

Stable rank one matrix completion is solved by two rounds of semidefinite programming relaxation.

Augustin Cosse¹ and Laurent Demanet²

¹ Département de Mathématiques et Applications,
Ecole Normale Supérieure, Ulm, Paris.
PSL Research University.

²Department of Mathematics, Massachusetts Institute of Technology, MA

December 2017

Abstract

This paper studies the problem of deterministic rank-one matrix completion. It is known that the simplest semidefinite programming relaxation, involving minimization of the nuclear norm, does not in general return the solution for this problem. In this paper, we show that in every instance where the problem has a unique solution, one can provably recover the original matrix through two rounds of semidefinite programming relaxation with minimization of the trace norm. We further show that the solution of the proposed semidefinite program is Lipschitz-stable with respect to perturbations of the observed entries, unlike more basic algorithms such as nonlinear propagation or ridge regression. Our proof is based on recursively building a certificate of optimality corresponding to a dual sum-of-squares (SOS) polynomial. This SOS polynomial is built from the polynomial ideal generated by the completion constraints and the monomials provided by the minimization of the trace. The proposed relaxation fits in the framework of the Lasserre hierarchy, albeit with the key addition of the trace objective function. We further show how to represent and manipulate the moment tensor in favorable complexity by means of a hierarchical low-rank decomposition.

Acknowledgement. Both authors were supported by a grant from the MISTI MIT-Belgium seed fund. AC was supported by the FNRS, FSMP, BAEF and Francqui Foundations. LD is supported by AFOSR grant FA9550-17-1-0316, ONR grant N00014-16-1-2122, and NSF grant DMS-1255203. AC is grateful to MIT Math, Harvard IACS, the University of Chicago as well as NYU Courant Institute and Center for Data Science for hosting him during this work.

1 Introduction

Low rank matrix completion has been studied extensively throughout the last few years, among other reasons because of its practical interest in machine learning and data science. Completion provides a useful tool to compress and manipulate large databases such as in genomics and finance, and to infer information from a few measurements such as in collaborative filtering or triangulation. Good introductions as well as recovery results for random designs and arbitrary ranks can be found in [16, 30].

The objective of this paper is to provide an algorithm that solves the rank one case in a stable and comprehensive way. Let $\mathcal{M}(1; m \times n)$ denote the set of rank-1 matrices of size $m \times n$; we consider the problem

of recovering an unknown rank one matrix $\mathbf{X}_0 \in \mathcal{M}(1; m \times n)$, $\mathbf{X}_0 = \mathbf{x}_0 \mathbf{y}_0^T$ when we are given $\mathcal{O}(m+n)$ entries from this matrix, possibly corrupted by an additive noise ε . We do not make any assumption on the noise. In the noiseless case, this problem reads

$$\begin{aligned} & \text{find } \mathbf{X} \in \mathbb{R}^{m \times n} \\ & \text{subject to } \text{rank}(\mathbf{X}) = 1 \\ & \mathbf{X}_{ij} = (\mathbf{X}_0)_{ij} \quad (i, j) \in \Omega, \end{aligned} \tag{1.1}$$

where Ω denotes the set of measurements. As a slight abuse, we will also speak of constraints $\{\mathbf{X}_{ij} - (\mathbf{X}_0)_{ij} = 0\}_{(i,j) \in \Omega}$ as belonging to the set Ω . In the noisy case, the data fit constraint is relaxed to $\|\mathbf{X}_{ij} - ((\mathbf{X}_0)_{ij} + \varepsilon_{ij})\| \leq \sigma$ in a standard fashion.

Clearly, one cannot always solve problem (1.1). For example, if no information is known on a given column (resp. row), it becomes impossible to recover the entries corresponding to this column (resp. row). Another limitation occurs when the rank-1 matrix has a zero entry; then the corresponding row or column will be zero, and the completion problem will generically lack injectivity. As an illustration of the issue with zero entries, consider the problem where the first row and last column are known and are both trivial. The number of measurements is $(m+n-1)$. However in this case, any matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ of the form $\mathbf{X} = \mathbf{v} \mathbf{w}^*$ with $v_1 = w_n = 0$ is a valid solution of the problem. For this reason, we consider the completion problem on $\mathcal{M}^*(1, m \times n)$, where $\mathcal{M}^*(1, m \times n)$ denotes the restriction of $\mathcal{M}(1; m \times n)$ to matrices for which none of the entries are zero.

To formalize the notion of injectivity, we introduce the mapping $\mathcal{R}_\Omega : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{|\Omega|}$ that corresponds to extracting the observed entries of the matrix. We let \mathcal{R}_Ω^1 denote the restriction of \mathcal{R}_Ω to matrices of rank-1 that have no zero rows/columns. Invertibility of this restriction \mathcal{R}_Ω^1 corresponds to asking whether one can uniquely recover the matrix \mathbf{X} from the knowledge of $\mathcal{R}_\Omega(\mathbf{X})$ and the fact that \mathbf{X} has rank 1. Let us denote by $\mathcal{V}_1, \mathcal{V}_2$ the sets of row and column indices of \mathbf{X} . We consider the bipartite graph $\mathcal{G}(\mathcal{V}_1, \mathcal{V}_2, \mathcal{E})$ associated to problem (1.1), where the set of edges is defined by $(i, j) \in \mathcal{E}$ iff $(i, j) \in \Omega$. The vertices of the bipartite graph \mathcal{G} corresponding to \mathbf{X} are labeled by the corresponding row and column indices. The conditions for the recovery of the matrix \mathbf{X} from the set Ω are related to the properties of this bipartite graph as expressed by the following lemma which can be found, for example, in [33]:

Lemma 1 (Rank-1 completion). *The mask \mathcal{R}_Ω is injective on $\mathcal{M}^*(1; m \times n)$ if and only if \mathcal{G} is connected.*

Lemma (1) has an interesting consequence. Within the noiseless framework, rank one matrix completion can be solved exactly through a nonlinear propagation approach. To understand this, let us write $\mathbf{X} = \mathbf{x} \mathbf{y}^T \in \mathbb{R}^{m \times n}$ with $x_1 = 1$. Let us further use $\mathbf{z} \in \mathbb{R}^{m+n-1}$ to denote the concatenation of $[x_2, \dots, x_m] \in \mathbb{R}^{m-1}$ and $\mathbf{y} \in \mathbb{R}^n$, $\mathbf{z} = (\mathbf{x}, \mathbf{y})$. When we deal with the rank one case, an implication of lemma 1 is that for all x_n, y_m , the bipartite graph corresponding to the mask Ω always contains at least one connected path starting with an edge corresponding to an element of the first row and for which the series of existing edges corresponds to running through \mathbf{X} according to chains of constraints such as

$$y_{i_1} \rightarrow y_{i_1} x_{i_2} \rightarrow x_{i_2} y_{i_3} \rightarrow \dots y_{i_{L+1}} x_n \quad (\text{to reach } x_n) \tag{1.2}$$

$$y_{i'_1} \rightarrow y_{i'_1} x_{i'_2} \rightarrow x_{i'_2} y_{i'_3} \rightarrow \dots x_{i'_{L'+1}} y_m \quad (\text{to reach } y_m) \tag{1.3}$$

More generally, the two chains (1.2) and (1.3) can read, using the vector $\mathbf{z} \in \mathbb{R}^{m+n-1}$,

$$z_{i_1} \rightarrow z_{i_1} z_{i_2} \rightarrow z_{i_2} z_{i_3} \rightarrow \dots z_{i_{L+1}} z_n. \tag{1.4}$$

In other words, each of the entries of \mathbf{x} and \mathbf{y} can always be related to an element of the first row whose value is known because of the normalization $x_1 = 1$. Each of the elements making up the bilinear constraints can then be obtained in the absence of noise by iteratively propagating the value of the elements of the first row through (1.4). As we explain in the sequel, such a propagation scheme however lacks robustness to noise, especially, when the magnitude of the entries is on the order of the magnitude of the noise.

When the measurements are corrupted by noise, a popular approach is to turn to minimization of the nuclear norm as a proxy for the rank (see [23, 46] for early references). However, the nuclear norm does not always guarantee recovery of the rank one matrix \mathbf{X}_0 when the noise vanishes. An important gap regarding rank-one matrix completion has thus been the lack of an algorithm providing a proper (deterministic) stability estimate of the form

$$\|\mathbf{X} - \mathbf{X}_0\| \leq \omega(\|\mathcal{R}_\Omega(\mathbf{X}) - \mathcal{R}_\Omega(\mathbf{X}_0)\|). \quad (1.5)$$

for some Lipschitz function $\omega(\tau)$ obeying $\omega(\tau) \rightarrow 0$ when $\tau \rightarrow 0$.

We can now state the main contributions of the paper.

- First, we show that rank-one matrix completion can be solved through two rounds of semidefinite programming relaxation. Our result is sharp in terms of measurements; recovery is always possible as soon as the nonlinear problem has a unique solution. This is in contrast to previous results that required a random measurement set [14]. This result also confirms that there exist instances of rank minimization problems that are solved in a comprehensive way (without constraints of incoherence and/or randomness) using a fixed, higher (> 1) number of rounds of semidefinite programming relaxation.
- Second, we show that when the measurements are corrupted by noise, the solution to the semidefinite relaxation remains proportional to the noise level. In particular, this solution is shown to be Lipschitz-stable with respect to the noise level. This is in contrast with nonlinear approaches such as [33, 34].
- Finally, our proof system, based on constructing a dual polynomial, incidentally reveals two important facts: First, minimization of the trace norm helps certify recovery because it provides additional squares of monomials that are useful in constructing the dual polynomial. Second, recovery can be related to the possibility of propagating known information through the graph by means of polynomial equations.

The next sections discuss the limitations of propagation, minimization of the nuclear norm, and ridge regression. We illustrate these limitations on the simple problem of completing the rank-one matrix \mathbf{X}_0 with a small parameter δ ,

$$\mathbf{X}_0 = \begin{pmatrix} 1 & ? \\ \delta & 1 \end{pmatrix}, \quad (1.6)$$

for which, given the rank one constraint, the only missing entry is obviously given by $1/\delta$.

1.1 Propagation is unstable

We start by discussing the simple propagation scheme. In the noiseless framework, this scheme can be efficiently applied by writing \mathbf{X}_0 as

$$\mathbf{X}_0 = \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix}, \quad (1.7)$$

and simply deriving X_{12} as $X_{12} = \frac{X_{22}X_{11}}{X_{21}}$. Now assume that the entries are corrupted by a noise ϵ so that the measurements are now given by $\tilde{X}_{11} = (X_0)_{11} + \epsilon_{11}$, $\tilde{X}_{21} = (X_0)_{21} + \epsilon_{21}$ and $\tilde{X}_{22} = (X_0)_{22} + \epsilon_{22}$ with $\epsilon = (\epsilon_{11}, \epsilon_{21}, \epsilon_{22}) \in \mathbb{R}^3$. Taking a noise ϵ with $\|\epsilon\|$ on the order of δ , such as for example $\epsilon_{21} = -.9\delta$, will result in important errors when using propagation as shown below,

$$\tilde{X}_{12} = \frac{\tilde{X}_{22}\tilde{X}_{11}}{\tilde{X}_{21}} = \frac{1 + \mathcal{O}(\delta)}{\delta + \epsilon_{21}} \sim 10\frac{1}{\delta} \quad \frac{\tilde{X}_{12} - (X_0)_{12}}{(X_0)_{12}} = 900\% \quad (1.8)$$

The estimate derived through the propagation algorithm are thus unreliable when the entries are corrupted by an unknown noise ε of magnitude comparable to the smallest entries in the matrix. In addition, there is no effective, general method to select the propagation path optimally.

Another elementary method consists in taking the logarithm of the constraints, and solving the resulting system to obtain the logarithm of the unknowns. It is a very reasonable method for some positive matrices, although it is easy to see that it suffers from a similar kind of instability as the propagation scheme.

1.2 Nuclear norm fails

In this section, we briefly study how nuclear norm minimization would perform in the same framework of problem (1.6) as before. Nuclear norm minimization was first formalized in [23, 46] and guarantees were given, for the matrix completion problem, in a probabilistic framework, in [15]. Nuclear norm minimization relies on solving the convex program

$$\text{minimize } \|\mathbf{X}\|_* \tag{1.9a}$$

$$\text{subject to } \mathbf{X}_{ij} = (\mathbf{X}_0)_{ij}, \quad (i, j) \in \Omega. \tag{1.9b}$$

Where $\|\mathbf{X}\|_*$ is used as a proxy for the rank. In the case of (1.6), for a sufficiently small δ , if we let

$$\mathbf{X} = \begin{pmatrix} 1 & 2/\delta \\ \delta & 1 \end{pmatrix} \tag{1.10}$$

it can be easily verified that $\|\mathbf{X}\|_* < \|\mathbf{X}_0\|_*$, $\text{rank}(\mathbf{X}) = 2$, and the nuclear norm minimization (1.9) thus doesn't return the unknown matrix \mathbf{X}_0 despite the fact that a sufficient amount of measurements are provided.

In the case of a symmetric, positive semidefinite matrix \mathbf{X} , program (1.9) becomes minimization of the trace under the constraint $\mathbf{X} \succeq 0$. When the diagonal is fully measured, it is known that this formulation succeeds at recovering the original matrix \mathbf{X}_0 when the completion problem is well-posed [22, 27].

1.3 Ridge regression has local minima

For the sake of completeness, we briefly discuss ridge regression (a.k.a Tikhonov regularization) on the rank one factorization. This approach has gained in popularity over the last years and is equivalent to solving the quartic regularized problem. In fact it is natural to wonder whether the semidefinite programming formulation of this paper which relies on the minimization of the trace norm is not simply a form of ridge regression. This section precisely refutes this idea. The ridge regression problem reads

$$\min \|\mathcal{R}_\Omega(\mathbf{x}\mathbf{y}^T) - \mathcal{R}_\Omega(\mathbf{X}_0 + \varepsilon)\|_F^2 + \lambda(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2) \tag{1.11}$$

The most popular way to solve problem (1.11) is through gradient descent. However, when several measurements are given and no convergence guarantee is known, it is not clear how to initialize the algorithm. We choose to follow standard practice and initialize it with the singular value decomposition of the matrix $\sum_{ij} (\mathcal{P}_\Omega(\mathbf{X}))_{ij} e_i e_j^T$ and taking the outerproduct of the corresponding top singular vectors weighted by their singular value. As shown by Fig. 1, the landscape underlying this formulation suffers from a lack of convexity. As a consequence, even in the absence of noise when the matrix size is sufficiently large, and the number of measurements is sufficiently close to the recoverability limit, ridge regression will fail to return the global minimizer.




Figure 1: Representation of the ridge regression energy landscape for a rank-one matrix completion problem with $m = n = 1$. We consider the minimization problem $f(x, y) = \|xy - 1\|^2 + \|y - \delta\|^2$, plus regularization terms such as in (1.11). Without regularization, there is one infimum at $(-\infty, 0)$, and one minimum at $(\delta, 1/\delta)$ (In the figure above, we take $\delta = .1$). With regularization the infimum located at $(-\infty, 0)$ becomes a minimum and moves closer to the origin such as shown above. The evolution of the loss function on the line joining the two minima is displayed by the figure on the right.

1.4 Algorithm: two rounds of semidefinite relaxation

When minimizing the nuclear norm of rank one matrices, one only enforces constraints on monomials of degree at most two on the entries of the generating vectors \mathbf{x} and \mathbf{y} . The nuclear norm was shown in [23] to be equivalent to the following semidefinite program,

$$\begin{aligned}
 & \text{minimize} && \text{Tr}(\mathbf{W}) \\
 & \text{subject to} && \mathbf{W} = \begin{pmatrix} \mathbf{X}_{11} & \mathbf{X}_{12} \\ \mathbf{X}_{21} & \mathbf{X}_{22} \end{pmatrix} \\
 & && (\mathbf{X}_{12})_{ij} = (\mathbf{X}_0)_{ij}, \quad (i, j) \in \Omega, \\
 & && \mathbf{W} \succeq 0.
 \end{aligned} \tag{1.12}$$

When $\mathbf{X}_0 = \mathbf{x}_0 \mathbf{y}_0^T$, the matrix \mathbf{W} is a proxy for the rank one matrix

$$\mathbf{W}_0 = \begin{pmatrix} \mathbf{x}_0 \mathbf{x}_0^T & \mathbf{x}_0 \mathbf{y}_0^T \\ \mathbf{y}_0 \mathbf{x}_0^T & \mathbf{y}_0 \mathbf{y}_0^T \end{pmatrix} \tag{1.13}$$

The positive semidefinite constraint on \mathbf{X} is thus used in combination with the trace norm, as a convex relaxation of the rank one constraint. It is interesting to note that formulation (1.12) only optimizes over monomials of bidegree (1, 1). The key idea of the "second round of lifting" is to extend this type of formulation to monomials of higher degree in the original unknowns. Introducing $\mathbf{z}_0 = (\mathbf{x}_0, \mathbf{y}_0)$ and $\mathbf{z}_0^{(2)} = \text{vec}(\mathbf{z}_0^{\otimes 2}) = \text{vec}(\mathbf{z}_0 \otimes \mathbf{z}_0)$, we consider the larger matrix \mathbf{M}_0 defined as

$$\mathbf{M}_0 = \begin{pmatrix} 1 & \mathbf{z}_0 & \mathbf{z}_0^{(2)} \\ \mathbf{z}_0 & \mathbf{z}_0^{\otimes 2} & \mathbf{z}_0 \otimes \mathbf{z}_0^{(2)} \\ \mathbf{z}_0^{(2)} & \mathbf{z}_0^{(2)} \otimes \mathbf{z}_0 & \mathbf{z}_0^{(2)} \otimes \mathbf{z}_0^{(2)} \end{pmatrix}, \tag{1.14}$$

In (1.14), we thus have $\mathbf{W}_0 = \mathbf{z}_0^{\otimes 2}$. At order two, the semidefinite relaxation considers as unknowns all the entries of a positive semi-definite proxy \mathbf{M} of the same structure as \mathbf{M}_0 , but without the explicit link to a

vector \mathbf{z}_0 . Instead, two categories of linear constraints are intended to force the matrix \mathbf{M} to inherit the structure of \mathbf{M}_0 :

- *Structural constraints/ total symmetry.* Due to the additional monomials that appear in (1.14), there now exist corresponding additional relations between the entries of \mathbf{M}_0 (and thus \mathbf{M} as well). In particular, all monomials in $\mathbf{z}_0^{(2)}$ find an exact match in the block $\mathbf{z}_0^{\otimes 2}$. Within the block $\mathbf{z}_0^{(2)} \otimes \mathbf{z}_0^{(2)}$, one must also list all the total symmetry constraints of a tensor of order 4. More generally, the structural constraints enforce equality of the entries that are identical in the rank-one matrix \mathbf{M}_0 .
- *Higher-order affine constraints.* Similarly, for any of the original affine constraints applying on the elements of $\mathbf{z}_0^{\otimes 2}$, one can now define higher order constraints that are jointly enforced on the elements of \mathbf{z}_0 and the elements of the block $\mathbf{z}_0^{(2)} \otimes \mathbf{z}_0$. As an example, consider that one is given the constraint $X_{ij} = x_i y_j = (X_0)_{ij}$. It is now possible to enforce the constraints $x_i y_i x_k - (X_0)_{ij} x_k$ for any monomial x_k . More generally, the higher-order constraints are obtained by multiplying the original constraints by any product of the entries of x and y of degree at most two.

The second-order formulation in this paper consists in combining all these constraints with $\mathbf{M} \succeq 0$, and with minimization of $\text{Tr}(\mathbf{M})$. The point of this paper is to prove that this formulation, of order 4 in \mathbf{x} and \mathbf{y} , is enough to recover every rank-one matrix in the completion problem, and to provide a scalable algorithm for it.

This idea is not new and can be found through various formulations in the work of Parrilo [41, 42], Shor [49, 47, 48], Nesterov [38], and Lasserre [35]. For now, we simply write this semidefinite programming relaxation in the following general form. More details on the constraints will be provided in section 2.1.

$$\begin{aligned} & \text{minimize} && \text{Tr}(\mathbf{M}) \\ & \text{subject to} && \mathcal{A}(\mathbf{M}) = \mathbf{b}, \\ & && \mathbf{M} \succeq 0. \end{aligned} \tag{1.15}$$

At this point, we just note that the linear map \mathcal{A} now encodes the original constraints from (1.1) together with the additional structural and higher order constraints mentioned above.

To define the corresponding stable formulation for the semidefinite relaxation (1.15), we first introduce a decomposition of the linear map into the structural part \mathcal{A}_S and the remaining part $\tilde{\mathcal{A}}$. The motivation behind such a decomposition comes from the fact that structural constraints, unlike the original measurement constraints and their higher order extensions, are not affected by noise. The higher-order constraint in the Lasserre hierarchy use (noisy) data in an essential manner in their expression, not just in a right-hand-side, hence we use the notation $\tilde{\mathcal{A}}$ as a shorthand for those constraints. Following this decomposition, the stable version of the relaxation (1.15) can be posed generally as,

$$\begin{aligned} & \text{minimize} && \text{Tr}(\mathbf{M}) \\ & \text{subject to} && \mathcal{A}_S(\mathbf{M}) = 0 \\ & && \|\tilde{\mathcal{A}}(\mathbf{M}) - \tilde{\mathbf{b}}\| \leq \sigma \\ & && \mathbf{M} \succeq 0. \end{aligned} \tag{1.16}$$

The vector $\tilde{\mathbf{b}}$ encodes the noisy measurements $(X_0)_{ij} + \varepsilon_{ij}$ and the RHS corresponding to their higher order extensions, $(i, j) \in \Omega$. The next section introduces the main result of this paper.

1.5 Main Result

The main result of this paper only requires the necessary and sufficient conditions of lemma 1 to be satisfied. Our first theorem states that the noiseless semidefinite program (1.15), for a linear map \mathcal{A} encoding both the original constraints as well as the additional higher order and structural constraints, solves the rank one completion problem exactly under minimization of the Trace norm.

Theorem 1. *Consider problem (1.1) in the context of lemma 1, with $\mathbf{X} \in \mathcal{M}^*(1; m \times n)$. Then this problem can be solved exactly through two rounds of semidefinite programming relaxation with minimization of the trace norm.*

The interest of semidefinite programs lies in their robustness vis-a-vis corruption of the data. This is what Theorem 2 below makes precise. It shows that when considering observations that are corrupted by a noise ε , so that $(\tilde{X}_0)_{ij} = (X_0)_{ij} + \varepsilon_{ij}$, the solution to the semidefinite programming relaxation (1.16) remains within the noise level.

Theorem 2. *Let \mathbf{M}_0 denote the rank one matrix introduced in (1.14) for $\mathbf{z}_0 \in \mathbb{R}^{m+n-1}$. Assume that the necessary and sufficient conditions of lemma 1 are satisfied. Let $\|\varepsilon\|_2 = \sqrt{\sum_{(i,j) \in \Omega} \varepsilon_{ij}^2}$. Let \mathbf{M} denote the solution to the semidefinite program (1.16). This solution satisfies*

$$\frac{\|\mathbf{M} - \mathbf{M}_0\|_F}{\|\mathbf{M}_0\|_F} \leq C_0(m+n)^{7/2} \|\varepsilon\|_2. \quad (1.17)$$

The constant C_0 depends on the entries of \mathbf{X}_0 , but not on Ω , m , n , or ε .

Once \mathbf{M} is found, one can read off \mathbf{X} from the entries of \mathbf{M} corresponding to $|\alpha| = 1$. Note that for the propagation and log-system algorithms, a similar error bound can only be expected to hold provided $\|\varepsilon\|_2 \leq c \min_{i,j} |(X_0)_{ij}|$ for some $c < 1$, and would otherwise become unbounded.

Most of the $(m+n)^{7/2}$ multiplicative factor in Theorem 2 arises because of the propagation of noise through the certificate (i.e. the fact that the certificate relies on a chain of length $m+n$). As we don't make prior assumption regarding propagation of information through the graph, the noise along the path is constrained by a global bound of the form $\|\varepsilon\| \sqrt{(1 + \|\mathbf{m}_0\|_1)}$, instead of a bound in $\mathcal{O}(\|\varepsilon\|)$ which could be enforced if the path was explicitly given. This global rather than path specific bound results in additional $\mathcal{O}(\|\varepsilon\|(m+n))$ multiplicative factors when a path specific bound, would lead to better $\mathcal{O}(\|\varepsilon\|)$. The remaining factor affecting the constant is the size of the moments matrix. In this regard, a second possible simplification is to restrict the set of second-order moments to moments that appear in the constraints only. In this case, the first column of the moments matrix (1.14) reduces to $(1, \mathbf{z}_0, (x_0)_i(y_0)_j)$ for $(i, j) \in \Omega$ and the semidefinite program (1.16) becomes as scalable as the traditional SDP relaxation, or the first round of the Lasserre hierarchy since the matrix of unknowns is now on the order of $(m+n)^2$

Those ideas are summarized through Corollary 3 below.

Corollary 3. *Assume that the paths in the bipartite graph relating each of the unknown vertices x_i (resp. y_j) to the root node y_ℓ are explicitly given. Then we have the following stability estimate,*

$$\frac{\|\mathbf{M} - \mathbf{M}_0\|_F}{\|\mathbf{M}_0\|_F} \leq C_0(m+n)^2 \|\varepsilon\|_2. \quad (1.18)$$

The constant C_0 depends on the entries of \mathbf{X}_0 , but not on Ω , m , n , or ε .

Before introducing the proofs of Theorems 1 and 2 as well as Corollary 3, and describing how to make numerical schemes scalable, the next section lists the most relevant connections of this work with the developments of the last few years.

1.6 Connections with existing work

Low rank matrix completion and semidefinite programming relaxations have both attracted a lot of attention from various communities over the past few years. Low-rank matrix completion is a problem that has been extensively studied in the literature and has led to numerous successful approaches. One of the most famous, nuclear norm minimization, led to the derivation of important recovery guarantees [12, 14, 16, 18, 45]. In [14] in particular, the authors derive (probabilistic) bounds on the recovery error for low rank matrices, when the measurements are corrupted by noise. Other notable progress on this question includes the results of Keshavan et al. [30] who certify recovery with high probability (w.h.p.) given $\mathcal{O}((m+n)\log(m+n))$ measurements and some incoherence conditions in a noiseless framework. In [28, 29], the same authors derive a probabilistic bound that can be used in the presence of noise and improve the results obtained in [28] to a recovery w.h.p. that scales linearly in the noise as $\mathcal{O}(m+n)$ provided that both the magnitude of the entries as well as the number of measurements are sufficiently large. The noise is assumed to be i.i.d random with zero mean and sub-gaussian tail.

Singer et al. [50] investigate matrix completion with a non random sampling mask based on the structure of the measurements. Their paper is interested in determining whether completion is possible or not in the general rank- r case using rigidity theory. Other papers that focus on characterizing the sampling patterns enabling matrix completion include [32, 34] where the authors extend the idea which was studied for rank one matrices in [33] to the more general case of low rank matrices. Those papers show that feasibility and uniqueness of the completion only depends on the structure of the measurements. They propose an algorithm based on completion of the k -by- k minors (circuit polynomials) to complete the matrix in the noiseless case. This algorithm lacks robustness to noise and requires an additional step averaging the values over different paths across the bipartite graph to cancel out the noise. The idea therefore cannot be applied in the case of $\mathcal{O}(m+n)$ measurements. Neither can it be used with deterministic noise. The question of completability patterns is also discussed in the more recent [43] by Pimentel et al. In [44], the same authors study reconstruction under an i.i.d. zero mean noise with covariance matrix $\sigma^2\mathbf{I}$. This last paper is based on the concept of recommender systems where only a given number of measurements ($r+1$) are allowed per column. They show asymptotic convergence of their estimator for sufficiently large matrices, when a sufficiently large number of columns are used to generate the measurements.

Since the pioneering work of Goemans and Williamson [24], which started popularizing the use of semidefinite programs as an approximation to hard quadratic optimization problems, semidefinite programming has gained a reputation as a potentially powerful candidate to derive interesting approximations to hard/nonconvex problems. This activity culminated in the now famous Unique Games Conjecture [31] in complexity theory. Examples of successful developments based on semidefinite programming or nuclear norm relaxations of nonlinear problems can be found in [16, 17, 22, 13, 2]. In [22, 27], one of the authors solves the symmetric rank one matrix completion problem when the diagonal entries are given. The proof system in this paper relies on spectral graph theory, and use the fact that the eigenvector of the exact solution \mathbf{X}_0 is also an eigenvector of the data weighted graph Laplacian, to derive a bound on the recovery. Finally, the noiseless result of this paper was presented in the introductory note [19].

As stated earlier, the semidefinite program (1.14) of this paper is in fact an instance of the more general Lasserre and sum-of-squares (SOS) hierarchies of semidefinite programs [37] which were introduced through the work of Parrilo [41, 42], Shor [49, 47, 48], Nesterov [38], and Lasserre [35] as an extension of the basic semidefinite programming relaxation. Those hierarchies are based on making semidefinite programming relaxations gradually tighter by adding more variables and constraints, resulting in optimization on gradually larger subspaces.

Semidefinite programming hierarchies have received a lot of attention over the last few years, both positively as a potential extension to the traditional semidefinite programming relaxations, and negatively because of their practical intractability resulting from the gradually higher dimension of their successive rounds.

Another important drawback associated to those hierarchies has been the lack of convincing instances for which rounds higher than one were leading to noticeable improvements. For the most recent developments on the convergence of those hierarchies, see for example the papers by Barak et al. [3, 5, 4].

A few improvements have however been made over the last few years. On the first point, in a paper which is very related to this one [52], Tang et al. show that the tensor decomposition problem can be solved through a semidefinite programming relaxation with minimal number of rounds. In [4], Barak et al. certify using the Rademacher complexity, that tensor completion can be solved with high probability with $\tilde{O}(n^{3/2})$ measurements through 6 rounds of semidefinite programming relaxation. Finally, other results along that line recently appeared in papers by Nie et al. In a first paper tackling asymptotic convergence to the minimum [39] for general polynomial problems, these authors introduce an updated formulation based on the Jacobian of the polynomial constraints for which convergence of the hierarchy at a sufficiently large order is certified. In a second monograph [21] which is discussed further in the last section of this paper, the same author shows that computation of the real eigenvalues of symmetric tensors can be achieved through a finite number of semidefinite programming relaxation rounds. No upper bound is provided on the number of rounds required.

Convergence of the sum-of-squares and Lasserre hierarchies are also discussed by Gouveia et al. in [25]. This paper relates the sequence of theta bodies of an ideal \mathcal{I} and the Lasserre hierarchy and shows that under some assumptions, the k^{th} theta body of an ideal is equal to the set of solutions resulting from the k^{th} round of the Lasserre hierarchy and that for real radical ideal the k^{th} theta body corresponds to the closure of the convex hull of the variety of the ideal \mathcal{I} as soon as every polynomial of degree one that is non negative on the variety can be represented by sum of squares of degree at most k modulo the ideal.

A few papers adress simplifications of higher rounds of semidefinite programming relaxations by means of sparsity of the polynomial constraints. Among those papers, Lasserre [36] as well as Nie et al. [40] introduce tailored relaxations for problems where sparsity occurs in the constraints and the objective function. This adapted relaxation enables a significant reduction in the size of the matrices whenever a property known as *the running intersection property* is satisfied together with some independence between the sets of variables used by the constraints. Ahmadi [1] also discusses possible reduction in the complexity of semidefinite programming hierarchies by means of the chordal extension of the graph whose cliques are defined from the polynomial constraints.

The Lasserre and sum-of-squares hierarchies are built upon the resolution of systems of polynomial equations. For this reason we also briefly address another important line of work following from computational algebraic geometry. When looking for the solution to a system of polynomial equations (in particular when the underlying ideal is zero-dimensional), the very first question one want to ask is whether it is possible to compute a Gröbner basis for this system (see for example [20] as well as [51]). When such a Gröbner basis can be found, the solutions can be computed as the vectors of joint eigenvalues of the companion matrices (see Theorem 2.6 in [51]).

Computing a Gröbner basis is at least NP-complete in the general case (see for example [6]). In fact, the notion of Gröbner basis is somehow complementary to proving the tightness of the SOS/Lasserre hierarchies. Finding one helps understand the other and vice versa. The degree of the Gröbner basis is unknown before the computation and bases with higher degree are more difficult to compute. Moreover the numerical computations involved are known to be numerically unstable (see for example [11, 26] and the discussion therein) and do not scale well with the dimension. The dual polynomial that we build in this paper is in fact equivalent to showing that such a Gröbner basis can be constructed (modulo the Trace) with degree at most 4 from the polynomial ideal generated by the completion constraints.

1.7 Notations

Let $\mathbf{z} = (\mathbf{x}, \mathbf{y}) \in \mathbb{R}^{m+n-1}$. When dealing with algebraic problems like (1.1), it will be useful to write those problems as general polynomial optimization problems (POP) of the form

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} \quad p_0(\mathbf{z}), \quad \mathbf{z} \in \mathbb{R}^n \quad (1.19a)$$

$$\text{subject to} \quad p_1(\mathbf{z}) \geq 0, p_2(\mathbf{z}) \geq 0, \dots, p_m(\mathbf{z}) \geq 0. \quad (1.19b)$$

For some polynomials $p_0(\mathbf{z}), p_1(\mathbf{z}), \dots, p_m(\mathbf{z}) \in \mathbb{R}[\mathbf{z}]$ where $\mathbb{R}[\mathbf{z}]$ is used to denote the ring of multivariate polynomials in the optimization variable $\mathbf{z} \in \mathbb{R}^{m+n-1}$. We will sometimes use the compact notation $\{p_j(\mathbf{z})\}_{j \in [J]}$ to denote the set of polynomial constraints. This set of constraints defines a semialgebraic set K of feasible points, which we write as

$$K = \{\mathbf{z} \in \mathbb{R}^{m+n-1}, \mid p_1(\mathbf{z}) \geq 0, \dots, p_m(\mathbf{z}) \geq 0\} \quad (1.20)$$

For $\mathbf{z} \in \mathbb{R}^{m+n-1}$ and $\boldsymbol{\alpha} \in \mathbb{N}^{m+n-1}$, we introduce the multi-index notation $\mathbf{z}^\alpha = z_1^{\alpha_1} z_2^{\alpha_2} \dots z_n^{\alpha_n}$ with $|\boldsymbol{\alpha}| = \alpha_1 + \alpha_2 + \dots + \alpha_n$, the degree of the monomial \mathbf{z}^α . We will use \mathbf{z}_B to denote the sequence of all monomials in $\mathbf{z} \in \mathbb{R}^{m+n-1}$ for some standard ordering (standard monomial basis). Hence $\mathbf{z}_B := (1, z_1, z_2, \dots, z_1^2, z_1 z_2, \dots, \mathbf{z}^\alpha, \dots)$. Similarly, let \mathbf{z}_B^t denote the vector of all monomials \mathbf{z}^α from the standard basis with degree bounded by t : $|\boldsymbol{\alpha}| \leq t$.

In this paper, polynomials will be alternatively be denoted through either of the representations below,

- Weigthed sums of monomials $p_i(\mathbf{z}) = \sum_{\boldsymbol{\alpha}} (p_i)_{\boldsymbol{\alpha}} \mathbf{z}^\alpha$, where the $(p_i)_{\boldsymbol{\alpha}}$ thus denotes the coefficient of the monomial \mathbf{z}^α in $p_i(\mathbf{z})$.
- Vectors/sequences of coefficients $p_i = ((p_i)_{\boldsymbol{\alpha}})_{\boldsymbol{\alpha} \in \mathbb{N}^n}$ as $p_i(\mathbf{z}) = p_i^T \mathbf{z}_B$
- Matrices of coefficients, P_i , such that $p_i(\mathbf{z}) = \langle P_i, \mathbf{z}_B \mathbf{z}_B^T \rangle = \mathbf{z}_B^T P_i \mathbf{z}_B$. In this case we will use the notation \mathbf{Z}_B (resp. \mathbf{Z}_B^d) to represent the matrix generated from the standard basis as $\mathbf{Z}_B = \mathbf{z}_B \mathbf{z}_B^T$ (resp. $\mathbf{Z}_B^{2d} = \mathbf{z}_B^d (\mathbf{z}_B^d)^T$). The moments matrix \mathbf{M}_0 encountered earlier is simply $\int \mathbf{Z}_B^{2d} d\mu$ when $d = 2$ and for the measure $d\mu = \delta(\mathbf{z} - \mathbf{z}_0) d\mathbf{z}$.

For the general set of polynomials $p_1(\mathbf{z}), p_2(\mathbf{z}), \dots, p_m(\mathbf{z})$, we let \mathcal{I} denote the ideal generated by those polynomials. This set is defined as from all the combinations that are generated by multiplying the polynomials $p_i(\mathbf{z})$ by any other polynomials $h_i(\mathbf{z}) \in \mathbb{R}[\mathbf{z}]$.

$$\mathcal{I}(p_1, \dots, p_m) := \left\{ \sum_{j \in [m]} p_j(\mathbf{z}) h_j(\mathbf{z}), \quad \text{for polynomials } h_j(\mathbf{z}) \in \mathbb{R}[\mathbf{z}] \right\} \quad (1.21)$$

Equivalently, we will use \mathcal{I}_d to denote the *truncated* ideal, whose maximal degree is bounded by d ,

$$\mathcal{I}_d(p_1, \dots, p_m) := \left\{ \sum_{j \in [m]} p_j(\mathbf{z}) h_j(\mathbf{z}), \quad h_j(\mathbf{z}) \in \mathbb{R}[\mathbf{z}], \deg(h_j) \leq d - \deg(p_j) \right\} \quad (1.22)$$

Given the matrix $\mathbf{Z}_B^d = \mathbf{z}_B^d (\mathbf{z}_B^d)^T$ used to represent monomials of degree at most $2d$, when writing polynomials in matrix form, we will need to access monomials of a given degree. As an example, consider the univariate monomial basis $(1, z, z^2, z^3, \dots)$. The corresponding matrix for the monomial basis truncated at degree 4 reads

$$\mathbf{Z}_B^2 = \begin{pmatrix} 1 & z & z^2 \\ z & z^2 & z^3 \\ z^2 & z^3 & z^4 \end{pmatrix} \quad (1.23)$$

Now consider the polynomial $p(z) := z^2 - 1$. This polynomial can be applied on the matrix \mathbf{Z}_B^2 by introducing appropriate matrices to access the monomials. Those matrices are simply assembled from the product of two canonical basis vectors. That is, for any degree γ , one access the monomial z^γ in \mathbf{Z} by means of the matrices $e_\alpha e_\beta^T$ for any α, β such that $\gamma = \alpha + \beta$. For any such matrix, we have

$$\langle \mathbf{Z}_B^2, e_\alpha e_\beta^T \rangle = z^\gamma.$$

In particular, using those matrices, the polynomial $p(z)$, $z \in \mathbb{R}$ reads

$$p(z) = \langle \mathbf{P}, \mathbf{Z}_B^2 \rangle = \sum_{\gamma \in \mathbb{N}^{\deg(p)}} \sum_{\alpha + \beta = \gamma} C_\gamma p_\gamma \langle \mathbf{Z}_B^2, e_\alpha e_\beta^T \rangle. \quad (1.24)$$

C_γ is a normalizing constant defined from each degree γ as

$$C_\gamma = \frac{1}{\#\{(\alpha, \beta) \mid \alpha + \beta = \gamma\}}$$

To write expression (1.24) compactly, we introduce auxiliary matrices B_γ relative to each of the monomials z^γ , defined as

$$\mathbf{B}_\gamma = \sum_{\alpha + \beta = \gamma} e_\alpha e_\beta^T. \quad (1.25)$$

Using those matrices, the polynomial $p(z)$ can now read directly as $p(z) = \sum_{\gamma \in \mathbb{N}^{\deg(p)}} C_\gamma p_\gamma \langle \mathbf{Z}_B^2, B_\gamma \rangle$. Moreover, the coefficients of this polynomial can be obtained via $p_\gamma = \langle \mathbf{P}, B_\gamma \rangle$. The constant C_γ can now also be expressed more simply as $C_\gamma = \frac{1}{\|\mathbf{B}_\gamma\|_F^2}$.

As an example, the matrix \mathbf{B}_2 used to access the monomial z^2 in (1.23) reads

$$\mathbf{B}_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad (1.26)$$

When dealing with polynomials on \mathbb{R}^{m+n-1} , the same idea applies and we will denote the corresponding matrices as \mathbf{B}_γ where γ is the multi-index used to access the monomial \mathbf{z}^γ .

2 Proof of Theorem 1.

To ensure unique recovery of the matrix \mathbf{M}_0 from the semidefinite program (1.15), traditional convex optimization proofs are based on satisfying first order optimality conditions¹ by exhibiting a dual vector λ such that $-\mathcal{A}^* \lambda - I \in \partial \iota_{\mathcal{K}}(\mathbf{M}_0)$ where $\iota_{\mathcal{K}}$ denotes the indicator function of the positive semidefinite (PSD) cone (see for example [16]). In section 2.1 below, we start by giving a more detailed characterization of problem (1.15) in terms of the constraints. We then provide the general conditions for the existence of such a certificate. In section 2.2, we show how satisfiability of these conditions can be reduced to the construction of a dual polynomial with particular structure. Section 2.3 finally shows how such a dual polynomial can be constructed.

2.1 Dual certificate

In this section, we give an explicit expression for the general condition $-\mathcal{A}^* \lambda - I \in \partial \iota_{\mathcal{K}}(\mathbf{M}_0)$ on the dual vector λ certifying optimality in the case of problem (1.15). We then show how this condition can be made

¹Note that in the case of convex optimization those conditions are necessary and sufficient.

tighter to ensure uniqueness in addition to optimality at \mathbf{M}_0 . We start by giving the detailed expression of the constraints in (1.15). Note that each matrix \mathbf{B}_γ can be decomposed into a sum of elementary matrices $\mathbf{E}_{\gamma_1, \gamma_2}$ with only a single non zero entry, for multi-indices $\gamma_1, \gamma_2 \in \mathbb{N}^{m+n-1}$, i.e., $\mathbf{B}_\gamma = \sum_{\gamma_1 + \gamma_2 = \gamma} \mathbf{E}_{\gamma_1, \gamma_2}$ with $\mathbf{E}_{\gamma_1, \gamma_2} = \mathbf{e}_{\gamma_1} \mathbf{e}_{\gamma_2}^T$.

$$\begin{aligned}
& \text{minimize} && \text{Tr}(\mathbf{M}) \\
& \text{subject to} && \sum_{\zeta} \frac{(h_\ell)_\zeta}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{M}, \mathbf{B}_{\zeta+\kappa} \rangle = 0. \\
& && \text{for } \kappa \in \mathbb{N}_{2(t-d_{h_\ell})}^K, 1 \leq \ell \leq L \\
& && \mathbf{M} \succeq 0, \quad \mathbf{M}_{11} = 1, \\
& && \langle \mathbf{M}, \mathbf{E}_{\delta_1, \delta_2} - \mathbf{E}_{\gamma_1, \gamma_2} \rangle = 0, \\
& && \text{for all } (\delta_1, \delta_2), (\gamma_1, \gamma_2) \text{ s.t. } \delta_1 + \delta_2 = \gamma_1 + \gamma_2 \leq 2t.
\end{aligned} \tag{2.1}$$

In the proof of Theorem 1, we will write the last constraint of (2.1) together with the normalization constraint $\mathbf{M}_{11} = 1$ compactly as $\mathbf{M} = \sum_{\gamma} m_\gamma \mathbf{B}_\gamma + \mathbf{e}_1 \mathbf{e}_1^T$ by introducing additional variables m_γ . This enables us to get rid of the last structural constraint in (2.1). The resulting structure of \mathbf{M} is Hankel-type (and would be exactly Hankel in the one-dimensional case as we saw earlier). The first sum in (2.1) is taken over all the coefficients of each constraint $h_\ell(x) = 0$, $h_\ell(x) = \sum_{\zeta} (h_\ell)_\zeta \mathbf{x}^\zeta$.

We now derive the first order optimality conditions $-\mathcal{A}^* \lambda - \mathbf{I} \in \partial \iota_{\mathcal{K}}(\mathbf{M}_0)$ for problem (2.1) in terms of the Lagrangian dual function \mathcal{L} . Introducing multipliers for each of the polynomial constraints, the Lagrangian can be written as

$$\begin{aligned}
\mathcal{L}(\mathbf{M}, \mathbf{m}, \boldsymbol{\lambda}, \boldsymbol{\xi}) &= \text{Tr}(\mathbf{M}) + \langle \mathbf{M} - \sum_{\gamma} m_\gamma \mathbf{B}_\gamma - \mathbf{e}_1 \mathbf{e}_1^T, \boldsymbol{\xi} \rangle \\
&+ \sum_{\ell} \sum_{\kappa \in \mathbb{N}_{2(t-d_{h_\ell})}^K} \lambda_{\ell, \kappa} \left(\sum_{\zeta} \frac{(h_\ell)_\zeta}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right) \\
&+ \iota_{\mathcal{K}}(\mathbf{M}).
\end{aligned} \tag{2.2}$$

The multipliers $\lambda_{\ell, \kappa}$ correspond to each of the original and shifted polynomial constraints while $\boldsymbol{\xi}$ encode the Hankel-type structure of the matrix \mathbf{M} . Usual convex optimization theory states that $\mathbf{M}_0 = \mathbf{m}_0 \mathbf{m}_0^T$ is a minimizer for problem (2.1) if and only if one can find dual vectors $(\boldsymbol{\xi}, \boldsymbol{\lambda})$ such that $0 \in \partial \mathcal{L}(\mathbf{M}_0, \mathbf{m}_0, \boldsymbol{\lambda}, \boldsymbol{\xi})$. The dual variables $\boldsymbol{\xi}, \boldsymbol{\lambda}$ combine into a dual certificate Z , and must obey the following three conditions. Let

$$T = \left\{ \mathbf{m}_0 \mathbf{v}^T + \mathbf{v} \mathbf{m}_0^T, \quad \mathbf{v} \in \mathbb{R}^{|\mathbb{N}_2^K|} \right\}, \tag{2.3}$$

T^\perp being its orthogonal complement, and let Y_T denote the projection of the matrix Y onto the subspace T .

- 1) $Y = \mathbf{I} - \boldsymbol{\xi} - \sum_{\ell} \sum_{\kappa \in \mathbb{N}_{2(t-d_{h_\ell})}^K} \lambda_{\ell, \kappa} \left(\sum_{\zeta} \frac{(h_\ell)_\zeta}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \mathbf{B}_{\zeta+\kappa} \right)$
- 2) $Y_T = 0, \quad Y_{T^\perp} \succeq 0$
- 3) $\langle \mathbf{B}_\gamma, \boldsymbol{\xi} \rangle = 0, \quad \forall \gamma \neq 0.$

Conditions 1 and 2 are obtained by requiring the derivative of this Lagrangian with respect to the moments matrix \mathbf{M} belongs to the normal cone (subdifferential of the indicator of the PSD cone) at \mathbf{M}_0 , and Condition 3 is obtained by requiring that the derivative of the Lagrangian with respect to the vector of moments, \mathbf{m} , vanishes.

The following proposition guarantees unique recovery in addition to the optimality ensured by the satisfiability of conditions 1) to 3).

Proposition 1. *To ensure unique recovery of \mathbf{M}_0 , in addition to the conditions 1), 2), and 3) mentioned above, it is sufficient to require $Y_{T^\perp} \succ 0$ as well as injectivity on T of all the linear constraints $\mathcal{A}(\mathbf{M}) = \mathbf{b}$ arising from the measure version of the polynomial constraints $h_\ell(\mathbf{z}) = 0$ as well as from the structure of the moments matrix.*

Proof. We will now use the decompositions $Y = \mathbf{I} - Y_2 = I - Y_1 - \xi$

$$\text{Tr}(\mathbf{M}_0) = \langle I, \mathbf{M}_0 \rangle = \langle I_T, \mathbf{M}_0 \rangle = \langle (Y_2)_T, \mathbf{M}_0 \rangle \quad (2.4)$$

$$= \langle Y_2, \mathbf{M}_0 - \mathbf{M} \rangle + \langle Y_2, \mathbf{M} \rangle = \langle Y_2, \mathbf{M} \rangle \quad (2.5)$$

$$= \langle I_T, \mathbf{M}_T \rangle + \langle (Y_2)_{T^\perp}, \mathbf{M} \rangle \quad (2.6)$$

$$= \text{Tr}(\mathbf{M}_T) + \langle (Y_2)_{T^\perp}, \mathbf{M} \rangle \quad (2.7)$$

$$< \text{Tr}(\mathbf{M}) \quad \text{for } \mathbf{M}_{T^\perp} \neq 0 \quad (2.8)$$

In (2.5), we use $\langle \xi, \mathbf{M} - \mathbf{M}_0 \rangle = 0$ as well as the fact that Y_1 belongs to the range of \mathcal{A}^* and $\mathcal{A}(\mathbf{M}) = \mathcal{A}(\mathbf{M}_0)$. The last inequality follows from $(Y_2)_{T^\perp} \prec I_{T^\perp}$ which since $\mathbf{M} \succeq 0$ implies $\langle (Y_2)_{T^\perp}, \mathbf{M} \rangle < \text{Tr}(\mathbf{M}_{T^\perp})$ for $\mathbf{M}_{T^\perp} \neq 0$. This last inequality thus implies $\mathbf{M}_{T^\perp} = 0$. Finally $\mathbf{M}_T = (\mathbf{M}_0)_T$ by injectivity of the constraints on T . \square

Note that, to satisfy $Y_{T^\perp} \succ 0$ and $Y_T = 0$, it is sufficient to ask for $\mathbf{m}_0 \in \text{Null}(Y)$ and to require Y to be positive semidefinite and exact rank $|\mathbb{N}_2^K| - 1$. In the next section, we show how the duality between sum-of-squares polynomials and positive semidefinite matrices can help us construct a dual certificate satisfying those conditions.

2.2 Sum-of-squares and positive semidefinite matrices

We call sum-of-squares (SOS) polynomial, any polynomial $p(z)$ for which there exists a decomposition $p(z) = \sum_{j=1}^m s_j^2(z)$ for some polynomials $s_j \in \mathbb{R}[\mathbf{z}]$. Introducing a polynomial version of proposition 1 requires the following lemma from [37] relating SOS and semidefinite programming (SDP). For completeness we also provide a proof.

Proposition 2 (Equivalence between SOS and SDP). *Let \mathbb{N}_{2t}^K denote the set of K -tuples $\alpha \in \mathbb{N}^K$ such that $\sum_i \alpha_i \leq 2t$ and let $\mathbf{x}^\alpha = x_1^{\alpha_1} x_2^{\alpha_2} \dots x_K^{\alpha_K}$. Let $p(\mathbf{z}) \in \mathbb{R}[\mathbf{z}]$ with $p(\mathbf{z}) = \sum_{\alpha \in \mathbb{N}_{2t}^K} p_\alpha \mathbf{x}^\alpha$ be a polynomial of degree $\leq 2t$, the following assertions are equivalent,*

1) $p(\mathbf{z})$ is a sum-of-squares polynomial

2) There exists a positive semidefinite matrix \mathbf{A} such that

$$p(\mathbf{z}) = \mathbf{z}_B^T \mathbf{A} \mathbf{z}_B, \quad (2.9)$$

Proof. If $p(\mathbf{x})$ is SOS then $p(\mathbf{x}) = \sum_j s_j^2(\mathbf{x})$ for some polynomial $s_j(\mathbf{x})$. Let $d = \max\{\deg(s_j)\}$ denote the maximum degree of the $s_j(\mathbf{x})$. Further let $\delta = \lceil d/2 \rceil$. For each of the $s_j(\mathbf{x})$, for some ordering of the monomials, construct the corresponding vector of coefficients $\mathbf{s}_j \in \mathbb{R}^\delta$ with $s_j(\mathbf{x}) = \mathbf{s}_j^T \mathbf{x}_B$, then the positive semidefinite matrix $\mathbf{A} = \sum_j \mathbf{s}_j \mathbf{s}_j^T$ satisfies $\mathbf{x}^T \mathbf{A} \mathbf{x} = p(\mathbf{x}) \geq 0$ for all \mathbf{x} . Conversely, let \mathbf{A} be a matrix such that $p(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x}$. Since $\mathbf{A} \succeq 0$, it has the spectral decomposition $\mathbf{A} = \sum_j \mu_j \mathbf{v}_j \mathbf{v}_j^T$ for some $\mu_j \geq 0$. Then we write $p(\mathbf{x}) = \sum_j \mu_j (\mathbf{x}^T \mathbf{v}_j)^2$, which is a sum of squares. □

It is important to notice that proposition 2 doesn't provide a strict equivalence between a matrix certificate and a polynomial certificate. Observe that the existence of a sum-of-squares polynomial $p(\mathbf{x})$ such that $p(\mathbf{x}) = \mu_{\mathbf{x}}^T \mathbf{A} \mu_{\mathbf{x}}$ doesn't imply that \mathbf{A} is positive semidefinite. In other words, not all matrices encoding sum-of-squares polynomials are PSD. As an illustration, consider the following example:

Example 1 (sum-of-squares and positive semidefiniteness).

$$\mathbf{A} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 0 & 0 & 1/2 \\ 0 & 0 & 0 \\ 1/2 & 0 & 0 \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 0 & 0 & 1/4 \\ 0 & 1/2 & 0 \\ 1/4 & 0 & 0 \end{pmatrix}.$$

For the vector of monomials $\mathbf{z}_B = (1 \quad z \quad z^2)$. All those matrices are encoding the same SOS polynomial $p(z)$

$$\mathbf{z}_B^T \mathbf{A} \mathbf{z}_B = \mathbf{z}_B^T \mathbf{B} \mathbf{z}_B = \mathbf{z}_B^T \mathbf{C} \mathbf{z}_B = p(z) = z^2.$$

However, only one of them is positive semidefinite. The second and third ones can therefore not be used as the matrix form of a SOS-type certificate. However, note that there exists a matrix $\boldsymbol{\xi}$ such that $\langle \boldsymbol{\xi}, \mathbf{B}_\gamma \rangle = 0$ for all γ , satisfying $\mathbf{C} + \boldsymbol{\xi} = \mathbf{A}$ or equivalently $\mathbf{B} + \boldsymbol{\xi} = \mathbf{A}$. Indeed, for \mathbf{C} it suffices to take

$$\boldsymbol{\xi} = \begin{pmatrix} 0 & 0 & -1/4 \\ 0 & 1/2 & 0 \\ -1/4 & 0 & 0 \end{pmatrix}$$

This is the point of the following lemma which formalizes and closes the gap between matrix and polynomial certificate.

The following lemma proves equivalence of the matrix certificates up to a $\boldsymbol{\xi}$ provided that the corresponding polynomials are the same.

Lemma 2. Let Y_1 and Y_2 be two matrices such that $\mathbf{z}_B^T Y_1 \mathbf{z}_B = \mathbf{z}_B^T Y_2 \mathbf{z}_B$ for all z , i.e., the polynomials corresponding to Y_1 and Y_2 are identical. Then there exists a matrix $\boldsymbol{\xi}$ with $\langle \boldsymbol{\xi}, \mathbf{B}_\gamma \rangle = 0$ for all γ and such that $Y_1 = Y_2 + \boldsymbol{\xi}$.

Proof. $\mathbf{z}_B^T (Y_1 - Y_2) \mathbf{z}_B = 0 \quad \forall z \Leftrightarrow \langle Y_1 - Y_2, \mathbf{z}_B \mathbf{z}_B^T \rangle = 0 \quad \forall z \Rightarrow \langle Y_1 - Y_2, \mathbf{B}_\gamma \rangle = 0 \quad \forall \gamma$. This last implication holds in the reverse direction: if a polynomial $p(\mathbf{z})$ has all zero coefficients, then it must be the zero polynomial. □

The conditions of proposition 1, together with proposition 2 and lemma 2 imply the following result, arising from the polynomial nature of problem (1.1),

Proposition 3 (Polynomial Form). To ensure unique recovery of $\mathbf{M}_0 = \mathbf{m}_0 \mathbf{m}_0^T$ with $(\mathbf{m}_0)_\gamma = \mathbf{z}_0^\gamma$, in addition to the injectivity of the constraints on T , it is sufficient to find a sum of $(|\mathbb{N}_2^K| - 1)$ linearly

independent squares $s_j^2(z)$ of degree less than or equal to 4, polynomials $\lambda_\ell(z)$ of degree less than or equal to $4 - 2d_{h_\ell}$ and constant ρ such that

$$q(z) = \sum_j s_j^2(z) = \sum_{\alpha \in \mathbb{N}_2^K} \mathbf{z}^{2\alpha} - \rho + \sum_\ell h_\ell(z) \lambda_\ell(z), \quad (2.10)$$

and such that $q(z_0) = 0$.

Proof. The form of the polynomial $q(z)$ in (2.10) implies the existence of a matrix Y_1 in the range of \mathcal{A}^* such that $\mathbf{z}_B^T(I - Y_1)\mathbf{z}_B = \sum_j s_j^2(z)$. By lemma 2, we can then add a matrix $\boldsymbol{\xi}$ to $I - Y_1$ to get the positive semidefinite matrix $\sum_j \mathbf{s}_j \mathbf{s}_j^T = I - Y_2 = I - Y_1 - \boldsymbol{\xi} \succeq 0$ which now satisfies the condition $Y_T^\perp \succeq 0$ of section 2.1. Note that such a $\boldsymbol{\xi}$ always exists, by lemma 2, as we have

$$q(z) = \mathbf{z}_B^T \sum_j \mathbf{s}_j \mathbf{s}_j^T \mathbf{z}_B = \mathbf{z}_B^T (I - Y_2) \mathbf{z}_B = \mathbf{z}_B^T (I - Y_1 - \boldsymbol{\xi}) \mathbf{z}_B = \mathbf{z}_B^T (I - Y_1) \mathbf{z}_B$$

Finally, as indicated by proposition 3, since $q(z)$ is SOS, to satisfy the last condition, $Y_T = 0$, it suffices to require $s_j(z_0) = 0$. Indeed, for $(\mathbf{m}_0)_\gamma = \mathbf{z}_0^\gamma$, we have

$$\{s_j(z_0) = 0, \forall j\} \iff \left\{ \left\langle \sum_j \mathbf{s}_j \mathbf{s}_j^T, \mathbf{m}_0 \mathbf{y}^T + \mathbf{y} \mathbf{m}_0^T \right\rangle, \forall \mathbf{y} \in \mathbb{R}^{|\mathbb{N}_2^K|} \right\} = 0.$$

The value of the constant ρ , which derives from the one degree of freedom of \mathbf{B}_0 , is fixed by enforcing $q(z_0) = 0$. The last term on the RHS of (2.10) is a contribution of degree ≤ 4 from the ideal $\mathcal{I} := \{\sum_{j=1}^L u_j(z) h_j(z) \mid u_1, \dots, u_L \in \mathbb{R}[\mathbf{z}]\}$ generated from the constraints $h_j(z)$. \square

Because of proposition 2 and lemma 2, we can now just focus on finding a (dual) polynomial $q(z)$ with the structure (2.10).

2.3 Construction of the dual polynomial

In this section we show how to construct the dual polynomial satisfying the decomposition (2.10). As explained above, such a polynomial implies the existence of a matrix Y satisfying the conditions 1) to 3) and proposition 1 and serves as the first part of the proof of Theorem 1. We then prove injectivity on T to conclude this proof.

Remember that \mathbf{z} is given by the concatenation $\mathbf{z} = (\mathbf{x}, \mathbf{y})$ of all first order monomials arising in problem (1.1). Our construction of the certificate is based on choosing the squares on the LHS of (2.10) to be the canonical polynomials $(\mathbf{z}^\alpha - \mathbf{z}_0^\alpha)^2$ for all $|\alpha| \leq 2$ and to show that those canonical squares can be obtained from the ideal; the squared monomials arising from the trace norm and the constant $\rho = \sum_\alpha \mathbf{z}_0^{2\alpha}$. The resulting expression for the certificate is simply

$$q(z) = \sum_{|\gamma| \leq 2} (\mathbf{z}^\gamma - \mathbf{z}_0^\gamma)^2.$$

First, let us show that for all monomials \mathbf{z}^α with $|\alpha| = 1$ one can build the polynomial $-2\mathbf{z}^\alpha \mathbf{z}_0^\alpha + 2(\mathbf{z}_0^\alpha)^2$ by using a decomposition from the ideal of degree at most 3.

- Either the constraint $\mathbf{z}^\alpha = \mathbf{z}_0^\alpha$ is present explicitly ($\mathbf{z}^\alpha = y_\ell$ corresponds to an element of the first row of \mathbf{X} and $h_\ell(z) \equiv y_\ell - (y_0)_\ell$ is a constraint in Ω) and one can then just multiply this constraint by $-2(\mathbf{z}_0^\alpha)$ to get the desired polynomial $-2(\mathbf{z}_0)^\alpha \mathbf{z}^\alpha + 2(\mathbf{z}_0^\alpha)^2$

- Or, since the bipartite graph is connected, the first order monomial \mathbf{z}^α , $|\alpha| = 1$, appears in a chain like (1.4), such that if we denote the corresponding numerical values by $(z_0)_{i_1}, (z_0)_{i_1}(z_0)_{i_2}, \dots, (z_0)_{i_{\ell-1}}(z_0)_\ell$, the constraints $z_{i_1} - (z_0)_{i_1}, \dots, z_{i_{\ell-1}}z_\ell - (z_0)_\ell(z_0)_{i_{\ell-1}}$ belong to Ω and thus to the ideal \mathcal{I} . Using (1.4), one can thus recursively combine the elements of the chain in the following way,

$$\begin{aligned} (z_0)_{i_{\ell-2}}(z_0)_{i_{\ell-1}}(z_\ell - (z_0)_\ell) &= (z_\ell z_{i_{\ell-1}} - (z_0)_\ell(z_0)_{i_{\ell-1}})z_{i_{\ell-2}} \\ &\quad - (z_{i_{\ell-2}}z_{i_{\ell-1}} - (z_0)_{i_{\ell-2}}(z_0)_{i_{\ell-1}})z_\ell \\ &\quad + (z_0)_\ell(z_0)_{i_{\ell-1}}(z_{i_{\ell-2}} - (z_0)_{i_{\ell-2}}). \end{aligned} \quad (2.11)$$

This telescoping relation holds for all ℓ throughout the chain until the second element, (z_{i_2}) , for which we have $(z_0)_{i_1}(z_{i_2} - (z_0)_{i_2}) = (z_{i_2}z_{i_1} - (z_0)_{i_2}(z_0)_{i_1}) - z_{i_2}(z_{i_1} - (z_0)_{i_1}) \in \mathcal{I}$. The key here is that one can make use of the bilinear constraints to get a propagation argument which remains degree-3 since the multiplicative factor $(z_0)_\ell(z_0)_{i_{\ell-1}}$ in front of the propagation term $(z_{i_{\ell-2}} - (z_0)_{i_{\ell-2}})$ remains constant. In particular, note that we never use the third order constraints $z^\alpha(z_{i_1} - (z_0)_{i_1})$ for $|\alpha| = 2$, namely the highest degree of the monomials multiplying the first order constraints is one. This will be important later when establishing the stability result.

Now that we can build the polynomials $-2(z_0)_k z_k + 2(z_0)_k^2$ for all k as degree-3 decompositions from the ideal \mathcal{I} , one can just add those polynomials to the trace and constant ρ contributions $z_k^2 - (z_0)_k^2$ in order to get the squares $(z_k - (z_0)_k)^2$. We thus get $|\mathbb{N}_1^K| - 1$ of the required squares. The remaining $\binom{K}{2}$ decompositions for the second order squared polynomials $(\mathbf{z}^\alpha - \mathbf{z}_0^\alpha)^2$ for $|\alpha| = 2$, are built from the first order decompositions, the trace, and constant ρ as follows. $\forall \alpha, \beta$ with $|\alpha|, |\beta| = 1$,

$$\begin{aligned} (\mathbf{z}^\alpha \mathbf{z}^\beta - \mathbf{z}_0^\alpha \mathbf{z}_0^\beta)^2 &= (\mathbf{z}^\alpha \mathbf{z}^\beta)^2 - (\mathbf{z}_0^\alpha \mathbf{z}_0^\beta)^2 \\ &\quad - 2\mathbf{z}_0^\alpha \mathbf{z}_0^\beta (\mathbf{z}^\alpha \mathbf{z}^\beta - \mathbf{z}_0^\alpha \mathbf{z}_0^\beta), \end{aligned} \quad (2.12)$$

where the first two terms arise from the contribution of the trace and ρ , and the third one can be expressed from the ideal \mathcal{I} with degree at most 4, as

$$\begin{aligned} -2\mathbf{z}_0^\alpha \mathbf{z}_0^\beta (\mathbf{z}^\alpha \mathbf{z}^\beta - \mathbf{z}_0^\alpha \mathbf{z}_0^\beta) &= (-2\mathbf{z}^\alpha \mathbf{z}_0^\alpha + 2(\mathbf{z}_0^\alpha)^2)(\mathbf{z}_0^\beta)^2 \\ &\quad + (\mathbf{z}^\beta - \mathbf{z}_0^\beta)(-2\mathbf{z}^\alpha \mathbf{z}_0^\alpha) \mathbf{z}_0^\beta \end{aligned} \quad (2.13)$$

The first term is of degree at most 3 and the second one is of degree at most 4.

To conclude the proof of Theorem 1, we show that the linear map \mathcal{A} grouping the linear constraints derived from the polynomials h_ℓ and the structure of the moments matrix, is injective on T . For this purpose, let us show that the nullspace of \mathcal{A} is empty on T . Let us consider any $\mathbf{H} = \mathbf{m}_0 \mathbf{v}^T + \mathbf{v} \mathbf{m}_0^T$. Normalization of \mathbf{H}_{11} implies $v_1 = 0$ and reduces \mathbf{H} to a matrix for which the first column equals the first row and is given by $(v_2 \dots, v_{|\mathbb{N}_2^K|})$. Then recall that there is a least one constraint setting to zero one of the elements of the first column. So there exists ℓ s.t. $v_\ell = 0$. Accordingly the whole corresponding row and column reduce to $((z_0)_\ell v_k)_{k \leq |\mathbb{N}_2^K|}$. Since $(z_0)_\ell \neq 0^2$, one can then apply the next constraint $z_\ell z_m = 0$ which implies $v_m = 0$. By recursively applying this idea, one can show that the first block of \mathbf{H} corresponding to the monomials of degree at most two is zero. The remaining part of the matrix can then be set to 0 as well through the structural constraints (equality of corresponding monomials) for the first row/column and then using the fact that \mathbf{H} is defined as $\mathbf{m}_0 \mathbf{v}^T + \mathbf{v} \mathbf{m}_0^T$.

²Recall that we assumed $(\mathbf{X}_0)_{ij} \neq 0$ for all (i, j)

3 Stability

In this section, we prove Theorem 2 and Corollary 3. We let the noisy measurements be given by $(\tilde{X}_0)_{ij} = (X_0)_{ij} + \varepsilon_{ij}$ for $(i, j) \in \Omega$. We further let \tilde{h}_ℓ denote the corresponding noisy constraints. If $h_\ell(\mathbf{z}) := \mathbf{z}^\alpha - \mathbf{z}_0^\alpha$ denotes a constraint in Ω with either $\mathbf{z}^\alpha = x_{i+1}y_j$ and $(i, j) \in \Omega$ or $\mathbf{z}^\alpha = y_\ell$, $(1, \ell) \in \Omega$, we let $\tilde{h}_\ell(\mathbf{z}) := \mathbf{z}^\alpha - \tilde{\mathbf{z}}_0^\alpha$ denote the corresponding noisy constraint with $\tilde{\mathbf{z}}_0^\alpha = \mathbf{z}_0^\alpha + \varepsilon_{ij}$. Hence, $\tilde{h}_\ell(\mathbf{z}) = h_\ell(\mathbf{z}) - \varepsilon_{ij}$ with (i, j) relative to the constraint indexed by ℓ , or with a slight abuse of notation, $\tilde{h}_\ell(\mathbf{z}) = h_\ell(\mathbf{z}) - \varepsilon_\ell$.

Let $\eta \geq \|\varepsilon\|_2 \sqrt{(1 + \|\mathbf{z}_0\|_1^2)}$ with $\|\varepsilon\|_2 = \sqrt{\sum_{(i,j) \in \Omega} \varepsilon_{ij}^2}$. The stable version of (2.1) reads,

$$\begin{aligned}
& \text{minimize} && \text{Tr}(\mathbf{M}) \\
& \text{subject to} && \sqrt{\sum_{\kappa} \sum_{\ell} \left| \sum_{\zeta} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{M}, \mathbf{B}_{\zeta+\kappa} \rangle \right|^2} \leq \eta. \\
& && \text{for } \kappa \in \mathbb{N}_{2(t-d_{h_\ell})}^K, 1 \leq \ell \leq L \\
& && \mathbf{M} \succeq 0, \quad \mathbf{M}_{11} = 1, \\
& && \langle \mathbf{M}, \mathbf{E}_{\delta_1, \delta_2} - \mathbf{E}_{\gamma_1, \gamma_2} \rangle = 0, \\
& && \text{for all } (\delta_1, \delta_2), (\gamma_1, \gamma_2) \text{ s.t. } \delta_1 + \delta_2 = \gamma_1 + \gamma_2 \leq 2t.
\end{aligned} \tag{3.1}$$

The first constraint in formulation (3.1) is simply the ℓ_2 norm of the constraints appearing in (2.1). In this first constraint, the first sum is taken over the different noisy polynomials h_ℓ and the second is taken over all the ‘‘shifts’’ of those polynomials. For a given κ , the corresponding shifted polynomial is simply obtained by multiplying h_ℓ by the corresponding monomial \mathbf{z}^κ .

It is worth pointing out that formulation (3.1) is not unit-independent, since the moment matrix M mixes different powers of the original variables. This can be remedied by assigning dimensional weights w_γ to the B_γ matrices – an operation that modifies the numerics and the theory in an obvious way. Formulation (3.1) leads to the recovery result of Theorem 2 which is restated below for clarity.

Theorem 2. *Let \mathbf{M}_0 denote the rank one matrix introduced in (1.14) for $\mathbf{z}_0 \in \mathbb{R}^{m+n-1}$. Assume that the necessary and sufficient conditions of lemma 1 are satisfied. Let $\|\varepsilon\|_2 = \sqrt{\sum_{(i,j) \in \Omega} \varepsilon_{ij}^2}$. Let \mathbf{M} denote the solution to the semidefinite program (1.16). This solution satisfies*

$$\frac{\|\mathbf{M} - \mathbf{M}_0\|_F}{\|\mathbf{M}_0\|_F} \leq C_0(m+n)^{7/2} \|\varepsilon\|_2. \tag{1.17}$$

The constant C_0 depends on the entries of \mathbf{X}_0 , but not on Ω , m, n , or ε .

The stability result of Theorem 2 can be improved if a path is known that relates one entry to all the others. In this last case, the scalings can be reduced from $\mathcal{O}((m+n)^{7/2})$ to $\mathcal{O}((m+n)^2)$. This is the point of Corollary 3 which is proved in section 3.2,

Corollary 3. *Assume that the paths in the bipartite graph relating each of the unknown vertices x_i (resp. y_j) to the root node y_ℓ are explicitly given. Then we have the following stability estimate,*

$$\frac{\|\mathbf{M} - \mathbf{M}_0\|_F}{\|\mathbf{M}_0\|_F} \leq C_0(m+n)^2 \|\varepsilon\|_2. \tag{1.18}$$

The constant C_0 depends on the entries of \mathbf{X}_0 , but not on Ω , m, n , or ε .

Let $\mathcal{P}_1, \dots, \mathcal{P}_P$ denote the sets of constraints that appear along each path between the root node and the leaf nodes in the sense of (1.4). As explained in section 2.3, the dual certificate only relies on the monomials appearing along each of the paths multiplied either by the previous missing variable or the next one. For each path \mathcal{P}_i , let \mathcal{K}_i denote the subset of multi-indices corresponding to variables that are multiplying the constraints in the chain in the expression of the certificate (2.11). Let $\eta' \geq \|\varepsilon\|_2 \sqrt{1 + \sup_{|\alpha| \leq 1} \mathbf{z}_0^{2\alpha}}$. The formulation for Corollary 3 is obtained by replacing the ℓ_2 constraint in (3.1) by a corresponding ℓ_2 term minimizing the noise along the paths,

$$\sqrt{\sum_{\kappa \in \mathcal{K}_i} \sum_{\ell \in \mathcal{P}_i} \left| \sum_{\zeta} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{M}, \mathbf{B}_{\zeta+\kappa} \rangle \right|^2} \leq \eta', \quad i = 1, \dots, P. \quad (3.2)$$

The improvement in the prefactors of Theorem 2 essentially arises from the tighter bound η' on the ℓ_2 constraints in (3.2). This tighter bound is due to the fact that along a given path, the constraints are always distinct and that following the discussion in section (2.3), one can express each first order monomial that appear in the path from the constraints along the path multiplied by either the first or the previous or next missing degree one monomial. This idea is expressed through section 3.2.

3.1 Proof of Theorem 2

Let $\|\mathbf{M}\|_p$ to denote the Schatten p -norm of \mathbf{M} ,

$$\|\mathbf{M}\|_p = \left(\sum_k \sigma_k^p \right)^{1/p}.$$

We therefore have $\|\mathbf{M}\|_1 = \|\mathbf{M}\|_*$ which denotes the nuclear norm of \mathbf{M} , $\|\mathbf{M}\|_2 = \|\mathbf{M}\|_F$ which is used to denote the Frobenius norm of \mathbf{M} and $\|\mathbf{M}\|_\infty = \|\mathbf{M}\|$ which denotes the operator norm of \mathbf{M} .

Any solution \mathbf{M} to (3.1) reads $\mathbf{M} = \mathbf{H} + \mathbf{M}_0$. To prove stability of the recovery, we first highlight the following,

- $\text{Tr}(\mathbf{M}_0 + \mathbf{H}) \leq \text{Tr}(\mathbf{M}_0)$ and therefore $\text{Tr}(\mathbf{H}) \leq 0$.
- Both \mathbf{M} and \mathbf{M}_0 are feasible points for (3.1), and hence both satisfy the normalization constraint $\mathbf{M}_{11} = (\mathbf{M}_0)_{11}$ which can be exactly enforced. As a consequence, $\mathbf{H}_{11} = (\mathbf{M}_0)_{11} - \mathbf{M}_{11} = 0$, and all degree zero terms in the constraints $\tilde{h}_1, \dots, \tilde{h}_L$ vanish when those constraints are applied to \mathbf{H} . We have

$$\frac{(\tilde{h}_\ell)_0}{\|\mathbf{B}_0\|_F^2} \langle \mathbf{H}, \mathbf{B}_0 \rangle = 0, \quad \ell = 1, \dots, L. \quad (3.3)$$

More generally, for both \mathbf{M} and \mathbf{M}_0 , as η is bounding the vector $(\varepsilon_\ell \mathbf{z}_0^\kappa)_{\ell, |\kappa| \leq 2}$ of weighted residuals, we must have,

$$\begin{aligned} \sum_{\ell, \kappa} \left| \sum_{\zeta} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M}_0 \rangle \right|^2 &= \sum_{\ell, \kappa} \left| \sum_{\zeta} (h_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M}_0 \rangle - \varepsilon_\ell \mathbf{z}_0^\kappa \right|^2 \leq \sum_{\ell, \kappa} (\varepsilon_\ell \mathbf{z}_0^\kappa)^2 \leq \eta^2 \\ \sum_{\ell, \kappa} \left| \sum_{\zeta} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right|^2 &\leq \sum_{\ell, \kappa} (\varepsilon_\ell \mathbf{z}_0^\kappa)^2 \leq \eta^2 \end{aligned}$$

From those relations we can derive a similar bound on \mathbf{H} ,

$$\sqrt{\sum_{\kappa} \sum_{\ell} \left| \sum_{\zeta} \frac{(\tilde{h}_{\ell})_{\zeta}}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{H}, \mathbf{B}_{\zeta+\kappa} \rangle \right|^2} = \sqrt{\sum_{\kappa} \sum_{\ell} \left| \sum_{\zeta} \frac{(\tilde{h}_{\ell})_{\zeta}}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{M} - \mathbf{M}_0, \mathbf{B}_{\zeta+\kappa} \rangle \right|^2} \leq 2\eta.$$

- Finally, note that $\mathbf{H} + \mathbf{M}_0 \succeq 0$ implies $\langle \mathbf{H} + \mathbf{M}_0, \mathbf{W} \rangle \geq 0$ for all $\mathbf{W} \succeq 0$ including all $\mathbf{W} \in T^{\perp}$ which implies $\mathbf{H}_{T^{\perp}} \succeq 0$.

The polynomial form of \mathbf{Y}_2 belongs to the range of \mathcal{A}^* (i.e, its polynomial form belongs to the ideal \mathcal{I}) modulo a $\boldsymbol{\xi}$ and \mathbf{Y}_2 is written as

$$\mathbf{Y}_2 = \boldsymbol{\xi} + \sum_{\ell} \sum_{\kappa \in \mathbb{N}_{2(t-d_{h_{\ell}})}^{\kappa}} \lambda_{\ell\kappa} \left(\sum_{\zeta} \frac{(h_{\ell})_{\zeta}}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \mathbf{B}_{\zeta+\kappa} \right), \quad (3.4)$$

where $\boldsymbol{\xi}$ is orthogonal to the \mathbf{B}_{γ} .

The next section derives a bound on $\|\mathbf{H}_{T^{\perp}}^{\perp}\|$. For this, we start by bounding $|\langle \mathbf{H}, \mathbf{Y}_1 \rangle|$. Note that $|\langle \mathbf{Y}_2, \mathbf{H} \rangle| = |\langle \mathbf{Y}_1, \mathbf{H} \rangle|$, as $\langle \boldsymbol{\xi}, \mathbf{H} \rangle = 0$.

3.1.1 Bound on $\mathbf{H}_{T^{\perp}}^{\perp}$

The certificate (3.4) is built from the noiseless constraints h_{ℓ} , while the solutions \mathbf{M} , \mathbf{M}_0 and thus \mathbf{H} are bounded with respect to the corrupted constraints from (3.1). As we saw earlier, the noisy constraints relate to the noiseless constraints as

$$\tilde{h}_{\ell}(z) = \sum_{\zeta} (\tilde{h}_{\ell})_{\zeta} \mathbf{z}^{\zeta} = \sum_{\zeta \neq 0} (h_{\ell})_{\zeta} \mathbf{z}^{\zeta} + (h_{\ell})_0 - \varepsilon_{\ell} = h_{\ell}(z) - \varepsilon_{\ell}. \quad (3.5)$$

Let $\mathbf{Y}_1^{(1)}$ and $\mathbf{Y}_1^{(2)}$ denote the contributions to \mathbf{Y}_1 corresponding to the first and second order squares in the sos certificate of section 2.3 respectively. For any constraint h_{ℓ} , using the recursion (2.11), the difference between h_{ℓ} and \tilde{h}_{ℓ} will only affect the entries in \mathbf{H} corresponding to first and zero order moments. Since $\mathbf{H}_{1,1} = 0$ (see the discussion above), this discrepancy will thus only affect first order entries. Let W_{ℓ} denote the number of times that each constraint is used in the construction of $\mathbf{Y}_1^{(1)}$

$$|\langle \mathbf{Y}_1^{(1)}, \mathbf{H} \rangle| = \left| \sum_{\kappa} \sum_{\ell} W_{\ell,\kappa} \sum_{\zeta} (h_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{H} \rangle \right| \quad (3.6)$$

$$= \left| \sum_{\kappa} \sum_{\ell} W_{\ell,\kappa} \sum_{\zeta} (\tilde{h}_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle - \sum_{|\kappa| \leq 1} \sum_{\ell} W_{\ell,\kappa} \varepsilon_{\ell} \langle \mathbf{B}_{\kappa}, \mathbf{M} \rangle \right| \quad (3.7)$$

$$|\langle \mathbf{Y}_1^{(1)}, \mathbf{H} \rangle| \leq \left| \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (\tilde{h}_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| + \left| \sum_{|\kappa| \leq 1} \sum_{\ell} W_{\ell, \kappa} \varepsilon_{\ell} \langle \mathbf{B}_{\kappa}, \mathbf{M} \rangle \right| \quad (3.8)$$

$$\leq \left| \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (\tilde{h}_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| + \mathcal{O}((m+n)^{3/2}) \|\mathbf{M}_{11}\|_F \|\varepsilon\|_{\infty} \quad (3.9)$$

$$\leq \left| \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (\tilde{h}_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| + \mathcal{O}((m+n)^{3/2}) \|\mathbf{m}_0\| \|\varepsilon\|_{\infty} \quad (3.10)$$

$$\leq \mathcal{O}((m+n)^{3/2}) \eta + \mathcal{O}((m+n)^{3/2}) \|\mathbf{m}_0\| \|\varepsilon\|_{\infty} \quad (3.11)$$

In (3.11) we use the fact that $\mathbf{Y}_1^{(1)} \in \mathcal{Ran}(\mathcal{A}^*)$ and $\{\mathcal{A}(\mathbf{M}_0)\}_{\ell \neq 0} = 0$ (the first constraint is simply $(\mathbf{M}_0)_{11} = 1$ and does not appear in \mathbf{Y}_1). Let us introduce the following decomposition for \mathbf{M}/\mathbf{M}_0 ,

$$\mathbf{M} = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}. \quad (3.12)$$

For both \mathbf{M} and \mathbf{M}_0 , because of the structural constraints and PSD constraint, one can write $\text{Tr}(\mathbf{M}) \geq \|\mathbf{M}_{11}\|_F^2 + \text{Tr}(\mathbf{M}_{11})$. Moreover, we have $\text{Tr}(\mathbf{M}_0) \geq \text{Tr}(\mathbf{M})$ so in particular, we have

$$\|\mathbf{M}_{11}\|_F \leq \sqrt{\text{Tr}(\mathbf{M}) - \text{Tr}(\mathbf{M}_{11})} \leq \sqrt{\text{Tr}(\mathbf{M})} \leq \sqrt{\text{Tr}(\mathbf{M}_0)} = \|\mathbf{m}_0\|.$$

In (3.8), since all the \mathbf{B}_{κ} are accessing moments of order at most one in \mathbf{M} , the sum $\sum_{\ell} \sum_{\kappa} \mathbf{B}_{\kappa}$ in the second term of (3.8) has the form $v e_1^* + e_1 v^*$ where $v_i \leq \mathcal{O}(m+n) \|\varepsilon\|_{\infty}$ and the norm $\|\sum_{\ell} \sum_{\kappa} \mathbf{B}_{\kappa}\|_F$ can thus be bounded as $\|\sum_{\ell} \sum_{\kappa} \mathbf{B}_{\kappa}\|_F = \mathcal{O}((m+n)^{3/2})$. We can also replace \mathbf{M} by $\mathbf{M}_{1,1}$. Equations (3.9) and (3.10) then follow from Cauchy-Schwarz. For (3.11) simply note that

- In the sum, every constraint appears at most $\mathcal{O}(m+n)$ times (as an example, the first constraint $z_{i_1} - (z_0)_{i_1} = 0$ will appear exactly $m+n$ times as it is used to express every square in the chain), i.e, $W_{\ell, \kappa} = \mathcal{O}(m+n)$ if κ corresponds to either of the two monomials multiplying the constraints in (2.11) and 0 otherwise.
- Each of the constraints is multiplied by at most two different monomials leading to two distinct entries in the vector $(\varepsilon_{\ell} \mathbf{z}_0^{\alpha})_{\ell, |\alpha| \leq 2}$ whose ℓ_2 norm is bounded by η .

Equation (3.11) follows from Cauchy-Schwarz, noting that the first term in (3.9) can be written as $|\langle \tilde{\mathcal{A}}(\mathbf{M}), \boldsymbol{\lambda} \rangle|$ where $\|\boldsymbol{\lambda}\|_{\infty} = \mathcal{O}(m+n)$ and the ℓ_2 norm of $\tilde{\mathcal{A}}(\mathbf{H})$ is bounded from the constraints in (3.1).

We now bound the second order contributions gathered in $\mathbf{Y}_1^{(2)}$. From (2.12), this contribution can be decomposed as $\mathbf{Y}_2^{(2)} = \mathbf{V}_1 + \mathbf{V}_2$, where \mathbf{V}_1 only involves the decomposition of first order monomials (first term on the RHS of (2.12)), and \mathbf{V}_2 denotes the higher order contributions (second term on the RHS of (2.12)). The contribution of \mathbf{V}_1 , corresponding to the first term in (2.12) is identical to the contribution from $\mathbf{Y}_1^{(1)}$ except that it is now also summed $|\{\boldsymbol{\beta} \mid |\boldsymbol{\beta}| \leq 1\}|$ times. We thus have

$$|\langle \mathbf{H}, \mathbf{V}_1 \rangle| = \left| \sum_{|\boldsymbol{\beta}| \leq 1} \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (h_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{H} \rangle \right| \quad (3.13)$$

$$= \left| \sum_{|\boldsymbol{\beta}| \leq 1} \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (\tilde{h}_{\ell})_{\zeta} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| + \left| \sum_{|\boldsymbol{\beta}| \leq 1} \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta=0} \varepsilon_{\ell} \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| \quad (3.14)$$

$$\leq \max \{ \mathcal{O}((m+n)^2) \eta, \mathcal{O}((m+n)^2) \|\mathbf{m}_0\| \|\varepsilon\|_{\infty} \} \quad (3.15)$$

In (3.15), we use the fact that every constraint of the form $h_j(\mathbf{z})\mathbf{z}^\beta$ corresponds to a distinct entry in η . The sum over the multi-indices β is thus included into $\mathcal{O}(m+n)$ entries of η which have to be multiplied by $\mathcal{O}(m+n)$ as each constraint $h_j(\mathbf{z})$ appears at most $\mathcal{O}(m+n)$.

For the second term \mathbf{V}_2 , it suffices to note that this term corresponds to multiplying all the polynomials appearing in $\mathbf{Y}_1^{(1)}$ by \mathbf{z}^β and summing up all the resulting polynomials over all possible first order multi-indices β . In terms of \mathbf{V}_2 , for the term which is multiplying ε , this means shifting the first column in $\sum_\ell \sum_\kappa \mathbf{B}_\kappa$ into $\sum_\ell \sum_\kappa \mathbf{B}_{\kappa+\alpha}$ and summing all resulting matrices over α . If we let $\mathcal{C}(\beta)$ denote the set of pairs (ℓ, κ) representing the constraints $h_\ell(z)z^\kappa$ that appear in the expression of $\mathbf{z}^\beta - \mathbf{z}_0^\beta$ following the decomposition (2.11), for the second order contribution \mathbf{V}_2 , using the decomposition given by the second term on the RHS of (2.12), we can write

$$|\langle \mathbf{H}, \mathbf{V}_2 \rangle| = \left| \sum_{|\alpha| \leq 1} \sum_{|\beta| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \sum_{\zeta} (h_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa+\alpha}, \mathbf{H} \rangle \right| \quad (3.16)$$

$$= \left| \sum_{|\alpha| \leq 1} \sum_{|\beta| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \sum_{\zeta \neq 0} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa+\alpha}, \mathbf{M} \rangle \right| \quad (3.17)$$

$$+ \left| \sum_{|\alpha| \leq 1} \sum_{|\beta| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \sum_{\zeta=0} ((\tilde{h}_\ell)_\zeta - \varepsilon_\ell) \langle \mathbf{B}_{\zeta+\kappa+\alpha}, \mathbf{M} \rangle \right| \quad (3.18)$$

$$\leq \left| \sum_{|\alpha| \leq 1} \sum_{|\beta| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \sum_{\zeta} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa+\alpha}, \mathbf{M} \rangle \right| \quad (3.19)$$

$$+ \left| \sum_{|\alpha| \leq 1} \sum_{|\beta| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \sum_{\zeta=0} \varepsilon_\ell \langle \mathbf{B}_{\kappa+\alpha}, \mathbf{M} \rangle \right| \quad (3.20)$$

$$\leq \left| \sum_{|\alpha| \leq 1} \sum_{|\beta| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \sum_{\zeta} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa+\alpha}, \mathbf{M} \rangle \right| \quad (3.21)$$

$$+ \mathcal{O}((m+n)^2) \|\mathbf{m}_0\| \|\varepsilon\|_\infty \quad (3.22)$$

In (3.21) we use the discussion above and the fact that, again since \mathbf{B}_κ only targets monomials of order at most 2, we can focus on the submatrix $\mathbf{M}_{1,1}$ from the decomposition (3.12). Moreover, the norm $\|\sum_{|\beta| \leq 1} \sum_{|\alpha| \leq 1} \sum_{(\ell, \kappa) \in \mathcal{C}(\beta)} \varepsilon_\ell \mathbf{B}_{\kappa+\alpha}\|_F = \|\varepsilon\|_\infty \mathcal{O}(m+n) \|\mathbf{1}\mathbf{1}^*\|_F = \|\varepsilon\|_\infty \mathcal{O}((m+n)^2)$. Equation (3.22) follows from Cauchy-Schwarz. To bound (3.21), simply use the result of (3.11) (first term), noting that in each $\mathbf{z}^\beta - \mathbf{z}_0^\beta$ each polynomial from the ideal appears at most $\mathcal{O}(m+n)$ times. Then use the fact that every α in $(\mathbf{z}^\beta - \mathbf{z}_0^\beta)\mathbf{z}^\alpha$ gives a different constraint in (3.1) so that the sum over α can be included within η . Apply Cauchy-Schwarz to $|\langle \lambda_2, \tilde{\mathcal{A}}(M) \rangle|$ with the bound on $\tilde{\mathcal{A}}(H)$ given by (3.1) and $\|\lambda\| = \mathcal{O}(m+n)$. This gives the following bound $|\langle \mathbf{H}, \mathbf{V}_2 \rangle|$

$$|\langle \mathbf{H}, \mathbf{V}_2 \rangle| \leq \max \{ (m+n)^2 \eta, \mathcal{O}((m+n)^2) \|\mathbf{m}_0\|_2 \|\varepsilon\|_\infty \} \quad (3.23)$$

Consider the sum-of-squares certificate of section (2.3). In polynomial form, we have seen that this certificate reads $\sum_{|\gamma| \leq 2} (\mathbf{z}^\gamma - \mathbf{z}_0^\gamma)^2 = \sum_{|\gamma| \leq 2} (\mathbf{z}^\gamma - (\mathbf{m}_0)_\gamma)^2$. One possible matrix representation³ of this certificate is thus given by $\bar{Z} = \sum_j \mathbf{s}_j \mathbf{s}_j^T$ where each \mathbf{s}_j denote a vector of the form $-(m_0)_\gamma \mathbf{e}_1 + \mathbf{e}_\gamma$. Using this form, we get

$$\bar{Z} = \mathbf{I} - \mathbf{m}_0 \mathbf{e}_1^T - \mathbf{e}_1 \mathbf{m}_0^T + \|\mathbf{m}_0\|^2 \mathbf{e}_1 \mathbf{e}_1^T = \mathbf{I} - \mathbf{Y}_2. \quad (3.24)$$

³An alternative representation would be given by the decomposition Trace + ideal of section (2.3) and encoded as $\mathbf{I} - \mathbf{Y}_1$.

For $\mathbf{m} \perp \mathbf{m}_0$, with $\|\mathbf{m}\| = 1$ and $m_1 = \mathbf{e}_1^T \mathbf{m}$, we have

$$\langle \mathbf{Y}_2, \mathbf{m} \mathbf{m}^T \rangle = -m_1^2 \|\mathbf{m}_0\|^2 \leq 0.$$

This last equation implies that $(\mathbf{Y}_2)_{T^\perp} \preceq 0$.

Let $\mathbf{I} - \mathbf{Y}_1$ denote the matrix form of the polynomial certificate constructed in section 2.3. As we have $\mathbf{z}_B^* (\mathbf{I} - \mathbf{Y}_1) \mathbf{z}_B = \mathbf{z}_B (\mathbf{I} - \mathbf{Y}_2) \mathbf{z}_B = \sum_{|\gamma| \leq 2} (\mathbf{z}^\gamma - \mathbf{z}_0^\gamma)^2 = \sum_j s_j(\mathbf{z})^2$, proposition (2) applies and there exists a matrix $\boldsymbol{\xi}$ satisfying $\boldsymbol{\xi} + \mathbf{Y}_1 = \mathbf{Y}_2$. Recall that the certificate reads

$$\mathbf{Y} = \mathbf{I} - \mathbf{Y}_2 = \mathbf{I} - \boldsymbol{\xi} - \mathbf{Y}_1 = \mathbf{I} - \boldsymbol{\xi} - \sum_{\ell} \sum_{\boldsymbol{\kappa} \in \mathbb{N}_{2(t-d_{h_\ell})}^K} \lambda_{\ell \boldsymbol{\kappa}} \left(\sum_{\zeta} \frac{(h_\ell)_{\zeta}}{\|\mathbf{B}_{\zeta+\boldsymbol{\kappa}}\|_F^2} \mathbf{B}_{\zeta+\boldsymbol{\kappa}} \right) \quad (3.25)$$

Now using $\text{Tr}(\mathbf{H}) \leq 0$, we can write,

$$0 \geq \text{Tr}(\mathbf{H}_T) + \text{Tr}(\mathbf{H}_{T^\perp}) \quad (3.26)$$

$$= \langle \mathbf{H}, \mathbf{I}_T \rangle + \langle \mathbf{H}, \mathbf{I}_{T^\perp} \rangle \quad (3.27)$$

$$= \langle \mathbf{H}, \mathbf{I}_T \rangle - \langle \mathbf{H}, \mathbf{Y}_2 \rangle + \langle \mathbf{H}, \mathbf{Y}_2 \rangle + \langle \mathbf{H}, \mathbf{I}_{T^\perp} \rangle \quad (3.28)$$

$$= \langle \mathbf{H}_T, \mathbf{I}_T - (\mathbf{Y}_2)_T \rangle - \langle \mathbf{H}_{T^\perp}, (\mathbf{Y}_2)_{T^\perp} \rangle + \langle \mathbf{H}, \mathbf{Y}_2 \rangle + \langle \mathbf{H}, \mathbf{I}_{T^\perp} \rangle \quad (3.29)$$

$$\geq -\langle \mathbf{H}_{T^\perp}, (\mathbf{Y}_2)_{T^\perp} \rangle - |\langle \mathbf{H}, \mathbf{Y}_2 \rangle| + \langle \mathbf{H}, \mathbf{I}_{T^\perp} \rangle \quad (3.30)$$

$$\geq -\langle \mathbf{H}_{T^\perp}, (\mathbf{Y}_2)_{T^\perp} \rangle - |\langle \mathbf{H}, \mathbf{Y}_1 \rangle| + \langle \mathbf{H}, \mathbf{I}_{T^\perp} \rangle \quad (3.31)$$

$$\geq -|\langle \mathbf{H}, \mathbf{Y}_1 \rangle| + \text{Tr}(\mathbf{H}_{T^\perp}) \quad (3.32)$$

As explained above, \mathbf{Y}_1 is used to denote the component of the dual certificate which is in the range of \mathcal{A}^* , i.e. $\mathbf{Y}_1 = \mathcal{A}^* \boldsymbol{\lambda}$. In (3.31), we use the fact that $\mathbf{Y}_1 = \mathbf{Y}_2 + \boldsymbol{\xi}$ and $\langle \boldsymbol{\xi}, \mathbf{B}_\gamma \rangle = 0$, for all γ . Since both \mathbf{M} and \mathbf{M}_0 are solutions to problem (3.1). Both of these matrices thus satisfy the structural constraints exactly, and read $\mathbf{M} = \sum_{\gamma} m_{\gamma} \mathbf{B}_\gamma$, $\mathbf{M}_0 = \sum_{\gamma} (m_0)_{\gamma} \mathbf{B}_\gamma$ for some m_{γ} . Together with lemma 2, this implies $\langle \mathbf{H}, \boldsymbol{\xi} \rangle = \langle \mathbf{M} - \mathbf{M}_0, \boldsymbol{\xi} \rangle = 0$. Finally, in (3.32), we use the fact that for a positive semidefinite matrix \mathbf{H}_{T^\perp} , and a matrix $(\mathbf{Y}_2)_{T^\perp}$ such that $(\mathbf{Y}_2)_{T^\perp} \preceq 0$, $\langle \mathbf{H}_{T^\perp}, (\mathbf{Y}_2)_{T^\perp} \rangle \leq 0$.

We also use the fact that both \mathbf{M} and \mathbf{M}_0 satisfies the structural constraints so that $\langle \boldsymbol{\xi}, \mathbf{H} \rangle = 0$. The last line implies

$$\text{Tr}(\mathbf{H}_{T^\perp}) \leq |\langle \mathbf{H}, \mathbf{Y}_1 \rangle| \leq \max \{ (m+n)^2 \eta, \mathcal{O}((m+n)^2) \|\mathbf{m}_0\|_2 \|\varepsilon\|_\infty \} \quad (3.33)$$

Finally $|\langle \mathbf{H}, \mathbf{Y}_1 \rangle|$ is bounded from (3.11), (3.15) and (3.23).

3.1.2 Bound on \mathbf{H}_T

We now use a more quantitative version of injectivity of the linear map \mathcal{A} , encoding the polynomial constraints, on T to derive a bound on \mathbf{H}_T . Let \mathbf{H}_T be expressed as $\mathbf{H}_T = \mathbf{y} \mathbf{m}_0^T + \mathbf{m}_0 \mathbf{y}^T$ for some $\mathbf{y} \in \mathbb{R}^{\mathbb{N}_2^K}$ (see (2.3)).

Using this decomposition for \mathbf{H}_T , and letting $h_{i_1} \rightarrow h_{i_2} \rightarrow \dots$ denote the ordered series of constraints

making the chain (1.4), we have

$$\begin{aligned}
y_{i_1} + y_1(\mathbf{z}_0)_{i_1} &= \left\{ \tilde{\mathcal{A}}(\mathbf{H}_T) \right\}_1 = \sum_{|\zeta|>0} \frac{(\tilde{h}_{i_1})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \\
y_{i_1}(\mathbf{z}_0)_{i_2} + y_{i_2}(\mathbf{z}_0)_{i_1} &= \left\{ \tilde{\mathcal{A}}(\mathbf{H}_T) \right\}_2 = \sum_{|\zeta|>0} \frac{(\tilde{h}_{i_2})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \\
y_{i_2}(\mathbf{z}_0)_{i_3} + y_{i_3}(\mathbf{z}_0)_{i_2} &= \left\{ \tilde{\mathcal{A}}(\mathbf{H}_T) \right\}_3 = \sum_{|\zeta|>0} \frac{(\tilde{h}_{i_3})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \\
y_{i_3}(\mathbf{z}_0)_{i_4} + y_{i_4}(\mathbf{z}_0)_{i_3} &= \dots
\end{aligned}$$

To derive a bound for \mathbf{H}_T , we then isolate each of the entries in \mathbf{y} as,

$$\begin{aligned}
y_{i_1} &= \sum_{|\zeta|>0} \frac{(\tilde{h}_{i_1})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle - y_1(\mathbf{z}_0)_{i_1} \\
y_{i_2} &= \frac{1}{(z_0)_{i_1}} \left[\sum_{|\zeta|>0} \frac{(\tilde{h}_{i_2})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle - (z_0)_{i_2} y_{i_1} \right] \\
y_{i_3} &= \frac{1}{(z_0)_{i_2}} \left[\sum_{|\zeta|>0} \frac{(\tilde{h}_{i_3})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle - y_{i_2}(\mathbf{z}_0)_{i_3} \right] \\
y_{i_4} &= \dots
\end{aligned} \tag{3.34}$$

(3.34) thus gives a general expression for every first order entry y_{i_ℓ} of \mathbf{y} , as a weighted combination of the constraints which can be considered as a noisy version of (2.10) or (2.11). Generally, every first order y_{i_ℓ} can thus be expressed as the weighted combination

$$y_{i_\ell} = \frac{1}{(z_0)_{i_{\ell-1}}} \sum_{|\zeta|>0} \frac{(\tilde{h}_{i_\ell})_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \pm \sum_{j=1}^{\ell-1} \frac{(z_0)_{i_\ell}}{z_0(i_{\ell-j-1})z_0(i_{\ell-j})} \sum_{|\zeta|>0} \frac{(\tilde{h}_j)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \pm (z_0)_{i_\ell} (\mathbf{H}_{T^\perp})_{11}. \tag{3.35}$$

The last term in (3.35) follows from $y_1(\mathbf{z}_0)_1 = (\mathbf{H}_T)_{11} = H_{11} - (\mathbf{H}_T^\perp)_{1,1} = -(\mathbf{H}_T^\perp)_{1,1}$. Let $(\mathbf{y}_\alpha)_{|\alpha|\leq 1}$ denote the entries in \mathbf{y} corresponding to the multi-indices that give rise to degree one monomials. Let C_3 bound each of the weights appearing in front of the constraints making up the chain in (3.35). The first order part of \mathbf{y} , $(\mathbf{y}_\alpha)_{|\alpha|\leq 1}$, has length $m+n$ and each of its entry is bounded by at most a sum of all the constraints making the connected path in the bipartite graph. One can thus write

$$\begin{aligned}
\|(\mathbf{y}_\alpha)_{|\alpha|\leq 1}\|_2 &\leq C_3(m+n)^{1/2} \left(\sum_\ell \left| \sum_{|\zeta|>0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \right| + \|\mathbf{H}_{T^\perp}\|_1 \right) \\
&\leq C_3(m+n)^{1/2} \left(\sum_\ell \left| \sum_{|\zeta|>0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \right| + \|\mathbf{H}_{T^\perp}\|_1 \right). \tag{3.36}
\end{aligned}$$

The constant C_3 depends on the entries of \mathbf{X}_0 as

$$C_3 = \mathcal{O} \left(\sup_{(k,\ell)\in\Omega} \frac{|z_0(m)|}{|z_0(k)z_0(\ell)|} \right) = \mathcal{O}(1).$$

To bound the second order components of \mathbf{y} , we use the structural constraints $\langle \mathbf{M}, \mathbf{E}_{\delta_1, \delta_2} - \mathbf{E}_{\gamma_1, \gamma_2} \rangle = 0$, as those are not affected by the noise. Those constraints are enforcing equality between the entries

$y_\ell m_0(k) + y_k m_0(\ell)$ and the second order entries of the first column of \mathbf{H}_T , namely $y(\ell, k)m_0(1) + y_1 m_0(\ell, k)$. For multiindices $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that $|\boldsymbol{\alpha}| = |\boldsymbol{\beta}| = 1$, using those relations, we therefore have

$$y_{\boldsymbol{\alpha}+\boldsymbol{\beta}}(\mathbf{z}_0)_1 + (y)_1(\mathbf{z}_0)^{\boldsymbol{\alpha}+\boldsymbol{\beta}} = y_{\boldsymbol{\alpha}}(m_0)_{\boldsymbol{\beta}} + y_{\boldsymbol{\beta}}(m_0)_{\boldsymbol{\alpha}} + \langle \mathbf{H}_T, \mathbf{E}_{\boldsymbol{\alpha}+\boldsymbol{\beta},1} - \mathbf{E}_{\boldsymbol{\alpha},\boldsymbol{\beta}} \rangle, \quad (3.37)$$

$$= y_{\boldsymbol{\alpha}}(m_0)_{\boldsymbol{\beta}} + y_{\boldsymbol{\beta}}(m_0)_{\boldsymbol{\alpha}} + \langle \mathbf{H} - \mathbf{H}_{T^\perp}, \mathbf{E}_{\boldsymbol{\alpha}+\boldsymbol{\beta},1} - \mathbf{E}_{\boldsymbol{\alpha},\boldsymbol{\beta}} \rangle, \quad (3.38)$$

$$= y_{\boldsymbol{\alpha}}(m_0)_{\boldsymbol{\beta}} + y_{\boldsymbol{\beta}}(m_0)_{\boldsymbol{\alpha}} - \langle \mathbf{H}_{T^\perp}, \mathbf{E}_{\boldsymbol{\alpha}+\boldsymbol{\beta},1} - \mathbf{E}_{\boldsymbol{\alpha},\boldsymbol{\beta}} \rangle. \quad (3.39)$$

The last line follows from the fact that $\langle \mathbf{M} - \mathbf{M}_0, \mathbf{E}_{\delta_1, \delta_2} - \mathbf{E}_{\gamma_1, \gamma_2} \rangle = 0$ for any $\delta_1 + \delta_2 = \gamma_1 + \gamma_2$. Using (3.36) as well as Hölder's inequality and the fact that, for a constant C_4 , $\|\mathbf{E}_{\delta_1, \delta_2} - \mathbf{E}_{\gamma_1, \gamma_2}\|_\infty \leq C_4$ for any $(\boldsymbol{\delta}_1, \boldsymbol{\delta}_2), (\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2)$, one can write,

$$\|y_{(\boldsymbol{\alpha}, \boldsymbol{\beta})}\|_2 \lesssim (m+n) \left(\left\| \left(\sum_{|\zeta|>0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \right)_{\ell|_1} \right\| + \|\mathbf{H}_{T^\perp}\|_1 + |y_1| \right) \quad (3.40)$$

$$\lesssim (m+n) \left(\left\| \left(\sum_{|\zeta|>0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \right)_{\ell|_1} \right\| + \|\mathbf{H}_{T^\perp}\|_1 \right) \quad (3.41)$$

In (3.41), we again use $|y_1(m_0)_1| = |(\mathbf{H}_T)_{11}| = |\mathbf{H}_{11} - (\mathbf{H}_{T^\perp})_{11}| = |(\mathbf{H}_{T^\perp})_{11}|$. Combining (3.36) and (3.41), we get the following bound on \mathbf{H}_T ,

$$\begin{aligned} \|\mathbf{H}_T\|_F &\leq \|\mathbf{m}_0\|_2 \|\mathbf{y}\|_2 \\ &\leq C_5 \|\mathbf{m}_0\|_2 (m+n) \left(\sum_{\ell} \left| \sum_{|\zeta|>0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_T, \mathbf{B}_\zeta \rangle \right| + \|\mathbf{H}_{T^\perp}\|_1 \right). \end{aligned} \quad (3.42)$$

The second term can be bounded by (3.33). To bound the first term, note that we have $\mathbf{H}_T = \mathbf{H} - \mathbf{H}_{T^\perp}$ and that the expression obtained by considering the first term above and substituting \mathbf{H} for \mathbf{H}_T can be bounded through the ℓ_2 constraint in (3.1). We can thus focus on bounding this term when replacing \mathbf{H}_T by \mathbf{H}_T^\perp . In each first order constraint \tilde{h}_ℓ , there is only one non zero coefficient $(\tilde{h}_\ell)_\zeta$ for $|\zeta| > 0$ and each \mathbf{B}_γ only has $\mathcal{O}(1)$ non zero entries. Moreover none of the constraints in the chain are targeting the same entry in \mathbf{H}_T^\perp . Let Ω_P denote the moments matrix defined as $(\Omega_P)_\zeta = \text{sign}(h_\ell)_\zeta$, for all ζ such that there exists a ℓ with $(h_\ell)_\zeta \neq 0$, and 0 otherwise. This matrix has Frobenius norm at most $\mathcal{O}(\sqrt{m+n})$. Using this matrix, we can write,

$$\sum_{\ell} \left| \sum_{|\zeta|>0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}_{T^\perp}, \mathbf{B}_\zeta \rangle \right| \leq C_6 |\langle \Omega_P, \mathbf{H}_T^\perp \rangle| \quad (3.43)$$

$$\leq C_6 (m+n)^{1/2} \|\mathbf{H}_{T^\perp}\|_F \quad (3.44)$$

$$\leq C_6 (m+n)^{1/2} \|\mathbf{H}_{T^\perp}\|_1 \quad (3.45)$$

In (3.43), we use Hölder's inequality, together with the fact that the chain has length $\mathcal{O}(m+n)$. Using the trace bound (3.33), and substituting (3.45) into (3.42), we finally get the bound on \mathbf{H}_T as,

$$\|(\mathbf{H}_T)\|_F \leq C_7 (m+n) \|\mathbf{m}_0\|_2 \left((m+n)^{1/2} \|\mathbf{H}_{T^\perp}\|_1 + \sum_{\ell} \left| \sum_{\zeta} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{H}, \mathbf{B}_\zeta \rangle \right| \right) \quad (3.46)$$

$$\leq C_7 (m+n) \|\mathbf{m}_0\|_2 \left((m+n)^{1/2} |\langle \mathbf{H}, \mathbf{Y}_1 \rangle| + 2\eta (m+n)^{1/2} \right) \quad (3.47)$$

$$\leq C_7 (m+n)^{3/2} \|\mathbf{m}_0\|_2 \max \{ (m+n)^2 \eta, \mathcal{O}((m+n)^2) \|\mathbf{m}_0\|_2 \|\varepsilon\|_\infty \} \quad (3.48)$$

$$\leq C_7 \|\mathbf{m}_0\|_2^2 (m+n)^{5/2} \eta. \quad (3.49)$$

In (3.49), we use $\|\mathbf{m}_0\|_2 = \mathcal{O}(m+n)$. Using this last bound together with (3.33), we finally get,

$$\|\mathbf{H}\|_F \leq \|\mathbf{H}_T\|_F + \|\mathbf{H}_{T^\perp}\|_F \leq \|\mathbf{H}_T\|_F + \|\mathbf{H}_{T^\perp}\|_1 \lesssim (m+n)^{5/2} \eta \|\mathbf{M}_0\|_F.$$

By definition of η , we also have $\eta = \mathcal{O}(m+n)\|\varepsilon\|$ which enables to concludes.

The next section shows how the scaling factor can be reduced to $(m+n)^2\|\varepsilon\|_2$ when paths are known between any root node and the corresponding leaf nodes in the bipartite graph, and the noise can be constrained along those paths.

3.2 Proof of Corollary 3

The proof of corollary 3 follows the idea of section 3.1 with the difference that we now constrain the noise along the path and consider a reduced SDP. When considering the reduced (sparse) formulation, we only consider monomials of order 2 that are appearing in the constraints. There are $\mathcal{O}(m+n)$ such monomials. The moments matrix has now size $\mathcal{O}(m+n) \times \mathcal{O}(m+n)$. The part of the certificate expressing first order squares remain unchanged. The second order squares can be written directly from the constraints, trace and constant ρ without the need for any propagation, i.e. $(\mathbf{z}^\alpha \mathbf{z}^\beta - \mathbf{z}_0^\alpha \mathbf{z}_0^\beta)^2 = -2\mathbf{z}_0^\alpha \mathbf{z}_0^\beta (\mathbf{z}^\alpha \mathbf{z}^\beta - \mathbf{z}_0^\alpha \mathbf{z}_0^\beta) + (\mathbf{z}^\alpha \mathbf{z}^\beta)^2 - (\mathbf{z}_0^\alpha \mathbf{z}_0^\beta)^2$. The certificate thus becomes much sparser. Conditions 1) to 3) still hold for this certificate as it still has the exact same structure $\sum_j \mathbf{s}_j \mathbf{s}_j^*$ as before except that the number of such squares is reduced. The squared polynomials are now given by $(\mathbf{z}_0^\alpha - \mathbf{z}^\alpha)^2$, for all $|\alpha| \leq 1$ and $(\mathbf{z}_0^\gamma - \mathbf{z}^\gamma)^2$ for all $|\gamma| = 2$ such that $\mathbf{z}^\gamma - \mathbf{z}_0^\gamma$ appears in the constraints. As the matrix has now size $|\{\alpha, |\alpha| \leq 1\}| + |\{\gamma \mid \gamma \in \Omega\}|$, the rank condition in 2) still holds as well.

Moreover, from (3.2), we now have for each path \mathcal{P}_i , $i = 1, \dots, P$,

$$\sqrt{\sum_{\kappa \in \mathcal{K}_i} \sum_{\ell \in \mathcal{P}_i} \left| \sum_{\zeta} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_{\zeta+\kappa}\|_F^2} \langle \mathbf{H}, \mathbf{B}_{\zeta+\kappa} \rangle \right|^2} \leq 2\eta'. \quad (3.50)$$

As for the proof of Theorem 1, we let $\mathbf{Y}_1^{(1)}$ and $\mathbf{Y}_1^{(2)}$, with $\mathbf{Y}_1 = \mathbf{Y}_1^{(1)} + \mathbf{Y}_1^{(2)}$ denote the contributions of first and second order squares to the certificate \mathbf{Y}_1 . To bound the inner product $|\langle \mathbf{H}, \mathbf{Y} \rangle|$, we once again replace the noiseless constraints appearing in the expression of the SOS certificate with the noiseless constraints that are bounded through (3.2). For the first order contribution $\mathbf{Y}_1^{(1)}$, we have,

$$|\langle \mathbf{Y}_1^{(1)}, \mathbf{H} \rangle| \leq \left| \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| + \left| \sum_{|\kappa| \leq 1} \sum_{\ell} W_{\ell, \kappa} \langle \mathbf{B}_\kappa, \mathbf{M} \rangle \right| \quad (3.51)$$

$$\leq \left| \sum_{\kappa} \sum_{\ell} W_{\ell, \kappa} \sum_{\zeta} (\tilde{h}_\ell)_\zeta \langle \mathbf{B}_{\zeta+\kappa}, \mathbf{M} \rangle \right| + \mathcal{O}((m+n)^{3/2}) \|\mathbf{m}_0\| \|\varepsilon\|_\infty \quad (3.52)$$

$$\leq \mathcal{O}((m+n)^{3/2}) \eta' + \mathcal{O}((m+n)^{3/2}) \|\mathbf{m}_0\| \|\varepsilon\|_\infty \quad (3.53)$$

The bound (3.52) on the second term follows the exact same reasoning as (3.10). The difference is for the first term, for which we now use the bound on the given path \mathcal{P} . For the second term, following the proof of Theorem (1), noting that we now only use second order moments appearing in the constraints, and using (3.2), we can write

$$|\langle \mathbf{Y}_1^{(2)}, \mathbf{H} \rangle| \leq \max \left\{ \eta', \mathcal{O}((m+n)^{3/2}) \|\mathbf{m}_0\| \|\varepsilon\|_\infty \right\} \quad (3.54)$$

So that $|\langle \mathbf{H}, \mathbf{Y} \rangle| \leq (m+n) \max \{ \mathcal{O}((m+n)^{3/2})\eta', \mathcal{O}((m+n)^{3/2})\|\mathbf{m}_0\|\|\varepsilon\|_\infty \}$. In a similar way, the expression for the y_{i_k} in (3.34) also relies on the first order constraints making up the path from the root node to y_{i_k} so that the relations and bound in (3.34) and (3.36) can now be reduced to

$$\|\mathbf{y}_{|\alpha| \leq 1}\|_2 \lesssim \sqrt{m+n} \sup_{\mathcal{P}_i} \sum_{\ell \in \mathcal{P}_i} \left| \left(\sum_{|\zeta| > 0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{B}_\zeta, \mathbf{H}_T \rangle \right) \right|. \quad (3.55)$$

Relation (3.39) still holds. We still only need to account for $(m+n)$ second order monomials corresponding to the $\mathcal{O}(m+n)$ constraints and one can thus simply bound the second order part of \mathbf{y} as

$$\|y_{\alpha+\beta}\|_2 \lesssim \sqrt{m+n} \sup_{\mathcal{P}_i} \sum_{\ell \in \mathcal{P}_i} \left| \left(\sum_{|\zeta| > 0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{B}_\zeta, \mathbf{H}_T \rangle \right) \right| + \sqrt{m+n} \|\mathbf{H}_{T^\perp}\|_1 \quad (3.56)$$

The square root in (3.56) comes from the problem-depedent formulation. Grouping (3.55) and (3.56), the bound on \mathbf{H}_T can therefore read,

$$\|\mathbf{H}_T\|_F \lesssim \sqrt{m+n} \|\mathbf{m}_0\|_2 \left(\sup_{\mathcal{P}_i} \sum_{\ell \in \mathcal{P}_i} \left| \left(\sum_{|\zeta| > 0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{B}_\zeta, \mathbf{H}_T \rangle \right) \right| + \|\mathbf{H}_{T^\perp}\|_1 \right) \quad (3.57)$$

$$\lesssim \sqrt{m+n} \|\mathbf{m}_0\|_2 \left(\sup_{\mathcal{P}_i} \sum_{\ell \in \mathcal{P}_i} \left| \left(\sum_{|\zeta| > 0} \frac{(\tilde{h}_\ell)_\zeta}{\|\mathbf{B}_\zeta\|_F^2} \langle \mathbf{B}_\zeta, \mathbf{H} \rangle \right) \right| + (m+n)^{1/2} \|\mathbf{H}_{T^\perp}\|_1 \right) \quad (3.58)$$

$$(3.59)$$

Using (3.2), we have

$$\|(\mathbf{H}_T)\|_F \lesssim \|\mathbf{m}_0\|_2 (m+n) \max \left\{ \mathcal{O}((m+n)^{3/2})\eta', \mathcal{O}((m+n)^{3/2})\|\mathbf{m}_0\|\|\varepsilon\|_\infty \right\} \quad (3.60)$$

Noting that in the reduced formulation, $\|\mathbf{m}_0\|_2 = \mathcal{O}(\sqrt{m+n})$, $\|\mathbf{M}_0\|_F = \mathcal{O}(m+n)$ and using $\eta' = \mathcal{O}(\|\varepsilon\|)$ gives the desired result.

4 Numerical methods

Section 4.1 starts by providing a comparison of the stability and recovery guarantees of the convex formulation against traditional approaches such as nuclear norm minimization, nonlinear propagation, and ridge regression.

Sections 4.2 through 4.5 discuss scalable numerical schemes. Simply listing the moments up to order 4 has complexity $\mathcal{O}(N^4)$ where $N = m+n$, hence is not a scalable representation of the moments matrix. The traditional remedy is the factorized gradient approach due to Burer and Monteiro [9, 10], but our first numerical observation will not be a surprise to the specialist: difficult instances of matrix completion lead to the presence of spurious local minimizers. With adequate compression of the variables and constraints, and provided convergence is to the global minimizer, we show how the problem can be solved in an empirical $\mathcal{O}(N^2)$ complexity.

The conclusions of section 4 can be summarized as follows.

- Factorization approaches sometimes introduce spurious minimizers for sufficiently difficult (small δ) problems. When convergence to such minimizers occur, it is sometimes possible to add an additional rounding step and to extract the solution from the second order block rather than considering the whole matrix.

- Factorizing the moment matrix in low rank form still has storage complexity $\mathcal{O}(N^2)$, hence is not fully scalable. We propose to instead view the moment matrix as a tensor, and upgrade to a more efficient hierarchical low-rank factorization with storage complexity $\mathcal{O}(N)$. This factorization seems to always work when the simpler factorized gradient works.
- The hierarchical factorization is in itself not sufficient to guarantee scalability, as formulation (2.1), and in particular total symmetry, still requires encoding a combinatorial ($\mathcal{O}(N^4)$) number of constraints. Section 4.4 then introduces three different trace relations, which are derived from the third and fourth order total symmetry constraints. Enforcing those relations in place of the original total symmetry constraints reduces the computational cost required to enforce these constraints from $\mathcal{O}(N^4)$ to $\mathcal{O}(1)$ in the best case. This compression of the total symmetry constraint thus reduces the global complexity to a factor $\mathcal{O}(N^3)$. Empirically, we again observe that those trace relations can be used as a substitute for the more expensive total symmetry constraints as soon as the traditional factorized gradient method works.
- Given the $\mathcal{O}(1)$ trace relations and the hierarchical low rank factorization of the moments tensor, a last bottleneck that prevent reducing the global computational cost from $\mathcal{O}(N^3)$ to $\mathcal{O}(N^2)$ is given by the Higher Order Affine constraints which enforce the moments constraints derived from multiplying any of the original constraint by any monomial of degree at most two, to be satisfied. Encoding those constraints requires storing matrices of size $\mathcal{O}(N^3)$ (i.e product of $\mathbf{b} \in \mathbb{R}^N$ by $\mathbf{m}_0 \in \mathbb{R}^{N^2}$). We propose to encode these constraints through random sampling, minimizing over distinct batches of size $\mathcal{O}(N)$ iteratively. Such formulation does not seem to modify the convergence properties and enables us to apply the semidefinite program (2.1) to matrices \mathbf{X} of sizes up to 100×100 without making use of the reduced sparsity based formulation of corollary 3. Dealing with such matrices is not practical in the original framework of the Lasserre hierarchy with two rounds of lifting.

4.1 Lipschitz stability

To illustrate how the noise can affect a nonlinear reconstruction in the propagation framework, we conduct the following experiments. We consider a noise vector $\boldsymbol{\varepsilon} = \gamma \mathbf{n} / \|\mathbf{n}\|$ for $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_3)$. We gradually increase the amplitude γ of the noise vector. For those noise vectors, we let \mathbf{M}_P denote the solution obtained through propagation and \mathbf{M}_L the solution obtained through the stable semidefinite relaxation (1.16). We consider the matrix of example (1.6) for which we let $\delta = .01$. The numerical experiments are then repeated as follows.

- We randomly draw the noise vector $n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_3)$.
- The noise vector is multiplied by the scaling coefficient γ taking values between .001 and .01, so that the corruption is at most 100% of the signal. The noise is added to the entries $(X_0)_{11}$, $(X_0)_{22}$ and $(X_0)_{21}$ of \mathbf{X}_0 to define the (noisy) measurements.
- Our semidefinite programming relaxation is then solved with CVX⁴ for the noisy measurements. We compute the difference between the returned solution \mathbf{M}_L and the optimal solution to the noiseless problem \mathbf{M}_0 through the Frobenius norm as $\|\mathbf{M}_L - \mathbf{M}_0\|_F / \|\mathbf{M}_0\|_F$.
- The equivalent solution obtained through nonlinear propagation is computed and compared to M_0 as $\|\mathbf{M}_P - \mathbf{M}_0\|_F / \|\mathbf{M}_0\|_F$.

Those various steps are repeated for the various noise levels and for a collection of random vectors n . Note that, because we consider example (1.6), nuclear norm fails even in the absence of noise. For each choice of

⁴<http://cvxr.com/about/>

γ the relative errors $\|\mathbf{M}_P - \mathbf{M}_0\|/\|\mathbf{M}_0\|$ and $\|\mathbf{M}_L - \mathbf{M}_0\|_F/\|\mathbf{M}_0\|_F$ are averaged over all the noise vectors. The results are shown in Fig. 2. This figure thus illustrates the evolution of the averaged relative errors $\mathbb{E}_n\|\mathbf{M}_L - \mathbf{M}_0\|/\|\mathbf{M}_0\|$ and $\mathbb{E}_n\|\mathbf{M}_P - \mathbf{M}_0\|_F/\|\mathbf{M}_0\|_F$ for our semidefinite programming relaxation, as well as for nonlinear propagation in an instance where nuclear norm minimization fails. The Figure on the Right is truncated above to enable the comparison between both figures. The relative errors corresponding to low signal to noise ratio (SNR) were otherwise rising above 200%. The SNR is measured in [dB] as $20 \log(\varepsilon/\gamma)$.

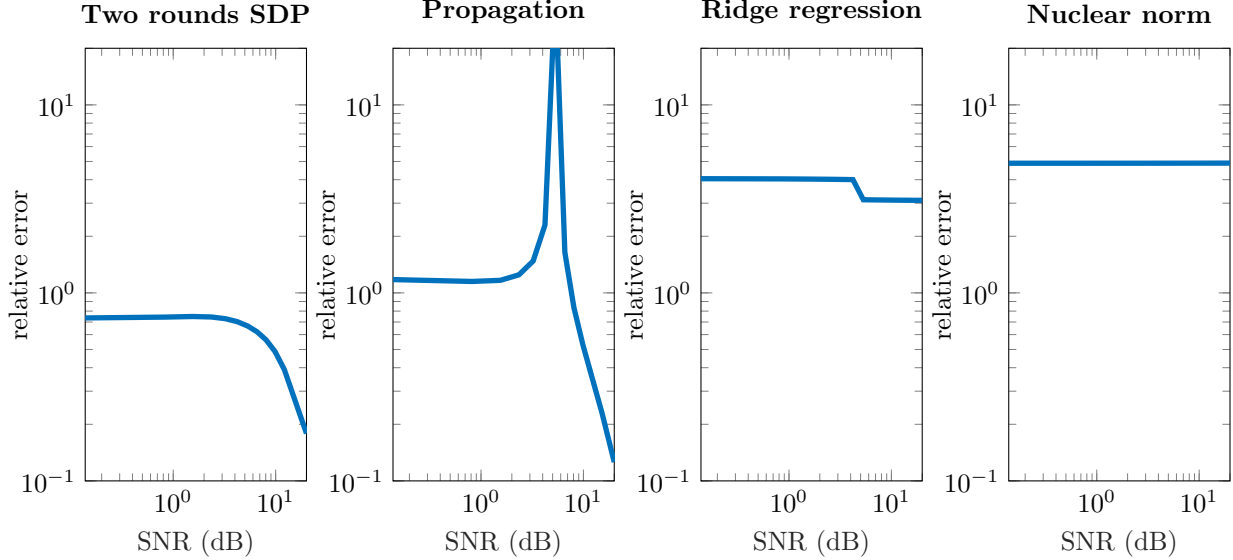


Figure 2: Evolution of the relative error $\|\mathbf{M} - \mathbf{M}_0\|_F/\|\mathbf{M}_0\|_F$ as a function of the noise level (SNR [dB]) for the semidefinite program; nonlinear propagation; ridge regression; and nuclear norm minimization. Blowup can occur for nonlinear propagation whenever the noise takes on values that are close, yet opposite in sign to the small entries in the matrix. Both ridge regression and nuclear norm minimization are known to fail, even in the absence of noise.

4.2 Toward scalability: low-rank factorization

Despite its interest in terms of stability, the semidefinite program (3.1) remains difficult to implement for practical problems because of the size of the second order moments matrix involved. Solving the completion problem on a matrix of size $N \times N$ through (3.1) requires storing a matrix of size N^4 which is often out of reach for typical numerical solvers, on sufficiently interesting instances. In this section we introduce and discuss more scalable numerical methods based on low rank factorizations of the moment matrix (1.14). As is usual in semidefinite programming, the recovery guarantees are however lost when passing to such formulations. This phenomenon is illustrated by Fig 3 to 5.

Among the most popular approaches of the last few years, one of the most efficient, popularized by [8] encodes the unknown positive semidefinite matrix \mathbf{M} from (2.1) as a low rank factorization $\mathbf{M} \approx \mathbf{T}\mathbf{T}^*$, with \mathbf{T} of size $N + N^2$ by r for small r , and then minimizes the augmented Lagrangian over the factor \mathbf{T} . Note that in our case \mathbf{T} is of the form $\mathbf{T} = [\boldsymbol{\alpha}^{(k)}, \mathbf{R}^{(k)}, \boldsymbol{\Pi}^{(k)}]$ where $\boldsymbol{\alpha}^{(k)} \in \mathbb{R}^r$ encodes the normalizing constant, $\mathbf{R}^{(k)} \in \mathbb{R}^{n \times r}$ and $\boldsymbol{\Pi} \in \mathbb{R}^{n^2 \times r}$. We further let $\boldsymbol{\Pi}^{(k)} = [\boldsymbol{\Pi}_1^{(k)}, \dots, \boldsymbol{\Pi}_r^{(k)}]$ where we use $\boldsymbol{\Pi}_j^{(k)} \in \mathbb{R}^{N^2}$ to denote each of the full rank matrices of size $N \times N$ making up the factors in the low rank factorization of the

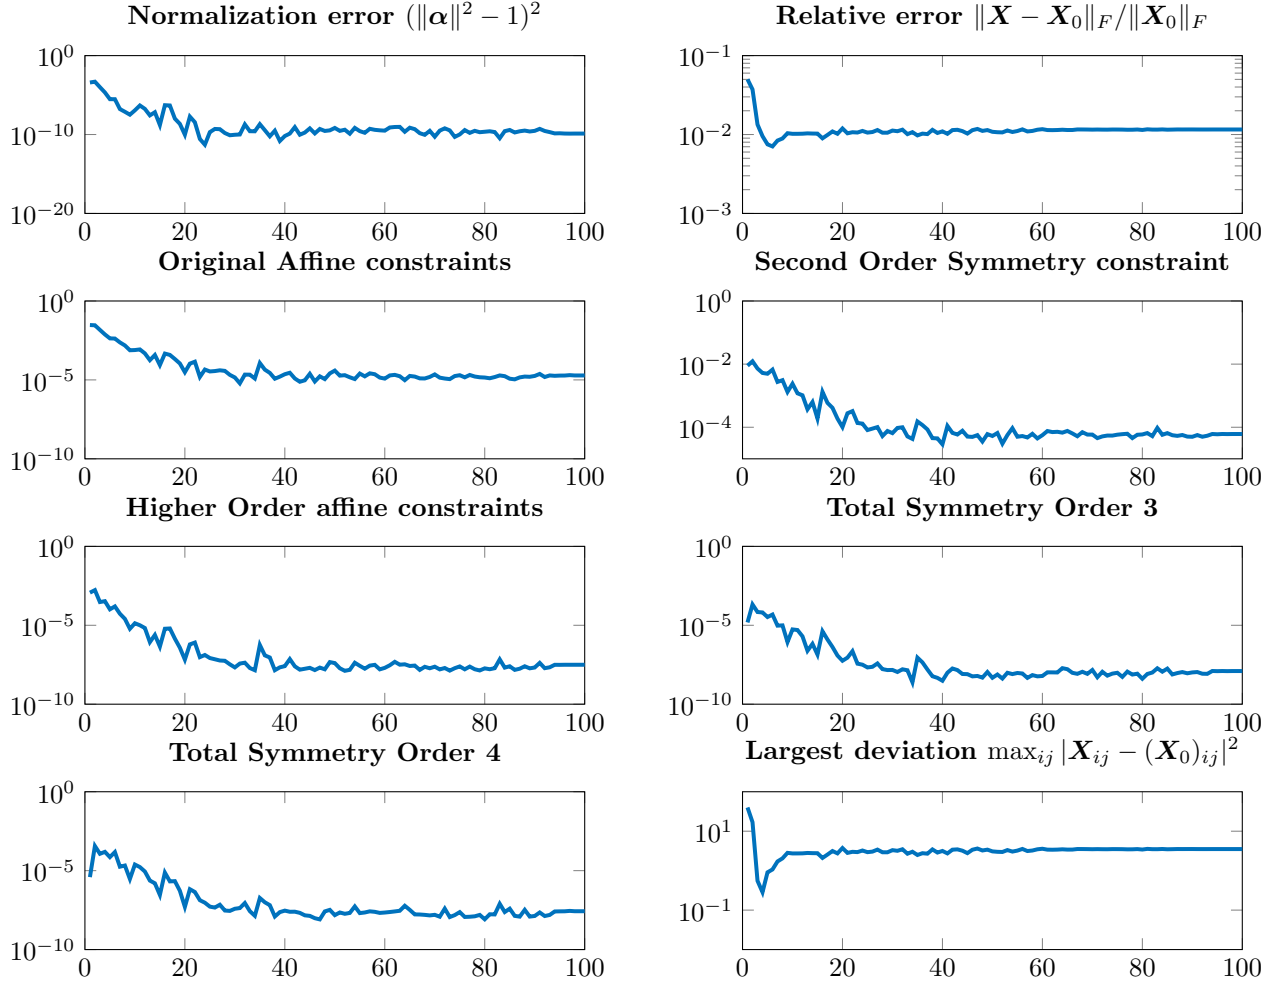


Figure 3: Evolution of each of the error terms appearing in the augmented Lagrangian formulation (4.2) on the benchmark problem (1.6) with $\delta = 0.3$. The relative error (top right) together with the largest deviation highlight the convergence to a local minimizer, while all the other error terms, certifying feasibility, have already reached small thresholds. Those plots can also be compared to the evolution of the trace and global misfit shown in Fig. 4 and the comparison of the structures of the global minimizer with the returned local minimizer shown in Fig. 5.

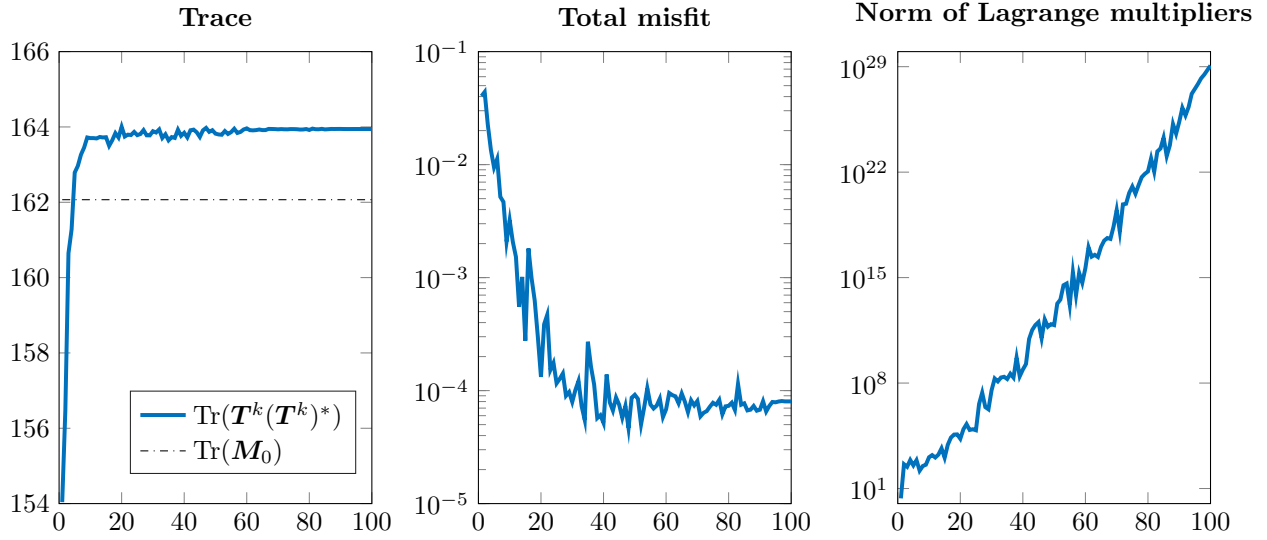


Figure 4: Evolution of the Trace, misfit and norm of the dual Lagrange multipliers. The slow blowup in the norm of those multipliers confirms that we leave the regime in which Theorem 5.4 in [10] works, and thus loses the recovery guarantees for the factorized gradient formulation that are following from this theorem.




Figure 5: Despite the existence of local minima for sufficiently difficult formulations (small δ), the solutions returned by the minimization of the factorized augmented Lagrangian, or hierarchical low rank Lagrangian sometimes remain relatively close to the global solution except for the highest order moments. This figure illustrates this phenomenon. The exact moments matrix \mathbf{M}_0 for the simple example of (1.6) with $\delta = 0.3$ is shown on the left. The moments matrix \mathbf{M} recovered through low rank factorization is shown in the center and the difference $\mathbf{M} - \mathbf{M}_0$ is shown on the right. In practice, for a problem that is not too difficult (i.e. a sufficiently large value of δ in (1.6)) it is thus empirically possible to recover the global solution by simply extracting the lower order moments.

matricization of the fourth order moments tensor. The moments matrix \mathbf{M} then reads,

$$\mathbf{M} = \begin{bmatrix} \boldsymbol{\alpha} \\ \mathbf{R} \\ \boldsymbol{\Pi} \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha}^T & \mathbf{R}^T & \boldsymbol{\Pi}^T \end{bmatrix} = \mathbf{T}\mathbf{T}^T, \quad (4.1)$$

In the factorization above, \mathbf{X} and \mathbf{X}^T are meant to appear as off-diagonal blocks of $\mathbf{R}\mathbf{R}^T$. The rank of each of the factors can be constrained, and is increased, when reaching local minimizers. If we let r denote the rank of the compressed matrix M , such a formulation thus results in a reduction of the number of unknowns from $\mathcal{O}(N^2)$ to only $\mathcal{O}(Nr)$ unknowns. For a set of constraints defined as $\langle \mathbf{A}_i, \mathbf{X} \rangle = b_i$ and encoded in the linear map $\mathcal{A} : \mathbf{X} \mapsto \mathcal{A}(\mathbf{X}) = \{\langle \mathbf{A}_i, \mathbf{X} \rangle\}_{i=1}^m$, a vector of multipliers $\lambda \in \mathbb{R}^m$ and penalty term $\sigma \in \mathbb{R}$, the augmented Lagrangian function corresponding to a minimization of the trace under the linear constraints $\mathcal{A}(\mathbf{T}\mathbf{T}^*) = b$ reads

$$\mathcal{L}(\mathbf{T}, \lambda, \sigma) = \|\mathbf{T}\|_F^2 - 2 \sum_{i=1}^m \lambda_i (\langle \mathbf{A}_i, \mathbf{T}\mathbf{T}^* \rangle - b_i) + \sigma \sum_{i=1}^m |\langle \mathbf{A}_i, \mathbf{T}\mathbf{T}^* \rangle - b_i|^2 \quad (4.2)$$

For some initial guess $\mathbf{T}^{(0)}$, we let v_0 be initialized as $v_0 = \rho(\mathbf{T}^{(0)}(\mathbf{T}^{(0)})^*) = \sum_{i=1}^m (\langle \mathbf{A}_i, \mathbf{T}^{(0)}(\mathbf{T}^{(0)})^T \rangle - b_i)^2$. Finally set k to 0. The augmented Lagrangian algorithm iteratively minimizes the Lagrangian over \mathbf{T} (step 1) and updates the multipliers (step 2) according to the following rule (see [8]). Let ρ denote the norm of the vector of residuals following from step 1, $\rho := \sum_{i=1}^m (\langle \mathbf{A}_i, \mathbf{T}^{(k)}(\mathbf{T}^{(k)})^T \rangle - b_i)^2$. We set $\gamma = 2$ (when dealing with more difficult cases, this update parameter should be increased) and $\eta = 0.25$. If $\rho < \eta v_k$, $y^{k+1} \leftarrow y^k - \sigma^k (\mathcal{A}(\mathbf{T}\mathbf{T}^*) - b)$, $\sigma^{k+1} \leftarrow \sigma^k$, and $v_{k+1} \leftarrow \rho$. Otherwise, $y^{k+1} \leftarrow y^k$, $\sigma^{k+1} \leftarrow \gamma \sigma^k$, $v_{k+1} \leftarrow v_k$. Finally set $k \leftarrow k + 1$ and repeat step 1.

When dealing with problems (2.1) and (3.1) we should favor penalty formulations over Lagrangian formulations, as the number of symmetry constraints is combinatorial in the dimension and therefore requires large vectors of multipliers. The convex formulation (1.15) then turns into

$$\min \quad \|\mathbf{T}\|_F^2 + \sigma \|\mathcal{A}(\mathbf{T}\mathbf{T}^*) - \mathbf{b}\|^2 \quad (4.3)$$

In difficult cases (e.g., when δ is sufficiently small), convergence of iterative methods can suffer for (4.3) and the Lagrangian formulation is thus more appropriate. A hybrid formulation, intermediate between the penalty formulation (4.3) and the more expensive Lagrangian (4.2) is to consider an incomplete set of Lagrange multipliers. In the rest of this section, we will focus on making formulation 4.2 more tractable.

4.3 From low rank to hierarchical low rank

When considering large matrices, such as used by the Lasserre hierarchy, even when using rank constrained factorization, an optimization framework such as (4.2) with the factorization (4.1) still requires storing $\mathcal{O}(N^2r)$ unknowns. It is however possible to factorize in low-rank form the higher order blocks in (1.14). This leads to a multi-level or *hierarchical encoding* of the moments matrix underlying the Lasserre/sos hierarchies. Within the completion framework, it means that each of the factors $\boldsymbol{\Pi}_\ell^{(k)}$, $\ell = 1, \dots, r_1$ can be encoded as a symmetric low rank factorization. This idea is known as hierarchical Tucker decomposition in tensor analysis.

The hierarchical factorization thus relies on two dynamic ranks. The first rank r_1 controls the factorization of the moments matrix as a whole. The second rank, r_2 controls the factorization of the fourth order tensors $\boldsymbol{\Pi}_\ell$ which are thus stored as the tuples $\{\mathbf{S}_{k,\ell}\}_{(k,\ell) \in [r_2] \times [r_1]}$, i.e.,

$$\boldsymbol{\Pi}_\ell = \sum_{k=1}^{r_2} \mathbf{S}_{k,\ell} \mathbf{S}_{k,\ell}^T, \quad \ell = 1, \dots, r_1. \quad (4.4)$$

Optimization is then performed on the augmented Lagrangian obtained by substituting this nested low rank factorization. The power of the hierarchical low-rank idea lies in its scalability, and the fact that it can be applied recursively to higher-degree moment matrices, thus potentially enabling scalable optimization over higher rounds of semidefinite programming hierarchies. Function and gradient derivation are given for the hierarchical factorization on the penalty formulation (4.3) in appendix A. The derivations on the augmented Lagrangian formulation follow the exact same idea. The next section discusses how the combinatorial total symmetry constraints can be enforced efficiently.

4.4 Replacing total symmetry with trace relations

In this section, we discuss three trace relations whose linearizations can be used as scalable substitutes to the more computationally expensive third and fourth order total symmetry constraints. We provide numerical evidence that whenever the factorized gradient method works, enforcing those trace relations in place of total symmetry works just as well, yet reduces the computational (combinatorial) cost of those constraints from $\mathcal{O}(N^4)$ to $\mathcal{O}(1)$. Those relations seem to work best when applied to the multilevel low rank decomposition introduced in section 4.3.

The total symmetry constraints are used to encode correspondence of the entries of \mathbf{M} that correspond to the same moments (see the discussion in section 1.4). When applied on the third and fourth moments tensors, those constraints enforce equality between any permutation of the multi-index. I.e if $\mathbf{M}^{(3)}$ and $\mathbf{M}^{(4)}$ encode the third and fourth order blocks in \mathbf{M} , then those constraints require that for any 3-tuple (i, j, k) and permutation π , $(\mathbf{M}^{(3)})_{i,j,k} = (\mathbf{M}^{(3)})_{\pi(i,j,k)}$. Similarly, on the fourth order block, for any 4-tuple (i, j, k, ℓ) and any permutation π , the moments matrix must satisfy $(\mathbf{M}^{(4)})_{i,j,k,\ell} = (\mathbf{M}^{(4)})_{\pi(i,j,k,\ell)}$.

One of the implications of total symmetry constraints is that the contraction of any fourth order block does not depend on the indices over which this contraction is taken. In other words, the sum $\sum_{i=1}^N \mathbf{M}_{iijk}^{(4)}$ is the same as the sum $\sum_{i=1}^N \mathbf{M}_{jii k}^{(4)}$, and so is it for any of the sums $\sum_{i=1}^N \mathbf{M}_{\pi(iijk)}^{(4)}$ for any permutation operator $\pi : [N] \mapsto [N]$. When assuming that the tensor $\mathbf{M}^{(4)}$ is rank one, that is $\mathbf{M}_0^{(4)} = \text{vec}(\mathbf{z}_0 \otimes \mathbf{z}_0) \text{vec}(\mathbf{z}_0 \otimes \mathbf{z}_0)^T$, those constraints can be used to derive interesting trace relations on the second order tensor $\mathbf{M}^{(2)}$. For $\mathbf{M}^{(4)} = \mathbf{M}^{(2)} \otimes \mathbf{M}^{(2)}$, $\sum_{i=1}^N \mathbf{M}_{iijk}^{(4)} = \sum_{i=1}^N \mathbf{M}_{jii k}^{(4)}$ in particular implies the following trace relation on $\mathbf{M}^{(2)}$,

$$\text{Tr}(\mathbf{M}^{(2)})\mathbf{M}^{(2)} = (\mathbf{M}^{(2)})^2. \quad (4.5)$$

Linearizing this trace relation brings us back to enforcing equality of the contractions $\sum_{i=1}^N \mathbf{M}_{iijk}^4 - \sum_{i=1}^N \mathbf{M}_{jii k}^4 = 0$. Moreover this first contraction can be enforced very efficiently on the (hierarchical) low rank factorization of $\mathbf{M}^{(4)}$, $\mathbf{M}^{(4)} = \sum_{k=1}^r \mathbf{S}_k \mathbf{S}_k$,

$$\sum_{k=1}^{r_1} \text{Tr}(\text{Mat}(\mathbf{\Pi}_k)) \text{Mat}(\mathbf{\Pi}_k) = \sum_{k=1}^r \text{Mat}(\mathbf{\Pi}_k) \text{Mat}(\mathbf{\Pi}_k) \quad (4.6)$$

The natural extension to (4.5) is to go one step further and take a second contraction with respect to the indices remaining in this first constraint. This gives a second trace relation that requires the trace of the squared matrix to match the square of this matrix trace,

$$\text{Tr}(\mathbf{M}^{(2)})^2 = \text{Tr}((\mathbf{M}^{(2)})^2). \quad (4.7)$$

This last relation reduces the set of $\mathcal{O}(N^4)$ symmetry constraints to a single constraint that can be enforced efficiently on the hierarchical low rank factors. Note that when enforced on positive semidefinite matrices, (4.7) is in fact equivalent to enforcing an exact rank one constraint, as it requires $(\sum_i \lambda_i)^2 = \sum_i \lambda_i^2$, for $\lambda_i \geq 0$,

$$\{X \in \mathbb{S}_N^+ : \text{Tr}(\mathbf{X})^2 = \text{Tr}(\mathbf{X}^2)\} = \{X \in \mathbb{S}_N^+ : \text{rank}(\mathbf{X}) \leq 1\} \quad (4.8)$$

Again, Equation (4.7) can be written compactly for the low rank as well as for the hierarchical low rank formulations. For this last factorization, we get

$$\sum_r \|\mathbf{S}_r\|_F^2 \mathbf{S}_r \mathbf{S}_r^* = \sum_r \mathbf{S}_r (\mathbf{S}_r^* \mathbf{S}_r) \mathbf{S}_r \quad (4.9)$$

$$= \sum_r \mathbf{S}_r \mathbf{\Delta}_{r,r} \mathbf{S}_r^* \quad (4.10)$$

where we let $\mathbf{\Delta}_{r,r}$ denote the r by r matrix encoding the products $\mathbf{S}_r^* \mathbf{S}_r$.

As the fourth order symmetry constraints are not as important as a means to express the first order monomials as the third order symmetry constraints (especially in the reduced framework of Corollary 3), one could argue that replacing those constraints with a simpler contraction does not have a significant impact on the outcome of the relaxation. It is in fact possible to consider a trace contraction for third order total symmetry constraints as well. At order 3, following from the constraints $\mathbf{M}_{ijk}^{(3)} = \mathbf{M}_{\pi_1(i,j,k)}^{(3)}$, one possible contraction can be taken over the first 2 indices in the third order tensor $\mathbf{M}^{(3)}$, for any permutation π . This gives the following relation

$$\text{Tr}(\mathbf{M}^{(2)}) \mathbf{M}^{(1)} = \mathbf{M}^{(2)} \mathbf{M}^{(1)} \quad (4.11)$$

where we again let $\mathbf{M}^{(2)}$ denote the matrix encoding the second order moments $\mathbf{M}^{(2)} \approx \mathbf{M}^{(1)} \otimes \mathbf{M}^{(1)}$ and $\mathbf{M}^{(1)}$ denotes the vector of first order moments. Again, this third relation can be expressed compactly for both the low rank and hierarchical low rank formulations. For this last factorization, we can write

$$\sum_{r=1}^{r_1} \text{Tr}(\text{Mat}(\mathbf{\Pi}_r)) (\mathbf{R}_r) = \sum_{r=1}^{r_1} \text{Mat}(\mathbf{\Pi}_r) \mathbf{R}_r \quad (4.12)$$

In each of these examples, we provide the evolution of each of the error terms appearing in the Lagrangian (4.2) with the iterations. We also represent the global misfit, the trace and the evolution of the Lagrange multipliers as the main recovery guarantees provided so far on low rank factorization require those multipliers to remain bounded.

To study the result of replacing third and fourth order total symmetry constraints by the trace relations above, we apply those relations on the simple example (1.6) for $\delta = .5$ with the single low rank and hierarchical low rank factorization. The results are shown in Figs. 9 (low rank) and 10 (hierarchical low rank). When considering the simpler factorized gradient approach, it seems that replacing the full set of 4^{th} order symmetry constraints with the corresponding trace contraction can lead to a slight reduction in the accuracy. The total symmetry constraints are not entirely satisfied as highlighted by Fig. 9 and this results in a partial recovery of the fourth order tensor. It remains possible to extract the solution \mathbf{X}_0 from the second order moments. A comparison of the iterations of Fig. 9 and 10 seems to suggest that replacing total symmetry by the relations (4.5), (4.7) and (4.11) performs best when used on the hierarchical low rank factorization.

Generally speaking, it again seems that when the factorized gradient approach converges, which typically happens on problems that are not too difficult (i.e δ not too small), it always seems possible to replace the combinatorial Total Symmetry constraints by the more tractable trace contractions on both the 3^{rd} and 4^{th} order tensors, and to recover the solution for both the low rank and hierarchical low rank frameworks. As we don't have empirical evidence that choosing contraction (4.5) over contraction (4.7) will lead to better convergence properties, we will always favor the former over the latter, as this one reduces to a single equation. On the remaining large scale examples of this paper, we thus always replace total symmetry constraints with relation 4.11 (third order moments) and relation (4.7) (fourth order moments).

To illustrate the interest of the combination of a multi-level low rank decomposition and of the trace relations (4.7) and (4.11) for large rank one recovery problems, we now apply this combination on a first large

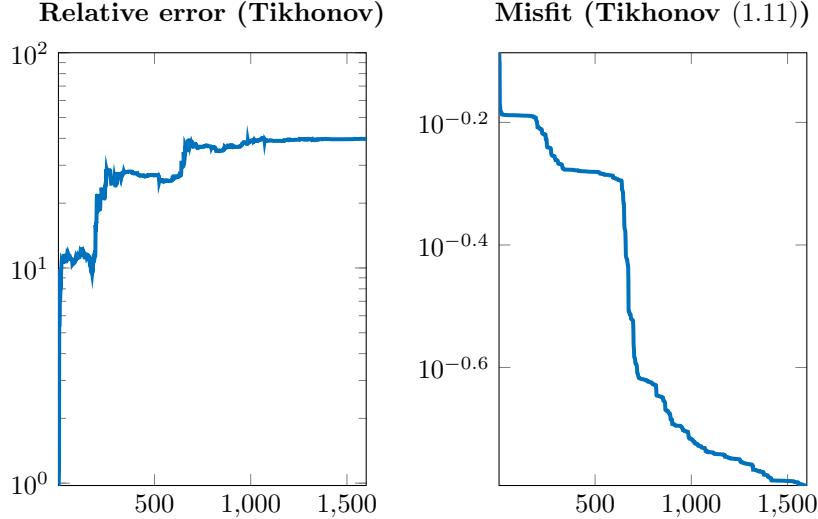


Figure 6: It is possible to find instances of the rank one matrix completion problem for which even the ridge regression formulation will not be able to return the global minimizer. Such instances are more frequent when considering matrices of large size. In this particular case, ridge regression is clearly shown to converge to the wrong minimizer. The higher semidefinite relation, on the other hand, returns the true solution, even when considering a hierarchical low rank factorization with strong constraint on the rank (2 in this case) (see Fig. 7). The matrix considered here is 20×20 with moments bounded as $\delta \leq (X_0)_{ij} \leq 1$ for $\delta = 0.01$ and a mask Ω whose underlying bipartite graph is shown in Fig. 8. The corresponding rank constrained iterations for the higher order relaxation are displayed in Fig. 7.

scale example. On this example, both nuclear norm and ridge regression fail at recovering the solution. For this example, we take the moments to be bounded as $\delta \leq \mathbf{z}_0^\alpha \leq 1$, for any $|\alpha| \leq 1$ and with $\delta = .1$ (i.e with a possibly larger gap between smallest and largest entries). The bipartite graph defining the measurements is represented in Fig. 8. This graph is generated at random while enforced to span the $(m + n)$ vertices with a minimal number of edges. For this particular problem, the solution returned by nuclear norm minimization gives a relative error $\|\mathbf{X} - \mathbf{X}_0\|_F / \|\mathbf{X}_0\|_F = 0.7382$. The iterates returned by the ridge regression formulation (1.11) are displayed in Fig. 6 (relative error and data misfit. In this case there is no need for any regularization as the problem is noiseless).

The iterations following from optimization over the hierarchical low rank augmented Lagrangian with the trace relations (4.7) and (4.11) is displayed in Fig. 7.

4.5 Subsampling the higher order affine constraints

A last computational bottleneck that hinders the application of the hierarchical formulation of section (4.3) to larger matrices comes from the higher-order affine constraints. Those constraints have the form $\mathbf{A}\mathbf{T}\mathbf{T}^* = 0$ where \mathbf{T} denotes the whole low rank factor of size $\mathcal{O}(N^2)$ and \mathbf{A} simply applies the affine constraints to the second order part of \mathbf{T} , columnwise. To further reduce the computational cost, we propose to draw smaller $\mathcal{O}(N)$ "batches" of moments $\mathcal{S}_i \subseteq [N] \times [N]$ from the full set of second order moments. We then minimize the resulting reduced augmented Lagrangian functions defined from each \mathcal{S}_i sequentially. Let \mathbf{T}_0 denote the initial iterate chosen at random. The procedure can be summarized as follows

1. Randomly select a subset \mathcal{S}_i of size $\mathcal{O}(N)$, without replacement, from the set of all second order

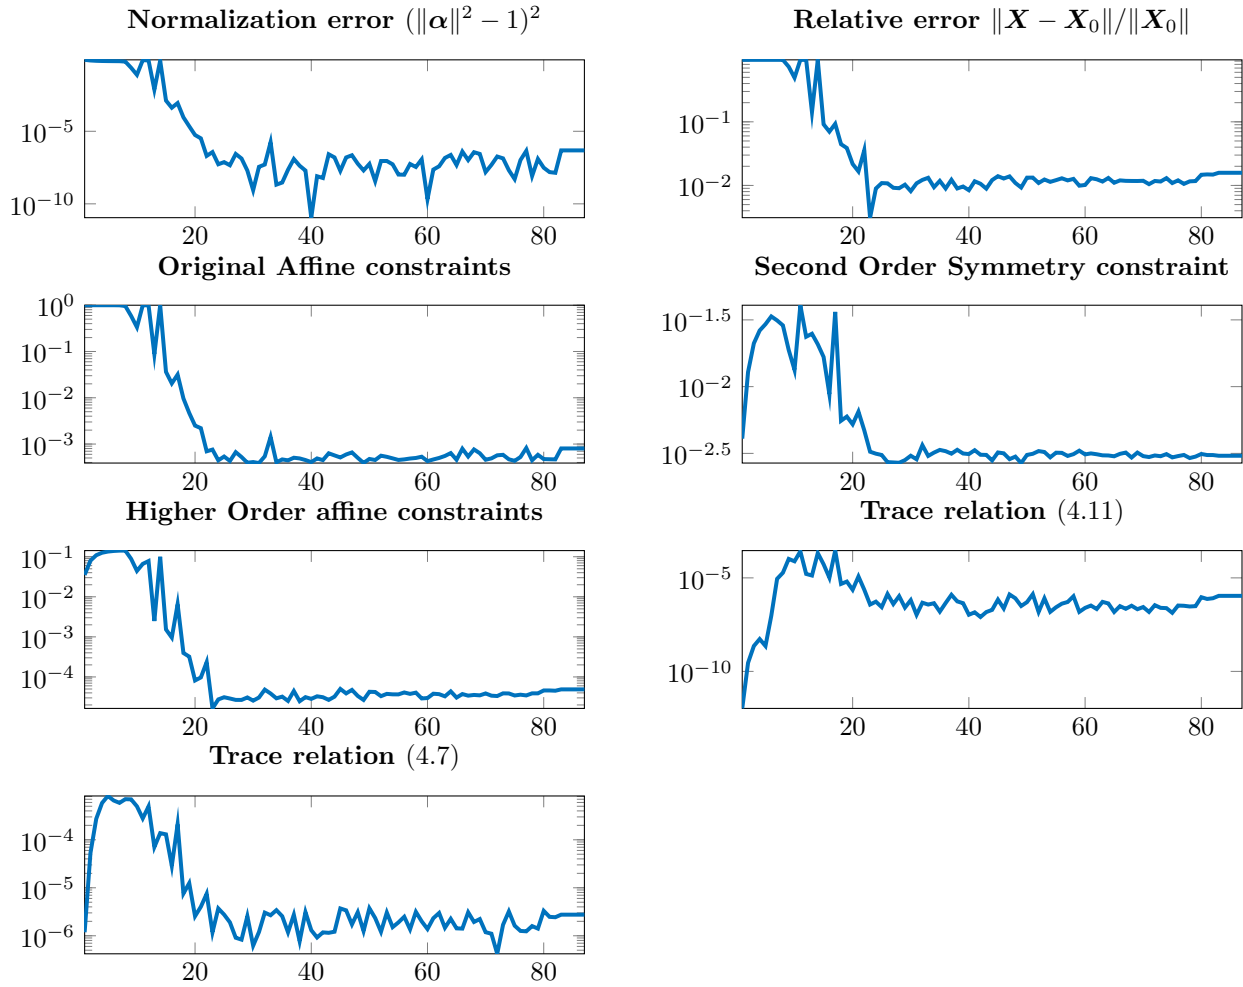


Figure 7: Evolution of the various constraints and error terms appearing in the augmented Lagrangian, for the hierarchical low rank factorization (here r_1 and r_2 are set to 4) with Trace contractions (4.7) and (4.11) for the 20×20 example of Fig. 6 with measurements defined from the bipartite graph of Fig 8.




Figure 8: Bipartite graph corresponding to the 20×20 example of Figs 6 and 7 used to illustrate the failure of nuclear norm minimization and ridge regression.

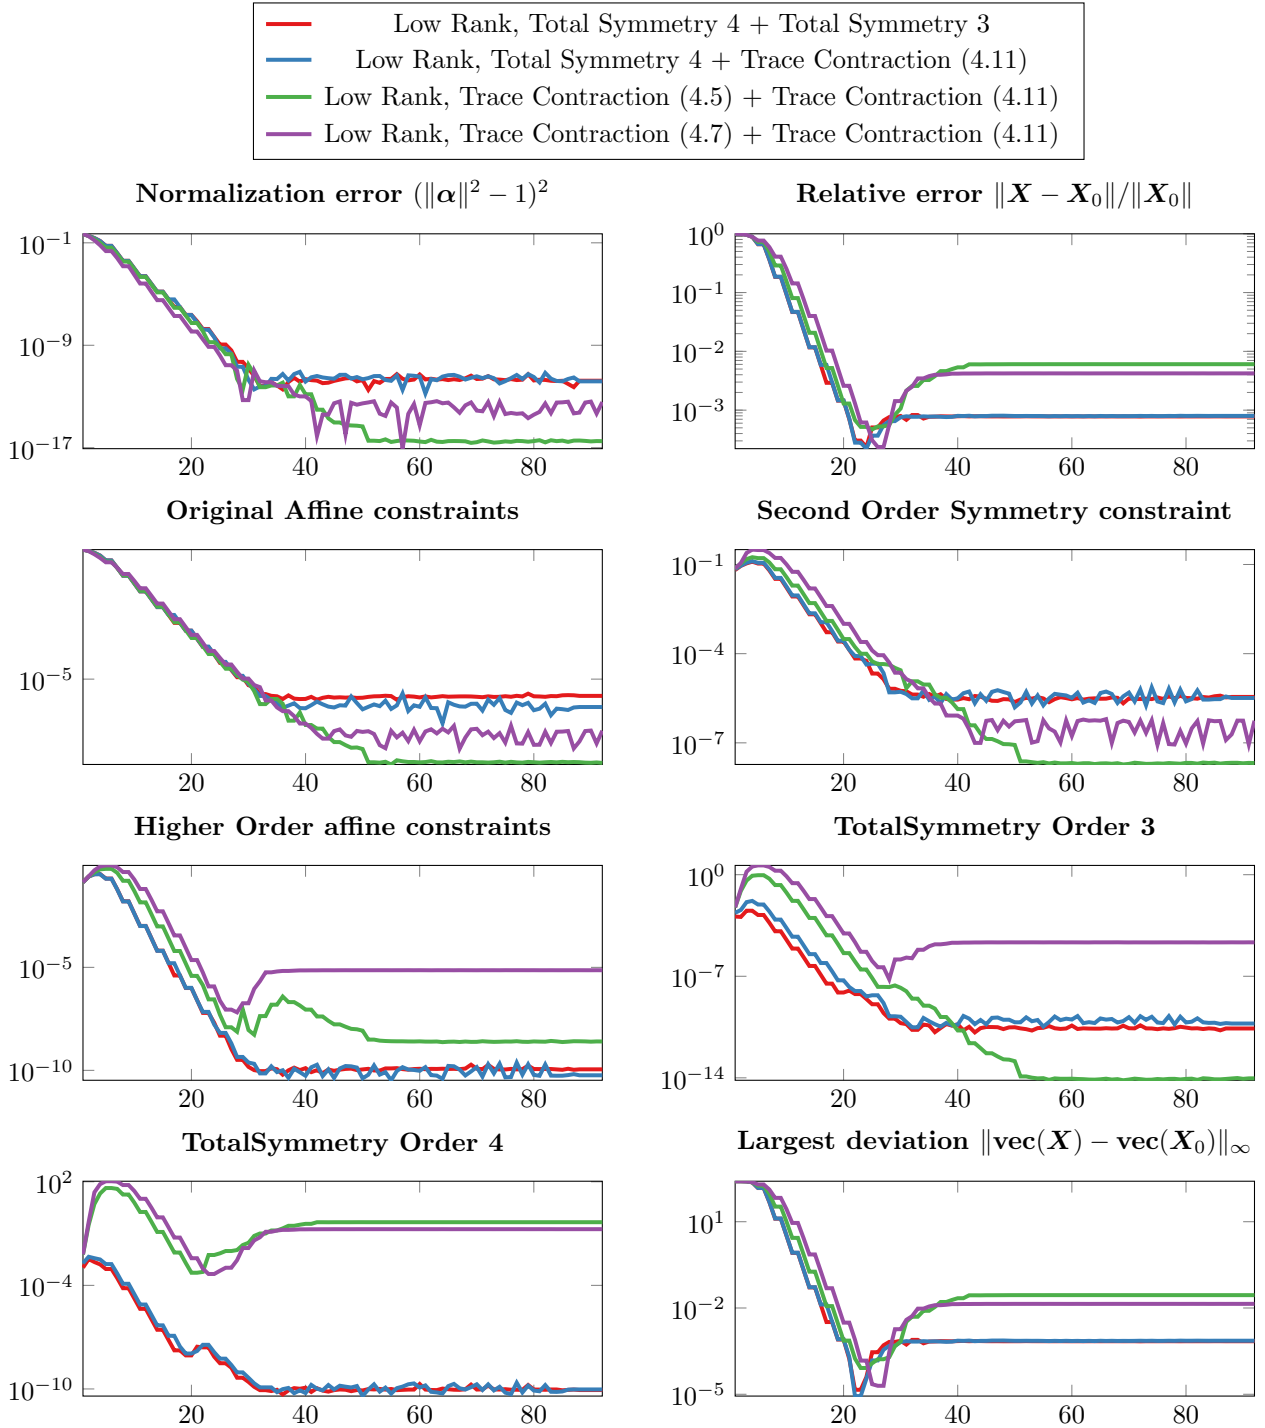


Figure 9: Comparison of the various Trace contractions (4.5), (4.7) and (4.11) on the simple low rank decomposition, as substitute for the combinatorial total symmetry constraints. As shown by the evolution of the relative error, all 4 approaches lead to estimates that are very close to the optimal solution M_0 . Adding some of the 4th order Total symmetry might help improving the estimate although the solution returned by the 4th order trace contraction already are sufficiently close to the ground truth to recover the solution.

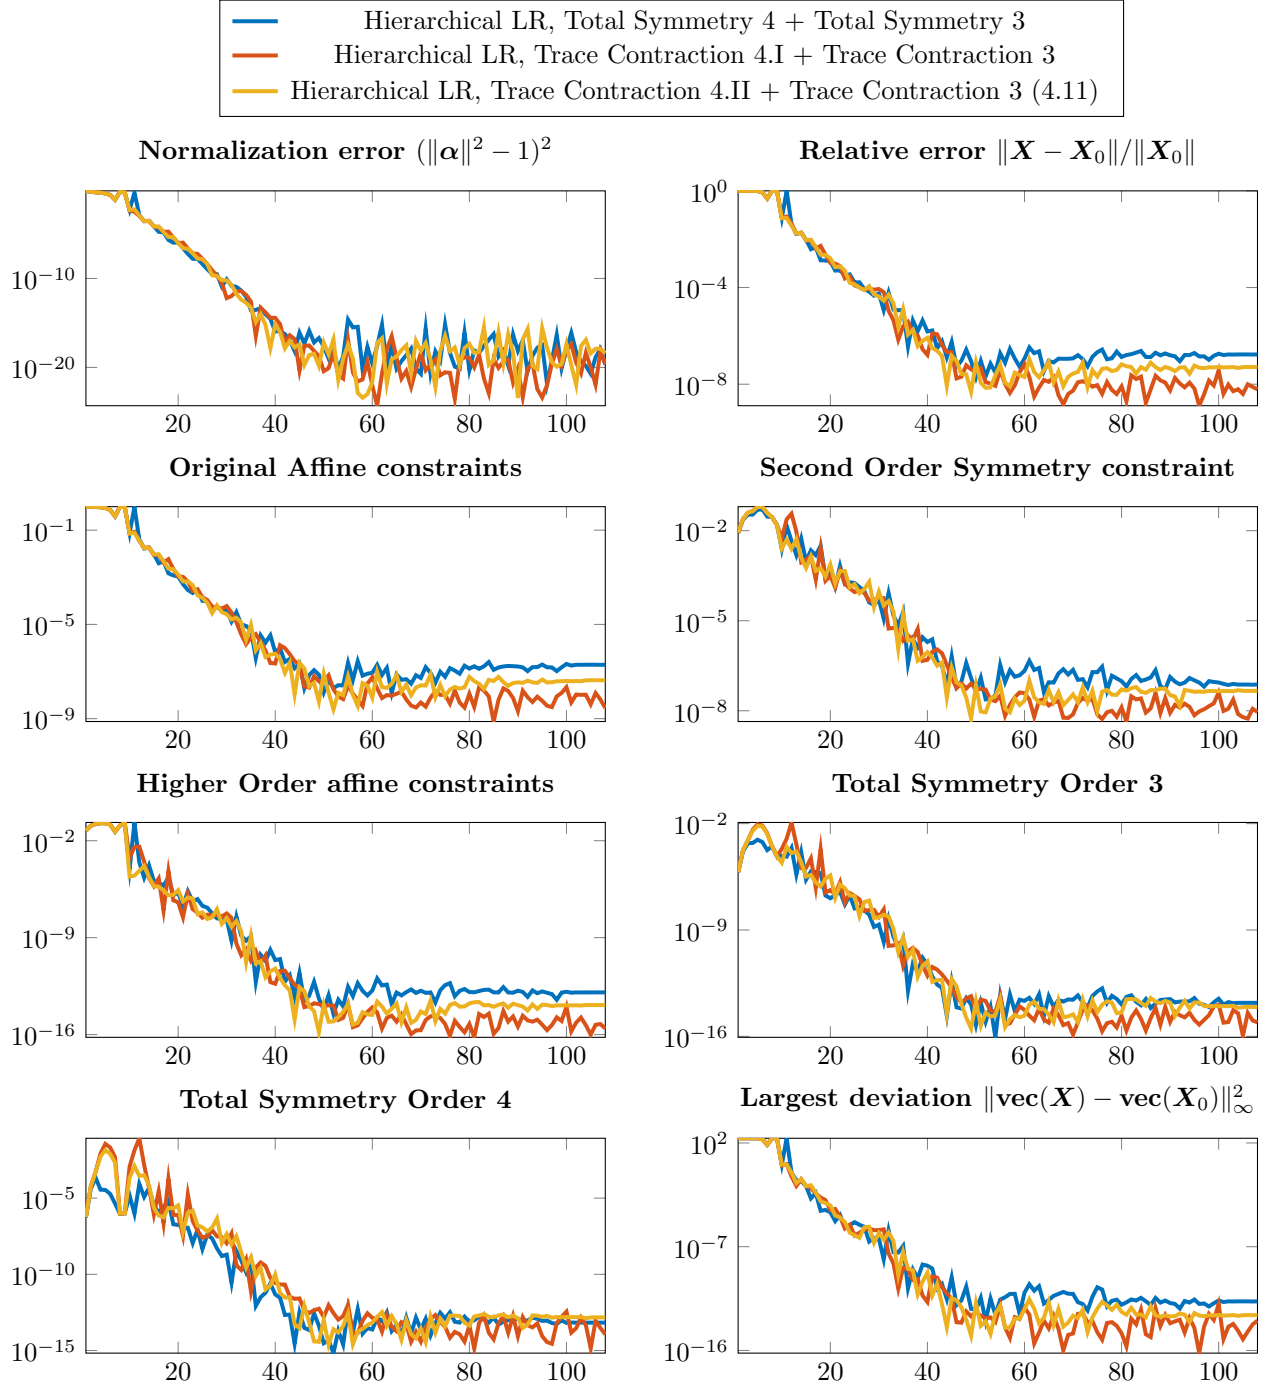


Figure 10: Minimization of the augmented Lagrangian, on the hierarchical low rank factorization (here we consider ranks 4 and 2) for various combinations of the total symmetry constraints and trace relations. Note that total symmetry and trace contractions are never imposed simultaneously, i.e., either we impose total symmetry or the corresponding trace contraction. The evolution of the total symmetry constraints shows that on the hierarchical factorization, enforcing the trace relations of section 4.4 is sometimes exactly equivalent to requiring total symmetry by means of the combinatorial constraints.

moments $[N] \times [N]$.

2. Minimize the augmented Lagrangian (4.2), considering only the higher order affine constraints of the form $\mathbf{A} \mathbf{T} \mathbf{T}_{\mathcal{S}_i}^* = 0$, where $\mathbf{T}_{\mathcal{S}_i} \equiv [\boldsymbol{\alpha}^*, \mathbf{R}^*, \boldsymbol{\Pi}_{\mathcal{S}_i}^*]$ resulting from the second order moments appearing in \mathcal{S}_i .
3. Let \mathbf{T}^\sharp denote the solution resulting from step 2. If $\|\mathcal{A}(\mathbf{T}^\sharp(\mathbf{T}^\sharp)^*) - \mathbf{b}\|$ is sufficiently small then stop. Otherwise, repeat step 1 with $\mathbf{T}_0 \leftarrow \mathbf{T}^\sharp$.

To illustrate this last algorithm, we provide numerical experiments on a 100×100 matrix. For the algorithm to be fully efficient, we combine the trace relations (4.7) (on the fourth order block) and (4.11) (on the third order block), and take advantage of the subsampling scheme discussed above. On a 100×100 matrix, the factorized gradient method would require storing matrices of size at least $\mathcal{O}(N^3) = 8e6$, thus leading to poor performance in terms of runtime. On convex solvers such as CVX or GLOTIPOLY, this example would require storing matrices of size $\mathcal{O}(N^4)$ ($\mathcal{O}(N^3)$ in the reduced formulation of corollary (3)). The iterations on this example are displayed in Fig. 11 and the corresponding bipartite graph used as mask is shown in Fig. 12. Solving this problem takes no more than 10 mins on a laptop with 2 GHz Intel Core i5.

5 Stable completion of rank-one tensors

Theorem 1 and 2 both have a direct extension to the tensor completion problem. Given a rank one d -tensor $T \in \mathbb{R}^{n^d}$, one can always write $T = x_1 \otimes x_2 \otimes \dots \otimes x_d$ with $(x_2)_1 = (x_3)_1 = \dots = (x_{d-1})_1 = 1$. To this tensor, one can associate a d -uniform hypergraph $\mathcal{H}(\mathcal{V}, \mathcal{E})$ whose set of vertices is given by the set of indices associated to each dimension and whose edges are defined from the measurements $\mathcal{P}_\Omega(T)$. For this hypergraph, we consider the following property.

Definition 1 (Definition 1 in [7]). *Let $\mathcal{H}(\mathcal{V}, \mathcal{E})$ be a d -uniform hypergraph on $N_1 \times N_2 \times \dots \times N_d = |\mathcal{V}|$ vertices. A sequence $E_1, \dots, E_K \in \mathcal{E}$ of hyperedges is called a propagation sequence if for any $1 \leq \ell \leq K$, $|E_{\ell+1} \cap_{i=1}^\ell E_\ell| = d-1$. If the hypergraph has a propagation sequence, then it is called propagation connected.*

Stable deterministic completion of rank-one, propagation-connected tensors directly follows from the definition of propagation connectivity and the normalization $(x_2)_1 = \dots = (x_{d-1})_1$ of the vectors in the decomposition. Indeed, as in section 1, let $\mathbf{z} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1})$ denote the concatenation of the monomials arising from the tensor decomposition. For every monomial z_L there always exists a sequence of hyperedges, each defined from its corresponding set of indices in $[N_1] \times [N_2] \times \dots \times [N_d]$, such that $E_1 = (1, 1, \dots, i_1)$, $E_2 = (1, \dots, 1, i_1, i_2)$, \dots and $|E_i \cap E_{i+1}| = d-1$. Let us denote by $z_{i_1}, z_{i_2}, \dots, z_L$ the variables defined from the chain as $i_1 \in E_1$ and $i_k \in E_{k-1} \cap E_k$ for $k = 2, \dots, L$ with $i_L = L$. Assume that we can express the polynomials $z_{i_k} - (z_0)_{i_k}$ for $k = 1, \dots, L-1$, then the k^{th} canonical polynomial $z_L - (z_0)_L$ can be expressed from $z_{i_{L-1}}$ and the constraints corresponding to the edges E_L and E_{L-1} as

$$\begin{aligned} (z_L - (z_0)_L) \prod_{i \in E_{L-1}} (z_0)_i &= z_{E_{L-1} \setminus (E_L \cap E_{L-1})} \left(\prod_{i \in E_L} z_i - \prod_{i \in E_L} (z_0)_i \right) - z_L \left(\prod_{i \in E_{L-1}} z_i - \prod_{i \in E_{L-1}} (z_0)_i \right) \\ &\quad + z_{E_{L-1} \setminus (E_L \cap E_{L-1})} \prod_{i \in E_L} (z_0)_i \end{aligned}$$

This discussion naturally leads to the following corollary (a corresponding stability result can be derived).

Corollary 4. *Let $T \in \mathbb{R}^{N_1 \times \dots \times N_d}$ denote an order d rank-one tensor. Assume that we are given the entries T_{i_1, \dots, i_d} for $(i_1, i_2, \dots, i_d) \in \Omega$. Further assume that the hypergraph corresponding to Ω is propagation connected. Then the tensor can be efficiently completed through d rounds of semidefinite programming relaxation with minimization of the trace norm of the moments matrix.*

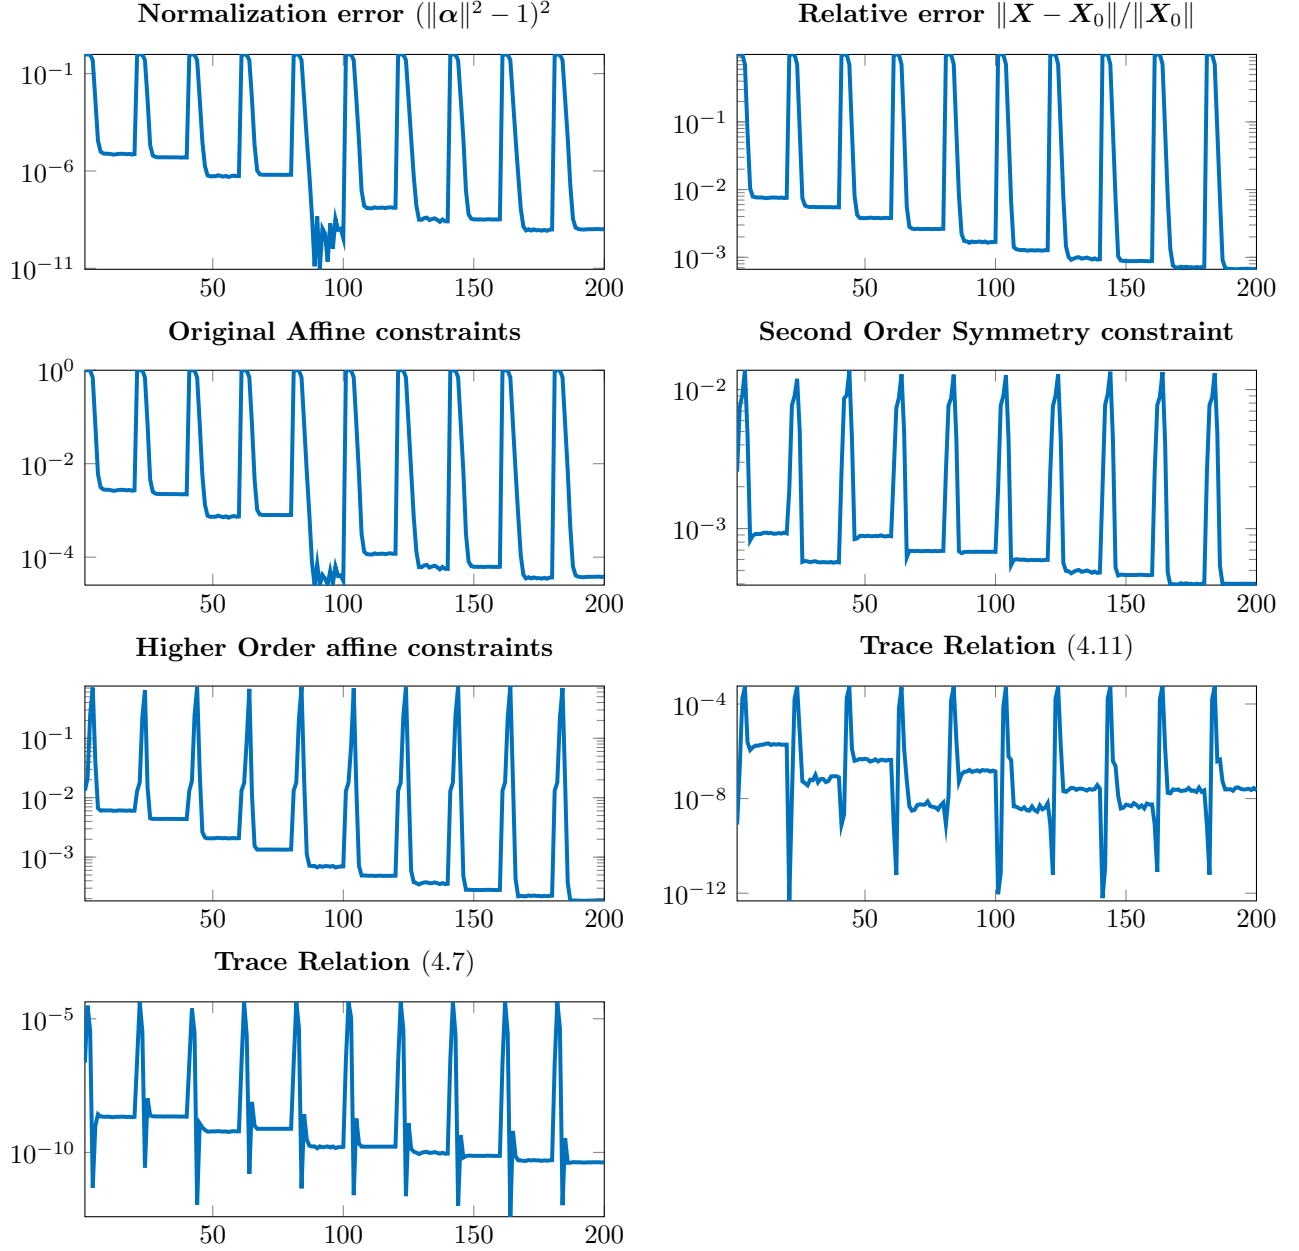


Figure 11: When considering sufficiently large (e.g. 100×100) completion problems, minimizing the whole set of higher order affine constraints is not efficient anymore because those constraints require storing matrices of size $\mathcal{O}(N^3)$ while the hierarchical low rank decomposition only requires storing matrices of size $\mathcal{O}(N)$. For this reason, we divide the vector of second order moments into smaller batches \mathcal{S}_i of size $\mathcal{O}(m+n)$ sampled at random, and minimize the resulting augmented Lagrangians sequentially as explained in section 4.5. Each jump in the figures above corresponds to a resampling of the moments. For each resampling of the moments, we reset the relative weight of the trace with respect to the misfit, whence the jump occurring at the transition between two batches. Here the moments of \mathbf{X}_0 are controlled as $\delta \leq (X_0)_{ij} \leq 1$ with $\delta = 0.25$ and the ranks of the hierarchical factorization are set as $r_1 = r_2 = 2$. For large matrices, only the trace relations (4.11) and (4.7) are represented as the total symmetry constraints are too expensive to compute and are thus not enforced.




Figure 12: Bipartite graph corresponding to the 100×100 example of Fig. 11. The number of edges is $\mathcal{O}(N) \approx 200$. The graph is generated at random while required to span the whole set of vertices minimally.

A Hierarchical low rank gradient and function

Before introducing the compressed function and gradient resulting from the multi-level low rank encoding of the moments matrix, recall that we normalize the first entry of \mathbf{x} and work with the matrix

$$\mathbf{X} = \mathbf{y}(1, \mathbf{x}^T) = \begin{pmatrix} y_1 & x_1 y_1 & \dots & x_n y_1 \\ y_2 & x_1 y_2 & \dots & x_n y_2 \\ \vdots & \vdots & \ddots & \vdots \\ y_m & x_1 y_m & \dots & y_m x_n \end{pmatrix} \quad (\text{A.1})$$

The trace can be computed efficiently as

$$\begin{aligned} \text{Trace}(\mathbf{M}) &= \|\boldsymbol{\alpha}\|^2 + \|\mathbf{R}\|_F^2 + \sum_{r=1}^{R_1} \langle \Pi_r, \Pi_r \rangle \\ &= \|\boldsymbol{\alpha}\|^2 + \|\mathbf{R}\|_F^2 + \sum_{r=1}^{R_1} \left\langle \sum_{r'=1}^{R_2} S_{r,r'} S_{r,r'}^T, \sum_{r'=1}^{R_2} S_{r,r'} S_{r,r'}^T \right\rangle \\ &= \|\boldsymbol{\alpha}\|^2 + \|\mathbf{R}\|_F^2 + \sum_{r=1}^{R_1} \sum_{k=1}^{R_2} \sum_{\ell=1}^{R_2} |\langle S_{r,k}, S_{r,\ell} \rangle|^2 \end{aligned}$$

The gradient for the trace can be computed efficiently as

$$\partial_{\boldsymbol{\alpha}} = 2\boldsymbol{\alpha}, \quad \partial_R = 2R \quad \text{and} \quad \partial_{S_{r,k}} = 4 \sum_{r'=1}^{R_2} S_{r,r'} S_{r,r'}^T S_{r,k}.$$

There are two sets of structural constraints. The first set enforces the equality between RR^T and $\sum_{r=1}^{R_1} \alpha_r \Pi_r = \sum_{r=1}^{R_1} \alpha_r \sum_k S_{r,k} S_{r,k}^T$. This first set can be expressed compactly as the following Frobenius contribution

$$\begin{aligned} \left\| RR^T - \sum_{r=1}^{R_1} \alpha_r \left(\sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T \right) \right\|_F^2 &= \|R^T R\|_F^2 + \sum_{r=1}^{R_1} \sum_{r'=1}^{R_1} \alpha_r \alpha_{r'} \left(\sum_{k=1}^{R_2} \sum_{k'=1}^{R_2} |\langle S_{r,k}, S_{r',k'} \rangle|^2 \right) \\ &\quad - 2 \sum_{r=1}^{R_1} \alpha_r \left(\sum_{k=1}^{R_2} \|S_{r,k}^T R\|_F^2 \right). \end{aligned}$$

The resulting gradient contribution can be computed efficiently as,

$$\begin{aligned}\partial_R &= 2 \left[(RR^T - \sum_{r=1}^{R_1} \alpha_r (\sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T)) R + R^T (RR^T - \sum_{r=1}^{R_1} \alpha_r (\sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T)) \right] \\ \partial_{S_{r,k}} &= -2\alpha_r \left(RR^T - \sum_{r=1}^{R_1} \alpha_r \left(\sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T \right) \right) S_{r,k} - 2\alpha_r S_{r,k}^T \left(RR^T - \sum_{r=1}^{R_1} \alpha_r \left(\sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T \right) \right) \\ \partial_{\alpha_r} &= -2 \left\langle RR^T - \sum_{r=1}^{R_1} \alpha_r \left(\sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T \right), \sum_{k=1}^{R_2} S_{r,k} S_{r,k}^T \right\rangle\end{aligned}$$

The second set of structural constraints enforces equality between corresponding third and fourth order monomials. As an example, we have $\Pi_{i,j} R_k^T = \Pi_{i,k} R_j^T$ or similarly $\Pi_{i,j} R_k^T = \Pi_{i,k} R_j^T$. Those constraints are first expressed through permutations of third and fourth indices. When dealing with third and fourth order monomials equivalences, only some of the permutation have to be explicitly enforced. The others are naturally encoded through the positive semidefinite constraint. The explicit ones are listed below. For third order monomials we have

$$(z_i z_j) z_k = (z_i z_k) z_j \quad (\text{A.2})$$

$$= (z_j z_k) z_i \quad (\text{A.3})$$

Equivalently, for fourth order monomials, we can only retain the following relations,

$$(z_i z_j)(z_k z_\ell) = (z_i z_k)(z_j z_\ell) \quad (\text{A.4})$$

$$= (z_i z_\ell)(z_j z_k) \quad (\text{A.5})$$

Let each of the permutations for third and fourth order monomials that encode the structural constraints be denoted as $\pi^{(3)}$ and $\pi^{(4)}$ where we let $\pi^{(3)} : (m+n-1)^3 \mapsto (m+n-1)^3$ with $\pi^{(3)}(i_1, i_2, i_3) = (\pi_1^{(3)}, \pi_2^{(3)}, \pi_3^{(3)})$ and similarly for $\pi^{(4)}$. The resulting constraints in the framework of formulation (4.1) can read as

$$\min \left\| \sum_{\ell=1}^{R_1} \Pi_\ell(i_1, i_2) R(i_3, \ell) - \sum_{\ell=1}^{R_1} \Pi_\ell(\pi_1^{(3)}, \pi_2^{(3)}) R(\pi_3^{(3)}, \ell) \right\|_F^2 \quad (\text{A.6})$$

Fourth order relations can be expressed in exactly the same way,

$$\min \left\| \sum_{\ell=1}^{R_1} \Pi_\ell(i_1, i_2) \Pi_\ell(i_3, i_4) - \sum_{\ell=1}^{R_1} \Pi_\ell(\pi_1^{(4)}, \pi_2^{(4)}) \Pi_\ell(\pi_3^{(4)}, \pi_4^{(4)}) \right\|_F^2 \quad (\text{A.7})$$

If we expand the second order low rank factorizations, the two structural contributions (A.6) and (A.7) respectively read as

$$\left\| \sum_{\ell=1}^{R_1} \left(\sum_{k=1}^{R_2} S_{\ell,k}[i_1] S_{\ell,k}[i_2] \right) R[i_3, \ell] - \sum_{\ell=1}^{R_1} \left(\sum_{k=1}^{R_2} S_{\ell,k}[\pi_1^{(3)}] S_{\ell,k}[\pi_2^{(3)}] \right) R[\pi_3^{(3)}, \ell] \right\|_F^2 \quad (\text{A.8})$$

as well as

$$\begin{aligned}\min & \left\| \sum_{\ell=1}^{R_1} \left(\sum_{k=1}^{R_2} S_{\ell,k}[i_1] S_{\ell,k}[i_2] \right) \left(\sum_{k'=1}^{R_2} S_{\ell,k'}[i_3] S_{\ell,k'}[i_4] \right) \right. \\ & \left. - \sum_{\ell=1}^{R_1} \left(\sum_{k=1}^{R_2} S_{\ell,k}[\pi_1^{(4)}] S_{\ell,k}[\pi_2^{(4)}] \right) \left(\sum_{k'=1}^{R_2} S_{\ell,k'}[\pi_3^{(4)}] S_{\ell,k'}[\pi_4^{(4)}] \right) \right\|_F^2\end{aligned} \quad (\text{A.9})$$

This equivalence of monomials is the most expensive step in the minimization. We let E^3 and E^4 be defined as

$$E_3 = \sum_{\ell=1}^{R_1} \Pi_\ell(i_1, i_2) R(i_3, \ell) - \sum_{\ell=1}^{R_1} \Pi_\ell(\pi_1^{(3)}, \pi_2^{(3)}) R(\pi_3^{(3)}, \ell) \quad (\text{A.10})$$

$$E_4 = \sum_{\ell=1}^{R_1} \Pi_\ell(i_1, i_2) \Pi_\ell(i_3, i_4) - \sum_{\ell=1}^{R_1} \Pi_\ell(\pi_1^{(4)}, \pi_2^{(4)}) \Pi_\ell(\pi_3^{(4)}, \pi_4^{(4)}) \quad (\text{A.11})$$

for which the contributions of (A.6) and (A.7) to the gradient accumulate for each ℓ and pairs of 3-tuple of indices $\{(i_1, i_2, i_3), (\pi_1^{(3)}, \pi_2^{(3)}, \pi_3^{(3)})\}$ appearing in the set of structural constraints as

$$\partial_R[i_3, \ell] \leftarrow \partial_R[i_3, \ell] + \Pi_\ell(i_1, i_2) E_3, \quad \partial_R[\pi_3^{(3)}, \ell] \leftarrow \partial_R[i_3, \ell] + \Pi_\ell(\pi_1^{(3)}, \pi_2^{(3)}) E_3 \quad (\text{A.12})$$

$$\partial_{\Pi_\ell}[i_1, i_2] \leftarrow \partial_{\Pi_\ell}[i_1, i_2] + R[i_3, \ell] E_3, \quad \partial_{\Pi_\ell}(\pi_1^{(3)}, \pi_2^{(3)}) \leftarrow \partial_{\Pi_\ell}(\pi_1^{(3)}, \pi_2^{(3)}) + R(\pi_3^{(3)}, \ell) E_3 \quad (\text{A.13})$$

Accordingly, for the fourth order contribution, we simply accumulate the contributions arising from each of the norms in (A.7) for each pair of four-tuples $\{(i_1, i_2, i_3, i_4), (\pi_1^{(4)}, \pi_2^{(4)}, \pi_3^{(4)}, \pi_4^{(4)})\}$ and each rank index ℓ as

$$\partial_{\Pi_\ell}[i_1, i_2] \leftarrow \partial_{\Pi_\ell}[i_1, i_2] + \Pi_\ell(i_3, i_4) E_4, \quad \partial_{\Pi_\ell}[\pi_1^{(4)}, \pi_2^{(4)}] \leftarrow \partial_{\Pi_\ell}[\pi_1^{(4)}, \pi_2^{(4)}] - \Pi_\ell(\pi_1^{(3)}, \pi_2^{(3)}) E_4 \quad (\text{A.14})$$

$$\partial_{\Pi_\ell}[i_3, i_4] \leftarrow \partial_{\Pi_\ell}[i_3, i_4] + \Pi_\ell(i_1, i_2) E_4, \quad \partial_{\Pi_\ell}[\pi_3^{(4)}, \pi_4^{(4)}] \leftarrow \partial_{\Pi_\ell}[\pi_3^{(4)}, \pi_4^{(4)}] - \Pi_\ell(\pi_1^{(3)}, \pi_2^{(3)}) E_4 \quad (\text{A.15})$$

The accumulations on the low rank factors Π_ℓ , $\ell = 1, \dots, R_1$ expand as accumulations on each of their low rank factorizations $\partial_{\Pi_\ell[i,j]} = \partial_{S_\ell[i]S_\ell[j]}$ from which we get each of the separate partials using the chain rule as $\partial_{S_\ell[i]} = \partial_{\Pi_\ell[i,j]} S_\ell[j]$ and equivalently $\partial_{S_\ell[j]} = \partial_{\Pi_\ell[i,j]} S_\ell[i]$.

Following the normalization (A.1), we can now express the original constraints together with their higher order extensions. Those sets of constraints read as follows. We first decompose the map A into the component $A_0 \in \mathbb{R}^{N \times m}$ acting on the first column and the remaining part $A_1 \in \mathbb{R}^{N \times n-1}$ acting on the matrix RR^T . Each of the constraint are encoded by means of appropriate matrices $A_{ij} = (A_0)_{ij} + (A_1)_{ij}$ as

$$\|\mathcal{A}(M) - b\|_2^2 = \sum_{k=1}^{|\Omega|} | \langle (A_0)_k, R_x \alpha^T \rangle + \langle (A_1)_k, R_x R_y^T b_k \rangle |^2 \quad (\text{A.16})$$

We will use corresponding linear maps \mathcal{A}_0 and \mathcal{A}_1 to encode the matrix constraints efficiently. We thus have $\mathcal{A}_0(x) = A_0 x$ and $\{\mathcal{A}_1(X)\}_k = \langle (A_1)_k, X \rangle$. Then for any vector $b \in \mathbb{R}^N$ $\mathcal{A}_0^* b = A_0^* b = \sum_{k=1}^N (a_0)_k b_k$ where $(a_0)_k$ denotes the transpose of the k^{th} row of A_0 and $\mathcal{A}_1^* b = \sum_{k=1}^N (A_1)_k b_k$. The gradient for those constraints reads

$$\begin{aligned} \partial_\alpha &= 2R_x^T \mathcal{A}_0^*(\mathcal{A}(M) - b) - 4\alpha(b^T(\mathcal{A}(M) - b)) \\ \partial_{R_x} &= 2\mathcal{A}_0^*(\mathcal{A}(M) - b)\alpha^T + 2\tilde{\mathcal{A}}^*(\mathcal{A}(M) - b)R_y \\ \partial_{R_y} &= R_x^* \tilde{\mathcal{A}}^*(\mathcal{A}(M) - b) \end{aligned}$$

We call higher order affine constraints the constraints derived from multiplying any of the constraints in $\{(\mathcal{A}(X) - b)_j\}_{j=1}^J = \{h_j(x, y)\}_{j=1}^J$ by any of the monomials of degree at most $t - \deg(h_j)$ for a relaxation of order t (a.k.a the t^{th} round of the hierarchy). Those higher order constraints can be encoded simply by multiplying the matrix of (pseudo)-moments by each of the constraints vectors of coefficients [37]. Since the moments matrix is low rank, for any vector $a \in \mathbb{R}^{\binom{N}{n}}$ this product reads very simply as $a^T M(m) = \langle a, [\alpha^T, R^T, \Pi^T] \rangle [\alpha^T, R^T, \Pi^T]^T = 0$. For a general set of affine constraints, the decomposition of X introduced in (A.1) and the decomposition of \mathcal{A} used in (A.16), stable minimization of the higher order

affine constraints reads

$$\begin{aligned}
\min & \left\| \sum_{\ell=1}^{R_1} (\mathcal{A}_0 R_\ell^x + \mathcal{A}(\Pi_\ell^{xy}) - b\alpha_\ell)(\alpha^T, R^T, \Pi^T)_\ell \right\|_F^2 \\
& = \left\langle \left(\mathcal{A}_0(R_k^x) + \mathcal{A} \left(\sum_{r=1}^{R_2} S_{k,\ell}^x (S_{k,\ell}^y)^T \right) - b\alpha_k \right)_{\ell=1}^{R_1}, \left(\mathcal{A}_0(R_k^x) + \mathcal{A} \left(\sum_{r=1}^{R_2} S_{k,\ell}^x (S_{k,\ell}^y)^T \right) - b\alpha_k \right)_{\ell=1}^{R_1} T^T T \right\rangle \\
& = \langle \Delta^T \Delta, T^T T \rangle \\
& = \sum_{\ell=1}^{R_1} \sum_{\ell'=1}^{R_1} (\Delta^T \Delta)[\ell, \ell'] \left(\left\langle \sum_{k=1}^{R_2} S_{\ell,k} S_{\ell,k}, \sum_{k'=1}^{R_2} S_{\ell',k'} S_{\ell',k'} \right\rangle \right)
\end{aligned}$$

Here we let $S_{k,\ell}^x, S_{k,\ell}^y$ denote the first and second blocks of the low rank factors $S_{k\ell} \in \mathbb{R}^{(m+n-1) \times R_2}$ of each matrix Π_ℓ . $\Delta^T \Delta \in \mathbb{R}^{R_1 \times R_1}$. We thus have $S_{\ell,k}^x \in \mathbb{R}^{m \times R_1}$, $S_{\ell,k}^y \in \mathbb{R}^{(n-1) \times R_1}$. The definition of Π^{xy} follows from those ideas,

$$\Pi_\ell^{xy} = \sum_{k=1}^{R_2} S_{\ell,k}^x (S_{\ell,k}^y)^T.$$

Finally, R_x derives from the decomposition of $R^{(m+n-1) \times R_1}$ into $R_x \in \mathbb{R}^{m \times R_1}$ and $R_y^{(n-1) \times R_1}$ with $R = [R_x^T, R_y^T]^T$. The contributions to the gradient are given by deriving each side and noting that

$$\langle T \Delta^T \Delta, T \rangle = \langle \alpha \Delta^T \Delta, \alpha \rangle + \langle R \Delta^T \Delta, R \rangle + \sum_{\ell=1}^{R_1} (\Pi \Delta^T \Delta)_\ell \Pi_\ell \quad (\text{A.17})$$

$$= \sum_{\ell=1}^{R_1} \left\langle \text{Mat}(\Pi \Delta^T \Delta)_\ell, \left(\sum_{k=1}^{R_2} S_{\ell,k} S_{\ell,k}^T \right) \right\rangle \quad (\text{A.18})$$

From which we have $\partial_{S_{\ell,k}^x} = 2 \text{Mat}((\Pi \Delta^T \Delta)_\ell) S_{\ell,k} + 2((S_{\ell,k})^T \text{Mat}((\Pi \Delta^T \Delta)_\ell))^T$, $\partial_R = 2R \Delta^T \Delta$ and $\partial_\alpha = 2\alpha \Delta^T \Delta$. For the $\Delta^T \Delta$ term, a similar approach yields $\partial_{R_x} = 2\mathcal{A}_0^* \Delta T^T T$, $\partial_\alpha = -2b^T \Delta T^T T$. For the partials with respect to $S_{\ell,k}$ we use

$$\langle \Pi^{xy}, \mathcal{A}^* \Delta T^T T \rangle = \sum_{\ell=1}^{R_1} \left\langle \sum_{k=1}^{R_2} S_{\ell,k}^x (S_{\ell,k}^y)^T, (\mathcal{A}^* \Delta T^T T)_\ell \right\rangle \quad (\text{A.19})$$

From which we can derive $\partial S_{\ell,k}^x = 2 \text{Mat}(\mathcal{A}^* \Delta T^T T)_\ell S_{\ell,k}^y$ as well as $\partial S_{\ell,k}^y = 2(S_{\ell,k}^x)^T \text{Mat}(\mathcal{A}^* \Delta T^T T)_\ell$. In the expressions above, we use $\text{Mat}(\mathbf{x})$ to denote the usual vector to matrix operator that turns the vector $\mathbf{x} = [x_1^T, \dots, x_n^T]^T$ into the matrix $M = [x_1, \dots, x_n]$. Note that the product $\Pi \Delta^T \Delta$ can be computed efficiently as

$$(\Pi \Delta^T \Delta)_{\bullet, \ell} = \sum_{\ell'=1}^{R_1} \left(\sum_{k=1}^{R_2} S_{\ell',k} S_{\ell',k}^T \right) (\Delta^T \Delta)_{\ell', \ell} \quad (\text{A.20})$$

From which, the expression of the partials follow as,

$$2 \text{Mat}((\Pi \Delta^T \Delta)_\ell) S_{\ell,k} = 2 \sum_{\ell'=1}^{R_1} \left(\sum_{k'=1}^{R_2} S_{\ell',k'} \langle S_{\ell',k'}^T, S_{\ell,k} \rangle \right) (\Delta^T \Delta)_{\ell', \ell} \quad (\text{A.21})$$

This expression just computes a projection of the low rank factors onto the subspace generated by each low rank decomposition and can be efficiently carried out by stacking all those low rank factors in a matrix of size $\mathcal{O}(m+n-1 \times R_2)$ premultiplying the matrix by the ℓ column of $\Delta^T \Delta$ and then applying the projector and summing over the ℓ' indices.

References

- [1] A. A. Ahmadi, G. Hall, A. Papachristodoulou, J. Saunderson, and Y. Zheng. Improving efficiency and scalability of sum of squares optimization: Recent advances and limitations. *arXiv preprint arXiv:1710.01358*, 2017.
- [2] A. S. Bandeira. Convex relaxations for certain inverse problems on graphs. 2015.
- [3] B. Barak, F. G. Brandao, A. W. Harrow, J. Kelner, D. Steurer, and Y. Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 307–326. ACM, 2012.
- [4] B. Barak and A. Moitra. Tensor prediction, Rademacher complexity and random 3-xor. *arXiv preprint arXiv:1501.06521*, 2015.
- [5] B. Barak and D. Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. *arXiv preprint arXiv:1404.5236*, 2014.
- [6] M. Bardet. On the complexity of a gröbner basis algorithm. In *Algorithms Seminar, 2002–2004*, page 85, 2005.
- [7] R. Berke and M. Onsjö. Propagation connectivity of random hypergraphs. In *Stochastic Algorithms: Foundations and Applications*, pages 117–126. Springer, 2009.
- [8] S. Burer and R. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [9] S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [10] S. Burer and R. D. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, 2005.
- [11] M. Byröd, K. Josephson, and K. Åström. Fast and stable polynomial equation solving and its application to computer vision. *International Journal of Computer Vision*, 84(3):237–256, 2009.
- [12] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM review*, 57(2):225–251, 2015.
- [13] E. J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.
- [14] E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [15] E. J. Candès and Y. Plan. Matrix completion with noise. *Proc. IEEE*, 98(6):925–936, 2010.
- [16] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [17] E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [18] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, 2010.

- [19] A. Cosse and L. Demanet. Rank-one matrix completion is solved by the sum-of-squares relaxation of order two. In *Proceedings of the 6th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP'15)*. IEEE, 2015.
- [20] D. Cox, J. Little, and D. O'shea. *Ideals, varieties, and algorithms*, volume 3. Springer, 1992.
- [21] C.-F. Cui, Y.-H. Dai, and J. Nie. All real eigenvalues of symmetric tensors. *SIAM Journal on Matrix Analysis and Applications*, 35(4):1582–1601, 2014.
- [22] L. Demanet and V. Jugnon. Convex recovery from interferometric measurements. *arXiv preprint arXiv:1307.6864*, 2013.
- [23] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, March 2002.
- [24] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [25] J. Gouveia, P. A. Parrilo, and R. R. Thomas. Theta bodies for polynomial ideals. *SIAM Journal on Optimization*, 20(4):2097–2118, 2010.
- [26] D. Heldt, M. Kreuzer, S. Pokutta, and H. Poulisse. Approximate computation of zero-dimensional polynomial ideals. *Journal of Symbolic Computation*, 44(11):1566–1591, 2009.
- [27] V. Jugnon, L. Demanet, et al. Interferometric inversion: a robust approach to linear inverse problems. In *Proceedings of SEG Annual Meeting, Houston*, pages 5180–5184, 2013.
- [28] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, pages 952–960, 2009.
- [29] R. H. Keshavan, A. Montanari, and S. Oh. Learning low rank matrices from $o(n)$ entries. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 1365–1372. IEEE, 2008.
- [30] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.
- [31] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 767–775. ACM, 2002.
- [32] F. Kiraly and L. Theran. Error-minimizing estimates and universal entry-wise error bounds for low-rank matrix completion. In *Advances in Neural Information Processing Systems*, pages 2364–2372, 2013.
- [33] F. Király and R. Tomioka. A combinatorial algebraic approach for the identifiability of low-rank matrix completion. *arXiv preprint arXiv:1206.6470*, 2012.
- [34] F. J. Király, L. Theran, and R. Tomioka. The algebraic combinatorial approach for low-rank matrix completion. *Journal of Machine Learning Research*, 16:1391–1436, 2015.
- [35] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [36] J. B. Lasserre. Convergent sdp-relaxations in polynomial optimization with sparsity. *SIAM Journal on Optimization*, 17(3):822–843, 2006.
- [37] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009.

- [38] Y. Nesterov. Squared functional systems and optimization problems. In *High performance optimization*, pages 405–440. Springer, 2000.
- [39] J. Nie. An exact jacobian sdp relaxation for polynomial optimization. *Mathematical Programming*, 137(1-2):225–255, 2013.
- [40] J. Nie and J. Demmel. Sparse sos relaxations for minimizing functions that are summations of small polynomials. *SIAM Journal on Optimization*, 19(4):1534–1558, 2008.
- [41] P. A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, Citeseer, 2000.
- [42] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.
- [43] D. L. Pimentel-Alarcón, N. Boston, and R. D. Nowak. A characterization of deterministic sampling patterns for low-rank matrix completion. *arXiv preprint arXiv:1503.02596*, 2015.
- [44] D. L. Pimentel-Alarcón and R. D. Nowak. Adaptive strategy for restricted-sampling noisy low-rank matrix completion.
- [45] B. Recht. A simpler approach to matrix completion. *The Journal of Machine Learning Research*, 12:3413–3430, 2011.
- [46] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [47] N. Shor. Class of global minimum bounds of polynomial functions. *Cybernetics and Systems Analysis*, 23(6):731–734, 1987.
- [48] N. Z. Shor. Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25(6):1–11, 1987.
- [49] N. Z. Shor. An approach to obtaining global extremums in polynomial mathematical programming problems. *Cybernetics*, 23(5):695–700, 1988.
- [50] A. Singer and M. Cucuringu. Uniqueness of low-rank matrix completion by rigidity theory. *SIAM Journal on Matrix Analysis and Applications*, 31(4):1621–1641, 2010.
- [51] B. Sturmfels. *Solving systems of polynomial equations*. Number 97. American Mathematical Soc., 2002.
- [52] G. Tang and P. Shah. Guaranteed tensor decomposition: A moment approach. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 1491–1500, 2015.