



HAL
open science

Compressive Recovery of Sparse Precision Matrices

Titouan Vayer, Etienne Lasalle, Rémi Gribonval, Paulo Gonçalves

► **To cite this version:**

Titouan Vayer, Etienne Lasalle, Rémi Gribonval, Paulo Gonçalves. Compressive Recovery of Sparse Precision Matrices. 2023. hal-04275341v2

HAL Id: hal-04275341

<https://hal.science/hal-04275341v2>

Preprint submitted on 11 Dec 2023 (v2), last revised 12 Dec 2023 (v3)

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.

Compressive Recovery of Sparse Precision Matrices

TITOUAN VAYER*, ETIENNE LASALLE, RÉMI GRIBONVAL AND PAULO GONÇALVES

Univ Lyon, Inria, CNRS, ENS de Lyon, UCB Lyon 1,
LIP UMR 5668, F-69342 Lyon, France

*Corresponding author: titouan.vayer@ens-lyon.fr

We consider the problem of learning a graph modeling the statistical relations of the d variables from a dataset with n samples $\mathbf{X} \in \mathbb{R}^{n \times d}$. Standard approaches amount to searching for a precision matrix Θ representative of a Gaussian graphical model that adequately explains the data. However, most maximum likelihood-based estimators usually require storing the d^2 values of the empirical covariance matrix, which can become prohibitive in a high-dimensional setting. In this work, we adopt a “compressive” viewpoint and aim to estimate a sparse Θ from a *sketch* of the data, *i.e.* a low-dimensional vector of size $m \ll d^2$ carefully designed from \mathbf{X} using non-linear random features. Under certain assumptions on the spectrum of Θ (or its condition number), we show that it is possible to estimate it from a sketch of size $m = \Omega((d + 2k) \log(d))$ where k is the maximal number of edges of the underlying graph. These information-theoretic guarantees are inspired by compressed sensing theory and involve restricted isometry properties and instance optimal decoders. We investigate the possibility of achieving practical recovery with an iterative algorithm based on the graphical lasso, viewed as a specific denoiser. We compare our approach and graphical lasso on synthetic datasets, demonstrating its favorable performance even when the dataset is compressed.

Keywords: inverse problem; graph inference; compressive learning; graphical lasso.

1. Introduction

Inferring a complex network from data is used in many applications, such as in neuroscience for the treatment of epilepsy [32], for biological networks [82] or in genomics to identify gene interactions [56, 88]. We consider in this paper the problem of estimating a certain graph representing the statistical relations between d variables from n observed signals $\mathbf{x}_1, \dots, \mathbf{x}_n$ where $\mathbf{x}_i \in \mathbb{R}^d$ follow a certain distribution μ , assumed to be centered for simplicity and with covariance matrix Σ . This graph is often associated with the so-called *precision matrix* $\Theta = \Sigma^{-1}$ which is sparse in many practical situations due to limited conditional dependencies [31]. However, simply inverting the empirical covariance matrix $\widehat{\Sigma}$ usually does not yield a sparse estimation of the precision matrix estimation. The challenge is thus to infer a sparse precision matrix from $\widehat{\Sigma}$. One of the most popular methods to do this graph estimation is probably the graphical lasso [2, 31]. Given the empirical covariance matrix

$$\widehat{\Sigma} \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top, \quad (1.1)$$

the graphical lasso estimator searches for a certain matrix $\widehat{\Theta}_{\text{GL}} \in \mathbb{R}^{d \times d}$, representative of the graph edges, that solves the optimization problem

$$\widehat{\Theta}_{\text{GL}} \triangleq \arg \min_{\Theta > 0} -\log \det(\Theta) + \langle \widehat{\Sigma}, \Theta \rangle + \lambda \|\Theta\|_{1, \text{off}}, \quad (1.2)$$




FIG. 1. Summary of our approach. The objective is to estimate a graph $\Theta \in \mathbb{R}^{d \times d}$ between the d^2 variables (i, j) of the data set from a sketch of the data, $\mathbf{s} \triangleq \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) \in \mathbb{R}^m$. We will show in the following sections that, under certain graph sparsity assumptions, it is theoretically possible to estimate Θ using a sketch of size $m \approx (d + 2k) \log(d) \ll d^2$ where k is the maximal number of edges of the underlying graph.

where the arg min is taken over the set of symmetric positive definite matrices and $\|\Theta\|_{1,\text{off}} \triangleq \sum_{i < j} |\Theta_{ij}|$ promotes sparsity. An interpretation is as follows: for a Gaussian model $\mu = \mathcal{N}(0, \Sigma = \Theta^{-1})$, equation (1.2) corresponds to an ℓ_1 -penalized maximum likelihood estimator [86]. In other words, the graphical lasso estimates a sparse graph that best fits the data. More precisely, in the Gaussian model, the pattern of zeros of the precision matrix Θ corresponds to conditional independencies among the variables, hence Θ encodes the statistical relations between them [25, 53]. Despite its many good properties, the graphical lasso suffers from a scaling problem with respect to the dimension. Indeed computing (1.2) requires to store in memory the entire empirical covariance matrix $\hat{\Sigma}$ and thus has a space complexity of $\mathcal{O}(d^2)$; this is problematic for applications where d is very large, such as high-resolution fMRI datasets or gene-microarray data where the number of genes d is typically around tens of thousands.

Graphical models estimation is a very active field of research and many alternatives to graphical lasso have been proposed in the literature. Numerous strategies aim to address the computational scalability of the optimization problem (1.2) through the use of approximation techniques, *e.g.* QUIC [41], BIG & QUIC [42], SQUIC [7]. Others are looking for estimators that have either better statistical guarantees or better algorithmic properties. We can mention some estimators that rely on non-convex penalties [28, 69, 71, 80], ℓ_2 penalties [50] or other estimators like CLIME (and Dantzig-like estimators) [13, 85], RobustCLIME [77], D-trace [87] or Elem-GGM [81]. Finally, some approaches make assumptions on the underlying graph (*e.g.* that it is chordal) in order to find fast algorithms or try to constrain the structure of the graph sought in the estimation [26, 51, 83].

1.1. Compressive learning of graphical models

In this work we take another path: that of compression. Inspired by the theories of compressed sensing [30] and sketching [35], we try to estimate a sparse precision matrix Θ associated to a covariance matrix $\Sigma = \Theta^{-1}$, not from the whole dataset, nor from the empirical covariance matrix $\hat{\Sigma}$, but from a *compressed version* of the data called *sketch*. More precisely, and given a well-chosen function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}^m$, we seek to estimate Θ from the vector $\mathbf{s} \in \mathbb{R}^m$ defined by (see Figure 1)

$$\mathbf{s} \triangleq \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) \in \mathbb{R}^m. \quad (1.3)$$

In sketching theory the Φ function is typically an adequately chosen *random function* and the reconstruction guarantees are usually based on ideas from compressed sensing [33]. The advantages of this

type of approach are multiple: 1) when $m \ll nd$ the dataset is massively compressed, with benefits for storage and transfer 2) as averages, the sketches can be computed online, in one pass on the data, or in a distributed manner 3) as the data get aggregated, sketches have also good properties regarding differential privacy [17]. The main objective of this paper is to use the sketching approach and to answer the following question:

Objective *Can we find Φ and a sketch dimension $m \ll d^2$ such that the sparse precision matrix Θ can be well estimated from a sketch of the data $\mathbf{s} = \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) \in \mathbb{R}^m$?*

In other words, can we obtain an estimate of the graph structure using a small-sized vector compared to the size of the empirical covariance matrix? We will shortly furnish a precise definition of what we consider to be “well estimated”. However, the underlying intuition guiding this inquiry is that, similar to the graphical lasso model, graphs estimated from data typically demonstrate sparsity. Therefore, there is hope that good approximations of these graphs can be obtained by accessing only a limited number of measurements compared to $\mathcal{O}(d^2)$, the total number of elements in the covariance matrix.

1.2. Contributions

The article presents the following key contributions:

1. We investigate, from an information-theoretic perspective, a sketching mechanism where we project against *i.i.d.* rank-one matrices. We show that when $m \gtrsim \|\Theta\|_0 \log(d)$, where $\|\Theta\|_0$ denotes the number of non-zero coefficients of Θ , the relevant information is preserved during the sketching phase. More precisely, we prove that the corresponding sketching operator satisfies a *Restricted Isometry Property* (RIP), ensuring that the precision matrix can be reconstructed accurately using a so-called robust decoder.
2. This robust decoder being *a priori* intractable, we propose a practical decoding scheme to estimate the sparse precision matrix from the sketch of the data. The algorithm relies on an iterative procedure that alternates between a gradient descent step associated to a sketch fidelity term and a “denoising step” performed by the graphical lasso.
3. For practical applications, we propose to deviate from the theoretical framework of 1) and leverage structured matrices (*e.g.*, Walsh-Hadamard matrices) to define a similar sketching mechanism but with boosted algorithmic efficiency. We demonstrate in the experiments that the combination of this sketching mechanism and our practical decoding scheme is able to significantly compress the data (*i.e.*, $m \ll d^2$) while still retaining the ability to recover the covariance and precision matrices.

1.3. Organization

The paper is organized as follows. Section 2 outlines the main information-theoretic results. The main theorem (Theorem 2.2) proves that the sketching operator satisfies a certain RIP with high probability. For an easier read, the technical part required to prove this theorem is deferred to the end of the paper, Section 5, where we obtain the control of covering numbers and a concentration inequality for the sketching operator. In Section 3, we present a random *structured* matrix approach for the sketching phase (*i.e.*, encoding phase) with better memory and computational efficiency. We also provide a practical algorithm to retrieve the precision matrix (*i.e.*, decoding phase) from the sketch. Section 4 presents the experiments on the encoding and decoding phases.

1.4. Notations and definitions.

For any integer m , $\llbracket m \rrbracket \triangleq \{1, \dots, m\}$. $S_d(\mathbb{R})$ is the linear space of symmetric matrices of size $d \times d$, while $S_d^{++}(\mathbb{R})$ (resp. $S_d^+(\mathbb{R})$) is the set of symmetric positive definite matrices (resp. positive semi-definite). We often write $\mathbf{A} \in S_d^{++}(\mathbb{R})$ as $\mathbf{A} \succ 0$ (resp. $\mathbf{A} \succeq 0$ for $S_d^+(\mathbb{R})$). We denote by $\|\cdot\|_{2 \rightarrow 2}$ the 2-operator norm or spectral norm for matrices i.e. $\|\mathbf{A}\|_{2 \rightarrow 2} \triangleq \sup_{\|\mathbf{x}\|_2 \leq 1} \|\mathbf{A}\mathbf{x}\|_2 = \sigma_{\max}(\mathbf{A})$ where $\sigma_{\max}(\mathbf{A})$ denotes the largest singular value of \mathbf{A} . The spectrum of $\mathbf{A} \in S_d(\mathbb{R})$ is denoted by $\text{spec}(\mathbf{A})$. The standard euclidean scalar product on matrices is denoted by $\langle \cdot, \cdot \rangle$ and its associated norm (the Frobenius norm) is denoted by $\|\cdot\|_{\text{Fro}}$. We denote by $\|\mathbf{A}\|_0$ the number of non-zeros components of \mathbf{A} , that is to say the ℓ_0 “norm” of \mathbf{A} . We also introduce the function $\text{inv}(\mathbf{A}) \triangleq \mathbf{A}^{-1}$.

2. Recovery guarantees

The sketching procedure proposed in this article is based on *rank-one projections* and (well-chosen) *randomness*. Recall that we are given a sample of d -dimensional vectors $(\mathbf{x}_1, \dots, \mathbf{x}_n)$, and we want to define a function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}^m$ such that the sketch is the average of the m -dimensional vectors $\Phi(\mathbf{x}_i)$ as in (1.3). Given a collection of d -dimensional random vectors $(\mathbf{a}_1, \dots, \mathbf{a}_m)$ we consider in this work¹

$$\Phi(\mathbf{x}) \triangleq \frac{1}{m} \left(|\mathbf{a}_1^\top \mathbf{x}|^2, |\mathbf{a}_2^\top \mathbf{x}|^2, \dots, |\mathbf{a}_m^\top \mathbf{x}|^2 \right)^\top. \quad (2.1)$$

Before proceeding, it is important to note that Φ can also be computed as $\Phi(\mathbf{x}) = \frac{1}{m} \{\mathbf{a}_j^\top \mathbf{x} \mathbf{x}^\top \mathbf{a}_j\}_{j \in \llbracket m \rrbracket} = \frac{1}{m} \{\langle \mathbf{a}_j \mathbf{a}_j^\top, \mathbf{x} \mathbf{x}^\top \rangle\}_{j \in \llbracket m \rrbracket}$. Therefore, by linearity of the inner product and equations (1.1) and (1.3), the induced sketch \mathbf{s} is the vector whose coordinates are the projections of the empirical covariance matrix $\widehat{\Sigma}$ against the rank-one matrices $\mathbf{a}_j \mathbf{a}_j^\top$: $\mathbf{s} = \frac{1}{m} \{\langle \mathbf{a}_j \mathbf{a}_j^\top, \widehat{\Sigma} \rangle\}_{j \in \llbracket m \rrbracket}$.

Consequently, the entry point of our theoretical analysis is to notice that the random function Φ defined in (2.1) involves a *linear operator* on symmetric matrices [24]. Indeed, our objective can be reformulated as finding $\Theta = \Sigma^{-1}$ from $\mathbf{s} = \mathcal{A}(\widehat{\Sigma})$ where $\mathcal{A}: S_d(\mathbb{R}) \rightarrow \mathbb{R}^m$ is defined by

$$\mathcal{A}(\mathbf{M}) \triangleq \frac{1}{m} \left\{ \langle \mathbf{a}_j \mathbf{a}_j^\top, \mathbf{M} \rangle \right\}_{j \in \llbracket m \rrbracket}. \quad (2.2)$$

We can therefore reformulate our problem as a compressed sensing problem: given a noisy linear measurement of the true covariance matrix $\mathbf{s} = \mathcal{A}(\Sigma) + \mathbf{e} \in \mathbb{R}^m$, where $\mathbf{e} = \mathcal{A}(\widehat{\Sigma} - \Sigma)$ is the noise, we want to recover the precision matrix $\Theta = \Sigma^{-1}$ that underlies certain sparsity assumptions. There are two difficulties here. First $m < d^2$, thus the problem is *a priori* ill-posed. Second, the sparsity assumption does not involve the “signal” Σ itself but a non-linear transformation of it, given by its inverse $\Theta = \Sigma^{-1}$.

Remark 1 (Related works) *This sketching mechanism appears in various works on compressed sensing, in particular for low-rank matrix completion [12, 14, 19, 47, 49], estimation of low-rank & positive semi-definite (PSD) matrices [55, 70], low-rank & sparse matrices [1] or in statistical learning for PCA [33] and in phase retrieval [46]. Here we propose to study it for the recovery of sparse precision matrices.*

¹ Readers familiar with sketching methods might be puzzled by the factor $1/m$ rather than the usual $1/\sqrt{m}$. In our case, as we will consider the ℓ_1 -norm on \mathbb{R}^m , it is in fact the right scaling to obtain averages.

Remark 2 (Dense Gaussian projections) *We want to mention that another compressive scheme, maybe more natural for theoretical analysis, could have been considered. Instead of projecting against rank-one matrices, one could project against symmetric matrices with independent Gaussian entries via $\Phi(\mathbf{x}) = \{\mathbf{x}^\top \mathbf{A}_j \mathbf{x}\}_{j \in [m]}$ where $\mathbf{A}_j \in \mathbb{R}^{d \times d}$ are dense Gaussian matrices. The theoretical analysis of this approach would be quite straightforward following classical compressed sensing works [30, 62]. However, this sketching mechanism would be far from efficient as the computation of $\widehat{\Phi}(\mathbf{x})$ would have an $\mathcal{O}(md^2)$ time and space complexity, which is even greater than directly storing $\widehat{\Sigma}$. Nevertheless, it is worth noting that optical computing, which relies on optical processing units, could be leveraged to compute these random features in constant time in any dimension [68], albeit constrained by the current hardware capabilities.*

2.1. Recipe for recovery guarantees

The general idea to obtain guarantees is the following: we will assume that Θ lies in a “low dimensional” set so that it is possible to estimate it from a very small number of measurements compared to the ambient dimension d^2 [9]. This hypothesis is formalized here by assuming that the true covariance matrix Σ belongs to a certain subset of $S_d^{++}(\mathbb{R})$, called *model set*. More precisely, and given $k \in \mathbb{N}, 0 < a \leq b$, we will study here the model set

$$\mathfrak{S}_{k,a,b} \triangleq \{\Sigma \in S_d^{++}(\mathbb{R}) : \Theta = \Sigma^{-1} \succ 0, \|\Theta\|_0 \leq d + 2k \text{ and } \text{spec}(\Theta) \subseteq [a, b]\}. \quad (2.3)$$

This set corresponds to covariance matrices associated to the $d + 2k$ -sparse precision matrices whose spectra are well localized in $[a, b]$. This constraint on the spectrum of the precision matrix can be relaxed to a constraint on its condition number as we will see later. Similar spectral conditions are usually assumed in the case of graphical lasso to derive sample complexities of the estimators [13, 57, 78]. The number k represents the maximal number of non-zero components of Θ on its upper triangular part, *i.e.*, the maximal number of edges of the graph corresponding to Θ , without counting self-loops. The sparsity of the graph is also quite natural in many applications where one wants to have a simple graph thus with few connections.

We will see later that $\mathfrak{S}_{k,a,b}$ is the image by inv of a “low dimensional” set and, by using a certain notion of stability of the inverse, we will be able to recover Θ from \mathbf{s} . We emphasize that while we have presented this specific model set, the majority of the results in this paper are general and can be extended to accommodate other model sets. Consequently, these results can apply to alternative assumptions concerning the underlying graph structure.

The central property for our guarantees is the *Restricted Isometric Property* (RIP) [15, 62], adapted to our context. In the following, we consider two general norms $\|\cdot\|_{S_d}$ and $\|\cdot\|_{\mathbb{R}^m}$ on S_d and \mathbb{R}^m , respectively.

Definition 1 (RIP) *A linear operator $\mathcal{A} : (S_d(\mathbb{R}), \|\cdot\|_{S_d}) \rightarrow (\mathbb{R}^m, \|\cdot\|_{\mathbb{R}^m})$ satisfies the RIP_δ for some $\delta \in [0, 1[$ on a set $\mathfrak{C} \subset S_d^{++}(\mathbb{R})$ if for every $(\Sigma_1, \Sigma_2) \in \mathfrak{C}^2$, it satisfies*

$$(1 - \delta)\|\Sigma_1 - \Sigma_2\|_{S_d} \leq \|\mathcal{A}(\Sigma_1 - \Sigma_2)\|_{\mathbb{R}^m} \leq (1 + \delta)\|\Sigma_1 - \Sigma_2\|_{S_d}. \quad (2.4)$$

We can show that, when the linear operator \mathcal{A} satisfies the RIP_δ , we can find a *decoder* that is robust to noise and that will allow us to recover Σ from \mathbf{s} . More precisely, following standard results in compressed sensing [30, 48], we have the following result (the proof is given in Appendix A.1.1 for

completeness): Let $\mathcal{A} : (S_d(\mathbb{R}), \|\cdot\|_{S_d}) \rightarrow (\mathbb{R}^m, \|\cdot\|_{\mathbb{R}^m})$ be a linear operator. Suppose that \mathcal{A} satisfies RIP_δ on a model set $\mathfrak{G} \subset S_d(\mathbb{R})$ and consider the decoder $\Delta : \mathbb{R}^m \rightarrow \mathfrak{G}$ defined by

$$\Delta[\mathbf{s}] \in \arg \min_{\Sigma \in \mathfrak{G}} \|\mathcal{A}(\Sigma) - \mathbf{s}\|_{\mathbb{R}^m}. \quad (2.5)$$

Suppose that $\mathbf{x}_1, \dots, \mathbf{x}_n \stackrel{i.i.d}{\sim} \mu$ where μ is a centered probability distribution with covariance $\Sigma = \mathbb{E}_{\mathbf{x} \sim \mu} [\mathbf{x}\mathbf{x}^\top] \succ 0$. Consider $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$ the empirical covariance matrix and $\mathbf{s} = \mathcal{A}(\widehat{\Sigma})$. If $\Sigma \in \mathfrak{G}$, then $\Sigma^* \triangleq \Delta[\mathbf{s}]$ satisfies

$$\|\Sigma^* - \Sigma\|_{S_d} \leq \frac{2}{1 - \delta} \|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma)\|_{\mathbb{R}^m}. \quad (2.6)$$

When $\Sigma \notin \mathfrak{G}$, the bound (2.6) still holds up to an additional ‘‘modeling error’’ term $d^\circ(\Sigma, \mathfrak{G})$ which quantifies the distance from Σ to \mathfrak{G} (see Appendix A.1.1). We always assume that the minimization problem (2.5) has at least one solution. The decoder can be adjusted as in [9], to handle the case where the argmin is only approximated to a certain accuracy. The estimator given by (2.5) has many desirable properties: 1) it is robust to noise as shown in [33, 48] 2) it allows a recovery of Σ from \mathbf{s} as the number of samples n grows. Indeed $\|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma)\|_{\mathbb{R}^m} \leq (\sup_{\|\mathbf{U}\|_{S_d}=1} \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m}) \|\widehat{\Sigma} - \Sigma\|_{S_d}$ and $\|\widehat{\Sigma} - \Sigma\|_{S_d}$ typically behaves as $\mathcal{O}(n^{-1/2})$ under standard sub-Gaussian assumptions [76, Section 6]. Thus, solutions of the optimization problem (2.5) theoretically yield good approximations of Σ . As an additional outcome, by leveraging the regularity of the inverse mapping for matrices with bounded spectra, this estimation of Σ also results in a reliable estimate of Θ . We refer to the discussion below Theorem 2.2 for a more precise statement.

With these remarks in mind, the strategy is now to obtain this RIP assumption, which enables the information-theoretic decoder and recovery guarantees. By the linearity of \mathcal{A} , we can observe that obtaining RIP_δ is equivalent to showing that, for $0 \leq \delta < 1$,

$$\sup_{\mathbf{U} \in S[\mathfrak{G}]} \left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1 \right| < \delta, \quad (2.7)$$

where we set

$$S[\mathfrak{G}] \triangleq \left\{ \frac{\Sigma_1 - \Sigma_2}{\|\Sigma_1 - \Sigma_2\|_{S_d}} : (\Sigma_1, \Sigma_2) \in \mathfrak{G}^2, \|\Sigma_1 - \Sigma_2\|_{S_d} > 0 \right\}. \quad (2.8)$$

This set is called the *normalized secant set* of \mathfrak{G} [62]. We will provide a more detailed description of the space $S[\mathfrak{G}]$ later on. For now, we present the classical framework that will enable us to prove the RIP_δ property of \mathcal{A} .

The operator \mathcal{A} being random we will show that, given a sufficient (but reasonable) dimension m , we can have a control of type (2.7) with high probability. In the following we denote by $\mathcal{N}(S[\mathfrak{G}], \|\cdot\|_{S_d}, \varepsilon)$ the *covering number* of $S[\mathfrak{G}]$ which, informally, quantifies the effective ‘‘dimension’’ of this set (see Section 5.1 for more details). The following theorem (whose proof is deferred to Appendix A.1.2) describes the main ingredients for establishing RIP_δ .

Consider a random sketching operator $\mathcal{A} : S_d(\mathbb{R}) \rightarrow \mathbb{R}^m$ and denote its operator norm by $\|\|\mathcal{A}\|\| \triangleq \sup_{\|\mathbf{U}\|_{S_d}=1} \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m}$. Suppose that we are given two functions $C_1, C_2 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$\forall t > 0, \mathbb{P}(\|\|\mathcal{A}\|\| > t) \leq C_1(t), \quad (2.9)$$

$$\forall \mathbf{U} \in \mathcal{S}[\mathfrak{S}], \forall t > 0, \mathbb{P}\left(\left|\|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1\right| > t\right) \leq C_2(t). \quad (2.10)$$

Then, for any $\varepsilon > 0$ and $\delta \in [0, 1[$,

$$\sup_{\mathbf{U} \in \mathcal{S}[\mathfrak{S}]} \left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1 \right| < \delta, \quad (2.11)$$

with probability at least $1 - \mathcal{N}(\mathcal{S}[\mathfrak{S}], \|\cdot\|_{S_d}, \varepsilon) C_2(\frac{\delta}{2}) - C_1(\frac{\delta}{2\varepsilon})$. Consequently with the same probability the operator \mathcal{A} satisfies RIP_δ on \mathfrak{S} .

Remark 3 *The above theorem is presented in its most usual form, with both controls (2.9) and (2.10). However, (2.10) is sometimes easier to obtain (by leveraging classical concentration inequalities). Fortunately, we can obtain (2.9) from (2.10), as for $\varepsilon' > 0$, defining C_1 by*

$$C_1(t) = \mathcal{N}(B_{S_d}, \|\cdot\|_{S_d}, \varepsilon') C_2((1 - \varepsilon')t - 1) \quad \forall t > 1/(1 - \varepsilon'), \quad (2.12)$$

where $\mathcal{N}(B_{S_d}, \|\cdot\|_{S_d}, \varepsilon')$ is the covering number of the unit sphere $B_{S_d} = \{\mathbf{U} \in S_d : \|\mathbf{U}\|_{S_d} = 1\}$, yields a valid upper-bound of $\mathbb{P}(\|\mathcal{A}\| > t)$. See Appendix A.1.3 for the proof.

2.2. Main theoretical result

Guided by Theorem 2.1 we need to control the covering number of $\mathcal{S}[\mathfrak{S}_{k,a,b}]$ as well as the concentration of the random operator \mathcal{A} to obtain the desired RIP_δ with high probability. To do so, the randomness related to the vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ needs to be specified. We consider in this section two probability distributions for the vectors $\mathbf{a}_j \stackrel{i.i.d.}{\sim} \Lambda$, one being $\Lambda = \Lambda_G \stackrel{\Delta}{=} \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ the multivariate Gaussian distribution with covariance matrix $d^{-1}\mathbf{I}_d$ and the other $\Lambda = \Lambda_U \stackrel{\Delta}{=} \mathcal{U}(\mathbb{S}^{d-1})$ the uniform distribution on the hypersphere. The scaling of the random vectors \mathbf{a}_j is such that they have unit expected square Euclidean norm under both distributions. These distributions are classic choices when it comes to random projections. We envision that our results remain valid for any choice of sub-Gaussian rotational-invariant distribution. Both norms $\|\cdot\|_{S_d}$ and $\|\cdot\|_{\mathbb{R}^m}$ in Theorem 2.1 must also be defined. For the former, we choose a norm that is *specific* to how our random operator \mathcal{A} is drawn. Indeed, the probability distribution $\Lambda \in \{\Lambda_G, \Lambda_U\}$ induces a norm on $S_d(\mathbb{R})$ defined by

$$\forall \mathbf{M} \in S_d(\mathbb{R}), \|\mathbf{M}\|_\Lambda \stackrel{\Delta}{=} \mathbb{E}_{\mathbf{a} \sim \Lambda} \left[\left| \langle \mathbf{a}\mathbf{a}^\top, \mathbf{M} \rangle \right| \right] = \mathbb{E}_{\mathbf{a} \sim \Lambda} \left[\left| \mathbf{a}^\top \mathbf{M} \mathbf{a} \right| \right]. \quad (2.13)$$

This choice is adapted to our problem when $\|\cdot\|_{\mathbb{R}^m} = \|\cdot\|_1$ since $\mathbb{E}_{\mathbf{a}_j \sim \Lambda} [\|\mathcal{A}(\mathbf{M})\|_1] = \|\mathbf{M}\|_\Lambda$. Therefore, if for every $\mathbf{U} \in \mathcal{S}[\mathfrak{S}]$, $\|\mathcal{A}(\mathbf{U})\|_1$ well concentrates around its expectation, then it has a high probability of being close to its expectation $\|\mathbf{U}\|_\Lambda = 1$, and thus (2.11) will hold.

From these choices, we show that with a reasonable value of m the rank-one operators \mathcal{A} satisfy the RIP on $\mathfrak{S}_{k,a,b}$. It is formalized in the theorem below which is the main theoretical result of the paper. We advise readers that are interested in the proof to refer to Section 5 where the results on the covering numbers (Section 5.1) and the concentration of \mathcal{A} (Section 5.2) are derived.

Let $\mathcal{A} : (S_d(\mathbb{R}), \|\cdot\|_\Lambda) \rightarrow (\mathbb{R}^m, \|\cdot\|_1)$ be a rank-one projection operator as defined in (2.2), with $(\mathbf{a}_j)_j$ either Gaussian or uniform ($\Lambda \in \{\Lambda_G, \Lambda_U\}$). For all $\delta, \rho \in]0, 1[$, there exists $C = C(\delta, \rho, b/a)$, independent of m, k and d , such that, whenever

$$m \geq C(d + 2k) \log d, \quad (2.14)$$

the operator \mathcal{A} satisfies RIP_δ on $\mathfrak{S}_{k,a,b}$ with probability at least $1 - \rho$. In particular the following holds uniformly on $\Sigma \in \mathfrak{S}_{k,a,b}$ with probability at least $1 - \rho$: Let $\mathbf{x}_1, \dots, \mathbf{x}_n \sim \mu$ with μ a centered probability

distribution with covariance Σ , $\widehat{\Sigma}$ the empirical covariance matrix and $\mathbf{s} = \mathcal{A}(\widehat{\Sigma})$ a sketch of the data. The estimator $\Sigma^* = \Delta[\mathbf{s}]$ defined in (2.5) satisfies

$$\|\Sigma^* - \Sigma\|_{\Lambda} \leq \frac{2}{1-\delta} \|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma)\|_1. \quad (2.15)$$

A more precise condition on m can be found in the proof of this theorem in Appendix A.3.3 (in particular see (A.46)). This theorem shows that our sketching operators satisfies the RIP with high probability with a sketching dimension $m \gtrsim (d+2k)\log d$, which is much smaller than the d^2 required to recover a matrix of size $d \times d$. Let us mention that although in (2.15) the error between the estimator Σ^* and the true covariance matrix Σ is measured with the unusual Λ -norm, we can have the same type of control for $\|\Sigma^* - \Sigma\|_{\text{Fro}}$, at the expense of a multiplicative constant $1/c_{\text{Fro}} = \frac{9\sqrt{15}}{2}d \leq 18d$ (see Proposition 4). In addition, (2.15) also provides guarantees for the recovery of the precision matrix Θ . By making use of the bounded spectra of $(\Sigma^*)^{-1}$ and Θ , we can exploit the regularity of the inverse map to obtain $\|(\Sigma^*)^{-1} - \Theta\|_{\text{Fro}} \leq b^2 \|\Sigma^* - \Sigma\|_{\text{Fro}}$. Hence, whenever (2.15) is verified, we also have

$$\|(\Sigma^*)^{-1} - \Theta\|_{\text{Fro}} \leq \frac{9\sqrt{15} b^2 d}{(1-\delta)} \|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma)\|_1. \quad (2.16)$$

Remark 4 (A similar result with an unbounded model set.) *We want to point out that in (A.46), the condition on m depends on a and b only through the ratio b/a . This might suggest that the fundamental quantity to consider is the ratio between the largest and smallest eigenvalues, i.e., the condition number of the matrix, instead of the eigenvalues themselves. In Appendix A.6, we introduce a model set $\mathfrak{S}_{k,\kappa_0}$, where precision matrices have a condition number bounded by some $\kappa_0 \geq 1$. Even though $\mathfrak{S}_{k,\kappa_0}$ is much “bigger” than $\mathfrak{S}_{k,a,b}$, as the latter is bounded and the former is not, we obtain an upper-bound on the covering number of its normalized secant $S[\mathfrak{S}_{k,\kappa_0}]$ (see Theorem A.6.3 in Appendix A.6). This is sufficient to derive a result similar to Theorem 2.2 with the same $m \gtrsim (d+2k)\log d$ condition, yielding guarantees on the recovery of Σ with only bounded condition numbers rather than bounded spectra.*

2.3. Connection to prior works

Theorem 2.2 indicates that it is theoretically possible to recover Θ , with a $\mathcal{O}(n^{-1/2})$ error, by keeping in memory only a single sketch of size $m \gtrsim (d+2k)\log(d)$. To the best of our knowledge this is the first result for compressive recovery of sparse precision matrices from empirical covariance measurements. It can however be put in perspective with [58] which tries to find the support of Θ by observing, in an adaptive manner, only a small fraction of the entries of the *true* covariance Σ . The authors show that it is possible to find Θ from $\mathcal{O}(d \text{polylog}(d))$ elements of Σ with some assumptions on the graph (small treewidth). Interestingly enough, we are able to obtain the same type of guarantees but with the big difference that, in our case, we observe a non-adaptive compressed version of the empirical covariance that takes the whole matrix into account, which is more in line with concrete applications.

2.4. Practical limitations of Theorem 2.2

Theorem 2.2 indicates that the sketch contains the necessary information to recover the true covariance matrix using the decoder

$$\Delta[\mathbf{s}] \in \arg \min_{\Sigma \in \mathfrak{S}_{k,a,b}} \|\mathcal{A}(\Sigma) - \mathbf{s}\|_{\mathbb{R}^m}. \quad (2.17)$$

However, even though this theorem holds theoretical significance, its practical application encounters two significant challenges when employing the decoder (2.17) and the sketching operator with *independent* rank-one projections.

- To retrieve Θ from the sketch, one needs to use \mathcal{A} and thus to store the m d -dimensional vectors $(\mathbf{a}_j)_{j \in \llbracket m \rrbracket}$. Hence, the overall memory cost comprises a reasonable $\mathcal{O}(m)$ expense to store the sketch \mathbf{s} but also an additional $\mathcal{O}(md)$ cost to store the $(\mathbf{a}_j)_{j \in \llbracket m \rrbracket}$. Unfortunately, according to Theorem 2.2, the sketch size must be $m \gtrsim d \log(d)$, yielding a total memory cost larger than $\mathcal{O}(d^2)$, similar to the cost of storing the empirical covariance matrix.
- Directly solving the optimization problem (2.17) is probably intractable as the constraint $\Sigma \in \mathfrak{S}_{k,a,b}$ implies to search among the matrices Σ such that $\|\Sigma^{-1}\|_0 \leq (d + 2k)$ which is a highly non-convex constraint.

3. Towards practical recovery

As a remedy to the above-mentioned issues, we propose in this section to use structured matrices for efficient sketching and a tractable approximate decoder as a more practical alternative to (2.17).

3.1. Structured matrices for efficient sketching

Using structured random matrices is a recurrent idea in compressive learning [18, 54, 84]. It reduces the degrees of freedom while mimicking the behavior of random matrices with *i.i.d.* columns. In the following, we assume that $d = 2^K$ is a power of 2 and that $m = B \times d$ is a multiple of d . Padding strategies can be implemented when these requirements are not met, but we leave this technicality out of the scope of this paper for the sake of simplicity and refer the interested reader to [17]. Here, we adopt the same approach as [18, 84]. Instead of independent random vectors, we define $(\mathbf{a}_1, \dots, \mathbf{a}_m)$ as the columns of the random matrix $\mathbf{A} = (\mathbf{B}_1 | \dots | \mathbf{B}_B) \in \mathbb{R}^{d \times m}$ made of B independent structured blocks $\mathbf{B}_l \in \mathbb{R}^{d \times d}$ defined as triple-Rademacher matrices:

$$\mathbf{B}_l \triangleq \frac{1}{d^{3/2}} \mathbf{H} \mathbf{D}_l^{(1)} \mathbf{H} \mathbf{D}_l^{(2)} \mathbf{H} \mathbf{D}_l^{(3)}. \quad (3.1)$$

The matrix \mathbf{H} denotes the Walsh-Hadamard matrix with entries in $\{\pm 1\}$ and $\mathbf{D}_l^{(k)}$ are independent random diagonal “sign-flipping” matrices, *i.e.*, random diagonal matrices with independent Rademacher entries. The scaling factors $d^{3/2}$ yields $\|\mathbf{a}_j\|_2 = 1$. We emphasize that this strategy results in the use of *dependent* random vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$.

This approach has a dual advantage: it improves both memory and computational efficiency. Indeed, notice that $\Phi(\mathbf{x})$ can be expressed as $\Phi(\mathbf{x}) = \{(\mathbf{A}^\top \mathbf{x})_j^2\}_{j \in \llbracket m \rrbracket} = (\mathbf{A}^\top \mathbf{x}) \odot (\mathbf{A}^\top \mathbf{x})$ where \odot denotes the Hadamard product. Thus the memory and computational bottleneck *a priori* lies in the storage of \mathbf{B}_l and the computations of $\mathbf{B}_l^\top \mathbf{x}$ for $l \in \llbracket B \rrbracket$. Fortunately, we benefit from the properties of the fast Walsh-Hadamard transform: 1) it computes $\mathbf{H}^\top \mathbf{y}$ for any vector \mathbf{y} without storing \mathbf{H} as it is hard-coded into the algorithm, 2) the matrix-vector multiplication $\mathbf{H}^\top \mathbf{y}$ requires only $\mathcal{O}(d \log d)$ operations. As a consequence for our sketching mechanism, only the diagonal matrices $\mathbf{D}_l^{(k)}$ must be stored. This reduces the space complexity of storing the dense matrix $\mathbf{A} \in \mathbb{R}^{d \times m}$ from $\mathcal{O}(md)$ to $\mathcal{O}(m)$ [29, 84]. Moreover, point 2) implies that $\Phi(\mathbf{x})$ can be computed with a time complexity of $\mathcal{O}(m \log d)$ as each of the B matrix-vector multiplications $\mathbf{B}_l^\top \mathbf{x}$ requires $\mathcal{O}(d \log d)$ operations with the fast Walsh-Hadamard transform.

Algorithm 1 Algorithm for solving the decoding problem

-
- 1: Input: Sketching operator \mathcal{A} and sketch of the data $\mathbf{s} = \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) = \mathcal{A}(\widehat{\Sigma}) \in \mathbb{R}^m$.
 - 2: Initial guess $\Sigma_0 \succ 0$, regularization parameter $\lambda > 0$, step size $\gamma > 0$, and maximum number of iterations t_{\max} .
 - 3: **for** $t \in 0, 1 \dots, t_{\max} - 1$ **do**
 - 4: $\Sigma_{t+\frac{1}{2}} \leftarrow \Sigma_t - \gamma \mathcal{A}^*(\mathcal{A}(\Sigma_t) - \mathbf{s})$
 - 5: $\Sigma_{t+1} \leftarrow \text{GLASSO}_{\lambda, \gamma}[\Sigma_{t+\frac{1}{2}}]$ // with standard graphical lasso solver
 - 6: **end for**
 - 7: **return** estimated covariance $\Sigma_{t_{\max}}$ and precision $\Sigma_{t_{\max}}^{-1}$.
-

Theoretical examination of recovery guarantees for these structured sketching operators is deferred to future research. To achieve the same reconstruction guarantees as in the independent case, the key lies in verifying the concentration of \mathcal{A} in the structured case, which amounts to proving a result similar to Proposition 3. We expect that guarantees comparable to those in the independent case can be obtained.

3.2. Algorithmic solution to decoding

Finally, we present a heuristic algorithmic approach to obtain the precision matrix from a sketch of the data. We will demonstrate that a computationally simpler alternative decoder to (2.17) works effectively in practice.

The graphical lasso as a denoiser Our algorithmic solution is inspired by the connections between proximal operators and denoisers in the context of inverse problems. Numerous studies have indeed demonstrated that proximal operators can be regarded as efficient denoisers [21, 22, 44, 45]. The core concept behind our approach thus revolves around utilizing a decoder that relies on the graphical lasso (1.2), which not only benefits from efficient algorithms but also carries the interpretation of a proximal operator [6]. Subsequently, we introduce

$$\text{GLASSO}_{\lambda}[\mathbf{Z}] \stackrel{\Delta}{=} \underset{\Sigma \succ 0}{\text{argmin}} \quad -\log \det \Sigma^{-1} + \langle \mathbf{Z}, \Sigma^{-1} \rangle + \lambda \|\Sigma^{-1}\|_{1, \text{off}}, \quad (3.2)$$

where we recall that $\|\mathbf{M}\|_{1, \text{off}} = \sum_{i < j} |M_{ij}|$. This operator is equivalent to the one introduced in (1.2), which computes the precision matrix instead of the covariance matrix. Indeed, although (3.2) is non-convex, a solution can be computed by solving the convex graphical lasso problem (1.2) with the change of variable $\Theta = \Sigma^{-1}$. We also emphasize that most graphical lasso solvers such as [2, 31, 59] compute both the covariance Σ and the precision $\Theta = \Sigma^{-1}$, without calculating the inverse, by relying instead on duality theory. We argue that the operator (3.2), when applied to the empirical covariance of the data $\widehat{\Sigma}$, can be interpreted as a “denoiser” of $\widehat{\Sigma}$ in the sense that it returns a covariance matrix whose inverse is sparse. This interpretation is motivated by the fact that the graphical lasso can be seen as a specific proximal operator in the framework of *Bregman divergences* [11, 16] which we now briefly describe (we refer to [72] for a more precise discussion). Let \mathcal{H} be a Hilbert space and $h : \mathcal{H} \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, convex and differentiable on its open domain $\text{dom}(h)$. The Bregman divergence associated to h is given by

$$\forall \mathbf{x}, \mathbf{y} \in \mathcal{H} \times \mathcal{H}, D_h(\mathbf{x}|\mathbf{y}) = \begin{cases} h(\mathbf{x}) - h(\mathbf{y}) - \langle \nabla h(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle & \text{if } \mathbf{y} \in \text{dom}(h), \\ +\infty & \text{otherwise.} \end{cases} \quad (3.3)$$

The function D_h measures a similarity or “distance” between the points in \mathcal{H} . For instance if $\mathcal{H} = \mathbb{R}^d$ and $h = \frac{1}{2} \|\cdot\|_2^2$ then D_h is simply $D_h(\mathbf{x}|\mathbf{y}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2$. For a function $\varphi : \mathcal{H} \rightarrow \mathbb{R}$, the so-called (*left*) *Bregman proximal operator* of φ is defined as $\text{prox}_\varphi^h(\mathbf{z}) \triangleq \arg \min_{\mathbf{x}} \varphi(\mathbf{x}) + D_h(\mathbf{x}|\mathbf{z})$ [72, Definition 2.3]. It generalizes the standard Euclidean proximal operator that can be computed with $h = \frac{1}{2} \|\cdot\|_2^2$. To relate with the context of the graphical lasso, we introduce the function $h(\mathbf{X}) = -\log \det \mathbf{X}$ if $\mathbf{X} \succ 0$ and $h(\mathbf{X}) = +\infty$ otherwise. It is strictly convex, continuously differentiable over its domain, and $\nabla h(\mathbf{X}) = -\mathbf{X}^{-1}$ [10, Appendix A.4.1]. The associated Bregman divergence writes (see *e.g.* [61, 64])

$$\forall \mathbf{X}, \mathbf{Y} \in S_d^{++}(\mathbb{R}), D_h(\mathbf{X}|\mathbf{Y}) \triangleq -\log \det(\mathbf{X}) + \log \det \mathbf{Y} + \langle \mathbf{Y}^{-1}, \mathbf{X} \rangle - d. \quad (3.4)$$

Based on the definition of D_h we can rewrite the operator (3.2) as $\text{GLASSO}_\lambda[\mathbf{Z}] = \arg \min_{\substack{\Sigma \\ \Sigma \succ 0}} \lambda \|\Sigma^{-1}\|_{1,\text{off}} + D_h(\mathbf{Z}|\Sigma)$. The graphical lasso can thus be interpreted as a proximal Bregman operator of the function $\varphi : \Sigma \mapsto \lambda \|\Sigma^{-1}\|_{1,\text{off}}$ but by operating on the *right variable* of the Bregman divergence. Although not previously studied with the graphical lasso, this type of operators, known as *right proximity operator* [4, 5, 52], has nevertheless been considered in the context of Poisson inverse problems [3], for image restoration problems [79, Section 4] or in [36] where it was shown to admit a characterization in terms of gradient of a convex function. This discussion highlights that $\text{GLASSO}_\lambda[\widehat{\Sigma}]$ computes a covariance matrix candidate which is both close to $\widehat{\Sigma}$ in the sense of D_h and whose inverse tends to be sparse.

Algorithm Inspired by this interpretation of the graphical lasso as a denoiser, we present an iterative algorithm for estimating the true covariance from the sketch. In the following, we consider the data fidelity term $f(\Sigma) \triangleq \frac{1}{2} \|\mathcal{A}(\Sigma) - \mathbf{s}\|_2^2$ whose gradient is $\nabla f(\Sigma) = \mathcal{A}^*(\mathcal{A}(\Sigma) - \mathbf{s})$ where

$$\mathcal{A}^* : \mathbf{y} \in \mathbb{R}^m \mapsto \frac{1}{m} \sum_{j=1}^m y_j \mathbf{a}_j \mathbf{a}_j^\top \in S_d \quad (3.5)$$

is the adjoint operator of \mathcal{A} . Given a initial estimate $\Sigma_0 \succ 0$ and a step-size $\gamma > 0$, our proposed algorithmic solution computes

$$\forall t \in \{0, \dots, t_{\max}\}, \Sigma_{t+1} = \text{GLASSO}_{\gamma\lambda}[\Sigma_t - \gamma \nabla f(\Sigma_t)]. \quad (3.6)$$

In other words, we alternate between a gradient step $\Sigma_{t+\frac{1}{2}} = \Sigma_t - \gamma \nabla f(\Sigma_t)$ in the direction of minimizing f and a denoising/proximal Bregman step $\Sigma_{t+1} = \text{GLASSO}_{\gamma\lambda}[\Sigma_{t+\frac{1}{2}}]$ (with parameter $\gamma\lambda$) in the vein of standard forward-backward algorithms. The overall procedure is summarized in Algorithm 1. The computational bottleneck of this algorithm is the graphical lasso step that has a computational $\mathcal{O}(d^3)$ complexity with standard graphical lasso solvers such as [2, 31, 59], which rely on the dual formulation of the graphical lasso, or [66] which rely on an iterative thresholding procedure. Although more efficient solvers exist [7, 41, 42] we choose to solve graphical lasso with the `scikit-learn` implementation [60]. It relies on the block coordinate procedure described in [59] and finds the solution of the graphical lasso with cubic complexity. It is noteworthy that the iterations of our algorithm, as described in (3.6), can be linked to Bregman Proximal Gradient (BPG) descent [3] with the Bregman divergence D_h . In Appendix A.4, we demonstrate that the iterations in (3.6) correspond to those of a BPG algorithm, albeit with a Riemannian gradient instead of the usual Euclidean gradient in (3.6). In practice, we observe that this algorithm always converges when $\gamma > 0$ is sufficiently small to ensure that the iterates $(\Sigma_{t+\frac{1}{2}})_t$ remain positive definite (we give a safe step-size strategy ensuring this condition in Appendix A.5). The

intriguing questions regarding the convergence rate of this algorithm and the minimizers associated with it are left for future research. We envision the use of guarantees inspired by Plug-and-Play literature, as described *e.g.* in [67]. However, we experimentally demonstrate in the next section that the associated estimator effectively recovers a precision matrix from a sketch of the data.

4. Experiments

In this section we provide experiments for assessing the efficiency of Algorithm 1. The following experiments are conducted using *structured sketching* as described in Section 3.1. The objectives of these experiments is to answer the following questions:

1. Does the decoder described in Algorithm 1 give qualitatively coherent results?
2. What is the impact of the sketch size m on the final estimation? More precisely, can we obtain a good estimate of the true precision matrix with a number of measurements m close to the theoretical $m_0 = C(d + 2k) \log(d)$ obtained in Theorem 2.2?
3. How does the sketching approach combined with the decoder described in Algorithm 1 compare to classical methods such as the graphical lasso in terms of performance?

In all experiments, the setting is as follows: we generate a true sparse precision matrix $\Theta \in S_d^{++}(\mathbb{R})$ and we consider $\mathbf{x}_1, \dots, \mathbf{x}_n \sim \mathcal{N}(0, \Theta^{-1})$ *i.e.* n i.i.d. samples of a multivariate Gaussian with covariance $\Sigma = \Theta^{-1}$. In the experiments, we explore two methods for generating sparse precision matrices Θ . For each method, the matrix Θ first consists of L blocks, each having a size of $M \times M$ with $L \times M = d$. The generation of each block follows the distribution of a random graph. In the first method, referred to as `ERDOS`, each block follows the distribution of an Erdős-Rényi graph [27] where the probability of connection is set to $p = 0.2$. In the second method, referred to as `POWERLAW`, the random graph is a tree with a power-law degree distribution. Both methods utilize the `Networkx` library [37] for generating these distributions. After fixing the support with the above strategy, the value of each coefficient is arbitrarily set to εu where $u \sim \text{Unif}[1, 4]$ and $\varepsilon = 1$ with probability 0.5 and -1 with probability 0.5. The matrix Θ is then symmetrized and made positive definite by adding a sufficiently large diagonal $(0.1 + \lambda_{\min})\mathbf{I}$. Finally a random permutation permutes the rows and columns of the matrix. Two examples of matrices Θ generated according to these procedures are shown on the left side of Figure 2 (for $d = 64$).

In the experiments, we consider two performance measures. The first one is the relative error computed as

$$\text{RE} \triangleq \frac{\|\Theta_{\text{true}} - \Theta_{\text{esti}}\|_{\text{Fro}}}{\|\Theta_{\text{true}}\|_{\text{Fro}}} \quad (4.1)$$

and the second one is the $F_1 \in [0, 1]$ score between the true matrix and the estimated one. It is calculated as $F_1 = \frac{2\text{tp}}{2\text{tp} + \text{fp} + \text{fn}}$ where true positive `tp` stands for the case when there is an actual edge and the algorithm detects it; false positive `fp` stands for the case when there is no actual edge but the algorithm detects one, and false negative `fn` stands for the case when the algorithm failed to detect an actual edge.

4.1. First illustration

First, we qualitatively illustrate the behavior of our Algorithm 1. In this experiment, we set $d = 64$ and generate $n = 8000$ samples from Θ using the `ERDOS` and `POWERLAW` procedures. The number of non-zero elements are respectively 512 for `POWERLAW` and 386 for `ERDOS`. We consider one draw of structured sketching operator as described in Section 3.1. We compute a sketch of the dataset $\mathbf{s} = \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) = \mathcal{A}(\widehat{\Sigma}) \in \mathbb{R}^m$ and set the number of measurements for the sketch to $m = 1536$. In other




FIG. 2. Illustrative example of estimation using the decoding procedure described in Algorithm 1. The dimension is $d = 64$ and the sketch size $m = 1536$. We set $\lambda = 0.008$ for the ℓ_1 penalty parameter for all solvers. RE stands for the relative error between the true precision matrix and the estimated one (4.1). **(First row)** With a precision matrix Θ generated from the Erdos process. **(Second row)** With a matrix Θ generated from the PowerLaw process. The columns are (from left to right): the empirical covariance $\hat{\Sigma}$ for $n = 8000$, the true precision matrix Θ and the decoder Algorithm 1 based on a sketch of the data with structured rank-one projections.

words, the final sketch has a number of measurements that is approximately 75% the degrees of freedom of $d \times d$ symmetric matrices, which is $d(d+1)/2 = 2080$. We fix the number of iterations to $t_{\max} = 3500$, the step size to $\gamma = 0.005$ and $\lambda = 0.008$. The results are depicted in Figure 2. We can observe that in both cases, our method provides a visually consistent estimation of the precision matrix.

4.2. Impact of the number of measurements on the estimation - asymptotic regime

In order to quantitatively evaluate the impact of the number of measurements m on the final estimation we consider the asymptotic regime where $n = +\infty$ or equivalently we sketch the true covariance matrix as $\mathbf{s} = \mathcal{A}(\Sigma)$ instead of the empirical covariance matrix. We take $d = 256$ and we draw three precision matrices from the settings PowerLaw and Erdos. For each setting PowerLaw and Erdos, each $m \in \{256, 512, 1024, 2048, 4096, 8192, 16384, 32768 = d^2/2\}$, and each score function (RE and F_1 -score) we choose the parameter $\lambda \in \{1e-4, 5e-4, 1e-3, \dots, 5e-1\}$ that gives the best average score. We report the results in Figure 3.




FIG. 3. Impact of the number of measurements on the final estimation for the datasets Erdos and PowerLaw. **(Left)** the best average relative error on the three draws of the true precision matrix for the PowerLaw and Erdos settings. **(Right)** the best average F_1 score. The 10-th percentile for these scores are reported in shaded line. Vertical colored dashed lines are located at $\hat{k} \log(d)$ where \hat{k} is the average number of non-zero elements $\|\Theta\|_0$ of the precision matrices for each setting.

We observe that for Erdos the estimator based on the sketching procedure and Algorithm 1 gives a recovery with a relative error that is below 10% starting from $m \approx 4096$ that is $\frac{m}{d^2/2} \approx 12\%$. This corresponds to a compression rate of approximately 88%. This result is also consistent with the theoretical bound $m \approx \|\Theta\|_0 \log(d)$ (represented by the vertical dashed lines) found in Theorem 2.2: we can see that the best average relative error is below 10% starting from this limit. From $\frac{m}{d^2/2} = 50\%$, the recovery is also nearly perfect. The F_1 score is in agreement with the relative error: starting from $m \approx 4096$, the F_1 score is above 0.8, indicating that our estimator captures the correct statistical relationships between the variables.

For the PowerLaw setting, the results are slightly less favorable, and a relative error below 10% is only achieved for $\frac{m}{d^2/2} \approx 50\%$, resulting in a compression rate of approximately 50%. This can be explained by the fact that the precision matrices in this case are less sparse (vertical dashed orange line), and the theory also involves constants (e.g., the constant $C(\delta, \rho, b/a)$ in Theorem 2.2) that can have a significant impact on the actual number of measurements required for reconstruction.

In all cases, this experiment indicates that the proposed sketching method preserves information, and the decoding method allows for a good estimation in an optimistic scenario with a large number of samples and an optimally-chosen regularization parameter.

4.3. Comparison with graphical lasso type estimators - finite sample regime

In a last experiment we compare quantitatively our approach with the graphical lasso estimator (computed with the Scikit-learn implementation [60]) and we investigate the influence of the sample size on the final estimation. We consider $d = 256$ and n samples $\mathbf{x}_1, \dots, \mathbf{x}_n \sim \mathcal{N}(0, \Theta^{-1})$ with three draws of precision matrices for each Erdos and Powerlaw settings. We use a sketch of the data for a number of samples n varying from 50 (in order to have the small sample regime) to 500000. We compute the sketch of the data for sketch size $m \in \{1024, 2048, 4096, 8192, 16384, 32768 = d^2/2\}$. As in the previous experiment, for each method, each n and each score function we pick the regularization parameter $\lambda \in \{1e-4, 5e-4, 1e-3, \dots, 5e-1\}$ that leads to the best average score.

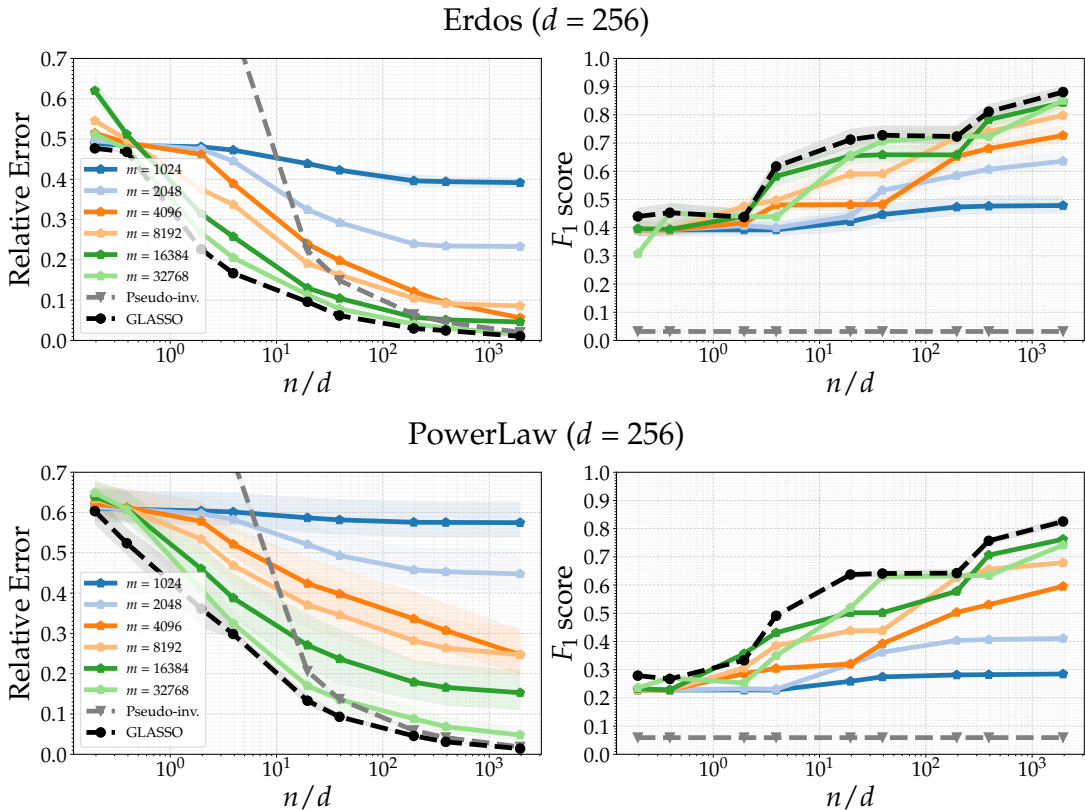


FIG. 4. Comparison between our estimator and graphical lasso type estimators with the Erdos (**top row**) and PowerLaw settings (**bottom row**). The 10-th percentile for these scores are reported in shaded line.

We also consider a simple baseline, namely the pseudo-inverse of the empirical covariance matrix $(\hat{\Sigma})^\dagger$ as an estimate for Θ_{true} . Note that when $n > d$ the matrix $\hat{\Sigma}$ is almost surely invertible to that this baseline corresponds to the maximum likelihood estimator of Θ_{true} without ℓ_1 regularization (same as graphical lasso with $\lambda = 0$ in this case). We report these results in Figure 4.

In both settings, the results show that our estimator improves as m and n increase, as expected. Moreover, for $m = 32768 = d^2/2$, the results of the graphical lasso and our estimator are nearly identical, demonstrating that our estimator is consistent with the graphical lasso when the number of measurements reaches the number of degrees of freedom of the empirical covariance matrix. Furthermore, we notice that for the Erdos setting, the performances are very similar to those of the graphical lasso even for $m \approx 8192$. This result indicates that even in the case of a finite number of samples, we are able to accurately estimate the precision matrices with a limited number of measurements (8192 corresponds to a compression rate of $\approx 75\%$). For the PowerLaw setting, the results in terms of relative error are more mixed. However, the F_1 score remains comparable to that of the graphical lasso, indicating that our approach captures the correct statistical dependencies, but may struggle to accurately estimate the intensity of these dependencies. In a rather reassuring way, in the small sample regime, our estimator

outperforms the MLE estimator $(\widehat{\Sigma})^\dagger$ in terms of relative error and consistently outperforms it in terms of the F_1 score (this is reasonable because $(\widehat{\Sigma})^\dagger$ unlike Θ_{true}).

5. Covering numbers bounds and concentration inequalities for Theorem 2.2

In this section, we provide the necessary ingredients to prove Theorem 2.2. Guided by Theorem 2.1, we start by giving general results to control covering numbers before applying them to the one considered in this article. Then, we provide the concentration inequality needed on the sketching operator \mathcal{A} .

5.1. Controlling covering numbers

We take the general point of view of normed vector spaces E, F with norms $\|\cdot\|_E, \|\cdot\|_F$ and a function $f : \Omega \subseteq E \rightarrow F$ defined on some domain $\Omega = \text{dom}(f)$ of E . For a subset $\mathfrak{X} \subseteq E$ the covering number of \mathfrak{X} w.r.t. $\|\cdot\|_E$ with radius $\varepsilon > 0$, denoted by $\mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \varepsilon)$, is the minimal number of closed balls of radius ε (w.r.t. $\|\cdot\|_E$) required to entirely cover \mathfrak{X} and whose centers are in \mathfrak{X} (see Appendix A.2.1 for a formal definition).

Remark 5 (Link with the precision matrix estimation problem) *This section is presented in a general case, but its results will ultimately be applied to control the covering number of the normalized secant of our model set $\mathfrak{S}_{k,a,b}$ defined in (2.3). Therefore, one should keep in mind that our final application framework will consider $E = F = S_d$, $f = \text{inv}$ with Ω the set of symmetric and invertible matrices and $\mathfrak{X} = \mathfrak{S}_{k,a,b}^{-1} = \{\Theta = \Sigma^{-1} : \Sigma \in \mathfrak{S}_{k,a,b}\}$.*

The core of our reasoning is based on the following sets, which generalize definition (2.8).

Definition 2 (Normalized secant sets) *We define the normalized secant set of some $\mathfrak{X} \subseteq E$ as*

$$S[\mathfrak{X}] \triangleq \left\{ \frac{x-y}{\|x-y\|_E} : (x,y) \in \mathfrak{X}^2, \|x-y\|_E > 0 \right\}. \quad (5.1)$$

We introduce also the normalized secant set of $\mathfrak{X} \subseteq \Omega = \text{dom}(f)$ embedded by f as

$$S[f(\mathfrak{X})] \triangleq \left\{ \frac{f(x)-f(y)}{\|f(x)-f(y)\|_F} : (x,y) \in \mathfrak{X}^2, \|f(x)-f(y)\|_F > 0 \right\}. \quad (5.2)$$

The main contribution of this section is the control of the covering number of a normalized secant $S[f(\mathfrak{X})] \subseteq F$ by the covering number of \mathfrak{X} and $S[\mathfrak{X}]$, for a sufficiently smooth f . The intuition is the following: if we consider a signal $x \in \mathfrak{X}$ in a “low-dimensional” space, then its image by f is also “low-dimensional” if f is sufficiently regular on \mathfrak{X} .

To carry out the analysis of $S[f(\mathfrak{X})]$ we introduce the notions of *long* and *short chords*. These objects are inspired by results in compressive independent component analysis and the theory of random projections on manifolds [20].

Long and short chords The set $S[f(\mathfrak{X})]$ can be divided in two subsets that are more analytically tractable. First, we introduce, for $\eta > 0$, the following sets

$$\begin{aligned} \mathbf{C}_\eta^+(\mathfrak{X}, f) &\triangleq \{(x, y) \in \mathfrak{X}^2 : \|f(x) - f(y)\|_F > \eta\}, \\ \mathbf{C}_\eta^-(\mathfrak{X}, f) &\triangleq \{(x, y) \in \mathfrak{X}^2 : 0 < \|f(x) - f(y)\|_F \leq \eta\}. \end{aligned} \quad (5.3)$$

With these notations, $S[f(\mathfrak{X})]$ can be decomposed into long and short chords $S[f(\mathfrak{X})] = S_\eta^+[f(\mathfrak{X})] \cup S_\eta^-[f(\mathfrak{X})]$, where the long and short chords are respectively defined by

$$\begin{aligned} S_\eta^+[f(\mathfrak{X})] &\triangleq \left\{ \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} : (x, y) \in \mathbf{C}_\eta^+(\mathfrak{X}, f) \right\}, \\ S_\eta^-[f(\mathfrak{X})] &\triangleq \left\{ \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} : (x, y) \in \mathbf{C}_\eta^-(\mathfrak{X}, f) \right\}. \end{aligned} \quad (5.4)$$

In order to control the covering number of $S[f(\mathfrak{X})]$ we will exploit its decomposition via the spaces $S_\eta^\pm[f(\mathfrak{X})]$ whose covering numbers are easier to obtain, assuming some regularity of the function f on \mathfrak{X} .

Definition 3 (Bi-Lipschitz assumption) *Given positive constants α and β , we say that a function f is (α, β) -bi-Lipschitz if for all $(x, y) \in \mathfrak{X}^2$ it verifies*

$$\alpha \|x - y\|_E \leq \|f(x) - f(y)\|_F \quad (5.5)$$

$$\|f(x) - f(y)\|_F \leq \beta \|x - y\|_E. \quad (5.6)$$

In order to define the notion of differential we assume that the model set is contained in the interior of $\Omega = \text{dom}(f)$, i.e., $\mathfrak{X} \subseteq \text{int}(\Omega)$. In particular when f is differentiable on \mathfrak{X} , (5.6) implies that $\sup_{x \in \mathfrak{X}} \|\mathbf{D}f_x\|_{\text{op}} \leq \beta < +\infty$, where $\|\cdot\|_{\text{op}}$ is the operator norm defined as $\|\mathbf{D}f_x\|_{\text{op}} \triangleq \sup_{\|h\|_E \leq 1} \|\mathbf{D}f_x(h)\|_F$.

We further consider the following regularity assumptions:

Assumption A-1 (Second order approximation) *f is differentiable on $\text{int}(\Omega)$ and there exists $L > 0$ such that*

$$\forall (x, y) \in \mathfrak{X}^2, \|f(x) - f(y) - \mathbf{D}f_y(x - y)\|_F \leq L \|x - y\|_E^2. \quad (5.7)$$

Assumption A-2 (Bounded curvature) *f is differentiable on $\text{int}(\Omega)$ and ζ -smooth on \mathfrak{X} , where $\zeta > 0$, i.e.*

$$\forall (x, y) \in \mathfrak{X}^2, \|\mathbf{D}f_x - \mathbf{D}f_y\|_{\text{op}} \leq \zeta \|x - y\|_E. \quad (5.8)$$

As described in Lemma 4 (see Appendix A.2.2) the condition **A-2** implies **A-1** when \mathfrak{X} is a *convex* model set.

Covering numbers of $S_\eta^+[f(\mathfrak{X})]$ controlling the covering with long chords is relatively easy under some assumptions on f as proven in the following proposition (proof in Appendix A.2.3).

Proposition 1 *Assume that f is β -Lipschitz as in (5.6). Then for any $\eta > 0$ and $\varepsilon > 0$,*

$$\mathcal{N}(S_\eta^+[f(\mathfrak{X})], \|\cdot\|_F, \varepsilon) \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\eta}{16\beta}\varepsilon)^2. \quad (5.9)$$

In other words, the ‘‘dimension’’ of the long chords can be controlled by that of the model set itself, assuming only that f is Lipschitz-continuous.

Covering numbers of $S_\eta^-[f(\mathfrak{X})]$ controlling the covering number of short chords is a little more delicate and generally requires a fine analysis of \mathfrak{X} [34]. However, by adding the hypothesis **A-1** and **A-2** we are able to prove the following result (proof in Appendix A.2.4).

Proposition 2 *Assume that f is (α, β) -bi-Lipschitz and satisfies assumptions **A-1** and **A-2**. Then for any $\varepsilon > 0$ and $\eta > 0$,*

$$\mathcal{N}(S_\eta^-[f(\mathfrak{X})], \|\cdot\|_F, 2(\varepsilon + \frac{L}{\alpha^2}\eta)) \leq \frac{\zeta + \beta\alpha}{\alpha^2\varepsilon} \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\alpha\varepsilon}{\zeta + \beta\alpha}) \times \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \frac{\alpha^2\varepsilon}{2(\zeta + \beta\alpha)}). \quad (5.10)$$

Intuition of the proof The idea is to show that each element of $S_\eta^-[f(\mathfrak{X})]$ can be described by a ‘‘tangent’’ vector to $S[\mathfrak{X}]$ and that the set of tangent vectors has a covering number which can be controlled by those of \mathfrak{X} and $S[\mathfrak{X}]$. \square

By combining the two propositions we are now ready to state the main theorem of this section: Assume that f is (α, β) -bi-Lipschitz and satisfies assumptions **A-1** and **A-2** Then for all $\eta > 0, \varepsilon > 0$, we have

$$\begin{aligned} \mathcal{N}(S[f(\mathfrak{X})], \|\cdot\|_F, 2\left[\varepsilon + \frac{L}{\alpha^2}\eta\right]) &\leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\eta}{8\beta}\left[\varepsilon + \frac{L}{\alpha^2}\eta\right])^2 \\ &\quad + \frac{\zeta + \beta\alpha}{\alpha^2\varepsilon} \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\alpha\varepsilon}{\zeta + \beta\alpha}) \times \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \frac{\alpha^2\varepsilon}{2(\zeta + \beta\alpha)}). \end{aligned} \quad (5.11)$$

Proof We use that for any $\eta > 0$ $S[\mathfrak{X}] = S_\eta^+[f(\mathfrak{X})] \cup S_\eta^-[f(\mathfrak{X})]$ thus the covering number of $S[f(\mathfrak{X})]$ is less than the sum of the coverings of $S_\eta^+[f(\mathfrak{X})]$ and $S_\eta^-[f(\mathfrak{X})]$. Then we apply Proposition 1 and Proposition 2. \square

Putting everything together in the case of sparse precision matrix we apply the previous results to our framework, that is $E = F = S_d$, $f = \text{inv}$ and $\mathfrak{X} = \mathfrak{S}_{k,a,b}^{-1} = \{\Theta = \Sigma^{-1} : \Sigma \in \mathfrak{S}_{k,a,b}\}$. We set the norms as follow : $\|\cdot\|_E = \|\cdot\|_{\text{Fro}}$ the Frobenius norm and $\|\cdot\|_F = \|\cdot\|_\Lambda$ as defined in (2.13). We have the following lemma which shows that $f = \text{inv}$ satisfies all the necessary regularity assumptions (the proof can be found in Appendix A.2.5).

Lemma 1 *Assume that there exist constants c_{Fro} and C_{Fro} such that $c_{\text{Fro}}\|\mathbf{M}\|_{\text{Fro}} \leq \|\mathbf{M}\|_{\Lambda} \leq C_{\text{Fro}}\|\mathbf{M}\|_{\text{Fro}}$. Then the function $f = \text{inv}$ is (α, β) -bi-Lipschitz and satisfies assumptions **A-1** and **A-2** on $\mathfrak{S}_{k,a,b}^{-1}$ with $\alpha = \frac{c_{\text{Fro}}}{b^2}$, $\beta = \frac{C_{\text{Fro}}}{a^2}$, $L = \frac{C_{\text{Fro}}}{a^3}$, $\zeta = 2L$.*

Note that expressions for the constants c_{Fro} and C_{Fro} will be provided in Proposition 4. In order to control the covering number of $\mathcal{S}[\mathfrak{S}_{k,a,b}] = \mathcal{S}[\text{inv}(\mathfrak{S}_{k,a,b}^{-1})]$ we only have to check those of $\mathfrak{S}_{k,a,b}^{-1}$ and $\mathcal{S}[\mathfrak{S}_{k,a,b}^{-1}]$. It is done in the following lemma (the proof can be found in Appendix A.2.6).

Lemma 2 *For any $\varepsilon > 0$ we have*

$$\mathcal{N}(\mathfrak{S}_{k,a,b}^{-1}, \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{18\sqrt{d} \times b}{\varepsilon}\right)^{d+k} \text{ and } \mathcal{N}(\mathcal{S}[\mathfrak{S}_{k,a,b}^{-1}], \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \left(\frac{ed^2}{4k}\right)^{2k} \left(\frac{18}{\varepsilon}\right)^{d+2k}. \quad (5.12)$$

This result show that the *box counting dimensions* [65] (also called entropy dimensions) of $\mathfrak{S}_{k,a,b}^{-1}$ and $\mathcal{S}[\mathfrak{S}_{k,a,b}^{-1}]$ are smaller than $(d+k)\log(d)$ and $(d+2k)\log(d)$ and thus much smaller than d^2 , which will allow us to have the guarantees presented in the introduction.

Corollary 1 *Assuming the existence of the constant $C_{\text{Fro}}, c_{\text{Fro}}$ as in Lemma 1, there exist absolute constants $c_0, c_1 \geq 1$, such that, for any $\varepsilon > 0$, the covering number of $\mathcal{S}[\mathfrak{S}_{k,a,b}]$ verifies*

$$\mathcal{N}(\mathcal{S}[\mathfrak{S}_{k,a,b}], \|\cdot\|_{\Lambda}, \varepsilon) \leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(c_0 \frac{\sqrt{d}C_{\text{Fro}}^2 b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5}\right)^{2(d+k)} + \left(c_1 \frac{\sqrt{d}C_{\text{Fro}}^2}{\varepsilon^2 c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1\right) \frac{b^5}{a^5}\right)^{d+2k+1} \right]. \quad (5.13)$$

This is a direct consequence of Theorem 5.1, combined with Lemma 1 and 2. See Appendix A.2.7 for the proof. This corollary allows for a control of the covering number that is required to prove a RIP for rank-one projections.

5.2. Application to rank-one projections

The results of the previous section allowed us to control one of the three quantities of interest for establishing the RIP_{δ} : the covering number. We are left with studying the concentration properties of the operator \mathcal{A} (Theorem 2.1).

A natural choice for the norm $\|\cdot\|_{\mathbb{R}^m}$ would be the standard Euclidean norm $\|\cdot\|_2$. However, in order for the RIP to hold, one would have to choose $\|\mathbf{M}\|_{S_d} = \left(\mathbb{E}_{\mathbf{A} \sim \Lambda} \left[|\langle \mathbf{M}_j, \mathbf{M} \rangle|^2 \right]\right)^{1/2}$. Unfortunately, with these choices, we were unable to find sufficiently tight concentration functions C_1, C_2 that lead to better guaranties than $m \gtrsim d^2$. This phenomenom had already been observed, as written in [12] the rank-one projections lead to loose RIP constants for low-rank matrix completion problems (unless $m \gtrsim d^2$), and we envision that similar results hold in our case. The remedy found in [12] is to consider instead an ℓ_1 RIP (*i.e.* taking the ℓ_1 norm for $\|\cdot\|_{\mathbb{R}^m}$ in Definition 1 instead of the ℓ_2 norm), leading to the choice of $\|\cdot\|_{\Lambda}$ defined in (2.13) for $\|\cdot\|_{S_d}$. This next part shows that this remedy helps for the recovery of covariance matrices and thus of precision matrices.

Proposition 3 Consider the operator $\mathcal{A}(\Sigma) = \frac{1}{m} (\mathbf{a}_1^\top \Sigma \mathbf{a}_1, \dots, \mathbf{a}_m^\top \Sigma \mathbf{a}_m)^\top$, where the vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ are i.i.d. and either follow $\Lambda = \Lambda_G = \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ or $\Lambda = \Lambda_G = \mathcal{U}(\mathbb{S}^{d-1})$ with the associated norm $\|\Sigma\|_\Lambda = \mathbb{E}_{\mathbf{a} \sim \Lambda} [\|\mathbf{a}^\top \Sigma \mathbf{a}\|]$. Then, for any $\mathbf{U} \in S_d(\mathbb{R})$ satisfying $\|\mathbf{U}\|_\Lambda = 1$ and for any $t > 0$, we have

$$\mathbb{P}(|\|\mathcal{A}(\mathbf{U})\|_1 - 1| > t) \leq 2 \exp\left(-\frac{m}{8e^2} \min\left(\frac{t^2}{K_\Lambda^2}, \frac{t}{K_\Lambda}\right)\right), \quad (5.14)$$

where K_Λ is an absolute constant given by $K_{\Lambda_G} = \frac{76e^2\sqrt{15}}{\log 2}$ and $K_{\Lambda_U} = \frac{304e^3\sqrt{15}}{\log 2}$.

Intuition of the proof The proof is based on a Bernstein-type concentration inequality for sums of sub-exponential variables. The essential ingredient of the proof is thus to show that the centered random variable $|\mathbf{a}^\top \mathbf{U} \mathbf{a}| - 1$ involved in $\|\mathcal{A}(\mathbf{U})\|_1 - 1$ is subexponential with a subexponential norm bounded by an absolute constant K_Λ . The full proof can be found in Appendix A.3.1. \square

This result provides the function C_2 required to obtain a RIP with Theorem 2.1.

Remark 6 This is essentially the only result that would need to be adapted if one wants to provide information-theoretic guarantees to the sketching operator defined from random structured matrices. Indeed, we emphasize that all previous results on covering numbers are still valid in the structured case.

Finally to be able to control the covering number, the norm $\|\cdot\|_\Lambda$ needs to verify the hypothesis of Corollary 1. This is done in the following proposition (see Appendix A.3.2 for the proof).

Proposition 4 Let $\Lambda \in \{\Lambda_G, \Lambda_U\}$ be either the Gaussian or uniform distribution on \mathbb{R}^d and consider the associated Λ -norm defined in (2.13). Then,

$$\forall \mathbf{M} \in S_d(\mathbb{R}), \quad \frac{2}{9\sqrt{15}d} \|\mathbf{M}\|_{\text{Fro}} \leq \|\mathbf{M}\|_\Lambda \leq \frac{1}{\sqrt{d}} \|\mathbf{M}\|_{\text{Fro}}. \quad (5.15)$$

This gives $c_{\text{Fro}} = 2/(9\sqrt{15}d)$ and $C_{\text{Fro}} = 1/\sqrt{d}$ in Lemma 1.

In this section, we have established control over the covering numbers via Corollary 1 and introduced the concentration inequality for the sketching operator through Proposition 3. These components provide the necessary foundation for proving Theorem 2.2. The comprehensive proof of this theorem can be found in Appendix A.3.3.

6. Conclusion & perspectives

In this work, we have presented a compressive approach based on sketching to estimate sparse precision matrices. We have shown that it is possible to estimate a $(d + 2k)$ -sparse precision matrix from a data sketch of the order $(d + 2k) \log(d)$, which is significantly smaller than the typical memory complexity of d^2 associated with the graphical lasso. Our analysis is supported by information-theoretic guarantees, where we have established restricted isometries and instance optimality properties. Finally, we have proposed a practical algorithmic solution for computing an estimation of the precision matrix from the sketch.

Our work opens several new lines of research. Given the generality of the tools presented in this paper, it would be interesting to explore whether similar guarantees can be established for other model sets based on specific graph structures [51]. Also, as practitioners are not always interested in the graph associated with the precision matrix *per se* but rather in some of its properties (*e.g.*, a group structure among the nodes), it would be interesting to see how these properties can directly be inferred using a compressive learning approach and whether it can help further reduce the sketch’s dimension.

From an algorithmic point of view, our work also raises several questions. The proposed estimator is costly as it requires to solve several graphical lasso. An interesting further work would be to design a more efficient decoder that is sufficiently close to the optimal decoder given by the theory. In this context, the choice proposed in this paper is a first step toward practical recovery, but other algorithms could be used based on different precision matrix estimators like [81]. Finally, from an application point of view, an interesting perspective would be to use the sketching approach to learn, in an online way, a dynamic graph, in the way of the time varying graphical lasso [38].

A. Proofs

A.1. Proofs of Section 2.1

A.1.1. Proof of Theorem 2.1

Proof With the notations of the theorem $\Sigma^* = \Delta[\mathcal{A}(\widehat{\Sigma})] \in \arg \min_{\Sigma \in \mathfrak{S}} \|\mathcal{A}(\Sigma) - \mathcal{A}(\widehat{\Sigma})\|_{\mathbb{R}^m}$. Then for any $\Sigma_{\mathfrak{S}} \in \mathfrak{S}$:

$$\begin{aligned}
 \|\Sigma^* - \Sigma\|_{S_d} &\leq \|\Sigma - \Sigma_{\mathfrak{S}}\|_{S_d} + \|\Sigma_{\mathfrak{S}} - \Sigma^*\|_{S_d} \\
 &\leq \|\Sigma - \Sigma_{\mathfrak{S}}\|_{S_d} + \frac{1}{1-\delta} \|\mathcal{A}(\Sigma^*) - \mathcal{A}(\Sigma_{\mathfrak{S}})\|_{\mathbb{R}^m} \\
 &\leq \|\Sigma - \Sigma_{\mathfrak{S}}\|_{S_d} + \frac{1}{1-\delta} \left(\|\mathcal{A}(\Sigma^*) - \mathcal{A}(\widehat{\Sigma})\|_{\mathbb{R}^m} + \|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma_{\mathfrak{S}})\|_{\mathbb{R}^d} \right) \\
 &\leq \|\Sigma - \Sigma_{\mathfrak{S}}\|_{S_d} + \frac{2}{1-\delta} \|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma_{\mathfrak{S}})\|_{\mathbb{R}^m} \\
 &\leq \|\Sigma - \Sigma_{\mathfrak{S}}\|_{S_d} + \frac{2}{1-\delta} \|\mathcal{A}(\widehat{\Sigma}) - \mathcal{A}(\Sigma)\|_{\mathbb{R}^d} + \frac{2}{1-\delta} \|\mathcal{A}(\Sigma) - \mathcal{A}(\Sigma_{\mathfrak{S}})\|_{\mathbb{R}^m}.
 \end{aligned} \tag{A.1}$$

We introduce the following “distance” to the model set \mathfrak{S} :

$$d^\circ(\Sigma, \mathfrak{S}) = \inf_{\mathbf{M} \in \mathfrak{S}} \|\Sigma - \mathbf{M}\|_{S_d} + \frac{2}{1-\delta} \|\mathcal{A}(\Sigma) - \mathcal{A}(\mathbf{M})\|_{\mathbb{R}^m}. \tag{A.2}$$

Then we have $d^\circ(\Sigma, \mathfrak{S}) = 0$ if $\Sigma \in \mathfrak{S}$ and

$$\|\Sigma^* - \Sigma\|_{S_d} \leq d^\circ(\Sigma, \mathfrak{S}) + \frac{2}{1-\delta} \|\mathcal{A}(\Sigma) - \mathcal{A}(\mathbf{M})\|_{\mathbb{R}^m}. \tag{A.3}$$

□

A.1.2. Proof of Theorem 2.1

Proof Let us start by considering \bar{S}_ε , an ε -net of $S[\mathfrak{S}]$ with respect to $\|\cdot\|_{S_d}$, for some $\varepsilon > 0$. Then, for every $\mathbf{U} \in S[\mathfrak{S}]$, there exists $\bar{\mathbf{U}} \in \bar{S}_\varepsilon$ such that $\|\mathbf{U} - \bar{\mathbf{U}}\|_{S_d} \leq \varepsilon$. From the triangular inequality we have

$$\left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1 \right| \leq \left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} \right| + \left| \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1 \right|.$$

Focusing on the first term of the right-hand side, we obtain

$$\left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} \right| \leq \|\mathcal{A}(\mathbf{U} - \bar{\mathbf{U}})\|_{\mathbb{R}^m} \leq \|\mathcal{A}\| \cdot \|\mathbf{U} - \bar{\mathbf{U}}\|_{S_d} \leq \varepsilon \|\mathcal{A}\|. \quad (\text{A.4})$$

Hence for $\mathbf{U} \in S[\text{inv}(\mathfrak{S})]$:

$$\left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1 \right| \leq \max_{\bar{\mathbf{U}} \in \bar{S}_\varepsilon} \left| \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1 \right| + \varepsilon \|\mathcal{A}\|. \quad (\text{A.5})$$

So we have for any $0 < \delta < 1$:

$$\mathbb{P} \left(\sup_{\mathbf{U} \in S[\mathfrak{S}]} \left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1 \right| \leq \delta \right) \geq 1 - \mathbb{P}(\varepsilon \|\mathcal{A}\| > \frac{\delta}{2}) - \mathbb{P}(\max_{\bar{\mathbf{U}} \in \bar{S}_\varepsilon} \left| \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1 \right| > \frac{\delta}{2}). \quad (\text{A.6})$$

We will control these two terms. For the first one we have the concentration with C_1 . For the second one, using the union bound yields

$$\mathbb{P}(\max_{\bar{\mathbf{U}} \in \bar{S}_\varepsilon} \left| \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1 \right| > \frac{\delta}{2}) \leq \sum_{\bar{\mathbf{U}} \in \bar{S}_\varepsilon} \mathbb{P}(\left| \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1 \right| > \frac{\delta}{2}) \quad (\text{A.7})$$

Using the concentration given by C_2 , for $\bar{\mathbf{U}} \in \bar{S}_\varepsilon$ and $t \in]0, 1[$, we have $\mathbb{P}(\left| \|\mathcal{A}\bar{\mathbf{U}}\|_{\mathbb{R}^m} - 1 \right| > t) \leq C_2(t)$. Applying this with $t = \delta/2$ and using (A.7) gives

$$\mathbb{P}(\max_{\bar{\mathbf{U}} \in \bar{S}_\varepsilon} \left| \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1 \right| > \frac{\delta}{2}) \leq |\bar{S}_\varepsilon| C_2(\delta/2). \quad (\text{A.8})$$

As a result, we obtain for $\delta \in]0, 1[$ and $\varepsilon > 0$

$$\mathbb{P} \left(\sup_{\mathbf{U} \in S[\mathfrak{S}]} \left| \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} - 1 \right| \leq \delta \right) \geq 1 - C_1\left(\frac{\delta}{2\varepsilon}\right) - |\bar{S}_\varepsilon| C_2(\delta/2). \quad (\text{A.9})$$

□

A.1.3. Proof of Remark 3

Proof For some $\varepsilon' > 0$, consider $T_{\varepsilon'}$ an ε' -net of the sphere B_{S_d} . Then, for all $\mathbf{U} \in B_{S_d}$, there exists $\bar{\mathbf{U}} \in T_{\varepsilon'}$ such that

$$\|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} \leq \varepsilon' \|\mathcal{A}\| + \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m}.$$

Thus, taking the supremum over \mathbf{U} yields

$$\sup_{\mathbf{U} \in B_{S_d}} \|\mathcal{A}(\mathbf{U})\|_{\mathbb{R}^m} \leq \varepsilon' \|\mathcal{A}\| + \max_{\bar{\mathbf{U}} \in T_{\varepsilon'}} \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m},$$

therefore $\|\mathcal{A}\| \leq (1 - \varepsilon')^{-1} \max_{\bar{\mathbf{U}} \in T_{\varepsilon'}} \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m}$. As a consequence, for $t > 1/(1 - \varepsilon')$,

$$\begin{aligned} \mathbb{P}(\|\mathcal{A}\| > t) &\leq \mathbb{P}\left(\max_{\bar{\mathbf{U}} \in T_{\varepsilon'}} \|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} > (1 - \varepsilon')t\right) \leq \sum_{\bar{\mathbf{U}} \in T_{\varepsilon'}} \mathbb{P}(|\|\mathcal{A}(\bar{\mathbf{U}})\|_{\mathbb{R}^m} - 1| > (1 - \varepsilon')t - 1) \\ &\leq \mathcal{N}(B_{S_d}, \|\cdot\|_{S_d}, \varepsilon') C_2((1 - \varepsilon')t - 1). \end{aligned}$$

□

A.2. Proofs of Section 5.1

A.2.1. General results on covering numbers

Let (E, d) be a semi-metric space. The covering number of $\mathfrak{S} \subseteq E$ with radius ε with respect to d is defined as:

$$\mathcal{N}(\mathfrak{S}, d, \varepsilon) \triangleq \min \left\{ N \in \mathbb{N} : \exists x_1, \dots, x_N \in \mathfrak{S}, \mathfrak{S} \subseteq \bigcup_{i=1}^N B_d(x_i, \varepsilon) \right\}. \quad (\text{A.10})$$

If $N = \mathcal{N}(\mathfrak{S}, d, \varepsilon)$, then for any $x \in \mathfrak{S}$ there exists $i \in \llbracket N \rrbracket$ such that $d(x, x_i) \leq \varepsilon$. We recall the following Lemma regarding covering numbers that can be found in [34, Lemma A.3.].

Lemma 3 *Let Y, Z be two subset of a pseudo metric space (X, d) such that the following holds:*

$$\forall z \in Z, \exists y \in Y, d(z, y) \leq \delta, \quad (\text{A.11})$$

where $\delta \geq 0$. Then for all $\varepsilon > 0$

$$\mathcal{N}(Z, d, 2(\delta + \varepsilon)) \leq \mathcal{N}(Y, d, \varepsilon). \quad (\text{A.12})$$

A.2.2. Descent Lemma

Lemma 4 (Descent Lemma) *Let E, F be two normed vector spaces and Ω a subset of E and $\mathfrak{X} \subseteq \text{int}(\Omega)$. Consider $f : \Omega \rightarrow F$ a L -smooth function on \mathfrak{X} i.e.*

$$\forall (x, y) \in \mathfrak{X}^2, \|\text{D}f_x - \text{D}f_y\|_{\text{op}} \leq L\|x - y\|_E. \quad (\text{A.13})$$

Let $(x, y) \in \mathfrak{X}^2$ such that the segment $\llbracket x, y \rrbracket$ lies in \mathfrak{X} . Then,

$$\|f(x) - f(y) - \text{D}f_y(x - y)\|_F \leq L\|x - y\|_E^2. \quad (\text{A.14})$$

In particular if \mathfrak{X} is convex then **A-2** implies **A-1**.

Proof From [23, Corollary 3.3] f verifies

$$\|f(x) - f(y) - \text{D}f_y(x - y)\|_F \leq \sup_{z \in \llbracket x, y \rrbracket} \|\text{D}f_z - \text{D}f_x\|_{\text{op}} \|x - y\|_E.$$

Now take $z \in \llbracket x, y \rrbracket$ and write it as $z = (1 - t)x + ty \in \mathfrak{X}$ for some $t \in [0, 1]$. Since f is L -smooth on \mathfrak{X} we have:

$$\|\text{D}f_z - \text{D}f_x\|_{\text{op}} = \|\text{D}f_{(1-t)x+ty} - \text{D}f_x\|_{\text{op}} \leq L\|(1-t)x + ty - x\|_E = Lt\|x - y\|_E \leq L\|x - y\|_E$$

Hence $\sup_{z \in \llbracket x, y \rrbracket} \|\text{D}f_z - \text{D}f_x\|_{\text{op}} \leq L\|x - y\|_E$ and thus $\|f(x) - f(y) - \text{D}f_y(x - y)\|_F \leq L\|x - y\|_E^2$ □

A.2.3. Proof of Proposition 1

Proof In the following we note $\|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2} = \|x_1 - x_2\|_E + \|y_1 - y_2\|_E$. We introduce, for $\eta \geq 0$, the set

$$\mathfrak{X}_\eta^2 = \{(x, y) \in \mathfrak{X}^2 : \|f(x) - f(y)\|_F > \eta\}. \quad (\text{A.15})$$

First note that $\mathfrak{X}_\eta^2 \subset \mathfrak{X}^2$, and consequently

$$\mathcal{N}(\mathfrak{X}_\eta^2, \|\cdot\|_{\otimes 2}, \varepsilon) \leq \mathcal{N}(\mathfrak{X}^2, \|\cdot\|_{\otimes 2}, \frac{\varepsilon}{2}) \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\varepsilon}{4})^2. \quad (\text{A.16})$$

We consider $\eta > 0$ and define $g : \mathfrak{X}_\eta^2 \rightarrow S_\eta^+(f(\mathfrak{X}))$ by $g(x, y) = \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F}$ for $(x, y) \in \mathfrak{X}_\eta^2$. By definition, g is surjective. We will show that it is also Lipschitz. With $(x_1, y_1), (x_2, y_2) \in \mathfrak{X}_\eta^2 \times \mathfrak{X}_\eta^2$ we obtain

$$\begin{aligned} & \left\| \frac{f(x_1) - f(y_1)}{\|f(x_1) - f(y_1)\|_F} - \frac{f(x_2) - f(y_2)}{\|f(x_2) - f(y_2)\|_F} \right\|_F \\ & \leq \left\| \frac{f(x_1) - f(y_1)}{\|f(x_1) - f(y_1)\|_F} - \frac{f(x_1) - f(y_1)}{\|f(x_2) - f(y_2)\|_F} \right\|_F + \left\| \frac{f(x_1) - f(y_1)}{\|f(x_2) - f(y_2)\|_F} - \frac{f(x_2) - f(y_2)}{\|f(x_2) - f(y_2)\|_F} \right\|_F \\ & \leq \frac{1}{\|f(x_2) - f(y_2)\|_F} (\|f(x_1) - f(x_2)\|_F + \|f(y_1) - f(y_2)\|_F) \\ & \quad + \left| \frac{1}{\|f(x_1) - f(y_1)\|_F} - \frac{1}{\|f(x_2) - f(y_2)\|_F} \right| \|f(x_1) - f(y_1)\|_F \\ & \leq \frac{2\beta}{\|f(x_2) - f(y_2)\|_F} \|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2} \\ & \quad + \left| \|f(x_2) - f(y_2)\|_F - \|f(x_1) - f(y_1)\|_F \right| \times \frac{\|f(x_1) - f(y_1)\|_F}{\|f(x_1) - f(y_1)\|_F \|f(x_2) - f(y_2)\|_F} \\ & \leq \frac{2\beta}{\eta} \|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2} + \|f(x_2) - f(x_1) - f(y_2) + f(y_1)\|_F \times \frac{1}{\|f(x_2) - f(y_2)\|_F} \\ & \leq \frac{2\beta}{\eta} \|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2} + 2\beta \|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2} \times \frac{1}{\|f(x_2) - f(y_2)\|_F} \\ & \leq \frac{4\beta}{\eta} \|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2}. \end{aligned}$$

Consequently,

$$\|g(x_1, y_1) - g(x_2, y_2)\|_F \leq \frac{4\beta}{\eta} \|(x_1, y_1) - (x_2, y_2)\|_{\otimes 2}. \quad (\text{A.17})$$

So g is a surjective $\frac{4\beta}{\eta}$ -Lipschitz function from $(\mathfrak{X}_\eta^2, \|\cdot\|_{\otimes 2})$ to $(S_\eta^+(f(\mathfrak{X})), \|\cdot\|_F)$. Hence using [34, Lemma A.2.] we obtain

$$\mathcal{N}(S_\eta^+(f(\mathfrak{X})), \|\cdot\|_F, \varepsilon) \leq \mathcal{N}(\mathfrak{X}_\eta^2, \|\cdot\|_{\otimes 2}, \frac{\eta}{4\beta} \varepsilon) \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\eta}{16\beta} \varepsilon)^2. \quad (\text{A.18})$$

□

A.2.4. Proof of Proposition 2

In order to prove the result, the reasoning will be the following: 1) we will show that any element of $S_{\eta}^{-}[f(\mathfrak{X})]$ is close to an element of a certain ‘‘tangent space’’ $D_{\frac{x-y}{\|f(x)-f(y)\|_F}}$ 2) we will control the covering number of this space. We begin with the following result.

Lemma 5 *Assume that f is (α, β) -bi-Lipschitz and satisfies assumptions **A-1**. For $\varepsilon_0 > 0$, consider the following subset of \mathfrak{X}^2 :*

$$I_{\varepsilon_0} \triangleq \{(x, y) \in \mathfrak{X}^2 : 0 < \|x - y\|_E \leq \varepsilon_0\}. \quad (\text{A.19})$$

Then, for any $\varepsilon_0 > 0$ and $(x, y) \in I_{\varepsilon_0}$,

$$\left\| \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} - Df_y \frac{x - y}{\|f(x) - f(y)\|_F} \right\|_F \leq \frac{L}{\alpha} \varepsilon_0.$$

In particular, for any $\varepsilon_0 > 0$,

$$\forall (x, y) \in I_{\varepsilon_0}, \exists h \in \mathfrak{X}, \left\| \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} - Df_h \frac{x - y}{\|f(x) - f(y)\|_F} \right\|_F \leq \frac{L}{\alpha} \varepsilon_0.$$

Proof First, assumption **A-1** gives

$$\|f(x) - f(y) - Df_y(x - y)\|_F \stackrel{\text{A-1}}{\leq} L\|x - y\|_E^2 \leq L\varepsilon_0\|x - y\|_E. \quad (\text{A.20})$$

Now since $\|x - y\|_E > 0$, we have $\|f(x) - f(y)\|_F > 0$ using (5.5) (from the inverse Lipschitz property). Thus by dividing by $\|f(x) - f(y)\|_F$ in (A.20) we obtain

$$\left\| \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} - Df_y \frac{x - y}{\|f(x) - f(y)\|_F} \right\|_F \leq L\varepsilon_0 \frac{\|x - y\|_E}{\|f(x) - f(y)\|_F} \stackrel{(5.5)}{\leq} \frac{L}{\alpha} \varepsilon_0.$$

□

The above result induces the following corollary.

Corollary 2 *Assume that f is (α, β) -bi-Lipschitz and satisfies assumptions **A-1**. For $\eta > 0$ consider $S_{\eta}^{-}[f(\mathfrak{X})]$ as defined in (5.4) and the normalized secant set $S[\mathfrak{X}]$ (see Definition 2). Let*

$$C \triangleq \left\{ \lambda v : \lambda \in]0, \frac{1}{\alpha}], v \in S[\mathfrak{X}] \right\}, \quad (\text{A.21})$$

and

$$T_C \triangleq \{Df_h(c) : (h, c) \in \mathfrak{X} \times C\}. \quad (\text{A.22})$$

Then

$$\forall u \in S_{\eta}^{-}[f(\mathfrak{X})], \exists t \in T_C, \|u - t\|_F \leq \frac{L}{\alpha^2} \eta. \quad (\text{A.23})$$

Proof Take $u = \frac{f(x)-f(y)}{\|f(x)-f(y)\|_F} \in S_\eta^-[f(\mathfrak{X})]$ (thus with $0 < \|f(x) - f(y)\|_F \leq \eta$). By using that f is (α, β) -bi-Lipschitz, we have $0 < \|x - y\|_E \leq \eta/\alpha$. So we can apply the Lemma 5 with $\varepsilon_0 = \eta/\alpha$ to prove that there exists $h \in \mathfrak{X}$ such that

$$\left\| \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} - \mathbf{D}f_h \frac{x - y}{\|f(x) - f(y)\|_F} \right\|_F \leq \frac{L}{\alpha^2} \eta.$$

Rewrite $\frac{x-y}{\|f(x)-f(y)\|_F} = \frac{x-y}{\|x-y\|_F} \frac{\|x-y\|_E}{\|f(x)-f(y)\|_F}$ and define $\lambda = \frac{\|x-y\|_E}{\|f(x)-f(y)\|_F}$. Therefore, we have $\lambda > 0$ and $|\lambda| \leq \frac{1}{\alpha}$. It proves that there exists $\lambda \in]0, 1/\alpha]$ and $h \in \mathfrak{X}$, such that

$$\left\| \frac{f(x) - f(y)}{\|f(x) - f(y)\|_F} - \mathbf{D}f_h \lambda \frac{x - y}{\|x - y\|_E} \right\|_F \leq \frac{L}{\alpha^2} \eta.$$

Considering $v = \frac{x-y}{\|x-y\|_E} \in S[\mathfrak{X}]$ concludes the proof.

□

The last thing to do is to control the covering number of T_C in the previous result. This will be done using the following lemma.

Lemma 6 *Let $C \subseteq E$ be any set such that $\forall c \in C, \|c\|_E \leq \delta$ for some $\delta > 0$. Assume that f is (α, β) -bi-Lipschitz and satisfies A-2. Consider $T_C \triangleq \{\mathbf{D}f_h(c) : (h, c) \in \mathfrak{X} \times C\}$. Then for any $\varepsilon > 0$,*

$$\mathcal{N}(T_C, \|\cdot\|_F, \varepsilon) \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\varepsilon}{\zeta\delta + \beta}) \times \mathcal{N}(C, \|\cdot\|_E, \frac{\varepsilon}{\zeta\delta + \beta}).$$

Proof Take $\bar{\mathfrak{X}}$ a ε -net of \mathfrak{X} and \bar{C} a ε -net of C . Then take $t = \mathbf{D}f_h(u) \in T_C$ and consider $\bar{u}, \bar{h} \in \bar{\mathfrak{X}} \times \bar{C}$ such that $\|\bar{c} - c\|_E \leq \varepsilon$ and $\|\bar{h} - h\|_E \leq \varepsilon$. Then with $\bar{t} = \mathbf{D}f_{\bar{h}}(\bar{c}) \in T_C$

$$\begin{aligned} \|t - \bar{t}\|_F &= \|\mathbf{D}f_h(c) - \mathbf{D}f_{\bar{h}}(\bar{c})\|_F \leq \|\mathbf{D}f_h(c) - \mathbf{D}f_{\bar{h}}(c)\|_F + \|\mathbf{D}f_{\bar{h}}(c) - \mathbf{D}f_{\bar{h}}(\bar{c})\|_F \\ &\leq \|\mathbf{D}f_h - \mathbf{D}f_{\bar{h}}\|_{\text{op}} \|c\|_E + \|\mathbf{D}f_{\bar{h}}\|_{\text{op}} \|c - \bar{c}\|_E \\ &\stackrel{\text{A-2}}{\leq} \zeta \|h - \bar{h}\|_E \delta + \sup_{h \in \mathfrak{X}} \|\mathbf{D}f_h\|_{\text{op}} \varepsilon \\ &\leq \varepsilon (\zeta \delta + \sup_{h \in \mathfrak{X}} \|\mathbf{D}f_h\|_{\text{op}}) \stackrel{(5.6)}{\leq} \varepsilon (\zeta \delta + \beta). \end{aligned} \tag{A.24}$$

This gives $\mathcal{N}(T_C, \|\cdot\|_F, \varepsilon(\zeta\delta + \beta)) \leq |\bar{\mathfrak{X}}| \times |\bar{C}| \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \varepsilon) \times \mathcal{N}(C, \|\cdot\|_E, \varepsilon)$.

□

We can now prove Proposition 2 which we recall first.

Proposition 2 *Assume that f is (α, β) -bi-Lipschitz and satisfies assumptions A-1 and A-2. Then for any $\varepsilon > 0$ and $\eta > 0$,*

$$\mathcal{N}(S_\eta^-[f(\mathfrak{X})], \|\cdot\|_F, 2(\varepsilon + \frac{L}{\alpha^2} \eta)) \leq \frac{\zeta + \beta\alpha}{\alpha^2 \varepsilon} \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\alpha\varepsilon}{\zeta + \beta\alpha}) \times \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \frac{\alpha^2 \varepsilon}{2(\zeta + \beta\alpha)}). \tag{5.10}$$

Proof Consider C and T_C as defined in Corollary 2. We have for any $c \in C$, $\|c\|_E \leq \frac{1}{\alpha}$ since $\forall u \in S[\mathfrak{X}]$, $\|u\|_E = 1$. So by applying Lemma 6 with $\delta = \frac{1}{\alpha}$ we obtain

$$\mathcal{N}(T_C, \|\cdot\|_F, \varepsilon) \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\varepsilon}{\frac{\zeta}{\alpha} + \beta}) \times \mathcal{N}(C, \|\cdot\|_E, \frac{\varepsilon}{\frac{\zeta}{\alpha} + \beta}).$$

Also, by Corollary 2,

$$\forall u \in S_{\eta}^{-}[f(\mathfrak{X})], \exists t \in T_C, \|u - t\|_F \leq \frac{L}{\alpha^2} \eta. \quad (\text{A.25})$$

We then apply Lemma 3 (with $\delta = \frac{L}{\alpha^2} \eta$) to prove that, for any $\varepsilon > 0$,

$$\mathcal{N}(S_{\eta}^{-}[f(\mathfrak{X})], \|\cdot\|_F, 2(\varepsilon + \frac{L}{\alpha^2} \eta)) \leq \mathcal{N}(T_C, \|\cdot\|_F, \varepsilon). \quad (\text{A.26})$$

Consequently,

$$\mathcal{N}(S_{\eta}^{-}[f(\mathfrak{X})], \|\cdot\|_F, 2(\varepsilon + \frac{L}{\alpha^2} \eta)) \leq \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\varepsilon}{\frac{\zeta}{\alpha} + \beta}) \times \mathcal{N}(C, \|\cdot\|_E, \frac{\varepsilon}{\frac{\zeta}{\alpha} + \beta}). \quad (\text{A.27})$$

All we need now is to control $\mathcal{N}(C, \|\cdot\|_F, \frac{\varepsilon}{\frac{\zeta}{\alpha} + \beta})$. Take $\overline{S[\mathfrak{X}]}$ a ε -net of $S[\mathfrak{X}]$ with respect to $\|\cdot\|_E$ and $(0, \frac{1}{\alpha}]$ a (ε/α) -net of $(0, \frac{1}{\alpha}]$ with respect to $|\cdot|$. Consider $c = \lambda u \in C$ with $\lambda \in (0, \frac{1}{\alpha}]$ and $u \in S[\mathfrak{X}]$. Then there exists $\bar{u} \in \overline{S[\mathfrak{X}]}$ such that $\|u - \bar{u}\|_E \leq \varepsilon$ and there exists $\bar{\lambda} \in (0, \frac{1}{\alpha}]$ such that $|\lambda - \bar{\lambda}| \leq \varepsilon/\alpha$. We consider $\bar{c} = \bar{\lambda} \bar{u}$ which belongs to C . Then $\|c - \bar{c}\|_E = \|\lambda u - \bar{\lambda} \bar{u}\|_E \leq |\lambda - \bar{\lambda}| \|\bar{u}\|_E + |\lambda| \|\bar{u} - u\|_E \leq 2\varepsilon/\alpha$. Thus for any $\varepsilon > 0$ we have

$$\mathcal{N}(C, \|\cdot\|_E, 2\varepsilon/\alpha) \leq \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \varepsilon) \times \mathcal{N}((0, \frac{1}{\alpha}], |\cdot|, \varepsilon/\alpha) = \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \varepsilon) \times \mathcal{N}((0, 1], |\cdot|, \varepsilon). \quad (\text{A.28})$$

Equivalently with a change of variable $\varepsilon \leftarrow 2\varepsilon/\alpha$ we have $\forall \varepsilon > 0$, $\mathcal{N}(C, \|\cdot\|_E, \varepsilon) \leq \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \frac{\varepsilon\alpha}{2}) \times \mathcal{N}((0, 1], |\cdot|, \frac{\varepsilon\alpha}{2})$. Overall, as for all $\varepsilon > 0$, $\mathcal{N}((0, 1], |\cdot|, \varepsilon) \leq \frac{1}{2\varepsilon}$,

$$\begin{aligned} & \mathcal{N}(S_{\eta}^{-}[f(\mathfrak{X})], \|\cdot\|_F, 2(\varepsilon + \frac{L}{\alpha^2} \eta)) \\ & \leq \left(\frac{\zeta + \beta\alpha}{\alpha^2\varepsilon} \right) \times \mathcal{N}(\mathfrak{X}, \|\cdot\|_E, \frac{\varepsilon}{\frac{\zeta}{\alpha} + \beta}) \times \mathcal{N}(S[\mathfrak{X}], \|\cdot\|_E, \frac{\alpha\varepsilon}{2(\frac{\zeta}{\alpha} + \beta)}), \end{aligned} \quad (\text{A.29})$$

which concludes the proof. \square

A.2.5. Proof of Lemma 1

Proof The proof is based on various computations and the identity

$$\Theta_1^{-1} + \Theta_2^{-1} = \Theta_1^{-1}(\Theta_1 + \Theta_2)\Theta_2^{-1}, \quad \forall (\Theta_1, \Theta_2) \in (\mathfrak{S}_{k,a,b}^{-1})^2. \quad (\text{A.30})$$

We will also use

$$\|\mathbf{AB}\|_{\text{Fro}} \leq \|\mathbf{A}\|_{2 \rightarrow 2} \|\mathbf{B}\|_{\text{Fro}} \quad \text{and} \quad \|\mathbf{AB}\|_{\text{Fro}} \leq \|\mathbf{A}\|_{\text{Fro}} \|\mathbf{B}\|_{2 \rightarrow 2}, \quad \forall \mathbf{A}, \mathbf{B} \in S_d(\mathbb{R}).$$

For the rest of the proof, we take $\Theta_1, \Theta_2 \in \mathfrak{S}_{k,a,b}^{-1}$. Hence we have $\|\Theta_1\|_{2 \rightarrow 2} \leq b$ and $\|\Theta_1^{-1}\|_{2 \rightarrow 2} \leq \frac{1}{a}$ (the same inequalities hold for Θ_2).

Now we can prove that

$$\begin{aligned} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\Lambda} &\leq C_{\text{Fro}} \|\Theta_1^{-1}(\Theta_1 - \Theta_2)(-\Theta_2)^{-1}\|_{\text{Fro}} \leq C_{\text{Fro}} \|\Theta_1^{-1}\|_{2 \rightarrow 2} \|\Theta_2^{-1}\|_{2 \rightarrow 2} \|\Theta_1 - \Theta_2\|_{\text{Fro}} \\ &\leq C_{\text{Fro}} \frac{1}{a^2} \|\Theta_1 - \Theta_2\|_{\text{Fro}}. \end{aligned}$$

This proves (5.6) with $\beta = \frac{C_{\text{Fro}}}{a^2}$. The reverse inequality for (5.5) can be proven by using this time

$$\begin{aligned} \|\Theta_1 - \Theta_2\|_{\text{Fro}} &= \|\Theta_1(\Theta_1^{-1} - \Theta_2^{-1})\Theta_2\|_{\text{Fro}} \leq \|\Theta_1\|_{2 \rightarrow 2} \|\Theta_2\|_{2 \rightarrow 2} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\text{Fro}} \\ &\leq b^2 \frac{1}{C_{\text{Fro}}} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\Lambda}. \end{aligned}$$

Thus we have (5.5) with $\alpha = \frac{C_{\text{Fro}}}{b^2}$.

For assumption **A-1**, we use the formula of the differential of the inverse of a matrix $\text{D inv}_{\Theta}[\mathbf{H}] = -\Theta^{-1}\mathbf{H}\Theta^{-1}$. We have:

$$\begin{aligned} \|\Theta_1^{-1} - \Theta_2^{-1} - \text{D inv}_{\Theta_2}(\Theta_1 - \Theta_2)\|_{\Lambda} &= \|\Theta_1^{-1} - \Theta_2^{-1} - [-\Theta_2^{-1}(\Theta_1 - \Theta_2)\Theta_2^{-1}]\|_{\Lambda} \\ &= \|\Theta_1^{-1}(\Theta_2 - \Theta_1)\Theta_2^{-1} - \Theta_2^{-1}(\Theta_2 - \Theta_1)\Theta_2^{-1}\|_{\Lambda} \\ &= \|(\Theta_1^{-1}(\Theta_2 - \Theta_1) - \Theta_2^{-1}(\Theta_2 - \Theta_1))\Theta_2^{-1}\|_{\Lambda} \\ &= \|(\Theta_1^{-1} - \Theta_2^{-1})(\Theta_2 - \Theta_1)\Theta_2^{-1}\|_{\Lambda} \\ &\leq C_{\text{Fro}} \|\Theta_2^{-1}\|_{2 \rightarrow 2} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\text{Fro}} \|\Theta_1 - \Theta_2\|_{\text{Fro}} \\ &\leq C_{\text{Fro}} \|\Theta_2^{-1}\|_{2 \rightarrow 2} \|\Theta_1^{-1}\|_{2 \rightarrow 2} \|\Theta_2^{-1}\|_{2 \rightarrow 2} \|\Theta_1 - \Theta_2\|_{\text{Fro}}^2 \\ &\leq C_{\text{Fro}} \frac{1}{a^3} \|\Theta_1 - \Theta_2\|_{\text{Fro}}^2. \end{aligned}$$

This gives **A-1** with $L = \frac{C_{\text{Fro}}}{a^3}$.

Now take \mathbf{M} , such that $\|\mathbf{M}\|_{\text{Fro}} \leq 1$. We have

$$\begin{aligned} \|\text{D inv}_{\Theta_1}(\mathbf{M}) - \text{D inv}_{\Theta_2}(\mathbf{M})\|_{\Lambda} &= \|\Theta_1^{-1}\mathbf{M}\Theta_1^{-1} - \Theta_2^{-1}\mathbf{M}\Theta_2^{-1}\|_{\Lambda} \\ &\leq \|\Theta_1^{-1}\mathbf{M}\Theta_1^{-1} - \Theta_1^{-1}\mathbf{M}\Theta_2^{-1}\|_{\Lambda} + \|\Theta_1^{-1}\mathbf{M}\Theta_2^{-1} - \Theta_2^{-1}\mathbf{M}\Theta_2^{-1}\|_{\Lambda} \\ &= \|\Theta_1^{-1}\mathbf{M}(\Theta_1^{-1} - \Theta_2^{-1})\|_{\Lambda} + \|(\Theta_1^{-1} - \Theta_2^{-1})\mathbf{M}\Theta_2^{-1}\|_{\Lambda} \\ &\leq C_{\text{Fro}} \|\Theta_1^{-1}\mathbf{M}\|_{\text{Fro}} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\text{Fro}} + C_{\text{Fro}} \|\mathbf{M}\Theta_2^{-1}\|_{\text{Fro}} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\text{Fro}} \\ &\leq C_{\text{Fro}} (\|\Theta_1^{-1}\|_{2 \rightarrow 2} + \|\Theta_2^{-1}\|_{2 \rightarrow 2}) \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\text{Fro}} \\ &\leq 2C_{\text{Fro}} \frac{1}{a} \|\Theta_1^{-1} - \Theta_2^{-1}\|_{\text{Fro}} \\ &\leq 2C_{\text{Fro}} \frac{1}{a^3} \|\Theta_1 - \Theta_2\|_{\text{Fro}}. \end{aligned}$$

This gives **A-2** with $\zeta = 2\frac{C_{\text{Fro}}}{a^3} = 2L$. \square

A.2.6. Proof of Lemma 2

In order to prove the result we will use the following lemma.

Lemma 7 *Consider*

$$\mathfrak{W}_k = \{\Theta \in S_d(\mathbb{R}) : \|\Theta\|_0 \leq d + 2k, \|\Theta\|_{\text{Fro}} \leq 1\}. \quad (\text{A.31})$$

Then

$$\mathcal{N}(\mathfrak{W}_k, \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{9}{\varepsilon}\right)^{d+k}. \quad (\text{A.32})$$

Proof Take $\Theta \in \mathfrak{W}_k$, it can be written as $\Theta = \mathbf{D} + \mathbf{T} + \mathbf{T}^\top$ where \mathbf{D} is diagonal with d positive elements, \mathbf{T} is a strictly upper triangular matrix with at most k non zero elements. We have also that $\|\mathbf{D}\|_{\text{Fro}} \leq \|\Theta\|_{\text{Fro}} \leq 1$ and same for $\mathbf{T}, \mathbf{T}^\top$. Consider \bar{D} a $\varepsilon/3$ -net for the diagonal and \bar{T} a $\varepsilon/3$ -net for the upper triangle, both with respect to the $\|\cdot\|_{\text{Fro}}$ norm. Then with standard covering arguments $|\bar{D}| \leq (9/\varepsilon)^d$ and $|\bar{T}| \leq \binom{d(d-1)}{k} \left(\frac{9}{\varepsilon}\right)^k$ because it is included in the unit ball of k -sparse vector in dimension $\frac{d(d-1)}{2}$ (see e.g. [30]). Consider $\overline{\mathfrak{W}}_k = \{\mathbf{D}_* + \mathbf{T}_* + \mathbf{T}_*^\top, (\mathbf{D}_*, \mathbf{T}_*) \in \bar{D} \times \bar{T}\}$. Then $|\overline{\mathfrak{W}}_k| \leq \binom{d(d-1)}{k} \left(\frac{9}{\varepsilon}\right)^k \left(\frac{9}{\varepsilon}\right)^d = \left(\frac{d(d-1)}{k}\right) \left(\frac{9}{\varepsilon}\right)^{d+k}$. Also, for any $\Theta = \mathbf{D} + \mathbf{T} + \mathbf{T}^\top$ there exists $(\mathbf{D}_*, \mathbf{T}_*) \in \bar{D} \times \bar{T}$ such that $\|\mathbf{D} - \mathbf{D}_*\|_{\text{Fro}} \leq \varepsilon/3, \|\mathbf{T} - \mathbf{T}_*\|_{\text{Fro}} \leq \varepsilon/3$. Hence:

$$\|\mathbf{D} + \mathbf{T} + \mathbf{T}^\top - (\mathbf{D}_* + \mathbf{T}_* + \mathbf{T}_*^\top)\|_{\text{Fro}} \leq \|\mathbf{D} - \mathbf{D}_*\|_{\text{Fro}} + 2\|\mathbf{T} - \mathbf{T}_*\|_{\text{Fro}} \leq \varepsilon. \quad (\text{A.33})$$

Hence,

$$\mathcal{N}(\mathfrak{W}_k, \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \binom{d(d-1)}{k} \left(\frac{9}{\varepsilon}\right)^{d+k} \leq \left(\frac{ed(d-1)}{2k}\right)^k \left(\frac{9}{\varepsilon}\right)^{d+k} \leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{9}{\varepsilon}\right)^{d+k}, \quad (\text{A.34})$$

where in the last inequality we used the bound [30, Lemma C.5]. Note that we only considered the fact that \mathfrak{W} is the space of symmetric an $d + 2k$ sparse matrices. Restricting to positive definite matrices could be an avenue for further improvements. \square

As a consequence we can prove Lemma 2 as follows.

Proof of Lemma 2 Recall that

$$\mathfrak{S}_{k,a,b}^{-1} = \{\Theta \in S_d^{++}(\mathbb{R}) : \|\Theta\|_0 \leq d + 2k, \text{spec}(\Theta) \subseteq [a, b]\}.$$

Consider $\Theta \in \mathfrak{S}_{k,a,b}^{-1}$. We have $\|\Theta\|_{2 \rightarrow 2} \leq b$ which implies $\|\Theta\|_{\text{Fro}} \leq \sqrt{d}\|\Theta\|_{2 \rightarrow 2} \leq \sqrt{db}$. Thus $\mathfrak{S}_{k,a,b}^{-1} \subset \{\Theta \in S_d^{++}(\mathbb{R}) : \|\Theta\|_0 \leq d + 2k, \|\Theta\|_{\text{Fro}} \leq \sqrt{db}\} = \sqrt{db}\mathfrak{W}$. Consequently,

$$\mathcal{N}(\mathfrak{S}_{k,a,b}^{-1}, \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \mathcal{N}(\mathfrak{W}, \|\cdot\|_{\text{Fro}}, \varepsilon/2\sqrt{db}) \stackrel{\text{Lemma 7}}{\leq} \left(\frac{ed^2}{2k}\right)^k \left(\frac{18\sqrt{db}}{\varepsilon}\right)^{d+k}, \quad (\text{A.35})$$

which concludes the first part. For the second part we recall the definition of the normalized secant set

$$S[\mathfrak{S}_{k,a,b}^{-1}] = \left\{ \frac{\Theta_1 - \Theta_2}{\|\Theta_1 - \Theta_2\|_{\text{Fro}}} : (\Theta_1, \Theta_2) \in (\mathfrak{S}_{k,a,b}^{-1})^2, \|\Theta_1 - \Theta_2\|_{\text{Fro}} > 0 \right\}.$$

Moreover, if $\mathbf{U} = \Theta_1 - \Theta_2 \in \mathfrak{S}_{k,a,b}^{-1} - \mathfrak{S}_{k,a,b}^{-1}$ then $\|\mathbf{U}\|_0 \leq d + 4k$. Indeed, Θ_1 and Θ_2 have d nonzeros elements on the diagonal (since both are positive definite) and since these matrices are symmetric then

they have at most k nonzeros elements in the upper-triangular (*resp.* lower-triangular) part. Thus $\Theta_1 - \Theta_2$ has at most $2k$ nonzeros elements in the upper-triangular (*resp.* lower-triangular) part. Consequently,

$$S[\mathfrak{S}_{k,a,b}^{-1}] \subset \{\mathbf{M} \in S_d(\mathbb{R}) : \|\mathbf{M}\|_{\text{Fro}} \leq 1, \|\mathbf{M}\|_0 \leq d + 4k\} = \mathfrak{W}_{2k}.$$

This gives

$$\mathcal{N}(S[\mathfrak{S}_{k,a,b}^{-1}], \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \mathcal{N}(\mathfrak{W}_{2k}, \|\cdot\|_{\text{Fro}}, \varepsilon/2) \leq \left(\frac{ed^2}{4k}\right)^{2k} \left(\frac{18}{\varepsilon}\right)^{d+2k}. \quad (\text{A.36})$$

□

A.2.7. Proof of Corollary 1

Proof From Lemma 1, the assumptions of Theorem 5.1 hold for $f = \text{inv}$, $\mathfrak{S} = \mathfrak{S}_{k,a,b}^{-1}$, $\|\cdot\|_E = \|\cdot\|_{\text{Fro}}$ and $\|\cdot\|_F = \|\cdot\|_{\Lambda}$. Thus, for any $\eta, \varepsilon' > 0$ we have the inequality

$$\begin{aligned} & \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_{\Lambda}, 2\left[\varepsilon' + \frac{L}{\alpha^2}\eta\right]) \\ & \leq \mathcal{N}(\mathfrak{S}_{k,a,b}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\eta}{8\beta}\left[\varepsilon' + \frac{L}{\alpha^2}\eta\right])^2 \\ & \quad + \frac{\zeta + \beta\alpha}{\alpha^2\varepsilon'} \mathcal{N}(\mathfrak{S}_{k,a,b}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\alpha\varepsilon'}{\zeta + \beta\alpha}) \times \mathcal{N}(S[\mathfrak{S}_{k,a,b}^{-1}], \|\cdot\|_{\text{Fro}}, \frac{\alpha^2\varepsilon'}{2(\zeta + \beta\alpha)}). \end{aligned} \quad (\text{A.37})$$

Let $\varepsilon > 0$ be fixed. We define $\eta \triangleq \frac{\alpha^2}{4L}\varepsilon$ and $\varepsilon' \triangleq \varepsilon/4$. Then we have $\varepsilon = 2\left[\varepsilon' + \frac{L}{\alpha^2}\eta\right]$. With these ε', η and Lemma 2 we obtain

$$\begin{aligned} & \mathcal{N}(\mathfrak{S}_{k,a,b}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\eta}{8\beta}\left[\varepsilon' + \frac{L}{\alpha^2}\eta\right])^2 \leq \left(\frac{ed^2}{2k}\right)^{2k} \left(\frac{288\beta\sqrt{db}}{\eta\varepsilon}\right)^{2(d+k)} \\ & \mathcal{N}(\mathfrak{S}_{k,a,b}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\alpha\varepsilon}{4(\zeta + \beta\alpha)}) \leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{72\sqrt{d}(\zeta + \beta\alpha)}{\alpha\varepsilon}\right)^{d+k} \\ & \mathcal{N}(S[\mathfrak{S}_{k,a,b}^{-1}], \|\cdot\|_{\text{Fro}}, \frac{\alpha^2\varepsilon}{8(\zeta + \beta\alpha)}) \leq \left(\frac{ed^2}{4k}\right)^{2k} \left(\frac{144(\zeta + \beta\alpha)}{\alpha^2\varepsilon}\right)^{d+2k}. \end{aligned} \quad (\text{A.38})$$

Thus, inequality (A.37) yields

$$\begin{aligned} & \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_{\Lambda}, \varepsilon) \\ & \stackrel{(\text{A.38})}{\leq} \left(\frac{ed^2}{2k}\right)^{2k} \left(\frac{288\beta\sqrt{db}}{\eta\varepsilon}\right)^{2(d+k)} \\ & \quad + \frac{4(\zeta + \beta\alpha)}{\alpha^2\varepsilon} \left(\frac{ed^2}{2k}\right)^k \left(\frac{72\sqrt{d}(\zeta + \beta\alpha)}{\alpha\varepsilon}\right)^{d+k} \times \left(\frac{ed^2}{4k}\right)^{2k} \left(\frac{144(\zeta + \beta\alpha)}{\alpha^2\varepsilon}\right)^{d+2k} \\ & \leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{288\beta\sqrt{db}}{\eta\varepsilon}\right)^{2(d+k)} + \frac{4(\zeta + \beta\alpha)}{\alpha^2\varepsilon} \left(\frac{72\sqrt{d}(\zeta + \beta\alpha)}{\alpha\varepsilon}\right)^{d+k} \left(\frac{144(\zeta + \beta\alpha)}{\alpha^2\varepsilon}\right)^{d+2k} \right] \end{aligned}$$

Using the definition $\eta = \frac{\alpha^2}{4L} \varepsilon$ we obtain

$$\begin{aligned} & \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_\Lambda, \varepsilon) \\ & \leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152\sqrt{db}\beta L}{\alpha^2\varepsilon^2}\right)^{2(d+k)} + \frac{4(\zeta + \beta\alpha)}{\alpha^2\varepsilon} \left(\frac{72b\sqrt{d}(\zeta + \beta\alpha)}{\alpha\varepsilon}\right)^{d+k} \left(\frac{144(\zeta + \beta\alpha)}{\alpha^2\varepsilon}\right)^{d+2k} \right]. \end{aligned}$$

Now, from Lemma 1, we have $\beta = \frac{C_{\text{Fro}}}{a^2}$, $L = \frac{C_{\text{Fro}}}{a^3}$, $\alpha = \frac{c_{\text{Fro}}}{b^2}$, $\zeta = 2L = 2\frac{C_{\text{Fro}}}{a^3}$. So,

$$\begin{aligned} \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_\Lambda, \varepsilon) & \leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152b\sqrt{d}\frac{C_{\text{Fro}}}{a^2}\frac{C_{\text{Fro}}}{a^3}}{\frac{c_{\text{Fro}}^2}{b^4}\varepsilon^2}\right)^{2(d+k)} \right. \\ & \quad \left. + \frac{4(2\frac{C_{\text{Fro}}}{a^3} + \frac{C_{\text{Fro}}}{a^2}\frac{c_{\text{Fro}}}{b^2})}{\frac{c_{\text{Fro}}^2}{b^4}\varepsilon} \left(\frac{72b\sqrt{d}(2\frac{C_{\text{Fro}}}{a^3} + \frac{C_{\text{Fro}}}{a^2}\frac{c_{\text{Fro}}}{b^2})}{\frac{c_{\text{Fro}}}{b^2}\varepsilon}\right)^{d+k} \left(\frac{144(2\frac{C_{\text{Fro}}}{a^3} + \frac{C_{\text{Fro}}}{a^2}\frac{c_{\text{Fro}}}{b^2})}{\frac{c_{\text{Fro}}^2}{b^4}\varepsilon}\right)^{d+2k} \right] \\ & = \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152\sqrt{d}C_{\text{Fro}}^2b^5}{\varepsilon^2c_{\text{Fro}}^2a^5}\right)^{2(d+k)} \right. \\ & \quad \left. + \frac{4C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2b^4}{c_{\text{Fro}}a^3} + \frac{b^2}{a^2}\right) \left(\frac{72C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(2\frac{b^3}{a^3} + \frac{b}{a^2}c_{\text{Fro}}\right)\sqrt{d}\right)^{d+k} \left(\frac{144C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2b^4}{c_{\text{Fro}}a^3} + \frac{b^2}{a^2}\right)\right)^{d+2k} \right]. \end{aligned} \tag{A.39}$$

We will simplify this expression using the homogeneity of the normalized secant set. More precisely if $\frac{\Sigma_1 - \Sigma_2}{\|\Sigma_1 - \Sigma_2\|_\Lambda} \in S[\mathfrak{S}_{k,a,b}]$ then for any $t > 0$, $\frac{t\Sigma_1 - t\Sigma_2}{\|t\Sigma_1 - t\Sigma_2\|_\Lambda} \in S[\mathfrak{S}_{k,a,b}]$. This implies that $\forall t > 0$, $S[\mathfrak{S}_{k,a,b}] =$

$S[\mathfrak{S}_{k,t,a,t,b}]$. In particular for $t \triangleq \frac{a}{b^2}$ the previous expression (A.39) gives

$$\begin{aligned}
& \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_{\Lambda}, \varepsilon) \\
&= \mathcal{N}(S[\mathfrak{S}_{k,t,a,t,b}], \|\cdot\|_{\Lambda}, \varepsilon) \\
&\leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152\sqrt{d} C_{\text{Fro}}^2 b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5}\right)^{2(d+k)} \right. \\
&\quad \left. + \frac{4 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} t \frac{b^4}{a^3} + \frac{b^2}{a^2}\right) \left(\frac{72 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} (2 \frac{b^3}{a^3} + \frac{b}{ta^2} c_{\text{Fro}}) \sqrt{d}\right)^{d+k} \left(\frac{144 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} t \frac{b^4}{a^3} + \frac{b^2}{a^2}\right)\right)^{d+2k} \right] \\
&= \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152\sqrt{d} C_{\text{Fro}}^2 b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5}\right)^{2(d+k)} \right. \\
&\quad \left. + \frac{4 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1\right) \frac{b^2}{a^2} \left(\frac{72 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} (2 + c_{\text{Fro}}) \frac{b^3}{a^3} \sqrt{d}\right)^{d+k} \left(\frac{144 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1\right) \frac{b^2}{a^2}\right)^{d+2k} \right] \\
&\leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152\sqrt{d} C_{\text{Fro}}^2 b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5}\right)^{2(d+k)} \right. \\
&\quad \left. + \left(\frac{72 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} (2 + c_{\text{Fro}}) \frac{b^3}{a^3} \sqrt{d}\right)^{d+k} \left(\frac{144 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1\right) \frac{b^2}{a^2}\right)^{d+2k+1} \right].
\end{aligned}$$

Now, to simplify the expression, remark that $\frac{72 C_{\text{Fro}}}{\varepsilon c_{\text{Fro}}} (2 + c_{\text{Fro}}) \frac{b^3}{a^3} \sqrt{d} \geq 1$ as $C_{\text{Fro}}/c_{\text{Fro}} \geq 1$. Therefore we can increase its power from $d+k$ to $d+2k+1$ to match with the other multiplicative term. This yields

$$\begin{aligned}
& \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_{\Lambda}, \varepsilon) \\
&\leq \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(\frac{1152 C_{\text{Fro}}^2 \sqrt{d} b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5}\right)^{2(d+k)} + \left(\frac{72 \times 144 C_{\text{Fro}}^2 \sqrt{d}}{\varepsilon^2 c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1\right) \frac{b^5}{a^5}\right)^{d+2k+1} \right] \\
&= \left(\frac{ed^2}{2k}\right)^{4k} \left[\left(c_0 \frac{\sqrt{d} C_{\text{Fro}}^2 b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5}\right)^{2(d+k)} + \left(c_1 \frac{\sqrt{d} C_{\text{Fro}}^2}{\varepsilon^2 c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1\right) \frac{b^5}{a^5}\right)^{d+2k+1} \right],
\end{aligned}$$

where c_0, c_1 are absolute constants greater than 1. This concludes the proof. \square

A.3. Proofs of Section 5.2

A.3.1. Proofs of the rank-one projection operator properties

The goal of this section is to prove Proposition 3 and 4. Before that, we prove several results that will become handy afterwards. Firstly, we state a result that will be useful to leverage results from the Gaussian case to the uniform case.

Lemma 8 Let \mathbf{u} and ρ be independent variables with the following distributions : $\mathbf{u} \sim \mathcal{U}(\mathbb{S}^{d-1})$ is a uniform vector on the hyper-sphere and $\rho^2 \sim \chi^2(d)$ is a chi-square variable with d degrees of freedom. Then, $\rho \frac{\mathbf{u}}{\sqrt{d}} \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ is a standard normal vector.

Now, a lower bound is derived for the Λ -norm.

Proposition 5 For any $\mathbf{M} \in S_d(\mathbb{R})$ and $\Lambda \in \{\Lambda_G, \Lambda_U\}$, we have

$$\|\mathbf{M}\|_\Lambda \geq \frac{2}{9\sqrt{15d}} (\|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})|). \quad (\text{A.40})$$

Proof We use the fact that for any real random variable X , whenever its fourth moment exists, we have²

$$\mathbb{E}[|X|] \geq \sqrt{\frac{\mathbb{E}[X^2]^3}{\mathbb{E}[X^4]}}.$$

First, we focus on the Gaussian case. Using³ $X = \mathbf{a}^\top \mathbf{M} \mathbf{a} \stackrel{(d)}{=} \sum \lambda_k b_k^2$, where $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$, the (λ_k) are the eigenvalues of \mathbf{M} and the (b_k) are *i.i.d* Gaussian random variables of variance $1/d$, we can obtain the following bounds:

$$\begin{aligned} \mathbb{E}[X^2] &= \frac{2}{d^2} \|\mathbf{M}\|_{\text{Fro}}^2 + \frac{1}{d^2} \text{tr}(\mathbf{M})^2 \geq \frac{2}{3d} (\|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})|)^2, \\ \mathbb{E}[X^4] &= \sum_i \lambda_i^4 \mathbb{E}[b_i^8] + \sum_{i \neq j} \lambda_i \lambda_j^3 \mathbb{E}[b_i^2] \mathbb{E}[b_j^6] + \sum_{i \neq j} \lambda_i^2 \lambda_j^2 \mathbb{E}[b_i^4] \mathbb{E}[b_j^4] \\ &\quad + \sum_{i \neq j \neq k} \lambda_i \lambda_j \lambda_k^2 \mathbb{E}[b_i^2] \mathbb{E}[b_j^2] \mathbb{E}[b_k^4] + \sum_{i \neq j \neq k \neq l} \lambda_i \lambda_j \lambda_k \lambda_l \mathbb{E}[b_i^2] \mathbb{E}[b_j^2] \mathbb{E}[b_k^2] \mathbb{E}[b_l^2] \\ &= \frac{105}{d^4} \sum_i \lambda_i^4 + \frac{15}{d^4} \sum_{i \neq j} \lambda_i \lambda_j^3 + \frac{9}{d^4} \sum_{i \neq j} \lambda_i^2 \lambda_j^2 + \frac{3}{d^4} \sum_{i \neq j \neq k} \lambda_i \lambda_j \lambda_k^2 + \frac{1}{d^4} \sum_{i \neq j \neq k \neq l} \lambda_i \lambda_j \lambda_k \lambda_l \\ &= \frac{1}{d^4} \left[90 \|\lambda\|_4^4 + 12 \text{tr}(\mathbf{M}) \sum_i \lambda_i^3 + 6 \|\mathbf{M}\|_{\text{Fro}}^4 + 2 \text{tr}(\mathbf{M})^2 \|\mathbf{M}\|_{\text{Fro}}^2 + \text{tr}(\mathbf{M})^4 \right] \\ &\leq \frac{90}{d^4} (\|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})|)^4. \end{aligned}$$

This yields equation (A.40) for the Gaussian case.

For the uniform case, considering the independent random variables $\mathbf{u} \sim \mathcal{U}(\mathbb{S}^{d-1})$, $\rho^2 \sim \chi^2(d)$, from Lemma 8 we have that

$$\mathbb{E} \left[\left| \mathbf{a}^\top \mathbf{M} \mathbf{a} \right| \right] = \mathbb{E} \left[\left| \left(\rho \frac{1}{\sqrt{d}} \mathbf{u} \right)^\top \mathbf{M} \left(\rho \frac{1}{\sqrt{d}} \mathbf{u} \right) \right| \right] = \frac{1}{d} \mathbb{E}[\rho^2] \mathbb{E} \left[\left| \mathbf{u}^\top \mathbf{M} \mathbf{u} \right| \right] = \mathbb{E} \left[\left| \mathbf{u}^\top \mathbf{M} \mathbf{u} \right| \right].$$

So (A.40) also holds in the uniform case. \square

² It comes from applying the Hölder inequality to $\mathbb{E}[|X|^{2/3} |X|^{4/3}]$ with $1/p = 2/3$ and $1/q = 1/3$.

³ The last equality is obtained from the rotation invariance of the multivariate normal distribution.

The proof of Proposition 3 is based on a concentration inequality for subexponential variables. Here, we prove that the variables at play are indeed subexponential by providing an upper-bound on their subexponential norm. Recall that for a random variable X , its subexponential norm is defined by $\|X\|_{\psi_1} = \inf\{s > 0, \mathbb{E}[e^{|X|/s}] \leq 2\}$.

Proposition 6 *For any $\mathbf{M} \in S_d(\mathbb{R})$ and for $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{d}I_d)$ and $\mathbf{u} \sim \mathcal{U}(\mathbb{S}^{d-1})$, the following controls hold :*

$$\| |\mathbf{a}^\top \mathbf{M} \mathbf{a}| - \mathbb{E}[|\mathbf{a}^\top \mathbf{M} \mathbf{a}|] \|_{\psi_1} \leq \frac{2}{d \log 2} \left(\frac{76}{9} e^2 \cdot \|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})| \right), \quad (\text{A.41})$$

$$\| |\mathbf{u}^\top \mathbf{M} \mathbf{u}| - \mathbb{E}[|\mathbf{u}^\top \mathbf{M} \mathbf{u}|] \|_{\psi_1} \leq \frac{8e}{d \log 2} \left(\frac{76}{9} e^2 \cdot \|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})| \right). \quad (\text{A.42})$$

Proof of Equation (A.41) This proof revolves around the different characterizations of subexponentiality (indexed from (a) to (e)) presented in Proposition 2.7.1 of [74]. First from the centering Lemma (see Exercise 2.7.10 in [74]), we have the existence of a constant $C_1 > 0$ such that $\| |\mathbf{a}^\top \mathbf{M} \mathbf{a}| - \mathbb{E}[|\mathbf{a}^\top \mathbf{M} \mathbf{a}|] \|_{\psi_1} \leq C_1 \| |\mathbf{a}^\top \mathbf{M} \mathbf{a}| \|_{\psi_1} = C_1 \| \mathbf{a}^\top \mathbf{M} \mathbf{a} \|_{\psi_1}$. Let us denote by X our random variable of interest $X = \mathbf{a}^\top \mathbf{M} \mathbf{a}$, and Y its centered version $Y = X - \mathbb{E}[X]$. Working with Y , we now characterize the constant $K_5(Y)$ appearing in statement (e) of Proposition 2.7.1 in [74]. First, we can write $Y = \sum_k \lambda_k d^{-1} (z_k^2 - 1)$ where λ_k are the eigenvalues of \mathbf{M} and z_k is a standard normal variable. Remark that the centered $\chi^2(1)$ variables $z_k^2 - 1$ verify (e) with a certain constant K_{χ^2} . For $|t| \leq d/(K_{\chi^2} \cdot \|\mathbf{M}\|_{\text{Fro}})$ we have that

$$\mathbb{E}[e^{tY}] = \prod_{k=1}^d \mathbb{E}\left[e^{t \frac{\lambda_k}{d} (z_k^2 - 1)} \right] \leq \prod_{k=1}^d \mathbb{E}\left[e^{\frac{1}{d^2} K_{\chi^2}^2 \lambda_k^2 t^2} \right] = e^{\frac{1}{d^2} K_{\chi^2}^2 \|\mathbf{M}\|_{\text{Fro}}^2 t^2}.$$

This yield $K_5(Y) \leq \frac{1}{d} K_{\chi^2} \|\mathbf{M}\|_{\text{Fro}}$. Now, by considering statement (c), we have that $K_3(X) \leq K_3(Y) + |\mathbb{E}[X]| \leq C_{3,5} K_{\chi^2} \cdot \frac{1}{d} \|\mathbf{M}\|_{\text{Fro}} + \frac{1}{d} |\text{tr}(\mathbf{M})|$ (where $C_{3,5}$ is the universal constant allowing to pass from (c) to (e)). Finally, gathering up the pieces we have

$$\| |\mathbf{a}^\top \mathbf{M} \mathbf{a}| - \mathbb{E}[|\mathbf{a}^\top \mathbf{M} \mathbf{a}|] \|_{\psi_1} \leq \frac{C_1 C_{4,3}}{d} (C_{3,5} K_{\chi^2} \cdot \|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})|),$$

where $C_{4,3}$ is the constant allowing to pass from (d) (statement defining the ψ_1 -norm) to (c). To conclude the proof, it suffices to find the values of the various constants. This is completely general and does not depend on the rank-one projection considered here. The various constants can be set as follows :

$$\begin{aligned} C_1 &= 2, & C_{3,5} &= 4e^2, \\ C_{4,3} &= \frac{1}{\log 2}, & K_{\chi^2} &= \frac{19}{9}. \end{aligned}$$

Let us start by computing C_1 . Let X be a subexponential random variable. We have, for any $s > 0$,

$$\mathbb{E}\left[e^{\frac{||X| - \mathbb{E}[|X]|}{s}} \right] \stackrel{(\Delta \text{ ineq.})}{\leq} \mathbb{E}\left[e^{\frac{|X|}{s}} \right] e^{-\frac{\mathbb{E}[|X|]}{s}} \stackrel{(\text{Jensen})}{\leq} \mathbb{E}\left[e^{\frac{|X|}{s}} \right]^2 \stackrel{(\text{Jensen})}{\leq} \mathbb{E}\left[e^{\frac{2|X|}{s}} \right].$$

The last term is smaller than 2 for $s/2 \geq \|X\|_{\psi_1}$, so we have that $\| |X| - \mathbb{E}[|X]| \|_{\psi_1} \leq 2\|X\|_{\psi_1}$, yielding $C_1 = 2$.

For the constant $C_{3,5}$ we use that $C_{3,5} \leq C_{3,2}C_{2,5}$. In [74], the value of $C_{2,5}$ is given and equals $2e$. Let us focus on $C_{3,2}$ and assume that $K_2(X) = 1$. For any λ such that $0 \leq \lambda \leq 1/(2e)$, we have the following inequalities

$$\mathbb{E} \left[e^{\lambda|X|} \right] = 1 + \sum_{p \geq 1} \frac{\lambda^p \mathbb{E} [|X|^p]}{p!} \stackrel{K_2(X)=1}{\leq} 1 + \sum_{p \geq 1} \frac{\lambda^p p^p}{p!} \leq 1 + \sum_{p \geq 1} \frac{\lambda^p p^p}{(p/e)^p} = \frac{1}{1 - \lambda e} \leq e^{2e\lambda}.$$

Thus, $K_3(X) \leq 2e$. So we can take $C_{3,5} = 4e^2$.

For $C_{4,3}$ assume that $K_3(X) = 1$. For $\lambda \leq \log 2$ we have

$$\mathbb{E} \left[e^{\lambda|X|} \right] \stackrel{K_3(X)=1}{\leq} e^\lambda \leq 2.$$

So we can take $C_{4,3} = 1/\log 2$.

The value of K_{χ^2} can be obtain from the following computation. Let $b \sim \mathcal{N}(0, 1)$ be a standard normal distribution, then for $\lambda < 1/2$

$$\mathbb{E} \left[e^{\lambda(b^2-1)} \right] = \frac{e^{-\lambda}}{\sqrt{1-2\lambda}}.$$

We can show that for $K = \frac{2}{1-x_0}$ where x_0 is the smallest solution of $e^x = e^3 x$, we have for all λ such that $|\lambda| \leq 1/K$, $\mathbb{E} \left[e^{\lambda(b^2-1)} \right] \leq e^{K^2 \lambda^2}$. A numerical approximation gives $K \simeq 2.1107 \dots$. So we can take $K_{\chi^2} = 19/9$. This finishes the proof. \square

Proof of Equation (A.42). As in the proof of (A.41), we start by decentering : $\| |\mathbf{u}^\top \mathbf{M} \mathbf{u}| - \mathbb{E} [|\mathbf{u}^\top \mathbf{M} \mathbf{u}|] \|_{\psi_1} \leq C_1 \|\mathbf{u}^\top \mathbf{M} \mathbf{u}\|_{\psi_1}$. Then, recall (see [74]) that there exists an absolute constant $C_{4,2}$ such that $\|\mathbf{u}^\top \mathbf{M} \mathbf{u}\|_{\psi_1} \leq C_{4,2} K_2(\mathbf{u}^\top \mathbf{M} \mathbf{u})$ where $K_2(\mathbf{u}^\top \mathbf{M} \mathbf{u})$ is the smallest constant K such that for all $p \geq 1$, $\mathbb{E} [|\mathbf{u}^\top \mathbf{M} \mathbf{u}|^p] \leq K^p p^p$. From Lemma 8, we have that

$$\mathbb{E} [|\mathbf{a}^\top \mathbf{M} \mathbf{a}|^p] = \frac{1}{d^p} \mathbb{E} [\rho^{2p}] \mathbb{E} [|\mathbf{u}^\top \mathbf{M} \mathbf{u}|^p] = \frac{d(d+2) \dots (d+2(p-1))}{d^p} \mathbb{E} [|\mathbf{u}^\top \mathbf{M} \mathbf{u}|^p] \geq \mathbb{E} [|\mathbf{u}^\top \mathbf{M} \mathbf{u}|^p].$$

Hence, we have that $K_2(\mathbf{u}^\top \mathbf{M} \mathbf{u}) \leq K_2(\mathbf{a}^\top \mathbf{M} \mathbf{a})$. Then, using (A.41) and the existence of a constant $C_{2,4}$ such that $K_2(\mathbf{a}^\top \mathbf{M} \mathbf{a}) \leq C_{2,4} \|\mathbf{a}^\top \mathbf{M} \mathbf{a}\|_{\psi_1}$. We have

$$\|\mathbf{u}^\top \mathbf{M} \mathbf{u}\|_{\psi_1} \leq C_{4,2} C_{2,4} \|\mathbf{a}^\top \mathbf{M} \mathbf{a}\|_{\psi_1} \leq \frac{C_{4,2} C_{2,4} C_{4,3}}{d} (C_{3,5} K_{\chi^2} \cdot \|\mathbf{M}\|_{\text{Fro}} + |\text{tr}(\mathbf{M})|).$$

To finish the proof, we show that $C_{4,2}$ and $C_{2,4}$ can be chosen as

$$C_{4,2} = 2e, \quad C_{2,4} = 2.$$

For $C_{4,2}$, we need to prove that $\|X\|_{\psi_1} \leq C_{4,2}K_2(X)$, for any subexponential variable X . Without loss of generality, we can always assume that $K_2(X) = 1$. Thus, for $s > e$ we have

$$\begin{aligned} \mathbb{E} \left[e^{|X|/s} \right] &= 1 + \sum_{k=1}^{\infty} \frac{\mathbb{E} [|X|^k]}{k! s^k} \\ &\leq 1 + \sum_{k=1}^{\infty} \frac{k^k}{k! s^k} \stackrel{(*)}{\leq} \sum_{k=0}^{\infty} \left(\frac{e}{s} \right)^k = \frac{1}{1 - e/s}, \end{aligned}$$

where the $(*)$ inequality comes from the Stirling approximation $k! \geq (k/e)^k$. Thus, for $s \geq 2e$ we have $\mathbb{E} \left[e^{|X|/s} \right] \leq 2$. So we can take $C_{4,2} = 2e$.

For $C_{2,4}$, assume that $\|X\|_{\psi_1} = 1$. Then, for any $p \geq 1$,

$$\begin{aligned} \mathbb{E} [|X|^p] &= \int_0^{\infty} \mathbb{P} (|X|^p > u) du = \int_0^{\infty} \mathbb{P} (|X| > t) p t^{p-1} dt \\ &\leq \int_0^{\infty} \mathbb{E} \left[e^{|X|} \right] e^{-t} p t^{p-1} dt \leq 2 \int_0^{\infty} e^{-t} p t^{p-1} dt = 2p! \leq (2p)^p. \end{aligned}$$

Thus, we can take $C_{2,4} = 2$. \square

All this previous results allow now to prove Proposition 3, which is recalled below.

Proposition 3 Consider the operator $\mathcal{A}(\Sigma) = \frac{1}{m} (\mathbf{a}_1^\top \Sigma \mathbf{a}_1, \dots, \mathbf{a}_m^\top \Sigma \mathbf{a}_m)^\top$, where the vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ are i.i.d. and either follow $\Lambda = \Lambda_G = \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ or $\Lambda = \Lambda_G = \mathcal{U}(\mathbb{S}^{d-1})$ with the associated norm $\|\Sigma\|_\Lambda = \mathbb{E}_{\mathbf{a} \sim \Lambda} [|\mathbf{a}^\top \Sigma \mathbf{a}|]$. Then, for any $\mathbf{U} \in S_d(\mathbb{R})$ satisfying $\|\mathbf{U}\|_\Lambda = 1$ and for any $t > 0$, we have

$$\mathbb{P} \left(\left| \|\mathcal{A}(\mathbf{U})\|_1 - 1 \right| > t \right) \leq 2 \exp \left(-\frac{m}{8e^2} \min \left(\frac{t^2}{K_\Lambda^2}, \frac{t}{K_\Lambda} \right) \right), \quad (5.14)$$

where K_Λ is an absolute constant given by $K_{\Lambda_G} = \frac{76e^2\sqrt{15}}{\log 2}$ and $K_{\Lambda_U} = \frac{304e^3\sqrt{15}}{\log 2}$.

Proof of Proposition 3 From our previous results, we have :

$$\| |\mathbf{a}^\top \mathbf{U} \mathbf{a}| - \mathbb{E} [|\mathbf{a}^\top \mathbf{U} \mathbf{a}|] \|_{\psi_1} \stackrel{(A.41)}{\leq} \frac{2 \times 76 e^2}{9d \log 2} (\|\mathbf{U}\|_{\text{Fro}} + |\text{tr}(\mathbf{U})|) \stackrel{(A.40)}{\leq} \frac{76e^2\sqrt{15}}{\log 2} \|\mathbf{U}\|_\Lambda \stackrel{\|\mathbf{U}\|_\Lambda=1}{\leq} \frac{76e^2\sqrt{15}}{\log 2}. \quad (A.43)$$

Similarly, in the uniform case, we obtain $\| |\mathbf{u}^\top \mathbf{U} \mathbf{u}| - \mathbb{E} [|\mathbf{u}^\top \mathbf{U} \mathbf{u}|] \|_{\psi_1} \leq \frac{304e^3\sqrt{15}}{\log 2}$. Therefore, in both cases, the subexponential norm is bounded by an absolute constant that will be denoted by K_Λ in the following. Now, let us recall a Bernstein-type concentration inequality for sum of subexponential variables.

Lemma 9 (Proposition 5.16 [73]) *Let X_1, \dots, X_m be independent centered sub-exponential random variables, and $K = \max_i \|X_i\|_{\Psi_1}$. Then for every $\gamma = (\gamma_1, \dots, \gamma_m) \in \mathbb{R}^m$ and every $t \geq 0$, we have*

$$\mathbb{P}\left(\left|\sum_{i=1}^m \gamma_i X_i\right| \geq t\right) \leq 2 \exp\left(-c \min\left(\frac{t^2}{K^2 \|\gamma\|_2^2}, \frac{t}{K \|\gamma\|_\infty}\right)\right),$$

where $c = \frac{1}{8e^2}$. (The value of c can be tracked through the proofs of Lemma 5.15 and Proposition 5.16 in [73].)

Taking $X_i = |\mathbf{a}_i^\top \mathbf{U} \mathbf{a}_i| - \mathbb{E}[|\mathbf{a}_i^\top \mathbf{U} \mathbf{a}_i|] = |\mathbf{a}_i^\top \mathbf{U} \mathbf{a}_i| - 1$ (or $X_i = |\mathbf{u}_i^\top \mathbf{U} \mathbf{u}_i| - 1$ in the uniform case), and $\gamma_i = 1/m$ for all i , yields

$$\mathbb{P}(|\|\mathcal{A}(\mathbf{U})\|_1 - 1| > t) \leq 2 \exp\left(-\frac{m}{8e^2} \min\left(\frac{t^2}{K_\Lambda^2}, \frac{t}{K_\Lambda}\right)\right), \quad \forall t \geq 0.$$

□

A.3.2. Proof of Proposition 4

Proof The lower bound is a direct consequence of Proposition 5. Let us prove the upper bound. In the Gaussian case, the following inequalities hold

$$\mathbb{E}\left[|\mathbf{a}^\top \mathbf{M} \mathbf{a}|\right] = \mathbb{E}\left[\left|\sum_{k=1}^d \lambda_k b_k^2\right|\right] \leq \sum_{k=1}^d |\lambda_k| \mathbb{E}[b_k^2] = \frac{1}{d} \sum_{k=1}^d |\lambda_k| \leq \frac{\sqrt{d}}{d} \sqrt{\sum_{k=1}^d \lambda_k^2} = \frac{1}{\sqrt{d}} \|\mathbf{M}\|_{\text{Fro}},$$

where $\mathbf{a}, \mathbf{b} \sim \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$. Similarly, in the URO case,

$$\mathbb{E}\left[|\mathbf{u}^\top \mathbf{M} \mathbf{u}|\right] \leq \sum_{k=1}^d |\lambda_k| \mathbb{E}[v_k^2] = \frac{1}{d} \sum_{k=1}^d |\lambda_k| \leq \frac{1}{\sqrt{d}} \|\mathbf{M}\|_{\text{Fro}},$$

where $\mathbf{u}, \mathbf{v} \sim \mathcal{U}(\mathbb{S}^{d-1})$. □

A.3.3. Proof of Theorem 2.2

Proof From Theorem 2.1 and Remark 3, we know that the probability that the operator \mathcal{A} does not satisfy the RIP_δ is upper-bounded by

$$\mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_\Lambda, \varepsilon) C_2\left(\frac{\delta}{2}\right) + \mathcal{N}(B_\Lambda, \|\cdot\|_\Lambda, \varepsilon') C_2\left(\frac{\delta}{2\varepsilon'} - 1\right), \quad \forall \varepsilon, \varepsilon' > 0, \quad (\text{A.44})$$

where $B_\Lambda = \{\mathbf{U} \in S_d(\mathbb{R}) : \|\mathbf{U}\|_\Lambda = 1\}$ and $C_2(t) = 2 \exp(-\frac{m}{8e^2} \min(t/K_\Lambda, (t/K_\Lambda)^2))$ (see Proposition 3). In the following, we choose $\varepsilon' = 1/2$. Given ρ , the strategy is to find an ε small enough so that the right handside term is smaller than $\rho/2$. Then, a condition on m will be derived to ensure that the left handside term is also smaller than $\rho/2$.

First of all, notice that, from standard covering argument, $\mathcal{N}(B_\Lambda, \|\cdot\|_\Lambda, 1/2) \leq (3/(1/2))^{d(d+1)/2}$. Therefore, looking at the logarithm of the second term in (A.44), ε should verify

$$\frac{d(d+1)}{2} \log(6) + \log 2 - \frac{m}{8e^2} \min \left[\frac{1}{K_\Lambda} \left(\frac{\delta}{4\varepsilon} - 1 \right), \frac{1}{K_\Lambda^2} \left(\frac{\delta}{4\varepsilon} - 1 \right)^2 \right] \leq \log(\rho/2). \quad (\text{A.45})$$

Assuming that $\varepsilon \leq \frac{\delta}{4(K_\Lambda+1)}$ to ensure that the minimum in the above expression is $\frac{1}{K_\Lambda} \left(\frac{\delta}{4\varepsilon} - 1 \right)$, (A.45) is equivalent to

$$\varepsilon \leq \frac{\delta}{4} \left[\frac{8e^2 K_\Lambda}{m} \left(\frac{d(d+1)}{2} \log(6) + \log 2 + \log(2/\rho) \right) + 1 \right]^{-1}.$$

In order to remove the dependency in m and to simplify the expression, notice that it is sufficient to take

$$\varepsilon \leq \frac{\delta}{32e^2 K_\Lambda} [d^2 \log(6) + 4 \log(2/\rho)]^{-1} \triangleq \varepsilon_0.$$

In the following, we take $\varepsilon = \varepsilon_0$. Remark that it satisfies the assumption below (A.45). In particular, note that $\varepsilon \leq 1$. We now focus on the first term in (A.44). Notice that from Corollary 1 and Proposition 4 giving $c_{\text{Fro}} = 2/(9\sqrt{15}d)$ and $C_{\text{Fro}} = 1/\sqrt{d}$, the covering number is controlled as follows:

$$\begin{aligned} & \log \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_\Lambda, \varepsilon) \\ & \leq 4k \log\left(\frac{ed^2}{2k}\right) + \log \left[\left(c_0 \frac{\sqrt{d} C_{\text{Fro}}^2 b^5}{\varepsilon^2 c_{\text{Fro}}^2 a^5} \right)^{2(d+k)} + \left(c_1 \frac{\sqrt{d} C_{\text{Fro}}^2}{\varepsilon^2 c_{\text{Fro}}} \left(\frac{2}{c_{\text{Fro}}} + 1 \right) \frac{b^5}{a^5} \right)^{d+2k+1} \right] \\ & \leq 4k \log\left(\frac{ed^2}{2k}\right) + \log \left[\left(\frac{c'_0 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{2(d+k)} + \left(\frac{c'_1 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{d+2k+1} \right]. \end{aligned}$$

where c_0, c_1, c'_0, c'_1 are absolute constants greater than 1. As $\frac{c'_1 d^{3/2} b^5}{\varepsilon^2 a^5} \geq 1$, $\left(\frac{c'_0 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{2(d+k)} + \left(\frac{c'_1 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{d+2k+1} \leq \left(\frac{c'_0 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{2(d+k)} + \left(\frac{c'_1 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{2(d+k)} \leq \left(\frac{c'_0 d^{3/2} b^5}{\varepsilon^2 a^5} + \frac{c'_1 d^{3/2} b^5}{\varepsilon^2 a^5} \right)^{2(d+k)}$. This gives

$$\begin{aligned} \log \mathcal{N}(S[\mathfrak{S}_{k,a,b}], \|\cdot\|_\Lambda, \varepsilon) & \leq 4k \log\left(\frac{ed^2}{2k}\right) + 2(d+k) \log \left[(c'_0 + c'_1) \frac{d^{3/2} b^5}{\varepsilon^2 a^5} \right] \\ & = 4k \log\left(\frac{ed^2}{2k}\right) + 4(d+k) \log \left[c \frac{b}{a} \frac{d}{\varepsilon} \right], \end{aligned}$$

where

$$c \frac{b}{a} \triangleq \sqrt{(c'_0 + c'_1) \frac{b^{5/2}}{a^{5/2}}} \text{ only depends on } b/a.$$

Therefore, m needs to verify

$$4k \log\left(\frac{ed^2}{2k}\right) + 4(d+k) \log \left(c \frac{b}{a} \frac{d}{\varepsilon} \right) + \log 2 - \frac{m\delta^2}{32e^2 K_\Lambda^2} \leq \log(\rho/2),$$

which is equivalent to

$$\frac{m\delta^2}{32e^2K_\Lambda^2} \geq 4k \log\left(\frac{ed^2}{2k}\right) + 4(d+k) \log\left(c \frac{b}{a} \frac{d}{\varepsilon}\right) + \log 2 + \log(2/\rho).$$

To simplify the expression, we derive a sufficient condition on m given by

$$m \geq \frac{32e^2K_\Lambda^2}{\delta^2} \left[4k \log\left(\frac{ed^2}{2k}\right) + 4(d+k) \log\left(\frac{32e^2c_bK_\Lambda}{\delta} d [d^2 \log(6) + 4 \log(2/\rho)]\right) + 2 \log(2/\rho) \right]. \quad (\text{A.46})$$

This finishes the proof as we can find a constant $C = C(\delta, \rho, b/a)$ such that $m \geq C(d+2k) \log d$ implies (A.46). \square

A.4. Connection with Bregman proximal gradient

The iterations of our algorithm (3.6) can be related to the iterations of Bregman Proximal Gradient (BPG). Originally introduced in [3], BPG is a generalization of the classical proximal gradient method in which the proximal operator is replaced with a Bregman proximal operator. It aims at solving problems of the form $\min f + g$ where f, g are proper, convex and lower semi-continuous. For $\lambda > 0$, we consider the optimization problem

$$\min_{\Theta \succ 0} F(\Theta) + \lambda \|\Theta\|_{1,\text{off}} \text{ where } F(\Theta) \triangleq \frac{1}{2} \|\mathcal{A}(\Theta^{-1}) - \mathbf{s}\|_2^2 = f(\Theta^{-1}). \quad (\text{A.47})$$

Interestingly the function F is convex on the convex set $S := \{\Theta \succ 0 : \lambda_{\max}(\Theta) \leq 1/\lambda_{\max}(\widehat{\Sigma})\} \subset S_d^{++}$. Indeed as shown in [10, Example 3.4] the function $(\mathbf{a}, \Theta) \rightarrow \mathbf{a}^\top \Theta^{-1} \mathbf{a}$ is jointly convex on $\mathbb{R}^d \times S_d^{++}(\mathbb{R})$. Consequently $\Theta \rightarrow (\mathbf{a}_j^\top \Theta^{-1} \mathbf{a}_j - s_j)^2 = (\mathbf{a}_j^\top (\Theta^{-1} - \widehat{\Sigma}) \mathbf{a}_j)^2$ is convex on S as the composition of a convex function and $t \rightarrow t^2$ on \mathbb{R}_+ since $\mathbf{a}_j^\top (\Theta^{-1} - \widehat{\Sigma}) \mathbf{a}_j \geq 0$ on S . Hence $F(\Theta) = \frac{1}{2} \sum_{j=1}^m (\mathbf{a}_j^\top \Theta^{-1} \mathbf{a}_j - s_j)^2$ is also convex on S as the sum of convex functions. For a fixed step-size $\gamma > 0$ the BPG iterations with the Bregman divergence (3.4) are given by

$$\Theta_{t+1} = \underset{\Theta \succ 0}{\text{argmin}} \langle \gamma \nabla F(\Theta_t), \Theta \rangle + D_h(\Theta | \Theta_t) + \lambda \gamma \|\Theta\|_{1,\text{off}}. \quad (\text{A.48})$$

By expressing the divergence D_h and using that $\nabla F(\Theta_t) = -\Theta_t^{-1} \nabla f(\Theta_t^{-1}) \Theta_t^{-1}$ these iterations are equivalent to

$$\Theta_{t+1} = \underset{\Theta \succ 0}{\text{argmin}} \langle \Theta_t^{-1} - \gamma \Theta_t^{-1} \nabla f(\Theta_t^{-1}) \Theta_t^{-1}, \Theta \rangle - \log \det \Theta + \lambda \gamma \|\Theta\|_{1,\text{off}}. \quad (\text{A.49})$$

These iterations also correspond to a graphical lasso since (A.49) rewrites as $\Theta_{t+1}^{-1} = \text{GLASSO}_{\lambda\gamma}[\Theta_t^{-1} - \gamma \Theta_t^{-1} \nabla f(\Theta_t^{-1}) \Theta_t^{-1}]$. With the change of variable $\Sigma_t = \Theta_t^{-1}$, the BPG iterations for solving (A.47) equivalently write

$$\Sigma_{t+1} = \text{GLASSO}_{\lambda\gamma}[\Sigma_t - \gamma \Sigma_t \nabla f(\Sigma_t) \Sigma_t]. \quad (\text{A.50})$$

We can notice that these iterations are very similar to the one in (3.6) but with $\Sigma_t \nabla f(\Sigma_t) \Sigma_t$ instead of $\nabla f(\Sigma_t)$. In fact, the iterations (3.6) are equivalent to the BPG iterations (A.50) when using a Riemannian gradient instead of a Euclidean one. More precisely, when considering for $\mathbf{X} \succ 0$, the inner product

$\langle \mathbf{U}, \mathbf{V} \rangle_{\mathbf{X}} \triangleq \text{tr}(\mathbf{U}\mathbf{X}^{-1}\mathbf{V}\mathbf{X}^{-1})$ and computing the gradient w.r.t. $\langle \cdot, \cdot \rangle_{\mathbf{X}}$ at \mathbf{X} we get the formula $\text{grad} f(\mathbf{X}) = \mathbf{X}\nabla f(\mathbf{X})\mathbf{X}$ [39]. This corresponds to endowing the space $\mathcal{S}_d^{++}(\mathbb{R})$ with the affine-invariant geometry [8, Section 11.7]. In conclusion, if we consider our iterations (3.6) with the Riemannian gradient $\text{grad} f$ instead of the Euclidean gradient we get the BPG iterations (A.50). However, we observe in practice that the algorithm with the BPG has degraded performance compared to the one proposed in (3.6).

A.5. Safe step-size strategy

The goal of this section is to provide a step-size $\gamma > 0$ ensuring that the matrix $\Sigma_{t+\frac{1}{2}} := \Sigma_t - \gamma \nabla f(\Sigma_t)$ remains positive definite during the iterations. Recall that $\nabla f(\Sigma_t) = \mathcal{A}^*(\mathcal{A}(\Sigma_t) - \mathbf{s})$ where $\mathbf{s} = \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) = \frac{1}{m} (\mathbf{a}_1^\top \widehat{\Sigma} \mathbf{a}_1, \dots, \mathbf{a}_m^\top \widehat{\Sigma} \mathbf{a}_m)$ is the sketch of the data and $\widehat{\Sigma}$ is the empirical covariance matrix. The adjoint operator \mathcal{A}^* is given by $\mathbf{y} \rightarrow \mathcal{A}^*(\mathbf{y}) = \frac{1}{m} \sum_{j=1}^m y_j \mathbf{a}_j \mathbf{a}_j^\top$.

The matrix $\Sigma_{t+\frac{1}{2}}$ is positive definite when $\lambda_{\min}(\Sigma_t - \gamma \nabla f(\Sigma_t)) > 0$ that is when

$$\lambda_{\min} \left(\Sigma_t - \gamma \frac{1}{m} \sum_{j=1}^m \left(\frac{1}{m} \mathbf{a}_j^\top \Sigma_t \mathbf{a}_j - s_j \right) \mathbf{a}_j \mathbf{a}_j^\top \right) > 0. \quad (\text{A.51})$$

Moreover we have,

$$\begin{aligned} \lambda_{\min} \left(\Sigma_t - \gamma \frac{1}{m} \sum_{j=1}^m \left(\frac{1}{m} \mathbf{a}_j^\top \Sigma_t \mathbf{a}_j - s_j \right) \mathbf{a}_j \mathbf{a}_j^\top \right) &\geq \lambda_{\min}(\Sigma_t) - \gamma \frac{1}{m} \lambda_{\max} \left(\sum_{j=1}^m \left(\frac{1}{m} \mathbf{a}_j^\top \Sigma_t \mathbf{a}_j - s_j \right) \mathbf{a}_j \mathbf{a}_j^\top \right) \\ &= \lambda_{\min}(\Sigma_t) - \gamma \frac{1}{m^2} \lambda_{\max} \left(\sum_{j=1}^m [\mathbf{a}_j^\top (\Sigma_t - \widehat{\Sigma}) \mathbf{a}_j] \mathbf{a}_j \mathbf{a}_j^\top \right). \end{aligned} \quad (\text{A.52})$$

Using $\forall j \in \llbracket m \rrbracket, \mathbf{a}_j(\Sigma_t - \widehat{\Sigma})\mathbf{a}_j \leq \lambda_{\max}(\Sigma_t - \widehat{\Sigma}) \|\mathbf{a}_j\|_2^2$ we have, for any $\mathbf{z} \in \mathbb{R}^d, \|\mathbf{z}\|_2 = 1$,

$$\mathbf{z}^\top \left(\sum_{j=1}^m [\mathbf{a}_j^\top (\Sigma_t - \widehat{\Sigma}) \mathbf{a}_j] \mathbf{a}_j \mathbf{a}_j^\top \right) \mathbf{z} = \sum_{j=1}^m [\mathbf{a}_j^\top (\Sigma_t - \widehat{\Sigma}) \mathbf{a}_j] |\mathbf{z}^\top \mathbf{a}_j|^2 \leq \lambda_{\max}(\Sigma_t - \widehat{\Sigma}) \sum_{j=1}^m \|\mathbf{a}_j\|_2^2 |\mathbf{z}^\top \mathbf{a}_j|^2. \quad (\text{A.53})$$

By introducing the matrix $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_m) \in \mathbb{R}^{d \times m}$ the previous inequality leads to

$$\begin{aligned} \mathbf{z}^\top \left(\sum_{j=1}^m [\mathbf{a}_j^\top (\Sigma_t - \widehat{\Sigma}) \mathbf{a}_j] \mathbf{a}_j \mathbf{a}_j^\top \right) \mathbf{z} &\leq \lambda_{\max}(\Sigma_t - \widehat{\Sigma}) \max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \sum_{j=1}^m |\mathbf{z}^\top \mathbf{a}_j|^2 \\ &= \lambda_{\max}(\Sigma_t - \widehat{\Sigma}) \left(\max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \right) \mathbf{z}^\top \mathbf{A} \mathbf{A}^\top \mathbf{z} \\ &\leq \lambda_{\max}(\Sigma_t - \widehat{\Sigma}) \left(\max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \right) \lambda_{\max}(\mathbf{A} \mathbf{A}^\top) \\ &\leq [\lambda_{\max}(\Sigma_t) - \lambda_{\min}(\widehat{\Sigma})] \left(\max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \right) \lambda_{\max}(\mathbf{A} \mathbf{A}^\top). \end{aligned} \quad (\text{A.54})$$

Consequently,

$$\lambda_{\min} \left(\sum_{j=1}^m [\mathbf{a}_j^\top (\Sigma_t - \widehat{\Sigma}) \mathbf{a}_j] \mathbf{a}_j \mathbf{a}_j^\top \right) \leq [\lambda_{\max}(\Sigma_t) - \lambda_{\min}(\widehat{\Sigma})] \left(\max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \right) \lambda_{\max}(\mathbf{A} \mathbf{A}^\top). \quad (\text{A.55})$$

This shows that if $\lambda_{\max}(\Sigma_t) < \lambda_{\min}(\widehat{\Sigma})$ then the condition (A.51) is valid for any $\gamma > 0$ since in this case $\lambda_{\max}\left(\sum_{j=1}^m [\mathbf{a}_j^\top (\Sigma_t - \widehat{\Sigma}) \mathbf{a}_j] \mathbf{a}_j \mathbf{a}_j^\top\right) < 0$. On the other hand if $\lambda_{\max}(\Sigma_t) > \lambda_{\min}(\widehat{\Sigma})$ then, by (A.52) and (A.55), a step-size

$$\gamma < \frac{m^2}{\max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \sigma_{\max}^2(\mathbf{A})} \times \frac{\lambda_{\min}(\Sigma_t)}{\lambda_{\max}(\Sigma_t) - \lambda_{\min}(\widehat{\Sigma})}, \quad (\text{A.56})$$

where $\sigma_{\max}(\mathbf{A})$ is the largest singular value of \mathbf{A} , ensures that $\Sigma_{k+\frac{1}{2}}$ is positive definite. Overall a safe step-size strategy is given by

$$\gamma \in \begin{cases} (0, +\infty[& \text{if } \lambda_{\max}(\Sigma_t) < \lambda_{\min}(\widehat{\Sigma}), \\ (0, \frac{m^2}{\max_{j \in \llbracket m \rrbracket} \|\mathbf{a}_j\|_2^2 \sigma_{\max}^2(\mathbf{A})} \frac{\lambda_{\min}(\Sigma_t)}{\lambda_{\max}(\Sigma_t) - \lambda_{\min}(\widehat{\Sigma})}) & \text{if } \lambda_{\max}(\Sigma_t) > \lambda_{\min}(\widehat{\Sigma}). \end{cases} \quad (\text{A.57})$$

We emphasize that this strategy is quite conservative, and requires the computation of both the maximum and minimum eigenvalues of Σ_t at every iteration. In practical scenarios, we find that searching for γ in $\{1e-3, 1e-2, 1e-1\}$ is adequate for achieving convergence in our experimental setups.

A.6. Covering number of a model set with a condition number constraint.

We would like to consider a model set of covariance matrices that does not restrict their spectra but rather their condition numbers. For a square matrix \mathbf{M} , the condition number is defined as $\kappa(\mathbf{M}) \triangleq \|\mathbf{M}\|_{2 \rightarrow 2} \|\mathbf{M}^{-1}\|_{2 \rightarrow 2}$. In the special case of positive-definite matrices, it is the ratio between the largest and smallest eigenvalues. For some $\kappa_0 \geq 1$, we consider the following model set

$$\mathfrak{S}_{k, \kappa_0} \triangleq \left\{ \Sigma \in S_d^{++}(\mathbb{R}) : \Theta = \Sigma^{-1} \succ 0, \|\Theta\|_0 \leq d + 2k, \kappa(\Theta) \leq \kappa_0 \right\}.$$

Remark that the above definition implies that $\mathfrak{S}_{k, a, b} \subset \mathfrak{S}_{k, \kappa_0}$, for all $a, b > 0$ such that $b/a \leq \kappa_0$. In particular, $\mathfrak{S}_{k, 1/\kappa_0, 1} \subset \mathfrak{S}_{k, \kappa_0}$. However, $\mathfrak{S}_{k, \kappa_0}$ is a much ‘‘bigger’’ set than sets like $\mathfrak{S}_{k, a, b}$, as the latter are bounded sets while the former is not. Interestingly enough, we are able to upper-bound the covering number of the normalized secant of $\mathfrak{S}_{k, \kappa_0}$ that is

$$S[\mathfrak{S}_{k, \kappa_0}] = \left\{ \frac{\Theta_1^{-1} - \Theta_2^{-1}}{\|\Theta_1^{-1} - \Theta_2^{-1}\|_\Lambda} : (\Theta_1^{-1}, \Theta_2^{-1}) \in \mathfrak{S}_{k, \kappa_0}^2, \|\Theta_1^{-1} - \Theta_2^{-1}\|_\Lambda > 0 \right\}.$$

The key element to obtain this bound is that we are able to rewrite $S[\mathfrak{S}_{k, \kappa_0}]$ in term of matrices that are in

$$\mathfrak{S}_0 \triangleq \left\{ \Sigma \in S_d^{++}(\mathbb{R}) : \Theta = \Sigma^{-1} \succ 0, \|\Theta\|_0 \leq d + 2k, \text{spec}(\Theta) \subset [1/\kappa_0, 1], \|\Theta\|_{2 \rightarrow 2} = 1 \right\}.$$

Indeed, for an element $\frac{\tilde{\Theta}_1^{-1} - \tilde{\Theta}^{-1}}{\|\tilde{\Theta}_1^{-1} - \tilde{\Theta}^{-1}\|_\Lambda} \in S[\mathfrak{S}_{k, \kappa_0}]$, notice that for any $\lambda > 0$ the matrix $\frac{\lambda \tilde{\Theta}_1^{-1} - \lambda \tilde{\Theta}^{-1}}{\|\lambda \tilde{\Theta}_1^{-1} - \lambda \tilde{\Theta}^{-1}\|_\Lambda}$ still is on $S[\mathfrak{S}_{k, \kappa_0}]$. This implies that we can normalize the matrices involved in the secant. More precisely, by choosing $\lambda = \|\tilde{\Theta}_1\|_{2 \rightarrow 2}$, setting $\Theta_1 = \lambda^{-1} \tilde{\Theta}_1$, $\Theta_2 = \tilde{\Theta}_2 / \|\tilde{\Theta}_2\|_{2 \rightarrow 2}$ and $\tau = \|\tilde{\Theta}_1\|_{2 \rightarrow 2} / \|\tilde{\Theta}_2\|_{2 \rightarrow 2}$, we

have $\Theta_1, \Theta_2 \in \mathfrak{S}_0$, $\tau > 0$ and $\frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} \in S[\mathfrak{S}_{k, \kappa_0}]$. This shows that

$$S[\mathfrak{S}_{k, \kappa_0}] = \left\{ \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} : \tau > 0, (\Theta_1^{-1}, \Theta_2^{-1}) \in \mathfrak{S}_0^2, \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda > 0 \right\}.$$

Now, up to exchanging the role of $\tilde{\Theta}_1$ and $\tilde{\Theta}_2$, which result in changing the sign of $\frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda}$, we can always assume that $\|\tilde{\Theta}_1\|_{2 \rightarrow 2} \leq \|\tilde{\Theta}_2\|_{2 \rightarrow 2}$, meaning that $0 < \tau \leq 1$, which yields

$$\begin{aligned} S[\mathfrak{S}_{k, \kappa_0}] &= \left\{ \pm \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} : \tau \in (0, 1], (\Theta_1^{-1}, \Theta_2^{-1}) \in \mathfrak{S}_0^2 : \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda > 0 \right\} \\ &= \bar{S}[\mathfrak{S}_{k, \kappa_0}] \cup (-\bar{S}[\mathfrak{S}_{k, \kappa_0}]), \end{aligned}$$

where

$$\bar{S}[\mathfrak{S}_{k, \kappa_0}] \triangleq \left\{ \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} : \tau \in (0, 1], (\Theta_1^{-1}, \Theta_2^{-1}) \in \mathfrak{S}_0^2, \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda > 0 \right\}.$$

Remark that for any $\varepsilon > 0$, $\mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \leq 2\mathcal{N}(\bar{S}[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, \varepsilon)$, therefore we only have to control the covering number of $\bar{S}[\mathfrak{S}_{k, \kappa_0}]$. To do so, we follow the same line of proof as in the control of $S[\mathfrak{S}_{k, a, b}]$ by splitting our set of interests into long and short chords. The analysis of these chords is similar, although more technical and more computation-heavy. A slight complication in this new setting is the need to ensure that τ is bounded away from zero in the case of short chords. In the following, we briefly detail the analysis of the long and short chords given for some $\eta > 0$ by

$$\begin{aligned} \bar{S}_\eta^+[\mathfrak{S}_{k, \kappa_0}] &\triangleq \left\{ \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} : \tau \in (0, 1], (\Theta_1^{-1}, \Theta_2^{-1}) \in \mathfrak{S}_0^2, \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda > \eta \right\}, \\ \bar{S}_\eta^-[\mathfrak{S}_{k, \kappa_0}] &\triangleq \left\{ \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} : \tau \in (0, 1], (\Theta_1^{-1}, \Theta_2^{-1}) \in \mathfrak{S}_0^2, 0 < \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda \leq \eta \right\}. \end{aligned}$$

A.6.1. Control of the long chords

The strategy to control the covering number of long chords is to express $\bar{S}_\eta^+[\mathfrak{S}_{k, \kappa_0}]$ as the image of a Lipschitz-continuous function and control the covering number of the original set. Let us consider the set $\bar{\mathfrak{X}}_\eta \triangleq \{(\tau, \Theta_1, \Theta_2) \in (0, 1] \times \mathfrak{S}_0^{-1} \times \mathfrak{S}_0^{-1}, \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda > \eta\}$ equipped with the norm $\|(\tau, \mathbf{M}_1, \mathbf{M}_2)\|_\otimes = |\tau| + \|\mathbf{M}_1\|_{\text{Fro}} + \|\mathbf{M}_2\|_{\text{Fro}}$. Then we have the following lemma. For the sake of conciseness, its proof is not provided, but it is based on the one of Proposition 1 presented in Appendix A.2.3.

Lemma 10 *Let $g : (\bar{\mathfrak{X}}_\eta, \|\cdot\|_\otimes) \rightarrow (\bar{S}_\eta^+[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda)$ be the function defined by*

$$g(\tau, \Theta_1, \Theta_2) \triangleq \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda}.$$

Then, g is surjective and L_0/η -lipschitz continuous with $L_0 = 2C_{\text{Fro}}\kappa_0^2\sqrt{d}$.

As a consequence, we can control the covering number of $\bar{S}_\eta^+[\mathfrak{S}_{k,\kappa_0}]$ using the one of $\bar{\mathfrak{X}}_\eta$ which is easier to handle. This yields the following proposition which proof is also based on the one of Proposition 1.

Proposition 7 *For all $\varepsilon > 0$ and $\eta > 0$, we have*

$$\mathcal{N}(\bar{S}_\eta^+[\mathfrak{S}_{k,\kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \leq \mathcal{N}((0, 1], |\cdot|, \frac{\eta\varepsilon}{6L_0}) \times \mathcal{N}(\mathfrak{S}_0^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\eta\varepsilon}{6L_0})^2.$$

Note that $\mathcal{N}((0, 1], |\cdot|, \varepsilon)$ is bounded by ε^{-1} and the control of the covering of \mathfrak{S}_0^{-1} will be provided later by Lemma 13.

A.6.2. Control of the short chords

In order to control the covering number of the short chords we follow these two steps: 1) we show that any element of $\bar{S}_\eta^-[\mathfrak{S}_{k,\kappa_0}]$ is close to an element of a certain ‘‘tangent space’’ 2) we will control the covering number of this space.

Assumption: In all of this section we will assume that $0 < \eta \leq c_{\text{Fro}}/2$. This requirement will be useful for various simplification and will be met when we calibrate η for a good balance between the covering numbers of both long and short chords.

Point 1) is done through the following lemma which proof can be found in Appendix A.7.1.

Lemma 11 *For $\eta > 0$, consider the set of short chords $\bar{S}_\eta^-[\mathfrak{S}_{k,\kappa_0}]$ and the normalized secant set $S[\mathfrak{S}_{k,\kappa_0}^{-1}]$. Define $C \triangleq \{\lambda \mathbf{V} : \lambda \in]0, \lambda_0], \mathbf{V} \in S[\mathfrak{S}_{k,\kappa_0}^{-1}]\}$ and*

$$T_C \triangleq \left\{ \text{Dinv}_{\bar{\Theta}}(\mathbf{C}) : (\bar{\Theta}, \mathbf{C}) \in \mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1} \times C \right\},$$

with $\lambda_0 \triangleq \frac{2}{c_{\text{Fro}}}$. Defining $Z_0 \triangleq \frac{C_{\text{Fro}}\kappa_0^3}{c_{\text{Fro}}^2}$, we have

$$\forall \mathbf{U} \in \bar{S}_\eta^-[\mathfrak{S}_{k,\kappa_0}], \exists \mathbf{T} \in \mathcal{T}, \|\mathbf{U} - \mathbf{T}\|_\Lambda \leq Z_0 \eta.$$

As T_C is a good approximation of $\bar{S}_\eta^-[\mathfrak{S}_{k,\kappa_0}]$, we can bound the covering number of the latter by the covering of the former (with a different scale), see Lemma 3. Hence, we need to control the covering number of T_C .

Lemma 12 *For all $\varepsilon > 0$, we have*

$$\mathcal{N}(T_C, \|\cdot\|_\Lambda, \varepsilon) \leq \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}) \times \mathcal{N}(C, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}), \quad (\text{A.58})$$

with $C_0 = C_{\text{Fro}}(2\kappa_0^3\lambda_0 + \kappa_0^2)$.

Now, combining the above results, we are able to provide a control of the covering number of the short chords. See Appendix A.7.1 for the proof.

Proposition 8 (Similar to Proposition 2) *For any $\varepsilon > 0$ and $\eta > 0$, we have*

$$\mathcal{N}(\bar{S}_\eta^-[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, 2(\varepsilon + Z_0\eta)) \leq 2\lambda_0 C_0 \varepsilon^{-1} \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}) \times \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{2\lambda_0 C_0}),$$

$$\text{where } Z_0 = \frac{c_{\text{Fro}} \kappa_0^3}{c_{\text{Fro}}^2}, \lambda_0 = \frac{2}{c_{\text{Fro}}}, C_0 = c_{\text{Fro}}(2\kappa_0^3 \lambda_0 + \kappa_0^2).$$

A.6.3. Combining the results

To finish this section and obtain the covering of $\bar{S}_\eta^-[\mathfrak{S}_{k, \kappa_0}]$, we need the control of the covering of \mathfrak{S}_0^{-1} , $\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}$ and $S[\mathfrak{S}_{k, \kappa_0}^{-1}]$ as they appear in the control of the covering number of the long and short chords.

This is done in the following lemma which proof can be found in Appendix A.7.2.

Lemma 13 *For any $\varepsilon > 0$, and $0 < a \leq b$, we have*

$$\begin{aligned} \mathcal{N}(\mathfrak{S}_0^{-1}, \|\cdot\|_{\text{Fro}}, \varepsilon) &\leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{18\sqrt{d}}{\varepsilon}\right)^{d+k}, \\ \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \varepsilon) &\leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{2 \times 18\sqrt{d}}{\varepsilon}\right)^{d+k}, \\ \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \varepsilon) &\leq \left(\frac{ed^2}{4k}\right)^{2k} \left(\frac{18}{\varepsilon}\right)^{d+2k}. \end{aligned}$$

Gathering up all the pieces, we obtain the following theorem.

There exist absolute constants \tilde{c}_1 and \tilde{c}_2 such that for any ε such that $0 < \varepsilon \leq \frac{2\kappa_0^3}{c_{\text{Fro}}} \sqrt{d}$, we have

$$\mathcal{N}(S_\eta[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \leq 2 \left(\frac{ed^2}{2k}\right)^{3k} \left[\left(\frac{\tilde{c}_1 \kappa_0^5 d^2}{\varepsilon^2}\right)^{2(d+k)+1} + \left(\frac{\tilde{c}_2 \kappa_0^3 d^2}{\varepsilon}\right)^{2d+3k+1} \right].$$

See Appendix A.7.3 for the proof.

A.7. Proof for the coverings with a condition number hypothesis

Before diving into the control of the covering numbers of the long and short chords, let us claim various inequalities related to the inverse function on matrices that will be useful in the following.

Lemma 14 (Inverse function properties) *Assume that there exist constants c_{Fro} and C_{Fro} such that $c_{\text{Fro}} \|\mathbf{M}\|_{\text{Fro}} \leq \|\mathbf{M}\|_\Lambda \leq C_{\text{Fro}} \|\mathbf{M}\|_{\text{Fro}}$, for all $\mathbf{M} \in S_d$. Let \mathbf{M}_1 and \mathbf{M}_2 be two matrices in S_d^{++} . Then we*

have the following inequalities:

$$\|\mathbf{M}_1^{-1} - \mathbf{M}_2^{-1}\|_{\Lambda} \leq C_{\text{Fro}} \|\mathbf{M}_1^{-1}\|_{2 \rightarrow 2} \|\mathbf{M}_2^{-1}\|_{2 \rightarrow 2} \|\mathbf{M}_1 - \mathbf{M}_2\|_{\text{Fro}}, \quad (\text{A.59})$$

$$\|\mathbf{M}_1 - \mathbf{M}_2\|_{\text{Fro}} \leq \frac{1}{C_{\text{Fro}}} \|\mathbf{M}_1\|_{2 \rightarrow 2} \|\mathbf{M}_2\|_{2 \rightarrow 2} \|\mathbf{M}_1^{-1} - \mathbf{M}_2^{-1}\|_{\Lambda}, \quad (\text{A.60})$$

$$\|\mathbf{M}_1^{-1} - \mathbf{M}_2^{-1} - \text{D inv}_{\mathbf{M}_2}(\mathbf{M}_1 - \mathbf{M}_2)\|_{\Lambda} \leq C_{\text{Fro}} \|\mathbf{M}_1^{-1}\|_{2 \rightarrow 2} \|\mathbf{M}_2^{-1}\|_{2 \rightarrow 2}^2 \|\mathbf{M}_1 - \mathbf{M}_2\|_{\text{Fro}}^2, \quad (\text{A.61})$$

$$\begin{aligned} \|\text{D inv}_{\mathbf{M}_1}(\mathbf{M}) - \text{D inv}_{\mathbf{M}_2}(\mathbf{M})\|_{\text{op}} &\leq C_{\text{Fro}} (\|\mathbf{M}_1^{-1}\|_{2 \rightarrow 2} + \|\mathbf{M}_2^{-1}\|_{2 \rightarrow 2}) \\ &\quad \times \|\mathbf{M}_1^{-1}\|_{2 \rightarrow 2} \|\mathbf{M}_2^{-1}\|_{2 \rightarrow 2} \|\mathbf{M}_1 - \mathbf{M}_2\|_{\text{Fro}}. \end{aligned} \quad (\text{A.62})$$

Idea of the proof It follows the ideas of the proof of Lemma 1. \square

A.7.1. Control of the short chords

We begin with the following preliminary result that ensures that τ can not be too close to 0.

Lemma 15 *Assume that $(\tau, \Theta_1, \Theta_2) \in (0, 1] \times \mathfrak{S}_0^{-1} \times \mathfrak{S}_0^{-1}$ verifies $0 < \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_{\Lambda} \leq \eta$ and $\eta / \text{leq}_{C_{\text{Fro}}}/2$. Then τ is bounded away from 0 i.e.*

$$\tau \geq 1 - \frac{\eta}{C_{\text{Fro}}} \geq \frac{1}{2}.$$

Proof From the inverse-lipschitz property of the inverse function in (A.60), we have :

$$\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}} \leq \frac{1}{C_{\text{Fro}}} \|\Theta_1\|_{2 \rightarrow 2} \|\tau^{-1}\Theta_2\|_{2 \rightarrow 2} \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_{\Lambda} \leq \frac{\eta}{C_{\text{Fro}}\tau}.$$

Moreover,

$$\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}} \geq \|\Theta_1 - \tau^{-1}\Theta_2\|_{2 \rightarrow 2} \geq \tau^{-1} \|\Theta_2\|_{2 \rightarrow 2} - \|\Theta_1\|_{2 \rightarrow 2} = \frac{1}{\tau} - 1.$$

Combining the two inequalities yields $\tau \geq 1 - \frac{\eta}{C_{\text{Fro}}}$. \square

Let us now prove Lemma 11.

Proof of Lemma 11 Take $\mathbf{U} = \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_{\Lambda}} \in \overline{S}_{\eta}[\mathfrak{S}_{k, \kappa_0}]$ so we have $0 < \|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_{\Lambda} \leq \eta$. Note that by using (A.59) and (A.60), we have

$$0 \stackrel{(\text{A.59})}{<} \|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}} \stackrel{(\text{A.60})}{\leq} \frac{\eta}{C_{\text{Fro}}\tau} \|\Theta_1\|_{2 \rightarrow 2} \|\Theta_2\|_{2 \rightarrow 2} \leq \frac{\eta}{C_{\text{Fro}}\tau}. \quad (\text{A.63})$$

Now, from (A.61) with $\mathbf{M}_1 = \Theta_1$ and $\mathbf{M}_2 = \tau^{-1}\Theta_2$, we have

$$\begin{aligned} \|\Theta_1^{-1} - \tau\Theta_2^{-1} - \text{D inv}_{\tau^{-1}\Theta_2}(\Theta_1 - \tau^{-1}\Theta_2)\|_{\Lambda} &\stackrel{(\text{A.61})}{\leq} C_{\text{Fro}} \|\Theta_1^{-1}\|_{\text{Fro}} \|\tau\Theta_2^{-1}\|_{\text{Fro}}^2 \|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}^2 \\ &\leq C_{\text{Fro}} \kappa_0^3 \tau^2 \|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}^2. \end{aligned}$$

Dividing by $\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda > 0$ in the above inequality yields:

$$\begin{aligned} \left\| \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} - \text{D inv}_{\tau^{-1}\Theta_2} \frac{\Theta_1 - \tau^{-1}\Theta_2}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} \right\|_\Lambda &\leq C_{\text{Fro}} \kappa_0^3 \tau^2 \|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}} \frac{\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} \\ &\text{(applying (A.63) and (A.60))} \leq C_{\text{Fro}} \kappa_0^3 \tau^2 \frac{\eta}{C_{\text{Fro}} \tau} \frac{1}{C_{\text{Fro}}} \|\Theta_1\|_{2 \rightarrow 2} \|\tau^{-1}\Theta_2\|_{2 \rightarrow 2} \\ &= \frac{C_{\text{Fro}} \kappa_0^3}{C_{\text{Fro}}^2} \eta = Z_0 \eta. \end{aligned}$$

Moreover, as $\Theta_2 \in \mathfrak{S}_0^{-1} \subset \mathfrak{S}_{k, \frac{1}{\kappa_0}, 1}^{-1}$, setting $\tilde{\Theta} = \tau^{-1}\Theta_2$, we have $\tilde{\Theta} \in \mathfrak{S}_{k, \frac{1}{\tau\kappa_0}, \frac{1}{\tau}}^{-1} \subset \mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}$ (the inclusion comes from Lemma 15). Remark that the element in the differential can be expressed as

$$\frac{\Theta_1 - \tau^{-1}\Theta_2}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} = \frac{\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} \frac{\Theta_1 - \tau^{-1}\Theta_2}{\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}}.$$

So, if we define $\lambda = \frac{\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda}$, by (A.60) and Lemma 15 it satisfies $0 < \lambda \leq \frac{1}{C_{\text{Fro}} \tau} \leq \frac{2}{C_{\text{Fro}}} \stackrel{\Delta}{=} \lambda_0$.

Thus, there exists $\lambda \in]0, \lambda_0]$ and $\tilde{\Theta} \in \mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}$ such that:

$$\left\| \frac{\Theta_1^{-1} - \tau\Theta_2^{-1}}{\|\Theta_1^{-1} - \tau\Theta_2^{-1}\|_\Lambda} - \text{D inv}_{\tilde{\Theta}} \lambda \frac{\Theta_1 - \tau^{-1}\Theta_2}{\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}} \right\|_\Lambda \leq Z_0 \eta.$$

Now we set $\mathbf{V} = \frac{\Theta_1 - \tau^{-1}\Theta_2}{\|\Theta_1 - \tau^{-1}\Theta_2\|_{\text{Fro}}}$ and we have $\mathbf{V} \in \mathcal{S}[\mathfrak{S}_{k, \kappa_0}^{-1}]$, which finishes the proof. \square

Proof of Lemma 12 First observe that for all $\mathbf{C} \in \mathcal{C}$, we have $\|\mathbf{C}\|_{\text{Fro}} \leq \lambda_0$ since $\forall \mathbf{V} \in \mathcal{S}[\mathfrak{S}_{k, \kappa_0}^{-1}]$, $\|\mathbf{V}\|_{\text{Fro}} = 1$. Then, take $\tilde{\mathfrak{S}}_{k, \frac{1}{\kappa_0}, 2}^{-1}$ an ε -net of $\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}$ and $\bar{\mathcal{C}}$ an ε -net of \mathcal{C} . Take $\mathbf{T} = \text{D inv}_{\tilde{\Theta}}(\mathbf{C}) \in T_{\mathcal{C}}$ and consider $(\bar{\Theta}, \bar{\mathbf{C}}) \in \tilde{\mathfrak{S}}_{k, \frac{1}{\kappa_0}, 2}^{-1} \times \bar{\mathcal{C}}$ such that $\|\bar{\mathbf{C}} - \mathbf{C}\|_{\text{Fro}} \leq \varepsilon$ and $\|\bar{\Theta} - \tilde{\Theta}\|_{\text{Fro}} \leq \varepsilon$. Then with $\bar{\mathbf{T}} = \text{D inv}_{\bar{\Theta}}(\bar{\mathbf{C}}) \in T_{\bar{\mathcal{C}}}$,

$$\begin{aligned} \|\mathbf{T} - \bar{\mathbf{T}}\|_\Lambda &= \|\text{D inv}_{\tilde{\Theta}}(\mathbf{C}) - \text{D inv}_{\bar{\Theta}}(\bar{\mathbf{C}})\|_\Lambda \\ &\leq \|\text{D inv}_{\tilde{\Theta}}(\mathbf{C}) - \text{D inv}_{\bar{\Theta}}(\mathbf{C})\|_\Lambda + \|\text{D inv}_{\bar{\Theta}}(\mathbf{C}) - \text{D inv}_{\bar{\Theta}}(\bar{\mathbf{C}})\|_\Lambda \\ &\leq \|\text{D inv}_{\tilde{\Theta}} - \text{D inv}_{\bar{\Theta}}\|_{\text{op}} \|\mathbf{C}\|_{\text{Fro}} + \|\text{D inv}_{\bar{\Theta}}\|_{\text{op}} \|\mathbf{C} - \bar{\mathbf{C}}\|_{\text{Fro}}. \end{aligned}$$

According to (A.62),

$$\begin{aligned} &\|\text{D inv}_{\tilde{\Theta}} - \text{D inv}_{\bar{\Theta}}\|_{\text{op}} \\ &\leq C_{\text{Fro}} \left(\|\tilde{\Theta}^{-1}\|_{2 \rightarrow 2} + \|\bar{\Theta}^{-1}\|_{2 \rightarrow 2} \right) \|\tilde{\Theta}^{-1}\|_{2 \rightarrow 2} \|\bar{\Theta}^{-1}\|_{2 \rightarrow 2} \|\tilde{\Theta} - \bar{\Theta}\|_{\text{Fro}} \\ &\leq 2C_{\text{Fro}} \kappa_0^3 \|\tilde{\Theta} - \bar{\Theta}\|_{\text{Fro}}. \end{aligned}$$

We can also show that $\|\mathbf{D} \operatorname{inv}_{\bar{\Theta}}\|_{\text{op}} \leq C_{\text{Fro}} \kappa_0^2$. Therefore,

$$\begin{aligned} \|\mathbf{T} - \bar{\mathbf{T}}\|_{\Lambda} &\leq 2C_{\text{Fro}} \kappa_0^3 \|\bar{\Theta} - \bar{\Theta}\|_{\text{Fro}} \delta + \|\mathbf{D} \operatorname{inv}_{\bar{\Theta}}\|_{\text{op}} \varepsilon \\ &\leq 2C_{\text{Fro}} \kappa_0^3 \varepsilon \delta + C_{\text{Fro}} \kappa_0^2 \varepsilon \\ &= C_{\text{Fro}} (2\kappa_0^3 \lambda_0 + \kappa_0^2) \varepsilon. \end{aligned}$$

This gives $\mathcal{N}(T_C, \|\cdot\|_{\Lambda}, C_{\text{Fro}}(2\kappa_0^3 \lambda_0 + \kappa_0^2) \varepsilon) \leq \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \varepsilon) \times \mathcal{N}(C, \|\cdot\|_E, \varepsilon)$. Therefore, by setting $C_0 = C_{\text{Fro}}(2\kappa_0^3 \lambda_0 + \kappa_0^2)$, we have

$$\mathcal{N}(T_C, \|\cdot\|_{\Lambda}, \varepsilon) \leq \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}) \times \mathcal{N}(C, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}).$$

□

Proof of Proposition 8 Recall that from Lemma 12 we have

$$\mathcal{N}(T_C, \|\cdot\|_{\Lambda}, \varepsilon) \leq \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}) \times \mathcal{N}(C, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}),$$

with $C_0 = C_{\text{Fro}}(2\kappa_0^3 \lambda_0 + \kappa_0^2)$. Using Lemma 11, we also have the approximation of $\bar{S}_{\eta}^{-1}[\mathfrak{S}_{k, \kappa_0}]$ by the set T_C :

$$\forall \mathbf{U} \in \bar{S}_{\eta}^{-1}[\mathfrak{S}_{k, \kappa_0}], \exists \mathbf{T} \in T_C, \|\mathbf{U} - \mathbf{T}\|_{\Lambda} \leq Z_0 \eta.$$

Therefore, we can apply Lemma 3 (with $\delta = Z_0 \eta$) to prove that for any $\varepsilon > 0$:

$$\mathcal{N}(\bar{S}_{\eta}^{-1}[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_{\Lambda}, 2(\varepsilon + Z_0 \eta)) \leq \mathcal{N}(T_C, \|\cdot\|_{\Lambda}, \varepsilon). \quad (\text{A.64})$$

Thus, combining (A.58) and (A.64) yields

$$\mathcal{N}(\bar{S}_{\eta}^{-1}[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_{\Lambda}, 2(\varepsilon + Z_0 \eta)) \leq \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}) \times \mathcal{N}(C, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}).$$

All we need now is to control $\mathcal{N}(C, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0})$. Take $S[\mathfrak{S}_{k, \kappa_0}^{-1}]$ a ε -net of $S[\mathfrak{S}_{k, \kappa_0}^{-1}]$ and $(0, \lambda_0]$ an $(\lambda_0 \varepsilon)$ -net of $(0, \lambda_0]$. Take $\mathbf{C} = \lambda \mathbf{V} \in C$ with $\lambda \in (0, \lambda_0]$ and $\mathbf{V} \in S[\mathfrak{S}_{k, \kappa_0}^{-1}]$. Then there exists $\bar{\mathbf{V}} \in S[\mathfrak{S}_{k, \kappa_0}^{-1}]$ such that $\|\mathbf{V} - \bar{\mathbf{V}}\|_{\text{Fro}} \leq \varepsilon$ and $\bar{\lambda} \in (0, \lambda_0]$ such that $|\lambda - \bar{\lambda}| \leq \lambda_0 \varepsilon$. Define $\bar{\mathbf{C}} = \bar{\lambda} \bar{\mathbf{V}}$. We clearly have that $\bar{\mathbf{C}} \in C$ and

$$\begin{aligned} \|\bar{\mathbf{C}} - \mathbf{C}\|_{\text{Fro}} &= \|\bar{\lambda} \bar{\mathbf{V}} - \lambda \mathbf{V}\|_{\text{Fro}} \\ &\leq \|\bar{\lambda} \bar{\mathbf{V}} - \lambda \bar{\mathbf{V}}\|_{\text{Fro}} + \|\lambda \bar{\mathbf{V}} - \lambda \mathbf{V}\|_{\text{Fro}} \\ &\leq \|\bar{\mathbf{V}}\|_{\text{Fro}} |\lambda - \bar{\lambda}| + \lambda \|\bar{\mathbf{V}} - \mathbf{V}\|_{\text{Fro}} \leq 2\lambda_0 \varepsilon. \end{aligned}$$

Thus, for any $\varepsilon > 0$ we have $\mathcal{N}(C, \|\cdot\|_{\text{Fro}}, 2\lambda_0 \varepsilon) \leq \mathcal{N}((0, \lambda_0], |\cdot|, \lambda_0 \varepsilon) \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \varepsilon^{-1} \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \varepsilon)$ or equivalently for any $\varepsilon > 0$ we have $\mathcal{N}(C, \|\cdot\|_{\text{Fro}}, \varepsilon) \leq 2\lambda_0 \varepsilon^{-1} \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \varepsilon/(2\lambda_0))$. In conclusion we have

$$\mathcal{N}(\bar{S}_{\eta}^{-1}[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_{\Lambda}, 2(\varepsilon + Z_0 \eta)) \leq 2\lambda_{\eta} C_0 \varepsilon^{-1} \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{C_0}) \times \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \frac{\varepsilon}{2\lambda_0 C_0}),$$

which concludes the proof. □

A.7.2. Covering number for bounded sparse symmetric matrices and their secant

We now prove Lemma 13.

Proof of Lemma 13 Recall that Lemma 7 provides the following control of the covering number of the set of symmetric sparse matrices with bounded Frobenius norm $\mathfrak{W}_k = \{\Theta \in \mathcal{S}_d(\mathbb{R}) ; \|\Theta\|_0 \leq d + 2k, \|\Theta\|_{\text{Fro}} \leq 1\}$:

$$\mathcal{N}(\mathfrak{W}_k, \|\cdot\|_{\text{Fro}}, \varepsilon) \leq \left(\frac{ed^2}{2k}\right)^k \left(\frac{9}{\varepsilon}\right)^{d+k},$$

Noticing that $\mathfrak{S}_0^{-1} \subset \mathfrak{W}_k$, $\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1} \subset 2\mathfrak{W}_k$ and that $S[\mathfrak{S}_{k, \kappa_0}^{-1}] \subset \mathfrak{W}_{2k}$ yields the result. \square

A.7.3. Proof of Theorem A.6.3

Proof First recall that for any $\varepsilon' > 0$ and $\eta > 0$, by Proposition 8 we have

$$\mathcal{N}(\bar{S}_\eta[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, 2(\varepsilon' + Z_0\eta)) \leq 2\lambda_0 C_0 \varepsilon^{-1} \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon'}{C_0}) \times \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \frac{\varepsilon'}{2\lambda_0 C_0}),$$

Let us fix ε such that $0 < \varepsilon \leq \frac{2\kappa_0^3}{c_{\text{Fro}}} \sqrt{d} = c_{\text{Fro}} Z_0$, as in the hypothesis of Theorem A.6.3. We now set $\varepsilon' = \varepsilon/4$ and $\eta = \varepsilon/(4Z_0)$ such that $2(\varepsilon' + Z_0\eta) = \varepsilon$. Remark that these choices satisfy $\eta \leq c_{\text{Fro}}/2$ as desired in the previous section. Hence,

$$\begin{aligned} & \mathcal{N}(\bar{S}_\eta[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \\ & \leq \mathcal{N}(\bar{S}_\eta^+[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, \varepsilon) + \mathcal{N}(\bar{S}_\eta^-[\mathfrak{S}_{k, \kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \\ & \leq \mathcal{N}((0, 1], |\cdot|, \frac{\eta\varepsilon}{6L_0}) \times \mathcal{N}(\mathfrak{S}_0^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\eta\varepsilon}{6L_0})^2 \\ & \quad + 2\lambda_0 C_0 \varepsilon'^{-1} \mathcal{N}(\mathfrak{S}_{k, \frac{1}{\kappa_0}, 2}^{-1}, \|\cdot\|_{\text{Fro}}, \frac{\varepsilon'}{C_0}) \times \mathcal{N}(S[\mathfrak{S}_{k, \kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, \frac{\varepsilon'}{2\lambda_0 C_0}). \end{aligned}$$

Straightforward computations show that $\frac{\eta\varepsilon}{6L_0} = c_0 \varepsilon^2$ with $c_0 = \frac{1}{24Z_0 L_0}$, $\frac{\varepsilon'}{C_0 \eta} = c_1 \varepsilon$ where $c_1 = \frac{c_{\text{Fro}}}{4C_{\text{Fro}}(4\kappa_0^3 + c_{\text{Fro}}\kappa_0^2)}$ and $\frac{\varepsilon'}{2\lambda_0 C_0} = c_2 \varepsilon$ with $c_2 = \frac{c_{\text{Fro}}^2}{16C_{\text{Fro}}(4\kappa_0^3 + c_{\text{Fro}}\kappa_0^2)}$ ($= c_{\text{Fro}} c_1 / 4$). This allows us to continue

the computation as follows:

$$\begin{aligned}
& \mathcal{N}(\bar{S}_\eta[\mathfrak{S}_{k,\kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \\
& \leq \mathcal{N}((0, 1], |\cdot|, c_0 \varepsilon^2) \times N(\mathfrak{S}_0^{-1}, \|\cdot\|_{\text{Fro}}, c_0 \varepsilon^2)^2 \\
& \quad + \frac{1}{c_2} \varepsilon^{-1} \mathcal{N}(\mathfrak{S}_{k,\kappa_0}^{-1}, \|\cdot\|_{\text{Fro}}, c_1 \varepsilon) \times \mathcal{N}(S[\mathfrak{S}_{k,\kappa_0}^{-1}], \|\cdot\|_{\text{Fro}}, c_2 \varepsilon). \\
& \leq \frac{1}{c_0} \varepsilon^{-2} \times \left(\frac{ed^2}{2k}\right)^{2k} \left(\frac{18\sqrt{d}}{c_0 \varepsilon^2}\right)^{2(d+k)} \\
& \quad + \frac{1}{c_2} \varepsilon^{-1} \left(\frac{ed^2}{2k}\right)^k \left(\frac{2 \times 18\sqrt{d}}{c_1 \varepsilon}\right)^{d+k} \times \left(\frac{ed(d-1)}{4k}\right)^{2k} \left(\frac{18\sqrt{d}}{c_2 \varepsilon}\right)^{d+2k} \\
& \leq \frac{1}{c_0} \left(\frac{ed^2}{2k}\right)^{2k} \left(\frac{18\sqrt{d}}{c_0}\right)^{2(d+k)} \varepsilon^{-(4d+4k+2)} \\
& \quad + \frac{1}{c_2} \left(\frac{ed^2}{2k}\right)^k \left(\frac{2 \times 18\sqrt{d}}{c_1}\right)^{d+k} \times \left(\frac{ed^2}{4k}\right)^{2k} \left(\frac{18\sqrt{d}}{c_2}\right)^{d+2k} \varepsilon^{-(2d+3k+1)} \\
& \leq \left(\frac{ed^2}{2k}\right)^{3k} \left[\frac{(18\sqrt{d})^{2(d+k)}}{c_0^{2(d+k)+1}} \varepsilon^{-(4d+4k+2)} + \frac{2^{d+k} (18\sqrt{d})^{2d+3k}}{c_1^{d+k} c_2^{d+2k+1}} \varepsilon^{-(2d+3k+1)} \right].
\end{aligned}$$

Recall that c_0 , c_1 and c_2 depends on d (directly or from $c_{\text{Fro}} = 2/(9\sqrt{15d})$ and $C_{\text{Fro}} = 1/\sqrt{d}$) and κ_0 . Here are the explicit dependencies (\gtrsim meaning here \geq up to an absolute multiplicative constant): $c_0 \propto \kappa_0^{-5} d^{-3/2}$, $c_1 \gtrsim \kappa_0^{-3} d^{-1/2}$, $c_2 \gtrsim \kappa_0^{-3} d^{-3/2}$. So we get, that there exists absolute constants \tilde{c}_1 and \tilde{c}_2 such that

$$\begin{aligned}
& \mathcal{N}(\bar{S}_\eta[\mathfrak{S}_{k,\kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \\
& \leq \left(\frac{ed^2}{2k}\right)^{3k} \left[\tilde{c}_1^{d+k} \frac{d^{d+k} (\kappa_0^5 d^{3/2})^{2(d+k)+1}}{\varepsilon^{4d+4k+2}} + \tilde{c}_2^{d+k} \frac{\sqrt{d}^{2d+3k} (\kappa_0^3 d^{3/2})^{2d+3k+1}}{\varepsilon^{2d+3k+1}} \right] \\
& \leq \left(\frac{ed^2}{2k}\right)^{3k} \left[\left(\frac{\tilde{c}_1 \kappa_0^5 d^2}{\varepsilon^2}\right)^{2(d+k)+1} + \left(\frac{\tilde{c}_2 \kappa_0^3 d^2}{\varepsilon}\right)^{2d+3k+1} \right],
\end{aligned}$$

where we change the value of the absolute constant \tilde{c}_1 and \tilde{c}_2 at the last inequality. To finish the proof, recall that $\mathcal{N}(S_\eta[\mathfrak{S}_{k,\kappa_0}], \|\cdot\|_\Lambda, \varepsilon) \leq 2\mathcal{N}(\bar{S}_\eta[\mathfrak{S}_{k,\kappa_0}], \|\cdot\|_\Lambda, \varepsilon)$. \square

Acknowledgments

We gratefully acknowledge the support of the Centre Blaise Pascal's IT test platform at ENS de Lyon (Lyon, France) for Machine Learning facilities. The platform operates the SIDUS solution [63]. All the figures are generated using Matplotlib [43] and we also used Scipy and Numpy [40, 75] for computing our estimator. The authors are grateful to Mathurin Massias for insightful discussions.

Funding

This project was supported in part by the AllegroAssai ANR project ANR-19-CHIA-0009.

Data Availability Statements

The data and code underlying this article are available at the following link <https://github.com/tvayer/CompressedPrecisionMatrix>.

REFERENCES

1. Sohail Bahmani and Justin Romberg. Sketching for simultaneously sparse and low-rank covariance matrices. In *2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 357–360, 2015.
2. Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research*, 9(15):485–516, 2008.
3. Heinz H Bauschke, Jérôme Bolte, and Marc Teboulle. A descent lemma beyond lipschitz gradient continuity: first-order methods revisited and applications. *Mathematics of Operations Research*, 42(2):330–348, 2017.
4. Heinz H Bauschke, Patrick L Combettes, and Dominikus Noll. Joint minimization with alternating bregman proximity operators. *Pacific Journal of Optimization*, 2006.
5. Heinz H Bauschke, Minh N Dao, and Scott B Lindstrom. Regularizing with bregman–moreau envelopes. *SIAM Journal on Optimization*, 28(4):3208–3228, 2018.
6. Amir Beck. *First-Order Methods in Optimization*. SIAM-Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2017.
7. Matthias Bollhöfer, Aryan Eftekhari, Simon Scheidegger, and Olaf Schenk. Large-scale sparse inverse covariance matrix estimation. *SIAM Journal on Scientific Computing*, 41(1), 2019.
8. Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
9. Anthony Bourrier, Mike E. Davies, Tomer Peleg, Patrick Pérez, and Rémi Gribonval. Fundamental performance limits for ideal decoders in high-dimensional linear inverse problems. *IEEE Transactions on Information Theory*, pages 7928–7946, December 2014.
10. Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, March 2004.
11. Lev M Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967.
12. T. Tony Cai and Anru Zhang. Rop: Matrix recovery via rank-one projections. *The Annals of Statistics*, 43(1), Feb 2015.
13. Tony Cai, Weidong Liu, and Xi Luo. A constrained l1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.
14. Emmanuel Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Commun. ACM*, 55(6):111–119, jun 2012.
15. Emmanuel J Candes and Terence Tao. Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215, 2005.
16. Yair Censor and Stavros Andrea Zenios. Proximal minimization algorithm with d-functions. *Journal of Optimization Theory and Applications*, 73(3):451–464, 1992.
17. Antoine Chatalic. *Efficient and privacy-preserving compressive learning*. Theses, Université Rennes 1, November 2020.
18. Antoine Chatalic, Rémi Gribonval, and Nicolas Keriven. Large-Scale High-Dimensional Clustering with Fast Sketching. In *ICASSP 2018 - IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4714–4718, Calgary, Canada, April 2018. IEEE.

19. Yuxin Chen, Yuejie Chi, and Andrea J. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Transactions on Information Theory*, 61(7):4034–4059, 2015.
20. Kenneth L. Clarkson. Tighter bounds for random projections of manifolds. In *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry*, SCG '08, page 39–48, New York, NY, USA, 2008. Association for Computing Machinery.
21. Regev Cohen, Yochai Blau, Daniel Freedman, and Ehud Rivlin. It has potential: Gradient-driven denoisers for convergent solutions to inverse problems. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:18152–18164, 2021.
22. Regev Cohen, Michael Elad, and Peyman Milanfar. Regularization by denoising via fixed-point projection (red-pro). *SIAM Journal on Imaging Sciences*, 14(3):1374–1406, 2021.
23. R. Coleman. *Calculus on Normed Vector Spaces*. Universitext. Springer New York, 2012.
24. Rémi Delogne, Vincent Schellekens, and Laurent Jacques. Rop inception: signal estimation with quadratic random sketching. In *30th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, 2022.
25. Arthur P Dempster. Covariance selection. *Biometrics*, pages 157–175, 1972.
26. Hilmi E. Egilmez, Eduardo Pavez, and Antonio Ortega. Graph learning from data under laplacian and structural constraints. *IEEE Journal of Selected Topics in Signal Processing*, 11(6):825–841, 2017.
27. P. Erdős and A. Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6, 1959.
28. Jianqing Fan, Yang Feng, and Yichao Wu. Network exploration via the adaptive LASSO and SCAD penalties. *The Annals of Applied Statistics*, 3(2):521 – 541, 2009.
29. Fino and Algazi. Unified matrix treatment of the fast walsh-hadamard transform. *IEEE Transactions on Computers*, C-25(11):1142–1146, 1976.
30. Simon Foucart and Holger Rauhut. *A Mathematical Introduction to Compressive Sensing*. 2013.
31. Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, July 2008.
32. G. Frusque. *Inférence et décomposition modale de réseaux dynamiques en neurosciences*. PhD thesis, Université de Lyon, December 2020.
33. Rémi Gribonval, Gilles Blanchard, Nicolas Keriven, and Yann Traonmilin. Compressive Statistical Learning with Random Feature Moments. *Mathematical Statistics and Learning*, 3, 2021.
34. Rémi Gribonval, Gilles Blanchard, Nicolas Keriven, and Yann Traonmilin. Statistical Learning Guarantees for Compressive Clustering and Compressive Mixture Modeling. *Mathematical Statistics and Learning*, 3, 2021.
35. Rémi Gribonval, Antoine Chatalic, Nicolas Keriven, Vincent Schellekens, Laurent Jacques, and Philip Schniter. Sketching Data Sets for Large-Scale Learning: Keeping only what you need. *IEEE Signal Processing Magazine*, 38(5):12–36, September 2021.
36. Rémi Gribonval and Mila Nikolova. A characterization of proximity operators. *Journal of Mathematical Imaging and Vision*, 62(6-7):773–789, 2020.
37. Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
38. David Hallac, Youngsuk Park, Stephen Boyd, and Jure Leskovec. Network inference via the time-varying graphical lasso. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 205–213, New York, NY, USA, 2017. Association for Computing Machinery.
39. Andi Han, Bamdev Mishra, Pratik Kumar Jawanpuria, and Junbin Gao. On riemannian optimization over positive definite matrices with the bures-wasserstein geometry. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:8940–8953, 2021.
40. Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu

- Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, sep 2020.
41. Cho-Jui Hsieh, Mátyás A. Sustik, Inderjit S. Dhillon, and Pradeep Ravikumar. Quic: Quadratic approximation for sparse inverse covariance estimation. *Journal of Machine Learning Research*, 15(83):2911–2947, 2014.
 42. Cho-Jui Hsieh, Matyas A Sustik, Inderjit S Dhillon, Pradeep K Ravikumar, and Russell Poldrack. Big & quic: Sparse inverse covariance estimation for a million variables. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
 43. J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
 44. Samuel Hurault, Ulugbek Kamilov, Arthur Leclaire, and Nicolas Papadakis. Convergent bregman plug-and-play image restoration for poisson inverse problems. *arXiv preprint arXiv:2306.03466*, 2023.
 45. Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Proximal denoiser for convergent plug-and-play optimization with nonconvex regularization. In *International Conference on Machine Learning (ICML)*, pages 9483–9505. PMLR, 2022.
 46. Kishore Jaganathan, Yonina C Eldar, and Babak Hassibi. Phase retrieval: An overview of recent developments. *Optical Compressive Imaging*, pages 279–312, 2016.
 47. Maryia Kabanava, Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Stable low-rank matrix recovery via null space properties. *Information and Inference: A Journal of the IMA*, 5(4):405–441, 2016.
 48. Nicolas Keriven and Rémi Gribonval. Instance optimal decoding and the restricted isometry property. In *Journal of Physics: Conference Series*, volume 1131, page 012002. IOP Publishing, 2018.
 49. Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Low rank matrix recovery from rank one measurements. *Applied and Computational Harmonic Analysis*, 42(1):88–116, 2017.
 50. M. O. Kuismin, J. T. Kemppainen, and M. J. Sillanpää. Precision matrix estimation with rope. *Journal of Computational and Graphical Statistics*, 26(3):682–694, 2017.
 51. Sandeep Kumar, Jiaxi Ying, José Vinícius de M Cardoso, and Daniel P Palomar. A unified framework for structured graph learning via spectral constraints. *Journal of Machine Learning Research*, 21(22):1–60, 2020.
 52. Tim Tsz-Kit Lau and Han Liu. Bregman proximal langevin monte carlo via bregman-moreau envelopes. In *International Conference on Machine Learning*, pages 12049–12077. PMLR, 2022.
 53. Steffen L Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.
 54. Quoc Le, Tamás Sarló, Alex Smola, et al. Fastfood-approximating kernel expansions in loglinear time. In *Proceedings of the international conference on machine learning*, volume 85, page 8, 2013.
 55. Yuanxin Li, Yue Sun, and Yuejie Chi. Low-rank positive semidefinite matrix recovery from corrupted rank-one measurements. *IEEE Transactions on Signal Processing*, 65(2):397–408, 2017.
 56. Camilla Lingjærde, Tonje G Lien, Ørnulf Borgan, Helga Bergholtz, and Ingrid K Glad. Tailored graphical lasso for data integration in gene network reconstruction. *BMC bioinformatics*, 22(1):1–22, 2021.
 57. Weidong Liu and Xi Luo. Fast and adaptive sparse precision matrix estimation in high dimensions. *Journal of Multivariate Analysis*, 135:153–162, 2015.
 58. Gabor Lugosi, Jakub Trzuskowski, Vasiliki Velona, and Piotr Zwiernik. Learning partial correlation graphs and graphical models by covariance queries. *Journal of Machine Learning Research*, 22(203):1–41, 2021.
 59. Rahul Mazumder and Trevor Hastie. The graphical lasso: New insights and alternatives. *Electronic Journal of Statistics*, 6(none):2125 – 2149, 2012.
 60. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
 61. J-C Pesquet. Proximal approaches for matrix optimization problems: Application to robust precision matrix estimation. *Signal Processing*, 169:107417, 2020.
 62. Gilles Puy, Mike E. Davies, and Rémi Gribonval. Recipes for stable linear embeddings from Hilbert spaces to \mathbb{R}^m . *IEEE Transactions on Information Theory*, 2017. Submitted in 2015.

63. Emmanuel Quemener and Marianne Corvellec. Sidus—the solution for extreme deduplication of an operating system. *Linux Journal*, 2013(235):3, 2013.
64. Pradeep Ravikumar, Martin J. Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935 – 980, 2011.
65. James C. Robinson. *Dimensions, Embeddings, and Attractors*. Cambridge Tracts in Mathematics. Cambridge University Press, 2010.
66. Benjamin Rolfs, Bala Rajaratnam, Dominique Guillot, Ian Wong, and Arian Maleki. Iterative thresholding algorithm for sparse inverse covariance estimation. *Advances in Neural Information Processing Systems*, 25, 2012.
67. Ernest Ryu, Jialin Liu, Sicheng Wang, Xiaohan Chen, Zhangyang Wang, and Wotao Yin. Plug-and-play methods provably converge with properly trained denoisers. In *International Conference on Machine Learning*, pages 5546–5557. PMLR, 2019.
68. Alaa Saade, Francesco Caltagirone, Igor Carron, Laurent Daudet, Angélique Drémeau, Sylvain Gigan, and Florent Krzakala. Random projections through multiple optical scattering: Approximating kernels at the speed of light. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6215–6219. IEEE, 2016.
69. K. Sagar, S. Banerjee, J. Datta, and A. Bhadra. Precision matrix estimation under the horseshoe-like prior-penalty dual, 2021.
70. Martin Slawski, Ping Li, and Matthias Hein. Regularization-free estimation in trace regression with symmetric positive semidefinite matrices. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
71. Q. Sun, Kean M. Tan, H. Liu, and T. Zhang. Graphical nonconvex optimization via an adaptive convex relaxation. In *International Conference on Machine Learning*, volume 80, pages 4810–4817, 2018.
72. Marc Teboulle. A simplified view of first order methods for optimization. *Mathematical Programming*, 170(1):67–96, 2018.
73. Roman Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, page 210–268. Cambridge University Press, 2012.
74. Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
75. Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
76. Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
77. Lingxiao Wang and Quanquan Gu. Robust Gaussian graphical model estimation with arbitrary corruption. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3617–3626. PMLR, 06–11 Aug 2017.
78. Lingxiao Wang, Xiang Ren, and Quanquan Gu. Precision matrix estimation in high dimensional gaussian graphical models with faster rates. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 177–185, Cadiz, Spain, 09–11 May 2016. PMLR.
79. Hyenkyun Woo. A characterization of the domain of beta-divergence and its connection to bregman variational model. *Entropy*, 19(9):482, 2017.

80. P. Xu, J. Ma, and Q. Gu. Speeding up latent variable gaussian graphical model estimation via nonconvex optimization. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
81. Eunho Yang, Aurelie C Lozano, and Pradeep K Ravikumar. Elementary estimators for graphical models. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
82. Jieping Ye and Jun Liu. Sparse methods for biomedical data. *SIGKDD Explor. Newsl.*, 14(1):4–15, dec 2012.
83. J Ying, J. de Miranda Cardoso, and D. Palomar. Nonconvex sparse graph learning under laplacian constrained graphical model. In *Advances in Neural Information Processing Systems*, volume 33, pages 7101–7113, 2020.
84. Felix Xinnan X Yu, Ananda Theertha Suresh, Krzysztof M Choromanski, Daniel N Holtmann-Rice, and Sanjiv Kumar. Orthogonal random features. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
85. Ming Yuan. High dimensional inverse covariance matrix estimation via linear programming. *Journal of Machine Learning Research*, 11(79):2261–2286, 2010.
86. Ming Yuan and Yi Lin. Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.
87. Teng Zhang and Hui Zou. Sparse precision matrix estimation via lasso penalized D-trace loss. *Biometrika*, 101(1):103–120, 02 2014.
88. Yiming Zuo, Yi Cui, Guoqiang Yu, Ruijiang Li, and Habtom W Resson. Incorporating prior biological knowledge for network-based differential gene expression analysis using differentially weighted graphical lasso. *BMC bioinformatics*, 18(1):1–14, 2017.