} \nabla \varphi_{\omega}^{\top} \right) \mathrm{d}\Lambda = \mathrm{i} \int \operatorname{Im} \left(\overline{\nabla \varphi_{\omega}} \nabla \varphi_{\omega}^{\top} \right) \mathrm{d}\Lambda = \mathrm{i} \cdot \operatorname{Im} \left(\nabla_{1} \nabla_{2} K(x, x) \right) \end{aligned}$$

which proves the first claim. The second claim is immediate by noticing from (13) that $\nabla_p \alpha = \text{Im}(\nabla_2 K(x, x))$, which cancels when K(x, x') is real (in particular in a neighborhood around x = x').

Since in this paper the limit covariance kernel (10) is assumed real-valued, the previous Lemma justifies the definition $\mathfrak{g}_x = \nabla_1 \nabla_2 K(x, x)$ that we adopt in the rest of the paper. For two vectors $u, v \in \mathbb{C}^d$, we define the corresponding inner product

$$\langle u, v \rangle_x \stackrel{\text{def.}}{=} u^* \mathfrak{g}_x v \quad \text{and} \quad \|u\|_x \stackrel{\text{def.}}{=} \sqrt{\langle u, u \rangle_x}$$
(14)

As described in the introduction, this induces a geodesic distance on \mathcal{X} :

$$\mathfrak{d}_{\mathfrak{g}}(x,x') \stackrel{\text{\tiny def.}}{=} \inf\left\{\int_{0}^{1} \|\gamma'(t)\|_{\gamma(t)} \,\mathrm{d}t \; ; \; \gamma: [0,1] \to \mathcal{X} \text{ smooth, } \gamma(0) = x, \; \gamma(1) = x'\right\}$$
(15)

and in the case where $\varphi_{\omega}(x)$ is real-valued, this coincides with the "Fisher-Rao" geodesic distance [47] which is used extensively in information geometry for estimation and learning problems on parametric families of distributions [3]. Remark 1 (As a distance on the feature space). The geodesic distance induced by \mathfrak{g} is the natural distance between the random features $\varphi_{\cdot}(x)$. Indeed, as discussed in [11], the manifold $(\mathcal{X},\mathfrak{g})$ as an embedded submanifold of the sphere in Hilbert space $L_2(d\Lambda)$ with embedding $x \mapsto \varphi_{\cdot}(x)$, and given any $x, x' \in \mathcal{X}$, we have

$$\inf_{\gamma \in \Gamma_{x,x'}} \int_0^1 \|\gamma'(t)\|_{L_2(\mathrm{d}\Lambda)} \,\mathrm{d}t = \mathfrak{d}_\mathfrak{g}(x,x'),\tag{16}$$

where $\Gamma_{x,x'}$ consists of all piecewise smooth paths $\gamma : [0,1] \to \{\varphi(x) ; x \in \mathcal{X}\}$ with $\gamma(0) = \varphi(x)$ and $\gamma(1) = \varphi(x')$.

Remark 2 (Fisher metric and invariances). The Fisher-Rao metric $\mathfrak{d}_{\mathfrak{g}}$ is "canonical" in the sense that it is the only (up to scalar multiples) geodesic distance which satisfies the natural invariances of the BLASSO problem. Indeed, the solutions to $(\mathcal{P}_{\lambda}(y))$, in the large sample limit $m \to +\infty$, are (i) invariant by the multiplication of $\varphi(x) \stackrel{\text{def.}}{=} (\varphi_{\omega}(x))_{\omega \in \Omega}$ by an arbitrary orthogonal transform U (orthogonality on $L_2(d\Lambda)$), i.e. invariance to $\varphi(x) \mapsto U\varphi(x)$, (ii) covariance under any change of variable $\varphi \mapsto \varphi \circ h$ where h is a diffeomorphism between two d-dimensional parameter spaces. The covariance (ii) means that if $\mu = \sum_i a_i \delta_{x_i}$ is a solution associated to φ , then the push-forward measure $(h^{-1})_{\sharp} \mu \stackrel{\text{def.}}{=} \sum_i a_i \delta_{h^{-1}(x_i)}$ is a solution associated to $\varphi \circ h$. Note that the invariance (i) is different from the usual invariance under "Markov morphisms" considered in information theory [18, 13]. When considering $\mathfrak{d}_{\mathfrak{g}} = \mathfrak{d}_{\mathfrak{g}_{\varphi}}$ as a Riemannian distance depending solely on φ , the invariance under any diffeomorphism h reads

$$\mathfrak{d}_{\mathfrak{g}_{\varphi}}(x,x') = \mathfrak{d}_{\mathfrak{g}_{\varphi \circ h}}(h^{-1}(x),h^{-1}(x')). \tag{17}$$

Assuming for simplicity that φ is injective, this invariance (17) is equivalent to the fact that the formula

$$\forall (q,q') \in \mathcal{M}^2, \quad d_{\mathcal{M}}(q,q') \stackrel{\text{\tiny def.}}{=} \mathfrak{d}_{\mathfrak{g}_{\varphi}}(\varphi^{-1}(q),\varphi^{-1}(q'))$$

defines a proper (i.e. parameterization-independent) Riemannian distance $d_{\mathcal{M}}$ on the embedded manifold $\mathcal{M} \stackrel{\text{def.}}{=} (\varphi(x))_x \subset L_2(\mathrm{d}\Lambda)$. Among all possible such Riemannian metrics on \mathcal{M} , the only ones being invariant by orthogonal transforms $\varphi \mapsto U\varphi$ are scalar multiples of the hermitian positive tensor $\partial \varphi(x)^* \partial \varphi(x) \in \mathbb{C}^{d \times d}$, which is equal to \mathfrak{g}_{φ} (here $\partial \varphi(x)^*$ refers to the adjoint in $L_2(\mathrm{d}\Lambda)$ for the inner product defined by the measure $\Lambda(\omega)$).

Remark 3 (Tangent spaces). Formally, in Riemannian geometry, one would use the notion of tangent space \mathcal{T}_x , and for instance the inner product $\langle \cdot, \cdot \rangle_x$ would only be defined between vectors belonging to \mathcal{T}_x . However, in our case, since the considered ambient "manifold" is just \mathbb{R}^d , in the sense that \mathcal{X} is not a low-dimensional sub-manifold of \mathbb{R}^d but an open set of \mathbb{R}^d , each tangent space can be identified with \mathbb{R}^d , and we extend the definitions to complex vectors for our needs.

3.1.2 Optimal Transport metric

In order to state quantitative performance bounds, one needs to consider a geometric distance between measures. The canonical way to "lift" a ground distance $\mathfrak{d}_{\mathfrak{g}}(x, x')$ between parameter to a distance between measure is to use optimal transport distances [49].

Definition 1 (Wasserstein distance). Given $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ with $|\mu|(\mathcal{X}) = |\nu|(\mathcal{X})$, the Wasserstein distance between μ and ν , relative to the metric \mathfrak{d} on \mathcal{X} is defined by

$$W^2_{\mathfrak{d}}(\mu,\nu) \stackrel{\text{def.}}{=} \inf_{\gamma \in \Pi(\mu,\nu)} \int_{\mathcal{X}^2} \mathfrak{d}(x,x') \mathrm{d}\gamma(x,x'),$$

where $\Gamma(\mu,\nu) \subset \mathcal{M}_+(\mathcal{X}^2)$ is the set of all transport plans with marginals μ and ν . Given $\mu,\nu \in \mathcal{M}_+(\mathcal{X})$ (not necessarily of equal total mass), the optimal partial distance between μ and ν is defined as

$$\mathcal{T}_{\mathfrak{d}}^{2}(\mu,\nu) \stackrel{\text{\tiny def.}}{=} \inf_{\tilde{\mu},\tilde{\nu}} \{ W_{\mathfrak{d}}^{2}(\tilde{\mu},\tilde{\nu}) + |\mu - \tilde{\mu}| \left(\mathcal{X}\right) + |\tilde{\nu} - \nu| \left(\mathcal{X}\right) \}$$

Note that the distance $W_{\mathfrak{d}}(\mu,\nu)$ is actually an hybridation (an inf-convolution) between the classical Wasserstein distance between probability distributions and the total variation norm. It is often called "partial optimal transport" in the literature (see for instance [12]), and belongs to the larger class of unbalanced optimal transport distances [43, 21].

3.2 Nondegenerate certificates, uniqueness and stability for sparse measures

We now introduce the notion of a dual certificate and prove recovery guarantees under certain nondegeneracy conditions, which is the first step toward our main result.

3.2.1 Dual certificates

The minimisation problem $(\mathcal{P}_{\lambda}(y))$ is a convex optimisation problem and a natural way of studying their solutions are via their corresponding Fenchel-dual problems. It is well known that, in the limit as $\lambda \to 0$, its solutions cluster in a weak-* sense around minimisers of

$$\min_{\mu \in \mathcal{M}(\mathcal{X})} |\mu| (\mathcal{X}) \text{ subject to } \Phi \mu = y , \qquad (\mathcal{P}_0(y))$$

and that properties of the dual solutions to $(\mathcal{P}_0(y))$ with $y = \Phi \mu_0$ can be used to derive stability estimates for $(\mathcal{P}_{\lambda}(y))$ under noisy measurements. In this section, we recall some of these results (see [9, 32] for further details). The (pre)dual of $(\mathcal{P}_{\lambda}(y))$ is

$$\sup\left\{\langle p, y\rangle_2 - \frac{\lambda}{2} \|p\|_2^2 ; p \in \mathbb{C}^m, \|\Phi^*p\|_\infty \leqslant 1\right\}$$
 $(\mathcal{D}_\lambda(y))$

where we remark that the adjoint operator $\Phi^* : \mathbb{C}^m \to \mathscr{C}(\mathcal{X})$ is defined by $(\Phi^* p)(x) = \frac{1}{\sqrt{m}} \sum_{i=1}^m p_i \varphi_{\omega_i}(x)$. Note that for $\lambda > 0$, this is the projection of y/λ onto the closed convex set $\{p : \|\Phi^*p\|_{\infty} \leq 1\}$ and the solution p_{λ} is hence unique. The dual solution p_{λ} is related to any primal solution μ_{λ} of $(\mathcal{P}_{\lambda}(y))$ by the condition

$$\Phi^* p_{\lambda} \in \partial |\mu_{\lambda}| (\mathcal{X}) \text{ and } p_{\lambda} = \frac{1}{\lambda} (y - \Phi \mu_{\lambda}).$$
 (18)

Conversely, any pair p_{λ} and μ_{λ} which satisfy this equation (18) are necessarily dual and primal solutions of $(\mathcal{D}_{\lambda}(y))$ and $(\mathcal{P}_{\lambda}(y))$ respectively. In the case where $\lambda = 0$, a dual solution need not be unique, although existence is guaranteed (since in our setting, the dual variable belongs to a finite dimensional space). In this case, p_0 and μ_0 solve $(\mathcal{D}_{\lambda}(y))$ with $\lambda = 0$ and $(\mathcal{P}_0(y))$, respectively, if and only if

$$\Phi \mu_0 = y \quad \text{and} \quad \Phi^* p_0 \in \partial |\mu_0| (\mathcal{X}). \tag{19}$$

Following the literature, we call any element $\eta \in \operatorname{Im}(\Phi^*) \cap \partial |\mu_0|(\mathcal{X})$ a dual certificate for μ_0 . For $\mu_0 = \sum_{j=1}^s a_j \delta_{x_j}$, the condition $\eta \in \partial |\mu_0|(\mathcal{X})$ imposes that $\eta(x_j) = \operatorname{sign}(a_j)$ and $\|\eta\|_{\infty} \leq 1$. Furthermore, it is known that in the noiseless case, μ_0 is the unique solution to $(\mathcal{P}_0(y))$ if: the operator $\Phi_x : \mathbb{C}^s \to \mathbb{C}^m$ defined by $\Phi_x b = \sum_{j=1}^s b_j \Phi \delta_{x_j}$ is injective, and there exists $\eta \in \operatorname{Im}(\Phi^*) \cap \partial |\mu_0|(\mathcal{X})$ such that $|\eta(x)| < 1$ for all $x \notin \{x_j\}$. In order to quantify the latter constraint and provide quantitative stability bounds, we impose even stronger conditions on η and make the following definition.

Definition 2 (Non-degenerate dual certificate). Given $(a_i, x_i)_{i=1}^s$, we say that $\eta \in \text{Im}(\Phi^*)$ is an $(\varepsilon_0, \varepsilon_2, r)$ nondegenerate dual certificate if:

- (*i*) $\eta(x_i) = \text{sign}(a_i)$ for all i = 1, ..., s,
- (*ii*) $|\eta(x)| \leq 1 \varepsilon_0$ for all $x \in \mathcal{X}^{\text{far}}$,
- (iii) $|\eta(x)| \leq 1 \varepsilon_2 \mathfrak{d}_{\mathfrak{g}}(x, x_i)^2$ for all $x \in \mathcal{X}_i^{\text{near}}$,

where $\mathcal{X}_{i}^{\text{near}} \stackrel{\text{def.}}{=} \mathcal{B}_{\mathfrak{d}_{\mathfrak{q}}}(x_{i}; r)$ and $\mathcal{X}^{\text{far}} \stackrel{\text{def.}}{=} \mathcal{X} \setminus \bigcup_{i} \mathcal{X}_{i}^{\text{near}}$.

In other words, there are neighborhoods of the x_j such that, outside of these neighborhoods, η is strictly bounded away from 1, and inside, $|\eta|$ has quadratic decay. In the next section we prove stable recovery results from the existence of non-degenerate dual certificates.

3.2.2 Stable recovery bounds

The following two propositions describe stability guarantees under the nondegeneracy condition. Proposition 1 quantifies how the recovered measure is approximated by a sparse measure supported on $\{x_j\}_j$, and Proposition 2 describes the error in measure around small neighbourhoods of the points $\{x_j\}_j$.

Proposition 1 (Stability away from the sparse support). Suppose that there exists $\varepsilon_0, \varepsilon_2 > 0, \eta \stackrel{\text{def.}}{=} \Phi^* p$ for some $p \in \mathbb{C}^m$ such that η is $(\varepsilon_0, \varepsilon_2, r)$ -nondegenerate. Assuming the measurement model (1), any minimiser $\hat{\mu}$ of $(\mathcal{P}_{\lambda}(y))$, with $\|w\| \leq \delta$ and $\lambda \sim \delta/\|p\|$ is approximately sparse: by defining $\hat{A}_j = |\hat{\mu}|(\mathcal{X}_j^{\text{near}})$, we have

$$\mathcal{T}^{2}_{\mathfrak{d}_{\mathfrak{g}}}\left(\left|\hat{\mu}\right|, \sum_{j=1}^{s} \hat{A}_{j} \delta_{x_{j}}\right) \lesssim \frac{1}{\min\left(\varepsilon_{0}, \varepsilon_{2}\right)} \left(\left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) + \delta \left\|p\right\|\right).$$

$$(20)$$

Proof. To prove this proposition, we first establish the following bound

$$\varepsilon_{0} \left| \hat{\mu} \right| \left(\mathcal{X}^{\text{far}} \right) + \varepsilon_{2} \sum_{i=1}^{s} \int_{\mathcal{X}_{i}^{\text{near}}} \mathfrak{d}_{\mathfrak{g}}(x, x_{i})^{2} \mathrm{d} \left| \hat{\mu} \right| (x) \lesssim \delta \left\| p \right\| + \left| \tilde{\mu}_{0} \right| (\mathcal{X}).$$

$$(21)$$

As we will see, the optimal partial transport bound above is then a consequence of this bound.

For i = 1, ..., s, let $\mathcal{X}_i^{\text{near}} \subset \mathcal{X}$ and $\mathcal{X}^{\text{far}} = \mathcal{X} \setminus \bigcup_{j=1}^s \mathcal{X}_j^{\text{near}}$ be as in Definition 2. Recall the measurement model $y = \Phi(\mu_0 + \tilde{\mu}_0) + w$, and define $\bar{\mu}_0 = \mu_0 + \tilde{\mu}_0$ for simplicity. We first adapt the proof of [10, Thm. 2] to derive an upper bound for $|\hat{\mu}| - |\bar{\mu}_0| - \text{Re}(\langle \eta, \hat{\mu} - \bar{\mu}_0 \rangle_{\mathcal{M}})$. By minimality of $\hat{\mu}$ and since $||w|| \leq \delta$,

$$\lambda \left| \hat{\mu} \right| \left(\mathcal{X} \right) + \frac{1}{2} \left\| \Phi \hat{\mu} - y \right\|^2 \leqslant \lambda \left| \bar{\mu}_0 \right| \left(\mathcal{X} \right) + \frac{1}{2} \left\| \Phi \bar{\mu}_0 - y \right\|^2 \leqslant \lambda \left| \bar{\mu}_0 \right| \left(\mathcal{X} \right) + \frac{\delta^2}{2}$$

Using $\eta = \Phi^* p$, and by adding and subtracting $\operatorname{Re}(\langle \eta, \hat{\mu} - \bar{\mu}_0 \rangle_{\mathcal{M}}) = \operatorname{Re}(\langle p, \Phi \hat{\mu} - y \rangle_2) + \operatorname{Re}(\langle p, w \rangle_2)$, we obtain

$$\lambda\left(\left|\hat{\mu}\right|\left(\mathcal{X}\right)-\left|\bar{\mu}_{0}\right|\left(\mathcal{X}\right)-\operatorname{Re}\left(\langle\eta,\,\hat{\mu}-\bar{\mu}_{0}\rangle_{\mathcal{M}}\right)\right)+\operatorname{Re}\left(\langle\lambda p,\,\Phi(\hat{\mu}-\bar{\mu}_{0})\rangle_{2}\right)+\frac{1}{2}\left\|\Phi\hat{\mu}-y\right\|^{2}\leqslant\frac{\delta^{2}}{2}$$

$$\implies\lambda\left(\left|\hat{\mu}\right|\left(\mathcal{X}\right)-\left|\bar{\mu}_{0}\right|\left(\mathcal{X}\right)-\operatorname{Re}\left(\langle\eta,\,\hat{\mu}-\bar{\mu}_{0}\rangle_{\mathcal{M}}\right)\right)+\frac{1}{2}\left\|\Phi\hat{\mu}-y+\lambda p\right\|^{2}\leqslant\frac{\delta^{2}}{2}+\frac{\lambda^{2}\left\|p\right\|^{2}}{2}-\operatorname{Re}\left(\langle\lambda p,\,w\rangle_{2}\right) \quad (22)$$

$$\implies\left|\hat{\mu}\right|\left(\mathcal{X}\right)-\left|\bar{\mu}_{0}\right|\left(\mathcal{X}\right)-\operatorname{Re}\left(\langle\eta,\,\hat{\mu}-\bar{\mu}_{0}\rangle_{\mathcal{M}}\right)\leqslant\frac{1}{2\lambda}\left(\delta+\lambda\left\|p\right\|\right)^{2}\lesssim\delta\left\|p\right\|$$

using $\lambda \sim \delta / \|p\|$. We now derive a lower bound for $|\hat{\mu}| - |\mu_0| - \operatorname{Re}(\langle \eta, \hat{\mu} - \bar{\mu}_0 \rangle_{\mathcal{M}})$. Since η is a dual certificate, we have $\langle \eta, \bar{\mu}_0 \rangle_{\mathcal{M}} = |\mu_0|(\mathcal{X})$ and $|\langle \eta, \mu \rangle_{\mathcal{M}}| \leq |\mu|(\mathcal{X})$. By further exploiting the nondegeneracy assumptions (ii) and (iii) on η , we have

$$\begin{split} |\hat{\mu}|\left(\mathcal{X}\right) - |\bar{\mu}_{0}|\left(\mathcal{X}\right) - \operatorname{Re}\left(\langle\eta, \,\hat{\mu} - \bar{\mu}_{0}\rangle\right) \geqslant |\hat{\mu}|\left(\mathcal{X}\right) - \operatorname{Re}\left(\langle\eta, \,\hat{\mu}\rangle\right) - 2\left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) \\ \geqslant |\hat{\mu}|\left(\mathcal{X}\right) - \sum_{i} \int_{\mathcal{X}_{i}^{\operatorname{near}}} |\eta| \,\mathrm{d}\left|\hat{\mu}\right| - \int_{\mathcal{X}^{\operatorname{far}}} |\eta| \,\mathrm{d}\left|\hat{\mu}\right| - 2\left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) \\ \geqslant |\hat{\mu}|\left(\mathcal{X}\right) - \sum_{i} \int_{\mathcal{X}_{i}^{\operatorname{near}}} \left(1 - \varepsilon_{2} \mathfrak{d}_{\mathfrak{g}}(x, x_{i})^{2}\right) \,\mathrm{d}\left|\hat{\mu}\right|\left(x\right) - (1 - \varepsilon_{0})\left|\hat{\mu}\right|\left(\mathcal{X}^{\operatorname{far}}\right) - 2\left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) \\ = \varepsilon_{0}\left|\hat{\mu}\right|\left(\mathcal{X}^{\operatorname{far}}\right) + \varepsilon_{2} \sum_{i} \int_{\mathcal{X}_{i}^{\operatorname{near}}} \mathfrak{d}_{\mathfrak{g}}(x, x_{i})^{2} \,\mathrm{d}\left|\hat{\mu}\right|\left(x\right) - 2\left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) \end{split}$$

which proves (21). Note also that by combining this with (22), we obtain the following bound that we will use later:

$$\left\|\Phi\hat{\mu} - y + \lambda p\right\|^{2} \leq \left(\delta + \lambda \left\|p\right\|\right)^{2} + 4\lambda \left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) \implies \left\|\Phi\hat{\mu} - y\right\| \leq \delta + 2\lambda \left\|p\right\| + 2\sqrt{\lambda \left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right)}$$
(23)

It remains to show that the bound (21) yields an upper bound on the partial optimal transport distance between the recovered measure $|\hat{\mu}|$ and $\rho \stackrel{\text{def.}}{=} \sum_{i} |\hat{\mu}| (\mathcal{X}_{i}^{\text{near}}) \delta_{x_{i}}$, its "projection" onto the positions $\{x_{j}\}_{j}$. To see this, first note that the Kantorovich dual formulation [49] of the Wasserstein distance in Def. 1 is

$$\sup\left\{\int_{\mathcal{X}}\varphi d\mu + \int_{\mathcal{X}}\psi d\nu \; ; \; \varphi, \psi \in C_b(\mathcal{X}), \; \forall x, y \in \mathcal{X}, \; \varphi(x) + \psi(y) \leqslant \mathfrak{d}_{\mathfrak{g}}(x, y)^2\right\}$$

Given any $\varphi, \psi \in C_b(\mathcal{X})$ satisfying $\varphi(x) + \psi(y) \leq \mathfrak{d}_{\mathfrak{g}}(x, y)^2$ for all $x, y \in \mathcal{X}$, we have

$$\begin{split} W_{\mathfrak{g}}^{2}(\rho,|\hat{\mu}|_{\mathcal{X}^{\text{near}}}) &\leqslant \int \varphi \mathbf{d} \, |\hat{\mu}|_{\mathcal{X}^{\text{near}}} + \int \psi \mathbf{d}\rho \\ &= \sum_{j} \left(\int_{\mathcal{X}_{j}^{\text{near}}} (\varphi(x) + \psi(x_{j})) \mathbf{d} \, |\hat{\mu}| \, (x) - \psi(x_{j}) \int_{\mathcal{X}_{j}^{\text{near}}} \mathbf{d} \, |\hat{\mu}| \, (x) + \psi(x_{j}) \, |\hat{\mu}| \, (\mathcal{X}_{j}^{\text{near}}) \right) \\ &= \sum_{j} \int_{\mathcal{X}_{j}^{\text{near}}} (\varphi(x) + \psi(x_{j})) \mathbf{d} \, |\hat{\mu}| \, (x) \leqslant \sum_{j} \int_{\mathcal{X}_{j}^{\text{near}}} \mathfrak{d}(x, x_{j})^{2} \mathbf{d} \, |\hat{\mu}| \, (x) \end{split}$$

So,

$$\varepsilon_2 W_{\mathfrak{g}}^2(
ho, |\hat{\mu}|_{\mathcal{X}^{near}}) \lesssim |\tilde{\mu}_0| (\mathcal{X}) + \delta \|p\|$$

So, since $\varepsilon_0 |\hat{\mu}|_{\mathcal{X}^{\mathrm{far}}}(\mathcal{X}) \lesssim |\tilde{\mu}_0|(\mathcal{X}) + \delta ||p||$, we have

$$\mathcal{T}_{\mathfrak{g}}^{2}(\left|\hat{\mu}\right|,\rho) \lesssim \frac{1}{\min\left(\varepsilon_{0},\varepsilon_{2}\right)}\left(\left|\tilde{\mu}_{0}\right|\left(\mathcal{X}\right) + \delta\left\|p\right\|\right).$$

We now give stability bounds around the sparse support, under some additional assumptions.

Proposition 2 (Stability around the sparse support). Under the assumptions of Proposition 1, let $\hat{\mu}$ be a solution of $(\mathcal{P}_{\lambda}(y))$, and let $\hat{a} = (\hat{\mu}(\mathcal{X}_{j}^{\text{near}}))_{j=1}^{s}$. Suppose in addition that for $j = 1, \ldots, s$, there exists $\eta_{j} = \Phi^{*}p_{j}$ which satisfies

(i)
$$\eta_j(x_j) = 1$$
 and $\eta_j(x_\ell) = 0$ for all $\ell \neq j$

(*ii*)
$$|1 - \eta_j(x)| \leq \varepsilon_2 \mathfrak{d}_{\mathfrak{g}}(x, x_j)^2$$
 for all $x \in \mathcal{X}_j^{\text{near}}$

(iii) $|\eta_j(x)| \leq \varepsilon_2 \mathfrak{d}_{\mathfrak{g}}(x, x_\ell)^2$ for all $x \in \mathcal{X}_\ell^{\text{near}}$ and $\ell \neq j$,

(iv)
$$|\eta_j(x)| \leq 1 - \varepsilon_0$$
 for all $x \in \mathcal{X}^{\text{far}}$.

Then

$$\forall j = 1, \dots, s, \quad |\hat{a}_j - a_j| \lesssim \|p_j\| \left(\delta + \lambda \|p_j\|\right) + \varepsilon_0^{-1} \left(\delta \|p\| + |\tilde{\mu}_0| \left(\mathcal{X}\right)\right) \tag{24}$$

where p is as in Proposition 1.

Proof. First observe that writing $\nu = \hat{\mu} - \mu_0$, we have

$$\begin{split} |\hat{a}_{j} - a_{j}| &= \left| \int_{\mathcal{X}_{j}^{\text{near}}} \mathrm{d}\nu(x) \right| = \left| \int_{\mathcal{X}} \eta_{j}(x) \mathrm{d}\nu(x) + \int_{\mathcal{X}_{j}^{\text{near}}} (1 - \eta_{j}(x)) \mathrm{d}\nu(x) - \sum_{\ell \neq j} \int_{\mathcal{X}_{\ell}^{\text{near}}} \eta_{j}(x) \mathrm{d}\nu(x) - \int_{\mathcal{X}^{\text{far}}} \eta_{j}(x) \mathrm{d}\nu(x) \right| \\ &\leq \left| \int_{\mathcal{X}} \eta_{j}(x) \mathrm{d}\nu(x) \right| + \varepsilon_{2} \sum_{j=1}^{s} \left| \int_{\mathcal{X}_{j}^{\text{near}}} \mathfrak{d}_{\mathfrak{g}}(x, x_{j})^{2} \mathrm{d}\nu(x) \right| + (1 - \varepsilon_{0}) \left| \nu \right| (\mathcal{X}^{\text{far}}). \end{split}$$

Using (21), we have $|\nu|(\mathcal{X}^{\mathrm{far}}) = |\hat{\mu}|(\mathcal{X}^{\mathrm{far}}) \lesssim \varepsilon_0^{-1} \left(\delta \|p\| + |\tilde{\mu}_0|(\mathcal{X})\right)$ and

$$\varepsilon_{2} \sum_{j=1}^{s} \left| \int_{\mathcal{X}_{j}^{\text{near}}} \mathfrak{d}_{\mathfrak{g}}(x, x_{j})^{2} \mathrm{d}\nu(x) \right| = \varepsilon_{2} \sum_{j=1}^{s} \left| \int_{\mathcal{X}_{j}^{\text{near}}} \mathfrak{d}_{\mathfrak{g}}(x, x_{j})^{2} \mathrm{d}\hat{\mu}(x) \right| \leq \delta \left\| p \right\| + \left| \tilde{\mu}_{0} \right| (\mathcal{X})$$

Finally, by (23),

$$\begin{aligned} \left| \int_{\mathcal{X}} \eta_{j}(x) \mathrm{d}\nu(x) \right| &\leq |\langle \eta_{j}, \, \hat{\mu} - \bar{\mu}_{0} \rangle_{\mathcal{M}}| + |\tilde{\mu}_{0}| \left(\mathcal{X}\right) \leq \|p_{j}\| \left\| \Phi(\hat{\mu} - \bar{\mu}_{0}) \right\| + |\tilde{\mu}_{0}| \left(\mathcal{X}\right) \\ &\leq \|p_{j}\| \left(\delta + \|\Phi\hat{\mu} - y\|\right) + |\tilde{\mu}_{0}| \left(\mathcal{X}\right) \leq \|p_{j}\| \left(2\delta + 2\lambda \|p\| + 2\sqrt{\lambda |\tilde{\mu}_{0}| \left(\mathcal{X}\right)}\right) + |\tilde{\mu}_{0}| \left(\mathcal{X}\right) \\ &\leq 2\delta \|p_{j}\| + 2\lambda \|p\| \|p_{j}\| + \lambda \|p_{j}\|^{2} + 2 |\tilde{\mu}_{0}| \left(\mathcal{X}\right) \end{aligned}$$

using $\sqrt{ab} \leq (a+b)/2$. Therefore, we obtain

$$|\hat{a}_j - a_j| \lesssim \|p_j\| \left(\delta + \lambda \|p_j\|\right) + \varepsilon_0^{-1} \left(\delta \|p\| + |\tilde{\mu}_0| \left(\mathcal{X}\right)\right)$$

Additional certificates. Proposition 2 assumes the construction of additional functions $\eta_j \in \text{Im}(\Phi^*)$, which are essentially similar to non-degenerate certificates but with all "signs" to interpolate put to 0 except for one. As we will see, they are even simpler to construct than η : indeed, the reason one has to resort to the random signs assumption (as in [54]) or to the golfing scheme (as in this paper) is that the Euclidean norm of the vector of signs $(\text{sign}(a_i))_{i=1}^s$ appears in the proof, which results in a spurious term \sqrt{s} . When constructing the η_j , this problem does not occur, since only one sign is non-zero.

Relation to previous works. Note that (21) and (24), without the inexact sparsity term $|\tilde{\mu}_0|(\mathcal{X})$, were previously presented in [37] in the context of sampling Fourier coefficients and in a more general setting in [4]. However, the statement in [4] is given in terms of orthonormal systems, and the so-called Bernstein Isolation Property which imposes that $|P'(x)| \leq Cm^2 ||P||_{\infty}$ for all $P \in \text{Im}(\Phi^*)$. These conditions can be difficult to check in our setting of random sampling and were imposed only to ensure the existence of nondegenerate dual certificates, and to have explicit control on the constant C. For completeness, we still present the proof of (21) under nondegeneracy assumptions, and we later establish that these nondegeneracy assumptions hold, under appropriate separation conditions imposed via $\mathfrak{d}_{\mathfrak{g}}$.

In [14], one could also obtain bounds $\sum_{j=1}^{s} |\hat{a}_j - a_j| \leq \delta$ in the case of Fourier sampling, however, to prove such a statement, one is required to construct a trigonometric function (a dual certificate) which interpolates arbitrary sign patterns. In the case of subsampling, such an approach cannot lead to sharp dependency on s, since in the real setting, one is then required to show the existence of 2^s random polynomials corresponding to all possible sign patterns. We therefore settle for the bound (24) in this paper. We remark that being able to construct dual functions which interpolate arbitrary signs patterns lead to Wasserstein-1 error bounds, as opposed to Wasserstein-2 error bounds presented here.

Finally, we mention the more recent work of [34] which presents stability bounds for the sparse spikes problem where one restricts to positive measures and where the sampling functions form a T-systems. Under a positivity constraint (rather than total variation penalisation), they derive stability bounds in terms of optimal partial transport distances. We stress that since we consider more general measurement operators than T-systems in this work, we consider transport distances under the Fisher metric as opposed to the Euclidean metric. Moreover, another difference is that our error bounds use the Wasserstein-2 distance, whereas they use the Wasserstein-1 distance – the reason is that since they do not consider random subsampling, their proofs in fact follow the work of [14] to construct dual certificates which interpolate arbitrary sign patterns.

4 Nondegenerate limit certificates

In this section, we provide the second building block of our main theorem: a generic way to ensure the existence and construct non-degenerates dual certificates, when $m \to \infty$ and the sought-after Diracs satisfy a minimal separation condition with respect to the metric $\mathfrak{d}_{\mathfrak{g}}$.

4.1 Notions of differential geometry

We start with additional definitions in differential Riemannian geometry. All these notions can be found in the textbook [1], to which we refer the reader for further details. In many instances, we extend classical definitions to the complex case in a natural way.

Riemannian gradient and Hessian. Let $f : \mathbb{R}^d \to \mathbb{C}$ be a smooth function. The Riemannian gradient $\operatorname{grad} f(x) \in \mathbb{C}^d$ and Riemannian Hessian $\operatorname{Hess} f(x) : \mathbb{C}^d \to \mathbb{C}^d$, which is a linear mapping, can be defined as:

$$\operatorname{grad} f(x) = \mathfrak{g}_x^{-1} \nabla f(x)$$
$$\langle \operatorname{Hess} f(x)[e_i], e_j \rangle_x = \partial_i \partial_j f(x) - \Gamma_{ij}(x)^\top \nabla f(x)$$

where ∇, ∂_i are the classical Euclidean gradient and partial derivatives, and the $\{e_i\}$ are the canonical basis of \mathbb{R}^d . The $\Gamma_{ij}(x) = [\Gamma_{ij}^k(x)]_k \in \mathbb{R}^d$ are the Christoffel symbols, here equal to:

$$\Gamma_{ij}^k(x) = \frac{1}{2} \sum_{\ell} g^{k\ell}(x) \left(\partial_i g_{\ell j}(x) + \partial_j g_{\ell i}(x) - \partial_\ell g_{ij}(x) \right) \,,$$

where $g_{ij}(x) = [\mathfrak{g}_x]_{ij}$ and $g^{ij}(x) = [\mathfrak{g}_x^{-1}]_{ij}$. Finally we denote by $\mathrm{H}f(x) \in \mathbb{C}^{d \times d}$ the matrix that contains these terms: $\mathrm{H}f(x) \stackrel{\text{def.}}{=} \left(\langle \mathrm{Hess}f(x)[e_i], e_j \rangle_x \right)_{ij}$.

For r = 0, 1, 2, the "covariant derivative" $D_r[f](x) : (\mathbb{C}^d)^r \to \mathbb{C}$ are mappings (or scalar in the case r = 0) defined as:

$$\begin{aligned} \mathbf{D}_0\left[f\right](x) &\stackrel{\text{def.}}{=} f(x) \\ \mathbf{D}_1\left[f\right](x)[v] &\stackrel{\text{def.}}{=} \langle v, \operatorname{grad} f(x) \rangle_x = v^* \nabla f(x) \\ \mathbf{D}_2\left[f\right](x)[v,v'] &\stackrel{\text{def.}}{=} \langle \operatorname{Hess} f(x)[v], v' \rangle_x = v^* \mathrm{H} f(x) v' \end{aligned}$$

We define associated operator norms

$$\begin{split} \|\mathbf{D}_{1}\left[f\right](x)\|_{x} &\stackrel{\text{def.}}{=} \sup_{\|v\|_{x}=1} \mathbf{D}_{1}\left[f\right](x)[v] = \left\|\mathbf{g}_{x}^{-\frac{1}{2}}\nabla f(x)\right\|_{2} \\ \|\mathbf{D}_{2}\left[f\right](x)\|_{x} &\stackrel{\text{def.}}{=} \sup_{\|v\|_{x}=1, \|v'\|_{x}=1} \mathbf{D}_{2}\left[f\right](x)[v,v'] = \left\|\mathbf{g}_{x}^{-\frac{1}{2}}\mathbf{H}f(x)\mathbf{g}_{x}^{-\frac{1}{2}}\right\|_{2} \end{split}$$

where we recall that $\|\cdot\|_x$ is defined by (14).

Covariant derivatives of the kernel. Recall the definition of the limit covariance kernel (10). Given $0 \leq i, j \leq 2$, let $K^{(ij)}(x, x')$ be a "bi"-multilinear map, defined for $Q \in (\mathbb{C}^d)^i$ and $V \in (\mathbb{C}^d)^j$ as

$$[Q]K^{(ij)}(x,x')[V] \stackrel{\text{def.}}{=} \mathbb{E}[\overline{\mathcal{D}_i[\varphi_\omega](x)[Q]}\mathcal{D}_j[\varphi_\omega](x')[V]]].$$
(25)

In the case $i, j \leq 1$, note that these admits simplified expressions: $K^{(00)}(x, x') = K(x, x')$, $[v]K^{(10)}(x, x') = v^{\top} \nabla_1 K(x, x')$ and $[v]K^{(11)}(x, x')[v'] = v^{\top} \nabla_1 \nabla_2 K(x, x')\overline{v'}$. Define the operator norm of $K^{(ij)}(x, x')$ as

$$\left\| K^{(ij)}(x,x') \right\|_{x,x'} \stackrel{\text{\tiny def.}}{=} \sup_{Q,V} \left| [Q] K^{(ij)}(x,x')[V] \right|$$
(26)

where the supremum is over all $V = [v_1, \ldots, v_i]$ with $||v_\ell||_x \leq 1$ for all $\ell \in [i]$, and all $Q = [q_1, \ldots, q_j]$ with $||q_\ell||_{x'} \leq 1$ for all $\ell \in [j]$. We will sometimes overload the notations and write $||\cdot||_x$ when the dependence is only on x, i.e. for $K^{(ij)}$ where j = 0. Note that, in particular,

$$\left\| K^{(10)}(x,x') \right\|_{x} = \left\| \mathfrak{g}_{x}^{-\frac{1}{2}} \nabla_{1} K(x,x') \right\|_{2}, \quad \left\| K^{(11)}(x,x') \right\|_{x,x'} = \left\| \mathfrak{g}_{x}^{-\frac{1}{2}} \nabla_{1} \nabla_{2} K(x,x') \mathfrak{g}_{x'}^{-\frac{1}{2}} \right\|_{2}$$
 (27)
and
$$\left\| K^{(20)}(x,x') \right\|_{x} = \left\| \mathfrak{g}_{x}^{-\frac{1}{2}} \mathrm{H}[K(\cdot,x')](x) \mathfrak{g}_{x}^{-\frac{1}{2}} \right\|_{2}$$

All these definitions are naturally extended to the covariance kernel \hat{K} by replacing the expectation \mathbb{E} in (25) by an empirical expectation over $\omega_1, \ldots, \omega_m$.

4.2 Non-degenerate dual certificate with $m \to \infty$

Recall the definition of the covariance kernel (9). Following [15], a natural approach towards constructing a dual certificate is by interpolating the sign vector $\operatorname{sign}(a_j)$ using the functions $\hat{K}(x_j, \cdot)$ and $\hat{K}^{(10)}(x_j, \cdot)$, since we have

$$\left\{ \eta \stackrel{\text{\tiny def.}}{=} \sum_{j=1}^{s} \alpha_{1,j} \hat{K}(x_j, \cdot) + \sum_{j=1}^{s} [\alpha_{2,j}] \hat{K}^{(10)}(x_j, \cdot) \; ; \; \alpha_{1,j} \in \mathbb{C}, \; \alpha_{2,j} \in \mathbb{C}^d \right\} \subset \operatorname{Im}(\Phi^*) \, .$$

Using the gradients of the kernel allows to additionally impose that $\nabla \eta(x_i) = 0$, which is a necessary (but not sufficient) condition for the dual certificate to reach its maximum amplitude in x_i . Usual proofs then show that, under minimal separation, applying this strategy indeed yields a non-degenerate dual certificate.

We first consider the case where one has access to arbitrarily many measurements $(m \to \infty)$, and to this end, we consider the limit covariance kernel K defined in (10). Let us introduce some handy notations that will be particularly useful in later proofs (Section 6). Our aim is to find coefficients $(\alpha_{1,j})_{j=1}^s \in \mathbb{C}^s$ and $(\alpha_{2,j})_{j=1}^s \in (\mathbb{C}^d)^s$ such that

$$\eta \stackrel{\text{\tiny def.}}{=} \sum_{j=1}^{s} \alpha_{1,j} K(x_j, \cdot) + \sum_{j=1}^{s} [\alpha_{2,j}] K^{(10)}(x_j, \cdot)$$
(28)

satisfies $\eta(x_j) = \text{sign}(a_j)$ and $\nabla \eta(x_j) = 0$ for all j = 1, ..., s. Note that these s(d+1) constraints can be written as the linear system

$$\Upsilon\begin{pmatrix}\alpha_1\\\alpha_2\end{pmatrix} = \begin{pmatrix}(\operatorname{sign}(a_i))_{i=1}^s\\0_{sd}\end{pmatrix} \stackrel{\text{\tiny def.}}{=} \mathbf{u}_s , \qquad (29)$$

where $\Upsilon \in \mathbb{R}^{s(d+1) \times s(d+1)}$ is a real symmetric matrix defined as

$$\Upsilon \stackrel{\text{\tiny def.}}{=} \mathbb{E}_{\omega}[\gamma(\omega)\gamma(\omega)^*] \in \mathbb{C}^{s(d+1)\times s(d+1)},\tag{30}$$

with the vector $\gamma(\omega) \in \mathbb{C}^{s(d+1)}$ defined as

$$\gamma(\omega) \stackrel{\text{\tiny def.}}{=} \left(\left(\varphi_{\omega}(x_i) \right)_{i=1}^s, \left(\nabla \varphi_{\omega}(x_i)^\top \right)_{i=1}^s \right)^\top.$$
(31)

Assuming that Υ is invertible, we can therefore rewrite (28) as $\eta(x) = (\Upsilon^{-1}\mathbf{u}_s)^{\top}\mathbf{f}(x)$, where

$$\mathbf{f}(x) \stackrel{\text{def.}}{=} \mathbb{E}_{\omega}[\overline{\gamma(\omega)}\varphi_{\omega}(x)] = \left(\left(K(x_i, x)\right)_{i=1}^s, \left(\nabla_1 K(x_i, x)^{\top}\right)_{i=1}^s \right)^{\top} \in \mathbb{R}^{s(d+1)}.$$
(32)

We also define the block diagonal normalisation matrix $D_{\mathfrak{g}} \in \mathbb{R}^{s(d+1) \times s(d+1)}$ as

$$D_{\mathfrak{g}} \stackrel{\text{def.}}{=} \begin{pmatrix} \mathrm{Id}_{s} & & & \\ & \mathfrak{g}_{x_{1}}^{-\frac{1}{2}} & & \\ & & \ddots & \\ & & & \mathfrak{g}_{x_{s}}^{-\frac{1}{2}} \end{pmatrix}$$
(33)

so that $\Upsilon = D_{\mathfrak{g}} \Upsilon D_{\mathfrak{g}}$ has constant value 1 along its diagonal.

We will prove in Theorem 2 below that η of the form (28) is indeed nondegenerate, provided that there is sufficient curvature on $K(x, \cdot)$ in a small neighbourhood around x and $\min_{k \neq j} \mathfrak{d}_{\mathfrak{g}}(x_j, x_k) \ge \Delta$ where Δ is the distance at which the kernel and its partial derivatives are sufficiently small (to allow for interpolation with $K(\cdot, x_j)$ with minimal inference between the point sources). To do so we need the following definition.

Definition 3. Given r > 0, the local curvature constants $\bar{\varepsilon}_0(r)$ and $\bar{\varepsilon}_2(r)$ of K are defined as

$$\bar{\varepsilon}_{0}(r) \stackrel{\text{def.}}{=} \sup \left\{ \varepsilon \; ; \; K(x,x') \leqslant 1 - \varepsilon, \; \forall x, x' \in \mathcal{X} \; s.t. \; \mathfrak{d}_{\mathfrak{g}}(x,x') \geqslant r \right\}$$
$$\bar{\varepsilon}_{2}(r) \stackrel{\text{def.}}{=} \sup \left\{ \varepsilon \; ; \; -K^{(02)}(x',x)[v,v] \geqslant \varepsilon \left\| v \right\|_{x}^{2}, \; \forall x, x' \in \mathcal{X} \; s.t. \; \mathfrak{d}_{\mathfrak{g}}(x,x') < r, \forall v \in \mathbb{R}^{d} \right\}$$

Given h > 0 and $s \in \mathbb{N}$, the kernel width of K is defined as

$$\Delta(h,s) \stackrel{\text{\tiny def.}}{=} \inf \left\{ \Delta \; ; \; \sum_{k=2}^{s} \left\| K^{(ij)}(x_1, x_k) \right\|_{x_1, x_k} \leqslant h, \; (i,j) \in \{0,1\} \times \{0,2\}, \; \{x_k\}_{k=1}^s \in \mathcal{S}_\Delta \right\}$$

where $\mathcal{S}_{\Delta} \stackrel{\text{def.}}{=} \{(x_k)_{k=1}^s \in \mathcal{X}^s ; \mathfrak{d}(x_k, x_\ell) \ge \Delta, \forall k \neq \ell\}$ is the set of k-tuples of Δ -separated points. We define $\inf \emptyset \stackrel{\text{def.}}{=} +\infty$.

Intuitively, these notions are similar to those appearing in the definition of non-degenerate dual certificates (and will ultimately serve in the proof of existence of such certificates): r is a neighborhood size, $\bar{\varepsilon}_0$ represents the distance to 1 of the kernel away from x = x', and $\bar{\varepsilon}_2$ is the "curvature" of the kernel when $x \approx x'$. Finally, Δ is the "minimal separation" under which s Diracs have minimal interference between them, or, in other words, the covariance kernel and its derivatives have low value. We formalize this in the following assumption.

Assumption 1 (Assumptions on the kernel.). Suppose that K is a real-valued kernel. For $i, j \leq 2$ and $i+j \leq 3$, assume that $B_{ij} \stackrel{\text{def.}}{=} \sup_{x,x' \in \mathcal{X}} \left\| K^{(ij)}(x,x') \right\|_{x,x'} < \infty$ and denote $B_i \stackrel{\text{def.}}{=} B_{0i} + B_{1i} + 1$. Assume that K has positive curvature constants $\bar{\varepsilon}_0$ and $\bar{\varepsilon}_2$ at radius $0 < r_{\text{near}} < B_{02}^{-\frac{1}{2}}$. Let $s \in \mathbb{N}$ be such that $\Delta \stackrel{\text{def.}}{=} \Delta(h, s) < \infty$ with $h \leq \frac{1}{64} \min\left(\frac{\bar{\varepsilon}_0}{B_0}, \frac{\bar{\varepsilon}_2}{B_2}\right)$.

Under this assumption, the following theorem, which is the main result of this section, proves that a limit non-degenerate dual certificate can be constructed under minimal separation.

Theorem 2. Under Assumption 1, for all $\{x_k\}_{k=1}^s$ with $\min_{k \neq \ell} \mathfrak{d}_{\mathfrak{g}}(x_k, x_\ell) \ge \Delta$, there exists a unique function η of the form (28) which is $(\frac{\overline{\varepsilon}_0}{2}, \frac{\overline{\varepsilon}_2}{4}, r_{\text{near}})$ -nondegenerate. Moreover,

$$\left\|\overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\eta\right](x) - K^{(02)}(x_j, x)\right\|_x \leqslant \frac{\bar{\varepsilon}_2}{16} \qquad \forall x \in \mathcal{B}_{\mathfrak{d}_g}(x_j; r_{\operatorname{near}}).$$

We delay the (slightly lengthy) proof of this result to the next subsection. Before that, we make a few comments.

Dependency on s. As we have seen in the examples of Section 2, for a constant h we generally let the minimal separation $\Delta = W(h, s)$ depend on s. Indeed, in dimension d, it is well known one can pack C^d Δ -separated points in a ball of radius 2Δ for some constant C (this is known as the kissing number). Hence, there exist s Δ -separated points such that

$$\sum_{k=2}^{\circ} \left\| K^{(ij)}(x_1, x_k) \right\|_{x_1, x_k} \ge \min\left(C^d, s\right) \sup_{\mathfrak{d}(x, x') \ge \Delta} \left\| K^{(ij)}(x, x') \right\|_{x, x'}.$$

Therefore, while the kernel width can be independent of s in low dimensions (and the trick is then to upper bound this by a constant bound $s \to \infty$, assuming the sum on the l.h.s. converges), as d increases, the dependence on s will become inevitable, otherwise Δ generally depends exponentially on d. **Babel function.** The attentative reader might recognise the similarity of definition of kernel width W(h, s) with the Babel function from compressed sensing [56], if we restrict the definition to (i, j) = (0, 0) and recall that $K(x, x') = \mathbb{E}_{\omega}[\overline{\varphi_{\omega}(x)}\varphi_{\omega}(x')]$. The Babel function of a $m \times N$ matrix A with columns \mathbf{a}_{j} is defined as

$$\mu(s) = \max_{i \in [N]} \max\left\{ \sum_{j \in S} |\langle \mathbf{a}_i, \, \mathbf{a}_j \rangle| \ ; \ S \subset [N], |S| = s, i \neq S \right\},\$$

and small value of $\mu(s)$ ensure that the sub-matrix $A_S^*A_S$, where A_S is the matrix A restricted to index set S with $|S| \leq s$, is well conditioned and invertible. Furthermore, recovery guarantees for Basis Pursuit and Orthogonal Matching Pursuit can be stated in terms of $\mu(s)$. In Theorem 2, sufficient kernel width also ensures that $\Phi_x^*\Phi_x$ is well conditioned and thereby provide performance guarantees for the BLASSO.

4.3 Proof of Theorem 2

Before proving Theorem 2, we illustrate the link between *curvature of the kernel* as represented by $\bar{\varepsilon}_2$ in Def. 3 and the *quadratic decay* condition $|\eta| \leq 1 - \varepsilon \mathfrak{d}_{\mathfrak{g}}(x_i, \cdot)^2$ that we used in the definition of non-degenerate certificates (Def. 2). The resulting condition (35) is the one that we are actually going to prove in practice. The following Lemma is based on a generalized second-order Taylor expansion.

Lemma 2. Let $x_0 \in \mathcal{X}$ and $a \in \mathbb{C}$ with |a| = 1. Suppose that for some $\varepsilon > 0$, B > 0 and $0 < r \leq B^{-\frac{1}{2}}$ we have: for all $x \in \mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_0; r)$ and $v \in \mathbb{C}^d$, it holds that $-K^{(02)}(x_0, x)[v, v] \ge \varepsilon ||v||_x^2$ and $||K^{(02)}(x_0, x)||_x \le B$. Let $\eta : \mathcal{X} \to \mathbb{C}$ be a smooth function.

(i) If
$$\eta(x_0) = 0$$
, $\nabla \eta(x_0) = 0$ and

$$\left\| \mathcal{D}_{2}\left[\eta\right](x) \right\|_{x} \leqslant \delta \qquad \forall x \in \mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_{0}; r) \tag{34}$$

then $|\eta(x)| \leq \delta \mathfrak{d}_{\mathfrak{g}}(x_0, x)^2$ for all $x \in \mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_0; r)$.

(*ii*) If $\eta(x_0) = a$, $\nabla \eta(x_0) = 0$ and

$$\left\| \overline{a} \mathcal{D}_{2} \left[\eta \right] (x) - K^{(02)}(x_{0}, x) \right\|_{x} \leq \delta \qquad \forall x \in \mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_{0}; r)$$
(35)

for some $\delta < \frac{\varepsilon}{2}$, then, for all $x \in \mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_0; r)$ we have $|\eta(x)| \leq 1 - \varepsilon' \mathfrak{d}_{\mathfrak{g}}(x_0, x)^2$ with $\varepsilon' = \frac{\varepsilon - 2\delta}{2}$.

Proof. We prove (*ii*), the proof for (*i*) is similar and simpler. Using (35) and the assumption on $K^{(02)}$, we can deduce that for all $v \in \mathbb{R}^d$ we have

$$\operatorname{Re}\left(\overline{a}\operatorname{D}_{2}\left[\eta\right](x)[v,v]\right) \leqslant -(\varepsilon-\delta)\left\|v\right\|_{x}^{2} \quad \text{and} \quad \left|\operatorname{Im}\left(\overline{a}\operatorname{D}_{2}\left[\eta\right](x)[v,v]\right)\right| \leqslant \delta\left\|v\right\|_{x}^{2}$$

Given a geodesic $\gamma : [0,1] \to \mathbb{R}^d$, it is a well-known property that $\ddot{\gamma} + \sum_{i,j} \Gamma_{ij}(\gamma) \dot{\gamma}_i \dot{\gamma}_j = 0$ where we recall that $\Gamma_{ij} \in \mathbb{R}^d$ are the Christoffel symbols. Therefore, we have

$$\begin{aligned} \frac{\mathrm{d}^2}{\mathrm{d}t^2} \eta(\gamma(t)) &= \dot{\gamma}(t)^\top \nabla^2 \eta(\gamma(t)) \dot{\gamma}(t) + \nabla \eta(t)^\top \ddot{\gamma}(t) \\ &= \dot{\gamma}(t)^\top \nabla^2 \eta(\gamma(t)) \dot{\gamma}(t) - \nabla \eta(t)^\top \left(\sum_{ij} \Gamma_{ij}(\gamma(t)) \dot{\gamma}_j(t) \dot{\gamma}_k(t) \right) \\ &= \dot{\gamma}(t)^\top \mathrm{H}\eta(\gamma(t)) \dot{\gamma}(t) = \mathrm{D}_2 \left[\eta \right] (\gamma(t)) [\dot{\gamma}(t), \dot{\gamma}(t)] \end{aligned}$$

So, given any geodesic path with $\gamma(0) = x_0$ and $\gamma(1) = x$, since of course we have $\mathfrak{d}_{\mathfrak{g}}(x_0, \gamma(t)) \leq \mathfrak{d}_{\mathfrak{g}}(x_0, x) \leq r$, by applying the inequalities above:

$$\operatorname{Re}\left(\overline{a}\eta(x)\right) = \operatorname{Re}\left(\overline{a}\left(\eta(x_{0}) + \nabla\eta(x_{0})^{\top}\dot{\gamma}(0) + \frac{1}{2}\int_{0}^{1}(1-t)\frac{\mathrm{d}^{2}}{\mathrm{d}t^{2}}\eta(\gamma(t))dt\right)\right)$$
$$= 1 + \frac{1}{2}\int_{0}^{1}(1-t)\operatorname{Re}\left(\overline{a}\operatorname{D}_{2}\left[\eta\right](\gamma(t))[\dot{\gamma}(t),\dot{\gamma}(t)]\right)\mathrm{d}t$$
$$\leq 1 - (\varepsilon - \delta)\int_{0}^{1}(1-t)\left\|\dot{\gamma}(t)\right\|_{\gamma(t)}^{2}\mathrm{d}t = 1 - \frac{(\varepsilon - \delta)}{2}\mathfrak{d}_{\mathfrak{g}}(x_{0}, x)^{2}.$$
(36)

where the last line follows because $\|\dot{\gamma}(t)\|_{\gamma(t)}$ is constant for all $t \in [0,1]$. Similarly, we can show that $\operatorname{Re}\left(\overline{a}\eta(x)\right) \geq 1 - \frac{B+\delta}{2}\mathfrak{d}_{\mathfrak{g}}(x_0,x)^2 \geq 0$ since $r \leq B^{-\frac{1}{2}}$, and $|\operatorname{Im}\left(\overline{a}\eta(x)\right)| \leq \frac{\delta}{2}\mathfrak{d}_{\mathfrak{g}}(x_0,x)^2$, from which we got $|\eta(x)| \leq \operatorname{Re}\left(\overline{a}\eta(x)\right) + |\operatorname{Im}\left(\overline{a}\eta(x)\right)| \leq 1 - \frac{\varepsilon-2\delta}{2}\mathfrak{d}_{\mathfrak{g}}(x_0,x)^2$.

We can now proceed with the proof of Theorem 2.

Proof of Theorem 2. Recall the block diagonal matric $D_{\mathfrak{g}}$ from (33). The system (29) is equivalent to

$$\tilde{\Upsilon} \begin{pmatrix} \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \end{pmatrix} = \mathbf{u}_s. \tag{37}$$

where $\tilde{\Upsilon} = D_{\mathfrak{g}} \Upsilon D_{\mathfrak{g}}$ and $\tilde{\alpha} = D_{\mathfrak{g}}^{-1} \alpha$. So, if $\tilde{\Upsilon}$ is invertible, then we can write $\eta = (\tilde{\Upsilon}^{-1} \mathbf{u}_s)^{\top} D_{\mathfrak{g}} \mathbf{f} = (\Upsilon^{-1} \mathbf{u}_s)^{\top} \mathbf{f}$. Therefore, we will proceed as follows: First, prove that $\tilde{\Upsilon}$ is invertible. Second, bound the coefficients α_1 and α_2 . Third, prove that η is nondegenerate.

We first prove that the matrix $\tilde{\Upsilon}$ is invertible. To this end, we decompose it into blocks

$$\tilde{\Upsilon} = \begin{pmatrix} \Upsilon_0 & \Upsilon_1^\top \\ \Upsilon_1 & \Upsilon_2 \end{pmatrix}$$
(38)

where $\Upsilon_0 \in \mathbb{C}^{s \times s}$, $\Upsilon_1 \in \mathbb{C}^{sd \times s}$ and $\Upsilon_2 \in \mathbb{C}^{sd \times sd}$ are defined as

$$\Upsilon_0 \stackrel{\text{def.}}{=} (K(x_i, x_j))_{i,j=1}^s, \qquad \Upsilon_1 \stackrel{\text{def.}}{=} (\mathfrak{g}_{x_i}^{-\frac{1}{2}} \nabla_1 K(x_i, x_j))_{i,j=1}^s, \quad \Upsilon_2 \stackrel{\text{def.}}{=} (\mathfrak{g}_{x_i}^{-\frac{1}{2}} \nabla_1 \nabla_2 K(x_i, x_j) \mathfrak{g}_{x_j}^{-\frac{1}{2}})_{i,j=1}^s.$$

To prove the invertibility of $\tilde{\Upsilon}$, it suffices to prove that Υ_2 and its Schur complement $\Upsilon_S \stackrel{\text{def.}}{=} \Upsilon_0 - \Upsilon_1 \Upsilon_2^{-1} \Upsilon_1^{\top}$ are both invertible. To show that Υ_2 is invertible, we define $A_{ij} = \mathfrak{g}_{x_i}^{-\frac{1}{2}} \nabla_1 \nabla_2 K(x_i, x_j) \mathfrak{g}_{x_j}^{-\frac{1}{2}}$, such that Υ_2 has the form:

$$\Upsilon_2 = \begin{pmatrix} \operatorname{Id} & A_{12} & \dots & A_{1s} \\ A_{21} & \operatorname{Id} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ A_{s1} & \dots & \dots & \operatorname{Id} \end{pmatrix}$$

and by Lemma 5 in Appendix A.1, Assumption 1 and (27), we have

$$\| \mathrm{Id} - \Upsilon_2 \|_{\mathrm{block}} \leq \max_i \sum_{j \neq i} \|A_{ij}\|_2 = \max_i \sum_{j \neq i} \left\| K^{(11)}(x_i, x_j) \right\|_{x_i, x_j} \leq h \leq 1/32$$

Since $\|\mathrm{Id} - \Upsilon_2\|_{\mathrm{block}} < 1$, Υ_2 is invertible, and we have $\|\Upsilon_2^{-1}\|_{\mathrm{block}} \leq \frac{1}{1 - \|\mathrm{Id} - \Upsilon_2\|_{\mathrm{block}}} \leq \frac{4}{3}$. Next, again with Lemma 5, we can bound

$$\|\operatorname{Id} - \Upsilon_0\|_{\infty} = \max_i \sum_{j \neq i} |K(x_i, x_j)| \leq h$$
$$\|\Upsilon_1\|_{\infty \to \operatorname{block}} \leq \max_i \sum_j \left\|\mathfrak{g}_{x_i}^{-\frac{1}{2}} \nabla_1 K(x_i, x_j)\right\|_2 = \max_i \sum_j \left\|K^{(10)}(x_i, x_j)\right\|_{x_i} \leq h$$

since $K^{(10)}(x, x) = 0$. Hence, we have

$$\|\mathrm{Id} - \Upsilon_S\|_{\infty} \leqslant \|\mathrm{Id} - \Upsilon_0\|_{\infty} + \|\Upsilon_1^{\top}\|_{\mathrm{block} \to \infty} \|\Upsilon_2^{-1}\|_{\mathrm{block}} \|\Upsilon_1\|_{\infty \to \mathrm{block}} \leqslant h + \frac{4}{3}h^2 \leqslant 2h \stackrel{\mathrm{\tiny def.}}{=} h' < 1.$$
(39)

Therefore the Schur complement of $\tilde{\Upsilon}$ is invertible and so is $\tilde{\Upsilon}$. Moreover, $\|\Upsilon_S^{-1}\|_{\infty} \leq \frac{1}{1-h'}$. We can now define:

$$\tilde{\alpha} = \tilde{\Upsilon}^{-1} \mathbf{u}_s = \begin{pmatrix} \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \end{pmatrix}$$

and, as described above, $\alpha = D_{\mathfrak{g}}^{-1}\tilde{\alpha}$. The Schur's complement of $\tilde{\Upsilon}$ allows us to express α_1 and α_2 as

$$\begin{pmatrix} \tilde{\alpha}_1\\ \tilde{\alpha}_2 \end{pmatrix} = \begin{pmatrix} \Upsilon_S^{-1} \operatorname{sign}(a)\\ -\Upsilon_2^{-1} \Upsilon_1 \Upsilon_S^{-1} \operatorname{sign}(a) \end{pmatrix}$$
(40)

and therefore we can bound

$$\begin{aligned} \|\alpha_1\|_{\infty} \leqslant \left\|\Upsilon_S^{-1}\right\|_{\infty} \leqslant \frac{1}{1-h'} \\ \max_i \|\alpha_{2,i}\|_{x_i} &= \|\tilde{\alpha}_2\|_{\text{block}} \leqslant \left\|\Upsilon_2^{-1}\right\|_{\text{block}} \left\|\Upsilon_1\right\|_{\infty \to \text{block}} \left\|\Upsilon_S^{-1}\right\|_{\infty} \leqslant 4h \end{aligned}$$

Moreover, we have

$$\|\alpha_1 - \operatorname{sign}(a)\|_{\infty} \leq \|\operatorname{Id} - \Upsilon_S^{-1}\|_{\infty} \leq \|\Upsilon_S^{-1}\|_{\infty} \|\operatorname{Id} - \Upsilon_S\|_{\infty} \leq \frac{h'}{1 - h'}$$
(41)

We can now prove that η is non-degenerate. For any x such that $\mathfrak{d}_{\mathfrak{g}}(x, x_i) \ge r_{\text{near}}$ for all x_i 's, there exists at most one index i such that $\mathfrak{d}_{\mathfrak{g}}(x, x_i) < \Delta/2$ and so, for all $j \neq i$, we have $\mathfrak{d}_{\mathfrak{g}}(x, x_j) \ge \Delta/2$. Therefore,

$$\begin{aligned} |\eta(x)| &= \left| \alpha_{1,i} K(x_i, x) + \sum_{j \neq i} \alpha_{1,j} K(x_j, x) + [\alpha_{2,i}] K^{(10)}(x_i, x) + \sum_{j \neq i} [\alpha_{2,j}] K^{(10)}(x_j, x) \right| \\ &\leqslant \left\| \alpha_1 \right\|_{\infty} \left(|K(x_i, x)| + \sum_{j \neq i} |K(x_j, x)| \right) + \max_i \| \alpha_{2,i} \|_{x_i} \left(\left\| K^{(10)}(x_i, x) \right\|_{x_i} + \sum_{j \neq i} \left\| K^{(10)}(x_j, x) \right\|_{x_j} \right) \\ &\leqslant \frac{1}{1 - h'} \left(1 - \bar{\varepsilon}_0 + h \right) + 4h \left(B_{10} + h \right) \leqslant 1 - \frac{\bar{\varepsilon}_0}{2}. \end{aligned}$$

Now, let x be such that $\mathfrak{d}_{\mathfrak{g}}(x_i, x) \leq r_{\text{near}}$. Similarly, for all $j \neq i$ we have $\mathfrak{d}_{\mathfrak{g}}(x, x_j) \geq \Delta/2$. Observe that

$$\overline{\operatorname{sign}(a_i)} D_2[\eta](x) = K^{(02)}(x_i, x) + \left(\overline{\operatorname{sign}(a_i)}\alpha_{1,i} - 1\right) K^{(02)}(x_i, x) + \overline{\operatorname{sign}(a_i)} \left[\sum_{j \neq i} \alpha_{1,j} K^{(02)}(x_j, x) + [\alpha_{2,i}] K^{(12)}(x_i, x) + \sum_{j \neq i} [\alpha_{2,j}] K^{(12)}(x_j, x) \right]$$

So,

$$\begin{split} \left\| \overline{\operatorname{sign}(a_{i})} \mathcal{D}_{2}\left[\eta\right](x) - K^{(02)}(x_{i}, x) \right\|_{x} \\ \leqslant \left\| \left(\overline{\operatorname{sign}(a_{i})} \alpha_{1,i} - 1 \right) K^{(02)}(x_{i}, x) + \overline{\operatorname{sign}(a_{i})} \left[\sum_{j \neq i} \alpha_{1,j} K^{(02)}(x_{j}, x) + [\alpha_{2,i}] K^{(12)}(x_{i}, x) + \sum_{j \neq i} [\alpha_{2,j}] K^{(12)}(x_{j}, x) \right] \right\|_{x} \\ \leqslant \frac{h'}{1 - h'} B_{02} + h \|\alpha_{1}\|_{\infty} + \max_{i} \|\alpha_{2,i}\|_{x_{i}} (B_{12} + h) \leqslant \frac{h'}{1 - h'} B_{02} + \frac{h}{1 - h'} + 4h B_{12} + 4h^{2} \leqslant \frac{\bar{\varepsilon}_{2}}{16} \end{split}$$
We conclude using Lemma 2 and $\frac{\bar{\varepsilon}_{2} - 2\bar{\varepsilon}_{2}/16}{\bar{\varepsilon}_{2}} \geq \bar{\varepsilon}_{2}/4$

We conclude using Lemma 2 and $\frac{c_2}{2}$ $\stackrel{--}{=} \geq \varepsilon_2/4.$

5 Sparse recovery

In this section, we formulate our main contribution, Theorem 3, which is a detailed version of Theorem 1. In previous sections, we have shown that the existence of a non-degenerate dual certificates implies sparse recovery guarantees, and that in the limit case $m \to \infty$, a minimal separation assumption implies the existence of a dual certificate. Our main theorem is obtained by bounding the deviations from the limit case when m is finite. We do so by extending the celebrated golfing scheme [41] to the infinite-dimensional case. We first begin by our assumptions on the feature functions φ_{ω} .

5.1 Almost bounded random features

In order to bound the variation between K and \hat{K} , we would ideally like the features φ_{ω} and their derivatives to be uniformly bounded for all ω . However this may not be the case: think of $e^{i\omega^{\top}x}$, which does not have a uniformly bounded gradient when the support of the distribution Λ is not bounded. On the other hand, if $\Lambda(\omega)$ has sufficient decay as $||\omega||$ increases, one could argue that the selected random features and their derivatives are uniformly bounded with high probability. For $r \in \{0, 1, 2\}$, we define the random variables

$$L_{r}(\omega) \stackrel{\text{def.}}{=} \sup_{x \in \mathcal{X}} \left\| \mathbf{D}_{r} \left[\varphi_{\omega} \right](x) \right\|_{x}.$$

$$\tag{42}$$

Note that $L_r(\omega) < \infty$ for each ω since \mathcal{X} is a bounded domain and φ_{ω} is smooth.

Since $|\varphi_{\omega}(x) - \varphi_{\omega}(x')| = \left|\int_{0}^{1} \frac{d}{dt}\varphi_{\omega}(\gamma(t))dt\right| = \left|\int_{0}^{1} D_{1}[\varphi_{\omega}](\gamma(t))[\dot{\gamma}(t)]dt\right|$ for a smooth path from x to x', it is easy to see that

$$|\varphi_{\omega}(x) - \varphi_{\omega}(x')| \leqslant L_1(\omega)\mathfrak{d}_{\mathfrak{g}}(x, x') \tag{43}$$

We will also require $D_2[\varphi_{\omega}](x)$ to be Lipschitz, to this end, we assume that for all $x, x' \in \mathcal{X}$, there exists $\tau_{x \to x'} : \mathbb{C}^d \to \mathbb{C}^d$ an isometric isomorphism with respect to \mathfrak{g}_x , that is, such that $\langle u, v \rangle_x = \langle \tau_{x \to x'} u, \tau_{x \to x'} v \rangle_{x'}$, such that for all ω :

$$L_{3}(\omega) \stackrel{\text{\tiny def.}}{=} \inf \left\{ L > 0 \; ; \; \sup_{\mathfrak{d}_{\mathfrak{g}}(x,x') \leqslant r_{\text{near}}} \frac{\| \mathbf{D}_{2} \left[\varphi_{\omega} \right](x) - \mathbf{D}_{2} \left[\varphi_{\omega} \right](x') [\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot] \|_{x}}{\mathfrak{d}_{\mathfrak{g}}(x,x')} \leqslant L \right\} < \infty.$$

where naturally

$$\left\| \mathcal{D}_{2}\left[\varphi_{\omega}\right](x) - \mathcal{D}_{2}\left[\varphi_{\omega}\right](x')\left[\tau_{x \to x'}, \tau_{x \to x'}\right] \right\|_{x} = \sup_{\left\|u\right\|_{x} \leqslant 1, \left\|v\right\|_{x} \leqslant 1} \mathcal{D}_{2}\left[\varphi_{\omega}\right](x)\left[u, v\right] - \mathcal{D}_{2}\left[\varphi_{\omega}\right](x')\left[\tau_{x \to x'}u, \tau_{x \to x'}v\right]$$

and r_{near} comes from Assumption 1. One possible choice of $\tau_{x \to x'}$ is to choose the parallel transport along the unique geodesic connecting x and x'. Another possible choice is to simply choose $\tau_{x \to x'} : v \mapsto \mathfrak{g}_{x'}^{-\frac{1}{2}} \mathfrak{g}_x^{\frac{1}{2}} v$. The latter choice implies

$$\left\| \mathcal{D}_{2}\left[\varphi_{\omega}\right](x) - \mathcal{D}_{2}\left[\varphi_{\omega}\right](x')\left[\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot\right] \right\|_{x} = \left\| \mathfrak{g}_{x'}^{-\frac{1}{2}} \mathcal{H}\varphi_{\omega}(x') \mathfrak{g}_{x'}^{-\frac{1}{2}} - \mathfrak{g}_{x}^{-\frac{1}{2}} \mathcal{H}\varphi_{\omega}(x) \mathfrak{g}_{x}^{-\frac{1}{2}} \right\|.$$
(44)

which is a more convenient expression that we will use in the examples.

Finally, we let $F_r: [0,\infty) \to [0,1]$ be decaying tail functions such that

$$\mathbb{P}_{\omega}\left(L_r(\omega) > t\right) \leqslant F_r(t). \tag{45}$$

Our sampling complexity will depend on the decay of these tail distributions so that the derivatives of the selected random features are bounded with high probability. A similar idea of stochastic incoherence was exploited in [16] for deriving compressed sensing bounds.

5.2 Main result

Our main result is valid under the following assumption, which links the tail probabilities of the bounds on the feature functions and the final number of measurements m.

Assumption 2 (Assumption on the features and the sample complexity). For $\rho > 0$, suppose that $m \in \mathbb{N}$ and some constant $\{\bar{L}_i\}_{i=0}^3 \in \mathbb{R}^4_+$ are chosen such that

$$\sum_{j=0}^{3} F_j(\bar{L}_j) \leqslant \frac{\min(\bar{\varepsilon}_0, \bar{\varepsilon}_2, \rho)}{m} \quad and \quad \max_{j=0}^{3} \left(\bar{L}_j^2 \sum_{i=0}^{3} F_i(\bar{L}_i) + 6 \int_{\bar{L}_j}^{\infty} tF_j(t) \mathrm{d}t \right) \leqslant \frac{\min(\bar{\varepsilon}_0, \bar{\varepsilon}_2)}{m} \tag{46}$$

and

$$m \gtrsim s\left(C_1\log(s)\log\left(\frac{s}{\rho}\right) + C_2\log\left(\frac{(sN)^d}{\rho}\right)\right)$$
(47)

where $N \stackrel{\text{def.}}{=} \frac{\mathcal{R}_{\mathcal{X}}\bar{L}_{1}}{\bar{\varepsilon}_{0}} + \frac{r_{\text{near}}\bar{L}_{3}\bar{L}_{0} + \bar{L}_{2}}{\bar{\varepsilon}_{2}}, C_{1} \stackrel{\text{def.}}{=} (\bar{L}_{0}^{2} + \bar{L}_{1}^{2}) \sum_{r=0,2} \frac{B_{r}^{2}}{\bar{\varepsilon}_{r}^{2}}, and C_{2} \stackrel{\text{def.}}{=} \frac{B_{22}\bar{L}_{01}^{2}}{B_{2}^{2}} + \sum_{r=0,2} \left(\frac{\bar{L}_{r}^{2}}{\bar{\varepsilon}_{r}^{2}} + \frac{\bar{L}_{01}\bar{L}_{r}}{\bar{\varepsilon}_{r}}\right)$ with $\bar{L}_{ij} = \sqrt{\bar{L}_{i}^{2} + \bar{L}_{j}^{2}}.$

The constants \bar{L}_r play the role of "stochastic" Lipschitz constant: for r = 0, 1, 2, with high probability on ω_j , $D_r [\varphi_{\omega}](x)$ will be \bar{L}_r -bounded and \bar{L}_{r+1} -Lipschitz. The condition (46) ensures that this is true with probability $1 - \rho$, that is, with the same desired probability of failure. Then, the entire proof is done conditionally on these bounds to hold.

Note also that, generally, $\{\bar{L}_r\}$ depend on m, through (46). However, all our examples fall under two categories (see Sec. 2):

- (i) either $\|D_r[\varphi_{\omega}](x)\|_x$ is already uniformly bounded, in which case \bar{L}_r can be chosen independently of ρ and m, this is for instance the case of discrete Fourier sampling;
- (ii) or the $F_r(t)$ are exponentially decaying, in which case we can show that $\bar{L}_r = \mathcal{O}\left(\log\left(\frac{m}{\rho}\right)^p\right)$ for some p > 0, which only incurs additional logarithmic terms in the bound (47). This occurs in the case of sampling the Laplace transform or sampling the Fourier transform with respect to a Gaussian distribution.

We are now ready to state the detailed version of Theorem 1, which is the main result of this paper.

Theorem 3. Suppose that Assumptions 1 and 2 hold. Let y be as in (2) with $\min_{i\neq j} \mathfrak{d}_{\mathfrak{g}}(x_i, x_j) \geq \Delta$ and $||w|| \leq \delta$. Then, with probability at least $1 - \rho$, any solution $\hat{\mu}$ of $(\mathcal{P}_{\lambda}(y))$ with $\lambda \sim \frac{\delta}{\sqrt{s}}$ satisfies

$$\mathcal{T}^2_{\mathfrak{d}_{\mathfrak{g}}}\left(\left|\hat{\mu}\right|, \sum_{j=1}^s \hat{A}_i \delta_{x_j}\right) \leqslant e \quad and \quad \max_{i=1}^s \left|\hat{a}_i - a_i\right| \leqslant e$$

where $\hat{A}_i = |\hat{\mu}| \left(\mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_i; r_{\text{near}}) \right), \ \hat{a}_i = \hat{\mu} \left(\mathcal{B}_{\mathfrak{d}_{\mathfrak{g}}}(x_i; r_{\text{near}}) \right) \ and \ e \lesssim \frac{1}{\min(\bar{\varepsilon}_0, \bar{\varepsilon}_2)} \left(|\tilde{\mu}_0| \left(\mathcal{X} \right) + \delta \cdot \sqrt{s} \right).$

The next section is dedicated to the proof of Theorem 3 using an infinite-dimensional golfing scheme. Appendix A is dedicated to the proof of some technical Lemmas. Appendix B gathers all the concentration inequalities that we use in the golfing scheme, which are essentially many variants of Bernstein's inequality. Finally, Appendices C, D and E are dedicated to the computation of all the constants in Assumptions 1 and 2 for the examples described in Section 2, which can be quite verbose.

6 Proof of Theorem 3

The main step towards proving Theorem 3 is to prove the existence of a dual certificate satisfying the properties described in Proposition 1. More precisely, we are going to prove the following theorem.

Theorem 4. Suppose that Assumptions 1 and 2 hold. Let $\{x_j\}_{j=1}^s$ be such that $\min_{i\neq j} \mathfrak{d}_{\mathfrak{g}}(x_i, x_j) \geq \Delta$. Then, with probability at least $1 - \rho$, there exists $p \in \mathbb{C}^m$ with $\|p\| \lesssim \sqrt{s}$ such that $\hat{\eta} = \Phi^* p$ is $(\frac{\overline{\varepsilon}_0}{8}, \frac{3\overline{\varepsilon}_2}{8}, r_{\text{near}})$ -nondegenerate.

Outline of the proof. The construction of the non-degenerate certificate includes several intermediate steps. As usual in this type of proof, we will first prove these properties on a finite ε -net that covers \mathcal{X} , then extend them to the whole space by regularity. Here we work with several nets $\mathcal{G}_j^{\text{near}} \subset \mathcal{X}_j^{\text{near}}$ and $\mathcal{G}^{\text{far}} \subset \mathcal{X}^{\text{far}}$ whose precision will be adjusted later. The principle of the golfing scheme is to work with an "approximate" dual certificate η^{app} (which is actually not a dual certificate at all), then "correct" it to obtain the desired true certificate. In details, we will go through the following steps:

1. First, show that with probability at least $1 - \rho$, there is an approximate certificate $\eta^{\text{app}} \in \text{Im}(\Phi^*)$ such that, for some constant c_0 that will be adjusted later,

$$\begin{cases} \sum_{j=1}^{s} |\eta^{\text{app}}(x_{j}) - \text{sign}(a_{j})|^{2} + \|D_{1}[\eta^{\text{app}}](x_{j})\|_{x_{j}}^{2} \leqslant c_{0}^{2} \quad \text{for all } j = 1, \dots, s \\ |\eta^{\text{app}}(x)| \leqslant 1 - \frac{\bar{\varepsilon}_{0}}{4} \qquad \qquad \text{for all } x \in \mathcal{G}^{\text{far}} \\ \|\overline{\text{sign}(a_{j})}D_{2}[\eta^{\text{app}}](x) - K^{(02)}(x_{j}, x)\|_{x} \leqslant \frac{7\bar{\varepsilon}_{2}}{64} \qquad \qquad \text{for all } j = 1, \dots, s, x \in \mathcal{G}_{j}^{\text{near}} \end{cases}$$
(48)

In other words, we relax the condition $\eta(x_j) = \text{sign}(a_j)$, $\nabla \eta(x_j) = 0$, and replace it with the first equation above.

2. Second, correct the approximate certificate to obtain a function $\hat{\eta} \in \text{Im}(\Phi^*)$ such that:

$$\begin{cases} \hat{\eta}(x_j) = \operatorname{sign}(a_j) \quad \text{and} \quad \nabla \hat{\eta}(x_j) = 0 & \text{for all } j = 1, \dots, s \\ |\hat{\eta}(x)| \leqslant 1 - \frac{3\bar{\varepsilon}_0}{16} & \text{for all } x \in \mathcal{G}^{\operatorname{far}} \\ \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\hat{\eta}\right](x) - K^{(02)}(x_j, x) \right\|_x \leqslant \frac{15\bar{\varepsilon}_2}{128} & \text{for all } j = 1, \dots, s, \, x \in \mathcal{G}_j^{\operatorname{near}} \end{cases}$$
(49)

That is, $\hat{\eta}$ satisfy all the properties we want, but on the finite nets $\mathcal{G}_{i}^{\text{far}}, \mathcal{G}_{i}^{\text{near}}$.

- 3. Third, bound the norm of the $p \in \mathbb{C}^m$ corresponding to $\hat{\eta} = \Phi^* p$.
- 4. Then, use Assumption 2 on the feature functions and the bound on ||p|| to show that actually, the $\hat{\eta}$ constructed above satisfy:

$$\begin{cases} \hat{\eta}(x_j) = \operatorname{sign}(a_j) \quad \text{and} \quad \nabla \hat{\eta}(x_j) = 0 & \text{for all } j = 1, \dots, s \\ |\hat{\eta}(x)| \leqslant 1 - \frac{\bar{\varepsilon}_0}{8} & \text{for all } x \in \mathcal{X}^{\operatorname{far}} \\ \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\hat{\eta}\right](x) - K^{(02)}(x_j, x) \right\|_x \leqslant \frac{\bar{\varepsilon}_2}{8} & \text{for all } j = 1, \dots, s, \ x \in \mathcal{X}_j^{\operatorname{near}} \end{cases}$$
(50)

which, by Lemma 2, will imply that $\hat{\eta}$ is non-degenerate with the desired constants and conclude the proof of Theorem 4.

¹Here we write $\hat{\eta}$ to distinguish from the "limit" certificate η that we built in the case $m \to \infty$.

5. In a fifth and final step, prove the existence of s additional certificates $\hat{\eta}_j$ as appear in Prop. 2. Combined with the existence of $\hat{\eta}$ and Prop. 1 and 2, it concludes the proof of Theorem 3.

We dedicate a subsection to each step of the proof. Before that, we start in the next subsection with some technical preliminaries and notations.

6.1 Preliminaries

Let us introduce some notations and show some technical bounds that will be handy. Recall the definitions of the sign vector \mathbf{u}_s from (29), γ , Υ and \mathbf{f} from (31), (30) and (32), and $D_{\mathfrak{g}}$ from (33). We have the following additional bounds, whose proof, in Appendix A.2, follows similar arguments to that of Theorem 2.

Lemma 3. Under Assumption 1, Υ and **f** defined as in (30) and (32) satisfy the following.

(i) Υ is invertible and satisfies

$$\|\operatorname{Id} - D_{\mathfrak{g}}\Upsilon D_{\mathfrak{g}}\|_{2} \leqslant \frac{1}{2} \quad and \quad \|\operatorname{Id} - D_{\mathfrak{g}}\Upsilon D_{\mathfrak{g}}\|_{\operatorname{Block}} \leqslant \frac{1}{2}.$$
 (51)

(ii) For any vector $q \in \mathbb{C}^{s(d+1)}$ and any $x \in \mathcal{X}^{\text{far}}$, we have

$$\|D_{\mathfrak{g}}\mathbf{f}(x)\|_{2} \leqslant B_{0} \quad and \quad \left|q^{\top}\mathbf{f}(x)\right| \leqslant B_{0} \left\|D_{\mathfrak{g}}^{-1}q\right\|_{\text{Block}}$$
(52)

(iii) For any vector $q \in \mathbb{C}^{s(d+1)}$ and any $x \in \mathcal{X}^{\text{near}}$ we have the bound:

$$\left\| \mathbf{D}_{2} \left[q^{\mathsf{T}} \mathbf{f}(.) \right] (x) \right\|_{x} \leq \left\| D_{\mathfrak{g}}^{-1} q \right\| B_{2} \quad and \quad \left\| \mathbf{D}_{2} \left[q^{\mathsf{T}} \mathbf{f}(.) \right] (x) \right\|_{x} \leq \left\| D_{\mathfrak{g}}^{-1} q \right\|_{\text{Block}} B_{2} \tag{53}$$

Now, for $\omega_1, \ldots, \omega_m$, denote the empirical versions of Υ and **f** by:

$$\hat{\Upsilon} \stackrel{\text{\tiny def.}}{=} \frac{1}{m} \sum_{k=1}^{m} \gamma(\omega_k) \gamma(\omega_k)^* \quad \text{and} \quad \hat{\mathbf{f}}(x) \stackrel{\text{\tiny def.}}{=} \frac{1}{m} \sum_{k=1}^{m} \overline{\gamma(\omega_k)} \varphi_{\omega_k}(x).$$
(54)

Recall the definition of $L_j(\omega)$ and \bar{L}_j in Assumption 2. Let the event \bar{E} be defined by

$$\bar{E} \stackrel{\text{\tiny def.}}{=} \bigcap_{k=1}^{m} E_{\omega_k} \quad \text{where} \quad E_{\omega} \stackrel{\text{\tiny def.}}{=} \left\{ L_j(\omega) \leqslant \bar{L}_j \; ; \; j = 0, 1, 2, 3 \right\}.$$
(55)

Since by Assumption 2, eq. (46), we have $\mathbb{P}(\bar{E}^c) \leq \rho$, a nondegenerate dual certificate can be constructed with probability at least $(1 - \rho)^2 \geq 1 - 2\rho$ provided that, *conditional on event* \bar{E} , a nondegenerate dual certificate can be constructed with probability at least $1 - \rho$.

We therefore assume for the rest of this proof that event \overline{E} holds and establish the probability conditional on \overline{E} that a nondegenerate dual certificate exists. To control this probability, we will need to control the deviation of $\hat{\mathbf{f}}$ and $\hat{\Upsilon}$ from their conditional expectations $\mathbf{f}_{\overline{E}} = \mathbb{E}_{\overline{E}}[\hat{\mathbf{f}}]$ and $\Upsilon_{\overline{E}} \stackrel{\text{def.}}{=} \mathbb{E}_{\overline{E}}[\hat{\Upsilon}]$, where we denote $\mathbb{E}_{\overline{E}}[\cdot] \stackrel{\text{def.}}{=} \mathbb{E}[\cdot|\overline{E}]$. The following Lemma, proved in Appendix A.3, bounds the deviations between these.

Lemma 4. Under Assumption 1 and 2, we have:

(i)
$$\|D_{\mathfrak{g}}(\Upsilon - \Upsilon_{\bar{E}})D_{\mathfrak{g}}\|_{2} \leq 4 \frac{(s+1)\min(\bar{\varepsilon}_{0},\bar{\varepsilon}_{2})}{m} \text{ and } \|D_{\mathfrak{g}}(\Upsilon - \Upsilon_{\bar{E}})D_{\mathfrak{g}}\|_{\operatorname{Block}} \leq 8 \frac{(s+1)\min(\bar{\varepsilon}_{0},\bar{\varepsilon}_{2})}{m}$$

(ii) for all $x \in \mathcal{X}^{\operatorname{far}}, \|D_{\mathfrak{g}}(\mathbf{f}(x) - \mathbf{f}_{\bar{E}}(x))\|_{2} \leq \frac{(B_{0}+2\sqrt{s})\min(\bar{\varepsilon}_{0},\bar{\varepsilon}_{2})}{m}$

(iii) for all $x \in \mathcal{X}^{\text{near}}$, $\sup_{\|q\|_2 \leq 1} \left\| \mathbf{D}_2 \left[(\mathbf{f} - \mathbf{f}_{\bar{E}})^\top D_{\mathfrak{g}} q \right] (x) \right\|_x \leq \frac{(B_2 + 2\sqrt{s}) \min(\bar{\varepsilon}_0, \bar{\varepsilon}_2)}{m}$

6.2 Step 1: construction of an approximate certificate with the golfing scheme

The first step is to construct an approximate certificate η^{app} using the so-called "golfing scheme". The golfing scheme was introduced in [41] and successfully used in compressed sensing for instance in [16]. It can be intuitively explained as follows. Recall that the certificate constructed in Theorem 2 in the case $m \to \infty$ is of the form $\eta = (\Upsilon^{-1}\mathbf{u})^{\top}\mathbf{f}$. It is therefore natural to try to show directly that $\hat{\eta} \stackrel{\text{def.}}{=} (\hat{\Upsilon}^{-1}\mathbf{u})^{\top}\hat{\mathbf{f}}$ is also nondegenerate by bounding the variation between η and $\hat{\eta}$. This is the strategy adopted by Tang et al [54] and in our previous work [45]. However, as mentioned before, this proof technique requires the random signs assumption, otherwise a sub-optimal bound on m is obtained. To solve this, the golfing scheme starts by writing the following Neumann expansion: assuming that $\hat{\Upsilon}$ is invertible, we have

$$\hat{\boldsymbol{\gamma}} = (\hat{\boldsymbol{\Upsilon}}^{-1} \mathbf{u})^{\top} \hat{\mathbf{f}} = (\boldsymbol{\Upsilon}^{-1} (\hat{\boldsymbol{\Upsilon}} \boldsymbol{\Upsilon}^{-1})^{-1} \mathbf{u})^{\top} \hat{\mathbf{f}}$$
$$= \sum_{\ell=1}^{\infty} \left(\boldsymbol{\Upsilon}^{-1} \left(\mathrm{Id} - \hat{\boldsymbol{\Upsilon}} \boldsymbol{\Upsilon}^{-1} \right)^{\ell-1} \mathbf{u} \right)^{\top} \hat{\mathbf{f}} = \sum_{\ell=1}^{\infty} (\boldsymbol{\Upsilon}^{-1} q_{\ell-1})^{\top} \hat{\mathbf{f}}$$
(56)

where $q_{\ell} \stackrel{\text{def.}}{=} \left(\text{Id} - \hat{\Upsilon} \Upsilon^{-1} \right) q_{\ell-1}, q_0 \stackrel{\text{def.}}{=} \mathbf{u}$. By cutting the sum above to a finite number of terms, one effectively obtains an approximate certificate that must be later corrected. However, there is an additional difficulty in analysing the sum, which comes from the fact that for each summand, $\hat{\mathbf{f}}$ and $\Upsilon^{-1}q_{\ell-1}$ are random variables which are not mutually independent. The idea of [41, 16] is to *decouple* the random variables by partitioning the indices $\{1, \ldots, m\}$ into J disjoint blocks \mathcal{B}_{ℓ} of size m_{ℓ} with $\sum_{\ell=1}^{J} m_{\ell} = m$, for some J and m_{ℓ} that are adjusted below. Denote by $\hat{\Upsilon}_{\ell}$ and $\hat{\mathbf{f}}_{\ell}$ the empirical versions of Υ and \mathbf{f} over the m_{ℓ} random variables included in \mathcal{B}_{ℓ} , that is:

$$\hat{\Upsilon}_{\ell} \stackrel{\text{\tiny def.}}{=} \frac{1}{m_{\ell}} \sum_{k \in \mathcal{B}_{\ell}} \gamma(\omega_k) \gamma(\omega_k)^* \quad \text{and} \quad \hat{\mathbf{f}}_{\ell}(x) \stackrel{\text{\tiny def.}}{=} \frac{1}{m_{\ell}} \sum_{k \in \mathcal{B}_{\ell}} \overline{\gamma(\omega_k)} \varphi_{\omega_k}(x).$$

Then, instead of (56), we consider

$$\eta^{\mathrm{app}} = \sum_{\ell=1}^{J} (\Upsilon^{-1} q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell}$$

where $q_{\ell} \stackrel{\text{def.}}{=} \left(\text{Id} - \hat{\Upsilon}_{\ell} \Upsilon^{-1} \right) q_{\ell-1}, q_0 \stackrel{\text{def.}}{=} \mathbf{u}$. Note that this can be rewritten as:

$$q_{\ell} = \mathbf{u}_s - \sum_{p=1}^{\ell} \hat{\Upsilon}_p \Upsilon^{-1} q_{p-1}$$
(57)

Now, the idea is that one can control each term $q_{\ell-1}^{\top} \hat{\mathbf{f}}_{\ell}$ conditional on $q_{\ell-1}$ and for appropriate choices of the blocksizes m_{ℓ} , η^{app} can be shown to be approximately nondegenerate with high probability. Each additional term in the sum brings the certificate "closer" to its desired properties, hence the term "golfing" scheme.

Parameters and intermediate assumptions. We set the error c_0 that appears in (48) as

$$c_0 = C_0 \min\left(\frac{\bar{\varepsilon}_0}{B_0}, \frac{\bar{\varepsilon}_2}{B_2}, 1\right)$$

for some universal constant C_0 . We define the parameters of our golfing scheme as follows:

$$J = \lceil \log(s) \rceil + 2,$$

$$c_1 = c_2 = \frac{c_0}{4\sqrt{\log(s)}}$$
 and $\forall \ell = 3, \dots, J, \quad c_\ell = c_0,$

$$t_1 = 1 - \frac{\bar{\varepsilon}_0}{2} + \frac{\bar{\varepsilon}_0}{8} \quad t_2 = 4B_0\sqrt{\log(s)},$$
 and $\forall \ell = 3, \dots, J, \quad t_\ell = 4B_0\log(s),$

$$b_1 = \frac{3\bar{\varepsilon}_2}{32}, \quad b_2 = 4B_2\sqrt{\log(s)},$$
 and $\forall \ell = 3, \dots, J, \quad b_\ell = 4B_2\log(s).$

We now formulate an intermediate set of assumptions, and proceed to show that: first, they imply the desired properties on η^{app} , and second, they are valid with high probability. For $1 \leq \ell \leq J$, we define:

- (I_{ℓ}) $||D_{\mathfrak{g}}q_{\ell}||_{\text{Block}} \leq c_{\ell} ||D_{\mathfrak{g}}q_{\ell-1}||_{\text{Block}},$
- (II_{ℓ}) For all $x \in \mathcal{G}^{\text{far}}$, $\left| (\Upsilon^{-1}q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell}(x) \right| \leq t_{\ell} \|D_{\mathfrak{g}}q_{\ell-1}\|_{\text{Block}}$,

(III_{$$\ell$$}) If $\ell = 1$: for all $j = 1, ..., s$, $x \in \mathcal{G}_j^{\text{near}}$, $\left\|\overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[(\Upsilon^{-1}\mathbf{u}_s)^{\top} \hat{\mathbf{f}}_1\right](x) - K^{(02)}(x_j, x)\right\|_x \leqslant b_1$; and if $\ell \ge 2$: for all $x \in \mathcal{G}^{\text{near}}$, $\left\|\mathcal{D}_2\left[(\Upsilon^{-1}q_{\ell-1})^{\top} \hat{\mathbf{f}}_\ell\right](x)\right\|_x \leqslant b_\ell \left\|\mathcal{D}_{\mathfrak{g}}q_{\ell-1}\right\|_{\text{Block}}$.

Let us now assume that (I_{ℓ}) , (II_{ℓ}) and (III_{ℓ}) are true for all ℓ , and show that η^{app} satisfy the desired properties. We define $\Psi : \mathscr{C}(\mathcal{X}) \to \mathbb{C}^{s(d+1)}$ by

$$\Psi f \stackrel{\text{\tiny def.}}{=} \left[f(x_1), \dots, f(x_s), \nabla f(x_1)^\top, \dots, \nabla f(x_s)^\top \right]^\top.$$
(58)

In words, Ψ evaluates a function and its first derivative at the points $\{x_j\}_{j=1}^s$. Note that for any vector $v \in \mathbb{C}^{s(d+1)}$, by definition we have $\Psi(v^{\top} \hat{\mathbf{f}}_{\ell}) = \hat{\Upsilon}_{\ell} v$. Using this, we have

$$\begin{split} & \left\| \sum_{j=1}^{s} \left\| \eta^{\operatorname{app}}(x_{j}) - \operatorname{sign}(a_{j}) \right\|^{2} + \left\| \operatorname{D}_{1} \left[\eta^{\operatorname{app}} \right](x_{j}) \right\|_{x_{j}}^{2} \\ &= \left\| \operatorname{\mathbf{u}}_{s} - D_{\mathfrak{g}} \Psi \eta^{\operatorname{app}} \right\| \leqslant \sqrt{2s} \left\| D_{\mathfrak{g}} \left(\operatorname{\mathbf{u}}_{s} - \Psi \eta^{\operatorname{app}} \right) \right\|_{\operatorname{Block}} = \sqrt{2s} \left\| D_{\mathfrak{g}} \left(\operatorname{\mathbf{u}}_{s} - \Psi \left(\sum_{\ell=1}^{J} (\Upsilon^{-1} q_{\ell-1})^{\top} \widehat{\mathbf{f}}_{\ell} \right) \right) \right\|_{\operatorname{Block}} \\ &= \sqrt{2s} \left\| D_{\mathfrak{g}} \left(\operatorname{\mathbf{u}}_{s} - \sum_{\ell=1}^{J} \hat{\Upsilon}_{\ell} \Upsilon^{-1} q_{\ell-1} \right) \right\|_{\operatorname{Block}} \stackrel{(57)}{=} \sqrt{2s} \left\| D_{\mathfrak{g}} q_{J} \right\|_{\operatorname{Block}} \leqslant \sqrt{s} \prod_{\ell=1}^{J} c_{\ell} \stackrel{(1)}{\leqslant} \frac{\sqrt{2s} c_{0}^{J}}{16 \log(s)} \leqslant c_{0} \,, \end{split}$$

since by adjusting C_0 we can have $c_0 \leq \left(\frac{1}{\sqrt{6}}\right)^{\frac{1}{\log(3)-1}} \leq \left(\frac{1}{\sqrt{2s}}\right)^{\frac{1}{\log(s)-1}}$ where the last inequality is valid for all s and results from a simple function study. It proves the first part of (48). Next, for all $x \in \mathcal{G}^{\text{far}}$,

$$\begin{aligned} |\eta^{\mathrm{app}}(x)| &\leq \sum_{\ell=1}^{J} \left| (\Upsilon^{-1}q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell}(x) \right| \stackrel{(\mathrm{II})}{\leq} \sum_{\ell=1}^{J} t_{\ell} \left\| D_{\mathfrak{g}} q_{\ell-1} \right\|_{\mathrm{Block}} \stackrel{(\mathrm{I})}{\leq} \sum_{\ell=1}^{J} t_{\ell} \prod_{p=1}^{\ell-1} c_{p} \\ &\leq 1 - \frac{\bar{\varepsilon}_{0}}{2} + \frac{\bar{\varepsilon}_{0}}{8} + B_{0}c_{0} + \frac{B_{0}}{4} \sum_{\ell=2}^{J-1} c_{0}^{\ell} \leq 1 - \frac{\bar{\varepsilon}_{0}}{2} + \frac{\bar{\varepsilon}_{0}}{8} + B_{0}c_{0} + \frac{B_{0}c_{0}^{2}}{4(1-c_{0})} \leq 1 - \frac{\bar{\varepsilon}_{0}}{4}. \end{aligned}$$

since by our choice of c_0 and adjusting C_0 , $B_0c_0 + \frac{B_0c_0^2}{4(1-c_0)} \leq \frac{\overline{\varepsilon}_0}{8}$. Similarly, for all $x \in \mathcal{G}_j^{\text{near}}$,

$$\begin{split} \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\eta^{\operatorname{app}} \right](x) - K^{(02)}(x_j, x) \right\|_x \\ &\leqslant \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[(\Upsilon^{-1} \mathbf{u}_s)^\top \hat{\mathbf{f}}_1 \right](x) - K^{(02)}(x_j, x) \right\|_x + \sum_{\ell=1}^J \left\| \mathcal{D}_2\left[(\Upsilon^{-1} q_{\ell-1})^\top \hat{\mathbf{f}}_\ell \right](x) \right\|_x \\ &\leqslant \frac{3\bar{\varepsilon}_2}{32} + \sum_{\ell=2}^J b_\ell \prod_{p=1}^{\ell-1} c_p = \frac{3\bar{\varepsilon}_2}{32} + B_2 c_0 + \frac{B_2}{4} \sum_{\ell=2}^{J-1} c_0^\ell \leqslant \frac{3\bar{\varepsilon}_2}{32} + B_2 c_0 + \frac{B_2 c_0^2}{4(1-c_0)} \leqslant \frac{7\bar{\varepsilon}_2}{64} \end{split}$$

since similarly, $B_2 c_0 + \frac{B_2 c_0^2}{4(1-c_0)} \leq \frac{\bar{c}_2}{64}$. Hence (I_ℓ) , (II_ℓ) , (III_ℓ) indeed implies (48). Next we derive a condition on m under which they are true with probability $1 - \rho$ (conditional on event \bar{E}).

Probability of successful construction. Let us now prove that (I_{ℓ}) , (II_{ℓ}) and (III_{ℓ}) are indeed valid with the desired probability. Let $p_1(\ell)$, $p_2(\ell)$ and $p_3(\ell)$ be the probabilities conditional on event \bar{E} that $(I_{\ell}), (II_{\ell})$ and (III_{ℓ}) fail, respectively. By a union bound, our goal is to derive a bound on m such that $\sum_{k=1}^{3} \sum_{\ell=1}^{J} p_k(\ell) \leq \rho$. We do so by applying variants of Bernstein's concentration inequality, that are all detailed in Appendix B. As we mentioned before, a crucial construction of the golfing scheme is that, at each step, $q_{\ell-1}$ and \mathbf{f}_{ℓ} are mutually independent, such that we can reason conditionally on $q_{\ell-1}$ and treat it as a fixed vector when bounding the probabilities w.r.t. $\hat{\mathbf{f}}_{\ell}$ and $\hat{\Upsilon}_{\ell}$.

We define $\bar{q}_{\ell} \stackrel{\text{def.}}{=} D_{\mathfrak{g}}^{-1} \Upsilon^{-1} q_{\ell}$ for short. To bound $p_1(\ell)$, we first observe the recurrence relation $D_{\mathfrak{g}} q_{\ell} =$ $\begin{array}{l} D_{\mathfrak{g}}(\mathrm{Id} - \hat{\Upsilon}_{\ell} \Upsilon^{-1}) q_{\ell-1} &= D_{\mathfrak{g}}(\Upsilon - \hat{\Upsilon}_{\ell}) D_{\mathfrak{g}} \bar{q}_{\ell-1}. \quad \text{Moreover, by Lemma 3 we have } \left\| D_{\mathfrak{g}}^{-1} \Upsilon^{-1} D_{\mathfrak{g}}^{-1} \right\|_{\mathrm{Block}} \leqslant \\ \frac{1}{1 - \|D_{\mathfrak{g}} \Upsilon D_{\mathfrak{g}}\|_{\mathrm{Block}}} \leqslant 2, \text{ and therefore } \|D_{\mathfrak{g}} q_{\ell-1}\|_{\mathrm{Block}} \geqslant \frac{1}{\|D_{\mathfrak{g}}^{-1} \Upsilon^{-1} D_{\mathfrak{g}}^{-1}\|_{\mathrm{Block}}} \|\bar{q}_{\ell-1}\|_{\mathrm{Block}} \leqslant \frac{1}{2} \|\bar{q}_{\ell-1}\|_{\mathrm{Block}}. \text{ Finally,} \\ \text{by Lemma 4 and our assumptions we have in particular that } \|D_{\mathfrak{g}}(\Upsilon_{\bar{E}} - \Upsilon) D_{\mathfrak{g}}\|_{\mathrm{Block}} \leqslant \min_{\ell} c_{\ell}/4. \text{ Therefore,} \end{array}$

$$p_{1}(\ell) = \mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}q_{\ell}\right\|_{\operatorname{Block}} \geqslant c_{\ell} \left\|D_{\mathfrak{g}}q_{\ell-1}\right\|_{\operatorname{Block}}\right) \leqslant \mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon - \hat{\Upsilon}_{\ell})D_{\mathfrak{g}}\bar{q}_{\ell-1}\right\|_{\operatorname{Block}} \geqslant \frac{c_{\ell}}{2} \left\|\bar{q}_{\ell-1}\right\|_{\operatorname{Block}}\right)$$
$$\leqslant \mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon_{\bar{E}} - \hat{\Upsilon}_{\ell})D_{\mathfrak{g}}\bar{q}_{\ell-1}\right\|_{\operatorname{Block}} \geqslant \frac{c_{\ell}}{4} \left\|\bar{q}_{\ell-1}\right\|_{\operatorname{Block}}\right)$$

Finally, applying Lemma 14, for some ρ_{ℓ} that we adjust later we obtain that

$$\mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon}_{\ell})D_{\mathfrak{g}}\bar{q}_{\ell-1}\right\|_{\mathrm{Block}} \geqslant \frac{c_{\ell}}{4} \|\bar{q}_{\ell-1}\|_{\mathrm{Block}}\right) \leqslant \rho_{\ell}$$

if $m_{\ell} \gtrsim \frac{s \bar{L}_{01}^2}{c_{\ell}^2} \log\left(\frac{s}{\rho_{\ell}}\right)$. For $p_2(\ell)$, we have

$$\begin{aligned} \left| (\Upsilon^{-1}q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell}(x) \right| &= \left| (\bar{q}_{\ell-1})^{\top} D_{\mathfrak{g}} \hat{\mathbf{f}}_{\ell}(x) \right| \leqslant \left| (\bar{q}_{\ell-1})^{\top} D_{\mathfrak{g}} (\hat{\mathbf{f}}_{\ell}(x) - \mathbf{f}(x)) \right| + \left| (\bar{q}_{\ell-1})^{\top} D_{\mathfrak{g}} \mathbf{f}(x) \right| \\ &\leqslant \left| (\bar{q}_{\ell-1})^{\top} D_{\mathfrak{g}} (\hat{\mathbf{f}}_{\ell}(x) - \mathbf{f}(x)) \right| + \begin{cases} B_0 \| \bar{q}_{\ell-1} \|_{\text{Block}} & \ell \geqslant 2\\ 1 - \frac{\bar{\varepsilon}_0}{2} & \ell = 1 \end{cases} \end{aligned}$$

by Lemma 3 for the case $\ell \ge 2$ and Theorem 2 for the case $\ell = 1$. Hence,

$$p_{2}(\ell) = \mathbb{P}_{\bar{E}} \left(\exists x \in \mathcal{G}^{\mathrm{far}}, \ \left| (\Upsilon^{-1}q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell}(x) \right| > t_{\ell} \left\| D_{\mathfrak{g}} q_{\ell-1} \right\|_{\mathrm{Block}} \right)$$

$$\leq \mathbb{P}_{\bar{E}} \left(\exists x \in \mathcal{G}^{\mathrm{far}}, \ \left| (\Upsilon^{-1}q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell}(x) \right| > \frac{t_{\ell}}{2} \left\| \bar{q}_{\ell-1} \right\|_{\mathrm{Block}} \right)$$

$$\leq \mathbb{P}_{\bar{E}} \left(\exists x \in \mathcal{G}^{\mathrm{far}}, \ \left| (\bar{q}_{\ell-1})^{\top} D_{\mathfrak{g}} (\hat{\mathbf{f}}_{\ell}(x) - \mathbf{f}(x)) \right| > \tilde{t}_{\ell} \left\| \bar{q}_{\ell-1} \right\|_{\mathrm{Block}} \right) \quad \text{where} \quad \tilde{t}_{\ell} \stackrel{\mathrm{def.}}{=} \begin{cases} \left(\frac{t_{\ell}}{2} - B_{0} \right) & \ell \geqslant 2\\ \frac{\bar{\xi}_{0}}{16} & \ell = 1 \end{cases}$$

Since by Lemma 4 we have in particular

$$\left| \left(\bar{q}_{\ell-1} \right)^{\top} D_{\mathfrak{g}} (\mathbf{f}_{\bar{E}}(x) - \mathbf{f}(x)) \right| \leqslant \sqrt{2s} \left\| \bar{q}_{\ell-1} \right\|_{\mathrm{Block}} \left\| D_{\mathfrak{g}} (\mathbf{f}_{\bar{E}}(x) - \mathbf{f}(x)) \right\| \leqslant \frac{t_{\ell}}{2} \left\| \bar{q}_{\ell-1} \right\|_{\mathrm{Block}} \,,$$

by Lemma 8 and a union bound we have

$$p_2(\ell) \leqslant \mathbb{P}_{\bar{E}}\left(\exists x \in \mathcal{G}^{\mathrm{far}}, \ \left| (\bar{q}_{\ell-1})^\top D_{\mathfrak{g}}(\hat{\mathbf{f}}_{\ell}(x) - \mathbf{f}_{\bar{E}}(x)) \right| > \frac{\tilde{t}_{\ell}}{2} \left\| \bar{q}_{\ell-1} \right\|_{\mathrm{Block}} \right) \leqslant \rho_{\ell}$$

provided that $m_{\ell} \gtrsim s\left(\frac{\bar{L}_0^0}{\bar{t}_{\ell}^2} + \frac{\bar{L}_{01}\bar{L}_0}{\bar{t}_{\ell}}\right)\log\left(\frac{|\mathcal{G}^{\text{far}}|}{\rho_{\ell}}\right)$. For $p_3(\ell)$, fix j, for any $x \in \mathcal{G}_j^{\text{near}}$: in the case $\ell \ge 2$, by Lemma 3,

$$\begin{split} \left\| \mathbf{D}_{2} \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell} \right] (x) \right\|_{x} &\leq \left\| \left(\mathbf{D}_{2} \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^{\top} (\hat{\mathbf{f}}_{\ell} - \mathbf{f}) \right] (x) \right) \right\|_{x} + \left\| \mathbf{D}_{2} \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^{\top} \mathbf{f} \right] (x) \right\|_{x} \\ &\leq \left\| \mathbf{D}_{2} \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^{\top} (\hat{\mathbf{f}}_{\ell} - \mathbf{f}) \right] (x) \right\|_{x} + B_{2} \left\| \bar{q}_{\ell-1} \right\|_{\mathrm{Block}} \end{split}$$

and for $\ell = 1$, by Theorem 2,

$$\begin{split} \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2 \left[(D_{\mathfrak{g}} \bar{q}_0)^\top \hat{\mathbf{f}}_1 \right] (x) - K^{(02)}(x_j, x) \right\|_x \\ &\leqslant \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2 \left[(D_{\mathfrak{g}} \bar{q}_0)^\top \mathbf{f} \right] (x) - K^{(02)}(x_j, x) \right\|_x + \left\| \mathcal{D}_2 \left[(D_{\mathfrak{g}} \bar{q}_0)^\top (\hat{\mathbf{f}}_1 - \mathbf{f}) \right] (x) \right\|_x \\ &\leqslant \frac{\bar{\varepsilon}_2}{16} + \left\| \mathcal{D}_2 \left[(D_{\mathfrak{g}} \bar{q}_0)^\top (\hat{\mathbf{f}}_1 - \mathbf{f}) \right] (x) \right\|_x. \end{split}$$

Therefore, by the same computation as before,

$$p_3(\ell) \leqslant \mathbb{P}_{\bar{E}} \left(\exists x \in \mathcal{G}^{\text{near}}, \left\| \mathbf{D}_2 \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^\top (\hat{\mathbf{f}}_{\ell} - \mathbf{f}) \right] (x) \right\|_x > \tilde{b}_\ell \left\| \bar{q}_{\ell-1} \right\|_{\text{Block}} \right), \text{ where } \tilde{b}_\ell \stackrel{\text{def.}}{=} \begin{cases} \left(\frac{b_\ell}{2} - B_2 \right) & \ell \geqslant 2\\ \frac{\bar{c}_2}{64} & \ell = 1. \end{cases}$$

Again using Lemma 4 we bound $\left\| \mathbf{D}_2 \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^\top (\mathbf{f}_{\bar{E}} - \mathbf{f}) \right] (x) \right\|_x \leq \frac{\tilde{b}_\ell}{2} \left\| \bar{q}_{\ell-1} \right\|_{\text{Block}}$ and

$$p_{3}(\ell) \leq \mathbb{P}_{\bar{E}}\left(\exists x \in \mathcal{G}^{\text{near}}, \left\| \mathbf{D}_{2} \left[(D_{\mathfrak{g}} \bar{q}_{\ell-1})^{\top} (\hat{\mathbf{f}}_{\ell} - \mathbf{f}_{\bar{E}}) \right] (x) \right\|_{x} > \frac{\tilde{b}_{\ell}}{2} \left\| \bar{q}_{\ell-1} \right\|_{\text{Block}} \right) \leq \rho_{\ell}$$

by Lemma 10 and a union bound, provided that $m_{\ell} \gtrsim s \left(\frac{\bar{L}_2^2}{\bar{b}_{\ell}^2} + \frac{\bar{L}_2\bar{L}_{01}}{\bar{b}_{\ell}}\right) \log\left(\frac{|\mathcal{G}^{\text{near}}|}{\rho_{\ell}}\right)$. Choosing $\rho_1 = \rho_2 = \rho/9$ and $\rho_{\ell} = \rho/(9J)$ for $\ell \ge 3$, recalling that obviously $\bar{\varepsilon}_r \leqslant B_r$ for r = 1, 2 and denoting $N_0 = |\mathcal{G}^{\text{far}}|$ and $N_2 = |\mathcal{G}^{\text{near}}|$ for short, we have $\sum_{k=1}^3 \sum_{\ell=1}^J p_k(\ell) \leqslant \rho$ provided that

$$m_1 = m_2 \gtrsim s \sum_{r=0,2} \left(\bar{L}_{01}^2 \frac{B_r^2}{\bar{\varepsilon}_r^2} \log(s) \log\left(\frac{s}{\rho}\right) + \left(\frac{\bar{L}_r^2}{\bar{\varepsilon}_r^2} + \frac{\bar{L}_{01}\bar{L}_r}{\bar{\varepsilon}_r}\right) \log\left(\frac{N_r}{\rho}\right) \right)$$

and for $\ell \ge 3$,

$$m_{\ell} \gtrsim s \sum_{r=0,2} \left(\bar{L}_{01}^2 \frac{B_r^2}{\bar{\varepsilon}_r^2} \log\left(\frac{s\log(s)}{\rho}\right) + \left(\frac{\bar{L}_r^2}{B_r^2\log^2(s)} + \frac{\bar{L}_{01}\bar{L}_r}{B_r\log(s)}\right) \log\left(\frac{N_r\log(s)}{\rho}\right) \right)$$

Therefore, conditionally on \bar{E} , η^{app} can be constructed with probability at least $1 - \rho$ if $m \gtrsim m_1 + m_2 + Jm_3$, for which it is sufficient that

$$m \gtrsim s \sum_{r=0,2} \left(\bar{L}_{01}^2 \frac{B_r^2}{\bar{\varepsilon}_r^2} \log(s) \log\left(\frac{s}{\rho}\right) + \left(\frac{\bar{L}_r^2}{\bar{\varepsilon}_r^2} + \frac{\bar{L}_{01}\bar{L}_r}{\bar{\varepsilon}_r}\right) \log\left(\frac{N_r \log(s)}{\rho}\right) \right)$$
(59)

6.3 Step 2: correcting the approximate certificate

The second step of our proof is to "correct" the previously constructed approximate certificate η^{app} to obtain a certificate $\eta \in \text{Im}(\Phi^*)$ satisfying (49). Recalling the definition (58) of Ψ , let $e \stackrel{\text{def.}}{=} \Psi \eta^{\text{app}} - \mathbf{u}_s$ be the error made by η^{app} and define

$$\hat{\eta} \stackrel{\text{\tiny def.}}{=} \eta^{\text{app}} - \eta^{\text{e}}, \text{ where } \eta^{\text{e}} \stackrel{\text{\tiny def.}}{=} (\hat{\Upsilon}^{-1}e)^{\top} \hat{\mathbf{f}}.$$

Then,

$$\Psi \hat{\eta} = \Psi \eta^{\rm app} - e = \mathbf{u}_s \,,$$

and we have indeed that $\hat{\eta}(x_i) = \operatorname{sign}(a_i)$ and $\nabla \hat{\eta}(x_i) = 0$. We will now bound the deviations of $\hat{\eta}$ on the grids \mathcal{G}^{far} and $\mathcal{G}^{\text{near}}$, using the fact that e has a small norm. Note that there is a subtlety here: e itself is random, and not independent of $\hat{\mathbf{f}}$ or $\hat{\Upsilon}$. So we must use "uniform" concentration bounds.

Using Lemma 3 in combination with Lemma 4 and Lemma 12, we have that with probability at least $1 - \rho$:

$$\left\| \operatorname{Id} - D_{\mathfrak{g}} \hat{\Upsilon} D_{\mathfrak{g}} \right\| \leq \left\| \operatorname{Id} - D_{\mathfrak{g}} \Upsilon D_{\mathfrak{g}} \right\| + \left\| D_{\mathfrak{g}} (\Upsilon - \Upsilon_{\bar{E}}) D_{\mathfrak{g}} \right\| + \left\| D_{\mathfrak{g}} (\Upsilon_{\bar{E}} - \hat{\Upsilon}) D_{\mathfrak{g}} \right\| \leq \frac{1}{2} + \frac{1}{8} + \frac{1}{8} = \frac{3}{4}$$
(60)

and therefore

$$\left\| D_{\mathfrak{g}}^{-1} \hat{\Upsilon}^{-1} D_{\mathfrak{g}}^{-1} \right\| \leqslant 4.$$
(61)

By Lemma 3, 4, 9 and a union bound to respectively bound each term in the following triangular inequality, with probability $1 - \rho$ we have

$$\forall x \in \mathcal{G}^{\text{far}}, \quad \left\| D_{\mathfrak{g}} \hat{\mathbf{f}}(x) \right\| \leq \left\| D_{\mathfrak{g}} \mathbf{f}(x) \right\| + \left\| D_{\mathfrak{g}} (\mathbf{f}_{\bar{E}}(x) - \mathbf{f}(x)) \right\| + \left\| D_{\mathfrak{g}} (\hat{\mathbf{f}}(x) - \mathbf{f}_{\bar{E}}(x)) \right\| \leq 2B_0$$

if $m \gtrsim B_0^{-2} \log\left(\frac{|\mathcal{G}^{\text{far}}|}{\rho}\right) (s\bar{L}_{01}^2 + \sqrt{s}\bar{L}_{01}\bar{L}_0)$. Then, for all $x \in \mathcal{G}^{\text{far}}$, since by adjusting C_0 we can have in particular $\|D_{\mathfrak{g}}e\| \leq \frac{c_0}{\sqrt{s}} \leq c_0 \leq \frac{1}{128} \min\left(\frac{\bar{\varepsilon}_2}{B_2}, \frac{\bar{\varepsilon}_0}{B_0}\right)$, we have

$$|\hat{\eta}(x)| \leqslant |\eta^{\mathrm{app}}(x)| + \left\| D_{\mathfrak{g}} \hat{\mathbf{f}}(x) \right\| \left\| D_{\mathfrak{g}}^{-1} \hat{\Upsilon}^{-1} D_{\mathfrak{g}}^{-1} \right\| \left\| D_{\mathfrak{g}} e \right\| \leqslant 1 - \frac{3\bar{\varepsilon}_0}{16},$$

Similarly, by Lemma 3, 4, with probability $1 - \rho$ we have for all $x \in \mathcal{G}^{\text{near}}$ and $q \in \mathbb{C}^{s(d+1)}$,

$$\begin{split} \left\| \mathbf{D}_{2} \left[\hat{\mathbf{f}}^{\top} D_{\mathfrak{g}} q \right](x) \right\|_{x} &\leq \left\| \mathbf{D}_{2} \left[\mathbf{f}^{\top} D_{\mathfrak{g}} q \right](x) \right\|_{x} + \left\| \mathbf{D}_{2} \left[(\mathbf{f}_{\bar{E}} - \mathbf{f})^{\top} D_{\mathfrak{g}} q \right](x) \right\|_{x} + \left\| \mathbf{D}_{2} \left[(\mathbf{f}_{\bar{E}} - \hat{\mathbf{f}})^{\top} D_{\mathfrak{g}} q \right](x) \right\|_{x} \\ &\leq \left(B_{2} + B_{2}/2 \right) \left\| q \right\| + \left\| q \right\| \sup_{\|v\|_{x} \leq 1} \left\| \frac{1}{m} \sum_{k=1}^{m} D_{\mathfrak{g}} \overline{\gamma(\omega_{k})} g_{\omega_{k}}(v) - \mathbb{E}_{\bar{E}} D_{\mathfrak{g}} \overline{\gamma(\omega)} g_{\omega}(v) \right\| \end{split}$$

where $g_{\omega}(v) \stackrel{\text{\tiny def.}}{=} D_2[\varphi_{\omega}](x)[v,v]$. By Lemma 11 and and a union bound, for all $x \in \mathcal{G}^{\text{near}}$,

$$\sup_{\|v\|_x \leqslant 1} \left\| \frac{1}{m} \sum_{k=1}^m D_{\mathfrak{g}} \overline{\gamma(\omega_k)} g_{\omega_k}(v) - \mathbb{E}_{\bar{E}} D_{\mathfrak{g}} \overline{\gamma(\omega)} g_{\omega}(v) \right\| \leqslant B_2$$

if $m \gtrsim \frac{sB_{22}\bar{L}_{01}^2 + \sqrt{s}\bar{L}_{01}\bar{L}_2B_2}{B_2^2} \left(\log\left(\frac{|\mathcal{G}^{\text{near}}|}{\rho}\right) + d\log\left(\frac{s\bar{L}_{01}\bar{L}_2}{B_2}\right) \right)$. Using this property with $q \stackrel{\text{def.}}{=} D_{\mathfrak{g}}^{-1}\hat{\Upsilon}^{-1}e$ such that $\|q\| \leqslant 4c_0$, and by adjusting C_0 , we obtain: for all $x \in \mathcal{G}_j^{\text{near}}$,

$$\begin{split} \left|\overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\hat{\eta}\right](x) - K^{(02)}(x_j, x)\right\|_x &\leqslant \left\|\overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\eta^{\operatorname{app}}\right](x) - K^{(02)}(x_j, x)\right\|_x + \left\|\mathcal{D}_2\left[\hat{\mathbf{f}}^{\top} D_{\mathfrak{g}}q\right](x)\right\|_x \\ &\leqslant \frac{7\bar{\varepsilon}_2}{64} + \frac{\bar{\varepsilon}_2}{128} = \frac{15\bar{\varepsilon}_2}{128} \end{split}$$

which concludes the second step of our proof. By combining the bounds on m that we obtained with (59), after simplification we still obtain

$$m \gtrsim s \sum_{r=0,2} \left(\bar{L}_{01}^2 \frac{B_r^2}{\bar{\varepsilon}_r^2} \log(s) \log\left(\frac{s}{\rho}\right) + \left(\frac{\bar{L}_r^2}{\bar{\varepsilon}_r^2} + \frac{\bar{L}_{01}\bar{L}_r}{\bar{\varepsilon}_r} + \frac{B_{22}}{B_2^2}\bar{L}_{01}^2\right) \log\left(\frac{N_r'\log(s)}{\rho}\right) \right)$$
(62)

with $N'_0 = N_0 = \left| \mathcal{G}^{\text{far}} \right|$ but $N'_2 = \left| \mathcal{G}^{\text{near}} \right| + (s \bar{L}_{01} \bar{L}_2 / B_2)^d$.

6.4 Step 3: Bounding the norm ||p||

In this section we upper bound $\|p\|$ where $\Phi^* p = \hat{\eta}$, for the $\hat{\eta}$ that we have constructed in the previous section. We recall that $\Phi^* p = \frac{1}{\sqrt{m}} \sum_{k=1}^m p_k \varphi_{\omega_k}(\cdot)$, and

$$\eta^{\mathrm{app}} = \sum_{\ell=1}^{J} (\Upsilon^{-1} q_{\ell-1})^{\top} \hat{\mathbf{f}}_{\ell} = \frac{1}{\sqrt{m}} \sum_{\ell} \frac{\sqrt{m}}{m_{\ell}} \sum_{k \in \mathcal{B}_{\ell}} (\Upsilon^{-1} q_{\ell-1})^{\top} \overline{\gamma(\omega_k)} \varphi_{\omega_k} = \Phi^* p^{\mathrm{app}},$$

where $p^{\text{app}} \stackrel{\text{def.}}{=} (p_{\ell})_{\ell=1}^{J} \in \mathbb{C}^{m}$ and $p_{\ell} \stackrel{\text{def.}}{=} \frac{\sqrt{m}}{m_{\ell}} (\gamma(\omega_{k})^{*} \Upsilon^{-1} q_{\ell-1})_{k \in \mathcal{B}_{j}} \in \mathbb{C}^{m_{\ell}}$. So, $\|p^{\text{app}}\|^{2} = \sum_{\ell=1}^{J} \|p_{\ell}\|_{2}^{2}$. To upper bound this, for each $\ell = 1, \ldots, J$,

$$\begin{split} \frac{m_{\ell}}{m} \|p_{\ell}\|_{2}^{2} &= \frac{1}{m_{\ell}} \sum_{k \in \mathcal{B}_{\ell}} q_{\ell-1}^{*} \Upsilon^{-1} \gamma(\omega_{k}) \gamma(\omega_{k})^{*} \Upsilon^{-1} q_{\ell-1} = q_{\ell-1}^{*} \Upsilon^{-1} \hat{\Upsilon}_{\ell} \Upsilon^{-1} q_{\ell-1} \\ &= q_{\ell-1}^{*} \Upsilon^{-1} (\hat{\Upsilon}_{\ell} \Upsilon^{-1} - \operatorname{Id}) q_{\ell-1} + q_{\ell-1}^{*} \Upsilon^{-1} q_{\ell-1} = q_{\ell-1}^{*} \Upsilon^{-1} q_{\ell} + q_{\ell-1}^{*} \Upsilon^{-1} q_{\ell-1} \\ &\leq \|D_{\mathfrak{g}}^{-1} \Upsilon^{-1} D_{\mathfrak{g}}^{-1}\| \|D_{\mathfrak{g}} q_{\ell-1}\| \left(\|D_{\mathfrak{g}} q_{\ell-1}\| + \|D_{\mathfrak{g}} q_{\ell}\|\right) \\ &\leq 4s \|D_{\mathfrak{g}} q_{\ell-1}\|_{\operatorname{Block}} \left(\|D_{\mathfrak{g}} q_{\ell}\|_{\operatorname{Block}} + \|D_{\mathfrak{g}} q_{\ell-1}\|_{\operatorname{Block}}\right) \leq 4s \left(c_{\ell} + 1\right) \prod_{i=1}^{\ell-1} c_{i}^{2}. \end{split}$$

where we have used $\|D_{\mathfrak{g}}^{-1}\Upsilon^{-1}D_{\mathfrak{g}}^{-1}\| \leq 2$ by Lemma 3, $\|\cdot\| \leq \sqrt{2s} \|\cdot\|_{\text{Block}}$, and the computation that precedes for $\|D_{\mathfrak{g}}q_{\ell}\|_{\text{Block}}$. For $\ell = 1, 2 \frac{m}{m_{\ell}} = \mathcal{O}(1)$ and $\frac{m}{m_3} = \mathcal{O}(\log(s))$. Also, for $\ell \geq 3$,

$$(c_{\ell}+1)\prod_{i=1}^{\ell-1}c_i^2 = (1+c_0)\frac{c_0^{\ell-1}}{16\log(s)}$$

Therefore,

$$\|p^{\mathrm{app}}\|^2 \lesssim 4s \left(1 + \frac{c_0}{4\sqrt{\log(s)}} + \frac{c_0^2}{16\log(s)} + (1+c_0)\frac{c_0^2}{16(1-c_0)} \right) \lesssim s.$$

On the other hand, $\eta^{e} = \Phi^{*} p^{e}$ where $p^{e} = \left(\gamma(\omega_{k})^{*} \Upsilon^{-1} e\right)_{k=1}^{m}$. So,

$$\|p^{\mathbf{e}}\|^{2} = e^{*} \Upsilon^{-1} \hat{\Upsilon} \Upsilon^{-1} e \leqslant 8 \|D_{\mathfrak{g}} e\|^{2} \lesssim 1.$$

Therefore, $\hat{\eta} = \Phi^* p$ with $\|p\|^2 \lesssim s$.

6.5 Step 4: Nondegeneracy on the entire domain

We conclude by showing that the $\hat{\eta}$ constructed in the previous sections is indeed nondegenerate on the entire domain. For this we simply need to control the Lipschitz constants of $\hat{\eta}$ and its Hessian, which are in fact directly related to $\|p\|$. Let any $x \in \mathcal{X}^{\text{far}}$, and $x' \in \mathcal{G}^{\text{far}}$ be the point in the grid closest to it. Under \bar{E} , we have

$$\begin{aligned} |\hat{\eta}(x)| &\leq 1 - \frac{3\bar{\varepsilon}_0}{16} + |\hat{\eta}(x) - \hat{\eta}(x')| = 1 - \frac{3\bar{\varepsilon}_0}{16} + |(\Phi^*p)(x) - (\Phi^*p)(x')| \\ &\leq 1 - \frac{3\bar{\varepsilon}_0}{16} + \|p\| \sqrt{\frac{1}{m} \sum_{k=1}^m |\varphi_{\omega_k}(x) - \varphi_{\omega_k}(x')|^2} \leq 1 - \frac{3\bar{\varepsilon}_0}{16} + \bar{L}_1 \|p\| \,\mathfrak{d}_{\mathfrak{g}}(x, x') \end{aligned}$$

Hence we prove the first part of (50) by choosing \mathcal{G}^{far} such that $\mathfrak{d}_{\mathfrak{g}}(x, x') \leq \frac{\overline{\varepsilon}_0}{16\overline{L}_1 \|p\|}$, which results in

$$\left|\mathcal{G}^{\mathrm{far}}\right| = \left(\frac{C\mathcal{R}_{\mathcal{X}}\bar{L}_{1} \|p\|}{\bar{\varepsilon}_{0}}\right)^{d}$$

for an appropriate constant C.

Now, for any $x \in \mathcal{X}_j^{\text{near}}$, and $x' \in \mathcal{G}_j^{\text{near}}$ closest to it, we write

$$\begin{aligned} \left\| \overline{\operatorname{sign}(a_{j})} \mathcal{D}_{2}\left[\hat{\eta}\right](x) - K^{(02)}(x_{j}, x) \right\|_{x} &\leq \left\| \mathcal{D}_{2}\left[\hat{\eta}\right](x) - \mathcal{D}_{2}\left[\hat{\eta}\right](x') [\tau_{x \to x'}, \tau_{x \to x'}] \right\|_{x} \\ &+ \left\| \overline{\operatorname{sign}(a_{j})} \mathcal{D}_{2}\left[\hat{\eta}\right](x') [\tau_{x \to x'}, \tau_{x \to x'}] - K^{(02)}(x_{j}, x') [\tau_{x \to x'}, \tau_{x \to x'}] \right\|_{x} \\ &+ \left\| K^{(02)}(x_{j}, x') [\tau_{x \to x'}, \tau_{x \to x'}] - K^{(02)}(x_{j}, x) \right\|_{x} \end{aligned}$$
(63)

We bound each of these terms. For the first, under \bar{E} we have

$$\begin{aligned} \|\mathbf{D}_{2}\left[\hat{\eta}\right](x) - \mathbf{D}_{2}\left[\hat{\eta}\right](x')[\tau_{x \to x'}, \tau_{x \to x'}\cdot]\|_{x} \\ \leqslant \|p\| \sqrt{\frac{1}{m} \sum_{k=1}^{m} \|\mathbf{D}_{2}\left[\varphi_{\omega_{k}}\right](x) - \mathbf{D}_{2}\left[\varphi_{\omega_{k}}\right](x')[\tau_{x \to x'}, \tau_{x \to x'}\cdot]\|_{x}^{2}} \leqslant \bar{L}_{3} \|p\| \,\mathfrak{d}_{\mathfrak{g}}(x, x') \end{aligned}$$

For the second term in (63), we have

$$\left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\hat{\eta}\right](x') [\tau_{x \to x'}, \tau_{x \to x'}] - K^{(02)}(x_j, x') [\tau_{x \to x'}, \tau_{x \to x'}] \right\|_x$$
$$= \left\| \overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\hat{\eta}\right](x') - K^{(02)}(x_j, x') \right\|_{x'} \leqslant \frac{15\bar{\varepsilon}_2}{128}$$

from what we have proved in the previous section.

Finally, for the third term in (63) we naturally introduce $K_{\bar{E}}^{(ij)}$ defined as $K^{(ij)}$ in (25), but by replacing \mathbb{E} with the conditional $\mathbb{E}_{\bar{E}}$. From Lemma 4 the deviation between $K^{(02)}$ and $K_{\bar{E}}^{(02)}$ can be bounded by

$$\forall x \in \mathcal{X}^{\text{near}}, \ \left\| K_{\bar{E}}^{(02)}(x_j, x) - K^{(02)}(x_j, x) \right\|_x = \left\| \mathbf{D}_2 \left[(\mathbf{f}_{\bar{E}} - \mathbf{f})^\top D_{\mathfrak{g}} u_j \right](x) \right\|_x \leqslant \frac{\bar{\varepsilon}_2}{512}$$

where u_j is the jth canonical vector of $\mathbb{C}^{s(d+1)}$. Moreover, by Assumption 2 it is easy to see that

$$\left\| K_{\bar{E}}^{(02)}(x_j, x') [\tau_{x \to x'}, \tau_{x \to x'}] - K_{\bar{E}}^{(02)}(x_j, x) \right\|_x \leqslant \bar{L}_0 \bar{L}_3 \mathfrak{d}_{\mathfrak{g}}(x, x')$$

Hence by a triangular inequality we have

$$\begin{split} \left\| K^{(02)}(x_j, x') [\tau_{x \to x'}, \tau_{x \to x'}] - K^{(02)}(x_j, x) \right\|_x &\leq \left\| K^{(02)}(x_j, x') [\tau_{x \to x'}, \tau_{x \to x'}] - K^{(02)}_{\bar{E}}(x_j, x') [\tau_{x \to x'}, \tau_{x \to x'}] \right\|_x \\ &+ \left\| K^{(02)}_{\bar{E}}(x_j, x') [\tau_{x \to x'}, \tau_{x \to x'}] - K^{(02)}_{\bar{E}}(x_j, x) \right\|_x \\ &+ \left\| K^{(02)}_{\bar{E}}(x_j, x) - K^{(02)}(x_j, x) \right\|_x \leq \frac{\bar{\varepsilon}_2}{256} + \bar{L}_0 \bar{L}_3 \mathfrak{d}_{\mathfrak{g}}(x, x') \end{split}$$

Therefore, (63) becomes

$$\left\|\overline{\operatorname{sign}(a_j)} \mathcal{D}_2\left[\hat{\eta}\right](x) - K^{(02)}(x_j, x)\right\|_x \leqslant \bar{L}_3(\bar{L}_0 + \|p\|) \mathfrak{d}_{\mathfrak{g}}(x, x') + \frac{15\bar{\varepsilon}_2}{128} + \frac{\bar{\varepsilon}_2}{256}$$
(64)

We prove the desired property on $D_2[\hat{\eta}]$ by choosing $\mathfrak{d}_{\mathfrak{g}}(x, x') \leq \frac{\overline{\varepsilon}_2}{256\overline{L}_3(\overline{L}_0 + \|p\|)}$, which yields

$$\left|\mathcal{G}^{\text{near}}\right| = s \left|\mathcal{G}_{j}^{\text{near}}\right| = s \left(\frac{Cr_{\text{near}}\bar{L}_{0}\bar{L}_{3}\left\|p\right\|}{\bar{\varepsilon}_{2}}\right)^{d}$$

for an appropriate constant C. Gathering everything with (62), we finally obtain

$$m \gtrsim s \sum_{r=0,2} \left(\bar{L}_{01}^2 \frac{B_r^2}{\bar{\varepsilon}_r^2} \log(s) \log\left(\frac{s}{\rho}\right) + \left(\frac{\bar{L}_r^2}{\bar{\varepsilon}_r^2} + \frac{\bar{L}_{01}\bar{L}_r}{\bar{\varepsilon}_r}\right) \log\left(\frac{\bar{N}_r^d}{\rho}\right) \right)$$
(65)
$$\bar{v}_r = s(\tau_{\text{regr}}\bar{L}_0\bar{L}_r + \bar{L}_2)$$

with $\bar{N}_0 = \frac{s\mathcal{R}_{\mathcal{X}}\bar{L}_1}{\bar{\varepsilon}_0}, \ \bar{N}_2 = \frac{s(r_{\text{near}}\bar{L}_0\bar{L}_3 + \bar{L}_2)}{\bar{\varepsilon}_2}.$

6.6 Step 5: additional certificates

Nondegeneracy of $\hat{\eta}$ directly allows us to apply Proposition 1 to deduce stability away from the sparse support in the reconstructed measure. In order to apply Proposition 2, we need to construct an additional s certificates η_j , which are however "simpler" to construct since they need to interpolate a "sign vector" that has only one non-zero coordinate, and do not require the golfing scheme to do so.

For each j = 1, ..., s, let u_j be the vector of length s(d+1) whose j^{th} entry is one and all other entries are zero. Define the functions

$$\eta_j^+ \stackrel{\text{\tiny def.}}{=} \left(\Upsilon^{-1} \begin{pmatrix} 1_s \\ 0_{sd} \end{pmatrix}\right)^\top \mathbf{f}(x) \quad \text{and} \quad \eta_j^- \stackrel{\text{\tiny def.}}{=} \left(\Upsilon^{-1} \left(2u_j - \begin{pmatrix} 1_s \\ 0_{sd} \end{pmatrix}\right)\right)^\top \mathbf{f}(x),$$

and

$$\eta_j \stackrel{\text{\tiny def.}}{=} \frac{1}{2} (\eta_j^+ + \eta_j^-) = \left(\Upsilon^{-1} u_j\right)^\top \mathbf{f}(x).$$

By Theorem 2, η_j^+ and η_j^- are nondegenerate (limit) dual certificates with respect to signs 1_s and $-1_s + 2u_j$ respectively, and η_j satisfies, for all $\ell \neq j$:

$$\begin{aligned} \eta_{j}(x_{j}) &= 1, \quad \nabla \eta_{j}(x_{j}) = 0 \quad \text{and} \quad \eta_{j}(x_{\ell}) = 0, \quad \nabla \eta_{j}(x_{\ell}) = 0 \\ |\eta_{j}(x)| &\leq \frac{1}{2} \left(\left| \eta_{j}^{+}(x) \right| + \left| \eta_{j}^{-}(x) \right| \right) \leq 1 - \frac{\bar{\varepsilon}_{0}}{4}, \quad \forall x \in \mathcal{X}^{\text{far}} \\ \left\| \mathbf{D}_{2} \left[\eta_{j} \right](x) - K^{(02)}(x_{j}, x) \right\|_{x} &\leq \frac{\bar{\varepsilon}_{2}}{16}, \quad \forall x \in \mathcal{X}^{\text{near}}_{j} \\ \| \mathbf{D}_{2} \left[\eta_{j} \right](x) \|_{x} \\ &\leq \frac{1}{2} \left(\left\| \mathbf{D}_{2} \left[\eta_{j}^{+} \right](x) - K^{(02)}(x_{\ell}, x) \right\|_{x} + \left\| -\mathbf{D}_{2} \left[\eta_{j}^{-} \right](x) - K^{(02)}(x_{\ell}, x) \right\|_{x} \right) \leq \frac{\bar{\varepsilon}_{2}}{16}, \quad \forall x \in \mathcal{X}^{\text{near}}_{\ell} \end{aligned}$$
(66)

Thus, using Lemma 2 to translate the last two conditions into quadratic decay, we conclude that η_j satisfies the conditions of Proposition 2.

To conclude, we will show that

$$\hat{\eta}_j \stackrel{\text{\tiny def.}}{=} \left(\hat{\Upsilon}^{-1} u_j \right)^\top \hat{\mathbf{f}} \in \operatorname{Im} \Phi^*$$

does not deviate too much from η_j and satisfies the conditions of Proposition 2. Note that by construction, $\hat{\eta}_j(x_j) = 1$, $\hat{\eta}_j(x_\ell) = 0$ for all $\ell \neq j$, and $\nabla \hat{\eta}_j(x_\ell) = 0$ for all ℓ . It therefore remains to control the deviation of $\hat{\eta}_j$ from η_j on \mathcal{X}^{far} and $D_2[\hat{\eta}_j]$ from $D_2[\eta_j]$ on $\mathcal{X}^{\text{near}}$.

Proposition 3. Under Assumption 1 and 2, suppose that $\min_{i\neq j} \mathfrak{d}_{\mathfrak{g}}(x_i, x_j) \ge \Delta$. Then, with probability at least $1 - \rho$, for all $j = 1, \ldots, s$, there exists $\hat{\eta}_j = \Phi^* p_j$ where $\|p_j\| \le 4$ which satisfies, for all $\ell \neq j$:

$$\begin{aligned} \hat{\eta}_{j}(x_{j}) &= 1, \quad \nabla \hat{\eta}_{j}(x_{j}) = 0 \quad and \quad \hat{\eta}_{j}(x_{\ell}) = 0, \quad \nabla \hat{\eta}_{j}(x_{\ell}) = 0 \\ |\hat{\eta}_{j}(x)| &\leq 1 - \frac{\bar{\varepsilon}_{0}}{8}, \quad \forall x \in \mathcal{X}^{\text{far}} \\ \left\| \mathbf{D}_{2}\left[\eta_{j} \right](x) - K^{(02)}(x_{j}, x) \right\|_{x} &\leq \frac{\bar{\varepsilon}_{2}}{8}, \quad \forall x \in \mathcal{X}^{\text{near}}_{j}, \quad \left\| \mathbf{D}_{2}\left[\eta_{j} \right](x) \right\|_{x} \leq \frac{\bar{\varepsilon}_{2}}{8}, \quad \forall x \in \mathcal{X}^{\text{near}}_{\ell} \end{aligned}$$
(67)

The proof controls the deviation between $\hat{\eta}_j$ and η_j on a fine grid using Bernstein's concentration inequalities, and extend the bound to the entire domain using Lipschitz properties of $\hat{\eta}_j$. As we mentioned above, the proof of this result is conceptually simpler than the deviation bounds on η^{app} since $||u_j|| = 1$. We therefore defer its proof to Appendix B.5. Using Lemma 2, we have therefore constructed the additional certificates to apply Proposition 2 and conclude the proof of Theorem 3.

7 Conclusion and outlooks

In this paper, we have presented an unifying geometric view on the problem of sparse measures recovery from random measurements. This theoretical analysis highlights the key role played by the invariant Fisher metric to define a precise notion of Rayleigh limit in the case of possibly non-translation invariant measurement kernels. We analyzed several examples including Laplace measurements in imaging, and left partially open some other important examples such as one-hidden-layer neural networks. Analyzing the super-resolution regime (going below the Rayleigh limit) requires stringent assumptions, such as positivity of the measures. Beyond the 1-D case, this is still mostly an open question, and we refer to [46] for some partial results.

References

- [1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, Optimization Algorithms on Matrix Manifolds, 2014.
- [2] M. AKKOUCHI, On the convolution of exponential distributions, Journal of the Chungcheong Mathematical Society, 21 (2008), pp. 501–510.
- [3] S.-I. AMARI AND H. NAGAOKA, Methods of information geometry, vol. 191, American Mathematical Soc., 2007.
- [4] J.-M. AZAIS, Y. DE CASTRO, AND F. GAMBOA, Spike detection from inaccurate samplings, Applied and Computational Harmonic Analysis, 38 (2015), pp. 177–195.
- [5] F. BACH, Breaking the curse of dimensionality with convex neural networks, Journal of Machine Learning Research, 18 (2017), pp. 1–53.
- [6] T. BENDORY AND Y. C. ELDAR, Recovery of sparse positive signals on the sphere from low resolution measurements, IEEE Signal Processing Letters, 22 (2015), pp. 2383–2386.
- [7] A. BEURLING, Sur les intégrales de fourier absolument convergentes et leur application à une transformation fonctionelle, in Ninth Scandinavian Mathematical Congress, 1938, pp. 345–366.
- [8] N. BOYD, G. SCHIEBINGER, AND B. RECHT, The alternating descent conditional gradient method for sparse inverse problems, SIAM Journal on Optimization, 27 (2017), pp. 616–639.
- K. BREDIES AND H. K. PIKKARAINEN, Inverse problems in spaces of measures, ESAIM: Control, Optimisation and Calculus of Variations, 19 (2013), pp. 190–218.
- [10] M. BURGER AND S. OSHER, Convergence rates of convex variational regularization, Inverse problems, 20 (2004), p. 1411.
- [11] C. J. BURGES, Geometry and invariance in kernel based methods, MIT Press, 1999.
- [12] L. CAFFARELLI AND R. J. MCCANN, Free boundaries in optimal transport and monge-ampere obstacle problems, Annals of mathematics, 171 (2010), pp. 673–730.
- [13] L. L. CAMPBELL, An extended čencov characterization of the information metric, Proceedings of the American Mathematical Society, 98 (1986), pp. 135–141.
- [14] E. J. CANDÈS AND C. FERNANDEZ-GRANDA, Super-resolution from noisy data, Journal of Fourier Analysis and Applications, 19 (2013), pp. 1229–1254.
- [15] —, Towards a mathematical theory of super-resolution, Communications on Pure and Applied Mathematics, 67 (2014), pp. 906–956.

- [16] E. J. CANDES AND Y. PLAN, A probabilistic and RIPless theory of compressed sensing, IEEE Transactions on Information Theory, 57 (2011), pp. 7235–7254.
- [17] E. J. CANDÈS, J. ROMBERG, AND T. TAO, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, IEEE Transactions on information theory, 52 (2006), pp. 489–509.
- [18] N. N. CENCOV, Statistical decision rules and optimal inference, no. 53, American Mathematical Soc., 1982.
- [19] S. S. CHEN, D. L. DONOHO, AND M. A. SAUNDERS, Atomic decomposition by basis pursuit, SIAM review, 43 (2001), pp. 129–159.
- [20] L. CHIZAT AND F. BACH, On the global convergence of gradient descent for over-parameterized models using optimal transport, in Advances in neural information processing systems, 2018, pp. 3036–3046.
- [21] L. CHIZAT, G. PEYRÉ, B. SCHMITZER, AND F.-X. VIALARD, Unbalanced optimal transport: Dynamic and kantorovich formulations, Journal of Functional Analysis, 274 (2018), pp. 3090–3123.
- [22] Y. CHO AND L. K. SAUL, Large-margin classification in infinite neural networks, Neural computation, 22 (2010), pp. 2678–2697.
- [23] —, Analysis and extension of arc-cosine kernels for large margin classification, arXiv preprint arXiv:1112.3712, (2011).
- [24] G. CORMODE, M. GAROFALAKIS, P. J. HAAS, AND C. JERMAINE, Synopses for Massive Data: Samples, Histograms, Wavelets, Sketches, Foundations and Trends in Databases, 4 (2011), pp. 1–294.
- [25] S. I. COSTA, S. A. SANTOS, AND J. E. STRAPASSON, Fisher information distance: a geometrical reading, Discrete Applied Mathematics, 197 (2015), pp. 59–69.
- [26] S. DASGUPTA AND A. GUPTA, An Elementary Proof of a Theorem of Johnson and Lindenstrauss, Random Structures and Algorithms, 22 (2003), pp. 60–65.
- [27] Y. DE CASTRO AND F. GAMBOA, Exact reconstruction using Beurling minimal extrapolation, Journal of Mathematical Analysis and applications, 395 (2012), pp. 336–354.
- [28] Y. DE CASTRO, F. GAMBOA, D. HENRION, AND J.-B. LASSERRE, Exact solutions to super resolution on semi-algebraic domains in higher dimensions, IEEE Transactions on Information Theory, 63 (2016), pp. 621–630.
- [29] Q. DENOYELLE, V. DUVAL, AND G. PEYRÉ, Support recovery for sparse super-resolution of positive measures, Journal of Fourier Analysis and Applications, 23 (2017), pp. 1153–1194.
- [30] Q. DENOYELLE, V. DUVAL, G. PEYRE, AND E. SOUBIES, *The Sliding Frank-Wolfe Algorithm and its Application to Super-Resolution Microscopy*, Inverse Problems, (2019).
- [31] D. L. DONOHO, Compressed sensing, IEEE Transactions on information theory, 52 (2006), pp. 1289– 1306.
- [32] V. DUVAL AND G. PEYRÉ, Exact support recovery for sparse spikes deconvolution, Foundations of Computational Mathematics, 15 (2015), pp. 1315–1355.
- [33] V. DUVAL AND G. PEYRÉ, Sparse spikes super-resolution on thin grids I: the LASSO, Inverse Problems, 33 (2017), p. 055008.
- [34] A. EFTEKHARI, J. TANNER, A. THOMPSON, B. TOADER, AND H. TYAGI, Sparse non-negative superresolution-simplified and stabilised, arXiv preprint arXiv:1804.01490, (2018).

- [35] C. EKANADHAM, D. TRANCHINA, AND E. P. SIMONCELLI, A unified framework and method for automatic neural spike identification, Journal of neuroscience methods, 222 (2014), pp. 47–55.
- [36] P. FACCHI, R. KULKARNI, V. I. MAN'KO, G. MARMO, E. C. SUDARSHAN, AND F. VENTRIGLIA, Classical and quantum Fisher information in the geometrical formulation of quantum mechanics, Physics Letters, Section A: General, Atomic and Solid State Physics, 374 (2010), pp. 4801–4803.
- [37] C. FERNANDEZ-GRANDA, Support detection in super-resolution, Proc. Proceedings of the 10th International Conference on Sampling Theory and Applications, (2013), pp. 145–148.
- [38] S. FOUCART AND H. RAUHUT, A mathematical introduction to compressive sensing, vol. 1, Birkhäuser Basel, 2013.
- [39] R. GRIBONVAL, G. BLANCHARD, N. KERIVEN, AND Y. TRAONMILIN, Compressive statistical learning with random feature moments, arXiv preprint arXiv:1706.07180, (2017).
- [40] D. GRIFFITHS, Introduction to Quantum Mechanics, Pearson Education, Inc., 2004.
- [41] D. GROSS, Recovering low-rank matrices from few coefficients in any basis, IEEE Transactions on Information Theory, 57 (2011), pp. 1548–1566.
- [42] W. LIAO AND A. FANNJIANG, MUSIC for single-snapshot spectral estimation: Stability and superresolution, Applied and Computational Harmonic Analysis, 40 (2016), pp. 33–67.
- [43] M. LIERO, A. MIELKE, AND G. SAVARÉ, Optimal entropy-transport problems and a new hellingerkantorovich distance between positive measures, Inventiones mathematicae, 211 (2018), pp. 969–1117.
- [44] S. MINSKER, On some extensions of bernstein's inequality for self-adjoint operators, Statistics & Probability Letters, 127 (2017), pp. 111–119.
- [45] C. POON, N. KERIVEN, AND G. PEYRÉ, Support localization and the fisher metric for off-the-grid sparse regularization, in Proc. AISTATS'19, 2019.
- [46] C. POON AND G. PEYRÉ, Multi-dimensional sparse super-resolution, SIAM Journal on Mathematical Analysis, 51 (2019), pp. 1–44.
- [47] C. R. RAO, Information and the accuracy attainable in the estimation of statistical parameters, Bull. Calcutta Math. Soc., 37 (1945), pp. 81–91.
- [48] R. ROY AND T. KAILATH, ESPRIT-estimation of signal parameters via rotational invariance techniques, IEEE Transactions on acoustics, speech, and signal processing, 37 (1989), pp. 984–995.
- [49] F. SANTAMBROGIO, Optimal transport for applied mathematicians, Birkäuser, NY, 55 (2015), pp. 58–63.
- [50] G. SCHIEBINGER, E. ROBEVA, AND B. RECHT, *Superresolution without separation*, arXiv preprint arXiv:1506.03144, (2015).
- [51] R. SCHMIDT, Multiple emitter location and signal parameter estimation, IEEE transactions on antennas and propagation, 34 (1986), pp. 276–280.
- [52] E. SOUBIES, L. BLANC-FÉRAUD, AND G. AUBERT, A continuous exact \ell_0 penalty (cel0) for least squares regularized problem, SIAM Journal on Imaging Sciences, 8 (2015), pp. 1607–1639.
- [53] G. TANG, B. N. BHASKAR, AND B. RECHT, Sparse recovery over continuous dictionaries-just discretize, in 2013 Asilomar Conference on Signals, Systems and Computers, IEEE, 2013, pp. 1043–1047.
- [54] G. TANG, B. N. BHASKAR, P. SHAH, AND B. RECHT, Compressed sensing off the grid, IEEE transactions on information theory, 59 (2013), pp. 7465–7490.

- [55] R. TIBSHIRANI, *Regression shrinkage and selection via the Lasso*, Journal of the Royal Statistical Society. Series B (Methodological), (1996), pp. 267–288.
- [56] J. A. TROPP, *Greed is good: Algorithmic results for sparse approximation*, IEEE Transactions on Information theory, 50 (2004), pp. 2231–2242.

A Preliminaries

In this Appendix, we provide the proofs to some technical lemmas in the paper, and give useful tools.

A.1 Linear algebra tools

We give the following simple lemma.

Lemma 5. For $1 \leq i, j \leq s$, take any scalars $a_{ij} \in \mathbb{C}$, vectors $Q_{ij}, R_{ij} \in \mathbb{C}^d$ and square matrices $A_{ij} \in \mathbb{C}^{d \times d}$.

- (i) For $q \in \mathbb{C}^{sd}$ and $M \in \mathbb{C}^{sd \times sd}$, we have $\|q\|_{\text{block}} \leq \|q\| \leq \sqrt{s} \|q\|_{\text{block}}$, and as a consequence $\|M\| \leq \sqrt{s} \|M\|_{\text{block}}$ and $\|M\|_{\text{block}} \leq \sqrt{s} \|M\|$. Similarly, for $q \in \mathbb{C}^{s(d+1)}$ and $M \in \mathbb{C}^{s(d+1) \times s(d+1)}$, we have $\|q\|_{\text{Block}} \leq \|q\| \leq \sqrt{2s} \|q\|_{\text{Block}}$, and as a consequence $\|M\| \leq \sqrt{2s} \|M\|_{\text{Block}}$ and $\|M\|_{\text{Block}} \leq \sqrt{2s} \|M\|$.
- (ii) Let $M \in \mathbb{C}^{sd \times sd}$ be a matrix formed by blocks :

$$M = \begin{pmatrix} A_{11} & \dots & A_{1s} \\ \vdots & \ddots & \vdots \\ A_{s1} & \dots & A_{ss} \end{pmatrix}$$

Then we have

$$\|M\|_{\text{block}} = \sup_{\|x\|_{\text{block}}=1} \|Mx\|_{\text{block}} \leqslant \max_{1 \leqslant i \leqslant s} \sum_{j=1}^{s} \|A_{ij}\|$$
(68)

Now, let $M \in \mathbb{R}^{sd \times s}$ be a rectangular matrix formed by stacking vectors $Q_{ij} \in \mathbb{R}^d$:

$$M = \begin{pmatrix} Q_{11} & \dots & Q_{1s} \\ \vdots & \ddots & \vdots \\ Q_{s1} & \dots & Q_{ss} \end{pmatrix}$$

Then,

$$\|M\|_{\infty \to \text{block}} \leqslant \max_{1 \leqslant i \leqslant s} \sum_{j=1}^{s} \|Q_{ij}\|_{2}, \quad \|M^{\top}\|_{\text{block} \to \infty} \leqslant \max_{1 \leqslant i \leqslant s} \sum_{j=1}^{s} \|Q_{ji}\|_{2}$$
(69)

(iii) Consider $M \in \mathbb{C}^{s(d+1) \times s(d+1)}$ decomposed as

$$M = \begin{pmatrix} a_{11} & \dots & a_{1s} & Q_{11}^{\top} & \dots & Q_{1s}^{\top} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ a_{s1} & \dots & a_{ss} & Q_{s1}^{\top} & \dots & Q_{ss}^{\top} \\ R_{11} & \dots & R_{1s} & A_{11} & \dots & A_{1s} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ R_{s1} & \dots & R_{ss} & A_{s1} & \dots & A_{ss} \end{pmatrix}$$

Then,

$$||M||^{2} \leq \max_{i} \left(\sum_{j=1}^{s} |a_{ij}| + ||Q_{ij}|| \right) \cdot \max_{j} \left(\sum_{i=1}^{s} |a_{ij}| + ||Q_{ij}|| \right) + \max_{i} \left(\sum_{j=1}^{s} ||R_{ij}|| + ||A_{ij}|| \right) \cdot \max_{j=1}^{s} \left(\sum_{i=1}^{s} ||R_{ij}|| + ||A_{ij}|| \right).$$

and

$$||M||_{\text{Block}} \leq \max_{i} \{\sum_{j} |a_{ij}| + ||Q_{ij}||, \sum_{j} ||R_{ij}|| + ||A_{ij}|| \}$$

Proof. The proof is simple linear algebra.

- (i) This is immediate by writing the definitions.
- (ii) Let x be a vector with $||x||_{\text{block}} \leq 1$ decomposed into blocks $x = [x_1, \ldots, x_s]$ with $x_i \in \mathbb{C}^d$, we have

$$\|Mx\|_{\text{block}} = \max_{1 \le i \le s} \left\| \sum_{j=1}^{s} A_{ij} x_j \right\| \le \max_{i} \sum_{j} \|A_{ij}\| \|x_j\| \le \max_{i} \sum_{j} \|A_{ij}\|$$

Similarly,

$$\left\|M^{\top}x\right\|_{\infty} = \max_{1 \leq i \leq s} \left\|\sum_{j=1}^{s} Q_{ji}^{\top}x_{j}\right\| \leq \max_{i} \sum_{j} \|Q_{ji}\| \|x_{j}\| \leq \max_{i} \sum_{j} \|Q_{ji}\|$$

Then, taking $x \in \mathbb{C}^s$ such that $||x||_{\infty} \leqslant 1$, we have

$$\|Mx\|_{\text{block}} = \max_{1 \le i \le s} \left\| \sum_{j=1}^{s} x_j Q_{ij} \right\| \le \max_i \sum_j \|Q_{ij}\|$$

(iii) Taking $x = [x_1, \ldots, x_s, X_1, \ldots, X_s] \in \mathbb{C}^{s(d+1)}$ with ||x|| = 1, we have

$$\|Mx\|^{2} = \sum_{i=1}^{s} \left(\sum_{j=1}^{s} a_{ij}x_{j} + Q_{ij}^{\mathsf{T}}X_{j} \right)^{2} + \left\| \sum_{j=1}^{s} R_{ij}x_{j} + A_{ij}X_{j} \right\|^{2}$$

$$\leq \sum_{i=1}^{s} \left(\sum_{j=1}^{s} |a_{ij}| x_{j}^{2} + ||Q_{ij}|| ||X_{j}||^{2} \right) \left(\sum_{j=1}^{s} |a_{ij}| + ||Q_{ij}|| \right)$$

$$+ \sum_{i=1}^{s} \left(\sum_{j=1}^{s} ||R_{ij}|| x_{j}^{2} + ||A_{ij}|| ||X_{j}||^{2} \right) \left(\sum_{j=1}^{s} ||R_{ij}|| + ||A_{ij}|| \right)$$

$$= \max_{i} \left(\sum_{j=1}^{s} |a_{ij}| + ||Q_{ij}|| \right) \cdot \max\left(\max_{j} \sum_{i=1}^{s} |a_{ij}|, \max_{j} \sum_{i=1}^{s} ||Q_{ij}|| \right) ||x||^{2}$$

$$+ \max_{i} \left(\sum_{j=1}^{s} ||R_{ij}|| + ||A_{ij}|| \right) \cdot \max\left(\max_{j=1}^{s} \sum_{i=1}^{s} ||R_{ij}||, \max_{j} \sum_{i=1}^{s} ||A_{ij}|| \right) ||x||^{2}.$$

Now, if $\|x\|_{\text{Block}} = 1$, we have

$$\|Mx\|_{\text{Block}} = \max_{i} \left(\left| \sum_{j=1}^{s} a_{ij} x_{j} + Q_{ij}^{\top} X_{j} \right|, \left\| \sum_{j=1}^{s} R_{ij} x_{j} + A_{ij} X_{j} \right\| \right)$$
$$\leq \max_{i} \left(\sum_{j=1}^{s} |a_{ij}| + \|Q_{ij}\|, \sum_{j=1}^{s} \|R_{ij}\| + \|A_{ij}\| \right)$$

r		
L		
L		

A.2 Proof of Lemma 3

The proof is similar to that of Theorem 2.

(i) We bound the spectral norm of Id $-D_{\mathfrak{g}}\Upsilon D_{\mathfrak{g}}$. By Lemma 5,

$$\begin{aligned} \| (\mathrm{Id} - D_{\mathfrak{g}} \Upsilon D_{\mathfrak{g}}) \|^{2} &\leq \max_{i} \left(\sum_{\substack{j=1\\j \neq i}}^{s} |K(x_{i}, x_{j})| + \sum_{j=1}^{s} \left\| K^{(10)}(x_{i}, x_{j}) \right\|_{x_{i}} \right)^{2} \\ &+ \max_{i} \left(\sum_{j=1,j \neq i}^{s} \left\| K^{(10)}(x_{j}, x_{i}) \right\|_{x_{j}} + \sum_{j=1}^{s} \left\| K^{(11)}(x_{i}, x_{j}) \right\|_{x_{i}, x_{j}} \right)^{2} \leq 8h^{2} \end{aligned}$$

by assumption on the kernel widths. Hence Υ is invertible. Similarly, by again applying Lemma 5, $\|D_{\mathfrak{g}}\Upsilon D_{\mathfrak{g}} - \mathrm{Id}\|_{\mathrm{Block}} \leq 2h.$

(ii) Let $x \in \mathcal{X}^{\text{far}}$, then we have

$$\|D_{\mathfrak{g}}\mathbf{f}(x)\| \leq \left(\sum_{i=1}^{s} |K(x_{i},x)|^{2} + \left\|K^{(10)}(x_{i},x)\right\|_{x_{i}}^{2}\right)^{\frac{1}{2}} \leq B_{00} + B_{10} + 2h \leq B_{0}$$

for which, similar to the proof above, we have used the fact that x is $\Delta/2$ -separated from at least s-1 points x_i . Similarly, for any vector $q = [q_1, \ldots, q_s, Q_1, \ldots, Q_s] \in \mathbb{C}^{s(d+1)}$ and any $x \in \mathcal{X}^{\text{far}}$, we have

$$\|q^{\top} \mathbf{f}(x)\| \leq \sum_{i=1}^{s} |q_{i}| |K(x_{i}, x)| + \|Q_{i}\|_{x_{i}} \|K^{(10)}(x_{i}, x)\|_{x_{i}}$$
$$\leq \|D_{\mathfrak{g}}^{-1}q\|_{\text{Block}} (B_{00} + B_{10} + 2h) \leq B_{0} \|D_{\mathfrak{g}}^{-1}q\|_{\text{Block}}$$

(iii) For any $x \in \mathcal{X}^{\text{near}}$ we have the bound:

$$\begin{aligned} \left\| \mathbf{D}_{2} \left[\boldsymbol{q}^{\top} \mathbf{f} \right] (\boldsymbol{x}) \right\|_{\boldsymbol{x}} &= \left\| \sum_{i=1}^{s} q_{i} K^{(02)}(\boldsymbol{x}_{i}, \boldsymbol{x}) + [Q_{i}] K^{(12)}(\boldsymbol{x}_{i}, \boldsymbol{x}) \right\|_{\boldsymbol{x}} \\ &\leqslant \left\| D_{\mathfrak{g}}^{-1} \boldsymbol{q} \right\| \left(\sum_{i=1}^{s} \left\| K^{(02)}(\boldsymbol{x}_{i}, \boldsymbol{x}) \right\|_{\boldsymbol{x}}^{2} + \left\| K^{(12)}(\boldsymbol{x}_{i}, \boldsymbol{x}) \right\|_{\boldsymbol{x}_{i}, \boldsymbol{x}}^{2} \right)^{\frac{1}{2}} \leqslant \left\| D_{\mathfrak{g}}^{-1} \boldsymbol{q} \right\| B_{2} \end{aligned}$$

and

$$\begin{split} \left\| \mathbf{D}_{2} \left[q^{\top} \mathbf{f} \right] (x) \right\|_{x} &= \left\| \sum_{i=1}^{s} q_{i} K^{(02)}(x_{i}, x) + [Q_{i}] K^{(12)}(x_{i}, x) \right\|_{x} \\ &\leqslant \left\| D_{\mathfrak{g}}^{-1} q \right\|_{\mathrm{Block}} \left(\sum_{i=1}^{s} \left\| K^{(02)}(x_{i}, x) \right\|_{x} + \left\| K^{(12)}(x_{i}, x) \right\|_{x_{i}, x} \right) \leqslant \left\| D_{\mathfrak{g}}^{-1} q \right\|_{\mathrm{Block}} B_{2} \end{split}$$

which concludes the proof.

A.3 Proof of Lemma 4

First note that, for $X = n^{-1} \sum_{k=1}^{m} f(\omega_k)$ any empirical average, since the ω_k are iid, we have $\mathbb{E}_{\bar{E}}[X] = \mathbb{E}_{E_{\omega}}[f(\omega)]$, and therefore $\mathbf{f}_{\bar{E}} = \mathbb{E}_{E_{\omega}}[\gamma(\omega)\gamma(\omega)^*]$, and so on.

We now prove a general bound, that we then implement for each item. Let $A = A_{\omega}$ be a random matrix that depends on ω , such that $||\mathbb{E}[A]|| \leq B$ and $||A|| \leq L(\omega)$, for any matrix norm $||\cdot||$. We have

$$\mathbb{E}[A] = \mathbb{E}[A1_{E_{\omega}}] + \mathbb{E}[A1_{E_{\omega}^{c}}] = \mathbb{E}_{\bar{E}}[A]\mathbb{P}(E_{\omega}) + \mathbb{E}[A1_{E_{\omega}^{c}}]$$

by Bayes' rule, and therefore,

$$\|\mathbb{E}[A] - \mathbb{E}_{\bar{E}}[A]\| \leqslant \frac{\|\mathbb{E}[A]\| \mathbb{P}(E_{\omega}^{c}) + \mathbb{E}[\|A\| \mathbf{1}_{E_{\omega}^{c}}]}{\mathbb{P}(E_{\omega})} \leqslant \frac{B\mathbb{P}(E_{\omega}^{c}) + \mathbb{E}[L(\omega)\mathbf{1}_{E_{\omega}^{c}}]}{\mathbb{P}(E_{\omega})}$$
(70)

Then, if we let $E_{\omega,q}$ be the event that $L_q(\omega) \leq \bar{L}_q$, so $E_\omega = \bigcap_{q=0}^3 E_{\omega,q}$, by the union bound we get $\mathbb{P}(E_\omega^c) \leq \sum_q \mathbb{P}(E_{\omega,q}^c) \leq \sum_q F_q(\bar{L}_q) \leq \frac{\min(\bar{\varepsilon}_0,\bar{\varepsilon}_2)}{m\max_j(\bar{L}_j^2)} \leq \frac{1}{2}$, and in particular $\mathbb{P}(E_\omega) \geq \frac{1}{2}$. In the following $L(\omega)$ will be a sum of some of the $L_q(\omega)^2$, so we bound $\mathbb{E}[L_q(\omega)^2 1_{E_\omega^c}] \leq \sum_j \mathbb{E}[L_q(\omega)^2 1_{E_{\omega,j}^c}]$ and we have

$$\begin{split} \mathbb{E}[L_q(\omega)^2 \mathbf{1}_{E_{\omega,j}^c}] &= \int_0^\infty \mathbb{P}(L_q(\omega)^2 \mathbf{1}_{E_{\omega,j}^c} \geqslant t) \mathrm{d}t = \int_0^\infty \mathbb{P}\left((L_q(\omega)^2 \geqslant t) \cap (L_j(\omega) \geqslant \bar{L}_j)\right) \mathrm{d}t \\ &\leqslant \bar{L}_q^2 F_j(\bar{L}_j) + \int_{\bar{L}_q^2}^\infty F_q(\sqrt{t}) \mathrm{d}t = \bar{L}_q^2 F_j(\bar{L}_j) + 2\int_{\bar{L}_q}^\infty t F_q(t) \mathrm{d}t \end{split}$$

where we have bounded $\mathbb{P}\left((L_q(\omega)^2 \ge t) \cap (L_j(\omega) \ge \bar{L}_j)\right)$ by respectively $\mathbb{P}(L_j(\omega) \ge \bar{L}_j) \le F_j(\bar{L}_j)$ in the first term and by $\mathbb{P}(L_q(\omega)^2 \ge t) \le F_q(\sqrt{t})$ in the second term. Hence by Assumption 2 we have

$$\mathbb{E}[L_q(\omega)^2 \mathbf{1}_{E_{\omega}^c}] \leqslant \frac{\min(\bar{\varepsilon}_0, \bar{\varepsilon}_2)}{m}$$
(71)

We can now obtain the desired results by combining (70) and (71) each time:

- (i) we let $A = D_{\mathfrak{g}}\gamma(\omega)\gamma(\omega)^*D_{\mathfrak{g}}$. We have $\|\mathbb{E}[A]\| \leq 2$ by Lemma 3, and $\|\gamma(\omega)\gamma(\omega)^*\| \leq sL_{01}^2(\omega)$. When applied with the norm $\|\cdot\|_{\text{Block}}$, we get $\|\mathbb{E}[A]\|_{\text{Block}} \leq 2$, and $\|\gamma(\omega)\gamma(\omega)^*\|_{\text{Block}} \leq 2sL_{01}^2(\omega)$ by Lemma 5 (*iii*).
- (ii) we let $A = D_{\mathfrak{g}}\gamma(\omega)\varphi_{\omega}(x)$. We have $\|\mathbb{E}[A]\| \leq B_0$ by Lemma 3, and $\|A\| \leq \sqrt{s}L_{01}(\omega)L_0(\omega) \leq \frac{1}{2}\sqrt{s}(L_{01}(\omega)^1 + L_0(\omega)^2)$.
- (iii) we let $A = (\tilde{\gamma}(\omega)^{\top}q)\mathfrak{g}_x^{-\frac{1}{2}}(\mathrm{H}\varphi_{\omega})(x)\mathfrak{g}_x^{-\frac{1}{2}}$. We have $\|\mathbb{E}[A]\| \leq B_2 \|q\|$ by Lemma 3, and $\|A\| \leq \sqrt{s}L_{01}(\omega)L_2(\omega)$.

B Concentration bounds

In this section, we detail the various Berstein concentration inequalities that we used in the golfing scheme. More precisely, we present some probabilistic bounds on deviation of $\hat{\mathbf{f}}$ and $\hat{\Upsilon}$ from their deterministic counterparts \mathbf{f} and $\hat{\Upsilon}$, conditional on event \bar{E} (recall their definitions in (30), (54) and (55)). Define the shorthands

$$L_{ij}(\omega) \stackrel{\text{def.}}{=} \sqrt{L_i(\omega)^2 + L_j(\omega)^2} \text{ and } \bar{L}_{ij} \stackrel{\text{def.}}{=} \sqrt{\bar{L}_i^2 + \bar{L}_j^2}.$$

Observe that conditional on \overline{E} , we have

$$\|D_{\mathfrak{g}}\gamma(\omega)\| \leqslant \sqrt{s(\bar{L}_0^2 + \bar{L}_1^2)} = \sqrt{s}\bar{L}_{01}$$

$$\tag{72}$$

All this section is done under the assumptions of Theorem 3, and we will use several times the following from Lemma 3 and 4:

$$\|D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}}\| \leq 1 + \|\mathrm{Id} - D_{\mathfrak{g}}\Upsilon D_{\mathfrak{g}}\| + \|D_{\mathfrak{g}}(\Upsilon - \Upsilon_{\bar{E}})D_{\mathfrak{g}}\| \leq 2$$
(73)

B.1 Elementary concentration inequalities

To begin, we first recall some elementary concentration inequalities.

Lemma 6 (Matrix Bernstein for complex matrices). Let Y_1, \ldots, Y_M be a sequence of $d_1 \times d_2$ complex random matrices with $\mathbb{E}[Y_\ell] = 0$, $\|Y_\ell\|_{2\to 2} \leq K$ for all $\ell = 1, \ldots, M$ and set

$$\sigma^{2} \stackrel{\text{\tiny def.}}{=} \max\left\{ \left\| \sum_{\ell=1}^{M} \mathbb{E}(Y_{\ell}Y_{\ell}^{*}) \right\|_{2 \to 2}, \left\| \sum_{\ell=1}^{M} \mathbb{E}(Y_{\ell}^{*}Y_{\ell}) \right\|_{2 \to 2} \right\}$$

Then,

$$\mathbb{P}\left(\left\|\frac{1}{M}\sum_{\ell=1}^{M}Y_{\ell}\right\|_{2\to 2} \ge t\right) \le 2(d_1+d_2)\exp\left(-\frac{Mt^2/2}{\sigma^2/M+Kt/3}\right).$$

Lemma 7 (Vector Bernstein for complex vectors [44]). Let $Y_1, \ldots, Y_M \in \mathbb{C}^d$ be a sequence of independent random vectors such that $\mathbb{E}[Y_i] = 0$, $||Y_i||_2 \leq K$ for $i = 1, \ldots, M$ and set

$$\sigma^2 \stackrel{\text{\tiny def.}}{=} \sum_{i=1}^M \mathbb{E} \left\| Y_i \right\|_2^2.$$

Then, for all $t \ge (2K + 6\sigma)/M$,

$$\mathbb{P}\left(\left\|\frac{1}{M}\sum_{i=1}^{M}Y_{i}\right\|_{2} \ge t\right) \le 28\exp\left(-\frac{Mt^{2}/2}{\sigma^{2}/M + tK/3}\right)$$

B.2 Deviation between $f_{\bar{E}}$ and \hat{f}

Lemma 8 (Bound against a fixed vector). Let $q \in \mathbb{C}^{s(d+1)}$ and $x \in \mathcal{X}$. For all u > 0 we have

$$\mathbb{P}_{\bar{E}}\left(\left|(\mathbf{f}_{\bar{E}}(x) - \hat{\mathbf{f}}(x))^{\top} D_{\mathfrak{g}}q\right| \ge u \, \|q\|\right) \le 4 \exp\left(\frac{-mu^2}{2\bar{L}_0^2 + 2\sqrt{s}\bar{L}_{01}\bar{L}_0 u/3)}\right).$$

As a corollary,

$$\mathbb{P}_{\bar{E}}\left(\left|(\mathbf{f}_{\bar{E}}(x) - \hat{\mathbf{f}}(x))^{\top} D_{\mathfrak{g}}q\right| \ge u \, \|q\|_{\mathrm{Block}}\right) \leqslant 4 \exp\left(\frac{-mu^2}{4s(2\bar{L}_0^2 + \sqrt{2}\bar{L}_{01}\bar{L}_0u/3)}\right).$$

Proof. Assume $\|q\|_2 = 1$ without lost of generality. We apply the classical (scalar) Bernstein inequality. By defining $Y_k \stackrel{\text{def.}}{=} \varphi_{\omega_k}(x)\gamma(\omega_k)^*D_{\mathfrak{g}}q - \mathbb{E}_E[\varphi_{\omega}(x)\gamma(\omega)^\top D_{\mathfrak{g}}q]$, we have $(\hat{\mathbf{f}}(x) - \mathbf{f}_{\bar{E}}(x))^\top D_{\mathfrak{g}}q = \frac{1}{m}\sum_{k=1}^m Y_k$. To apply Bernstein's inequality, observe that for each $k = 1, \ldots, m$, $\mathbb{E}_{\bar{E}}[Y_k] = 0$, and conditional on event \bar{E} , we have $|Y_k| \leq 2\sqrt{s}\bar{L}_{01}\bar{L}_0$ and $\mathbb{E}_E |Y_k|^2 = \mathbb{E}_E |\varphi_{\omega_k}(x)|^2 |\gamma(\omega_k)^* D_{\mathfrak{g}}q|^2 \leq \bar{L}_0^2 ||D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}}|| \leq 2\bar{L}_0^2$ by (73). Therefore,

$$\mathbb{P}\left(\left|\frac{1}{m}\sum_{k=1}^{m}Y_{k}\right| \ge u\right) \leqslant 4\exp\left(\frac{-mu^{2}}{2\bar{L}_{0}^{2}+2\sqrt{s}\bar{L}_{01}\bar{L}_{0}u/3}\right)$$

The last inequality follows because $\|q\|_{\text{Block}} \ge \|q\|_2 / \sqrt{2s}$.

Lemma 9 (Uniform bound). Fix $x \in \mathcal{X}$. For all $u > \frac{4\sqrt{s}\bar{L}_{01}\bar{L}_0}{m} + \frac{6\sqrt{s}\bar{L}_{01}}{\sqrt{m}}$ we have

$$\mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\mathbf{f}_{\bar{E}}(x)-\hat{\mathbf{f}}(x))\right\| \ge u\right) \leqslant 4\exp\left(\frac{-mu^2}{s\bar{L}_{01}^2+2\sqrt{s}\bar{L}_{01}\bar{L}_0u/3}\right).$$

		L
		L
		L
		L

Proof. We apply the vector Bernstein inequality (Lemma 7). By defining $Y_k \stackrel{\text{def.}}{=} D_{\mathfrak{g}} \overline{\gamma(\omega_k)} \varphi_{\omega_k}(x) - \mathbb{E}_E[D_{\mathfrak{g}} \overline{\gamma(\omega_k)} \varphi_{\omega_k}(x)]$, we have $D_{\mathfrak{g}}(\hat{\mathbf{f}}(x) - \mathbf{f}_{\bar{E}}(x)) = \frac{1}{m} \sum_{k=1}^{m} Y_k$. Observe that for each $k = 1, \ldots, m$, $\mathbb{E}_{\bar{E}}[Y_k] = 0$, and conditional on event \bar{E} , we have $|Y_k| \leq 2\sqrt{s}\bar{L}_{01}\bar{L}_0$ and $\mathbb{E}_E ||Y_k||^2 = \mathbb{E}_E |\varphi_{\omega_k}(x)|^2 ||D_{\mathfrak{g}}\gamma(\omega_k)||^2 \leq s\bar{L}_{01}^2$. Therefore, for all $u \geq \frac{4\sqrt{s}\bar{L}_{01}\bar{L}_0}{m} + \frac{6\sqrt{s}\bar{L}_{01}}{\sqrt{m}}$,

$$\mathbb{P}\left(\left\|\frac{1}{m}\sum_{k=1}^{m}Y_{k}\right\| \ge u\right) \leqslant 4\exp\left(\frac{-mu^{2}}{s\bar{L}_{01}^{2}+2\sqrt{s}\bar{L}_{01}\bar{L}_{0}u/3}\right).$$

The last inequality follows because $\|q\|_{\text{Block}} \ge \|q\|_2 / \sqrt{2s}$.

B.3 Deviation between $D_2[\mathbf{f}_{\bar{E}}]$ and $D_2[\hat{\mathbf{f}}]$

Lemma 10 (Bound against a fixed vector). Let $q \in \mathbb{C}^{s(d+1)}$ and $x \in \mathcal{X}$. For all u > 0 we have

$$\mathbb{P}_{\bar{E}}\left(\left\|\mathbf{D}_{2}\left[\left(\mathbf{f}_{\bar{E}}-\hat{\mathbf{f}}\right)^{\top}D_{\mathfrak{g}}q\right](x)\right\|_{x} \ge u \left\|q\right\|\right) \le 4d \exp\left(\frac{-mu^{2}}{2\bar{L}_{2}^{2}+2\sqrt{s}\bar{L}_{01}\bar{L}_{2}u/3)}\right).$$
(74)

as a corollary

$$\mathbb{P}_{\bar{E}}\left(\left\| \mathcal{D}_{2}\left[(\mathbf{f}_{\bar{E}} - \hat{\mathbf{f}})^{\top} D_{\mathfrak{g}} q \right](x) \right\|_{x} \ge u \left\| q \right\|_{\mathrm{Block}} \right) \le 4d \exp\left(\frac{-mu^{2}}{4s(2\bar{L}_{2}^{2} + \sqrt{2}\bar{L}_{01}\bar{L}_{2}u/3)}\right).$$
(75)

Proof. Assume ||q|| = 1 without lost of generality. Recalling the definitions of Sec. 4.1, we have

$$\left\| \mathbf{D}_2 \left[(\mathbf{f}_{\bar{E}} - \hat{\mathbf{f}})^\top D_{\mathfrak{g}} q \right] (x) \right\|_x = \left\| \mathbf{g}_x^{-\frac{1}{2}} \mathbf{H} \left((\mathbf{f}_{\bar{E}} - \hat{\mathbf{f}})^\top D_{\mathfrak{g}} q \right) (x) \mathbf{g}_x^{-\frac{1}{2}} \right\|$$

We now apply Lemma 6. Define

$$Y_{k} = (q^{\top} D_{\mathfrak{g}} \gamma(\omega_{k})) \mathfrak{g}_{x}^{-\frac{1}{2}} \mathrm{H}(\varphi_{\omega_{k}})(x) \mathfrak{g}_{x}^{-\frac{1}{2}} - \mathbb{E}_{\bar{E}}(q^{\top} D_{\mathfrak{g}} \gamma(\omega)) \mathfrak{g}_{x}^{-\frac{1}{2}} \mathrm{H}(\varphi_{\omega})(x) \mathfrak{g}_{x}^{-\frac{1}{2}}.$$

which are indeed symmetric matrices. We have $\mathbb{E}_E Y_k = 0$ and conditional on event E,

 $\|Y_k\| \leqslant 2\sqrt{s}\bar{L}_{01}\bar{L}_2.$

Furthermore, defining $A = (q^{\top}D_{\mathfrak{g}}\gamma(\omega_k))\mathfrak{g}_x^{-\frac{1}{2}}\mathrm{H}(\varphi_{\omega_k})(x)\mathfrak{g}_x^{-\frac{1}{2}}$ (which is symmetric), we have $0 \leq \mathbb{E}_{\bar{E}}[Y_jY_j^*] \leq \mathbb{E}_{\bar{E}}(AA^*) - \mathbb{E}_{\bar{E}}A\mathbb{E}_{\bar{E}}A^* \leq \mathbb{E}_{\bar{E}}(AA^*) \leq \bar{L}_2^2\mathbb{E}_{\bar{E}}|q^{\top}D_{\mathfrak{g}}\gamma(\omega)|^2\mathrm{Id} \leq \bar{L}_2^2||D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}}||\mathrm{Id} \leq 2\bar{L}_2^2\mathrm{Id}$ and thus $||\mathbb{E}_{\bar{E}}[Y_jY_j^*]|| \leq 2\bar{L}_2^2$. Therefore, the matrix Bernstein's inequality yields

$$\mathbb{P}\left(\left\|\frac{1}{m}\sum_{\ell=1}^{m}Y_{\ell}\right\|_{2} \ge u\right) \le 4d\exp\left(\frac{-mu^{2}}{2\bar{L}_{2}^{2}+2\sqrt{s}\bar{L}_{01}\bar{L}_{2}u/3}\right)$$

The last inequality follows because $\|q\|_{\text{Block}} \ge \|q\|_2 / \sqrt{2s}$.

Lemma 11 (Uniform bound). Let $x \in \mathcal{X}$. Let $\mathcal{B}_x \stackrel{\text{def.}}{=} \{v \in \mathbb{C}^d ; \|v\|_x \leq 1\}$ and given $v \in \mathcal{B}_x$, let $g_{\omega}(v) \stackrel{\text{def.}}{=} D_2[\varphi_{\omega}](x)[v,v] \in \mathbb{C}$. Then, for all $u > \frac{4\sqrt{s}\bar{L}_{01}\bar{L}_2}{m} + \frac{6\sqrt{2}\bar{L}_2^2}{\sqrt{m}}$,

$$\mathbb{P}\left(\sup_{v\in\mathcal{B}_x}\left\|\frac{1}{m}\sum_{k=1}^m D_{\mathfrak{g}}\overline{\gamma(\omega_k)}g_{\omega_k}(v) - \mathbb{E}_{\bar{E}}D_{\mathfrak{g}}\overline{\gamma(\omega)}g_{\omega}(v)\right\| \ge u\right) \le \exp\left(Cd\log\left(\frac{s\bar{L}_{01}\bar{L}_2}{u}\right) - \frac{mu^2}{s\bar{L}_{01}^2B_{22} + 2\sqrt{s}\bar{L}_{01}\bar{L}_2u/3)}\right)$$

for some constant C.

Proof. We use a covering net strategy: let $\mathcal{V} = \{\underline{v}_1, \ldots, v_N\}$ be a covering ε -net of \mathcal{B}_x , for $\varepsilon > 0$ that we will adjust later. Fix $v \in \mathcal{V}$, and define $Y_k = D_{\mathfrak{g}}\gamma(\omega_k)g_{\omega_k}(v) - \mathbb{E}_{\bar{E}}D_{\mathfrak{g}}\gamma(\omega_k)g_{\omega_k}(v) \in \mathbb{C}^{s(d+1)}$ centered *i.i.d.* variables. We have $\mathbb{E}_{\bar{E}}Y_k = 0$, $|Y_k| \leq 2\sqrt{s}\bar{L}_{01}\bar{L}_2$ and $\mathbb{E}_{\bar{E}} |Y_k||^2 \leq \mathbb{E}_{\bar{E}} |g_{\omega}(v)|^2 ||D_{\mathfrak{g}}\gamma(\omega)||^2 \leq s\bar{L}_{01}^2B_{22}$. Hence applying Lemma 7: for all $u \geq \frac{4\sqrt{s}\bar{L}_{01}\bar{L}_2}{m} + \frac{6\sqrt{B_{22}s}\bar{L}_{01}}{\sqrt{m}}$,

$$\mathbb{P}_{\bar{E}}\left(\left\|\frac{1}{m}\sum_{k=1}^{m}Y_{k}\right\| \ge u\right) \leqslant 4\exp\left(\frac{-mu^{2}}{s\bar{L}_{01}^{2}B_{22}+2\sqrt{s}\bar{L}_{01}\bar{L}_{2}u/3)}\right).$$

Next, we use the fact that for all ω

$$\left| D_{\mathfrak{g}} \overline{\gamma(\omega)} g_{\omega}(v) - D_{\mathfrak{g}} \overline{\gamma(\omega)} g_{\omega}(v') \right| \leq 2\sqrt{s} \overline{L}_{01} \overline{L}_2 \left\| v - v' \right\|_{\mathfrak{a}}$$

Hence by choosing

$$|\mathcal{V}| \sim \left(\frac{\sqrt{s}\bar{L}_{01}\bar{L}_2}{u}\right)^d$$

and using a union bound on $|\mathcal{V}|$ we conclude the proof.

B.4 Deviation between $\Upsilon_{\bar{E}}$ and $\hat{\Upsilon}$

Lemma 12 (Bound in spectral norm). For all u > 0, it holds that,

$$\mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon})D_{\mathfrak{g}}\right\| \ge u\right) \le 4(d+1)s\exp\left(-\frac{mu^2}{2s\bar{L}_{01}^2+2s\bar{L}_{01}^2u/3}\right).$$
(76)

Proof. To bound this probability, we apply Lemma 6 with $Y_k \stackrel{\text{def.}}{=} (D_g \gamma(\omega_k)) (D_g \gamma(\omega_k))^* - \Upsilon_{\bar{E}}$. We have, conditional on event E:

$$\mathbb{E}_{\bar{E}}[Y_k] = 0, \quad \|Y_k\| \stackrel{(72)}{\leqslant} 2s\bar{L}_{01}^2$$

Also,

$$0 \leq \mathbb{E}_{\bar{E}}[Y_k Y_k^*] = \mathbb{E}_{\bar{E}}[Y_k^* Y_k] = \mathbb{E}_{\bar{E}}[\|D_{\mathfrak{g}}\gamma(\omega_k)\|^2 (D_{\mathfrak{g}}\gamma(\omega_k))(D_{\mathfrak{g}}\gamma(\omega_k))^*] - (D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}})^2 \leq \mathbb{E}_{\bar{E}}[\|D_{\mathfrak{g}}\gamma(\omega_k)\|^2 (D_{\mathfrak{g}}\gamma(\omega_k))(D_{\mathfrak{g}}\gamma(\omega_k))^*] \leq s\bar{L}_{01}^2 \|D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}}\| \mathrm{Id}$$

So, $\|\mathbb{E}[Y_k^*Y_k]\| = \|\mathbb{E}[Y_kY_k^*]\| \leq s\bar{L}_{01}^2 \|D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}}\| \leq 2s\bar{L}_{01}^2$ by (73). By choosing $K = 2s\bar{L}_{01}^2$ and $\sigma^2 = ms\bar{L}_{01}^2 \|D_{\mathfrak{g}}\Upsilon_{\bar{E}}D_{\mathfrak{g}}\|$ in Lemma 6, we obtain

$$\mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon})D_{\mathfrak{g}}\right\| \ge u\right) \leqslant 4(d+1)s\exp\left(-\frac{mu^2}{2s\bar{L}_{01}^2+2s\bar{L}_{01}^2t/3}\right).$$

Lemma 13. For i = 1, ..., s, let $S_i = \{s + (i-1)d + 1, ..., s + id\}$, $q \in \mathbb{C}^{s(d+1)}$. Then, for all $u \ge \frac{4\sqrt{s}\bar{L}_{01}\bar{L}_1}{m} + \frac{6\sqrt{2}\bar{L}_1}{\sqrt{m}}$,

$$\mathbb{P}_{\bar{E}}\left(\left\| (D_{\mathfrak{g}}(\Upsilon_{\bar{E}} - \hat{\Upsilon})D_{\mathfrak{g}}q)_{S_{i}}\right\|_{2} > u \left\|q\right\|_{2}\right) \leqslant 28 \exp\left(\frac{-mu^{2}/2}{2\bar{L}_{1}^{2} + 2u\sqrt{s}\bar{L}_{01}\bar{L}_{1}/3}\right).$$

As a corollary, for all $u \ge \frac{4\sqrt{2}s\bar{L}_{01}\bar{L}_1}{m} + \frac{12\sqrt{s}\bar{L}_1}{\sqrt{m}}$, we have

$$\mathbb{P}_{\bar{E}}\left(\left\| (D_{\mathfrak{g}}(\Upsilon_{\bar{E}} - \hat{\Upsilon})D_{\mathfrak{g}}q)_{S_{i}}\right\|_{2} > u \left\|q\right\|_{\mathrm{Block}}\right) \leqslant 28 \exp\left(\frac{-mu^{2}}{4s\left(2\bar{L}_{1}^{2} + \sqrt{2}u\bar{L}_{01}\bar{L}_{1}/3\right)}\right)$$

Proof. Fix $i \in \{1, ..., s\}$. Without loss of generality, assume that $||q||_2 = 1$. The claim of this lemma follows by applying Lemma 7. Let

$$Y_k = \mathfrak{g}_{x_i}^{-\frac{1}{2}} \overline{\nabla \varphi_{\omega_k}(x_i)}(\gamma(\omega_k)^* D_{\mathfrak{g}}q) - \mathbb{E}_{\bar{E}}\left(\mathfrak{g}_{x_i}^{-\frac{1}{2}} \overline{\nabla \varphi_{\omega_k}(x_i)}(\gamma(\omega_k)^* D_{\mathfrak{g}}q)\right) \in \mathbb{C}^d,$$

and observe that $(D_{\mathfrak{g}}(\hat{\Upsilon} - \Upsilon_{\bar{E}})D_{\mathfrak{g}}q)_{S_i} = \frac{1}{m}\sum_k Y_k$. We apply Lemma 7. Observe that conditional on event \bar{E} , we have

$$||Y_k||_2 \leq 2 ||q||_2 ||D_{\mathfrak{g}}\gamma(\omega_k)||_2 ||D_1[\varphi_{\omega_k}](x_i)||_{x_i} \leq 2\sqrt{sL_{01}L_1}$$

and

$$\mathbb{E}_{\bar{E}} \left\| Y_k \right\|^2 \leqslant \mathbb{E}_{\bar{E}} \left| \gamma(\omega_k)^* D_{\mathfrak{g}} q \right|^2 \left\| \mathfrak{g}_{x_i}^{-\frac{1}{2}} \nabla \varphi_{\omega_k}(x_i) \right\|_2^2 \leqslant \bar{L}_1^2 q^* D_{\mathfrak{g}} \Upsilon_{\bar{E}} D_{\mathfrak{g}} q \leqslant \bar{L}_1^2 \left\| D_{\mathfrak{g}} \Upsilon_{\bar{E}} D_{\mathfrak{g}} \right\|_2 \leqslant 2\bar{L}_1^2$$

by (73). Therefore, for all $u \ge \frac{4\sqrt{s}\bar{L}_{01}\bar{L}_1}{m} + \frac{6\sqrt{2}\bar{L}_1}{\sqrt{m}}$,

$$\mathbb{P}_{\bar{E}}\left(\left\|\frac{1}{m}\sum_{i=1}^{m}Y_{i}\right\|_{2} \ge u\right) \le 28\exp\left(-\frac{mu^{2}/2}{2\bar{L}_{1}^{2}+2u\sqrt{s}\bar{L}_{01}\bar{L}_{1}/3}\right)$$

The last inequality follows because $\|q\|_{\text{Block}} \ge \|q\|_2 / \sqrt{2s}$.

Lemma 14 (Bound in block norm). Let $q \in \mathbb{C}^{s(d+1)}$ be any vector. For all $u \ge \frac{4\sqrt{2}s\overline{L}_{01}\overline{L}_{1}}{m} + \frac{12\sqrt{s}\overline{L}_{1}}{\sqrt{m}}$ we have

$$\mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon})D_{\mathfrak{g}}q\right\|_{\mathrm{Block}} \ge u \left\|q\right\|_{\mathrm{Block}}\right) \le 32s \exp\left(\frac{-mu^2}{4s\left(2\bar{L}_{01}^2+\sqrt{2}u\bar{L}_{01}\bar{L}_1/3\right)}\right)$$
(77)

Proof. Let $S_0 \stackrel{\text{def.}}{=} \{1, \ldots, s\}$ and $S_j \stackrel{\text{def.}}{=} \{s + (j-1)d + 1, \ldots, s + jd\}$ for $j = 1, \ldots, s$. By the union bound

$$\mathbb{P}_{\bar{E}}\left(\left\|D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon})D_{\mathfrak{g}}q\right\|_{\mathrm{Block}} \geqslant u \|q\|_{\mathrm{Block}}\right) \\
\leqslant \sum_{j=1}^{s} \mathbb{P}_{\bar{E}}\left(\left|(D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon})D_{\mathfrak{g}}q)_{j}\right| \geqslant u \|q\|_{\mathrm{Block}}\right) + \sum_{j=1}^{s} \mathbb{P}_{\bar{E}}\left(\left\|(D_{\mathfrak{g}}(\Upsilon_{\bar{E}}-\hat{\Upsilon})D_{\mathfrak{g}}q)_{S_{j}}\right\|_{2} \geqslant u \|q\|_{\mathrm{Block}}\right) \quad (78)$$

To bound the first sum, observe that for j = 1, ..., s, $(D_{\mathfrak{g}}(\Upsilon - \hat{\Upsilon})D_{\mathfrak{g}}q)_j = (D_{\mathfrak{g}}(\mathbf{f}(x_j) - \hat{\mathbf{f}}(x_j)))^{\top}q$ and apply Lemma 8. The second sum can be bounded by applying Lemma 13.

B.5 Proof of Proposition **3**

We fix a particular $j = 1, \ldots, s$, do the proof for $\hat{\eta}_j$, then use a union bound to conclude. As before, it is enough to establish the probability that $\hat{\eta}_j$ satisfies the properties of Proposition 3 conditional on event \bar{E} . We proceed in the same way as in the main proof of the golfing scheme: first we show that $\hat{\eta}_j$ satisfies the desired property on a finite grid, then we bound $||p_j||$, and finally we use the latter to extend the nondegeneracy to the whole space. As mentioned in the paper, the first step is considerably simpler and more direct than the golfing scheme, since the "sign" vector u_j is of norm 1.

Deviation bounds on a grid. Similar to our previous argument, we will bound the deviation between $\hat{\eta}_j$ and η_j on a finite grid $\mathcal{G}^{\text{far}} \subset \mathcal{X}^{\text{far}}$ whose precision we will later adjust, and between $D_2[\hat{\eta}_j]$ and $D_2[\eta_j]$ on $\mathcal{G}^{\text{near}} \subset \mathcal{X}^{\text{near}}$. We will show that

=

$$\begin{aligned} \forall x \in \mathcal{G}^{\text{far}}, \ |\eta_j(x) - \hat{\eta}_j(x)| &\leq \frac{\varepsilon_0}{16} \\ \forall x \in \mathcal{G}^{\text{near}}, \ \|D_2[\eta_j](x) - D_2[\hat{\eta}_j](x)\|_x &\leq \frac{\bar{\varepsilon}_2}{32} \end{aligned}$$

Let $\hat{q}_j \stackrel{\text{def.}}{=} D_{\mathfrak{g}}^{-1} \hat{\Upsilon}^{-1} u_j$ and $q_j \stackrel{\text{def.}}{=} D_{\mathfrak{g}}^{-1} \Upsilon^{-1} u_j$. Note that q_j is deterministic and $||q_j|| \leq 2$ for all j. Recall also that $\eta_j = q_j^\top D_{\mathfrak{g}} \mathbf{f}(x)$ and $\hat{\eta}_j = \hat{q}_j^\top D_{\mathfrak{g}} \hat{\mathbf{f}}(x)$. For $x \in \mathcal{G}^{\text{far}}$,

$$\begin{aligned} &|\eta_j(x) - \hat{\eta}_j(x)| \leqslant \left| q_j^\top D_{\mathfrak{g}}(\mathbf{f}(x) - \hat{\mathbf{f}}(x)) \right| + \left\| D_{\mathfrak{g}}^{-1} (\Upsilon^{-1} - \hat{\Upsilon}^{-1}) D_{\mathfrak{g}}^{-1} \right\| \left\| D_{\mathfrak{g}} \hat{\mathbf{f}}(x) \right\| \\ &\leqslant \left| q_j^\top D_{\mathfrak{g}}(\mathbf{f}(x) - \hat{\mathbf{f}}(x)) \right| + 8 \left\| D_{\mathfrak{g}}(\Upsilon - \hat{\Upsilon}) \right\| \left\| D_{\mathfrak{g}} \hat{\mathbf{f}}(x) \right\| \end{aligned}$$

where the last line is valid with probability $1 - \rho$ by Lemma 3 and (61). Similarly,

$$\begin{split} \|\mathbf{D}_{2}[\eta_{j}](x) - \mathbf{D}_{2}[\hat{\eta}_{j}](x)\|_{x} \\ &\leqslant \left\|\mathbf{D}_{2}\left[q_{j}^{\top}D_{\mathfrak{g}}(\mathbf{f}-\hat{\mathbf{f}})\right](x)\right\|_{x} + \left\|D_{\mathfrak{g}}^{-1}(\Upsilon^{-1}-\hat{\Upsilon}^{-1})D_{\mathfrak{g}}^{-1}\right\|\sup_{\|v\|_{x}\leqslant 1}\left\|\frac{1}{m}\sum_{k=1}^{m}D_{\mathfrak{g}}\overline{\gamma(\omega_{k})}\mathbf{D}_{2}[\varphi_{\omega_{k}}](x)[v,v]\right\| \\ &\leqslant \left\|\mathbf{D}_{2}\left[q_{j}^{\top}D_{\mathfrak{g}}(\mathbf{f}-\hat{\mathbf{f}})\right](x)\right\|_{x} + 8\left\|D_{\mathfrak{g}}(\Upsilon-\hat{\Upsilon})\right\|\sup_{\|v\|_{x}\leqslant 1}\left\|\frac{1}{m}\sum_{k=1}^{m}D_{\mathfrak{g}}\overline{\gamma(\omega_{k})}\mathbf{D}_{2}[\varphi_{\omega_{k}}](x)[v,v]\right\| \end{split}$$

where again, the last line is valid with probability $1 - \rho$ by Lemma 3 and (61). Therefore, we simply have to show that with probability at least $1 - \rho$,

(i) For
$$j = 1, ..., s$$
, $\left| q_j^\top D_{\mathfrak{g}}(\mathbf{f}_{\bar{E}}(x) - \hat{\mathbf{f}}(x)) \right| \leq \bar{\varepsilon}_0/32$ for all $x \in \mathcal{G}^{\text{far}}$.

(ii) For
$$j = 1, ..., s$$
, $\left\| \mathbb{D}_2 \left[q_j^\top D_{\mathfrak{g}}(\mathbf{f} - \hat{\mathbf{f}}) \right](x) \right\|_x \leq \bar{\varepsilon}_2/64$ for all $x \in \mathcal{G}^{\text{near}}$.

(iii)
$$\left\| D_{\mathfrak{g}} \hat{\mathbf{f}}(x) \right\| \leq 2B_0 \text{ for all } x \in \mathcal{G}^{\text{far}}.$$

(iv)
$$\sup_{\|v\|_{x} \leq 1} \left\| \frac{1}{m} \sum_{k=1}^{m} D_{\mathfrak{g}} \overline{\gamma(\omega_{k})} D_{2}[\varphi_{\omega_{k}}](x)[v,v] \right\| \leq 2B_{2} \text{ for all } x \in \mathcal{G}^{\text{near}}.$$

(v) $\left\| D_{\mathfrak{g}}(\Upsilon - \hat{\Upsilon}) D_{\mathfrak{g}} \right\| \leq \min\left(\frac{\bar{\varepsilon}_{0}}{512B_{0}}, \frac{\bar{\varepsilon}_{2}}{1024B_{2}}\right).$

By applying Lemma 4 and recalling our choice of m, (i) follows by Lemma 8, (ii) follows by Lemma 10, (iii) follows by Lemma 9, (iv) follows by Lemma 11, and (v) follows by Lemma 12.

Bound on p_j . By the same computations as in Section 6.5, we have $\hat{\eta}_j(x) = (\hat{\Upsilon}^{-1}u_j)^{\top}\hat{\mathbf{f}}(x) = \Phi^* p_j$ with $p_j = \frac{1}{\sqrt{m}} \left(\gamma(\omega_i)^* \hat{\Upsilon}^{-1}u_j\right)_{i=1}^m$. Therefore,

$$\begin{aligned} \|p_j\|_2^2 &= \frac{1}{m} \sum_{i=1}^m u_j^* \hat{\Upsilon}^{-1} \gamma(\omega_i) \gamma(\omega_i)^* \hat{\Upsilon}^{-1} u_j = u_j^* \hat{\Upsilon}^{-1} u_j \\ &\leqslant \left\| D_{\mathfrak{g}} \hat{\Upsilon}^{-1} D_{\mathfrak{g}}^{-1} \right\| \leqslant 4 \end{aligned}$$

with probability $1 - \rho$, by (61).

Extension to the whole domain. We proceed as in Section 6.4. By the same computations, we obtain: for any $x \in \mathcal{X}^{\text{far}}$ and $x' \in \mathcal{G}^{\text{far}}$,

$$|\hat{\eta}_j(x)| \leq |\eta_j(x')| + |\hat{\eta}_j(x') - \eta_j(x')| + |\hat{\eta}_j(x) - \hat{\eta}_j(x')| \leq 1 - \frac{\bar{\varepsilon}_0}{4} + \frac{\bar{\varepsilon}_0}{16} + \bar{L}_1 \|p_j\| \,\delta_{\mathfrak{g}}(x, x')$$

and therefore, we choose

$$\left|\mathcal{G}^{\mathrm{far}}\right| \sim \left(\frac{\mathcal{R}_{\mathcal{X}}\bar{L}_{1}}{\bar{\varepsilon}_{0}}\right)^{d}$$

For the second covariante derivative, as in Section 6.4 we get: for all $x \in \mathcal{X}_{i}^{\text{near}}$ and $x' \in \mathcal{G}_{i}^{\text{near}}$,

$$\begin{split} \left\| \overline{\operatorname{sign}(a_{j})} \mathcal{D}_{2}\left[\hat{\eta}_{j}\right](x) - K^{(02)}(x_{j}, x) \right\|_{x} \\ & \leqslant \left\| \mathcal{D}_{2}\left[\hat{\eta}_{j}\right](x) - \mathcal{D}_{2}\left[\hat{\eta}_{j}\right](x') [\tau_{x \to x'}, \tau_{x \to x'} \cdot] \right\|_{x} \\ & + \left\| \mathcal{D}_{2}\left[\hat{\eta}_{j}\right](x') [\tau_{x \to x'}, \tau_{x \to x'} \cdot] - \mathcal{D}_{2}\left[\eta_{j}\right](x') [\tau_{x \to x'}, \tau_{x \to x'} \cdot] \right\|_{x} \\ & + \left\| \overline{\operatorname{sign}(a_{j})} \mathcal{D}_{2}\left[\eta_{j}\right](x') [\tau_{x \to x'}, \tau_{x \to x'} \cdot] - K^{(02)}(x_{j}, x') [\tau_{x \to x'}, \tau_{x \to x'} \cdot] \right\|_{x} \\ & + \left\| K^{(02)}(x_{j}, x') [\tau_{x \to x'}, \tau_{x \to x'} \cdot] - K^{(02)}(x_{j}, x) \right\|_{x} \\ & \leqslant \frac{\tilde{\varepsilon}_{2}}{32} + \frac{\tilde{\varepsilon}_{2}}{16} + \frac{\tilde{\varepsilon}_{2}}{64} + \bar{L}_{3}(\bar{L}_{0} + \|p_{j}\|) \mathfrak{d}_{\mathfrak{g}}(x, x') \end{split}$$

and similarly for $\ell \neq j$, for all $x \in \mathcal{X}_{\ell}^{\text{near}}$ and $x' \in \mathcal{G}_{\ell}^{\text{near}}$,

$$\begin{split} \| \mathbf{D}_{2} \left[\hat{\eta}_{j} \right] (x) \|_{x} &\leq \| \mathbf{D}_{2} \left[\hat{\eta}_{j} \right] (x) - \mathbf{D}_{2} \left[\hat{\eta}_{j} \right] (x') [\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot] \|_{x} \\ &+ \| \mathbf{D}_{2} \left[\hat{\eta}_{j} \right] (x') [\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot] - \mathbf{D}_{2} \left[\eta_{j} \right] (x') [\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot] \|_{x} + \| \mathbf{D}_{2} \left[\eta_{j} \right] (x') [\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot] \|_{x} \\ &\leq \frac{\bar{\varepsilon}_{2}}{32} + \frac{\bar{\varepsilon}_{2}}{16} + \bar{L}_{3} \| p_{j} \| \, \mathfrak{d}_{\mathfrak{g}}(x, x') \end{split}$$

and therefore we conclude by setting

$$|\mathcal{G}^{\mathrm{near}}| \sim s \left(\frac{r_{\mathrm{near}} \bar{L}_3 \bar{L}_0}{\bar{\varepsilon}_2}\right)^d$$

The final bound on m is satisfied with the one we obtained previously (65).

C Application: Discrete Fourier sampling

In this section, we consider the case of sampling Fourier coefficients as described in [15]. Let $f \in \mathbb{N}$ and $\mathcal{X} \in \mathbb{T}^d$ the *d*-dimensional torus. Let $\Omega = \left\{ \omega \in \mathbb{Z}^d ; \|\omega\|_{\infty} \leq f \right\}$, $\varphi_{\omega}(x) \stackrel{\text{def.}}{=} e^{i2\pi\omega^\top x}$, and $\Lambda(\omega) = \prod_{j=1}^d g(\omega_j)$ where $g(j) = \frac{1}{f} \sum_{k=\max(j-f,-f)}^{\min(j+f,f)} (1 - |k/f|)(1 - |(j-k)/f|)$.

The kernel and Fisher metric The associated kernel is the multivariate Jackson kernel $K(x, x') = \prod_{i=1}^{d} \kappa(x_i - x'_i)$, where

$$\kappa(x) \stackrel{\text{\tiny def.}}{=} \left(\frac{\sin\left(\left(\frac{f}{2}+1\right)\pi x\right)}{\left(\frac{f}{2}+1\right)\sin(\pi x)} \right)^4$$

with constant metric tensor

$$\mathfrak{g}_x = C_f \mathrm{Id}$$
 and $\mathfrak{d}_\mathfrak{g}(x, x') = C_f^{\frac{1}{2}} \|x - x'\|_2$.

where $C_f \stackrel{\text{def.}}{=} -\kappa''(0) = \frac{\pi^2}{3}f(f+4) \sim f^2$. Note that $K^{(ij)} = \nabla_1^i \nabla_2^j K$ and $\left\|K^{(ij)}\right\|_{x,x'} = C_f^{-(i+j)/2} \left\|\nabla_1^i \nabla_2^j K\right\|$. Moreover, since the metric is constant, we have $\|\cdot\|_x = C_f^{\frac{1}{2}} \|\cdot\|$ for all x. The domain diameter is $\mathcal{R}_{\mathcal{X}} = C_f^{\frac{1}{2}} d^{\frac{1}{2}}$.

Sampling bounds Suppose that $f \ge 128$. The rest of this section consists of Lemmas which bound the parameters in Theorem 3: We show in Lemma 15 that by choosing $r_{\text{near}} = \frac{1}{8\sqrt{2}}$, for all $\mathfrak{d}_{\mathfrak{g}}(x, x') \le r_{\text{near}}$, we can set $\bar{\varepsilon}_2 = (1 - 6r_{\text{near}}^2)/(1 - r_{\text{near}}^2/(2 - r_{\text{near}}^2) - r_{\text{near}}^2) \ge 0.941$. In Lemma 16, we show that for all

 $\mathfrak{d}_{\mathfrak{g}}(x,x') \ge r_{\text{near}}, |K(x,x')| \le 1 - 1/(8^3 \cdot 2)$, so we can set $\bar{\varepsilon}_0 \stackrel{\text{def.}}{=} 0.00097$. Moreover, the uniform bounds given in Lemma 18 imply that

$$\frac{\min(\varepsilon_0, \varepsilon_2)}{32 \max_{i,j} B_{ij}} = \mathcal{O}(d^{-\frac{1}{2}}).$$

So, for $h = \mathcal{O}(d^{-\frac{1}{2}})$, by Lemma 17, we have $W(h, s) = \mathcal{O}(s^{\frac{1}{4}}d^{\frac{1}{2}})$. Gradient bounds are computed in Section C.6.

To summarise, Theorem 3 is applicable with:

- (i) $B_{00} = B_{02} = B_{12} = \mathcal{O}(1), B_{01} = \mathcal{O}(d^{\frac{1}{2}}), B_{22} = \mathcal{O}(d) \text{ and } C_{\mathfrak{g}} = 0.$
- (ii) $r_{\text{near}} = 1/(8\sqrt{2}), \ \bar{\varepsilon}_0 = 0.00097, \ \bar{\varepsilon}_2 = 0.941.$
- (iii) $\Delta = \mathcal{O}(d^{\frac{1}{2}}s_{\max}^{\frac{1}{4}}).$
- (iv) $\overline{L}_i = \mathcal{O}(d^{i/2}).$

and

$$m \gtrsim d^2 s \left(\log(s) \log\left(\frac{s}{\rho}\right) + \log\left(\frac{(fd)^d}{\rho}\right) \right).$$

C.1 Preliminaries: properties of the univariate kernel

We first summarise in Section C.1 some key properties of the univariate Jackson kernel κ when $f \ge 128$ which were derived in [15].

From [15, Equations (2.20)-(2.24) and (2.29)], for all $t \in [-1/2, 1/2]$ and $\ell = 0, 1, 2, 3$:

$$1 - \frac{C_f}{2}t^2 \leqslant \kappa(t) \leqslant 1 - \frac{C_f}{2}t^2 + 8\left(\frac{1+2/f}{1+2/(2+f)}\right)^2 C_f^2 t^4 \leqslant 1 - \frac{C_f}{2}t^2 + 8C_f^2 t^4$$
$$|\kappa'(t)| \leqslant C_f t, \quad |\kappa''(t)| \leqslant C_f, \quad |\kappa'''(t)| \leqslant 3\left(\frac{1+2/f}{1+2/(2+f)}\right)^2 C_f^2 t \leqslant 12C_f^2 t$$
$$\kappa'' \leqslant -C_f + \frac{3}{2}\left(\frac{1+2/f}{1+2/(2+f)}\right)^2 C_f^2 t^2 \leqslant -C_f + 6C_f^2 t^2.$$
(79)

By [15, Lemma 2.6],

$$\left|\kappa^{(\ell)}(t)\right| \leqslant \begin{cases} \frac{\pi^{\ell} H_{\ell}(t)}{(f+2)^{4-\ell}t^{4}}, & t \in \left[\frac{1}{2f}, \frac{\sqrt{2}}{\pi}\right] \\ \frac{\pi^{\ell} H_{\ell}^{\infty}}{(f+2)^{4-\ell}t^{4}}, & t \in \left[\frac{\sqrt{2}}{\pi}, \frac{1}{2}\right), \end{cases}$$

where $H_0^{\infty} \stackrel{\text{\tiny def.}}{=} 1$, $H_1^{\infty} \stackrel{\text{\tiny def.}}{=} 4$, $H_2^{\infty} \stackrel{\text{\tiny def.}}{=} 18$ and $H_3^{\infty} \stackrel{\text{\tiny def.}}{=} 77$, and $H_\ell(t) \stackrel{\text{\tiny def.}}{=} \alpha^4(t)\beta_\ell(t)$, with

$$\alpha(t) \stackrel{\text{\tiny def.}}{=} \frac{2}{\pi(1 - \frac{\pi^2 t^2}{6})}, \quad \bar{\beta}(t) \stackrel{\text{\tiny def.}}{=} \frac{\alpha(t)}{ft} = \frac{2}{ft\pi(1 - \pi^2 t^2/6)}$$

and $\beta_0(t) \stackrel{\text{def.}}{=} 1$, $\beta_1(t) \stackrel{\text{def.}}{=} 2 + 2\bar{\beta}(t)$, $\beta_2 \stackrel{\text{def.}}{=} 4 + 7\bar{\beta}(t) + 6\bar{\beta}(t)^2$ and $\beta_3(t) \stackrel{\text{def.}}{=} 8 + 24\bar{\beta} + 30\bar{\beta}(t)^2 + 15\bar{\beta}(t)^3$. Let us first remark that $\bar{\beta}$ is decreasing on $I \stackrel{\text{def.}}{=} [\frac{1}{2f}, \frac{\sqrt{2}}{\pi}]$, so $|\bar{\beta}(t)| \leq |\bar{\beta}(1/(2f))| \approx 1.2733$, and $a(t) \leq a(\sqrt{2}/\pi) = \frac{3}{\pi}$ on I. Therefore, on I, $H_0(t) \leq \frac{3}{\pi}$, $H_1(t) \leq 3.79$, $H_2(t) \leq 18.83$ and $H_3(t) \leq 98.26$, and we can conclude that on $[\frac{1}{2f}, \frac{1}{2})$, we have

$$\left|\kappa^{(\ell)}(t)\right| \leqslant \frac{\pi^{\ell} H_{\ell}^{\infty}}{(f+2)^{4-\ell} t^4}$$

where $H_0^{\infty} = 1$, $H_1^{\infty} \stackrel{\text{def.}}{=} 4$, $H_2^{\infty} \stackrel{\text{def.}}{=} 19$, $H_3^{\infty} \stackrel{\text{def.}}{=} 99$. Combining with (79), we have

$$\left\|\kappa^{(\ell)}\right\|_{\infty} \leqslant \kappa_{\ell}^{\infty} \tag{80}$$

where $\kappa_0^{\infty} \stackrel{\text{\tiny def.}}{=} 1, \, \kappa_2^{\infty} \stackrel{\text{\tiny def.}}{=} C_f,$

$$\kappa_1^{\infty} \stackrel{\text{\tiny def.}}{=} \sqrt{C_f} \max\left(\frac{2\pi^4}{(\frac{1}{2} + \frac{1}{f})^3} \frac{f}{\sqrt{C_f}}, \frac{\sqrt{C_f}}{2f}\right) = \mathcal{O}(\sqrt{C_f})$$
$$\kappa_3^{\infty} \stackrel{\text{\tiny def.}}{=} (C_f)^{3/2} \max\left(\frac{99\pi^3}{(\frac{1}{2} + \frac{1}{f})} \left(\frac{2f}{\sqrt{C_f}}\right)^4, \frac{6\sqrt{C_f}}{f}\right) = \mathcal{O}((C_f)^{3/2}).$$

Finally, given $p \in (0, 1)$,

$$(f+2)^4 t^4 \ge (1+p(f+2)^2 t^2)^2, \qquad \forall t \ge \frac{1}{\sqrt{(1-p)}(f+2)^2}$$

Choosing $p = \frac{1}{2}$ and using $(f+2)^2 = (\frac{3}{\pi^2}C_f + 4) \ge \frac{3}{\pi^2}C_f$, we have

$$\left|\kappa^{(\ell)}(t)\right| \leqslant \frac{\kappa_{\ell}^{\infty}}{(1+\frac{3}{2\pi^2}C_f t^2)^2}, \qquad \forall t^2 \geqslant \frac{2\pi^2}{3C_f},\tag{81}$$

In the following sections, we will repeatedly make use of (79), (80) and (81).

C.2 Notation

For notational convenience, write $t_i \stackrel{\text{\tiny def.}}{=} x_i - x'_i$, $\kappa_i \stackrel{\text{\tiny def.}}{=} \kappa(t_i)$, $\kappa'_i \stackrel{\text{\tiny def.}}{=} \kappa'(t_i)$, and so on. Let

$$K_i \stackrel{\text{def.}}{=} \prod_{\substack{k=1\\k\neq i}}^d \kappa_k, \quad K_{ij} \stackrel{\text{def.}}{=} \prod_{\substack{k=1\\k\neq i,j}}^d \kappa_k \quad \text{and} \qquad K_{ij\ell} \stackrel{\text{def.}}{=} \prod_{\substack{k=1\\k\neq i,j,\ell}}^d \kappa_k.$$

With this, we have:

$$\partial_{1,i}K(x,x') = \kappa'_i K_i$$

$$\partial_{1,i}\partial_{2,i}K(x,x') = -\kappa''_i K_i, \quad \text{and} \quad \forall i \neq j, \ \partial_{1,i}\partial_{2,j}K(x,x') = -\kappa'_i \kappa'_j K_{ij}.$$

Where convenient, we sometimes write $K(t) = K(x - x') \stackrel{\text{def.}}{=} K(x, x')$. Given a symmetric matrix M, we write $\lambda_{\min}(M)$ to denote the smallest eigenvalue of M.

C.3 Bounds when ||t|| is small

Lemma 15. Suppose that $C_f ||x - x'||_2^2 \leq c$ with c > 0 such that

$$\varepsilon \stackrel{\text{\tiny def.}}{=} (1 - 6c) \left(1 - \frac{c}{2 - c} \right) - c > 0$$

Then, $-\langle K^{(02)}(x-x')q, q \rangle \ge \varepsilon ||q||_x$.

Proof. Let $q \in \mathbb{R}^d$, and note that

$$-\langle \nabla_2^2 Kq, q \rangle = -\sum_i \left(q_i \kappa_i'' K_i - \kappa_i' \sum_{j \neq i} q_j \kappa_j' K_{ij} \right) q_i$$

$$= -\left(\sum_i q_i^2 \kappa_i'' K_i - \sum_i q_i \kappa_i \sum_{j \neq i} q_j \kappa_j K_{ij} \right)$$

$$\geqslant \|q\|_x^2 \frac{1}{C_f} \left(-\max_i \{\kappa_i'' K_i\} - \sum_j |\kappa_j'|^2 \right).$$
 (82)

We first consider $\kappa_i'' K_i$: By applying (79), we obtain

$$\kappa_i'' \leqslant -C_f + 6C_f^2 t_i^2,$$

$$K_i \geqslant \prod_{j \neq i} \left(1 - \frac{C_f}{2} t_i^2 \right) \geqslant 1 - \frac{C_f}{2} \|t\|_2^2 - \left(\frac{C_f}{2} \|t\|_2^2\right)^3 - \left(\frac{C_f}{2} \|t\|_2^2\right)^5 - \cdots$$

$$\geqslant 1 - \frac{C_f \|t\|_2^2}{2(1 - \frac{C_f}{2} \|t\|_2^2)}.$$

and hence,

$$\kappa_i'' K_i \leqslant \left(-C_f + 6C_f^2 \|t\|_2^2 \right) \left(1 - \frac{C_f \|t\|_2^2}{2(1 - \frac{C_f}{2} \|t\|_2^2)} \right)$$

For the second term in (82), again, by applying (79), we obtain

$$\sum_{j} \left| \kappa_{j}^{\prime} \right|^{2} \leqslant C_{f}^{2} \left\| t \right\|_{2}^{2}.$$

Therefore, for $\left\|q\right\|_{x} = 1$, we have

$$-\langle K^{(02)}(x-x')q, q \rangle \ge \left(1 - 6C_f \|t\|_2^2\right) \left(1 - \frac{C_f \|t\|_2^2}{2(1 - \frac{C_f}{2} \|t\|_2^2)}\right) - C_f \|t\|_2^2$$

Lemma 16. Assume that $\frac{1}{8\sqrt{C_f}} \ge ||t||_2$ Then,

$$K(t) \leq 1 - \frac{C_f}{4} \|t\|_2^2 + 16C_f^2 \|t\|_2^4.$$

Consequently, for all

$$0 < c \leqslant \frac{1}{8\sqrt{2C_f}},$$

 $and \ all \ t \ such \ that \ \|t\|_2 \geqslant c,$

$$|K(t)| \leqslant 1 - \frac{C_f}{8}c^2.$$

Proof. First note that by (79),

$$|\kappa(u)| \leq 1 - \frac{C_f}{2}u^2 + 32C_f^2u^4 = 1 - u^2g(u)$$

where

$$g(u) \stackrel{\text{\tiny def.}}{=} C_f\left(\frac{1}{2} - 32C_f u^2\right),$$

and note that $g(u) \in (0, \frac{C_f}{2})$ for $u \in (0, 1/(8\sqrt{C_f}))$. So, writing $t = (t_i)_{i=1}^d$ and $g_j \stackrel{\text{def.}}{=} g(t_j)$, we have

$$K(t) = \prod_{j=1}^{d} \kappa(t_i) \leqslant \prod_{j=1}^{d} \left(1 - t_j^2 \cdot g(t_j)\right)$$

= $1 - \sum_{j=1}^{d} t_j^2 g_j + \sum_{j \neq k} t_j^2 t_k^2 g_j g_k - \sum_{j \neq k \neq \ell} t_j^2 t_k^2 t_\ell^2 g_j g_k g_\ell + \cdots$
= $1 + \sum_{\ell=1}^{d} \sum_{j \in \mathcal{J}_\ell} (-1)^\ell \prod_{i=1}^{\ell} (t_{j_i}^2 g_{j_i}),$

where $\mathcal{J}_{\ell} \stackrel{\text{\tiny def.}}{=} \{ j \in \mathbb{N}^d ; j \leq d, \text{ all entries of } j \text{ are distinct} \}$. Note that for odd integers ℓ ,

$$-\sum_{j\in\mathcal{J}_{\ell}}\prod_{i=1}^{\ell}(t_{j_{i}}^{2}g_{j_{i}}^{2}) + \sum_{j\in\mathcal{J}_{\ell+1}}\prod_{i=1}^{\ell+1}(t_{j_{i}}^{2}g_{j_{i}}) \leqslant -\sum_{j\in\mathcal{J}_{\ell}}\prod_{i=1}^{\ell}(t_{j_{i}}^{2}g_{j_{i}}) + \left(\sum_{j\in\mathcal{J}_{\ell}}\prod_{i=1}^{\ell}(t_{j_{i}}^{2}g_{j_{i}}^{2})\right)\left(\sum_{k=1}^{d}t_{k}^{2}g_{k}\right)$$
$$\leqslant -\left(\sum_{j\in\mathcal{J}_{\ell}}\prod_{i=1}^{\ell}(t_{j_{i}}^{2}g_{j_{i}}^{2})\right)\left(1 - \frac{C_{f}}{2}\|t\|_{2}^{2}\right) < 0$$

since $\left(1 - \frac{C_f}{2} \|t\|_2^2\right) > 0$. Also,

$$\sum_{j=1}^{d} t_j^2 g_j \leqslant \frac{C_f}{2} \sum_{j=1}^{d} t_j^2 < 1$$

by assumption. So,

$$\begin{split} K(t) &\leqslant 1 - \sum_{j=1}^{d} t_{j}^{2} g_{j} + \sum_{j \neq k} t_{j}^{2} t_{k}^{2} g_{j} g_{k} \\ &\leqslant 1 - \sum_{j=1}^{d} t_{j}^{2} g_{j} + \frac{1}{2} \left(\sum_{j} t_{j}^{2} g_{j} \right)^{2} \leqslant 1 - \frac{1}{2} \sum_{j=1}^{d} t_{j}^{2} g_{j} \\ &\leqslant 1 - \frac{C_{f}}{2} \left(\frac{1}{2} \sum_{j=1}^{d} t_{j}^{2} - 32C_{f} \sum_{j=1}^{d} t_{j}^{4} \right) \leqslant 1 - \frac{C_{f}}{4} \| t \|_{2}^{2} + 16C_{f}^{2} \| t \|_{2}^{4} \end{split}$$

Finally, observe that the function

$$q(z) \stackrel{\text{\tiny def.}}{=} \frac{C_f}{4} z^2 - 16C_f^2 z^4$$

is positive and increasing on the interval $[0, \frac{1}{8\sqrt{2C_f}}]$. So, for t satisfing

$$c \leqslant \|t\|_2 \leqslant \frac{1}{8\sqrt{2C_f}},\tag{83}$$

we have $|K(t)| \leq 1 - q(c) \leq 1 - \frac{C_f}{8}c^2$. Finally, since |K(t)| is decreasing as t increases, we trivially have that $|K(t)| \leq 1 - q(c)$ for all t with $||t||_2 \geq c$.

C.4 Bounds when ||t|| is large

Lemma 17. Let $i, j \in \{0, 1, 2\}$ with $i + j \leq 3$. Let $\bar{A} \ge \sqrt{\frac{4\pi^2}{3}}$ and $||t||_2 \ge \bar{A}\sqrt{d}s_{\max}^{1/4}/\sqrt{C_f}$. Then, we have $||K^{(ij)}(t)||_{x,x'} \le d^{\frac{i+j-4}{2}}(\bar{A}^4s_{\max})^{-1}$.

Proof. Write $t = (t_j)_{j=1}^d$. To bound $K(t) = \prod_{j=1}^d \kappa(t_j)$, we want to make use of the bounds on κ_j^{∞} from (81). We can do this for each t_j such that $|t_j| \ge \sqrt{\frac{2\pi^2}{3C_f}}$. Note that there exists at least one such t_j since $||t||_{\infty} \ge ||t||_2 / \sqrt{d} \ge \bar{A} s_{\max}^{1/4} / \sqrt{C_f} \ge \sqrt{\frac{2\pi^2}{3C_f}}$. If $\{|t_j|\}_{j=1}^k \subset [0, \sqrt{\frac{2\pi^2}{3C_f}})$ for $k \le d-1$, then

$$k\frac{2\pi^2}{3C_f} + \sum_{j=k+1}^d t_j^2 \ge ||t||_2^2 \ge \frac{\bar{A}^2 ds_{\max}^{1/2}}{C_f},$$

which implies that $\sum_{j=k+1}^{d} t_j^2 \ge \frac{1}{C_f} \left(\bar{A}^2 ds_{\max}^{1/2} - \frac{2\pi^2(d-1)}{3} \right) \ge \frac{\bar{A}^2 ds_{\max}^{1/2}}{2C_f}$, by our assumptions on \bar{A} . Therefore, we may assume that we have some $d \ge p \ge 1$ such that $\{b_j\}_{j=1}^p \subseteq \{t_j\}$ with $|b_j| \ge \sqrt{\frac{2\pi^2}{3C_f}}$ and $\|b\|_2 \ge \frac{\bar{A}\sqrt{d}\sqrt[4]{s_{\max}}}{\sqrt{2C_f}}$. Observe that

$$\prod_{j=1}^{p} (1 + \frac{3C_f}{2\pi^2} b_j^2) \ge 1 + \frac{3C_f}{2\pi^2} \sum_{j=1}^{p} b_j^2 = 1 + \frac{3C_f}{2\pi^2} \|b\|_2^2 \ge 1 + \frac{3}{4\pi^2} \bar{A}^2 d\sqrt{s_{\max}}$$

So, by applying the fact that $|\kappa| \leqslant 1$, $\kappa_0^{\infty} = 1$ and (81), we have

$$|K(t)| \leq \prod_{j=1}^{p} |\kappa(b_j)| \leq \prod_{j=1}^{p} \frac{1}{\left(1 + \frac{3C_f}{2\pi^2} b_j^2\right)^2} \leq \frac{1}{\left(1 + \frac{3}{4\pi^2} \bar{A}^2 d\sqrt{s_{\max}}\right)^2}$$

For $|\kappa'_i K_i|$, if $i \notin \left\{ j \ ; \ |t_j| > \sqrt{\frac{2\pi^2}{3C_f}} \right\}$, then

$$|\kappa_i' K_i| \leq \|\kappa_i'\|_{\infty} \prod_{j=1}^{p} |\kappa(b_j)| \leq \frac{\|\kappa_i'\|_{\infty}}{\left(1 + \frac{3}{4\pi^2} \bar{A}^2 d\sqrt{s_{\max}}\right)^2}$$

and otherwise, we have $|\kappa'_i K_i| \leq |\kappa'(t_i)| \prod_{j \neq i} |\kappa(b_j)| \leq \frac{\kappa_1^{\infty}}{\left(1 + \frac{3}{4\pi^2} \bar{A}^2 d\sqrt{s_{\max}}\right)^2}$, In a similar manner, writing $V \stackrel{\text{def.}}{=} \left(1 + \frac{3}{4\pi^2} \bar{A}^2 d\sqrt{s_{\max}}\right)^{-2}$, we can deduce that

$$\begin{aligned} &|\kappa_i'K_i| \leqslant \kappa_1^{\max}V, \qquad |\kappa_i''K_i| \leqslant \kappa_2^{\infty}V, \qquad \left|\kappa_i'\kappa_j'K_{ij}\right|^2 \leqslant (\kappa_1^{\infty})^2V \\ &|\kappa_i'''K_i|^3 \leqslant \kappa_3^{\infty}V, \qquad \left|\kappa_i''\kappa_j'K_{ij}\right|^3 \leqslant \kappa_2^{\infty}\kappa_1^{\max}V, \qquad \left|\kappa_i'\kappa_j'\kappa_\ell'K_{ij\ell}\right| \leqslant (\kappa_1^{\max})^3V. \end{aligned}$$

Therefore,

$$\left\|K^{(10)}\right\|_{x,x'} = \frac{1}{\sqrt{C_f}} \left\|\nabla_1 K\right\| \leqslant \frac{1}{\sqrt{C_f}} \sqrt{\sum_{j=1}^d \left|\kappa'_j K_j\right|^2} \leqslant \frac{\kappa_1^{\infty}}{\sqrt{C_f}} V\sqrt{d} \lesssim \frac{1}{\bar{A}^4 d^{3/2} s_{\max}}.$$

Using Gershgorin theorem, we have

$$\left\|\nabla_{2}^{2}K(x,x')\right\| \leqslant \max_{1\leqslant i\leqslant d} \left\{\left|\kappa_{i}''K_{i}\right| + \left|\kappa_{i}'\right|\sum_{j\neq i}\left|\kappa_{j}'\right|\left|K_{ij}\right|\right\}$$

and hence,

$$\begin{split} \left\| K^{(02)} \right\|_{x'} &= \frac{1}{C_f} \left\| \nabla_2^2 K \right\| \leqslant \frac{1}{C_f} \max_{i=1}^d \{ |\kappa_i'' K_i| + |\kappa_i'| \sum_{j \neq i} |\kappa_j' K_{ij}| \} \\ &\leqslant \frac{1}{C_f} V \left(\kappa_2^\infty + (\kappa_1^\infty)^2 (d-1) \right) \leqslant \frac{\max\{\kappa_2^\infty, (\kappa_1^\infty)^2\}}{C_f} V d \lesssim \frac{1}{\bar{A}^4 ds_{\max}}. \end{split}$$

Note also that $\left\|K^{(11)}\right\|_{x,x'} = \left\|K^{(02)}\right\|_{x'}.$ Finally, since

$$\begin{aligned} \left\| \partial_{1,i} \nabla_{2}^{2} K(x,x') \right\| \leqslant \max \left\{ \left| \kappa_{i}^{\prime \prime \prime} K_{i} \right| + \left| \kappa_{i}^{\prime \prime} \right| \sum_{j \neq i} \left| \kappa_{j}^{\prime} \right| \left| K_{ij} \right|, \\ \max_{j \neq i} \left\{ \left| \kappa_{j}^{\prime \prime} \kappa_{i}^{\prime} K_{ij} \right| + \left| \kappa_{j}^{\prime} \kappa_{i}^{\prime \prime} K_{ij} \right| + \left| \kappa_{j}^{\prime} \right| \sum_{l \neq i,j} \left| \kappa_{l}^{\prime} \right| \left| K_{ij\ell} \right| \right\} \right\}, \end{aligned}$$

we have

$$\begin{split} \left| K^{(12)} \right\|_{x,x'} &= \frac{1}{C_f^{3/2}} \left\| \nabla_1 \nabla_2^2 K \right\| \\ &\leqslant \frac{1}{C_f^{3/2}} \sqrt{d} V \max\left(\kappa_3^\infty + \kappa_2^\infty \kappa_1^\infty (d-1), 2\kappa_2^\infty \kappa_1^\infty + (d-1)(\kappa_1^\infty)^3 \right) \\ &\leqslant d^{3/2} \max\{ \kappa_3^\infty, \kappa_1^\infty \kappa_2^\infty, (\kappa_1^\infty)^3 \} \frac{1}{C_f^{3/2}} V \lesssim \frac{1}{\bar{A}^4 d^{1/2} s_{\max}} \end{split}$$

L		
L		
L		
L		

C.5Uniform bounds

Lemma 18. If $r_{\text{near}} \sim 1/\sqrt{C_f}$, then $B_0 = \mathcal{O}(1)$, $B_{01} = \mathcal{O}(\sqrt{d})$, $B_{02} = B_{12} = \mathcal{O}(1)$ and $B_{22} = \mathcal{O}(d)$. *Proof.* We have $|K| \leq 1$, and

$$\left\|\nabla K\right\|^{2} \leqslant \sum_{i} \left|\kappa_{i}\right|^{2} \left|K_{i}\right|^{2} \leqslant d(\kappa_{1}^{\infty})^{2} \lesssim C_{f} d,$$

so $B_{01} = \mathcal{O}(\sqrt{d})$. From (82), for all ||q|| = 1,

$$\langle \nabla_2^2 K(t)q, q \rangle \leqslant \max_i |\kappa_i''| ||q||_2^2 + ||q||_2^2 \sum_i |\kappa_i|^2 \leqslant C_f + C_f^2 ||t||^2 = \mathcal{O}(C_f),$$

for $||t|| \lesssim 1/\sqrt{C_f}$. So, since $r_{\text{near}} \leqslant 2/\sqrt{C_f}$, $||K^{02}(t)|| \leqslant 2 \stackrel{\text{def.}}{=} B_{02}$. For the bound on B_{12} :

$$\begin{split} \left\| K^{(12)} \right\|_{x,x'} &= \sup_{\|q\| = \|p\| = 1} \frac{1}{C_f^{3/2}} \left(\sum_k \sum_{k \neq i} \partial_{1,i} \left(\partial_{2,k}^2 K p_i q_k^2 + \partial_{1,i} \partial_{2,i} \partial_{2,k} K p_i q_i q_k \right) \\ &+ \sum_i \sum_k \sum_j \partial_{1,i} \partial_{2,j} \partial_{2,k} p_i p_j p_k + \sum_i \sum_{j \neq i} \partial_{1,i} \partial_{2,j} K p_i q_i q_j + \sum_i \partial_{1,i} \partial_{2,j}^2 K p_i q_i^2 \right) \\ &= \sup_{\|q\| = \|p\| = 1} \frac{1}{C_f^{3/2}} \left(\sum_k \sum_{k \neq i} \kappa'_i \kappa''_k K_{ik} p_i q_k^2 + \kappa''_i \kappa'_k K_{ik} p_i q_i q_k \right) \\ &+ \sum_i \sum_k \sum_j \kappa'_i \kappa'_k \kappa'_j K_{ijk} p_i p_j p_k + \sum_i \sum_{j \neq i} \kappa''_i \kappa''_j K_{ij} p_i q_i q_j + \sum_i \kappa'_i \kappa''_j K_{ij} p_i q_i^2 \right) \\ &\leqslant \frac{1}{C_f^{3/2}} \left(3 \|\kappa''\|_{\infty} \sqrt{\sum_i |\kappa'_k|^2} + \left(\sum_i |\kappa'_k|^2 \right)^{3/2} + \|\kappa'\|_{\infty} \|\kappa''\|_{\infty} \right) \\ &\leqslant \frac{1}{C_f^{3/2}} \left(3 C_f^2 \|t\| + C_f^3 \|t\|^3 + \mathcal{O}(C_f^{3/2}) \right) = \mathcal{O}(1) \end{split}$$

for $||t|| \leq 1/C_f^{1/2}$.

C.6 Gradient bounds

The derivatives of the random features are uniformly bounded with

$$\left\|\nabla^{j}\varphi_{\omega}(x)\right\| = \|\omega\|^{j} \leqslant f^{j}d^{j/2} \sim C_{f}^{j/2}d^{j/2}$$

$$\tag{84}$$

So, we can set $\bar{L}_i = \mathcal{O}(d^{i/2})$ for i = 0, 1, 2. For \bar{L}_3 , the condition (44) is simply

$$C_{f}^{-1} \left\| \nabla^{2} \varphi_{\omega}(x) - \nabla^{2} \varphi_{\omega}(x') \right\| \leq \bar{L}_{3} C_{f}^{\frac{1}{2}} \left\| x - x' \right\|,$$

so $\bar{L}_3 = \mathcal{O}(d^{3/2})$ by (84).

D Application: Continuous Fourier sampling with the Gaussian kernel

In this section, we consider the case of continuous Fourier sampling with Gaussian frequencies, which may appear for instance in sketched Gaussian mixture learning [39]. Let $\mathcal{X} \subset \mathbb{R}^d$ be any bounded subset of \mathbb{R}^d . Let $\Omega = \mathbb{R}^d$, $\varphi_{\omega}(x) \stackrel{\text{def.}}{=} e^{i\omega^\top x}$, and $\Lambda(\omega) = \mathcal{N}(0, \Sigma^{-1})$, for a known covariance matrix Σ .

The kernel and Fisher metric The associated kernel is the Gaussian kernel

$$K(x, x') = \exp\left(-\frac{1}{2} \|x - x'\|_{\Sigma^{-1}}^2\right)$$

with constant metric tensor

$$\mathfrak{g}_x = \Sigma^{-1}$$
 and $\mathfrak{d}(x, x') = ||x - x'||_{\Sigma^{-1}} = \left\| \Sigma^{-\frac{1}{2}}(x - x') \right\|$

Sampling bounds The rest of this section consists of Lemmas which bound the parameters in Assumptions 1 and 2. We show that by choosing $r_{\text{near}} = \frac{1}{\sqrt{2}}$, we obtain $\bar{\varepsilon}_2 = \frac{1}{2}e^{-\frac{1}{4}}$ and $\bar{\varepsilon}_0 \stackrel{\text{def.}}{=} 1 - e^{-\frac{1}{4}}$. Moreover, Lemma 21 gives uniform bounds in $B_{ij} = \mathcal{O}(1)$ and, for $h = \mathcal{O}(1)$, $W(h, s) = \mathcal{O}(\sqrt{\log s} + 1)$. Gradient bounds are computed in Section D.5.

D.1 Properties of the kernel

Notations. For simplicity define t = x - x', b an abuse of notations $K_{\Sigma}(t) = \exp\left(-\frac{1}{2} \|t\|_{\Sigma^{-1}}^2\right)$ and for $u \in \mathbb{R}$, $\kappa(u) = \exp\left(-\frac{1}{2}u^2\right)$. Denote by $\{e_i\}$ the canonical basis of \mathbb{R}^d , and by $f_i = \Sigma^{-1}e_i$ the i^{th} row of Σ^{-1} .

Gradients of the kernel. We have the following:

$$\nabla K_{\Sigma}(t) = -\Sigma^{-1} t K_{\Sigma}(t)$$

$$\nabla^{2} K_{\Sigma}(t) = \left(-\Sigma^{-1} + \Sigma^{-1} t t^{\top} \Sigma^{-1}\right) K_{\Sigma}(t)$$

$$\partial_{i} \nabla^{2} K_{\Sigma}(t) = \left(\Sigma^{-1} t f_{i}^{\top} + f_{i} t^{\top} \Sigma^{-1}\right) K_{\Sigma}(t) - (t^{\top} f_{i}) \nabla^{2} K_{\Sigma}(t)$$

$$\partial_{ij} \nabla^{2} K_{\Sigma}(t) = \left(-\Sigma^{-1} ((t^{\top} f_{j}) t f_{i}^{\top} + (t^{\top} f_{i}) t f_{j}^{\top}\right) + (f_{i} f_{j}^{\top} + f_{j} f_{i}^{\top})\right) K_{\Sigma}(t) - f_{ij} \nabla^{2} K_{\Sigma}(t) - (t^{\top} f_{i}) \partial_{j} \nabla^{2} K_{\Sigma}(t)$$

Then we observe that for any $q \ge 1$ the function $f_q(r) = r^q e^{-\frac{1}{2}r^2}$ defined on \mathbb{R}_+ is increasing on $[0, \sqrt{q}]$ and decreasing after, and its maximum value is $f_q(\sqrt{q}) = \left(\frac{q}{e}\right)^{q/2}$. Furthermore, it is easy to see that we have $f_q(r) = r^q e^{-r^2/2} \le \left(\frac{2q}{2}\right)^{\frac{q}{2}} e^{-r^2/4}$ and therefore $f(r) \le \varepsilon$ if $r \ge 2 \left(\log\left(\frac{1}{\varepsilon}\right) + \frac{q}{2}\log\left(\frac{2q}{e}\right)\right)$.

D.2 Bounds when ||t|| is small

Lemma 19. For all $\mathfrak{d}_{\mathfrak{g}}(x,x') \leq r_{\text{near}} \stackrel{\text{def.}}{=} \frac{1}{\sqrt{2}}$ and all $v \in T_x \mathcal{M}$, we have $-K^{(02)}(x,x')[v,v] \geq \bar{\varepsilon}_2 \|v\|_x^2$ where $\bar{\varepsilon}_2 = \frac{1}{2}e^{-\frac{1}{4}}$.

Proof. From the derivations above we have $K^{(02)}(x, x')[v, v] = v^{\top} \nabla_2^2 K_{\Sigma}(t) v = (-1 + \mathfrak{d}_{\mathfrak{g}}(x, x')^2) \kappa(\mathfrak{d}_{\mathfrak{g}}(x, x')) \|v\|_x^2 \leq (r_{\text{near}}^2 - 1)\kappa(r_{\text{near}}) \|v\|_x.$

D.3 Bounds when ||t|| is large

Lemma 20. For all $\mathfrak{d}_{\mathfrak{g}}(x, x') \ge r_{\text{near}}$ we have $|K(x, x')| \le 1 - \overline{\varepsilon}_0$, where $\overline{\varepsilon}_0 \stackrel{\text{def.}}{=} 1 - e^{-\frac{1}{4}}$, and for $h = \mathcal{O}(1)$, $W(h, s) = \mathcal{O}(\sqrt{\log s} + 1)$.

Proof. For the first inequality we have $|K| \leq \kappa(r_{\text{near}}) = 1 - (1 - e^{-\frac{1}{4}})$.

Then, from (27), the fact that the metric tensor is constant, and the expressions for the derivatives of the kernel above, it is immediate that

$$\begin{split} \left\| K^{(10)}(x,x') \right\|_{x,x'} &= \left\| K^{(01)}(x,x') \right\|_{x,x'} = \left\| \Sigma^{\frac{1}{2}} \nabla K_{\Sigma}(t) \right\|_{1} \\ &= \mathfrak{d}(x,x')\kappa(\mathfrak{d}(x,x')) = f_{1}(\mathfrak{d}(x,x')) \\ \left\| K^{(02)}(x,x') \right\|_{x,x'} &= \left\| K^{(11)}(x,x') \right\|_{x,x'} = \left\| \Sigma^{\frac{1}{2}} \nabla^{2} K_{\Sigma}(t) \Sigma^{\frac{1}{2}} \right\|_{2} \\ &= (\mathfrak{d}(x,x')^{2} + 1)\kappa(\mathfrak{d}(x,x')) = f_{2}(\mathfrak{d}(x,x')) + f_{0}(\mathfrak{d}(x,x')) \end{split}$$

For $K^{(12)}$, again since the metric tensor \mathfrak{g} is constant, we observe that

$$[q]K^{(12)}(x,x')[v_1,v_2] = v_1^{\top} \left(\sum_{i=1}^d q_i \left(\partial_i \nabla^2 K_{\Sigma}(t) \right) \right) v_2$$

and

$$\begin{split} \left\| K^{(12)}(x,x') \right\| &= \sup_{\left\| \Sigma^{-\frac{1}{2}} q \right\|_{2} \leq 1, \left\| \Sigma^{-\frac{1}{2}} v_{i} \right\|_{2} \leq 1} \left| v_{1}^{\top} \left(\sum_{i=1}^{d} q_{i} \left(\partial_{i} \nabla^{2} K_{\Sigma}(t) \right) \right) v_{2} \right| \\ &= \sup_{\left\| q \right\|_{2} \leq 1, \left\| v_{i} \right\|_{2} \leq 1} \left| v_{1}^{\top} \left(\sum_{i=1}^{d} (\Sigma^{\frac{1}{2}} q)_{i} \Sigma^{\frac{1}{2}} \left(\partial_{i} \nabla^{2} K_{\Sigma}(t) \right) \Sigma^{\frac{1}{2}} \right) v_{2} \right| \\ &= \sup_{\left\| q \right\|_{1} = 1} \left\| \Sigma^{\frac{1}{2}} \left(\sum_{i=1}^{d} (\Sigma^{\frac{1}{2}} q)_{i} \partial_{i} \nabla^{2} K_{\Sigma}(t) \right) \Sigma^{\frac{1}{2}} \right\|_{2}. \end{split}$$

Using, $\sum_{i} (\Sigma^{\frac{1}{2}}q)_{i} f_{i} = \Sigma^{-\frac{1}{2}}q$, we observe that

$$\begin{split} \Sigma^{\frac{1}{2}} \left(\sum_{i} (\Sigma^{\frac{1}{2}} q)_{i} \Sigma^{-1} t f_{i}^{\top} \right) \Sigma^{\frac{1}{2}} &= \Sigma^{-\frac{1}{2}} t \left(\sum_{i} q^{\top} \Sigma^{\frac{1}{2}} e_{i} e_{i}^{\top} \Sigma^{-\frac{1}{2}} \right) = \Sigma^{-\frac{1}{2}} t q^{\top} \\ \Sigma^{\frac{1}{2}} \sum_{i} (\Sigma^{\frac{1}{2}} q)_{i} (t^{\top} f_{i}) \nabla^{2} K_{\Sigma}(t) \Sigma^{\frac{1}{2}} &= (q^{\top} \Sigma^{-\frac{1}{2}} t) \left(\Sigma^{\frac{1}{2}} \nabla^{2}_{2} K(x, x') \Sigma^{\frac{1}{2}} \right) \end{split}$$

Hence at the end of the day

$$\left\|K^{(12)}(x,x')\right\| \leq (3\mathfrak{d}(x,x') + \mathfrak{d}(x,x')^3)\kappa(\mathfrak{d}(x,x')) = 3f_1(\mathfrak{d}(x,x')) + f_3(\mathfrak{d}(x,x'))$$

Therefore, for h = O(1), using the properties of the functions f_q it is immediate that $W(h, s) = O(\sqrt{\log s} + 1)$.

D.4 Uniform bounds

Lemma 21. For $(i, j) \in \{0, 1, 2\}$, we have $B_{ij} = O(1)$.

Proof. The bounds for $i + j \leq 3$ are immediate using the identities in the proof of Lemma 20 and the properties of the functions f_q .

By the same reasoning we have

$$\left\| K^{(22)}(x,x') \right\| = \sup_{\|q_1\|=1, \|q_2\|=1} \left\| \Sigma^{\frac{1}{2}} \sum_{ij} (\Sigma^{\frac{1}{2}} q_1)_i (\Sigma^{\frac{1}{2}} q_2)_j \partial_{ij} \nabla^2 K_{\Sigma}(t) \Sigma^{\frac{1}{2}} \right\|$$

and we have

$$\begin{split} \Sigma^{\frac{1}{2}} \left(\sum_{ij} (\Sigma^{\frac{1}{2}} q_1)_i (\Sigma^{\frac{1}{2}} q_2)_j \Sigma^{-1} (t^{\top} f_j) (t f_i^{\top}) \right) \Sigma^{\frac{1}{2}} &= (q_2^{\top} \Sigma^{-\frac{1}{2}} t) \Sigma^{-\frac{1}{2}} t q_1^{\top} \\ \Sigma^{\frac{1}{2}} \left(\sum_{ij} (\Sigma^{\frac{1}{2}} q_1)_i (\Sigma^{\frac{1}{2}} q_2)_j f_i f_j^{\top} \right) \Sigma^{\frac{1}{2}} &= q_1 q_2^{\top} \\ \Sigma^{\frac{1}{2}} \left(\sum_{ij} (\Sigma^{\frac{1}{2}} q_1)_i (\Sigma^{\frac{1}{2}} q_2)_j f_{ij} \right) \Sigma^{\frac{1}{2}} &= q_1 q_2^{\top} \\ v_1 \Sigma^{\frac{1}{2}} \left(\sum_{ij} (\Sigma^{\frac{1}{2}} q_1)_i (\Sigma^{\frac{1}{2}} q_2)_j (t^{\top} f_i) \partial_j \nabla^2 K_{\Sigma}(t) \right) \Sigma^{\frac{1}{2}} v_2 &= (q_1^{\top} \Sigma^{-\frac{1}{2}} t) [q_2] K^{(12)}(x, x') [v_1, v_2] \end{split}$$

Hence

$$\left|K^{(22)}(x,x')\right\| \leq [3f_0 + 6f_2 + f_4](\mathfrak{d}(x,x'))$$

and $B_{22} = O(1)$.

D.5 Gradient bounds

For $j = \{0, 1, 2\}$, we have $D_j [\varphi_{\omega}](x)[q_1, \dots, q_j] = (\prod_i \omega^\top q_i) \varphi_{\omega}(x)$ and therefore

$$\left\| \mathbf{D}_{j} \left[\varphi_{\omega} \right] (x) \right\|_{x} \leqslant \left\| \Sigma^{\frac{1}{2}} \omega \right\|_{2}^{j}$$

And then, from (44), using $\tau_{x \to x'} = \text{Id}$,

$$\begin{split} \|\mathbf{D}_{2} [\varphi_{\omega}] (x) - \mathbf{D}_{2} [\varphi_{\omega}] (x') [\tau_{x \to x'} \cdot, \tau_{x \to x'} \cdot] \|_{x} &= \left\| \Sigma^{\frac{1}{2}} \left(\nabla_{2}^{2} \varphi(x') - \nabla_{2}^{2} \varphi(x) \right) \Sigma^{\frac{1}{2}} \right\|_{2} \\ &= \left\| \Sigma^{\frac{1}{2}} \omega \right\|_{2}^{2} |\varphi_{\omega}(x) - \varphi_{\omega}(x')| \\ &= \left\| \Sigma^{\frac{1}{2}} \omega \right\|_{2}^{2} |\omega^{\top}(x - x')| \leqslant \left\| \Sigma^{\frac{1}{2}} \omega \right\|_{2}^{3} \mathfrak{d}_{\mathfrak{g}}(x, x') \end{split}$$

Since $\omega \sim \mathcal{N}(0, \Sigma^{-1})$, $\left\| \Sigma^{\frac{1}{2}} \omega \right\|^{j} = W^{\frac{j}{2}}$ where W is a χ^{2} variable with d degrees of freedom. Then, we use the following Chernoff bound [26]: for $x \ge d$, we have

$$\mathbb{P}(W \ge x) \leqslant \left(\frac{ex}{d}e^{-\frac{x}{d}}\right)^{\frac{d}{2}} \leqslant \left(e\left(\sqrt{\frac{x}{d}}\right)^2 e^{-\frac{1}{2}\cdot\left(\sqrt{\frac{x}{d}}\right)^2} e^{-\frac{x}{2d}}\right)^{\frac{d}{2}} \leqslant 2^{\frac{d}{2}}e^{-\frac{x}{4}}$$

by using $x^2 e^{-\frac{x^2}{2}} \leq \frac{2}{e}$.

Hence we can define the F_j such that, for all $t \ge d^{j/2}$, $\mathbb{P}(L_j(\omega) \ge t) \le F_j(t) = 2^{\frac{d}{2}} \exp\left(-\frac{t^2_j}{4}\right)$, and $F_j(\bar{L}_j)$ is smaller than some δ if $\bar{L}_j \propto \left(d + \log \frac{1}{\delta}\right)^{\frac{j}{2}}$. Then we must choose the L_j such that $\int_{\bar{L}_j} tF_j(t)dt$ is bounded by some δ . Taking $\bar{L}_j \ge d^{j/2}$ in any case, we have

$$\begin{split} \int_{\bar{L}_{j}} tF_{j}(t)dt &= 2^{\frac{d}{2}} \int_{\bar{L}_{j}} t\exp\left(-\frac{t^{\frac{2}{j}}}{4}\right)dt = 2^{\frac{d}{2}} \int_{\bar{L}_{j}^{\frac{2}{j}}} (j/2)t^{j-1} \exp\left(-\frac{t}{4}\right)dt \\ &= 2^{\frac{d}{2}}(j/2) \int_{\bar{L}_{j}^{\frac{2}{j}}} \left(t^{j-1} \exp\left(-\frac{t}{8}\right)\right) \exp\left(-\frac{t}{8}\right)dt \leqslant 2^{\frac{d}{2}}(j/2) \left(\frac{8(j-1)}{e}\right)^{j-1} \int_{\bar{L}_{j}^{\frac{2}{j}}} \exp\left(-\frac{t}{8}\right)dt \\ &= 2^{\frac{d}{2}}j \left(\frac{8(j-1)}{e}\right)^{j-1} 8\exp\left(-\bar{L}_{j}^{\frac{2}{j}}/8\right) \end{split}$$

Hence this quantity is bounded by δ if $\bar{L}_j \propto \left(d + \log\left(\frac{1}{\delta}\right)\right)^{\frac{j}{2}}$. Then we have $\bar{L}_j^2 F_i(\bar{L}_i) = \bar{L}_j^2 2^{\frac{d}{2}} \exp\left(-\frac{\bar{L}_i^{\frac{2}{4}}}{4}\right)$ which is also bounded by δ if $\bar{L}_j \propto \left(d + \left(\log\frac{d}{\delta}\right)^2\right)^{\frac{j}{2}}$. At the end of the day, our assumptions are satisfied for

$$\bar{L}_j \propto \left(d + \left(\log \frac{dm}{\rho}\right)^2\right)^{\frac{j}{2}}$$

D.5.1 Gaussian mixture model learning

We apply the mixture model framework with the base distribution:

$$P_{\theta} = \mathcal{N}(\theta, \Sigma)$$

The random features on the data space are $\varphi'_{\omega}(x) = C e^{i\omega^{\top}x}$ with Gaussian distribution $\omega \sim \Lambda = \mathcal{N}(0, A)$ for some constant C and matrix A that we will choose later. Then, the features on the parameter space

are $\varphi_{\omega}(\theta) = \mathbb{E}_{x \sim P_{\theta}} \varphi'_{\omega}(x) = C e^{i\omega^{\top} \theta} e^{-\frac{1}{2} \|\omega\|_{\Sigma}^{2}}$ (that is, the characteristic function of Gaussians). Then, it is possible to show [39] that the kernel is

$$K(\theta, \theta') = C^2 \frac{|A^{-1}|^{\frac{1}{2}}}{|2\Sigma + A^{-1}|^{\frac{1}{2}}} e^{-\frac{1}{2} \|\theta - \theta'\|^2_{(2\Sigma + A^{-1})^{-1}}}$$

Hence we choose $A = c\Sigma^{-1}$, $C = (1+2c)^{\frac{d}{4}}$, and we come back to the previous case $K(\theta, \theta') = e^{-\frac{1}{2} \|\theta - \theta'\|_{\tilde{\Sigma}^{-1}}^2}$ with covariance $\tilde{\Sigma} = (2+1/c)\Sigma$. Hence $\bar{\varepsilon}_i = \mathcal{O}(1)$, $B_{ij} = \mathcal{O}(1)$, $\mathfrak{d}(\theta, \theta') = \|\theta - \theta'\|_{\tilde{\Sigma}^{-1}} = \frac{1}{\sqrt{2+1/c}} \|\theta - \theta'\|_{\Sigma^{-1}}$.

Admissible features. Unlike the previous case, the features are directly bounded and Lipschitz. We have

$$\begin{aligned} |\varphi_{\omega}(\theta)| &\leqslant C \stackrel{\text{def.}}{=} L_{0}, \\ \|\mathcal{D}_{j}\left[\varphi_{\omega}(\theta)\right]\| &= C \left\|\tilde{\Sigma}^{\frac{1}{2}}\omega\right\|^{j} e^{-\frac{\|\omega\|_{\Sigma}^{2}}{2}} = C\left(2+1/c\right)^{\frac{j}{2}} \left\|\Sigma^{\frac{1}{2}}\omega\right\|^{j} e^{-\frac{\|\omega\|_{\Sigma}^{2}}{2}} \leqslant C\left(2+1/c\right)^{\frac{j}{2}} \left(\frac{j}{e}\right)^{\frac{j}{2}} \stackrel{\text{def.}}{=} L_{j} \end{aligned}$$

Hence all constants L_j are in $\mathcal{O}\left(C(2+1/c)^{\frac{j}{2}}\right)$ by choosing $c = \frac{1}{d}$ they are in $\mathcal{O}\left(d^{\frac{j}{2}}\right)$.

E Application: Sampling the Laplace transform

Let $\alpha \in \mathbb{R}^d_+$ and let $\mathcal{X} = (0, R]^d \subset \mathbb{R}^d_+$ for some R > 0. Let $\Omega = \mathbb{R}^d_+$. Define for $x \in \mathcal{X}$ and $\omega \in \Omega$,

$$\varphi_{\omega}(x) \stackrel{\text{\tiny def.}}{=} \exp\left(-\langle x,\,\omega\rangle\right) \prod_{i=1}^{d} \sqrt{\frac{x_i + \alpha_i}{\alpha_i}} \quad \text{and} \quad \Lambda(\omega) = \exp\left(-\langle 2\alpha,\,\omega\rangle\right) \prod_{i=1}^{d} (2\alpha_i) \sum_{i=1}^{d} (2\alpha_i) \sum$$

The kernel and Fisher metric The associated kernel is $K(x, x') = \prod_{i=1}^{d} \kappa(x_i + \alpha_i, x'_i + \alpha_i)$ where

$$\kappa(u,v) \stackrel{\text{\tiny def.}}{=} 2\frac{\sqrt{uv}}{u+v}$$

The associated metric $\mathfrak{g}_x \in \mathbb{R}^{d \times d}$ is the diagonal matrix with diagonal $(h_{x_i + \alpha_i})_{i=1}^d$ where given $x \in \mathbb{R}_+$, $h_x \stackrel{\text{def.}}{=} \partial_x \partial_{x'} \kappa(x, x) = (2x)^{-2}$. The induced distance in dimension one is

$$\int_{\min\{s,t\}}^{\max\{s,t\}} (2x+2\alpha)^{-1} \mathrm{d}x = \left| \log\left(\frac{t+\alpha}{s+\alpha}\right) \right|$$
(85)

and hence,

$$\mathfrak{d}_{\mathfrak{g}}(x, x') = \sqrt{\sum_{i=1}^{d} \left| \log \left(\frac{x_i + \alpha_i}{x'_i + \alpha_i} \right) \right|^2}$$

is the Fisher distance between exponential distributions. The domain diameter is $\mathcal{R}_{\mathcal{X}} = \sqrt{\sum_{i} \left| \log \left(\frac{R + \alpha_i}{\alpha_i} \right) \right|}$.

The Christoffel symbol is $\Gamma_{jk}^i = -(x_i + \alpha_i)^{-1}$ when i = j = k and 0 otherwise, so the Riemannian Hessian of f at x is

$$\mathrm{H}f(x) = \nabla^2 f(x) + \mathrm{diag}(\mathfrak{g}_x^{\overline{2}} \nabla f(x))$$

Sampling bounds Assuming that the $\alpha_i \sim d$ and are all distinct, Theorem 3 is applicable with:

(i)
$$B_{00} = B_{01} = B_{02} = \mathcal{O}(1), B_{12} = \mathcal{O}(\sqrt{d}), B_{22} = \mathcal{O}(d).$$

(ii) $r_{\text{near}} = 0.2, \, \bar{\varepsilon}_0 = 0.005, \, \bar{\varepsilon}_2 = 0.7960.$

(iii)
$$\Delta = \mathcal{O}(d + \log(d^{3/2}s_{\max}))$$

(iv) $\bar{L}_j \propto d^j \left(\sqrt{d} + \left(\log(m) + \log\left(\frac{d}{\rho}\right)\right)\right)^j$

 $\quad \text{and} \quad$

$$m \gtrsim s\left(C\log(s)\log\left(\frac{s}{\rho}\right) + C^2\log\left(\frac{C^d}{\rho}\right)\right)$$
(86)

where $C \stackrel{\text{def.}}{=} d^2 \left(d + \log^2(m) + \log^2\left(\frac{d}{\rho}\right) \right)$. In the above, the implicit constant depends on R.

E.1 Preliminaries: properties of the univariate kernel

We first provide bounds for κ and its derivatives. In the following, let

$$\kappa^{(ij)}(u,v) \stackrel{\text{\tiny def.}}{=} h_u^{-i/2} h_v^{-j/2} \partial_u^i \partial_v^j \kappa(u,v).$$

We denote $\mathfrak{d}_{\kappa}(u,v) \stackrel{\text{\tiny def.}}{=} |\log(u/v)|$. Recall also the hyperbolic functions

$$\sinh(u) \stackrel{\text{\tiny def.}}{=} \frac{e^u - e^{-u}}{2}, \quad \cosh(u) \stackrel{\text{\tiny def.}}{=} \frac{e^u + e^{-u}}{2}, \quad \tanh(u) \stackrel{\text{\tiny def.}}{=} \frac{\sinh(u)}{\cosh(u)}, \quad \operatorname{sech}(u) \stackrel{\text{\tiny def.}}{=} \frac{1}{\cosh(u)}$$

Lemma 22. We have

$$\begin{aligned} (i) \ \kappa(u,v) &= \operatorname{sech}\left(\frac{\mathfrak{d}_{\kappa}(u,v)}{2}\right) \leqslant 2e^{-\frac{1}{2}\mathfrak{d}_{\kappa}(u,v)}.\\ (ii) \ \left|\kappa^{(10)}(u,v)\right| &= 2\left|\operatorname{tanh}\left(\frac{\mathfrak{d}_{\kappa}(u,v)}{2}\right)\kappa(u,v)\right|, \ and \ \left|\kappa^{(10)}(u,v)\right| \leqslant 2\left|\kappa(u,v)\right|.\\ (iii) \ \left|\kappa^{(11)}(u,v)\right| &\leq \left|\kappa(u,v)\right|^{3} + 4\left|\kappa(u,v)\right|\\ (iv) \ \left|\kappa^{(20)}(u,v)\right| &\leq 5\left|\kappa(u,v)\right| \ and \ -\kappa^{(20)}(u,v) \geqslant \kappa(u,v)\left(1 - 4\operatorname{tanh}\left(\frac{\mathfrak{d}_{\kappa}(u,v)}{2}\right)\right).\\ (v) \ \left|\kappa^{(12)}(u,v)\right| \leqslant 49\left|\kappa(u,v)\right|.\end{aligned}$$

Proof. We first state the partial derivatives of κ :

$$\begin{split} \kappa(u,v) &= \frac{2\sqrt{uv}}{u+v},\\ \partial_u \kappa(u,v) &= \frac{v(v-u)}{\sqrt{uv}(u+v)^2}\\ \partial_u \partial_v \kappa(u,v) &= \frac{-u^2 + 6uv - (v)^2}{2\sqrt{uv}(u+v)^3}\\ \partial_u^2 \kappa(u,v) &= -\frac{(v)^2 \left((u+v)^2 + 4u(v-u)\right)}{2 \left(uv\right)^{3/2} (u+v)^3}\\ &= -\frac{(v)^2}{2 \left(uv\right)^{3/2} (u+v)} - \frac{2v(v-u)}{(uv)^{1/2} (u+v)^3}\\ \partial_u \partial_v^2 \kappa(u,v) &= \frac{u^3 + 13u^2v - 33u(v)^2 + 3(v)^3}{4v(uv)^{1/2} (u+v)^4} \end{split}$$

We also make use of the following fact: For u > v,

$$\begin{aligned} \frac{v-u}{u+v} &= \left(\frac{1}{\frac{u}{v}+1} - \frac{1}{1+\frac{v}{u}}\right) \\ &= \left(\frac{1}{1+\exp(\mathfrak{d}_{\kappa}(u,v))} - \frac{1}{1+\exp(-\mathfrak{d}_{\kappa}(u,v))}\right) \\ &= \left(\frac{\exp(-\mathfrak{d}_{\kappa}(u,v)) - \exp(\mathfrak{d}_{\kappa}(u,v))}{2+\exp(\mathfrak{d}_{\kappa}(u,v)) + \exp(\mathfrak{d}_{\kappa}(u,v))}\right) \\ &= \frac{-\sinh(\mathfrak{d}_{\kappa}(u,v))}{1+\cosh(\mathfrak{d}_{\kappa}(u,v))} = -\tanh(\mathfrak{d}_{\kappa}(u,v)/2). \end{aligned}$$

(i)

$$\kappa(u,v) = 2\left(\sqrt{\frac{u}{v}} + \sqrt{\frac{v}{u}}\right)^{-1} = \frac{2}{e^{-\frac{\mathfrak{d}_{\kappa}(u,v)}{2}} + e^{\frac{\mathfrak{d}_{\kappa}(u,v)}{2}}} = \frac{1}{\cosh(\frac{\mathfrak{d}_{\kappa}(u,v)}{2})} \leqslant 2e^{-\frac{1}{2}\mathfrak{d}_{\kappa}(u,v)},$$

(ii) We have, assuming that u > v,

$$\kappa^{(10)}(u,v) = 2u\partial_u\kappa(u,v) = 2\frac{v-u}{u+v}\kappa(u,v) = -2\tanh(\mathfrak{d}_\kappa(u,v)/2)\kappa(u,v).$$

(iii)

$$\kappa^{(11)}(u,v) = 4uv\partial_v\partial_u\kappa(u,v) = 4uv\frac{4uv - (u-v)^2}{2\sqrt{uv}(u+v)^3}$$
$$= \kappa(u,v)\left(\kappa(u,v)^2 - \frac{(u-v)^2}{(u+v)^2}\right)$$
$$= \kappa(u,v)\left(\kappa(u,v)^2 - 4\tanh^2(\mathfrak{d}_\kappa(u,v)/2)\right)$$

so $\left| \kappa^{(11)} \right| \leq \left| \kappa \right|^3 + 4 \left| \kappa \right|.$ (iv)

$$\kappa^{(20)}(u,v) = 4u^2 \partial_u^2 \kappa(u,v) = -\frac{4(uv)^{1/2} \left((u+v)^2 + 4u(v-u)\right)}{2(u+v)^3}$$
$$= -\kappa(u,v) \left(1 + \frac{4u(v-u)}{(u+v)^2}\right)$$

so $\left|\kappa^{(20)}\right| \leqslant 5 \left|\kappa\right|$. Also,

$$-\kappa^{(20)} \ge \kappa(u,v) \left(1 - 4 \tanh(\mathfrak{d}_{\kappa}(u,v)/2)\right)$$

(v)

$$\kappa^{(12)}(u,v) = 2u(2v)^2 \partial_u \partial_v^2 \kappa(u,v)$$

= $\kappa(u,v) \left(1 + \frac{2v(5u^2 - 18uv + (v)^2)}{(u+v)^3} \right)$

so $\left|\kappa^{(12)}\right| \leqslant 49 \left|\kappa\right|$.

E.2 Kernel bounds

Theorem 5 (Kernel bounds). The following hold:

- 1. $1 \frac{1}{8}\mathfrak{d}(x, x')^2 \leq |K(x, x')| \leq \min\{2^d e^{-\frac{1}{2}\mathfrak{d}(x, x')}, \frac{8}{8 + \mathfrak{d}(x, x')^2}\}.$
- 2. $||K^{(10)}(x, x')|| \le \min\{2\sqrt{d} |K|, \sqrt{2}\}.$
- 3. $||K^{(11)}|| \leq \min\{9d |K|, 8\}$
- 4. $||K^{(20)}|| \leq \min\{9d | K|, 8\}$ and $\lambda_{\min}(-K^{(20)}) \geq (1 5\mathfrak{d}(x, x')^2) K$ when $\mathfrak{d}(x, x') \leq 1$.
- 5. $\|K^{(12)}\| \leq \min\{66 |K| d^{3/2}, 16\sqrt{d} + 49\}$ and $\|K^{(12)}(x, x')\| \leq 34$ if $\mathfrak{d}(x, x') \leq 1$. In particular, for $\mathfrak{d}(x, x') \geq 2d \log(2) + 2 \log\left(\frac{52d^{3/2}s_{\max}}{h}\right)$, we have $\|K^{(ij)}(x, x')\| \leq \frac{h}{s_{\max}}$.

Proof. Let $d_{\ell} \stackrel{\text{def.}}{=} \mathfrak{d}_{\kappa}(x_{\ell} + \alpha_{\ell}, x'_{\ell} + \alpha_{\ell})$ and note that $\mathfrak{d}_{\mathfrak{g}}(x, x') = \sqrt{\sum_{\ell} d_{\ell}^2}$. Define $g = \left(2 \tanh(\frac{d_{\ell}}{2})\right)_{\ell=1}^d$. We first prove that

- (i) $|K(x,x')| \leq \prod_{\ell=1}^{d} \operatorname{sech}(d_{\ell}/2) \leq \prod_{\ell=1}^{d} \frac{1}{1+d_{\ell}^2/8} \leq \frac{1}{1+\frac{1}{8}\mathfrak{d}(x,x')^2}$
- (ii) $||K^{(10)}(x, x')|| \leq ||g||_2 |K|.$
- (iii) $||K^{(11)}|| \leq |K| \left(||g||_2^2 + 5 \right)$
- (iv) $||K^{(20)}|| \leq |K| (||g||_2^2 + 5)$ and $\lambda_{\min} (-K^{(20)}) \geq K (1 5 ||g||_2^2)$.
- (v) $||K^{(12)}|| \leq |K| \left(||g||_2^3 + 16 ||g||_2 + 49 \right)$

The result would then follow because $|\tanh(x)| \leq \min\{x, 1\}$, so $||g|| \leq \min\{\mathfrak{d}(x, x'), 2\sqrt{d}\}$. For example, $||K^{(12)}|| \leq \frac{1}{1+\frac{1}{8}\mathfrak{d}(x, x')^2} \left(\mathfrak{d}(x, x')^3 + 16\mathfrak{d}(x, x') + 24\right) \leq 8\mathfrak{d}(x, x') + \frac{\sqrt{8}}{2} + 24 \leq 34$ when $\mathfrak{d}(x, x') \leq 1$.

In the following, we write

$$\kappa_{\ell}^{(ij)} \stackrel{\text{\tiny def.}}{=} \kappa^{(ij)}(x_{\ell} + \alpha_{\ell}, x_{\ell}' + \alpha_{\ell})$$

and $\kappa_{\ell} \stackrel{\text{def.}}{=} \kappa_{\ell}^{(00)}$ and $K_i \stackrel{\text{def.}}{=} \prod_{j \neq i} \kappa_j$. Moreover, we will make use of the inequalities for $\kappa^{(ij)}$ derived in Lemma 22.

(i) Note that $\operatorname{sech}(x) \leq 2e^{-x}$ and $\operatorname{sech}(x) \leq (1 + x^2/2)^{-1}$. So,

$$K(x,x')| \leqslant \prod_{\ell=1}^d \operatorname{sech}\left(\frac{\mathrm{d}_\ell}{2}\right) \leqslant \prod_{\ell=1}^d \left(1 + \frac{\mathrm{d}_\ell^2}{2}\right)^{-1} \leqslant \frac{1}{1 + \mathfrak{d}(x,x')^2}$$

Also, since $\operatorname{sech}(x) \ge 1 - \frac{x^2}{2}$, we also have $K(x, x') \ge \prod_{\ell=1}^d \left(1 - \frac{1}{8} \mathrm{d}_{\ell}^2\right) \ge 1 - \frac{1}{8} \mathfrak{d}(x, x')^2$. (ii) Note that $\|K^{(10)}(x, x')\| = \left\| \left(\kappa_{\ell}^{(10)} K_{\ell} \right)_{\ell=1}^d \right\|$, so by Lemma 22 (ii),

$$\left\| K^{(10)}(x,x') \right\| \le \|g\|_2 |K|.$$

(iii) For $i \neq j$

$$\left|\kappa_{i}^{(10)}\kappa_{j}^{(01)}K_{ij}\right| \leq 4 \tanh\left(\frac{\mathbf{d}_{i}}{2}\right) \tanh\left(\frac{\mathbf{d}_{j}}{2}\right) |K|,$$

and $\left|\kappa_{i}^{(11)}K_{i}\right|\leqslant5\left|K\right|$. So, given $p\in\mathbb{R}^{d}$ of unit norm,

$$\begin{aligned} \left\| K^{(11)} \right\| &= \sup_{\|p\|=1} \sum_{i=1}^{d} \sum_{j \neq i} \kappa_{i}^{(10)} \kappa_{j}^{(01)} K_{ij} p_{i} p_{j} + \sum_{i=1}^{d} p_{i}^{2} \kappa_{i}^{(11)} K_{i} \\ &\leqslant \sup_{\|p\|=1} |K| \left(\sum_{i=1}^{d} \sum_{j \neq i} 4 \tanh(d_{i}/2) \tanh(d_{j}/2) p_{i} p_{j} + 5 \sum_{i=1}^{d} p_{i}^{2} \right) \\ &\leqslant |K| \left(\left\| g \right\|_{2}^{2} + 5 \right). \end{aligned}$$

(iv) Note that

$$\left\|K^{(20)}\right\| = \sup_{\|p\|=1} \left|\sum_{i=1}^{d} \sum_{j \neq i} \kappa_i^{(10)} \kappa_j^{(10)} K_{ij} p_i p_j + \sum_{i=1}^{d} p_i^2 \kappa_i^{(20)} K_i + \sum_{i=1}^{d} \kappa_i^{(10)} K_i p_i^2\right|.$$

Observe that $\left|\kappa_{i}^{(20)}K_{i}\right| \leq 5|K|$ and $-\kappa_{i}^{(20)}K_{i} \geq K\left(1-4\tanh\left(\frac{d_{i}}{2}\right)\right)$.

$$\begin{split} \left\| K^{(20)} \right\| &\leq \sup_{\|p\| \leq 1} \left| \sum_{i=1}^{d} \sum_{j \neq i} \kappa_{i}^{(10)} \kappa_{j}^{(10)} K_{ij} p_{i} p_{j} + \sum_{i=1}^{d} p_{i}^{2} \kappa_{i}^{(20)} K_{i} \right| + \left\| g \right\|_{2} |K| \\ &\leq |K| \sup_{\|p\| \leq 1} \left(\sum_{i=1}^{d} \sum_{j \neq i} 4 \tanh(d_{i}/2) \tanh(d_{j}/2) p_{i} p_{j} + 5 \sum_{i=1}^{d} p_{i}^{2} \right) + \left\| g \right\|_{2} |K| \\ &\leq |K| \left(2 \left\| g \right\|_{2}^{2} + 5 \right), \end{split}$$

and given any p with $\left\|p\right\|_{x}=1,$

$$\langle -K^{(20)}p, p \rangle \ge K (1 - 4 \|g\|_{\infty})$$

(v)Note that $\|K^{12}\|_{x,x'} = \|A\|$, where $A = (A_{ij\ell})_{i,j,\ell=1}^d$ is defined as follows: For i, j, ℓ all distinct,

$$A_{ij\ell} = \kappa_i^{(10)} \kappa_j^{(01)} \kappa_\ell^{(01)} K_{ij\ell} \leqslant 8 \tanh\left(\frac{\mathbf{d}_i}{2}\right) \tanh\left(\frac{\mathbf{d}_j}{2}\right) \tanh\left(\frac{\mathbf{d}_\ell}{2}\right) K,$$

for all i, ℓ distinct,

$$A_{ii\ell} = 8\kappa_i^{(11)}\kappa_\ell^{(01)}K_{i\ell} \leqslant 10 \tanh\left(\frac{\mathrm{d}_\ell}{2}\right)K,$$
$$A_{i\ell i} = \kappa_i^{(11)}\kappa_\ell^{(01)}K_{ij} \leqslant 10 \tanh\left(\frac{\mathrm{d}_j}{2}\right)K,$$

and for $i \neq j$, $A_{ijj} = \kappa_i^{(10)} \kappa_j^{(02)} K_{ij} \leq 12 \tanh\left(\frac{\mathrm{d}_i}{2}\right) K$,

$$A_{ijj} = \kappa_i^{(10)} \kappa_j^{(02)} K_{ij} + \kappa_i^{(10)} \kappa_j^{(01)} K_{ij} \le 10 \tanh\left(\frac{d_i}{2}\right) K + 2 \tanh\left(\frac{d_i}{2}\right) K$$

and $A_{iii} = (\kappa_i^{(12)} + \kappa_i^{(02)})K_i \leqslant 54K$. So, for $p, q \in \mathbb{R}^d$ of unit norm,

$$\sum_{i} \sum_{j} \sum_{\ell} A_{ij\ell} p_{j} p_{\ell} q_{i} = \sum_{i} \left(\sum_{j \neq i} \sum_{\ell} A_{ij\ell} p_{j} p_{\ell} q_{i} + \sum_{\ell} A_{ii\ell} p_{i} p_{\ell} q_{i} \right)$$

=
$$\sum_{i} \sum_{j \neq i} \left(\sum_{\ell \notin \{i,j\}} A_{ij\ell} p_{j} p_{\ell} q_{i} + A_{iji} p_{j} p_{i} q_{i} + K_{ijj}^{(12)} p_{j}^{2} q_{i} \right)$$

+
$$\sum_{i} \sum_{\ell \neq i} A_{ii\ell} p_{i} p_{\ell} q_{i} + \sum_{i} A_{iii} p_{i}^{2} q_{i}$$

$$\leq |K| \left(||g||_{2}^{3} + 16 ||g||_{2} + 49 \right).$$

E.3 Gradient bounds

Theorem 6 (Stochastic gradient bounds). Assume that the α_i 's are all distinct. Then, $\bar{L}_0(\omega) \leq \bar{L}_0 \stackrel{\text{def.}}{=} \left(1 + \frac{R}{\min_i \alpha_i}\right)^d$ and

$$\mathbb{P}(L_j(\omega) \ge t) \leqslant F_j(t) \stackrel{\text{def.}}{=} \sum_{i=1}^d \beta_i \exp\left(-\alpha_i \left(\frac{1}{2(R+\|\alpha\|_{\infty})} \left(\frac{t}{\bar{L}_0}\right)^{1/j} - \sqrt{d}\right)\right), \qquad j \in \{1,2,3\}$$

and we have that $\sum_i F_j(\bar{L}_j) \leqslant \delta$ and $\bar{L}_j^2 \sum_i F_i(\bar{L}_i) + 2 \int_{\bar{L}_j}^{\infty} tF_j(t) dt \leqslant \delta$ provided that

$$\bar{L}_j \propto \bar{L}_0 (R + \|\alpha\|_\infty)^j \left(\sqrt{d} + \max_i \frac{1}{\alpha_i} \log\left(\frac{d\beta_i \bar{L}_0 (R + \|\alpha\|_\infty)}{\delta\alpha_i}\right)\right)^j, \qquad j \in \{1, 2, 3\}$$

where $\beta_i = \prod_{j \neq i} \frac{\alpha_j}{\alpha_j - \alpha_i}$. Note that $\alpha_i \sim d$ implies that $\bar{L}_0 \sim (1 + R/d)^d \sim e^R$.

Proof. Let $V_x \stackrel{\text{\tiny def.}}{=} (1 - 2(x_i + \alpha_i)\omega_i)_{i=1}^d \in \mathbb{R}^d$. Then,

$$\begin{aligned} \|V_x\| &= \sqrt{\sum_i (1 - 2(x_i + \alpha_i)\omega_i)^2} \\ &\leqslant \sqrt{\sum_i (1 + 4(x_i + \alpha_i)^2\omega_i^2)} \leqslant \sqrt{d + 4(R + \|\alpha\|_{\infty})^2 \|w\|^2} \\ &\leqslant \sqrt{d} + 2(R + \|\alpha\|_{\infty}) \|w\| \stackrel{\text{def.}}{=} \bar{V} \end{aligned}$$

We have the following bounds:

$$\begin{aligned} |\varphi_{\omega}(x)| &\leqslant \prod_{i=1}^{d} \sqrt{1 + \frac{x_i}{\alpha_i}} \leqslant \left(1 + \frac{R}{\min_i \alpha_i}\right)^d \stackrel{\text{def.}}{=} \bar{L}_0\\ \mathfrak{g}_x^{-\frac{1}{2}} \nabla \varphi_{\omega}(x) &= \varphi_{\omega}(x) V_x \implies \|\mathbf{D}_1 \left[\varphi_{\omega}\right](x)\|_x \leqslant \bar{L}_0 \bar{V} \end{aligned}$$

and

$$\mathfrak{g}_{x}^{-\frac{1}{2}}\mathrm{H}\varphi_{\omega}(x)\mathfrak{g}_{x}^{-\frac{1}{2}} = \mathfrak{g}_{x}^{-\frac{1}{2}}\nabla^{2}\varphi_{\omega}(x)\mathfrak{g}_{x}^{-\frac{1}{2}} + \mathrm{diag}\left(\mathfrak{g}_{x}^{-\frac{1}{2}}\varphi_{\omega}(x)\right)$$
$$= \varphi_{\omega}(x)(V_{x}V_{x}^{\top} - 2\mathrm{Id}) + \varphi_{\omega}(x)\mathrm{diag}(V_{x}).$$

which yields $\|D_2[\varphi_{\omega}](x)\|_x \leq \overline{L}_0(2+\overline{V}^2)$. Note that by the mean value theorem, $|x_i - x'_i| \leq (R + \alpha_i) |\log(x_i + \alpha_i) - \log(x'_i + \alpha_i)|$ and hence,

$$\|V_{x} - V_{x'}\|_{2} \leq 2 \|\omega\|_{2} \|x - x'\|_{2} \leq 2 \|\omega\|_{2} (R + \|\alpha\|_{\infty}) \mathfrak{d}_{\mathfrak{g}}(x, x')$$

Also, $|\varphi_{\omega}(x) - \varphi_{\omega}(x')| \leq \sup_{x} \|D_1[\varphi_{\omega}](x)\| \mathfrak{d}_{\mathfrak{g}}(x,x') \leq \overline{L}_0 \overline{V} \mathfrak{d}_{\mathfrak{g}}(x,x')$. Therefore,

$$\begin{split} \left\| \mathfrak{g}_{x}^{-\frac{1}{2}} \mathrm{H}\varphi_{\omega}(x) \mathfrak{g}_{x}^{-\frac{1}{2}} - \mathfrak{g}_{x'}^{-\frac{1}{2}} \mathrm{H}\varphi_{\omega}(x') \mathfrak{g}_{x'}^{-\frac{1}{2}} \right\| \\ &\leq \left| \varphi_{\omega}(x) - \varphi_{\omega}(x') \right| \left(2 + \bar{V} + \bar{V}^{2} \right) + \left| \varphi_{\omega}(x') \right| \left\| V_{x} - V_{x'} \right\| + \left| \varphi_{\omega}(x') \right| \left\| V_{x} V_{x}^{\top} - V_{x'} V_{x'}^{\top} \right\| \\ &\leq \bar{L}_{0} \bar{V} \left(2 + \bar{V} + \bar{V}^{2} \right) \mathfrak{d}_{\mathfrak{g}}(x, x') + (\bar{L}_{0} + 2\bar{L}_{0}\bar{V}) 2 \left\| \omega \right\|_{2} \left(R + \left\| \alpha \right\|_{\infty} \right) \mathfrak{d}_{\mathfrak{g}}(x, x') \end{split}$$

Define for j = 0, 1, 2, 3

$$G_{j}(\omega) \stackrel{\text{def.}}{=} \bar{L}_{0} \left(\sqrt{d} + 2(R + \|\alpha\|_{\infty}) \|w\| \right)^{j},$$

then, for $j = 0, 1, 2, L_j(\omega) \stackrel{\text{\tiny def.}}{=} \sup_x \left\| \mathcal{D}_j \left[\varphi_\omega \right](x) \right\|_x \lesssim G_j(\omega)$ and

$$L_{3}(\omega) \stackrel{\text{\tiny def.}}{=} \sup_{x,x'} \frac{\left\| \mathbf{g}_{x}^{-\frac{1}{2}} \mathrm{H} \varphi_{\omega}(x) \mathbf{g}_{x}^{-\frac{1}{2}} - \mathbf{g}_{x'}^{-\frac{1}{2}} \mathrm{H} \varphi_{\omega}(x') \mathbf{g}_{x'}^{-\frac{1}{2}} \right\|}{\mathbf{\mathfrak{d}}_{\mathfrak{g}}(x,x')} \lesssim G_{3}(\omega).$$

When all α_j are distinct, we have [2]:

$$\mathbb{P}(\|\omega\| \geqslant t) \leqslant \mathbb{P}(\|\omega\|_1 \geqslant t) = \sum_{i=1}^d \beta_i e^{-\alpha_i t}$$

where $\beta_i = \prod_{j \neq i} \frac{\alpha_j}{\alpha_j - \alpha_i}$, using the fact that $\|\omega\|_1$ is a sum of independent exponential random variable. Hence, for all $1 \leq j \leq 3$ and $t \geq d^{\frac{j}{2}}$ we have

$$\mathbb{P}(L_j(\omega) \ge t) \le \mathbb{P}\left(\|w\| \ge \frac{1}{2(R+\|\alpha\|_{\infty})} \left(\frac{t}{\bar{L}_0}\right)^{1/j} - \sqrt{d}\right)$$
$$\le F_j(t) \stackrel{\text{def.}}{=} \sum_{i=1}^d \beta_i \exp\left(-\alpha_i \left(\frac{1}{2(R+\|\alpha\|_{\infty})} \left(\frac{t}{\bar{L}_0}\right)^{1/j} - \sqrt{d}\right)\right) \le \delta$$

and $F_j(\bar{L}_j) \leq \delta$ if

$$\bar{L}_j \ge \bar{L}_0 \left(2^j (R + \|\alpha\|_{\infty})^j \left(\sqrt{d} + \max_i \frac{1}{\alpha_i} \log\left(\frac{d\beta_i}{\delta}\right) \right)^j \right)$$

Next, we compute

$$\begin{split} \int_{\bar{L}_{j}}^{\infty} tF_{j}(t)\mathrm{d}t &= \sum_{i=1}^{d} \beta_{i} \int_{\bar{L}_{j}}^{\infty} t \exp\left(-\alpha_{i} \left(\frac{1}{2(R+\|\alpha\|_{\infty})} \left(\frac{t}{\bar{L}_{0}}\right)^{1/j} - \sqrt{d}\right)\right) \mathrm{d}t \\ &= \bar{L}_{0}^{2} j \sum_{i=1}^{d} e^{\alpha_{i}\sqrt{d}} \beta_{i} \int_{(\bar{L}_{j}/\bar{L}_{0})^{1/j}}^{\infty} \exp\left(\frac{-\alpha_{i} u}{2(R+\|\alpha\|_{\infty})}\right) u^{2j-1} \mathrm{d}u \\ &\leqslant \left(\frac{(2j-1)4(R+\|\alpha\|_{\infty})}{e\alpha_{i}}\right)^{2j-1} \bar{L}_{0}^{2} j \sum_{i=1}^{d} e^{\alpha_{i}\sqrt{d}} \beta_{i} \int_{(\bar{L}_{j}/\bar{L}_{0})^{1/j}}^{\infty} \exp\left(\frac{-\alpha_{i} u}{4(R+\|\alpha\|_{\infty})}\right) \mathrm{d}u \\ &\leqslant \left(\frac{4(R+\|\alpha\|_{\infty})}{\alpha_{i}}\right)^{2j} \left(\frac{2j-1}{e}\right)^{2j-1} \bar{L}_{0}^{2} j \sum_{i=1}^{d} e^{\alpha_{i}\sqrt{d}} \beta_{i} \exp\left(\frac{-\alpha_{i}(\bar{L}_{j}/\bar{L}_{0})^{1/j}}{4(R+\|\alpha\|_{\infty})}\right). \end{split}$$

This is bounded from above by δ if for all $i = 1, \ldots, d$,

$$\frac{4(R+\|\alpha\|_{\infty})}{\alpha_i} \left(2j \log\left(\frac{4(2j-1)(R+\|\alpha\|_{\infty})}{e\alpha_i}\right) + \log(\bar{L}_0^2 j) + \alpha_i \sqrt{d} + \log\left(\frac{d\beta_i}{\delta}\right)\right) \leqslant \left(\frac{\bar{L}_j}{\bar{L}_0}\right)^{1/j}$$

that is,

$$\bar{L}_j \gtrsim \bar{L}_0 \left(2^j (R + \|\alpha\|_{\infty})^j \left(\sqrt{d} + \max_i \frac{1}{\alpha_i} \log\left(\frac{d\beta_i}{\delta}\right) \right)^j \right).$$

It remains to bound $\bar{L}_j F_\ell(\bar{L}_\ell)$ with $\ell, j \in \{0, 1, 2, 3\}$: Let $\bar{L}_\ell \ge \bar{L}_0 M^\ell$ for some M to be determined. Then,

$$\begin{split} \bar{L}_{j}F_{\ell}(\bar{L}_{\ell}) &\leqslant \bar{L}_{0}M^{j}\sum_{i=1}^{d}\beta_{i}\exp\left(\frac{-\alpha_{i}}{2(R+\|\alpha\|_{\infty})}M+\alpha_{i}\sqrt{d}\right) \\ &= \bar{L}_{0}\sum_{i=1}^{d}\beta_{i}M^{j}\exp\left(\frac{-\alpha_{i}}{4(R+\|\alpha\|_{\infty})}M\right)\exp\left(\frac{-\alpha_{i}}{4(R+\|\alpha\|_{\infty})}M\right)e^{\alpha_{i}\sqrt{d}} \\ &\leqslant \bar{L}_{0}e^{-j}\sum_{i=1}^{d}\left(\frac{4j(R+\|\alpha\|_{\infty})}{\alpha_{i}}\right)^{j}\beta_{i}\exp\left(\frac{-\alpha_{i}}{4(R+\|\alpha\|_{\infty})}M\right)e^{\alpha_{i}\sqrt{d}} \\ &\leqslant \bar{L}_{0}e^{-3}\sum_{i=1}^{d}\left(\frac{12(R+\|\alpha\|_{\infty})}{\alpha_{i}}\right)^{3}\beta_{i}\exp\left(\frac{-\alpha_{i}}{4(R+\|\alpha\|_{\infty})}M\right)e^{\alpha_{i}\sqrt{d}} \leqslant \delta \end{split}$$

if for each $i = 1, \ldots, d$

$$M \ge 4(R + \|\alpha\|_{\infty}) \left(\sqrt{d} + \max_{i} \frac{1}{\alpha_{i}} \log \left(\frac{\bar{L}_{0} d\beta_{i}}{\delta e^{3}} \left(\frac{12(R + \|\alpha\|_{\infty})}{\alpha_{i}} \right)^{3} \right) \right)$$

Therefore, the conclusion follows for $\bar{L}_0 = \left(1 + \frac{R}{\min_i \alpha_i}\right)^d$, and for j = 1, 2, 3,

$$\bar{L}_j \propto \bar{L}_0 (R + \|\alpha\|_{\infty})^j \left(\sqrt{d} + \max_i \frac{1}{\alpha_i} \log\left(\frac{d\beta_i \bar{L}_0 (R + \|\alpha\|_{\infty})}{\delta\alpha_i}\right)\right)^j.$$

	-