



HAL
open science

The basins of attraction of the global minimizers of non-convex inverse problems with low-dimensional models in infinite dimension

Yann Traonmilin, Jean-François Aujol, Arthur Leclaire

► **To cite this version:**

Yann Traonmilin, Jean-François Aujol, Arthur Leclaire. The basins of attraction of the global minimizers of non-convex inverse problems with low-dimensional models in infinite dimension. *Information and Inference*, 2022. hal-02941814v2

HAL Id: hal-02941814

<https://hal.science/hal-02941814v2>

Submitted on 18 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The basins of attraction of the global minimizers of non-convex inverse problems with low-dimensional models in infinite dimension

YANN TRAO NMILIN^{1,*}, JEAN-FRANÇOIS AUJOL¹ AND ARTHUR LECLAIRE¹

¹Univ. Bordeaux, Bordeaux INP, CNRS, IMB, UMR 5251, F-33400 Talence, France.

* Contact author : yann.traonmilin@math.u-bordeaux.fr

[Received on XX XX XXXX; revised on XX XX XXXX; accepted on XX XX XXXX]

Non-convex methods for linear inverse problems with low-dimensional models have emerged as an alternative to convex techniques. We propose a theoretical framework where both finite dimensional and infinite dimensional linear inverse problems can be studied. We show how the size of the basins of attraction of the minimizers of such problems is linked with the number of available measurements. This framework recovers known results about low-rank matrix estimation and off-the-grid sparse spike estimation, and it provides new results for Gaussian mixture estimation from linear measurements.

Keywords: low-dimensional models, non-convex methods, low-rank matrix recovery, off-the-grid sparse recovery, Gaussian mixture model estimation

1. Introduction

Many inverse problems can be modeled as follows. From m noisy linear measurements $y \in \mathbb{C}^m$ defined by a projection on functions $(\alpha_l)_{1 \leq l \leq m}$ (e.g. Fourier measurements):

$$y_l = \langle x_0, \alpha_l \rangle + e_l \quad (1.1)$$

where $e = (e_l)_{1 \leq l \leq m}$ is an additive noise with finite energy, we aim at recovering the unknown x_0 . This model is particularly used for imaging problems where the signal (e.g. a sound, an image, etc) must be recovered from digital measurements. The linear form $x \rightarrow \langle x, \alpha_l \rangle$ typically models the response of the l -th sensor for a signal x . Let \mathcal{D} be a space containing functions used to measure x_0 (e.g. a Banach space of smooth functions in infinite dimension or a set of vectors in finite dimension). The measurement described by Equation (1.1) makes sense for any signal x_0 living in the dual space \mathcal{D}^* of \mathcal{D} . The bracket $\langle \cdot, \cdot \rangle$ is then a duality product between \mathcal{D}^* and \mathcal{D} .

In our framework, the space \mathcal{D}^* is a locally convex topological vector space with weak-* topology (we will recall in Section 2 some tools that are relevant for our study). The measurement process is summarized

$$y = Ax_0 + e, \quad (1.2)$$

where the linear operator A is a weakly-* continuous linear measurement operator from \mathcal{D}^* to \mathbb{C}^m defined, for $l = 1, \dots, m$, by

$$(Ax_0)_l := \langle x_0, \alpha_l \rangle. \quad (1.3)$$

This weak topology is natural for the study of inverse problems in spaces of measures and distributions where many signals can be modeled (e.g. off-the-grid spikes [11]). We will see in particular that the related weak notion of differentiability is sufficient to study descent algorithms that we will consider in this article.

The theory of inverse problems with low-dimensional models has shown that it is possible to recover x_0 when it belongs to a low-dimensional model Σ with the procedure

$$x^* \in \operatorname{argmin}_{x \in \Sigma} \|Ax - y\|_2^2 \quad (1.4)$$

provided A is adequately chosen (e.g. fulfills a restricted isometry property (RIP) on the secant set $\Sigma - \Sigma$ [8], see Section 2.3). The estimation method (1.4) is called an ideal decoder for the considered inverse problem and low-dimensional model. It has been shown in very generic settings that it is possible to build compressive measurement operators having the required restricted isometry property for low-dimensional recovery [18, 22, 23, 30].

In imaging applications, the goal is often to guarantee that x^* is close to x_0 at a given precision. To describe this, we suppose that such guarantees can be described within a Hilbert space $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$ such that $\Sigma \subset \mathcal{H}$ (the Hilbert space assumption could be dropped to a metric space setting in our proofs but all our examples fall within the Hilbert space case). In other words, we want to ensure that the non-convex decoder (1.4) satisfies

$$\|x^* - x_0\|_{\mathcal{H}}^2 \leq C \|e\|_2^2, \quad (1.5)$$

where C is an absolute constant with respect to e and $x_0 \in \Sigma$.

We place ourselves in a context where the number of measurements, either deterministic or random, guarantees that (1.5) is obtained with the non-convex decoder (1.4), under a RIP assumption. The RIP is usually guaranteed by using a sufficient number of measurements with respect to the dimension of the low-dimensional model Σ . This assumption has been a cornerstone of the qualitative study of compressed sensing, sparse recovery [21] and general inverse problems with low-dimensional models [36].

Even if the decoder (1.4) is guaranteed to recover x_0 (up to the noise level), it is in general not convex and thus difficult to evaluate. To cope with that, one can try to find a convex regularized minimization problem with similar recovery guarantees. Although very successful in some examples (sparse recovery in finite dimension), this approach leads to algorithms that have computational scaling problems in some other examples (off-the-grid sparse recovery). Another general difficulty is the choice of the right convex regularization given a low-dimensional model [37, 38]. Another approach is to directly perform optimization (1.4) with a simple initialization followed by a descent algorithm procedure. This non-convex approach has been proposed for low-rank matrix factorization [5, 42], blind deconvolution [29], blind calibration [10], phase recovery [41] and off-the-grid sparse spike estimation [34, 35].

We propose a unified framework that follows the same idea. We consider inverse problems where the low-dimensional model can be parametrized by \mathbb{R}^d and we propose a general study of gradient descent in the parameter space (that can be easily extended to other descent algorithms).

1.1 Parametrization of the model set Σ

Let the low-dimensional model set $\Sigma \subset \mathcal{D}^*$ be a cone (an assumption we make throughout the whole article) and $x_0 \in \Sigma$. Cones are positively homogenous sets, i.e. for any $x \in \Sigma$ and $\lambda > 0$, $\lambda x \in \Sigma$.

We consider a particular (yet wide) class of inverse problems, where the low-dimensional model can be described by a (possibly constrained) parametrization in \mathbb{R}^d .

DEFINITION 1.1 (Parametrization of Σ) A parametrization of Σ is a function $\phi : \mathbb{R}^d \rightarrow \mathcal{D}^*$ such that $\Sigma \subset \phi(\mathbb{R}^d) = \{\phi(\theta) : \theta \in \mathbb{R}^d\}$.

Our goal is to study the optimization problem (1.4) in the parameter space.

DEFINITION 1.2 (Local minimum) The point $\theta \in \mathbb{R}^d$ is a local minimum of $g : \mathbb{R}^d \rightarrow \mathbb{R}$ if there is $\varepsilon > 0$ such that for any $\theta' \in \mathbb{R}^d$ such that $\|\theta - \theta'\|_2 \leq \varepsilon$, we have $g(\theta) \leq g(\theta')$.

We define the reciprocal image of Σ in the parameter space as

$$\Theta := \phi^{-1}(\Sigma) \quad (1.6)$$

and the parametrized functional

$$g(\theta) := \|A\phi(\theta) - y\|_2^2. \quad (1.7)$$

We consider the problem

$$\theta^* \in \arg \min_{\theta \in \Theta} \|A\phi(\theta) - y\|_2^2. \quad (1.8)$$

As we study descent algorithms in \mathbb{R}^d , we suppose in this article that Θ is an open set of \mathbb{R}^d . This guarantees that the gradient of g is 0 at θ^* even when $\Theta \subsetneq \mathbb{R}^d$ and that $\phi(\theta^*)$ is a minimizer of (1.4).

The model we have just described encompasses the following situations that will be studied in details within our framework in Section 3.

- Low-rank symmetric positive semi-definite (PSD) matrix estimation. We set $\mathcal{D} = \mathcal{D}^* = \mathbb{R}^{p \times p}$, $d = p \times r$, $\phi(Z) = ZZ^T$ (Θ is identified with the set of $p \times r$ matrices and $\Sigma = \Sigma_r$ is the set of PSD matrices of rank lower than r), see Section 3.1.
- Sparse off-the-grid estimation. The space \mathcal{D} is the set $\mathcal{C}_b^2(\mathbb{R}^p)$ of twice-differentiable bounded functions on \mathbb{R}^p with bounded derivatives. The space \mathcal{D}^* contains the set of compactly-supported distributions on \mathbb{R}^p of order less than 2. We have that $d = k(p+1)$, $\phi(a, t) = \sum_{i=1}^k a_i \delta_{t_i}$, $\Sigma = \Sigma_{k, \varepsilon}$ the set of ε -separated sums of k spikes, see Section 3.2.
- Gaussian mixture modeling from compressed data set. We have $\mathcal{D} = \mathcal{C}_b^2(\mathbb{R}^p)$ and \mathcal{D}^* contains the space of signed measures over \mathbb{R}^p , $d = k(\frac{p(p+1)}{2} + p + 1)$, $\phi(w, t, \Gamma) = \sum_{i=1}^k w_i \mu_{t_i, \Gamma_i}$ where $\mu_{t, \Gamma}$ is the Gaussian distribution with mean t and covariances $\Gamma = (\Gamma_1, \dots, \Gamma_k)$. The set $\Sigma = \Sigma_{k, \varepsilon, \rho, P}$ is the set of ε -separated (with respect to an appropriate metric) sums of k Gaussian distributions with eigenvalues of covariances bounded in (ρ, P) . See Section 3.3 for the study of Gaussian mixture models (GMM) with fixed covariance.

In order to link the gradient of g with properties of A , we must be able to apply a chain rule that uses the derivatives of ϕ . The weak-* topology permits to define such derivatives in a weak sense. The important features needed for ϕ will be to follow this weak differentiability assumption (see Section 2), a local Lipschitz behavior around the global minimum and local boundedness properties on its derivatives.

Note that ϕ is not injective in general. A consequence is that the differential of ϕ might have a non trivial kernel. This requires to adapt conventional convergence proofs to this generic setting (see Section 2.4).

1.2 Basin of attraction and descent algorithms

To perform the minimization (1.8), we consider the gradient descent with fixed step τ

$$\theta_{n+1} = \theta_n - \tau \nabla g(\theta_n) \quad (1.9)$$

where $\theta_0 \in \mathbb{R}^d$ is the initialization. Note that any descent algorithm could benefit from our framework (e.g. one can consider block coordinate descent to deal with the case of Gaussian mixtures with variable covariances). We choose the fixed step gradient descent for the simplicity of the analysis.

As only recovery of $x^* \in \Sigma$ matters to us, we propose the following definition of basin of attraction as we will work under conditions where any minimizer of g will lead to recovery guarantees (1.5).

DEFINITION 1.3 (Basin of attraction) We say that a set $\Lambda \subset \mathbb{R}^d$ is a g -basin of attraction of $\theta^* \in \Lambda$ if there exists $\tau_0 > 0$ such that for any $\tau \in (0, \tau_0]$, if $\theta_0 \in \Lambda$, then the sequence $g(\theta_n)$ with θ_n defined by (1.9) converges to $g(\theta^*)$.

This notion of basin of attraction is specific to this work in order to manage the potential indeterminacies of the parametrization. In terms of performance of the estimation, for any initialization in a g -basin of attraction of θ^* , we will have (Corollary 2.2)

$$\|\phi(\theta_n) - x_0\|_{\mathcal{H}}^2 \leq C\|e\|_2^2 + \mathcal{O}\left(\frac{1}{n}\right). \quad (1.10)$$

In other words, the gradient descent leads to an estimation of x_0 that verifies the recovery guarantee (1.5) with respect to the norm $\|\cdot\|_{\mathcal{H}}$ that was chosen to quantify the estimation performance of minimization (1.4), provided the initialization is in a basin of attraction of θ^* .

Following classical optimization results (see e.g. [3, 15]), an open set Λ containing θ^* is a g -basin of attraction of a global minimizer θ^* if

- g is differentiable with Lipschitz gradient;
- g is convex on Λ ;
- for all n , $\theta_n \in \Lambda$.

To deal with indeterminacy, we will show in our main theorem that the convexity property is only needed in relevant directions chosen between the current point and the closest minimizer (within equivalent parametrizations).

Finding a good initialization is a difficult problem that was solved practically and theoretically in selected applications such as phase recovery, and only practically for some others (e.g. of the grid super-resolution in 2D). We discuss possible leads for the systematic study of this problem in Section 4. Also, note that we focus on *convergence* of descent algorithms in this article, thus leading to the convergence speed (1.10). The study of faster convergence (i.e. geometric convergence) typically requires strong convexity, and thus would require a notion of “strong basin of attraction”, which is left for future work.

1.3 Related work

This work unifies recent results on descent algorithms for non-convex optimization for inverse problems with low-dimensional models in both finite and infinite dimension, such as low-rank matrix recovery, phase recovery, blind deconvolution (whose common properties are highlighted in [13]) and off-the-grid sparse spike estimation [34, 35]. We choose the point of view of optimization in the parameter space to keep things as simple as possible from a practical perspective. In [2], conditions on constraints in the parameter space were given in addition to a local convexity hypothesis on the studied functional to guarantee the success of projected gradient descent. We systematically link the measurement operator properties and the properties of the parametrization function to give explicit basins of attraction for

simple gradient descent. This requires to work in infinite dimensional spaces where the low dimensional model can live. The authors of [6] give properties of measurement operators that are sufficient for the success of iterative projections (which can be seen as a projected gradient descent) and an application for infinite union of subspaces (symmetric cones in real vector spaces). While of great interest in some of our examples (such as sparse spike estimation [35]), performing the projection step is in general not trivial, thus motivating the study of the simpler gradient descent.

Another approach is to define the descent algorithm directly on the manifold $\Sigma \subset \mathcal{D}^*$ [7] or even a lifted version of \mathcal{D}^* [14]. The main difficulty with this approach is to define the gradient on the manifold, since the tangent space of Σ might not stay in a “natural” ambient space. For example, for the case of recovery of separated Diracs on the space of measures, the “tangent space” includes distributions of order 1 (that is, distributions that only involve derivatives of order ≤ 1 of the test functions, see [24]) which are not measures. We define a minimal framework starting from the measurement process that allows us to study the non-convex optimization method (1.4). Such minimal structures for regularized inverse problems in Banach spaces have been mentioned in the case of off-the-grid spike recovery [17] and they have been studied precisely in [40]. As no particular metric is needed for the recovery process (only for measuring the success of recovery) on \mathcal{D}^* , we can give our result with only the weak-* topology on \mathcal{D}^* and the norm used to quantify estimation errors in Σ .

1.4 Contributions

We aim at giving a unified understanding at non-convex inverse problems with low-dimensional models frequently found in signal processing and machine learning in finite and infinite dimension. For low-dimensional models which can be parametrized in \mathbb{R}^d :

- We give a minimal framework where the gradient descent in the space of parameters can be described.
- We describe how basins of attraction of the global minimum can be studied, and how their size can be linked to the number of measurements in general under regularity conditions on the parametrization functional. This study is summarized by the general Theorem 2.8 and its Corollary 2.1 (which is used for our examples).
- We describe how this framework can be applied to the examples of low-rank matrix recovery and off-the-grid spike super-resolution, and we give new results for the estimation of Gaussian mixture models (such results were never given in the GMM case to the best of our knowledge).
- We present the general initialization technique by backprojection within our framework and we discuss its practical difficulties.

2. Explicit basins of attraction of the global minimizers

We define precisely a framework where commonly encountered linear inverse problems can be studied. In this framework, we can study the non-convex minimization problem (1.8). In particular, we give conditions which guarantee that explicit basins of attraction of the global minimizers of the function g can be given. Notations used in the article are summarized in Section 6.1.

2.1 Definitions

In our motivating examples, \mathcal{D} is a Banach space.

DEFINITION 2.1 Let $x_n \in \mathcal{D}^*$. The sequence $(x_n)_n$ converges to $x \in \mathcal{D}^*$ for the weak-* topology if for all $\alpha \in \mathcal{D}$

$$\langle x_n, \alpha \rangle \rightarrow \langle x, \alpha \rangle. \quad (2.1)$$

In this case, we denote $x_n \xrightarrow{*} x$.

By construction of \mathcal{D}^* and its weak-* topology, the operator A defined by (1.3) is a linear weak-* continuous operator over \mathcal{D}^* , which implies that for any $(x_n)_n$ such that $x_n \xrightarrow{*} x$, $Ax_n \rightarrow Ax$ (see Section 6.2 for the precise definition of weak-* continuity).

In selected examples, the considered objects are generally not Fréchet differentiable. We thus use the notion of weak-* Gateaux differentiability [19], which is based on directional derivatives.

DEFINITION 2.2 (Differential, directional derivative) In \mathcal{D}^* , a map $\phi : \mathbb{R}^d \rightarrow \mathcal{D}^*$ is weak-* Gateaux differentiable at θ if there exists a linear map $L_\theta(\phi) : \mathbb{R}^d \rightarrow \mathcal{D}^*$ such that for all $u \in \mathbb{R}^d$,

$$\frac{\phi(\theta + hu) - \phi(\theta)}{h} \xrightarrow[h \rightarrow 0]{*} L_\theta(\phi)u \quad (2.2)$$

We write $\partial_u \phi(\theta) = L_\theta(\phi)u$, and $\frac{\partial \phi(\theta)}{\partial \theta_i}$ the derivative in the direction of the i -th canonical vector of \mathbb{R}^d .

In the following sections, we shall assume that ϕ is twice weak-* Gateaux differentiable, i.e. ϕ is weak-* Gateaux differentiable and for any u , $\partial_u \phi$ is weak-* Gateaux differentiable. Note that we will not suppose weak-* continuity of the derivatives in our analysis.

We summarize in Figure 1 the objects and structures used in this article.

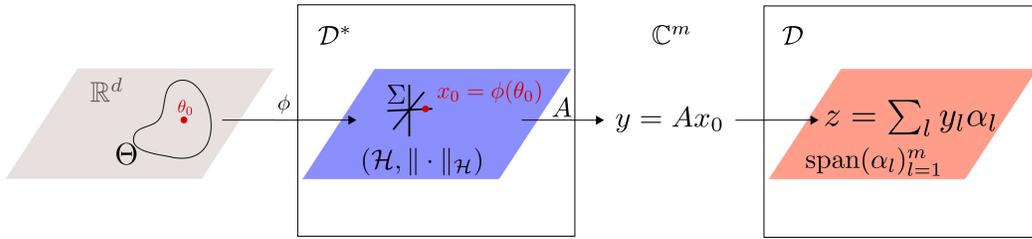


FIG. 1. A summary of the objects and structures of the framework. The cone Σ is a low-dimensional model set parametrized by Θ . The measurements $y = Ax_0$ can be projected in \mathcal{D} for initialization purpose (the ideal backprojection initialization z is described in Section 4).

2.2 Gradient and Hessian of the objective function

We calculate the gradient and Hessian of the function g (defined in (1.7)) in the two following propositions.

PROPOSITION 2.3 Let A be a linear weak-* continuous operator from \mathcal{D}^* to \mathbb{C}^m and ϕ a weak-* Gateaux differentiable function. Then for any $\theta \in \mathbb{R}^d$, the function g is Gateaux differentiable at θ and

$$\frac{\partial g(\theta)}{\partial \theta_i} = 2\Re e \langle A \frac{\partial \phi(\theta)}{\partial \theta_i}, A\phi(\theta) - y \rangle. \quad (2.3)$$

In the following we will denote $\nabla g(\theta) = \left(\frac{\partial g(\theta)}{\partial \theta_i} \right)_{1 \leq i \leq d}$ the Gateaux gradient.

Proof. See Section 6.3. \square

PROPOSITION 2.4 Let A be a linear weak-* continuous operator from \mathcal{D}^* to \mathbb{C}^m and ϕ a twice weak-* Gateaux differentiable function.

For any $\theta \in \mathbb{R}^d$, g is twice Gateaux differentiable at θ and

$$H_{i,j} := \frac{\partial^2 g(\theta)}{\partial \theta_i \partial \theta_j} = G_{i,j} + F_{i,j} \quad (2.4)$$

where

$$G_{i,j} := 2\Re e \left\langle A \frac{\partial \phi(\theta)}{\partial \theta_i}, A \frac{\partial \phi(\theta)}{\partial \theta_j} \right\rangle \quad (2.5)$$

and

$$F_{i,j} := 2\Re e \left\langle A \frac{\partial^2 \phi(\theta)}{\partial \theta_i \partial \theta_j}, A \phi(\theta) - y \right\rangle. \quad (2.6)$$

Proof. See Section 6.3. \square

2.3 Secant sets and the RIP

The following definitions allow to express the restricted isometry property. We then provide a fundamental lemma useful to make the connection between the RIP and the Hessian.

DEFINITION 2.5 (Secant) The secant set of the model set Σ is $\mathcal{S}(\Sigma) = \Sigma - \Sigma := \{x - y : x \in \Sigma, y \in \Sigma\}$. A secant is an element of the secant set.

DEFINITION 2.6 (Generalized secant) Suppose ϕ is weak-* Gateaux differentiable. A generalized secant is either a secant or a directional derivative $\partial_u \phi(\theta)$ with $\phi(\theta) \in \Sigma$. The generalized secant set $\overline{\mathcal{S}(\Sigma)}$ is the set of generalized secants.

In the context of manifolds, the generalized secant set is linked with the tangent space of the manifold Σ : it contains the directional derivatives of elements of Σ (with respect to their parametrization) because a tangent vector is a limit of secants. We suppose the existence of a Hilbert space $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$ such that $\Sigma \subset \mathcal{H}$. The following assumption will be needed.

ASSUMPTION 2.1 (Compatibility of $\|\cdot\|_{\mathcal{H}}$ with generalized secants) For all $x \in \overline{\mathcal{S}(\Sigma)}$ such that $x = \partial_u \phi(\theta)$, we have that $\left\| \frac{\phi(\theta + |h_n|u) - \phi(\theta)}{|h_n|} \right\|_{\mathcal{H}}$ converges for any $|h_n| \rightarrow 0$ to a limit that does not depend on the choice of the real sequence h_n . This limit is written $\|\partial_u \phi(\theta)\|_{\mathcal{H}}$.

With this assumption, we can extend $\|\cdot\|_{\mathcal{H}}$ to $\overline{\mathcal{S}(\Sigma)}$.

While trivial in finite dimension, the last assumption must be considered carefully in infinite dimension (see Section 3). We now have sufficient tools to define the RIP.

DEFINITION 2.7 (RIP) The operator A has the RIP on $\mathcal{S}(\Sigma)$ with respect to $\|\cdot\|_{\mathcal{H}}$ with constant γ if for all $x \in \mathcal{S}(\Sigma)$

$$(1 - \gamma)\|x\|_{\mathcal{H}}^2 \leq \|Ax\|_2^2 \leq (1 + \gamma)\|x\|_{\mathcal{H}}^2. \quad (2.7)$$

The RIP is very useful for the qualitative study of inverse imaging problems: measurement operators A are chosen such that the RIP constant γ improves (i.e. decreases) when the number of measurements

increases. In many compressed sensing examples, it can be guaranteed that appropriately chosen random operators A have the RIP with high probability as long as $m \geq O(d \text{polylog}(d))$, i.e. the number of measurements is of the order of the intrinsic dimension d of Σ , with a dependency on the dimension of \mathcal{H} (or the choice of $\|\cdot\|_{\mathcal{H}}$ in infinite dimension) in \log factors. Thanks to the compatibility assumption 2.1, we can extend the RIP to $\overline{\mathcal{S}(\Sigma)}$ which contains the directional derivatives of ϕ .

LEMMA 2.1 (RIP on the generalized secant set) Suppose A has the RIP on $\Sigma - \Sigma$ with constant γ and ϕ is weak-* Gateaux differentiable. Suppose A is weak-* continuous. Suppose that $\|\cdot\|_{\mathcal{H}}$ verifies compatibility assumption 2.1. Let $v \in \mathcal{S}(\Sigma)$ then

$$(1 - \gamma) \|v\|_{\mathcal{H}}^2 \leq \|Av\|_2^2 \leq (1 + \gamma) \|v\|_{\mathcal{H}}^2. \quad (2.8)$$

Proof. See Section 6.4. □

2.4 Indeterminacy of the parametrization

The parametrization function ϕ is not injective in general, leading to an indeterminacy in the parametrization. Theoretical complications appear especially when the set of equivalent parameters $\{\tilde{\theta} : \phi(\tilde{\theta}) = \phi(\theta)\}$ is not a set of isolated points, e.g. in the low-rank matrix recovery case when the factors can only be recovered up to a multiplication by an orthogonal matrix. While a basin of attraction exists, the function g might even not be locally convex [13] (e.g. in the low-rank matrix recovery case).

To cope with this indeterminacy, we study the Hessian of g in the directions u relevant to the proof of convergence of the gradient descent. To do this, we introduce the following notations. Let θ^* be a global minimizer of g on Θ . We define

$$d(\theta, \theta^*) := \min_{\substack{\tilde{\theta} \in \Theta \\ \phi(\tilde{\theta}) = \phi(\theta^*)}} \|\tilde{\theta} - \theta\|_2, \quad \text{and} \quad p(\theta, \theta^*) := \underset{\substack{\tilde{\theta} \in \Theta \\ \phi(\tilde{\theta}) = \phi(\theta^*)}}{\text{argmin}} \|\tilde{\theta} - \theta\|_2 \subset \Theta. \quad (2.9)$$

We will study basins of attraction having the shape

$$\Lambda_\beta := \{\theta \in \Theta : d(\theta, \theta^*) < \beta\}. \quad (2.10)$$

Notice that $\phi^{-1}(\{\phi(\theta^*)\})$ is a closed subset of Θ when ϕ is weak-* continuous, which allows to define $d(\theta, \theta^*)$ as a minimum. Actually, $d(\theta, \theta^*)$ is the distance to the closed set $\phi^{-1}(\{\phi(\theta^*)\})$ and $p(\theta, \theta^*)$ is the (set-valued) projection of θ on $\phi^{-1}(\{\phi(\theta^*)\})$. However, one should be warned that $d(\cdot, \cdot)$ is not a true distance function. In practice, in our examples, the set $p(\theta, \theta^*)$ is composed of a unique element in the basin of attraction.

This notion of distance plays an important role in the proof of our main result to show the stability of iterates of the gradient descent. The use of a ℓ^2 -based distance between parameters permits to use classical arguments relying on the scalar product.

2.5 Control of the Hessian with the restricted isometry property

We begin by giving a control on the Hessian of g around a minimizer θ^* .

LEMMA 2.2 Suppose A is weak-* continuous and has the RIP on $\Sigma - \Sigma$. Suppose ϕ is twice weak-* Gateaux differentiable. Suppose that $\|\cdot\|_{\mathcal{H}}$ verifies the compatibility assumption 2.1. Let θ^* be a global minimizer of g on the open set Θ . Let $\Lambda \subset \mathbb{R}^d$ be a set such that for all $\theta \in \Lambda$, we have $\phi(\theta) - \phi(\theta^*) \in \mathcal{S}(\Sigma)$. Let $\theta \in \Lambda$ and H be the Hessian of g at θ .

For all $u \in \mathbb{R}^d$, we have

$$u^T H u \geq 2(1 - \gamma) \|\partial_u \phi(\theta)\|_{\mathcal{H}}^2 - 2 \|A \partial_u^2 \phi(\theta)\|_2 (\sqrt{1 + \gamma} \|\phi(\theta) - \phi(\theta^*)\|_{\mathcal{H}} + \|e\|_2) \quad (2.11)$$

$$u^T H u \leq 2 \|A \partial_u \phi(\theta)\|_2^2 + 2 \|A \partial_u^2 \phi(\theta)\|_2 (\sqrt{1 + \gamma} \|\phi(\theta) - \phi(\theta^*)\|_{\mathcal{H}} + \|e\|_2) \quad (2.12)$$

Proof. See Section 6.5. \square

It is possible to control the Hessian of g on a set Λ with Lemma 2.2 in the directions which are relevant to guarantee convergence. For $\theta_1, \theta_2 \in \mathbb{R}^d$, let us define the line segment $[\theta_1, \theta_2] = \{t\theta_1 + (1 - t)\theta_2 : t \in [0, 1]\}$. We gather technical hypotheses in the following assumption.

ASSUMPTION 2.2 (Technical assumption on ϕ and radius β) Given $\theta^* \in \Theta$, $\beta > 0$, we say that the technical assumption on ϕ and on radius β are fulfilled with constants C_{ϕ, θ^*} , M_1 and M_2 if

1. $\theta \in \Lambda_{2\beta}$ implies $\phi(\theta) \in \Sigma$ (local stability of the model set);
2. there is $C_{\phi, \theta^*} > 0$ such that

$$\forall \theta \in \Lambda_{2\beta}, \quad \|\phi(\theta) - \phi(\theta^*)\|_{\mathcal{H}} \leq C_{\phi, \theta^*} d(\theta, \theta^*) \quad (\text{local control of } \|\cdot\|_{\mathcal{H}}); \quad (2.13)$$

3. the first-order derivatives of $A\phi$ are uniformly bounded on $\phi^{-1}(\phi(\theta^*))$:

$$M_1 := \sup_{\theta \in \phi^{-1}(\phi(\theta^*))} \sup_{u: \|u\|_2=1} \|A \partial_u \phi(\theta)\|_2 < +\infty; \quad (2.14)$$

4. the second-order derivatives of $A\phi$ are uniformly bounded on $\Lambda_{2\beta}$:

$$M_2 := \sup_{\theta \in \Lambda_{2\beta}} \sup_{u, v: \|u\|_2=1, \|v\|_2=1} \|A \partial_v \partial_u \phi(\theta)\|_2 < +\infty. \quad (2.15)$$

We propose the following generic theorem to show that a set Λ_β (defined by (2.10)) is a g -basin of attraction.

THEOREM 2.8 Consider the following two set of hypotheses.

Framework hypotheses: Let A be a weak-* continuous linear map from \mathcal{D}^* to \mathbb{C}^m . Suppose A has the RIP on $\mathcal{S}(\Sigma)$ with constant γ and ϕ is weak-* continuous and twice weak-* Gateaux differentiable. Suppose that $\|\cdot\|_{\mathcal{H}}$ verifies the compatibility assumption 2.1. Let θ^* be a global minimizer of g on the open set Θ .

Specific hypotheses: Assume that there exists $\beta > 0$ such that

1. the technical assumption 2.2 on ϕ and on radius β is fulfilled with constants C_{ϕ, θ^*} , M_1 and M_2 ;
2. for any $\theta \in \Lambda_\beta$, there exists $\tilde{\theta} \in p(\theta, \theta^*)$ such that

$$\forall z \in [\theta, \tilde{\theta}], \quad \frac{(1 - \gamma) \|\partial_{\tilde{\theta} - \theta} \phi(z)\|_{\mathcal{H}}^2}{\sqrt{1 + \gamma} \|A \partial_{\tilde{\theta} - \theta}^2 \phi(z)\|_2} \geq C_{\phi, \theta^*} \beta + \frac{1}{\sqrt{1 + \gamma}} \|e\|_2. \quad (2.16)$$

Then Λ_β is a g -basin of attraction of θ^* .

Proof. See Section 6.5. \square

This theorem highlights the regularity properties and the control on the derivatives of ϕ that we require to ensure convergence. However, it does not give an explicit expression of the radius of the basin of attraction at first sight. We propose a corollary that makes it more explicit in the case when the projection $p(\theta, \theta^*)$ is composed of a unique element when θ is in the basin, which is the case in all the examples covered in the next section.

COROLLARY 2.1 Under the framework hypotheses of Theorem 2.8, let $\beta_1 > 0$ such that

1. for any $\theta \in \Lambda_{2\beta_1}$, **there exists a unique** $\tilde{\theta} \in p(\theta, \theta^*)$;
2. the technical assumption 2.2 on ϕ and on radius β_1 is fulfilled with constants C_{ϕ, θ^*} , M_1 and M_2 ;
3. we have

$$\beta_2 := \frac{(1-\gamma)}{C_{\phi, \theta^*} \sqrt{1+\gamma}} \inf_{\theta \in \Lambda_{\beta_1}} \inf_{z \in [\theta, \tilde{\theta}]} \left(\frac{\|\partial_{\tilde{\theta}-\theta} \phi(z)\|_{\mathcal{H}}^2}{\|A \partial_{\tilde{\theta}-\theta}^2 \phi(z)\|_2} \right) - \frac{1}{C_{\phi, \theta^*} \sqrt{1+\gamma}} \|e\|_2 > 0. \quad (2.17)$$

Then $\Lambda_{\min(\beta_1, \beta_2)}$ is a g -basin of attraction of θ^* .

Proof. See Section 6.5. \square

REMARK 2.1 This technique for the study of basins of attraction yields results when we can guarantee

$$\inf_{\theta \in \Lambda_{\beta_1}} \inf_{z \in [\theta, \tilde{\theta}]} \left(\frac{\|\partial_{\tilde{\theta}-\theta} \phi(z)\|_{\mathcal{H}}^2}{\|A \partial_{\tilde{\theta}-\theta}^2 \phi(z)\|_2} \right) > 0 \quad \text{where } \tilde{\theta} \in p(\theta, \theta^*) \text{ is unique.} \quad (2.18)$$

We will see that we can verify this in all our examples. In the low-rank recovery example where indeterminacy causes problems, we control the second order derivatives in the relevant directions $u = \tilde{\theta} - \theta$. For Dirac and Gaussian estimation, we can bound uniformly the Hessian in all directions within the basin of attraction.

REMARK 2.2 The fact that regularity assumptions are on $\Lambda_{2\beta_1}$ instead of Λ_{β_1} is essentially a technical argument to guarantee the stability of the iterates in Λ_{β_1} in a general theorem. It could be reduced to an assumption on $\Lambda_{\beta_1+\eta}$ with η small, by reducing the step size. In our examples (Dirac and Gaussian estimation), it could be reduced to Λ_{β_1} by using a specific convergence proof.

Note that in the noisy case, a small noise assumption (which is linked with the smallest amplitudes in $\phi(\theta^*)$ in practice) guarantees the non-negativity of the Hessian with Lemma 2.2. In the next section we will present the results in the noiseless case for clarity purpose.

From the expression of β_2 (i.e. the size of the basin), we observe a general behavior that was already observed in the case of low-rank matrix recovery and spike estimation: when the RIP constant decreases (i.e. the number of measurements increases), the size of the basin increases (possibly not strictly).

In Hilbert spaces, when $\mathcal{S}(\Sigma)$ has finite dimension d (for the upper-box counting dimension), it was shown in [30] that it is possible to construct a random linear operator that have, with high probability, the RIP on $\mathcal{S}(\Sigma)$ with constant γ such that

$$m = O\left(\frac{d}{\gamma^2}\right). \quad (2.19)$$

With such operators, we can write $\gamma = \frac{1}{c\sqrt{m}}$ where c is a constant independent of m . This gives the explicit dependency on the number of measurements. Note that the constants involved are dependent on the model Σ and typically include log factors (see Section 3).

In all our examples, random operators following a similar relation can be constructed (even for sparse spike reconstruction and GMM estimation [22]).

Finally, we provide the following Corollary to show that the gradient descent leads to a solution that has the right estimation guarantees.

COROLLARY 2.2 Under the framework hypotheses and specific hypotheses of Theorem 2.8, let θ_n be the iterates (constructed in the proof of Theorem 2.8) of a gradient descent with sufficiently small fixed step size such that $g(\theta_n) \rightarrow g(\theta^*)$. Then

$$\|\phi(\theta_n) - x_0\|_{\mathcal{H}}^2 \leq \frac{4}{1-\gamma} \|e\|_2^2 + O\left(\frac{1}{n}\right). \quad (2.20)$$

Proof. See Section 6.5. □

3. Application of the framework

In this section, we apply our framework to three examples. We highlight how it relates to existing results and how it permits to give new ones. While suboptimal for the study of non-convex algorithms in general (e.g. for low-rank matrix recovery), we are able to give new results for non-convex recovery of low-dimensional models where such a study did not exist yet (Gaussian mixture models). For the sake of simplicity we study the three following examples in the noiseless case $\|e\|_2 = 0$.

3.1 Low-rank matrix recovery

The low-rank PSD matrix recovery problem falls into our analysis. We are able to give explicit basins of attraction of global minimizers. Stronger results involving the study of all critical points of the functional show that the global minimum is the only local minimum, thus justifying the use of stochastic descents which escape saddle points for global convergence results (see [13] for a complete overview). We still apply our framework to this case to check its validity against well known results. It is also a first step for the understanding of more complex models such as Gaussian mixture models with low-rank modeling of covariances (see Section 3.3).

The set-up is as follows:

- We measure matrices with projections on \mathbb{R}^m , i.e. $\mathcal{D} = \mathcal{D}^* = \mathbb{R}^{p \times p}$, the duality product is the scalar product associated to the Frobenius norm, which also defines the Hilbert structure. In the context of matrices the Euclidean norm is referred to as the Frobenius norm $\|\cdot\|_{\mathcal{H}} = \|\cdot\|_F$. The associated scalar product is $\langle X, Y \rangle_F = \text{tr}(X^T Y)$.
- The set of rank at most r PSD matrices is $\Sigma = \Sigma_r = \{ZZ^T : Z \in \mathbb{R}^{p \times r}\}$.
- We can simply parametrize the model set Σ with $\phi(\theta) = \phi(Z) = ZZ^T$ where we identified the d -dimensional parameter $\theta \in \mathbb{R}^d$ ($d = pr$) as a matrix $Z \in \mathbb{R}^{p \times r}$. Note that given Z_1, Z_2 such that $\phi(Z_1) = Z_1 Z_1^T = Z_2 Z_2^T = \phi(Z_2)$, then there exists an orthogonal matrix H such that $Z_2 H = Z_1$.
- Low-rank matrix recovery results (e.g. [12]) show that an appropriately chosen random A has a RIP on Σ_{2r} with constant γ with high probability provided $m \geq O\left(\frac{pr}{\gamma^2}\right)$, hence $\phi(\theta^*) = x_0$ (we set $x_0 = Z_0 Z_0^T$).

For low-rank matrix recovery, the parametrized minimization (1.8) is written

$$\min_{Z \in \mathbb{R}^{p \times r}} \|A(ZZ^T) - y\|_2^2. \quad (3.1)$$

This minimization is often called the Burer-Monteiro method [9].

To apply our framework, we calculate $\partial_U \phi(Z)$ the directional derivative in the direction $U \in \mathbb{R}^{p \times r}$. We have (see Section 6.6):

$$\begin{aligned} \partial_U \phi(Z) &= UZ^T + ZU^T \\ \partial_U^2 \phi(Z) &= 2UU^T. \end{aligned} \quad (3.2)$$

We see that the directional derivative $\partial_U \phi(Z)$ is 0 for (non-trivial) U such that $UZ^T + ZU^T = 0$. In the study of non-convex low-rank matrix recovery, most complications arise from this indeterminacy of the parametrization. To give a basin of attraction of $\theta^* = Z_0$, it is shown in the literature that the interesting directions of the Hessian are the solutions of an orthogonal Procrustes problem, i.e. it is sufficient to lower bound the Hessian in all directions $U \in \mathbb{R}^{p \times r}$ that can be written $U = Z'H_0 - Z_0$ with an arbitrary Z' and with $H_0 = \operatorname{argmin}_{H \in \mathcal{O}(r)} \|Z'H - Z_0\|_F^2$ to guarantee the success of the gradient descent [13] ($\mathcal{O}(r)$ is the orthogonal matrix group). Qualitatively, in the parameter space, the important directions of the Hessian are those between the unknown Z_0 and matrices Z minimizing the Frobenius distance to Z_0 within its class of equivalent parametrizations $\{\tilde{Z} : \tilde{Z}\tilde{Z}^T = ZZ^T\}$.

This idea can be used within our general framework to recover a g -basin of attraction of Z_0 . In our case we will instead study the Hessian in the direction $U = Z_0H_0 - Z$ where $H_0 = \operatorname{argmin}_{H \in \mathcal{O}(r)} \|Z_0H - Z\|_F^2$ for $Z \in \Lambda_\beta$. While we cannot expect the function g to be convex in a neighborhood of Z_0 [13], there is an underlying convexity property in these particular directions which allows to use Corollary 2.1. Within our framework, we have $d(Z, Z_0) = \|Z_0H_0 - Z\|_F^2$ and $p(Z, Z_0) = Z_0H_0$.

THEOREM 3.1 Let A be a linear map on $\mathbb{R}^{p \times p}$ with the RIP on Σ_{2r} with constant γ . Suppose $e = 0$ and let Z_0 be a rank- r global minimizer of (3.1). Let $\beta_{LR} := \frac{1}{8} \frac{(1-\gamma)(\sigma_{\min}(Z_0))^2}{(1+\gamma)\sigma_{\max}(Z_0)}$. Then $\Lambda_{\beta_{LR}} := \{Z : \inf_{H \in \mathcal{O}(r)} \|ZH - Z_0\|_F < \beta_{LR}\}$ is a g -basin of attraction of Z_0 .

Proof. See Section 6.6. \square

Our theorem gives a result that is similar to the one of [39]: the size of our basin of attraction depends on the smallest singular value of $\phi(\theta^*)$. It also shows explicitly the dependency on the number of measurements through the RIP constant. Our result has an added dependency on the conditioning of Z_0 . In return, we require the RIP only on Σ_{2r} , instead of Σ_{6r} .

REMARK 3.1 The descent algorithm for low-rank matrix recovery coined as Procrustes flow in the literature uses the expression of the gradient for measurements α_l obtained from symmetric measurement matrices A . In this case

$$\partial_U g(Z) = 2\Re e \langle A(UZ^T + ZU^T), A\phi(Z) - y \rangle = 4\Re e \langle AUZ^T, A\phi(Z) - y \rangle \quad (3.3)$$

Our analysis uses the true value of the gradient for any measurement operator A .

3.2 Off-the-grid sparse spike recovery

Off-the-grid sparse spike recovery is at the core of imaging problems in signal processing [11, 16, 17]. They can also be used to perform some machine learning tasks such as compressive clustering [26].

The size of basins of attraction is directly linked with the number of measurements through RIP constants [34]. The proof of the results of [34] is exactly the proof of Lemma 2.2 coupled with controls of the chosen norm $\|\cdot\|_{\mathcal{H}}$ and the explicit computation of the gradient and Hessian. We recall here the set-up leading to explicit basins of attraction for this specific case.

- The off-the-grid sparse signals supported on \mathbb{R}^p are measured by projections on twice differentiable functions with bounded derivatives (weighted Fourier measurements) $\alpha_l \in \mathcal{C}_b^2(\mathbb{R}^p) = \mathcal{D}$. Hence \mathcal{D}^* contains the set of compactly-supported distributions of order ≤ 2 on \mathbb{R}^p .
- The low-dimensional model is the subset of finite signed measures over \mathbb{R}^p defined by $\Sigma = \Sigma_{k,\varepsilon} := \{\sum_{i=1}^k a_i \delta_{t_i} : \|t_i - t_j\|_2 > \varepsilon, t_i \in \mathcal{B}_2(R)\}$, where $\mathcal{B}_2(R) := \{t \in \mathbb{R}^p : \|t\|_2 < R\}$.
- The parametrization function is defined for $\theta = (a_1, \dots, a_k, t_1, \dots, t_k) \in \mathbb{R}^{k(p+1)}$ by $\phi(\theta) = \sum_{i=1}^k a_i \delta_{t_i}$ ($a = (a_1, \dots, a_k)$ is the vector of amplitudes, $t = (t_1, \dots, t_k)$ defines the k positions in \mathbb{R}^p). Note that any parametrization is equivalent up to a permutation of the positions and amplitudes.
- In this case, the minimization (1.8) must be performed on the constrained set

$$\Theta_{k,\varepsilon} = \{(a_1, \dots, a_k, t_1, \dots, t_k) : a_i \in \mathbb{R}, t_i \in \mathcal{B}_2(R), \|t_i - t_j\|_2 > \varepsilon\} \quad (3.4)$$

The fact that we place ourselves within the hypotheses of Theorem 2.8 permits to guarantee that the gradient descent iterates stay in the constraint.

- The Hilbert norm for Σ can be induced by a kernel metric defined on the space of finite signed measures over \mathbb{R}^d and extended to distributions of order 2 of interest. Such a kernel metric takes the following form on a linear combination of Dirac masses

$$\begin{aligned} \left\| \sum_i a_i \delta_{t_i} \right\|_{\mathcal{H}}^2 &= \left\| \sum_i a_i \delta_{t_i} \right\|_K^2 := \int K(t, s) \left(\sum_i a_i d\delta_{t_i}(t) \right) \left(\sum_j a_j d\delta_{t_j}(s) \right) \\ &= \sum_i a_i a_j K(t_i - t_j) \end{aligned} \quad (3.5)$$

where $K(t, s) \propto e^{-\frac{\|t-s\|_2^2}{2\sigma^2}}$ is a Gaussian kernel with a variance σ^2 that defines the precision at which we measure distances between elements of Σ . It was shown in [34] that this kernel verifies the compatibility assumption 2.1.

- Either random or regular Fourier measurements A over \mathbb{R}^d can be considered. They have been shown to have a RIP on $\mathcal{S}(\Sigma_{k,\varepsilon})$ with constant γ with respect to the kernel metric $\|\cdot\|_{\mathcal{H}}$ as long as $m = O(\frac{k^2 d}{\gamma^2} (\log(k))^2 \log(kd/\varepsilon))$ (random Gaussian Fourier measurements) or $m = O(\frac{1}{\varepsilon^d})$ (regular Fourier measurements).

For off-the-grid sparse spike estimation, the parametrized minimization (1.8) is written

$$\min_{a_1, \dots, a_k \in \mathbb{R}; t_1, \dots, t_k \in \mathcal{B}_2(R); \forall i \neq j, \|t_i - t_j\|_2 > \varepsilon} \left\| A \left(\sum_{i=1}^k a_i \delta_{t_i} \right) - y \right\|_2^2. \quad (3.6)$$

The derivatives of ϕ are given by $\partial_u \phi(a, t) = \sum_i v_i \delta_{t_i} + a_i \partial_{w_i} \delta_{t_i}$ where $u = (v, w)$ (v is the direction for the derivative with respect to amplitudes and w is the direction for the derivative with respect

to positions) and $\partial_{w_i} \delta_{t_i}$ is a directional derivative of the Dirac in the distribution sense: for $\alpha \in \mathcal{D}$, $\langle \partial_{w_i} \delta_{t_i}, \alpha \rangle = -\partial_{w_i} \alpha(t_i)$.

Within this framework, the case of Dirac recovery is the one with the most complications as elements of the generalized secant set $\overline{\mathcal{S}(\Sigma)}$ are not elements of the set of finite signed measures which is naturally considered for off-the-grid Dirac recovery. In consequence, the considered kernel metric $\|\cdot\|_{\mathcal{H}}$ which is defined on the space of finite signed measures, must be extended to elements of $\overline{\mathcal{S}(\Sigma)}$ which are distributions of order 1 (the Dirac derivatives). Such an extension uses the smoothness of the Gaussian kernel and the fact that $\overline{\mathcal{S}(\Sigma)}$ acts on a bounded domain.

We obtain the following theorem in the noiseless case.

THEOREM 3.2 Suppose A has RIP with constant γ on $\mathcal{S}(\Sigma_{k, \frac{\varepsilon}{2}})$.

Let $\theta^* = (a_1, \dots, a_k, t_1, \dots, t_k) \in \Theta_{k, \varepsilon}$ be a result of constrained minimization (3.6) such that $0 < |a_1| \leq |a_2| \dots \leq |a_k|$. Then there is an explicit β_{spikes} depending on the a_i, γ, K, A such that $\Lambda_{\beta_{spikes}} := \{\theta : \|\theta - \theta^*\|_2 < \beta_{spikes}\}$ is a g -basin of attraction of θ^* .

We refer to [34, Corollary 3.1] for the proof and precise value of β_{spikes} . The value of β_{spikes} exhibits the behavior with respect to the RIP constant mentioned in Section 2.5. It also shows a dependency on the smallest amplitudes in θ^* . The strong RIP assumption (i.e with separation $\frac{\varepsilon}{2}$ instead of ε) is used to guarantee that the first condition of Theorem 2.8 is met (the same argument is used in the next section for GMM).

3.3 Gaussian mixture estimation from compressive measurements

Gaussian mixture model recovery from linear measurements can be used to model blind deconvolution problems. Additionally, the estimation of GMM from compressive measurements can be used to perform compressive statistical learning. Consider a database $x_1, \dots, x_n \in \mathbb{R}^p$ and represent it with its associated empirical probability measure $x = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$. Let A be a random Fourier measurement operator. For $m < np$, we obtain a compressed version $y = Ax$ of x called a sketch of the dataset (using random Fourier features). It was shown that the parameters of a Gaussian mixture (GMM) can be recovered by solving the minimization (1.4) with the appropriate design of A [22, 23, 25]. Recovery guarantees have been given for the case of a fixed known covariance PSD matrix Γ while practical results including the estimation of diagonal covariances were obtained using a heuristic based on orthogonal matching pursuit [25]. In such greedy methods, an unconstrained gradient descent step is used to refine the solution within the algorithm. We give the expression of an explicit basin of attraction in the fixed covariance case, which gives an understanding of the success of such descent algorithms and we discuss how the result can be extended to the case of variable covariance afterwards.

- The Gaussian mixtures on \mathbb{R}^d are measured by projections on bounded functions (Fourier measurements) $\alpha_i \in \mathcal{C}_b(\mathbb{R}^p) = \mathcal{D}$. Hence \mathcal{D}^* contains the set of finite signed measures on \mathbb{R}^p .
- The low-dimensional model is a subset of finite signed measures over \mathbb{R}^p defined by $\Sigma = \Sigma_{k, \varepsilon, \Gamma} := \{\sum_{i=1}^k a_i \mu_{t_i} : \|t_i - t_j\|_{\Gamma} > \varepsilon, t_i \in \mathcal{B}_2(R)\}$, where $d\mu_{t_i}(t) = e^{-\frac{1}{2}\|t-t_i\|_{\Gamma}^2} dt$, $\|u\|_{\Gamma}^2 = u^T \Gamma^{-1} u$ and Γ is the fixed known covariance matrix.
- The parametrization function is defined for $\theta = (a_1, \dots, a_k, t_1, \dots, t_k) \in \mathbb{R}^{k(p+1)}$ by $\phi(\theta) = \sum_{i=1}^k a_i \mu_{t_i}$ ($a = (a_1, \dots, a_n)$ is the vector of amplitudes, $t = (t_1, \dots, t_n)$ defines the means). As in the spike estimation problem, any parametrization is equivalent up to a permutation of the means and amplitudes.

- The minimization (1.8) is performed on the constrained set $\Theta_{k,\varepsilon} \subset \mathbb{R}^{(k+1)d}$.
- The Hilbert norm for Σ can be obtained from a kernel metric $\|\cdot\|_{\mathcal{H}} := \|\cdot\|_K$ defined on the space of finite signed measures over \mathbb{R}^d where the kernel is a Gaussian function with covariance proportional to Γ that defines the precision at which we measure distances between elements of Σ . The compatibility assumption 2.1 essentially comes from the fact that the directional derivatives of Gaussian measures with respect to amplitudes and means are finite signed measures (they are obtained by differentiating their \mathcal{C}^∞ densities).
- Adequately chosen random Fourier measurements with Gaussian frequency measurements have been shown to satisfy the RIP on the secant set $\mathcal{S}(\Sigma_{k,\varepsilon,\Gamma})$ with constant γ with respect to $\|\cdot\|_K$ as long as $\varepsilon \geq O(\sqrt{d(1+\log(k))})$ and $m = O(\frac{k^2 d}{\gamma^2} \text{polylog}(k, d))$ [22].

For GMM estimation, the parametrized minimization (1.8) is written

$$\min_{a_1, \dots, a_k \in \mathbb{R}; t_1, \dots, t_k \in \mathcal{B}_2(R); \forall i \neq j, \|t_i - t_j\|_\Gamma > \varepsilon} \left\| A \left(\sum_{i=1}^k a_i \mu_{t_i} \right) - y \right\|_2^2. \quad (3.7)$$

The derivatives of ϕ are given by $\partial_u \phi(a, t) = \sum_i v_i \mu_{t_i} + a_i \partial_{w_i} \mu_{t_i}$ where $u = (v, w)$ (v is the direction for the derivative with respect to amplitudes and w is the direction for the derivative with respect to means). As μ_{t_i} has a smooth density with respect to the Lebesgue measure its directional derivative $\partial_{w_i} \mu_{t_i}$ in the distribution sense is also a finite signed measure with density

$$t \rightarrow -w_i^T \Gamma^{-1} (t - t_i) e^{-\frac{1}{2} \|t - t_i\|_\Gamma^2}. \quad (3.8)$$

This makes the case of Gaussian mixtures slightly easier to manage than the Dirac recovery case as the kernel metric $\|\cdot\|_K$ is well defined on the space of finite signed measures.

We give an explicit uniform bound of the Hessian on a neighborhood of θ^* in this case. The main missing ingredient to give the bound is the mutual coherence of the kernel metric. It can be shown that an appropriately chosen Gaussian kernel K satisfies the following assumption [23] (by taking a kernel with small enough variance with respect to the separation).

ASSUMPTION 3.1 The kernel K follows this assumption if

- $\|\mu_{t_i}\|_K = 1$.
- There is a constant c_K such that for any k means $(t_i)_{i=1}^k$ that verifies for $i \neq j$, $\|t_i - t_j\|_\Gamma > \varepsilon$, we have $\|\sum_i v_i \mu_{t_i} + a_i \partial_{w_i} \mu_{t_i}\|_K^2 \geq (1 - c_K) \sum_i \|v_i \mu_{t_i} + a_i \partial_{w_i} \mu_{t_i}\|_K^2$.

In the noiseless case we get the following basin of attraction.

THEOREM 3.3 Suppose the Fourier measurement operator A has RIP γ on $\mathcal{S}(\Sigma_{k,\varepsilon,\Gamma})$. Suppose K is a Gaussian kernel with covariance proportional to Γ that follows Assumption 3.1. Suppose $e = 0$. Let $\theta^* = (a_1, \dots, a_k, t_1, \dots, t_k) \in \Theta_{k,\varepsilon}$ be a solution of the constrained minimization problem (3.7) such that $0 < |a_1| \leq |a_2| \leq \dots \leq |a_k|$. Let

$$\beta_{GMM} = \min \left(\frac{\varepsilon \sqrt{\lambda_{\min}(\Gamma)}}{8}, \frac{|a_1|}{2}, \frac{(1-\gamma)(1-c_K) \min(1, d_K |a_1|^2)}{8C_{\phi, \theta^*} \sqrt{1+\gamma}(\sqrt{1+\gamma}\sqrt{1+c_K\sqrt{D_K}}+2|a_k|D'_{A,K})} \right) \quad (3.9)$$

where $\lambda_{\min}(\Gamma)$ is minimum eigenvalue of Γ , c_K , d_K and D_K are constants depending only on the chosen kernel K , $D_{A,K}$ is an explicit constant depending on K and the acquisition operator A and $C_{\phi, \theta^*} > 0$. Additionally suppose that for all i , $\|t_i\|_2 \leq R - 2\beta_{GMM}$.

Then $\Lambda_{\beta_{GMM}} := \{\theta : \|\theta - \theta^*\|_2 < \beta_{GMM}\}$ is a g -basin of attraction of θ^* .

Proof. See Section 6.7. □

With Theorem 3.3, we are able to control the non-negativity of the Hessian over an explicit neighborhood of θ^* . Note that the assumption on the norm of the t_i is technical to guarantee the stability of iterates. It is not a very strong assumption in practice as β_{GMM} is generally small with respect to R . Also, it could be dropped with a dedicated convergence proof. The constants involved in the expression lead to a size of the basin of attraction having the same behavior with respect to parameters of the problem: it increases with respect to the number of measurements (RIP constant) and it decreases with respect to the minimum weight in x_0 . Note that with Fourier measurements proposed in [23] the constant $D_{A,K}$ is independent of m . A dependency on the minimum amplitude in θ^* is also observed similarly to the case of Dirac estimation.

OPEN QUESTIONS FOR THE EXTENSION TO UNKNOWN VARIABLE COVARIANCES WITH LOW-RANK CONSTRAINT For real world applications, the known covariance case can be too simplistic to perform compressive statistical learning with good results. In practice, the covariance matrices of Gaussian mixtures are also estimated. Often, a low-rank approximation is made (flat-tail) to reduce the number of estimated parameters. For example, an unknown low-rank covariances model was used in [32, 33] to learn an image patch model from a compressed database. This model is then used to perform patch based image denoising.

A low-dimensional model with unknown low-rank covariances can be defined as follows. The parameters are $\theta = (a_1, \dots, a_k, t_1, \dots, t_k, Z_1, \dots, Z_k)$ where the Z_i are $d \times r$ matrices used to model covariances $\Gamma_i = Z_i Z_i^T + \rho I$. In this case, the model is

$$\Sigma_{k,\varepsilon,r,\rho,P} := \left\{ \sum_{i=1}^k a_i \mu_{t_i, \Gamma_i} : a_i \in \mathbb{R}, \|t_i - t_j\|_2 > \varepsilon, t_i \in \mathcal{B}_2(R), \rho < \lambda_j(\Gamma_i) < P, \right. \\ \left. \text{rank}(\Gamma_i - \rho I) \leq r \right\} \quad (3.10)$$

where μ_{t_i, Γ_i} is the Gaussian measure of mean t_i and covariance Γ_i . Recovery guarantees of the ideal decoder have not yet been given in this more general case (and it is out of the scope of this article). To do so, one would need to show the existence of a linear measurement operator with a RIP on $\mathcal{S}(\Sigma_{k,\varepsilon,r,\rho,P})$.

For off-the-grid sparse recovery, a block coordinate descent with respect to amplitudes and means is easier to implement in practice. To consider the case of variable covariances, we suggest to follow this guideline and to perform a block coordinate descent step with respect to each covariance matrix. With block coordinate descent, we would just need to make sure the Hessian with respect to the coordinates of Z_i is positive over a neighborhood of θ^* [4] to calculate explicitly a basin of attraction.

The directional derivative of $a_i \mu_{t_i, Z_i Z_i^T + \rho I}$ in the direction (v, w, W) has density (proof in Section 6.7.1):

$$\sum_i (v_i - a_i w^T \Gamma_i^{-1} (t - t_i) + a_i \frac{1}{2} (t - t_i)^T \Gamma_i^{-1} (Z_i W_i^T + W_i Z_i^T) \Gamma_i^{-1} (t - t_i)) e^{-\frac{1}{2} \|t - t_i\|_{Z_i Z_i^T + \rho I}^2}. \quad (3.11)$$

Finding a lower bound of the kernel norm of this directional derivative is quite technical. We observe that we would need to handle the indeterminacy of the low-rank parametrization as in Section 3.1. We

propose the following lemma as a first step towards building a theorem for an explicit basin of attraction in the variable flat tail covariance case. The proof of this lemma shows the technicalities involved to deal with the full variable covariance case. With this lemma it is possible to show that the Hessian blocks corresponding to the covariances in the full rank diagonal case are positive over a neighborhood of the global minimum (in other words, the radius β_2 from Corollary 2.1 is strictly positive) as long as the precision of the kernel is good enough. We leave the full study of non-convex methods for the variable covariance case with low-rank approximation including establishing RIP recovery guarantees for future work.

LEMMA 3.1 Let K be a Gaussian kernel with covariance $\frac{1}{\lambda}I$. Suppose $\Gamma \in \mathbb{R}^{p \times p}$ is a diagonal definite positive matrix. Let $\partial_W \mu_{t,\Gamma}$ the directional derivative of $\mu_{t,\Gamma}$ with respect to Γ in (diagonal) direction W such that $\|W\|_F = 1$. There exists $L > 0$ such that if $\lambda \geq L$, it then holds:

$$\|\partial_W \mu_{t,\Gamma}\|_K^2 \geq D_{\lambda,\Gamma} \quad (3.12)$$

where $D_{\lambda,\Gamma} > 0$ is a constant that only depends on λ and Γ .

Proof. See Section 6.7.1. □

4. From ideal to practical backprojection initialization: the challenge of non-convex low-dimensional recovery?

A common approach for the initialization of non-convex low-dimensional recovery is the initialization by backprojection techniques. Such initialization is outdated for low-rank matrix recovery as global convergence of descent algorithms has been proven [27]. However, it is still necessary for other models such as sparse spikes, Gaussian mixtures and phase recovery (which is not developed in this article). Finding a practical initialization technique is essentially a case by case heuristic design problem. We investigate the usual backprojection method in the noiseless case $y = Ax_0$.

In finite dimension, backprojection techniques rely on the fact that the RIP of A implies $\|(A^H A - I)x_0\|_2^2 \leq \gamma$ (where A^H is the adjoint of A). In our general framework, the adjoint of A does not backproject in \mathcal{D}^* , but in \mathcal{D} . We formalize such backprojection within our framework.

DEFINITION 4.1 (Ideal backprojection) Given measurements $y = Ax_0 \in \mathbb{C}^m$, we define the ideal backprojection $z \in \mathcal{D}$ with

$$z := \sum_{l=1}^m y_l \alpha_l. \quad (4.1)$$

We immediately have the following Lemma (which is a direct consequence of the RIP):

LEMMA 4.1 Suppose A has the RIP with constant γ . Let $z \in \mathcal{D}$ the ideal backprojection for measurements $y = Ax_0$. Then

$$(1 - \gamma)\|x_0\|_{\mathcal{H}}^2 \leq \langle x_0, z \rangle \leq (1 + \gamma)\|x_0\|_{\mathcal{H}}^2. \quad (4.2)$$

Proof of Lemma 4.1.

$$\begin{aligned}
\langle x_0, z \rangle &= \langle x_0, \sum_{l=1}^m y_l \alpha_l \rangle \\
&= \langle x_0, \sum_{l=1}^m \langle x_0, \alpha_l \rangle \alpha_l \rangle \\
&= \sum_{l=1}^m \overline{\langle x_0, \alpha_l \rangle} \langle x_0, \alpha_l \rangle \\
&= \sum_{l=1}^m |\langle x_0, \alpha_l \rangle|^2 = \|Ax_0\|_2^2.
\end{aligned} \tag{4.3}$$

Hence, with the RIP,

$$(1 - \gamma) \|x_0\|_{\mathcal{H}}^2 \leq \langle x_0, z \rangle \leq (1 + \gamma) \|x_0\|_{\mathcal{H}}^2. \tag{4.4}$$

□

Lemma 4.1 shows that the ideal backprojection preserves the energy in x_0 (up to a RIP constant). Hence when the number of measurement increases, it can get arbitrarily close to the global optimum (looking through the duality product). The main challenge is to extract initial parameters θ_{init} from z . One possibility would be to separate this task in two steps: define θ_{init} minimizing a “distance” between $\phi(\theta)$ and z for theoretical study and then find practical ways to minimize such distance. Such distance could be built using the norm $\|\cdot\|_{\mathcal{H}}$ or even directly the duality product.

In phase retrieval [41], spectral initialization techniques consist in taking the leading eigenvectors of a matrix constructed as linear combination of backprojections of individual measurements.

In the case of spike super-resolution a heuristic based on a sampling on a grid of the ideal backprojection followed by dimension reduction was proposed to initialize a descent algorithm to perform spike recovery for Diracs supported on a low-dimensional domain \mathbb{R}^p (e.g. $p = 2$ and possibly 3,4) [35]. Qualitative evidence that a grid step size ε_g small enough permits the initialization in the basin of attraction are the only recovery results for sparse spike recovery with descent methods in the parameter space. Grid based initialization techniques also requires one algorithmic step (the backprojection) whose computational complexity ε_g^{-d} might be a drawback for high-dimensional d (curse of dimension). A practical alternative consists in greedy heuristics to initialize descent algorithms by greedily adding spikes one by one (general theoretical recovery guarantees for such methods are still being investigated beyond specific examples [20]).

For Gaussian mixture estimation, it is possible to imagine a grid based heuristic to estimate the means followed by clustering (to backproject onto the separation constraint). However, it is an open question to determine if the curse of dimension can be avoided in the initialization in order to outperform greedy methods [25].

5. Conclusion

We have described a generic framework to perform low-dimensional recovery from linear measurements using non-convex optimization. We showed how recent examples of the literature can be studied within this framework and we gave new results for the case of Gaussian mixture modeling.

The following open questions emerge from this work:

- Other applications: The objective of our framework is to be used for other applications where such results do not exist. For example, our results should apply to low-rank tensor recovery, where recent works study the RIP in this context [31] and study convergence of the Burer-Monteiro method [28]. It might be also possible to apply directly our framework for inverse problems using generative models with deep neural networks (e.g. in [1], a neural network is used to learn the parametrization function ϕ). We would then need to be able to show the existence of a RIP on such models first.
- RIP based guarantees, while useful qualitative tools, are often too strong to give useful quantitative guarantees. Can our framework be extended to a non-uniform recovery case to obtain more precise estimates?
- Is it possible in general to initialize within the basin of attraction, avoiding the curse of dimension (for parametrized infinite dimensional problems), and providing a full quantitative proof of convergence in the case of super-resolution and GMM estimation?

6. Annex

6.1 Summary of notations

Spaces and sets:

- \mathcal{D} : Banach space of functions used to measure the unknown x_0 .
- \mathcal{D}^* : ambient space where the objects of interest (e.g. the unknown x_0) belong, dual space of \mathcal{D} .
- Σ : low dimensional model set, i.e. a cone $\Sigma \subset \mathcal{D}^*$.
- $\mathcal{S}(\Sigma)$: secant set of Σ (differences of elements of Σ).
- Θ : open subset of \mathbb{R}^d which parametrizes the model Σ .
- \mathcal{H} : Hilbert space containing the low dimensional model. The reconstruction error is measured using the associated Hilbert norm.
- \mathbb{C}^m : complex finite dimensional vector space containing the m finite measurements.
- Λ_β : basin of attraction for the appropriate β (given by our theorems).
- $[\theta_1, \theta_2]$: for $\theta_1, \theta_2 \in \mathbb{R}^d$, denotes the line segment $\{t\theta_1 + (1-t)\theta_2 : t \in [0, 1]\}$.

Norms, functions and operators:

- $\langle \cdot, \cdot \rangle$: depending on context: duality product between \mathcal{D} and \mathcal{D}^* , conventional scalar Hermitian product in finite dimension.
- $\|\cdot\|_{\mathcal{H}}$: norm associated with \mathcal{H} .
- $\|\cdot\|_2$: usual ℓ^2 -norm.
- $d(\cdot, \theta)$: distance to the set of equivalent parametrizations of θ .
- $p(\cdot, \theta)$: projection on the set of equivalent parametrizations of θ .

- ϕ : parametrization function ($\phi(\Theta) = \Sigma$).
- α_l : measurement functions in \mathcal{D} for $1 \leq l \leq m$.
- A : linear measurement operator ($\mathcal{D}^* \rightarrow \mathbb{C}^m$) defined by $Ax = (\langle x, \alpha_l \rangle)_{1 \leq l \leq m}$.
- g : parametrized functional we want to minimize, $g(\theta) = \|A\phi(\theta) - y\|_2^2$.

6.2 Weak-* topology and continuity

PROPOSITION 6.1 (Neighborhood) For all $x_0 \in \mathcal{D}^*$, a basis of neighborhoods of x_0 for the weak-* topology is formed with the sets

$$V_{\varepsilon, (\alpha_i)_{1 \leq i \leq n}}(x_0) := \{x \in \mathcal{D}^* : \forall i = 1, \dots, n, |\langle x - x_0, \alpha_i \rangle| < \varepsilon\} \quad (6.1)$$

for all $\varepsilon > 0$ and all $\alpha_i \in \mathcal{D}^*$ and $n \in \mathbb{N}$.

Continuity is defined as follows.

DEFINITION 6.2 (Weak* continuity) Let $A : \mathcal{D}^* \rightarrow \mathbb{R}^d$. A is weak-* continuous at x_0 if for any $\varepsilon > 0$, there is a neighborhood $V(x_0)$ for the weak-* topology of x_0 such that

$$x \in V(x_0) \implies \|f(x) - f(x_0)\|_2 < \varepsilon. \quad (6.2)$$

6.3 Proofs for Section 2.2

Proof of Proposition 2.3. Let $z \in \mathbb{C}^m$, and let z^H be the Hermitian conjugate of z . Remark that $\theta \rightarrow A\phi(\theta)$ is a Gateaux differentiable function $\mathbb{R}^d \rightarrow \mathbb{C}^m$. We have

$$\begin{aligned} \frac{\partial g(\theta)}{\partial \theta_i} &= \left(\frac{\partial [A\phi](\theta)}{\partial \theta_i} \right)^H A\phi(\theta) + (A\phi(\theta))^H \frac{\partial [A\phi](\theta)}{\partial \theta_i} - 2\Re e \left\langle \frac{\partial [A\phi](\theta)}{\partial \theta_i}, y \right\rangle \\ &= 2\Re e \left(\frac{\partial [A\phi](\theta)}{\partial \theta_i} \right)^H A\phi(\theta) - 2\Re e \left\langle \frac{\partial [A\phi](\theta)}{\partial \theta_i}, y \right\rangle \\ &= 2\Re e \left\langle \frac{\partial [A\phi](\theta)}{\partial \theta_i}, A\phi(\theta) - y \right\rangle, \end{aligned} \quad (6.3)$$

By linearity and continuity of A : $\frac{\partial [A\phi](\theta)}{\partial \theta_i} = A \frac{\partial \phi(\theta)}{\partial \theta_i}$ and

$$\frac{\partial g(\theta)}{\partial \theta_i} = 2\Re e \left\langle A \frac{\partial \phi(\theta)}{\partial \theta_i}, A\phi(\theta) - y \right\rangle. \quad (6.4)$$

□

Proof of Proposition 2.4. Use Proposition 2.3 with the properties of the Hermitian product. □

6.4 Proof for Section 2.3

Proof of Lemma 2.1. Let $v = \partial_u \phi(\theta) \in \overline{\mathcal{S}(\Sigma)}$. Using the fact that Θ is an open set, let $v_n = \frac{\phi(\theta + |h_n|u) - \phi(\theta)}{|h_n|}$ with $|h_n| \rightarrow 0$ be a sequence of real numbers. Since Σ is a cone, $v_n \in \mathcal{S}(\Sigma)$. Thanks

to the compatibility assumption 2.1 on $\|\cdot\|_{\mathcal{H}}$, $\|v_n\|_{\mathcal{H}}$ converges to $\|\partial_u \phi(\theta)\|_{\mathcal{H}}$. Moreover, by definition $v_n \xrightarrow{*} \partial_u \phi(\theta) \in \overline{\mathcal{S}(\Sigma)}$, and by continuity of A , we have that $Av_n \rightarrow Av$ w.r.t to $\|\cdot\|_2$. Using the hypothesis,

$$(1 - \gamma) \|v_n\|_{\mathcal{H}}^2 \leq \|Av_n\|_2^2 \leq (1 + \gamma) \|v_n\|_{\mathcal{H}}^2. \quad (6.5)$$

Taking both inequalities to the limit yields the result. \square

6.5 Proofs for Section 2.5

Proof of Lemma 2.2. We write $H = F + G$ as in Proposition 2.4.

Let $u \in \mathbb{R}^d$, with the linearity of A and the linearity of the Gateaux differential, we have $\sum_i u_i \frac{\partial \phi(\theta)}{\partial \theta_i} = \partial_u \phi(\theta)$ and

$$\begin{aligned} u^T Gu &= \sum_{i,j} u_i u_j G_{i,j} = \sum_{i,j} u_i u_j 2 \Re e \langle A \frac{\partial \phi(\theta)}{\partial \theta_i}, A \frac{\partial \phi(\theta)}{\partial \theta_j} \rangle \\ &= 2 \Re e \langle A \sum_i u_i \frac{\partial \phi(\theta)}{\partial \theta_i}, A \sum_j u_j \frac{\partial \phi(\theta)}{\partial \theta_j} \rangle \\ &= 2 \left\| A \sum_i u_i \frac{\partial \phi(\theta)}{\partial \theta_i} \right\|_2^2 \\ &= 2 \|A \partial_u \phi(\theta)\|_2^2. \end{aligned} \quad (6.6)$$

By definition of the generalized secant, we have $\partial_u \phi(\theta) \in \overline{\mathcal{S}(\Sigma)}$. Using the RIP of A with Lemma 2.1, we get

$$2(1 - \gamma) \|\partial_u \phi(\theta)\|_{\mathcal{H}}^2 \leq u^T Gu \leq 2(1 + \gamma) \|\partial_u \phi(\theta)\|_{\mathcal{H}}^2. \quad (6.7)$$

Moreover, with the Cauchy-Schwarz and triangle inequalities,

$$\begin{aligned} |u^T Fu| &= |2 \langle A \partial_u^2 \phi(\theta), A \phi(\theta) - y \rangle| \\ &\leq 2 \|A \partial_u^2 \phi(\theta)\|_2 (\|A \phi(\theta) - A \phi(\theta^*)\|_2 + \|A \phi(\theta^*) - y\|_2). \end{aligned} \quad (6.8)$$

By definition of θ^* , $\|A \phi(\theta^*) - y\|_2 \leq \|Ax_0 - y\|_2 = \|e\|_2$. With the RIP,

$$|u^T Fu| \leq 2 \|A \partial_u^2 \phi(\theta)\|_2 (\sqrt{1 + \gamma} \|\phi(\theta) - \phi(\theta^*)\|_{\mathcal{H}} + \|e\|_2). \quad (6.9)$$

Combining (6.7) and (6.9) gives the lower bound. Combining the last equality of (6.6) and (6.9) gives the upper bound. \square

Proof of Theorem 2.8.

We start the proof by showing that given $\theta \in \Lambda_\beta$, there exists a global minimizer $\tilde{\theta} \in \Theta$ such that g is convex on the segment $[\tilde{\theta}, \theta]$. We then show the Lipschitz property of the gradient. These two facts lead to the proof of the stability of iterates in Λ_β and permit to show the convergence of the fixed step gradient descent.

Convexity property. Let $\theta \in \Lambda_\beta$ and $\tilde{\theta} \in p(\theta, \theta^*)$ (see Equation (2.9)) so that $u := \tilde{\theta} - \theta$ satisfies $\|u\|_2 < \beta$. By hypothesis, we can assume that $\tilde{\theta}$ satisfies (2.16). Let also $t \in [0, 1]$. We have $\|\theta + tu - \tilde{\theta}\|_2 = (1-t)\|\theta - \tilde{\theta}\|_2 < \beta$. This implies $\theta + tu \in \Lambda_\beta$ and $\phi(\theta + tu) \subset \Sigma$. Denoting by $H_{\theta+tu}$ the Hessian of g at $\theta + tu$, Lemma 2.2 and technical assumption 2.2 number 2 imply that

$$\begin{aligned} u^T H_{\theta+tu} u &\geq 2(1-\gamma) \|\partial_u \phi(\theta + tu)\|_{\mathcal{H}}^2 - 2 \|A \partial_u^2 \phi(\theta + tu)\|_2 (\sqrt{1+\gamma} \|\phi(\theta + tu) - \phi(\theta^*)\|_{\mathcal{H}} + \|e\|_2) \\ &\geq 2(1-\gamma) \|\partial_u \phi(\theta + tu)\|_{\mathcal{H}}^2 - 2 \|A \partial_u^2 \phi(\theta + tu)\|_2 (\sqrt{1+\gamma} C_{\phi, \theta^*} d(\theta + tu, \theta^*) + \|e\|_2) \\ &\geq 2(1-\gamma) \|\partial_u \phi(\theta + tu)\|_{\mathcal{H}}^2 - 2 \|A \partial_u^2 \phi(\theta + tu)\|_2 (\sqrt{1+\gamma} C_{\phi, \theta^*} \|\theta + tu - \tilde{\theta}\|_2 + \|e\|_2) \\ &\geq 2(1-\gamma) \|\partial_u \phi(\theta + tu)\|_{\mathcal{H}}^2 - 2 \|A \partial_u^2 \phi(\theta + tu)\|_2 (\sqrt{1+\gamma} C_{\phi, \theta^*} \beta + \|e\|_2). \end{aligned} \quad (6.10)$$

Therefore, using Hypothesis (2.16), we get that $u^T H_{\theta+tu} u \geq 0$ for all $t \in [0, 1]$, which guarantees the convexity of g on $[\tilde{\theta}, \theta]$. Since g is Gateaux differentiable, the properties of one-dimensional convex functions give

$$g(\tilde{\theta}) - g(\theta) \geq \langle \nabla g(\theta), \tilde{\theta} - \theta \rangle. \quad (6.11)$$

While g might not be convex in Λ_β due to indeterminacies, this generalized convexity inequality is enough to prove the convergence of the gradient descent, if we have the required Lipschitz nature of the gradient of g , which we address now.

Lipschitz gradient property. Let u be a vector such that $\|u\|_2 = 1$. Let $\theta \in \Lambda_{2\beta}$ and $\tilde{\theta} \in p(\theta, \theta^*)$ such that $\|\theta - \tilde{\theta}\|_2 < 2\beta$.

With Lemma 2.2 and technical assumption 2.2 number 4, we get the upper control

$$\begin{aligned} u^T H_\theta u &\leq 2 \|A \partial_u \phi(\theta)\|_2^2 + 2 \|A \partial_u^2 \phi(\theta)\|_2 (\sqrt{1+\gamma} C_{\phi, \theta^*} 2\beta + \|e\|_2). \\ &\leq 2 \|A \partial_u \phi(\theta)\|_2^2 + 2M_2 (\sqrt{1+\gamma} C_{\phi, \theta^*} 2\beta + \|e\|_2). \end{aligned} \quad (6.12)$$

Using again technical assumptions 2.2 number 4 and 3 give

$$\begin{aligned} \|A \partial_u \phi(\theta) - A \partial_u \phi(\tilde{\theta})\|_2 &\leq M_2 \|\theta - \tilde{\theta}\|_2 \\ \|A \partial_u \phi(\theta)\|_2 &\leq \|A \partial_u \phi(\tilde{\theta})\|_2 + M_2 \|\theta - \tilde{\theta}\|_2 \\ &\leq \|A \partial_u \phi(\tilde{\theta})\|_2 + M_2 2\beta \\ &\leq M_1 + 2M_2 \beta. \end{aligned} \quad (6.13)$$

We conclude that

$$u^T H_\theta u \leq 2(M_1 + 2M_2 \beta)^2 + 2M_2 (\sqrt{1+\gamma} C_{\phi, \theta^*} 2\beta + \|e\|_2) \quad (6.14)$$

and finally

$$L := \sup_{\theta \in \Lambda_{2\beta}} \sup_{u: \|u\|_2=1} u^T H_\theta u < +\infty. \quad (6.15)$$

Hence the gradient of g is L -Lipschitz on $\Lambda_{2\beta}$.

Stability of iterates. Let $(\theta_n)_{n \geq 0}$ be the sequence of iterates of the gradient descent with fixed step τ and $\theta_0 \in \Lambda_\beta$. Since ∇g is L -Lipschitz on Λ_β and $\|\nabla g(\tilde{\theta})\|_2 < \infty$ (using the fact that $\left| \frac{\partial g(\tilde{\theta})}{\partial \theta_i} \right| \leq$

$2\|A \frac{\partial \phi(\tilde{\theta})}{\partial \theta_i}\|_2 \|e\|_2 \leq 2M_1 \|e\|_2$), we have $\sup_{\Lambda_\beta} \|\nabla g(\theta)\|_2 < \infty$. We can thus set the descent step

$$\tau < \min \left(\frac{1}{L}, \frac{\beta}{\sup_{\Lambda_\beta} \|\nabla g(\theta)\|_2} \right). \quad (6.16)$$

For each n , we introduce $\tilde{\theta}_n \in p(\theta_n, \theta^*)$ which satisfies (2.16) in order to have the convexity property shown above. Suppose $\theta_n \in \Lambda_\beta$. We then have

$$d(\theta_{n+1}, \theta^*) \leq \|\theta_{n+1} - \tilde{\theta}_n\|_2 \leq \|\theta_n - \tilde{\theta}_n\|_2 + \|\tau \nabla g(\theta_n)\|_2 < \beta + \tau \sup_{\Lambda_\beta} \|\nabla g(\theta)\|_2 \leq 2\beta. \quad (6.17)$$

This proves that $\theta_{n+1} \in \Lambda_{2\beta}$ and, using technical assumption 2.2 number 1, that $\phi(\theta_{n+1}) \in \Sigma$. These last inequalities also give $[\theta_n, \theta_{n+1}] \subset \Lambda_{2\beta}$ (by replacing τ by τ' such that $0 \leq \tau' \leq \tau$).

As shown in the first part of the proof, g is convex on $[\theta_n, \tilde{\theta}_n]$, and thus

$$\begin{aligned} g(\theta_{n+1}) - g(\tilde{\theta}_n) &= g(\theta_{n+1}) - g(\theta_n) + g(\theta_n) - g(\tilde{\theta}_n) \\ &\leq g(\theta_{n+1}) - g(\theta_n) + \langle \nabla g(\theta_n), \theta_n - \tilde{\theta}_n \rangle \end{aligned} \quad (6.18)$$

Since g is twice Gateaux differentiable with L -Lipschitz gradient on $\Lambda_{2\beta}$, we can use the second-order Taylor inequality on the segment $[\theta_n, \theta_{n+1}] \subset \Lambda_{2\beta}$, which gives

$$g(\theta_{n+1}) - g(\theta_n) \leq -\tau \langle \nabla g(\theta_n), \nabla g(\theta_n) \rangle + \frac{L}{2} \|\tau \nabla g(\theta_n)\|_2^2 = \left(\frac{\tau^2 L}{2} - \tau \right) \|\nabla g(\theta_n)\|_2^2. \quad (6.19)$$

Since $\phi(\theta_{n+1}) \in \Sigma$ and $\tilde{\theta}_n$ is a global minimizer, we get

$$\langle \nabla g(\theta_n), \theta_n - \tilde{\theta}_n \rangle - \left(\tau - \frac{\tau^2 L}{2} \right) \|\nabla g(\theta_n)\|_2^2 \geq g(\theta_{n+1}) - g(\tilde{\theta}_n) \geq 0. \quad (6.20)$$

Now,

$$d(\theta_{n+1}, \theta^*)^2 \leq \|\theta_{n+1} - \tilde{\theta}_n\|_2^2 = \|\theta_n - \tilde{\theta}_n\|_2^2 - 2\langle \theta_n - \tilde{\theta}_n, \tau \nabla g(\theta_n) \rangle + \tau^2 \|\nabla g(\theta_n)\|_2^2. \quad (6.21)$$

Plugging the previous inequality gives

$$\begin{aligned} \|\theta_{n+1} - \tilde{\theta}_n\|_2^2 &\leq \|\theta_n - \tilde{\theta}_n\|_2^2 - (2\tau^2 - \tau^3 L) \|\nabla g(\theta_n)\|_2^2 + \tau^2 \|\nabla g(\theta_n)\|_2^2 \\ &= \|\theta_n - \tilde{\theta}_n\|_2^2 - c_0 \|\nabla g(\theta_n)\|_2^2 \end{aligned} \quad (6.22)$$

where $c_0 = \tau^2(1 - L\tau) > 0$. Because $\tau < 1/L$, we deduce that $d(\theta_{n+1}, \theta^*) < d(\theta_n, \theta^*)$ and $\theta_{n+1} \in \Lambda_\beta$. By induction, we get that the iterates stay in Λ_β because $\theta_0 \in \Lambda_\beta$.

Convergence of $g(\theta_n)$. We use the same notation $\theta_n, \tilde{\theta}_n$ as in the last paragraph. Using again (6.19), we get

$$\begin{aligned} g(\theta_{n+1}) - g(\theta^*) &= g(\theta_n) - g(\theta^*) + g(\theta_{n+1}) - g(\theta_n) \\ &\leq g(\theta_n) - g(\theta^*) - \left(\tau - \frac{L\tau^2}{2} \right) \|\nabla g(\theta_n)\|_2^2 \end{aligned} \quad (6.23)$$

Remark that $\tau - \frac{L\tau^2}{2} = \tau(1 - L\tau/2) \geq \tau/2 > 0$. Using the convexity on $[\theta_n, \tilde{\theta}_n]$ and the Cauchy-Schwarz inequality, we get

$$\begin{aligned} g(\theta_n) - g(\theta^*) &= g(\theta_n) - g(\tilde{\theta}_n) \leq \langle \nabla g(\theta_n), \theta_n - \tilde{\theta}_n \rangle \\ &\leq \|\nabla g(\theta_n)\|_2 \|\theta_n - \tilde{\theta}_n\|_2 \\ &\leq \|\nabla g(\theta_n)\|_2 \beta. \end{aligned} \quad (6.24)$$

We get

$$\begin{aligned} g(\theta_{n+1}) - g(\theta^*) &\leq g(\theta_n) - g(\theta^*) - \left(\tau - \frac{L\tau^2}{2}\right) \frac{(g(\theta_n) - g(\theta^*))^2}{\beta^2} \\ &= g(\theta_n) - g(\theta^*) - c_1 (g(\theta_n) - g(\theta^*))^2 \end{aligned} \quad (6.25)$$

with $c_1 := \left(\tau - \frac{L\tau^2}{2}\right) \frac{1}{\beta^2}$.

Let $d_n = g(\theta_n) - g(\theta^*) \geq 0$, using $d_{n+1} \leq d_n$, we have that

$$\begin{aligned} \frac{d_{n+1}}{d_n} &\leq 1 - c_1 d_n \\ \frac{1}{d_n} &\leq \frac{1}{d_{n+1}} - c_1 \frac{d_n}{d_{n+1}} \\ c_1 &\leq c_1 \frac{d_n}{d_{n+1}} \leq \frac{1}{d_{n+1}} - \frac{1}{d_n} \end{aligned} \quad (6.26)$$

We sum this inequality for $0, \dots, n-1$ and get

$$\begin{aligned} nc_1 &\leq \frac{1}{d_n} - \frac{1}{d_0} \leq \frac{1}{d_n} \\ d_n &\leq \frac{1}{c_1 n}. \end{aligned} \quad (6.27)$$

□

Proof of Corollary 2.1. Let $\beta = \min(\beta_1, \beta_2)$, as $\Lambda_{2\beta} \subset \Lambda_{2\beta_1}$, the whole technical assumption 2.2 required for Theorem 2.8 is verified.

Let $\theta \in \Lambda_\beta \subset \Lambda_{\beta_1}$ and let us consider the unique $\tilde{\theta} \in p(\theta, \theta^*)$. Using the last hypothesis, we have for all $z \in [\theta, \tilde{\theta}]$ that

$$\frac{1}{C_{\phi, \theta^*}} \frac{(1-\gamma) \|\partial_{\tilde{\theta}-\theta} \phi(z)\|_{\mathcal{H}}^2}{\sqrt{1+\gamma} \|A \partial_{\tilde{\theta}-\theta}^2 \phi(z)\|_2} - \frac{1}{C_{\phi, \theta^*} \sqrt{1+\gamma}} \|e\|_2 \geq \beta_2 \geq \beta. \quad (6.28)$$

which implies the last hypothesis of Theorem 2.8

$$\frac{(1-\gamma) \|\partial_{\tilde{\theta}-\theta} \phi(z)\|_{\mathcal{H}}^2}{\sqrt{1+\gamma} \|A \partial_{\tilde{\theta}-\theta}^2 \phi(z)\|_2} \geq C_{\phi, \theta^*} \beta + \frac{1}{\sqrt{1+\gamma}} \|e\|_2. \quad (6.29)$$

□

Proof of Corollary 2.2. Consider $x_0 \in \Sigma$, and $(\theta_n)_{n \geq 0}$ a sequence of iterates of the gradient descent converging to θ^* under the hypotheses of Theorem 2.8. Using the RIP, we have

$$\sqrt{1-\gamma} \|\phi(\theta_n) - x_0\|_{\mathcal{H}} \leq \|A\phi(\theta_n) - Ax_0\|_2 \leq \|A\phi(\theta_n) - y\|_2 + \|e\|_2 \quad (6.30)$$

Using the rate of convergence of $g(\theta_n)$ (inequality (6.27)), we have

$$\|A\phi(\theta_n) - y\|_2^2 \leq \|A\phi(\theta^*) - y\|_2^2 + O\left(\frac{1}{n}\right). \quad (6.31)$$

This gives, using the inequality $(a+b)^2 \leq 2(a^2 + b^2)$,

$$\begin{aligned} \|\phi(\theta_n) - x_0\|_{\mathcal{H}}^2 &\leq \frac{1}{1-\gamma} \left(\sqrt{\|A\phi(\theta^*) - y\|_2^2 + O\left(\frac{1}{n}\right)} + \|e\|_2 \right)^2 \\ &\leq \frac{2}{1-\gamma} \left(\|A\phi(\theta^*) - y\|_2^2 + O\left(\frac{1}{n}\right) + \|e\|_2^2 \right) \\ &\leq \frac{2}{1-\gamma} \left(\|Ax_0 - y\|_2^2 + O\left(\frac{1}{n}\right) + \|e\|_2^2 \right) = \frac{4\|e\|_2^2}{1-\gamma} + O\left(\frac{1}{n}\right). \end{aligned} \quad (6.32)$$

□

6.6 Proofs for Section 3.1

The expression of the directional derivative is deduced from the fact that

$$\begin{aligned} \partial_{ij} ZZ^T &= (\partial_{ij} Z) Z^T + Z (\partial_{ij} Z)^T = (\partial_{ij} Z) Z^T + ((\partial_{ij} Z) Z^T)^T \\ &= E_{ij} Z^T + Z E_{ji} \end{aligned} \quad (6.33)$$

where the family $\{E_{ij} : 1 \leq i \leq p, 1 \leq j \leq r\}$ is the canonical basis of $\mathbb{R}^{p \times r}$. Furthermore

$$\partial_{ij} \partial_U ZZ^T = U E_{ij}^T + E_{ij} U^T. \quad (6.34)$$

LEMMA 6.1 Let $U \in \mathbb{R}^{p \times r}$ and $M \in \mathbb{R}^{r \times r}$. Suppose M PSD, then

$$\text{tr}(UMU^T) \geq \sigma_{\min}(M) \|U\|_F^2. \quad (6.35)$$

Proof. Let U_i be the rows of U . We have

$$\begin{aligned} (UMU^T)_{i,i} &= \sum_{l=1,r} U_{i,l} (MU^T)_{l,i} \\ &= \sum_{l=1,r} U_{i,l} \sum_{s=1,r} M_{l,s} U_{i,s} \\ &= \sum_{l=1,r} \sum_{s=1,r} U_{i,l} M_{l,s} U_{i,s} \\ &= U_i M U_i^T \\ &\geq \sigma_{\min}(M) U_i U_i^T \\ &= \sigma_{\min}(M) \|U_i\|_2^2 \end{aligned} \quad (6.36)$$

Summing over i yields the result. \square

Proof of Theorem 3.1. For a rank r matrix M , we consider $\sigma_{\max}(M) = \sigma_1(M) \geq \dots \geq \sigma_r(M) = \sigma_{\min}(M)$ the singular values of M in decreasing order. We will use the following inequalities in this proof.

- $\sigma_{\max}(M) \leq \|M\|_F$
- $\|AB\|_F \leq \sigma_{\max}(A)\|B\|_F$

We verify the hypotheses of Corollary 2.1.

Technical assumption 2.2 number 1: First, remark that for any $Z \in \mathbb{R}^{p \times r}$, we have $\phi(Z) \in \Sigma_r$ (i.e. this assumption is verified for any β).

Hypothesis 1: Let $Z \in \Lambda_{\beta_{LR}}$. Consider H_0 the unique solution of the orthogonal Procrustes problem $\min_{H \in \mathcal{O}(r)} \|Z_0 H - Z\|_F$. Then it has been shown that $H_0 = QR^T$ where Q and R are obtained from the singular value decomposition $Z_0^T Z = Q\Delta R^T$ of $Z_0^T Z$ (see e.g. [39, Proof of Lemma 5.7]). It gives

$$Z^T Z_0 H_0 = R\Delta Q^T Q R^T = R\Delta R^T = RQ^T Q\Delta R^T = H_0^T Z_0^T Z. \quad (6.37)$$

We get the projection $p(Z, Z_0) = \{Z_0 H_0\}$.

Denoting $\tilde{Z} := Z_0 H_0$, we have $Z^T \tilde{Z} = \tilde{Z}^T Z$ and $Z^T \tilde{Z}$ PSD. Notice that for $U = \tilde{Z} - Z$, we also have $\tilde{Z}^T U = U^T \tilde{Z}$ and $Z^T U = U^T Z$. Note also that for $1 \leq i \leq r$, $\sigma_i(\tilde{Z}) = \sigma_i(Z_0)$.

Technical assumption 2.2 number 2: We calculate a constant C_{ϕ, Z_0} . Using the triangle inequality on the operator norm σ_{\max} , the fact that $Z \in \Lambda_{\beta_{LR}}$ and the typical matrix norm inequality $\sigma_{\max}(Z - \tilde{Z}) \leq \|Z - \tilde{Z}\|_F$, we have $\sigma_{\max}(Z) \leq \sigma_{\max}(\tilde{Z}) + \beta_{LR} \leq 2\sigma_{\max}(Z_0)$ and

$$\begin{aligned} \|ZZ^T - Z_0 Z_0^T\|_F &= \|ZZ^T - \tilde{Z}\tilde{Z}^T\|_F \\ &= \|Z(Z - \tilde{Z})^T + (Z - \tilde{Z})\tilde{Z}^T\|_F \\ &\leq \|Z(Z - \tilde{Z})^T\|_F + \|(Z - \tilde{Z})\tilde{Z}^T\|_F \\ &\leq (\sigma_{\max}(Z) + \sigma_{\max}(\tilde{Z}))\|Z - \tilde{Z}\|_F \\ &\leq 3\sigma_{\max}(Z_0)\|Z - \tilde{Z}\|_F. \end{aligned} \quad (6.38)$$

Hence we can set $C_{\phi, Z_0} = 3\sigma_{\max}(Z_0)$.

Technical assumption 2.2 number 3 and 4: They come from the fact that g is infinitely differentiable on a bounded domain.

Hypothesis 3: First note that, as $Z \in \Lambda_{\beta_{LR}}$, we have $\|Z - \tilde{Z}\|_F < \beta_{LR} \leq \frac{\sigma_{\min}(Z_0)}{8}$. Hence, using Weyl's perturbation inequality, $\sigma_r(Z) \geq \sigma_r(\tilde{Z}) - \sigma_1(Z - \tilde{Z}) \geq \sigma_r(Z_0) - \|Z - \tilde{Z}\|_F \geq \sigma_r(Z_0) - \frac{\sigma_r(Z_0)}{8} > 0$ and the rank of Z is r .

As $\partial_U^2 \phi(Z)$ is an element of Σ_r , with the RIP on Σ_{2r} ,

$$\|A\partial_U^2 \phi(Z)\|_2 = 2\|AUU^T\|_2 \leq 2\sqrt{1 + \gamma}\|UU^T\|_F. \quad (6.39)$$

Now, we bound the ratio $\frac{\|\partial_U \phi(Z+tU)\|_F^2}{\|A\partial_U^2 \phi(Z)\|_2} = \frac{\|U(Z+tU)^T + (Z+tU)U^T\|_F^2}{2\|UU^T\|_F}$ for $U = \tilde{Z} - Z$ with $0 \leq t \leq 1$.

We have

$$\begin{aligned} \partial_U \phi(Z+tU) &= U(Z+tU)^T + (Z+tU)U^T \\ &= UZ^T + ZU^T + 2tUU^T. \end{aligned} \quad (6.40)$$

For any $M_1, M_2 \in \mathbb{R}^{p \times r}$ with $M_1 \neq 0$, we have that $\|tM_1 + M_2\|_F^2$ is minimized for $t^* = -\frac{\langle M_1, M_2 \rangle_F}{\|M_1\|_F^2}$ and $\|t^*M_1 + M_2\|_F^2 = \|M_2\|_F^2 - \frac{\langle M_1, M_2 \rangle_F^2}{\|M_1\|_F^2}$. Hence if $t^* \geq 1$, $\|tM_1 + M_2\|_F^2$ is minimized over $[0, 1]$ at $t = 1$, at $t = 0$ if $t^* \leq 0$ and $t = t^*$ otherwise. Applying this to the case $M_1 = 2UU^T$ and $M_2 = UZ^T + ZU^T$.

Case 1: $t^* \geq 1$.

$$\min_{t \in [0, 1]} \|U(Z + tU)^T + (Z + tU)U^T\|_F^2 = \|U\tilde{Z}^T + \tilde{Z}U^T\|_F^2. \quad (6.41)$$

Case 2: $t^* \leq 0$.

$$\min_{t \in [0, 1]} \|U(Z + tU)^T + (Z + tU)U^T\|_F^2 = \|UZ^T + ZU^T\|_F^2. \quad (6.42)$$

We consider Case 1 and 2 together. Let $\tilde{Z} = Z$ or $\tilde{Z} = \tilde{Z}$. We have

$$\begin{aligned} \frac{\|U\tilde{Z}^T + \tilde{Z}U^T\|_F^2}{2\|UU^T\|_F} &= \frac{\|U\tilde{Z}^T\|_F^2 + \|\tilde{Z}U^T\|_F^2 + 2\langle U\tilde{Z}^T, \tilde{Z}U^T \rangle_F}{2\|UU^T\|_F} \\ &= \frac{2\|\tilde{Z}U^T\|_F^2 + 2\langle \tilde{Z}^T, U^T\tilde{Z}U^T \rangle_F}{2\|UU^T\|_F}. \end{aligned} \quad (6.43)$$

Using the fact that $U^T\tilde{Z} = \tilde{Z}^TU$,

$$\begin{aligned} \frac{\|U\tilde{Z}^T + \tilde{Z}U^T\|_F^2}{2\|UU^T\|_F} &= \frac{\|\tilde{Z}U^T\|_F^2 + \langle \tilde{Z}^T, \tilde{Z}^T UU^T \rangle_F}{\|UU^T\|_F} \\ &= \frac{\|\tilde{Z}U^T\|_F^2 + \langle \tilde{Z}\tilde{Z}^T, UU^T \rangle_F}{\|UU^T\|_F}. \end{aligned} \quad (6.44)$$

Let A_1, A_2 be two positive symmetric matrices (e.g. ZZ^T and UU^T), then $\langle A_1, A_2 \rangle_F \geq 0$. Indeed, let $A_2^{\frac{1}{2}}$ be a positive symmetric square root of A_2 . Remark that the matrix $A_2^{\frac{1}{2}}A_1^T A_2^{\frac{1}{2}} = A_2^{\frac{1}{2}}A_1^T (A_2^{\frac{1}{2}})^T$ is positive symmetric. Hence its trace is positive and we have

$$\langle A_1, A_2 \rangle_F = \text{tr}(A_1^T A_2) = \text{tr}(A_1^T A_2^{\frac{1}{2}} A_2^{\frac{1}{2}}) = \text{tr}(A_2^{\frac{1}{2}} A_1^T A_2^{\frac{1}{2}}) \geq 0. \quad (6.45)$$

We deduce

$$\frac{\|\tilde{Z}U^T + U\tilde{Z}^T\|_F^2}{2\|UU^T\|_F} \geq \frac{\|\tilde{Z}U^T\|_F^2}{\|UU^T\|_F} = \frac{\text{tr}(U\tilde{Z}^T \tilde{Z}U^T)}{\|UU^T\|_F}. \quad (6.46)$$

But $\tilde{Z}^T \tilde{Z}$ is a full rank PSD matrix, hence, with Lemma 6.1,

$$\text{tr}(U\tilde{Z}^T \tilde{Z}U^T) \geq \sigma_{\min}(\tilde{Z}^T \tilde{Z}) \text{tr}(UU^T) \geq (\sigma_{\min}(\tilde{Z}) - \beta_{LR})^2 \text{tr}(UU^T). \quad (6.47)$$

Let $\lambda_i(UU^T)$ denote the eigenvalues of UU^T , then

$$\begin{aligned}
\frac{\|\partial_U \phi(Z+tU)\|_F^2}{\|\partial_U^2 \phi(Z+tU)\|_F} &\geq \frac{\|\tilde{Z}U^T + U\tilde{Z}^T\|_F^2}{2\|UU^T\|_F} \\
&\geq (\sigma_{\min}(Z_0) - \beta_{LR})^2 \frac{\sum_i \lambda_i(UU^T)}{\sqrt{\sum_i (\lambda_i(UU^T))^2}} \\
&\geq (\sigma_{\min}(Z_0) - \beta_{LR})^2 \\
&\geq \left(\frac{7}{8}\sigma_{\min}(Z_0)\right)^2.
\end{aligned} \tag{6.48}$$

Case 3: $0 < t^* < 1$.

Using the definition of t^* , we have

$$\begin{aligned}
0 &< -\frac{\langle M_1, M_2 \rangle_F}{\|M_1\|_F^2} < 1, \\
0 &< -\frac{\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T - 2UU^T \rangle_F}{2\|UU^T\|_F^2} < 1, \\
0 &< -\frac{\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F}{2\|UU^T\|_F^2} + 1 < 1, \\
-1 &< -\frac{\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F}{2\|UU^T\|_F^2} < 0, \\
0 &< \langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F < 2\|UU^T\|_F^2.
\end{aligned} \tag{6.49}$$

We calculate

$$\begin{aligned}
\|M_2\|_2^2 - \frac{|\langle M_1, M_2 \rangle_F|^2}{\|M_1\|_2^2} &= \|UZ^T + ZU^T\|_F^2 - \frac{|\langle UU^T, UZ^T + ZU^T \rangle_F|^2}{\|UU^T\|_F^2} \\
&= \|U\tilde{Z}^T + \tilde{Z}U^T - 2UU^T\|_F^2 - \frac{|\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T - 2UU^T \rangle_F|^2}{\|UU^T\|_F^2} \\
&= \|U\tilde{Z}^T + \tilde{Z}U^T\|_F^2 - 4\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F + 4\|UU^T\|_F^2 \\
&\quad - \frac{|\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F - 2\langle UU^T, UU^T \rangle_F|^2}{\|UU^T\|_F^2} \\
&= \|U\tilde{Z}^T + \tilde{Z}U^T\|_F^2 - 4\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F + 4\|UU^T\|_F^2 \\
&\quad - \frac{|\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F|^2}{\|UU^T\|_F^2} - 4\|UU^T\|_F^2 + 4\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F \\
&= 2\|U\tilde{Z}^T\|_F^2 + 2\langle \tilde{Z}\tilde{Z}^T, UU^T \rangle_F - \frac{|\langle UU^T, U\tilde{Z}^T + \tilde{Z}U^T \rangle_F|^2}{\|UU^T\|_F^2}.
\end{aligned} \tag{6.50}$$

Using the last inequality of (6.49), we have

$$\|M_2\|_2^2 - \frac{|\langle M_1, M_2 \rangle_F|^2}{\|M_1\|_2^2} \geq 2\|U\tilde{Z}^T\|_F^2 + 2\langle \tilde{Z}\tilde{Z}^T, UU^T \rangle_F - 4\|UU^T\|_F^2. \tag{6.51}$$

Using Equation (6.48) and the fact that $\|UU^T\|_F = \sqrt{\sum_i \sigma_i(U^T)^4} \leq \sum_i \sigma_i(U^T)^2 = \|U^T\|_F^2 \leq \beta_{LR}^2$, we conclude

$$\begin{aligned}
 \frac{\|\partial_U \phi(Z+tU)\|_F^2}{\|\partial_U^2 \phi(Z+tU)\|_F} &\geq \frac{\|U\tilde{Z}^T\|_F^2 + \langle \tilde{Z}\tilde{Z}^T, UU^T \rangle_F - 2\|UU^T\|_F^2}{\|UU^T\|_F} \\
 &\geq \left(\frac{7}{8}\sigma_{\min}(Z_0)\right)^2 - 2\|UU^T\|_F \\
 &\geq \left(\frac{7}{8}\sigma_{\min}(Z_0)\right)^2 - 2\beta_{LR}^2 \\
 &\geq \left(\frac{49}{64} - \frac{2}{64}\right)(\sigma_{\min}(Z_0))^2 = \frac{47}{64}(\sigma_{\min}(Z_0))^2 > 0.
 \end{aligned} \tag{6.52}$$

This gives, for β_2 from Corollary 2.1,

$$\begin{aligned}
 \beta_2 &= \frac{1}{C_{\phi, Z_0}} \inf_{Z \in \Lambda_{\beta_1}} \inf_{Y \in [Z, \tilde{Z}]} \left(\frac{(1-\gamma)\|\partial_{\tilde{Z}-Z}\phi(Y)\|_{\mathcal{H}}^2}{\sqrt{1+\gamma}\|A\partial_{\tilde{Z}-Z}^2\phi(Y)\|_2} \right) \\
 &\geq \frac{47(1-\gamma)(\sigma_{\min}(Z_0))^2}{64 \cdot 3\sigma_{\max}(Z_0)(1+\gamma)} \\
 &\geq \frac{1(1-\gamma)(\sigma_{\min}(Z_0))^2}{8(1+\gamma)\sigma_{\max}(Z_0)} = \beta_{LR} > 0,
 \end{aligned} \tag{6.53}$$

which implies $\Lambda_{\beta_{LR}} \subset \Lambda_{\min(\beta_1, \beta_2)}$. Using Corollary 2.1 yields the final result: the set $\Lambda_{\beta_{LR}}$ is a g -basin of attraction of θ^* . \square

6.7 Proofs for the GMM example

We start by giving two lemma that bound derivatives of ϕ .

LEMMA 6.2 Suppose $K(t) \propto e^{-\frac{1}{2}\lambda^2\|t\|_F^2}$. Then there is a strictly positive constant d_K such that $\|\partial_w \mu_{t_0}\|_K^2 \geq d_K \|w\|_2^2$ where $\partial_w \mu_{t_0}$ is the derivative of μ_{t_0} with respect to t in the direction w .

Proof. Using Equation (3.8),

$$\|\partial_w \mu_{t_0}\|_K^2 \propto \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} e^{-\frac{1}{2}\lambda^2\|t-s\|_F^2} w^T \Gamma^{-1} s w^T \Gamma^{-1} t e^{-\frac{1}{2}\|s-t_0\|_F^2} e^{-\frac{1}{2}\|t-t_0\|_F^2} dt ds \tag{6.54}$$

Using the change of variables $t-t_0 \rightarrow t$ and $s-t_0 \rightarrow s$, we get

$$\|\partial_w \mu_{t_0}\|_K^2 \propto \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} e^{-\frac{1}{2}\lambda^2\|t-s\|_F^2} w^T \Gamma^{-1} s w^T \Gamma^{-1} t e^{-\frac{1}{2}\|s\|_F^2} e^{-\frac{1}{2}\|t\|_F^2} dt ds \tag{6.55}$$

Rewrite

$$\begin{aligned}
\lambda^2 \|s-t\|_F^2 + \|s\|_F^2 &= (1+\lambda^2)\|s\|_F^2 - 2\lambda^2 \langle s, t \rangle + \lambda^2 \|t\|_F^2 \\
&= (1+\lambda^2) \left(\|s\|_F^2 - 2\frac{\lambda^2}{1+\lambda^2} \langle s, t \rangle + \frac{\lambda^2}{1+\lambda^2} \|t\|_F^2 \right) \\
&= (1+\lambda^2) \left(\left\| s - \frac{\lambda^2}{1+\lambda^2} t \right\|_F^2 + \left(1 - \frac{\lambda^2}{1+\lambda^2}\right) \frac{\lambda^2}{1+\lambda^2} \|t\|_F^2 \right) \\
&= (1+\lambda^2) \left\| s - \frac{\lambda^2}{1+\lambda^2} t \right\|_F^2 + \frac{\lambda^2}{1+\lambda^2} \|t\|_F^2.
\end{aligned} \tag{6.56}$$

This leads to

$$\begin{aligned}
\|\partial_w \mu_0\|_K^2 &\propto \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} w^T \Gamma^{-1} s w^T \Gamma^{-1} t e^{-\frac{1}{2}(1+\lambda^2)\|s - \frac{\lambda^2}{1+\lambda^2} t\|_F^2} e^{-\frac{1}{2}(\|t\|_F^2 + \frac{\lambda^2}{1+\lambda^2} \|t\|_F^2)} dt ds \\
&= \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} w^T \Gamma^{-1} s w^T \Gamma^{-1} t e^{-\frac{1}{2}(1+\lambda^2)\|s - \frac{\lambda^2}{1+\lambda^2} t\|_F^2} e^{-\frac{1}{2} \frac{1+\lambda^2}{1+\lambda^2} \|t\|_F^2} dt ds.
\end{aligned} \tag{6.57}$$

Let $h(s) = e^{-\frac{1}{2}(1+\lambda^2)\|s - \frac{\lambda^2}{1+\lambda^2} t\|_F^2}$,

$$\begin{aligned}
\int_s w^T \Gamma^{-1} s e^{-\frac{1}{2}(1+\lambda^2)\|s - \frac{\lambda^2}{1+\lambda^2} t\|_F^2} ds &= \int_s w^T \Gamma^{-1} \left(s - \frac{\lambda^2}{1+\lambda^2} t \right) h(s) ds \\
&\quad + \frac{\lambda^2}{1+\lambda^2} w^T \Gamma^{-1} t \int_s h(s) ds \\
&= \int_s \partial_w h(s) ds + \frac{\lambda^2}{1+\lambda^2} w^T \Gamma^{-1} t \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right)
\end{aligned} \tag{6.58}$$

where $\mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right)$ is the normalization constant of the Gaussian of covariance $\frac{\Gamma}{1+\lambda^2}$. By linearity

$$\int_s \partial_w h(s) ds = \langle w, \int_s \nabla h(s) ds \rangle = 0. \tag{6.59}$$

We make the change of variable $u = \Gamma^{-1/2} t$

$$\|\partial_w \mu_0\|_K^2 \propto \frac{\lambda^2}{1+\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) |\det(\Gamma^{-1/2})| \int_u |w^T \Gamma^{-1/2} u|^2 e^{-\frac{1}{2} \frac{1+\lambda^2}{1+\lambda^2} \|u\|_2^2} du. \tag{6.60}$$

With the linearity of the integral, the change of variable $u \rightarrow \sqrt{\frac{1+\lambda^2}{1+2\lambda^2}} u$ and the fact that $\int_{\mathbb{R}} x^2 e^{-\frac{1}{2}x^2} dx =$

$\sqrt{2\pi}$,

$$\begin{aligned}
 \|\partial_w \mu_0\|_K^2 &\propto |\det(\Gamma^{-1/2})| \frac{\lambda^2}{1+\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) \sum_i \int_u |u_i(\Gamma^{-1/2}w)_i|^2 e^{-\frac{1}{2} \frac{1+2\lambda^2}{1+\lambda^2} \|u\|_2^2} dt \\
 &\quad + \sum_{i \neq j} \int_u u_i(\Gamma^{-1/2}w)_i u_j(\Gamma^{-1/2}w)_j e^{-\frac{1}{2} \frac{1+2\lambda^2}{1+\lambda^2} \|u\|_2^2} du \\
 &= |\det(\Gamma^{-1/2})| \frac{\lambda^2}{1+\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) \left(\sqrt{\frac{1+\lambda^2}{1+2\lambda^2}} \right)^p \\
 &\quad \times \sum_i \int_u \frac{1+\lambda^2}{1+2\lambda^2} |u_i(\Gamma^{-1/2}w)_i|^2 e^{-\frac{1}{2} \|u\|_2^2} du \\
 &= |\det(\Gamma^{-1/2})| \frac{\lambda^2}{1+2\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) \left(\sqrt{\frac{1+\lambda^2}{1+2\lambda^2}} \right)^p \sqrt{2\pi}^p \sum_i |(\Gamma^{-1/2}w)_i|^2 \\
 &= |\det(\Gamma^{-1/2})| \frac{\lambda^2}{1+2\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) \left(\sqrt{\frac{1+\lambda^2}{1+2\lambda^2}} \right)^p \sqrt{2\pi}^p \|\Gamma^{-1/2}w\|_2^2 \\
 &\geq |\det(\Gamma^{-1/2})| \frac{\lambda^2}{1+2\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) \left(\sqrt{\frac{1+\lambda^2}{1+2\lambda^2}} \right)^p \sqrt{2\pi}^p \lambda_{\min}(\Gamma) \|w\|_2^2.
 \end{aligned} \tag{6.61}$$

This gives the result with $d_K \propto |\det(\Gamma^{-1/2})| \frac{\lambda^2}{1+2\lambda^2} \mathbf{C} \left(\frac{\Gamma}{1+\lambda^2} \right) \left(\sqrt{\frac{1+\lambda^2}{1+2\lambda^2}} \right)^p \sqrt{2\pi}^p \lambda_{\min}(\Gamma) > 0$. \square

LEMMA 6.3 Suppose $K(t) \propto e^{-\frac{1}{2}\lambda^2 \|t\|_F^2}$. Then there is an explicit strictly positive constants D_K depending on K such that

$$\|\partial_w \mu_0\|_K^2 \leq D_K \|w\|_2^2 \tag{6.62}$$

and there is an explicit strictly positive constants $D'_{A,K}$ depending on A and K such that

$$\|A \partial_w^2 \mu_0\|_2 \leq D'_{A,K} (\|w\|_2^2 + \|w\|_2) \tag{6.63}$$

where $\partial_w^2 \mu_0$ is the second derivative of μ_0 with respect to t in the direction w and

$$\begin{aligned}
 D_K &:= \int_{s \in \mathbb{R}^p} \int_{t \in \mathbb{R}^p} K(t-s) \|\Gamma^{-1}s\|_2 \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2} \|s\|_F^2} e^{-\frac{1}{2} \|t\|_F^2} dt ds \\
 D'_{A,K} &:= \sqrt{\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2} \max \left(\frac{1}{\lambda_{\min}(\Gamma)} \mathbf{C}(\Gamma), \int_{t \in \mathbb{R}^p} \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2} \|t\|_F^2} \right).
 \end{aligned} \tag{6.64}$$

Proof.

For the first bound, we have

$$\begin{aligned}
 \|\partial_w \mu_0\|_{\mathcal{H}}^2 &= \int_{s \in \mathbb{R}^p} \int_{t \in \mathbb{R}^p} K(t-s) w^T \Gamma^{-1} s w^T \Gamma^{-1} t e^{-\frac{1}{2} \|s\|_F^2} e^{-\frac{1}{2} \|t\|_F^2} dt ds \\
 &\leq \|w\|_2^2 \int_{s \in \mathbb{R}^p} \int_{t \in \mathbb{R}^p} K(t-s) \|\Gamma^{-1}s\|_2 \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2} \|s\|_F^2} e^{-\frac{1}{2} \|t\|_F^2} dt ds \\
 &= D_K \|w\|_2^2
 \end{aligned} \tag{6.65}$$

where $D_K := \int_{s \in \mathbb{R}^p} \int_{t \in \mathbb{R}^p} K(t-s) \|\Gamma^{-1}s\|_2 \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2}\|s\|_F^2} e^{-\frac{1}{2}\|t\|_F^2} dt ds$.

For the second bound, we have (with $\mathbf{C}(\Gamma)$ the normalization constant of Gaussian of covariance Γ),

$$\begin{aligned} \|A \partial_w^2 \mu_{t_0}\|_2^2 &= \sum_{l=1}^m \left| \int \alpha_l(t) d\partial_w^2 \mu_{t_0}(t) \right|^2 \\ &\leq \left(\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2 \right) \left(\int_{t \in \mathbb{R}^p} |d\partial_w^2 \mu_{t_0}(t)| \right)^2 \\ &= \left(\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2 \right) \left(\int_{t \in \mathbb{R}^p} \left| (-w^T \Gamma^{-1}w + w^T \Gamma^{-1}(t-t_0)) e^{-\frac{1}{2}\|t-t_0\|_F^2} \right| dt \right)^2. \end{aligned} \quad (6.66)$$

This gives

$$\begin{aligned} \|A \partial_w^2 \mu_{t_0}\|_2 &\leq \sqrt{\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2} \left(\mathbf{C}(\Gamma) \|w\|_F^2 + \int_{t \in \mathbb{R}^p} \left| w^T \Gamma^{-1}(t-t_0) e^{-\frac{1}{2}\|t-t_0\|_F^2} \right| dt \right) \\ &\leq \sqrt{\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2} \left(\mathbf{C}(\Gamma) \|w\|_F^2 + \|w\|_2 \int_{t \in \mathbb{R}^p} \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2}\|t\|_F^2} dt \right) \\ &\leq \sqrt{\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2} \left(\mathbf{C}(\Gamma) \frac{1}{\lambda_{\min}(\Gamma)} \|w\|_2^2 + \|w\|_2 \int_{t \in \mathbb{R}^p} \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2}\|t\|_F^2} dt \right) \\ &\leq D'_{A,K} (\|w\|_2^2 + \|w\|_2) \end{aligned} \quad (6.67)$$

where $D'_{A,K} := \sqrt{\sum_{l=1}^m \sup_{t \in \mathbb{R}^p} |\alpha_l(t)|^2} \max \left(\frac{1}{\lambda_{\min}(\Gamma)} \mathbf{C}(\Gamma), \int_{t \in \mathbb{R}^p} \|\Gamma^{-1}t\|_2 e^{-\frac{1}{2}\|t\|_F^2} dt \right)$. \square

Proof of Theorem 3.3. We prove this theorem by verifying the hypotheses of Corollary 2.1. We take

$\beta_1 = \beta_{GMM} \leq \frac{\sqrt{\lambda_{\min}(\Gamma)} \varepsilon}{8}$ and $\Sigma = \Sigma_{k, \frac{\varepsilon}{2}}$. We recall that we write $\theta^* = (a_1, \dots, a_k, t_1, \dots, t_k)$.

Technical assumption 2.2 number 1: Let $\theta = (b_1, \dots, b_k, s_1, \dots, s_k) \in \Lambda_{2\beta_1}$. Similarly to [34], $\|s_i - s_j\|_F = \|s_i - t_i + t_i - t_j + t_j - s_j\|_F \geq \|t_i - t_j\|_F - \|t_i - s_i\|_F - \|t_j - s_j\|_F \geq \varepsilon - \frac{1}{\sqrt{\lambda_{\min}(\Gamma)}} (\|t_i - s_i\|_2 + \|t_j - s_j\|_2) > \varepsilon - 2\varepsilon/4 \geq \varepsilon/2$. Moreover, using the hypothesis on the t_i , $\|s_i\|_2 \leq \|t_i\|_2 + 2\beta_{GMM} \leq R$ and $\phi(\theta) \in \Sigma_{k, \frac{\varepsilon}{2}}$.

Hypothesis 1: The RIP and the properties of K guarantee the unicity of $\phi(\theta^*)$ as a minimizer of (1.4). Now for $\theta \in \Lambda_{2\beta_1}$, the set $p(\theta, \theta^*)$ is included in the set of all the possible orderings of amplitudes and positions. The fact that $\|\cdot\|_K$ increases with respect to the distance between positions and that $\|t_i - s_i\|_F \leq \frac{\varepsilon}{4}$ imply $p(\theta, \theta^*) = \{\theta^*\}$.

Technical assumption 2.2 number 2: This hypothesis comes from the following. For $\sum_i a_i \mu_{t_i} - \sum_i b_i \mu_{s_i} \in \Sigma - \Sigma$,

$$\left\| \sum_i a_i \mu_{t_i} - \sum_i b_i \mu_{s_i} \right\|_K \leq \sum_i \|a_i \mu_{t_i} - b_i \mu_{s_i}\|_K. \quad (6.68)$$

Moreover,

$$\|a_i \mu_{t_i} - b_i \mu_{s_i}\|_K^2 = (a_i - b_i)^2 + 2a_i b_i (1 - \langle \mu_{t_i}, \mu_{s_i} \rangle_K). \quad (6.69)$$

It was shown in [22] that for a well designed Gaussian kernel, there is an explicit constant C_K such that $1 - \langle \mu_{t_i}, \mu_{s_i} \rangle_K \leq C_K \|t_i - s_i\|_2^2$ for $\|t_i - s_i\|_F^2 \leq \varepsilon/4$. Hence for $\theta \in \Lambda_{2\beta_1}$, using the fact that $|b_i| \leq$

$|a_i| + 2\beta_1 \leq 2|a_i|$, we have

$$\begin{aligned} \|\phi(\theta) - \phi(\theta^*)\|_K^2 &\leq \sum_i (a_i - b_i)^2 + 4|a_i|^2 C_K \|t_i - s_i\|_2^2 \\ &\leq \max(1, 4C_K a_k^2) \|\theta - \theta^*\|_2^2 \end{aligned} \quad (6.70)$$

Hence we can set the constant $C_{\phi, \theta^*} = \max(1, 4C_K a_k^2)$.

Technical assumption 2.2 number 3 and 4: These hypotheses come from the fact that the function g is infinitely differentiable on the bounded domain $\Lambda_{2\beta_1}$.

Hypothesis 3: Let $\theta \in \Lambda_{\beta_{GMM}}$ such that $\phi(\theta) = \sum_{i=1}^k b_i \mu_{s_i}$ and $u = (v, w) \in \mathbb{R}^d$ such that $\|u\|_2 = 1$. We have, using the kernel assumption,

$$\begin{aligned} \|\partial_u \phi(\theta)\|_K^2 &\geq (1 - c_K) \sum_i \|v_i \mu_{s_i} + b_i \partial_{w_i} \mu_{s_i}\|_K^2 \\ &= (1 - c_K) \sum_i (|v_i|^2 \|\mu_{s_i}\|_K^2 + |b_i|^2 \|\partial_{w_i} \mu_{s_i}\|_K^2 + 2b_i v_i \langle \mu_{s_i}, \partial_{w_i} \mu_{s_i} \rangle_K). \end{aligned} \quad (6.71)$$

We calculate the cross-product

$$\langle \mu_{s_i}, \partial_{w_i} \mu_{s_i} \rangle_K = -2 \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} K(t, s) w_i^T \Gamma^{-1}(t - s_i) e^{-\frac{1}{2} \|s - s_i\|_F^2} e^{-\frac{1}{2} \|t - s_i\|_F^2} dt ds. \quad (6.72)$$

The kernel K is written $K(t, s) \propto e^{-\lambda^2 \|t - s\|_F^2}$. With the translational invariance,

$$\begin{aligned} \langle \mu_{s_i}, \partial_{w_i} \mu_{s_i} \rangle_K &\propto \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} e^{-\frac{1}{2} \lambda^2 \|t - s\|_F^2} w_i^T \Gamma^{-1}(t - s_i) e^{-\frac{1}{2} \|s - s_i\|_F^2} e^{-\frac{1}{2} \|t - s_i\|_F^2} dt ds \\ &= \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} e^{-\frac{1}{2} \lambda^2 \|t - s\|_F^2} w_i^T \Gamma^{-1} t e^{-\frac{1}{2} \|s\|_F^2} e^{-\frac{1}{2} \|t\|_F^2} dt ds. \end{aligned} \quad (6.73)$$

Using identity (6.56) gives

$$\int e^{-\frac{1}{2} \lambda^2 \|t - s\|_F^2} e^{-\frac{1}{2} \|s\|_F^2} ds = \mathbf{C} \left(\frac{\Gamma}{1 + \lambda^2} \right) e^{-\frac{1}{2} \frac{\lambda^2}{1 + \lambda^2} \|t\|_F^2}. \quad (6.74)$$

where $\mathbf{C}(X) := \int_{\mathbb{R}^p} e^{-\frac{1}{2} \|s\|_X^2} ds$. Using the fact that $t \rightarrow h(t) = w^T \Gamma^{-1} t e^{-\frac{1}{2} \|t\|_F^2} e^{-\frac{1}{2} \frac{\lambda^2}{1 + \lambda^2} \|t\|_F^2}$ is an odd function of t , the integral with respect to t is zero and

$$\langle \mu_{s_i}, \partial_{w_i} \mu_{s_i} \rangle_K = 0. \quad (6.75)$$

Hence, using the assumption that $\|\mu_{s_i}\|_K^2 = 1$ and $\|(v, w)\|_2 = 1$,

$$\begin{aligned} \|\partial_u \phi(\theta)\|_K^2 &\geq (1 - c_K) \left(\sum_i |v_i|^2 \|\mu_{s_i}\|_K^2 + |b_i|^2 \|\partial_{w_i} \mu_{s_i}\|_K^2 \right) \\ &\geq (1 - c_K) \left(\sum_i |v_i|^2 + |b_i|^2 d_K \|w_i\|_2^2 \right) \\ &\geq (1 - c_K) \min(1, d_K \min_i (|b_i|^2)). \end{aligned} \quad (6.76)$$

where d_K such that $\|\partial_w \mu_0\|_K^2 \geq d_K \|w\|_2^2$ is given by Lemma 6.2.

We now bound

$$\|A\partial_u^2 \phi(\theta)\|_{\mathcal{H}} = \left\| \sum_i A\partial_u^2 b_i \mu_{s_i} \right\|_{\mathcal{H}}. \quad (6.77)$$

We have

$$\begin{aligned} \partial_u^2 b_i \mu_{s_i} &= \partial_u (v_i \mu_{s_i} + b_i \partial_{w_i} \mu_{s_i}) = v_i \partial_{w_i} \mu_{s_i} + b_i \partial_{w_i}^2 \mu_{s_i} + v_i \partial_{w_i} \mu_{s_i} \\ &= 2v_i \partial_{w_i} \mu_{s_i} + b_i \partial_{w_i}^2 \mu_{s_i}. \end{aligned} \quad (6.78)$$

Hence

$$\begin{aligned} \|A\partial_u^2 \phi(\theta)\|_{\mathcal{H}} &= \left\| \sum_i 2Av_i \partial_{w_i} \mu_{s_i} + b_i \partial_{w_i}^2 \mu_{s_i} \right\|_2 \\ &\leq 2\sqrt{1+\gamma} \left\| \sum_i v_i \partial_{w_i} \mu_{s_i} \right\|_{\mathcal{H}} + \left\| \sum_i b_i A \partial_{w_i}^2 \mu_{s_i} \right\|_2 \\ &\leq 2 \sum_i \sqrt{1+\gamma} \sqrt{1+c_k} \sqrt{\sum_i v_i^2 \|\partial_{w_i} \mu_{s_i}\|_{\mathcal{H}}^2} + |b_k| \sum_i \|A \partial_{w_i}^2 \mu_{s_i}\|_2 \end{aligned} \quad (6.79)$$

With Lemma 6.3, we have

$$\begin{aligned} \|A\partial_u^2 \phi(\theta)\|_{\mathcal{H}} &\leq 2\sqrt{1+\gamma} \sqrt{1+c_k} \sqrt{D_K} \|v\|_2 + |b_k| D'_{A,K} \sum_i (\|w_i\|_2^2 + \|w_i\|_2) \\ &\leq 2\sqrt{1+\gamma} \sqrt{1+c_k} \sqrt{D_K} \|v\|_2 + |b_k| D'_{A,K} (\|w\|_2^2 + \|w\|_2) \\ &\leq 2\sqrt{1+\gamma} \sqrt{1+c_k} \sqrt{D_K} + 2|b_k| D'_{A,K}. \end{aligned} \quad (6.80)$$

Using the fact that $\beta_{GMM} \leq \frac{|a_1|}{2}$, we have $|b_i| \geq |a_i| - \beta_{GMM} \geq |a_i| - \beta_{GMM} \geq \frac{|a_1|}{2}$ and $|b_k| \leq 2|a_k|$. This gives, for β_2 from Corollary 2.1,

$$\begin{aligned} \beta_2 &= \frac{1}{C_{\phi, \theta^*}} \inf_{\theta \in \Lambda_{\beta_1}} \inf_{z \in [\theta, \theta^*]} \left(\frac{(1-\gamma) \|\partial_{\theta^* - \theta} \phi(z)\|_{\mathcal{H}}^2}{\sqrt{1+\gamma} \|A\partial_{\theta^* - \theta}^2 \phi(z)\|_2} \right) \\ &\geq \frac{1}{C_{\phi, \theta^*}} \inf_{\theta \in \Lambda_{\beta_1}} \inf_{u: \|u\|_2=1} \left(\frac{(1-\gamma) \|\partial_u \phi(\theta)\|_{\mathcal{H}}^2}{\sqrt{1+\gamma} \|A\partial_u^2 \phi(\theta)\|_2} \right) \\ &\geq \frac{(1-\gamma)(1-c_K) \min(1, d_K |a_1|^2)}{8C_{\phi, \theta^*} \sqrt{1+\gamma} (\sqrt{1+\gamma} \sqrt{1+c_k} \sqrt{D_K} + 2|a_k| D'_{A,K})} > 0 \end{aligned} \quad (6.81)$$

and $\beta_{GMM} \leq \beta_2$ which implies $\Lambda_{\beta_{GMM}} \subset \Lambda_{\min(\beta_1, \beta_2)}$. Finally the set $\Lambda_{\beta_{GMM}}$ is a g-basin of attraction of θ^* . \square

6.7.1 Proofs for GMM with variable covariances. We write $Z_{i,k,l}$ the coordinate k, l of matrix Z_i . We have that $\partial_{Z_{i,k,l}} \mu_{t_i, \Gamma_i}$ has density

$$-\frac{1}{2} (\partial_{Z_{i,k,l}} \|t - t_i\|_{(Z_i Z_i^T + \rho I)}^2) e^{-\frac{1}{2} \|t - t_i\|_{\tilde{\Gamma}_i}^2}. \quad (6.82)$$

We also have, using the inverse matrix differentiation formula $\partial(A^{-1}) = -A^{-1}\partial(A)A^{-1}$

$$\partial_{Z_i, k, l} \|t - t_i\|_{(Z_i Z_i^T + \rho I)}^2 = -(t - t_i)^T (Z_i Z_i^T + \rho I)^{-1} \partial_{Z_i, k, l} (Z_i Z_i^T) (Z_i Z_i^T + \rho I)^{-1} (t - t_i). \quad (6.83)$$

Using the fact that $\partial_W (Z_i Z_i^T) = W Z_i^T + Z_i W^T$, we get the expression of the directional derivative $\partial_W \mu_{t_i, \Gamma_i}$.

Proof of Lemma 3.1. We write $(X_i)_{i=1}^p$ the diagonal terms of a matrix $X \in \mathbb{R}^{p \times p}$. With the translational invariance of K we have $\|\partial_W \mu_{t, \Gamma}\|_K^2 = \|\partial_W \mu_{0, \Gamma}\|_K^2$.

The partial derivative in direction W of $\mu_{0, \Gamma}$ has density

$$\begin{aligned} h(s) &= \frac{1}{2} s^T \Gamma^{-1} W \Gamma^{-1} s e^{-\frac{1}{2} \|s\|_{\Gamma}^2} \\ &= \frac{1}{2} s^T \Gamma^{-2} W s e^{-\frac{1}{2} \|s\|_{\Gamma}^2}, \end{aligned} \quad (6.84)$$

because diagonal matrices commute, and thus

$$\|\partial_W \mu_{0, \Gamma}\|_K^2 = \int_{\mathbb{R}^p} \int_{\mathbb{R}^p} K(s, t) h(s) h(t) ds dt. \quad (6.85)$$

The aim of the following computations is to give a lower bound on $\|\partial_W \mu_{0, \Gamma}\|_K^2$.

The rest of the proof is quite technical, and we thus split it into four steps for ease of reading.

Step 1: rewriting $\|\partial_W \mu_{0, \Gamma}\|_K^2$.

We have

$$\begin{aligned} \|s\|_{\Gamma}^2 + \lambda \|s - t\|_2^2 &= s^T (\Gamma^{-1} + \lambda I) s - 2\lambda \langle s, t \rangle + \lambda \|t\|_2^2 \\ &= s^T (\Gamma^{-1} + \lambda I) s - 2s^T (\Gamma^{-1} + \lambda I) (\Gamma^{-1} + \lambda I)^{-1} \lambda t + \lambda \|t\|_2^2 \\ &= \|s - \lambda (\Gamma^{-1} + \lambda I)^{-1} t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2 + \lambda \|t\|_2^2 - \|\lambda (\Gamma^{-1} + \lambda I)^{-1} t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2. \end{aligned} \quad (6.86)$$

Hence

$$\begin{aligned} \int_{s \in \mathbb{R}^p} K(s, t) h(s) ds &= \frac{1}{2} e^{-\frac{1}{2} \lambda \|t\|_2^2 + \frac{1}{2} \|\lambda (\Gamma^{-1} + \lambda I)^{-1} t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} \\ &\quad \times \int_{s \in \mathbb{R}^p} s^T \Gamma^{-2} W s e^{-\frac{1}{2} \|s - \lambda (\Gamma^{-1} + \lambda I)^{-1} t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} ds. \end{aligned} \quad (6.87)$$

We calculate, using the change of variable $s - \lambda (\Gamma^{-1} + \lambda I)^{-1} t \rightarrow s$,

$$\begin{aligned} B &:= \int_{s \in \mathbb{R}^p} s^T \Gamma^{-2} W s e^{-\frac{1}{2} \|s - \lambda (\Gamma^{-1} + \lambda I)^{-1} t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} ds \\ &= \int_{s \in \mathbb{R}^p} (s + \lambda (\Gamma^{-1} + \lambda I)^{-1} t)^T \Gamma^{-2} W (s + \lambda (\Gamma^{-1} + \lambda I)^{-1} t) e^{-\frac{1}{2} \|s\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} ds \\ &= \int_{s \in \mathbb{R}^p} s^T \Gamma^{-2} W s e^{-\frac{1}{2} \|s\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} ds \\ &\quad + 2 \int_{s \in \mathbb{R}^p} (\lambda (\Gamma^{-1} + \lambda I)^{-1} t)^T \Gamma^{-2} W s e^{-\frac{1}{2} \|s\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} ds \\ &\quad + (\lambda (\Gamma^{-1} + \lambda I)^{-1} t)^T \Gamma^{-2} W (\lambda (\Gamma^{-1} + \lambda I)^{-1} t) \int_{s \in \mathbb{R}^p} e^{-\frac{1}{2} \|s\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} ds. \end{aligned} \quad (6.88)$$

The second term is 0 because $s \rightarrow se^{-\|s\|_{(\Gamma^{-1}+\lambda I)^{-1}}^2}$ is odd. With $\mathbf{C}(X)$ the normalization constant of the Gaussian of covariance matrix X , and

$$\mathbf{D}(X, Y) = \int_{s \in \mathbb{R}^p} s^T X s e^{-\frac{1}{2}\|s\|_Y^2} ds, \quad (6.89)$$

we have

$$\begin{aligned} B &= \mathbf{D}(\Gamma^{-2}W, (\Gamma^{-1} + \lambda I)^{-1}) + \lambda^2 t^T \Gamma^{-2} (\Gamma^{-1} + \lambda I)^{-2} W t \mathbf{C}(\Gamma^{-1} + \lambda I)^{-1} \\ &= \mathbf{D}(\Gamma^{-2}W, (\Gamma^{-1} + \lambda I)^{-1}) + \lambda^2 t^T (I + \lambda \Gamma)^{-2} W t \mathbf{C}(\Gamma^{-1} + \lambda I)^{-1}. \end{aligned} \quad (6.90)$$

Going back to the full integral, (6.87) and (6.90) yield

$$\begin{aligned} \|\partial_W \mu_{0, \Gamma}\|_K^2 &= \frac{1}{4} \int_{t \in \mathbb{R}^p} (\mathbf{D}(\Gamma^{-2}W, (\Gamma^{-1} + \lambda I)^{-1}) + \lambda^2 t^T (I + \lambda \Gamma)^{-2} W t \mathbf{C}(\Gamma^{-1} + \lambda I)^{-1}) \\ &\quad e^{-\frac{1}{2}\lambda \|t\|_2^2 + \frac{1}{2}\|\lambda(\Gamma^{-1} + \lambda I)^{-1}t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2} t^T \Gamma^{-2} W t e^{-\frac{1}{2}\|t\|_T^2} dt. \end{aligned} \quad (6.91)$$

We have

$$\begin{aligned} &-\lambda \|t\|_2^2 + \|\lambda(\Gamma^{-1} + \lambda I)^{-1}t\|_{(\Gamma^{-1} + \lambda I)^{-1}}^2 - \|t\|_T^2 \\ &= -t^T (\Gamma^{-1} + \lambda I)t + t^T (\Gamma^{-1} + \lambda I)\lambda^2 (\Gamma^{-1} + \lambda I)^{-1} (\Gamma^{-1} + \lambda I)^{-1}t \\ &= -t^T (\Gamma^{-1} + \lambda I - \lambda^2 (\Gamma^{-1} + \lambda I)^{-1})t \\ &= -\|t\|_{(\Gamma^{-1} + \lambda I - \lambda^2 (\Gamma^{-1} + \lambda I)^{-1})}^2. \end{aligned} \quad (6.92)$$

Let $Z = (\Gamma^{-1} + \lambda I - \lambda^2 (\Gamma^{-1} + \lambda I)^{-1})^{-1}$, we have

$$Z_i^{-1} = \Gamma_i^{-1} + \lambda - \frac{\lambda^2}{\Gamma_i^{-1} + \lambda} = \frac{\Gamma_i^{-2} + 2\lambda \Gamma_i^{-1}}{\Gamma_i^{-1} + \lambda} = \frac{(\Gamma_i^{-1} + 2\lambda)\Gamma_i^{-1}}{\Gamma_i^{-1} + \lambda}. \quad (6.93)$$

Hence, we have the following equivalent when $\lambda \rightarrow +\infty$:

$$Z_i \sim \frac{\Gamma_i}{2}. \quad (6.94)$$

Let us set

$$\mathbf{E}(X, Y, Z) := \int s^T X s s^T Y s e^{-\frac{1}{2}\|s\|_Z^2} ds. \quad (6.95)$$

We then have

$$\begin{aligned} \|\partial_W \mu_{0, \Gamma}\|_K^2 &= \frac{1}{4} \mathbf{D}(\Gamma^{-2}W, (\Gamma^{-1} + \lambda I)^{-1}) \mathbf{D}(\Gamma^{-2}W, (\Gamma^{-1} + \lambda I - \lambda^2 (\Gamma^{-1} + \lambda I)^{-1})^{-1}) \\ &\quad + \frac{\lambda^2}{4} \mathbf{C}(\Gamma^{-1} + \lambda I)^{-1} \mathbf{E}(\Gamma^{-2}W, (I + \lambda \Gamma)^{-2}W, (\Gamma^{-1} + \lambda I - \lambda^2 (\Gamma^{-1} + \lambda I)^{-1})^{-1}). \end{aligned} \quad (6.96)$$

Step 2: dependency of \mathbf{D} on W

We explicit the dependency of \mathbf{D} on W :

$$\begin{aligned}
 \mathbf{D}(\Gamma^{-2}W, Y) &= \int_{s \in \mathbb{R}^p} s^T \Gamma^{-2} W s e^{-\frac{1}{2} \|s\|_Y^2} ds \\
 &= \int_{s \in \mathbb{R}^p} \left(\sum_i s_i^2 \Gamma_i^{-2} W_i \right) e^{-\frac{1}{2} \|s\|_Y^2} ds \\
 &= \sum_i \Gamma_i^{-2} W_i \int_{s \in \mathbb{R}^p} s_i^2 e^{-\frac{1}{2} \|s\|_Y^2} ds \\
 &= \sum_i \Gamma_i^{-2} W_i \int_{s \in \mathbb{R}^p} s_i^2 \prod_j e^{-\frac{1}{2} \frac{|s_j|^2}{Y_j}} ds \\
 &= \sum_i \Gamma_i^{-2} W_i \sqrt{\pi}^{p-1} \sqrt{\prod_{j \neq i} Y_j} \int_{s_i \in \mathbb{R}} s_i^2 e^{-\frac{1}{2} \frac{|s_i|^2}{Y_i}} ds_i.
 \end{aligned} \tag{6.97}$$

We make the change of variable $s_i = \sqrt{Y_i} u$ and use the fact that $\int_{\mathbb{R}} e^{-\frac{1}{2} u^2} du = \int_{\mathbb{R}} u^2 e^{-\frac{1}{2} u^2} du = \sqrt{2\pi}$. We get

$$\begin{aligned}
 \mathbf{D}(\Gamma^{-2}W, Y) &= \sum_i \Gamma_i^{-2} W_i \left(\sqrt{2\pi} \right)^{p-1} \sqrt{\prod_{j \neq i} Y_j} \int_{u \in \mathbb{R}} Y_i^{\frac{3}{2}} u^2 e^{-\frac{1}{2} |u|^2} du \\
 &= \left(\sqrt{2\pi} \right)^p \sqrt{\prod_j Y_j} \sum_i Y_i \Gamma_i^{-2} W_i.
 \end{aligned} \tag{6.98}$$

Step 3: dependency of \mathbf{E} on W

We calculate $\mathbf{E}(\Gamma^{-2}W, (I + \lambda\Gamma)^{-2}W, Z)$

$$\mathbf{E}(\Gamma^{-2}W, (I + \lambda\Gamma)^{-2}W, Z) = \sum_{i,j} (1 + \lambda\Gamma_i)^{-2} \Gamma_j^{-2} W_i W_j \int_{\mathbb{R}} s_i^2 s_j^2 e^{-\frac{1}{2} \|s\|_Z^2} ds. \tag{6.99}$$

We make the change of variable $s_i = \sqrt{Z_i} u_i$. We get

$$\begin{aligned}
 &\sum_{i \neq j} (1 + \lambda\Gamma_i)^{-2} \Gamma_j^{-2} W_i W_j \int_{\mathbb{R}} s_i^2 s_j^2 e^{-\frac{1}{2} \|s\|_Z^2} ds \\
 &= \left(\sqrt{2\pi} \right)^2 \left(\sqrt{2\pi} \right)^{p-2} \sqrt{\prod_i Z_i} \sum_{i \neq j} (1 + \lambda\Gamma_i)^{-2} Z_i Z_j \Gamma_j^{-2} W_i W_j \\
 &= \left(\sqrt{2\pi} \right)^p \sqrt{\prod_i Z_i} \sum_{i \neq j} (1 + \lambda\Gamma_i)^{-2} Z_i Z_j \Gamma_j^{-2} W_i W_j
 \end{aligned} \tag{6.100}$$

and, using the fact that $\int_{u \in \mathbb{R}} u^4 e^{-\frac{1}{2} |u|^2} du = 3\sqrt{2\pi}$, we get

$$\begin{aligned}
 \sum_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} W_i^2 \int_{\mathbb{R}} s_i^4 e^{-\frac{1}{2} \|s\|_Z^2} ds &= 3\sqrt{2\pi} \sqrt{2\pi}^{p-1} \sqrt{\prod_i Z_i} \sum_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 W_i^2 \\
 &= 3 \left(\sqrt{2\pi} \right)^p \sqrt{\prod_i Z_i} \sum_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 W_i^2.
 \end{aligned} \tag{6.101}$$

Hence

$$\begin{aligned}
\mathbf{E}(\Gamma^{-2}W, (I + \lambda\Gamma)^{-2}W, Z) &= \sqrt{\prod_i Z_i} \left(3 \left(\sqrt{2\pi} \right)^p \sum_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 W_i^2 \right. \\
&\quad - \left(\sqrt{2\pi} \right)^p \sum_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 W_i^2 \\
&\quad \left. + \left(\sqrt{2\pi} \right)^p \left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \right) \\
&= \sqrt{\prod_i Z_i} \left(\sqrt{2\pi} \right)^p \left(2 \sum_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 W_i^2 \right. \\
&\quad \left. + \left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \right).
\end{aligned} \tag{6.102}$$

Step 4: computing a lower bound for $\|\partial_W \mu_{0,\Gamma}\|_K^2$

Using $\|W\|_{\tilde{F}}^2 = 1$, we have

$$\begin{aligned}
\mathbf{E}(\Gamma^{-2}W, (I + \lambda\Gamma)^{-2}W, Z) \\
\geq \left(\sqrt{2\pi} \right)^p \sqrt{\prod_i Z_i} \left(2 \inf_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 + \left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \right).
\end{aligned} \tag{6.103}$$

Also, we have

$$\text{sign} \left(\left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \right) = \text{sign} \left(\lambda^2 \left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \right) \tag{6.104}$$

and, using that $Z_i^{-1} \rightarrow_{\lambda \rightarrow \infty} 2\Gamma_i^{-1}$ (from (6.93)),

$$\begin{aligned}
\lambda^2 \left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) &= \left(\sum_i \lambda^2 (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \\
&\rightarrow_{\lambda \rightarrow \infty} \frac{1}{4} \left(\sum_j \Gamma_j^{-1} W_j \right)^2 > 0.
\end{aligned} \tag{6.105}$$

Hence, for λ large enough, $\left(\sum_i (1 + \lambda\Gamma_i)^{-2} Z_i W_i \right) \left(\sum_j Z_j \Gamma_j^{-2} W_j \right) \geq 0$. This implies

$$\mathbf{E}(\Gamma^{-2}W, (I + \lambda\Gamma)^{-2}W, Z) \geq \sqrt{\prod_i Z_i} \left(\sqrt{2\pi} \right)^p 2 \inf_i (1 + \lambda\Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 \tag{6.106}$$

Moreover, writing $Y_i = (\Gamma_i^{-1} + \lambda I)^{-1}$ we have

$$\mathbf{C}(\Gamma^{-1} + \lambda I)^{-1} = \sqrt{2\pi}^p \sqrt{\prod_i Y_i}. \tag{6.107}$$

Putting everything together, for λ large enough, we deduce from (6.98) that

$$\begin{aligned}
\|\partial_W \mu_{0,\Gamma}\|_K^2 &\geq \frac{1}{4} (\sqrt{2\pi})^{2p} \sqrt{\prod_i Y_i} \sum_i Y_i \Gamma_i^{-2} W_i \sqrt{\prod_i Z_i} \sum_i Z_i \Gamma_i^{-2} W_i \\
&\quad + \frac{\lambda^2}{4} (\sqrt{2\pi})^{2p} \sqrt{\prod_i Y_i} \sqrt{\prod_i Z_i} \cdot 2 \inf_i (1 + \lambda \Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2 \\
&= \frac{1}{4} (2\pi)^p \sqrt{\prod_i Y_i Z_i} ((\sum_i Y_i \Gamma_i^{-2} W_i)(\sum_i Z_i \Gamma_i^{-2} W_i) + 2\lambda^2 \inf_i (1 + \lambda \Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2) \quad (6.108) \\
&\geq \frac{1}{4} (2\pi)^p \sqrt{\prod_i Y_i Z_i} (O(\frac{1}{\lambda}) + 2\lambda^2 \inf_i (1 + \lambda \Gamma_i)^{-2} \Gamma_i^{-2} Z_i^2) \\
&\sim_{\lambda \rightarrow \infty} \frac{1}{8} (2\pi)^p \sqrt{\prod_i Y_i Z_i} \inf_i \Gamma_i^{-2}.
\end{aligned}$$

Hence there is λ large enough such that $\|\partial_W \mu_{0,\Gamma}\|_K^2$ is lower bounded by a positive constant that depends on λ and Γ . \square

Acknowledgement

Y. Traonmilin acknowledges the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-20-CE40-0001 EFFIREG. J-F Aujol acknowledges the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-18-CE92-0050 SUPREMATIM.

References

- [1] ASIM, M., DANIELS, M., LEONG, O., AHMED, A. & HAND, P. (2020) Invertible generative models for inverse problems: mitigating representation error and dataset bias. in *International Conference on Machine Learning*, pp. 399–409. PMLR.
- [2] BARBER, R. F. & HA, W. (2018) Gradient descent with non-convex constraints: local concavity determines convergence. *Inf. Inference*, **7**(4), 755–806.
- [3] BAUSCHKE, H. H. & COMBETTES, P. L. (2011) *Convex analysis and monotone operator theory in Hilbert spaces*, vol. 408. Springer.
- [4] BECK, A. & TETRUASHVILI, L. (2013) On the Convergence of Block Coordinate Descent Type Methods. *SIAM J. Optim.*, **23**(4), 2037–2060.
- [5] BHOJANAPALLI, S., NEYSHABUR, B. & SREBRO, N. (2016) Global optimality of local search for low rank matrix recovery. in *Advances in Neural Information Processing Systems*, pp. 3873–3881.
- [6] BLUMENSATH, T. (2011) Sampling and reconstructing signals from a union of linear subspaces. *IEEE Trans. Inform. Theory*, **57**(7), 4660–4671.
- [7] BOUMAL, N., ABSIL, P.-A. & CARTIS, C. (2018) Global rates of convergence for nonconvex optimization on manifolds. *IMA J. Numer. Anal.*, **39**(1), 1–33.

- [8] BOURRIER, A., DAVIES, M., PELEG, T., PEREZ, P. & GRIBONVAL, R. (2014) Fundamental Performance Limits for Ideal Decoders in High-Dimensional Linear Inverse Problems. *Information Theory, IEEE Transactions on*, **60**(12), 7928–7946.
- [9] BURER, S. & MONTEIRO, R. D. (2005) Local minima and convergence in low-rank semidefinite programming. *Math. Program.*, **103**(3), 427–444.
- [10] CAMBARERI, V. & JACQUES, L. (2018) Through the haze: a non-convex approach to blind gain calibration for linear random sensing models. *Inf. Inference*.
- [11] CANDÈS, E. J. & FERNANDEZ-GRANDA, C. (2013) Super-resolution from noisy data. . *Fourier Anal. Appl.*, **19**(6), 1229–1254.
- [12] CANDÈS, E. J. & PLAN, Y. (2011) Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Trans. Inform. Theory*, **57**(4), 2342–2359.
- [13] CHI, Y., LU, Y. M. & CHEN, Y. (2019) Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Trans. Signal Process.*, **67**(20), 5239–5269.
- [14] CHIZAT, L. & BACH, F. (2018) On the global convergence of gradient descent for over-parameterized models using optimal transport. in *Advances in neural information processing systems*, pp. 3036–3046.
- [15] CIARLET, P. G., MIARA, B. & THOMAS, J.-M. (1989) *Introduction to numerical linear algebra and optimisation*. Cambridge University Press.
- [16] DE CASTRO, Y., GAMBOA, F., HENRION, D. & LASSERRE, J.-B. (2016) Exact solutions to super resolution on semi-algebraic domains in higher dimensions. *IEEE Trans. Inform. Theory*, **63**(1), 621–630.
- [17] DUVAL, V. & PEYRÉ, G. (2015) Exact support recovery for sparse spikes deconvolution. *Found. Comput. Math.*, **15**(5), 1315–1355.
- [18] EFTEKHARI, A. & WAKIN, M. B. (2015) New analysis of manifold embeddings and signal recovery from compressive measurements. *Appl. Comput. Harmon. Anal.*, **39**(1), 67–109.
- [19] EKELAND, I. & TEMAM, R. (1999) *Convex analysis and variational problems*, vol. 28. SIAM.
- [20] ELVIRA, C., GRIBONVAL, R., SOUSSEN, C. & HERZET, C. (2019) OMP and continuous dictionaries: is k-step recovery possible?. in *ICASSP 2019 - IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 5546–5550.
- [21] FOU CART, S. & RAUHUT, H. (2013) *A mathematical introduction to compressive sensing*. Springer.
- [22] GRIBONVAL, R., BLANCHARD, G., KERIVEN, N. & TRAONMILIN, Y. (2021a) Compressive Statistical Learning with Random Feature Moments. *Math. Stat. Learn.*, *In press*.
- [23] GRIBONVAL, R., BLANCHARD, G., KERIVEN, N. & TRAONMILIN, Y. (2021b) Statistical Learning Guarantees for Compressive Clustering and Compressive Mixture Modeling. *Math. Stat. Learn.*, *In press*.

- [24] HÖRMANDER, L. (2015) *The analysis of linear partial differential operators I: Distribution theory and Fourier analysis*. Springer.
- [25] KERIVEN, N., BOURRIER, A., GRIBONVAL, R. & PÉREZ, P. (2018) Sketching for large-scale learning of mixture models. *Inf. Inference*, **7**(3), 447–508.
- [26] KERIVEN, N., TREMBLAY, N., TRAONMILIN, Y. & GRIBONVAL, R. (2017) Compressive K-means. in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pp. 6369–6373. IEEE.
- [27] LI, Q., ZHU, Z. & TANG, G. (2018) The non-convex geometry of low-rank matrix optimization. *Inf. Inference*, **8**(1), 51–96.
- [28] LI, S. & LI, Q. (2022) Local and Global Convergence of General Burer-Monteiro Tensor Optimizations. *arXiv preprint 2201.02298*.
- [29] LING, S. & STROHMER, T. (2017) Regularized gradient descent: a non-convex recipe for fast joint blind deconvolution and demixing. *Inf. Inference*.
- [30] PUY, G., DAVIES, M. E. & GRIBONVAL, R. (2017) Recipes for Stable Linear Embeddings From Hilbert Spaces to \mathbb{R}^m . *IEEE Trans. Inform. Theory*, **63**(4), 2171–2187.
- [31] RAUHUT, H., SCHNEIDER, R. & STOJANAC, Ž. (2015) Tensor completion in hierarchical tensor representations. in *Compressed sensing and its applications*, pp. 419–450. Springer.
- [32] SHI, H., TRAONMILIN, Y. & AUJOL, J.-F. (2021) Compressive learning for patch-based image denoising. *Preprint HAL <https://hal.archives-ouvertes.fr/hal-03429102>*.
- [33] SHI, H., TRAONMILIN, Y. & AUJOL, J.-F. (2021) Sketched learning for image denoising. in *The Eighth International Conference on Scale Space and Variational Methods in Computer Vision (SSVM)*, Cabourg, France.
- [34] TRAONMILIN, Y. & AUJOL, J.-F. (2020) The basins of attraction of the global minimizers of the non-convex sparse spike estimation problem. *Inverse Problems*, **36**(4), 045003.
- [35] TRAONMILIN, Y., AUJOL, J.-F. & LECLAIRE, A. (2020) Projected gradient descent for non-convex sparse spike estimation. *IEEE Signal Processing Letters*.
- [36] TRAONMILIN, Y. & GRIBONVAL, R. (2018) Stable recovery of low-dimensional cones in Hilbert spaces: One RIP to rule them all. *Appl. Comput. Harmon. Anal.*, **45**(1), 170 – 205.
- [37] TRAONMILIN, Y. & VAITER, S. (2018) Optimality of 1-norm regularization among weighted 1-norms for sparse recovery: a case study on how to find optimal regularizations. *Journal of Physics: Conference Series*, **1131**, 012009.
- [38] TRAONMILIN, Y., VAITER, S. & GRIBONVAL, R. (2018) Is the 1-norm the best convex sparse regularization?. in *iTWIST'18 - international Traveling Workshop on Interactions between low-complexity data models and Sensing Techniques*, Proceedings of iTWIST'18.
- [39] TU, S., BOCZAR, R., SIMCHOWITZ, M., SOLTANOLKOTABI, M. & RECHT, B. (2016) Low-rank Solutions of Linear Matrix Equations via Procrustes Flow. *Proceedings of The 33rd International Conference on Machine Learning*, **48**, 964–973.

- [40] UNSER, M. & FAGEOT, J. (2019) Native Banach spaces for splines and variational inverse problems. *arXiv preprint arXiv:1904.10818*.
- [41] WALDSPURGER, I. (2018) Phase retrieval with random Gaussian sensing vectors by alternating projections. *IEEE Trans. Inform. Theory*.
- [42] ZHAO, T., WANG, Z. & LIU, H. (2015) A nonconvex optimization framework for low rank matrix estimation. in *Advances in Neural Information Processing Systems*, pp. 559–567.