Structured nuclear norm matrix completion: Guaranteeing exact recovery via block-column scaling

Konstantin Usevich, Jonathan Gillard, Philippe Dreesen[‡]and Ivan Markovsky[§]

June 19, 2025

Abstract

The goal of low-rank matrix completion is to minimize the rank of a matrix while adhering to the constraint that known (non-missing) elements are fixed in the approximation. Minimizing rank is a difficult, non-convex, NP-hard problem, often addressed by substituting rank with the nuclear norm to achieve a convex relaxation. We focus on structured matrices for completion, where in addition to the constraints described earlier, matrices also adhere to a predefined structure. We propose a technique that ensures the exact recovery of missing entries by minimizing the nuclear norm of a matrix where the non-missing entries are first subject to block-column scaling. We provide the proofs for exact recovery and propose a way for choosing the scaling parameter to ensure exact recovery. The method is demonstrated in several numerical examples, showing the usefulness of the proposed technique.

1 Introduction

Low-rank matrix completion aims to reconstruct matrices with missing entries by exploiting an inherent low-rank structure. In the general case, low-rank matrix completion is an optimization problem with the objective of minimizing the rank of a matrix subject to constraints. Rank minimization problems are generally NP-hard and non-convex (see e.g., [Laz16]), so several authors propose to instead solve a convex relaxation based on the nuclear norm, which has well-grounded theory for the case of unstructured matrices [XJ12, CT10, CR12]. While there are other approaches to relax the nonconvexity, see e.g., [CGO22], we focus on the nuclear norm approach since it is pervasive. A natural question emerges — when does the solution of the convex relaxation coincide with the rank minimization problem? It is known that the nuclear norm performs well when the missing values are sampled in a random fashion, where perfect (or exact) recovery is ensured with high probability, see the seminal work [CT10, CR12] and numerous follow-ups.

In this paper, we consider the case of structured matrix completion. Informally speaking, structured matrices have dependencies between their elements and we wish to preserve the structure whilst making the completion. Examples of well-known structured matrices include Hankel, block-Hankel, Toeplitz, Sylvester and circulant; each of these structures is associated with particular problem domains, and we offer some examples shortly. The completion of structured matrices can be viewed as a special case of structured low-rank approximation [Mar08, Mar19]. In the applications we consider, and for structured matrices generally, missing data often appear in blocks as opposed to being randomly located throughout the matrix.

Specifically, we consider structured matrix completion based on minimizing the nuclear norm in the deterministic scenario (structure and missing-data-not-at-random). In this case, the use of the nuclear

^{*}CRAN UMR 7039, CNRS, Université de Lorraine, Nancy, France, corresponding author (konstantin.usevich@cnrs.fr)

[†]School of Mathematics, Cardiff University, UK

[‡]Department of Advanced Computing Sciences (DACS), Maastricht University, Maastricht, Netherlands

[§]Catalan Institution for Research and Advanced Studies and Centre Internacional de Mètodes Numèrics en Enginyeria

norm breaks down, as shown by several studies [Mar12]. For Hankel matrices, [DP15, UC16] suggest that an exponential rescaling (dampening) of the problem is possible, but it is difficult to provide good bounds for such a rescaling. In this paper, we propose a simple block-column scaling of the observed entries, which guarantees the exact recovery of missing entries. Our contribution in this paper is applicable to any affinely-structured matrix and provides theoretical guarantees for exact recovery as well as a method to choose the scaling parameter.

There are other methods available to solve some of the example problems described in this paper. Note however that our work applies to general affine structures (not just Hankel-type); for general structures parametric model equivalents may not be available. We propose an alternative to estimating parameters of a parametric model which is not only demonstrably useful for practical application, but also interesting mathematically, and complements the results in seminal work such as [CP10, CT10, CR12]. Hence our contributions are beyond just a method to solve a problem.

This paper has the following structure. In Section 2 we introduce the matrix completion problem and provide some examples of matrix structures and typical completions needed. In Section 3 we describe the means to scale the matrix structure, so that the nuclear norm relaxation achieves the correct recovery for the scaled matrix structure. Our result on scaling columns to achieve exact recovery of missing values is also given in Section 3. Before giving the proof of the main result in Section 5, we provide a roadmap in Section 4. Numerical examples are given in Section 6 before the paper is concluded in Section 7.

2 Affine structured matrix completion

2.1 Structured matrices

An affine matrix structure $S(\cdot)$ is an affine mapping $S : \mathbb{F}^N \to \mathbb{F}^{L \times K}$ where \mathbb{F} is \mathbb{R} or \mathbb{C} , thus parameterizing (an affine) set of structured matrices. An example of such matrices are given below and a formal construction is offered in Section 4.1.

Example 2.1 (running example, Hankel matrices). *The Hankel matrix structure* $\mathcal{H}_L : \mathbb{F}^N \to \mathbb{F}^{L \times K}$, K = N - L + 1 that maps a vector $\mathbf{p} \in \mathbb{F}^N$ to an $L \times K$ matrix with

$$(\mathcal{H}_L(\boldsymbol{p}))_{i,j} = \boldsymbol{p}_{i+j-1},$$

i.e., the values are constant on the antidiagonals and parameterized by the elements of p (see also (1)).

Note that the notation $\mathcal{H}_L(\mathbf{p})$ is used for a vector \mathbf{p} of any size $N \ge L$, and the number of columns of the matrix depends on the size of the input vector.

2.2 Matrix completion and rank minimization

The goal of low-rank matrix completion is to fill in the missing elements of the matrix $S(\mathbf{p})$ based on the low-rank assumption. From now on, we write N = n + m, and we assume that the first n values of \mathbf{p} (denoted $\mathbf{p}_{1:n}$) are known and the goal is to recover the last m missing values¹, $\mathbf{p}_{n+1:n+m}$ (so that N = n + m).

Formally, the exact structured low-rank matrix completion is formulated as the rank minimization

$$\min_{\boldsymbol{p} \in \mathbb{F}^{m+n}} \operatorname{rank} \{ \mathcal{S}(\boldsymbol{p}) \} \quad \text{subject to} \quad \boldsymbol{p}_{1:n} = \boldsymbol{p}_0, \tag{RMIN}$$

where $p_0 = (p_{0,1}, p_{0,2}, \dots, p_{0,n})$ is a given vector of known values. We denote an optimal solution to (RMIN) as $\hat{p}^{(\text{RMIN})}$.

Example 2.2 (Example 2.1, continued). Our running example stems from time series analysis, where we would like to fill in the last values of a scalar Hankel matrix. In (1), the grey-shaded elements are

¹Missing values in the middle of p can be also treated, see discussion in Section 3.3

known (fixed) and the remaining values are missing. The Hankel matrix in (1) has size $L \times K$ so that K + L - 1 = n + m.

$$\mathcal{H}_{L}(\boldsymbol{p}) \stackrel{\mathsf{def}}{=} \begin{bmatrix} p_{1} & p_{2} & \cdots & \cdots & p_{K} \\ p_{2} & p_{3} & \cdots & \ddots & \ddots & \vdots \\ \vdots & \cdots & \ddots & \ddots & \ddots & p_{n} \\ \vdots & \ddots & \ddots & \ddots & \ddots & p_{n+1} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ p_{L} & \cdots & p_{n} & p_{n+1} & \cdots & p_{n+m} \end{bmatrix}.$$
(1)

The Hankel low-rank completion problem is shown to be useful in the context of forecasting [GU18, BG17, GU23], where p corresponds to a scalar time series. The first n values of the time series p are known (observed) and we would like to forecast m steps forward.

This approach is motivated by the fact that time series which can be written as a sum of products of polynomial, exponential, and sinusoidal functions have low-rank Hankel matrices (see e.g., [GU23]). Such time series constitute an exceptionally rich class able to model complex trends and multiple modulated periodic components (see also [GNZ01]). Here, forming the forecasting problem as Hankel structured low-rank matrix completion is very attractive because it avoids parameter estimation necessary for classic model-based statistical forecasting methods [Mar17]. A more general class of problems with multivariate time series is described in Section 6.4. In practice, an approximation version of (RMIN) is often considered, when we allow for the error on the known values.

2.3 Convex relaxation and nuclear norm

In this manuscript, we investigate the performance of the nuclear norm relaxation, which for the rank minimization problem (RMIN) takes the form:

$$\hat{\boldsymbol{p}} = \operatorname*{argmin}_{\boldsymbol{p} \in \mathbb{F}^{(n+m)}} \| \mathcal{S}(\boldsymbol{p}) \|_{*}$$
 subject to $\boldsymbol{p}_{1:n} = \boldsymbol{p}_{0},$ (NNMIN)

where $||X||_*$ is the nuclear norm (sum of all singular values of X). The intuition behind this relaxation is the same as for using the ℓ_1 -norm in compressed sensing: the nuclear norm is expected to force all but a few singular values to be zero (a low-rank solution). In what follows, we will denote an optimal solution to (NNMIN) as $\hat{p}^{(NNMIN)}$.

In our applications, the pattern of missing data is fixed, as opposed to the random sampling in [CT10, CR12], and the conditions for perfect recovery are much more restrictive. In particular, it was shown in [DP15] and [UC16], that in Hankel matrix completion, the time series need to be sufficiently damped in order to allow for a perfect recovery. In this paper, we show that for a wide range of structures a simple block-wise scaling guarantees perfect recovery for the scaled matrix structure.

2.4 Alternatives to the nuclear norm

Nuclear norm minimization has been celebrated for its convex formulation, which offers strong theoretical recovery guarantees and robustness in the presence of noise. Nonetheless, several alternative methods for matrix completion exist that avoid its convex relaxation while offering complementary advantages. For instance, low-rank factorization-based approaches model the given matrix as a product UV and apply alternating minimization to optimize for U and V. Although these methods tend to be more computationally efficient, they often require careful initialization and might converge to local minima rather than the globally optimal solution [JONS13]. In addition, Riemannian optimization techniques directly optimize on the manifold of fixed-rank matrices to benefit from the intrinsic geometric structure of the solution space, typically resulting in rapid convergence; however, their performance can be sensitive to noise and depends on the quality of the starting point [Van13]. Iterative hard thresholding methods enforce a strict rank constraint via singular value truncation at each iteration, which can be highly efficient in practice, yet they often lack the universal convergence guarantees that come with convex nuclear norm methods [TW16]. Overall, while these alternatives provide efficient and scalable solutions in various settings, nuclear norm minimization remains a cornerstone approach due to its balanced trade-off of robust global optimality, provable error bounds, and practical effectiveness in a wide range of applications, as described in this paper.

3 Block-column scaling for nuclear norm completion

3.1 Block-column scaling

In this section, we describe the proposed scaling so nuclear norm minimization for fixed structure obtains exact recovery of missing values. Our proposal is to replace S(p) with the scaled matrix structure

$$S_{\varepsilon}(\boldsymbol{p}) := S(\boldsymbol{p}) \begin{bmatrix} \boldsymbol{I}_{K-\kappa} & \\ & \varepsilon \boldsymbol{I}_{\kappa} \end{bmatrix}, \qquad (2)$$

where $\varepsilon \ge 0$ is a small parameter, *i.e.* we scale the last κ columns of $S(\mathbf{p})$, with $0 < \kappa < K$, by a small number ε . The following remark is of prime importance:

Remark 3.1. For $\varepsilon > 0$, we have rank $\{S_{\varepsilon}(\mathbf{p})\} = \operatorname{rank}\{S(\mathbf{p})\}\$ therefore all the solutions of the exact matrix completion (RMIN) using either the original matrix $S(\mathbf{p})$ or its scaled version $S_{\varepsilon}(\mathbf{p})$ coincide.

Our proposal is to replace S(p) with $S_{\varepsilon}(p)$, in the nuclear norm minimization problem (NNMIN), and we seek to solve the modified problem

$$\min_{\boldsymbol{p} \in \mathbb{R}^{(n+m)}} \| \mathcal{S}_{\varepsilon}(\boldsymbol{p}) \|_{*} \quad \text{subject to} \quad \boldsymbol{p}_{(1:n)} = \boldsymbol{p}_{0}. \tag{NNMIN-}\varepsilon)$$

Under several natural assumptions, replacing S(p) with the scaled $S_{\varepsilon}(p)$ leads to guaranteed exact recovery of the missing values $p_{n+1:n+m}$, for suitably chosen ε . The goal is to choose such a scaling so that the solution $\hat{p}^{(\text{NNMIN}-\varepsilon)}$ of (NNMIN- ε) leads to exact recovery of $p_{n+1:n+m}$ and is therefore better than $\hat{p}^{(\text{NNMIN})}$.

Related work Scaling of rows and columns has already been proposed for the case of unstructured matrix completion. Such a scaling, equivalent to a weighted nuclear norm minimization [SS10] was shown to be effective and provides better recovery guarantees than the ordinary nuclear norm minimization [CBSW14]. There has been an increasing activity in analyzing non-uniform deterministic patterns of missing values [FNP⁺21]. However, the authors are unaware of any such works that address the case of structured matrix completion, except for the recent work [TL21] that mainly reveals connections between weighting and structured completion (rather than using scalings to improve the performance of the nuclear norm in the structured case). Our work therefore offers a novel contribution for ensuring exact recovery of missing entries of structured matrices.

3.2 Assumptions and the main result

We now describe the necessary assumptions before discussing them in more detail in the subsequent section.

Assumption 3.1. The solution $\hat{p}^{(RMIN)}$ of the rank minimization problem (RMIN) is unique.

Assumption 3.2. For any optimal solution $\widehat{p}^{(RMIN)}$ of (RMIN), we have

 $\operatorname{rank}\{\mathcal{S}(\widehat{\boldsymbol{p}}^{(RMIN)})\} = \operatorname{rank}\{\mathcal{S}_0(\widehat{\boldsymbol{p}}^{(RMIN)})\},\$

where $S_0(\mathbf{p})$ is the scaled structure (2) with $\varepsilon = 0$.

Assumption 3.3. The matrix structure $S_0(\mathbf{p})$ depends only on $\mathbf{p}_{1:n}$ (i.e., the first *n* elements of the vector \mathbf{p}); in such a case we write $S_0(\mathbf{p}) = S_0(\mathbf{p}_{1:n})$.

Before formulating the main results, we make some remarks about these technical assumptions.

Remark 3.2. Assumptions 3.2 to 3.3 concern the matrix $S_0(p)$

$$\mathcal{S}_0(\boldsymbol{p}) := \mathcal{S}(\boldsymbol{p}) \begin{bmatrix} \boldsymbol{I}_{K-\kappa} & \\ & \boldsymbol{0} \end{bmatrix}.$$
(3)

This is a limiting ($\varepsilon = 0$) case of $S_{\varepsilon}(\mathbf{p})$, which is special for the following reasons:

- In this case, Remark 3.1 (on the equivalence of rank minimization problems for the scaled matrix $S_{\varepsilon}(\mathbf{p})$ and $S(\mathbf{p})$) does not necessarily hold. Thus Assumption 3.2 ensures that an analogue of Remark 3.1 holds true for $\varepsilon = 0$;
- As seen in (3), S₀(**p**) contains the first K − κ columns of S(**p**) (which are left unscaled in S_ε(**p**) for any ε); thus Assumption 3.3 is equivalent to assuming that the parameters **p**_{n+1:n+m} do not appear in the first K − κ columns (i.e., all the missing elements **p**_{n+1:n+m} are scaled by ε);
- From a practical point of view, $S_0(\mathbf{p})$ is not important as it does not depend on the missing elements (so it does not provide a way to complete them); however, it is important for the theoretical analysis of the scaled problem (NNMIN- ε) which appears later.

We are also going to provide examples for the assumptions in the following subsections. Under Assumptions 3.1 to 3.3, the following result holds.

Theorem 3.3. Let $S_{\varepsilon}(\mathbf{p})$ be the scaled matrix structure (2) satisfying Assumptions 3.1 to 3.3, and $\hat{\mathbf{p}}^{(RMIN)}$ be the unique solution of (RMIN). Then there exists ε_0 such that for any $\varepsilon \in (0, \varepsilon_0)$, the minimizer $\hat{\mathbf{p}}^{(NNMIN-\varepsilon)}$ of the scaled problem (NNMIN- ε) is unique and is equal to $\hat{\mathbf{p}}^{(RMIN)}$.

For ease of reading, the proof of the theorem will be split into several stages and into dedicated subsections. In the rest of the section, we provide a discussion on the plausibility of the assumptions, show the key ideas behind the proof, and also indicate how we can find an estimate for ε_0 .

3.3 Discussion on the assumptions: the case of Hankel matrices

In this subsection, we give examples for the assumptions using our running example, that of completing Hankel matrices, which are linked to some well-known conditions in the literature.

Lemma 3.4. For completing the Hankel matrix $\mathcal{H}_L(\mathbf{p})$ given in (1), Assumption 3.3 is equivalent to $\kappa \geq m$.

Proof. The last elements $p_{n+1:n+m}$ appear only in the *m* last columns of $\mathcal{H}_L(p)$ and so the first $K - \kappa$ columns do not depend on $p_{n+1:n+m}$ if and only if $\kappa \ge m$.

To illustrate Lemma 3.4, consider a special case of a Hankel matrix (1) with n = 7, m = 2 and $\kappa = 3$ (i.e., the right-hand side block is scaled):

$$\mathcal{H}_3(oldsymbol{p}) = \left[egin{array}{cccccccccc} p_1 & p_2 & p_3 & p_4 & p_5 & p_6 & p_7 \ p_2 & p_3 & p_4 & p_5 & p_6 & p_7 & p_8 \ p_3 & p_4 & p_5 & p_6 & p_7 & p_8 & p_9 \end{array}
ight],$$

and it is seen that Assumption 3.3 is satisfied. We also note that a necessary condition for Assumption 3.1 is $r = \operatorname{rank}\{\mathcal{H}_L(p_0)\} < L$ (otherwise any completion of $\mathcal{H}_L(p_0)$ has rank L, and is thus nonunique). Under a stronger well-known condition, we can guarantee that both Assumption 3.1 and Assumption 3.2 are also satisfied.

Lemma 3.5. Let Assumption 3.3 be satisfied. If, in addition,

$$\boldsymbol{e}_L = (0, \ldots, 0, 1) \not\in \operatorname{span} \{ \mathcal{H}_L(\boldsymbol{p}_0) \},\$$

then both Assumption 3.1 and Assumption 3.2 are also satisfied.

Lemma 3.5 is well-known in Hankel low-rank approximation literature (see [HR84, Def. 5.9, p. 99],[GNZ01, §5.3], but also [GU23, Mar08]), and we summarize a proof in Appendix 7.1. In particular, Lemma 3.5 holds true for time series that are sums of complex exponentials and polynomials, and this is the example that follows.

Example 3.6 (see e.g., [GU23, Thm 5.1]). Let p_k , k = 1, ..., n be given as

$$p_k = \sum_{j=1}^{s} P_j(k) \lambda_j^k, \ k = 1, 2, \dots$$
 (4)

where $P_j(k)$ are complex polynomials of degrees at most $\nu_j - 1$ and $\lambda_j \in \mathbb{C}$. Then if $r = \nu_1 + \ldots + \nu_r \leq \min(\kappa - 1, L - 1, \frac{n+1}{2})$, then the rank minimization of $\mathcal{H}_L(\mathbf{p})$ is unique (is of rank r), and the unique completion is given by the same formula ($\hat{p}_k = p_k$, $k = n + 1, \ldots, n + m$ in (4)). Therefore, both Assumption 3.1 and Assumption 3.2 are satisfied.

We offer further commentary of the applicability of Assumptions 3.1 to 3.3 for some other problem settings.

Remark 3.7. The problem treated in Lemma 3.5 and Example 3.6 can be viewed as a special case of more general behavioral systems theory. In particular, Assumptions 3.1 to 3.3 are similar to Assumptions A1 - A3 in [MD22] which considers the same problem (for the special case of a so-called mosaic-Hankel structure) from a systems theoretic perspective. The main difference is the method of solution, which is not based on the nuclear norm relaxation.

Additionally, Assumptions 3.1 to 3.3 are intuitive for the problem of data-driven simulation, which is presented in depth in Section 6.4. For example, Assumption 3.2 corresponds to the condition of persistency of excitation [WRMD05]. We also note that a similar scaling was heuristically proposed for special case of data-driven simulation in [DM19], where it was proposed to scale the matrix as

$$egin{bmatrix} \gamma oldsymbol{I}_{K-\kappa} & \ & oldsymbol{I}_\kappa \end{bmatrix} \mathcal{S}(oldsymbol{p}),$$

where γ is a large number. Such a scaling is, in fact, equivalent to (2) if we take $\gamma = \varepsilon^{-1}$.

Finally we discuss on the applicability of Assumptions 3.1 to 3.3 when the missing values occur elsewhere in the parameter vector p.

Remark 3.8. Our results also apply when:

- missing data are located in the middle of the vector **p**;
- the columns in the middle of the matrix are scaled by ε .

Indeed, this case can be treated by rearranging the columns of the matrix and the elements of the parameter vector (this operation does not affect neither rank, nor the nuclear norm).

Consider the following example

$$\mathcal{H}_3(oldsymbol{p}) = egin{bmatrix} p_1 & p_2 & p_3 & p_4 & p_5 & p_6 & p_7 \ p_2 & p_3 & p_4 & p_5 & p_6 & p_7 & p_8 \ p_3 & p_4 & p_5 & p_6 & p_7 & p_8 & p_9 \end{bmatrix}$$

with missing elements p_4 and p_5 . We can rearrange the columns to get an equivalent matrix

$$\mathcal{S}(oldsymbol{p}) = egin{bmatrix} p_1 & p_6 & p_7 & p_2 & p_3 & p_4 & p_5 \ p_2 & p_7 & p_8 & p_3 & p_4 & p_5 & p_6 \ p_3 & p_8 & p_9 & p_4 & p_5 & p_6 & p_7 \end{bmatrix},$$

and by permuting the arguments of the vector, we obtain an equivalent matrix structure that satisfies Assumption 3.3.

4 Roadmap for the proof

4.1 Nuclear norm minimization: optimality conditions

In order to give the main ideas behind the proof of Theorem 3.3, we recall the optimality conditions for (NNMIN) from [UC16, GU18], and we cover both real- and complex-valued matrices.

Let \hat{p} be a parameter vector for which we wish to test optimality of (NNMIN) and $r = \operatorname{rank}\{S(\hat{p})\}$. Let us introduce the following notation for the compact SVD:

$$\mathcal{S}(\widehat{p}) = U\Sigma V^{\mathsf{H}}, \ U \in \mathbb{C}^{L imes r}, \ \Sigma \in \mathbb{R}^{r imes r}, \ V \in \mathbb{C}^{K imes r}.$$

The key quantities needed for formulating the optimality conditions are:

- the polar factor $B \stackrel{\text{def}}{=} UV^{\text{H}}$ (semi-unitary matrix in polar decomposition)
- and projectors on the left and right nullspace of $\mathcal{S}(\hat{p})$:

$$Q_1 = I_L - UU^{\mathsf{H}}, \quad Q_2 = I_K - VV^{\mathsf{H}}.$$

Note that the polar factor B and the projectors Q_1 and Q_2 depend on \hat{p} (which is fixed in the proofs), therefore we omit this dependence in the notation for the ease of reading.

Next, we recall notation for basis matrices, which are needed to consider structured matrix completion problems. An affine matrix structure [Mar19] can be parameterized as:

$$\mathcal{S}(\boldsymbol{p}) = \boldsymbol{S}_0 + \sum_{k=1}^{m+n} p_k \boldsymbol{S}_k.$$
(5)

For example, for the Hankel structure (1), the basis matrices are

$$\mathbf{S}_{1} = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdot & 0 & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \mathbf{S}_{2} = \begin{pmatrix} 0 & 1 & \cdots & 0 & 0 \\ 1 & 0 & \cdot & 0 & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdot & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \dots, \mathbf{S}_{m+n} = \begin{pmatrix} 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdot & 0 & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdot & 0 & 0 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix},$$

with S_0 being the zero matrix.

Using the notation introduced above, the necessary (and sufficient) optimality conditions for an optimal solution of (NNMIN) can be formulated as follows.

Lemma 4.1 (First-order optimality conditions, [UC16, Proposition 14], [GU18]). The point \hat{p} is a minimizer of (NNMIN) if and only if there exists a matrix $M \in \mathbb{C}^{L \times K}$ with spectral norm $||M||_2 \leq 1$, satisfying

$$\langle \boldsymbol{Q}_1 \boldsymbol{M} \boldsymbol{Q}_2 + \boldsymbol{B}, \boldsymbol{S}_k \rangle_F = 0, \quad \text{for all } k \in \{n+1, \dots, n+m\}.$$
 (6)

If, in addition, the norm of matrix M satisfies $||M||_2 < 1$, and the set of matrices $\{Q_1S_kQ_2\}_{k=n+1}^{n+m}$ is linearly independent, then \hat{p} is the unique minimizer of (NNMIN).

Lemma 4.1 is, in fact, based on the notion of subdifferential of the nuclear norm, which is parameterized by matrices M with $||M||_2 < 1$ (see Section 7.2). A particular matrix M satisfying the conditions of Lemma 4.1 (with $||M||_2 < 1$) is called a *dual certificate*². Our proof strategy relies on constructing a particular dual certificate for the given solution of (NNMIN), i.e., the matrix M satisfying linear constraints (6), for which we can guarantee $||M||_2 < 1$.

²It is a standard term in the nuclear norm minimization literature [CT10, CR12] as it refers to duality in convex optimization [BV04].

4.2 Candidate dual certificate and its norm

Instead of finding M with small spectral norm (as suggested by Lemma 4.1), we relax the problem and find M with small Frobenius norm (subject to constraints (6)), which has explicit solution. A matrix with minimal Frobenius norm will be called a *candidate dual certificate* (see also [CT10]), this is defined next.

Definition 4.2. A candidate dual certificate for \hat{p} in (NNMIN) is a minimizer of

$$M^* \stackrel{\mathsf{def}}{=} \operatorname{argmin} \|M\|_F$$
 subject to *constraints in* (6). (7)

Due to the standard inequality $||\mathbf{M}||_2 \leq ||\mathbf{M}||_F$, a candidate dual certificate becomes a dual certificate if $||\mathbf{M}_*||_F < 1$ (and thus will ensure optimality in (NNMIN)). Therefore, the key point of the proof would be to guarantee that such an inequality holds for a sufficiently scaled problem, as shown by the following proposition.

Proposition 4.3. Under Assumptions 3.1 to 3.3 let $\hat{p} = \hat{p}^{(RMIN)}$, and let $M^*(\varepsilon)$ denote the candidate dual certificate for the scaled problem (NNMIN- ε). Then, for small $\varepsilon > 0$, the candidate dual certificate exists, is unique, and its squared Frobenius norm has the expansion

$$\|\boldsymbol{M}^*(\varepsilon)\|_F^2 = m_0\varepsilon^2 + o(\varepsilon^2),$$

where m_0 is a constant depending only on $\mathcal{S}(\widehat{p}^{(RMIN)})$.

The precise form of the constant m_0 will be given in Section 5.3. Note that, in particular, Proposition 4.3 ensures that $\lim_{\varepsilon \to 0} || M^*(\varepsilon) ||_F^2 = 0$, which shows that for sufficiently small ε we have $|| M_*(\varepsilon) ||_F < 1$. We conclude this section with a remark on usefulness of Proposition 4.3; it will enable us to provide a practical estimate for the value of the scaling needed to obtain exact recovery.

Remark 4.4. For small ε , Proposition 4.3 suggests that we can replace $\|\mathbf{M}^*(\varepsilon)\|_F^2$ with its leading term approximation. Therefore, the value of ε , chosen such that $\|\mathbf{M}^*(\varepsilon)\|_F^2 \approx \varepsilon^2 m_0 = 1$ is a good estimate to enable exact recovery. It is given by

$$\tilde{\varepsilon} = \frac{1}{\sqrt{m_0}} \tag{8}$$

4.3 Sketch of the proof

To give a gist of the proof, we first note that the constraints (6) are linear in elements of M and thus can be equivalently vectorized as

$$\boldsymbol{A}\operatorname{vec}(\boldsymbol{M}) = \boldsymbol{b},\tag{9}$$

where $A \in \mathbb{C}^{m \times LK}$ and $b \in \mathbb{C}^m$ are the matrices obtained by vectorizing (6):

$$\boldsymbol{A} = \begin{bmatrix} \operatorname{vec}^{\mathsf{T}}(\boldsymbol{S}_{n+1})(\boldsymbol{Q}_{2}^{\mathsf{T}} \otimes \boldsymbol{Q}_{1}) \\ \vdots \\ \operatorname{vec}^{\mathsf{T}}(\boldsymbol{S}_{n+m})(\boldsymbol{Q}_{2}^{\mathsf{T}} \otimes \boldsymbol{Q}_{1}) \end{bmatrix},$$
(10)

$$\boldsymbol{b} = \begin{bmatrix} -\langle \boldsymbol{B}, \boldsymbol{S}_{n+1} \rangle \\ \vdots \\ -\langle \boldsymbol{B}, \boldsymbol{S}_{n+m} \rangle \end{bmatrix}, \qquad (11)$$

and $vec(\cdot)$ denotes the vectorization operation which stacks the columns of a matrix into a single column vector, \otimes is the Kronecker product, and the polar factor B, and the projectors Q_1, Q_2 are as in Lemma 4.1.

With such notation, the candidate dual certificate is more easily constructed. Indeed, in terms of vectorized constraints (9), the minimization (7) is equivalent to minimizing the 2-norm of the vector vec(M) subject to linear constraints A vec(M) = b. Therefore (see e.g., [BV04, Ch.6]), when A is full row-rank (i.e., rank $\{A\} = m$), the candidate dual certificate M^* is unique and has an explicit expression in terms of A and b in (9):

$$\operatorname{vec}(\boldsymbol{M}^*) = \boldsymbol{A}^{\mathsf{H}}(\boldsymbol{A}\boldsymbol{A}^{\mathsf{H}})^{-1}\boldsymbol{b};$$
(12)

in addition, the Frobenius norm of the optimal solution is given by

$$\|\boldsymbol{M}^*\|_F^2 = \|\operatorname{vec}(\boldsymbol{M}^*)\|_2^2 = \boldsymbol{b}^{\mathsf{H}}(\boldsymbol{A}\boldsymbol{A}^{\mathsf{H}})^{-1}\boldsymbol{b}.$$
 (13)

The proof of Theorem 3.3 can be summarized as follows:

- For a fixed p
 = p
 ^(RMIN), we test its optimality for (NNMIN-ε) using Lemma 4.1 and the vectorized form of constraints (9).
- 2. For the case of the scaled structure $S_{\varepsilon}(p)$, the matrices in (9) also depend on ε (i.e., they become $A(\varepsilon), b(\varepsilon)$).
- 3. Under Assumptions 3.1 to 3.3, we show $A(\varepsilon)$ is full row rank (hence (12) and (13) hold for the candidate dual certificate depending on $M^*(\varepsilon)$).
- 4. We then show that the norm $\|M^*(\varepsilon)\|_F < 1$ for small ε , and hence $\|M^*(\varepsilon)\|_2 \le \|M^*(\varepsilon)\|_F < 1$, which implies that $M^*(\varepsilon)$ is a valid dual certificate (in the sense of Lemma 4.1) for (NNMIN- ε).

5 Proof of the main result

This section is organized as follows. We first discuss in Section 5.1 the implications of the assumptions of Theorem 3.3 and derive the form of the perturbed projectors and basis matrices. Based on lemmas in Section 5.1, we provide in Section 5.2 expressions for $A(\varepsilon)$ and $b(\varepsilon)$ from (10) and (11). Finally, we Section 5.3 we give the proof of Theorem 3.3 and Proposition 4.3.

5.1 Scaled matrix structure: basis matrices and orthogonal projectors

In order to get analogous expressions for (12) and (13) for the scaled matrix structure $S_{\varepsilon}(p)$, we adapt (10)–(11) to the scaled problem (NNMIN- ε). Specifically, in this subsection, we provide details on the matrices B, Q_1, Q_2, S_k for the scaled matrix structure $S_{\varepsilon}(p)$. Note that the scaled matrix structure $S_{\varepsilon}(p)$ is expressed as

$$\mathcal{S}_{arepsilon}(oldsymbol{p}) = oldsymbol{S}_0(arepsilon) + \sum_{k=1}^N p_k oldsymbol{S}_k(arepsilon).$$

Following from the linearity of the scaling (2) where $S_k(\varepsilon)$ depend linearly on ε . We have the following lemma relating $S_k(\varepsilon)$ to S_k for the missing values.

Lemma 5.1. Under Assumption 3.3, the scaled basis matrices for the missing values can be expressed as:

$$S_k(\varepsilon) = \varepsilon S_k, \quad \text{for all } k \in \{n+1, \dots, n+m\}.$$
 (14)

Proof. Since the scaling (2) amounts just to multiplying by the matrix S(p) on the right and by linearity of (5) we have that, for all k,

$$oldsymbol{S}_k(arepsilon) = oldsymbol{S}_k egin{bmatrix} oldsymbol{I}_{K-\kappa} & oldsymbol{0} \ oldsymbol{0} & arepsilon oldsymbol{I}_\kappa \end{bmatrix},$$

Since, by Assumption 3.3, the limiting structure $S_0(p)$ (and thus the left block of S(p)) does not depend on p_{n+1}, \ldots, p_{n+m} , we obtain (14).

In what follows, we fix $\hat{p} = \hat{p}^{(\text{RMIN})}$ and assume that one of Assumptions 3.2 to 3.3 holds. For the scaled matrix structure $S_{\varepsilon}(\hat{p})$ we look at the corresponding projectors $Q_1(\varepsilon), Q_2(\varepsilon)$ (as defined as in Section 4.1), and show that these projectors have the following form.

Lemma 5.2. Let \hat{p} be fixed. Then:

1. Under Assumption 3.2, the projector $Q_1(\varepsilon)$ on the left nullspace of $S_{\varepsilon}(\hat{p})$ is constant for any ε (even, $\varepsilon = 0$)

$$\boldsymbol{Q}_1(\varepsilon) = \boldsymbol{Q}_1$$

where $Q_1 = Q_1(0)$ is the projector on the left nullspace of $S_0(\hat{p})$.

2. Under Assumption 3.3, the matrix $Q_2(0)$ (the projector on the right nullspace of $S_0(\hat{p})$) has the form

$$\boldsymbol{Q}_{2}(0) = \begin{bmatrix} * & \boldsymbol{0}_{(K-\kappa) \times \kappa} \\ \boldsymbol{0}_{\kappa \times (K-\kappa)} & \boldsymbol{I}_{\kappa} \end{bmatrix};$$

in particular, $\boldsymbol{S}_k \boldsymbol{Q}_2(0) = \boldsymbol{S}_k$ for $k \in \{n+1, \dots, n+m\}$

3. Under Assumptions 3.2 to 3.3, the right projector $Q_2(\varepsilon)$ has the expansion

$$\boldsymbol{Q}_2(\varepsilon) = \boldsymbol{Q}_2(0) + o(\varepsilon),$$

and thus is continuous in a neighborhood of $\varepsilon = 0$.

- *Proof.* 1. Due to Assumption 3.2, the range of $S_{\varepsilon}(p)$ does not change with ε and is equal to the range of $S_0(\hat{p})$. Therefore, the left nullspace also does not change with ε and $Q_1(\varepsilon) = Q_1(0)$.
 - 2. Note that we have

$$\mathcal{S}_0(\widehat{oldsymbol{p}}) = egin{bmatrix} * & oldsymbol{0}_{L imes\kappa} \end{bmatrix}$$

and therefore the right projector has the form

$$\boldsymbol{Q}_{2}(0) = \boldsymbol{I}_{K} - \left(\mathcal{S}_{0}(\boldsymbol{\hat{p}})^{\dagger} \mathcal{S}_{0}(\boldsymbol{\hat{p}}) \right) = \boldsymbol{I}_{K} - \begin{bmatrix} * & \boldsymbol{0}_{(K-\kappa)\times\kappa} \\ \boldsymbol{0}_{\kappa\times(K-\kappa)} & \boldsymbol{I}_{\kappa} \end{bmatrix}.$$
(15)

In particular, this implies that $S_k Q_2(0) = S_k$ for $k \in \{n+1, \ldots, n+m\}$ since by Assumption 3.3 only the last κ columns of such S_k are nonzero.

By Assumption 3.2, the rank of S_ε(p̂) is constant for all ε ∈ ℝ. Therefore, the right nullspace projector Q₂(ε) is continuous and analytic in a neighborhood of 0 (for example, due to the analyticity of the projectors [Kat95, Ch. II, Theorem 1.10]).

Finally, we return to the unscaled case, to demonstrate the implications of the assumptions on the solution of the rank minimization problem.

Lemma 5.3. Under Assumptions 3.2–3.3, a vector $\hat{p} = \hat{p}^{(RMIN)}$ is a minimizer of (RMIN) if and only if it is the solution of the following linear system of equations:

$$\begin{cases} \boldsymbol{Q}_1 \mathcal{S}(\hat{\boldsymbol{p}}) = 0, \\ \hat{\boldsymbol{p}}_{1:n} = \boldsymbol{p}_0, \end{cases}$$
(16)

where Q_1 is the projector on the left nullspace of $S_0(p_0)$.

Proof. By Assumption 3.2, $\hat{p} = \hat{p}^{(\text{RMIN})}$ is an optimal solution to (RMIN) if and only if

$$\operatorname{rank}\{\mathcal{S}(\widehat{\boldsymbol{p}})\} = \operatorname{rank}\{\mathcal{S}_0(\boldsymbol{p}_0)\}.$$
(17)

Next, we notice that the nonzero part of $S_0(p_0)$ is a submatrix of $S(\hat{p})$, and therefore (17) holds if and only if the column space (resp. left nullspace) of $S(\hat{p})$ coincides with the column space (resp. the left nullspace) of $S_0(p_0)$, which happens if and only if (16) is satisfied.

5.2 Candidate dual certificate for the scaled case

In order to get closed-form expressions of the matrices in (9), we require additional notation. Let S_k be the basis matrices in (5). We define the following matrix S':

$$\boldsymbol{S}' = \left[\operatorname{vec}(\boldsymbol{S}_{n+1}) \quad \cdots \quad \operatorname{vec}(\boldsymbol{S}_{n+m}) \right] \in \mathbb{R}^{LK \times m}.$$
(18)

Using such notation, the following lemma holds true.

Lemma 5.4. Under both Assumptions 3.2 and 3.3, we have that, for the scaled sturcture $S_{\varepsilon}(\mathbf{p})$ the matrices $\mathbf{A}(\varepsilon)$ and $\mathbf{b}(\varepsilon)$ defined in (9) have the form

$$\boldsymbol{A}(\varepsilon) = \varepsilon(\boldsymbol{S}')^{\mathsf{T}}(\boldsymbol{Q}_{2}^{\mathsf{T}}(\varepsilon) \otimes \boldsymbol{Q}_{1}) = \varepsilon\left((\boldsymbol{S}')^{\mathsf{T}}(\boldsymbol{I} \otimes \boldsymbol{Q}_{1}) + \varepsilon\right),\tag{19}$$

$$\boldsymbol{b}(\varepsilon) = -\varepsilon(\boldsymbol{S}')^{\mathsf{T}} \operatorname{vec}(\boldsymbol{B}(\varepsilon)).$$
(20)

Proof. We first note that

$$\boldsymbol{Q}_1(\varepsilon)\boldsymbol{S}_{n+k}(\varepsilon)\boldsymbol{Q}_2(\varepsilon) = \varepsilon \boldsymbol{Q}_1 \boldsymbol{S}_{n+k} \boldsymbol{Q}_2(\varepsilon) = \varepsilon \boldsymbol{Q}_1 \boldsymbol{S}_{n+k} (\boldsymbol{Q}_2(0) + o(\varepsilon)) = \varepsilon (\boldsymbol{Q}_1 \boldsymbol{S}_{n+k} + o(\varepsilon)),$$

where the last three equalities are due to Lemma 5.2. By vectorizing, rows of $A(\varepsilon)$ can be expressed as

$$\operatorname{vec}(\boldsymbol{Q}_1\boldsymbol{S}_{n+1}\boldsymbol{Q}_2(\varepsilon)) = \varepsilon(\boldsymbol{Q}_2^{\mathsf{T}}(\varepsilon)\otimes\boldsymbol{Q}_1)\operatorname{vec}(\boldsymbol{S}_{n+k}) = \varepsilon\left((\boldsymbol{I}\otimes\boldsymbol{Q}_1)\operatorname{vec}(\boldsymbol{S}_{n+k}) + o(\varepsilon)\right),$$

which, after transposition, yields the desired result for $A(\varepsilon)$. Similarly, for $b(\varepsilon)$ we get

$$(\boldsymbol{b}(\varepsilon))_k = -\operatorname{vec}(\boldsymbol{S}_{n+k}(\varepsilon))^{\mathsf{T}}\operatorname{vec}(\boldsymbol{B}(\varepsilon)) = -\varepsilon\operatorname{vec}(\boldsymbol{S}_{n+k})^{\mathsf{T}}\operatorname{vec}(\boldsymbol{B}(\varepsilon)),$$

which completes the proof.

In order to use (12) and (13), we require that $A(\varepsilon)$ is full row rank. The proof relies on Lemma 5.3 to show that the main term of (19) is full rank.

Proposition 5.5. Under Assumptions 3.1–3.3, rank $\{A(\varepsilon)\} = m$ in some neighborhood of 0 (for all $\varepsilon \in (0, \varepsilon_1)$ for some ε_1).

Proof of Proposition 5.5. The matrix $A(\varepsilon)$ has m rows, and therefore, by (19), it is sufficient to show that rank $\{S'^{\mathsf{T}}(I \otimes Q_1)\} = m$, which, by semi-continuity of the rank would imply rank $\{A(\varepsilon)\} = m$ in a neighborhood of 0.

To do that, we make use of Lemma 5.3. Let us take the parameterization of feasible \hat{p} in (NNMIN) $\hat{p} = \hat{p}(z) = (p_0, z)$, for $z \in \mathbb{C}^m$ and denote $x_0 = \text{vec } S(\hat{p}(0))$. Then, by applying vectorization, we have that

$$\operatorname{vec}(\mathcal{S}(\widehat{p})) = x_0 + S'z.$$

Therefore, the condition (16) can be equivalently rewritten as $(I \otimes Q_1)(S'z + x_0) = 0$, or after regrouping the terms,

$$(\boldsymbol{I} \otimes \boldsymbol{Q}_1)\boldsymbol{S}'\boldsymbol{z} = -(\boldsymbol{I} \otimes \boldsymbol{Q}_1)\boldsymbol{x}_0.$$
⁽²¹⁾

By Assumption 3.1, the solution of (21) exists and is unique, therefore, the matrix $(I \otimes Q_1)S'$ must be full column rank, otherwise there is another solution to (21) (and therefore, another solution to (16)). But the matrix $((I \otimes Q_1)S')$ is just the transpose of $(S')^{\mathsf{T}}(I \otimes Q_1)$, therefore, the proof is complete. \Box

We finish this section with a lemma on the form of $b(\varepsilon)$ from (20).

Lemma 5.6. Under Assumptions 3.1–3.3, the vector $\mathbf{b}(\varepsilon)$ is analytic in a small neighborhood of 0 and has an expansion:

$$\boldsymbol{b}(\varepsilon) = \varepsilon^2 (\boldsymbol{b}_2 + o(\varepsilon)). \tag{22}$$

Proof. We first show that $B(\varepsilon)$ is real-analytic in a neighborhood of ε . For this we first note that (thanks to Lemma 5.2) there is an semi-unitary (i.e., $W^{\mathsf{H}}W = I$) matrix $W \in \mathbb{C}^{L \times r}$ spanning the column space of $S_{\varepsilon}(\widehat{p})$ for any ε . Then we have that

$$S_{\varepsilon}(\widehat{p}) = WW^{\mathsf{H}}S_{\varepsilon}(\widehat{p}), \tag{23}$$

where $Z(\varepsilon) := W^{\mathsf{H}} S_{\varepsilon}(\hat{p})$ is full row rank for all $\varepsilon \in \mathbb{R}$. Therefore, the polar factor (for real ε) can be expressed as (see e.g., [GL17])

$$\boldsymbol{B}(\varepsilon) = \boldsymbol{W}(\boldsymbol{Z}(\varepsilon)\boldsymbol{Z}^{\mathsf{H}}(\varepsilon))^{-\frac{1}{2}}\boldsymbol{Z}(\varepsilon),$$

and therefore is real-analytic in ε in a neighborhood of 0:

$$\boldsymbol{B}(\varepsilon) = \boldsymbol{B}_0 + \varepsilon \boldsymbol{B}_1 + o(\varepsilon). \tag{24}$$

Now recall that the right block of $S_0(\hat{p})$ is zero by (3), hence,

$$\boldsymbol{B}_0 = \boldsymbol{W}(\boldsymbol{Z}(0)\boldsymbol{Z}^{\mathsf{H}}(0))^{-\frac{1}{2}}\boldsymbol{W}^{\mathsf{H}}\mathcal{S}_0(\widehat{p}) = \begin{bmatrix} * & \boldsymbol{0}_{L\times\kappa} \end{bmatrix}.$$

Therefore, by Lemma 5.2 (i.e., from Assumption 3.3), we have

 $\langle \boldsymbol{B}_0, \boldsymbol{S}_k \rangle = 0$ for all $k \in \{n+1, \dots, n+m\}$.

Combining (20) and (24), we get $\boldsymbol{b}(\varepsilon) = \varepsilon^2 (\boldsymbol{b}_2 + o(\varepsilon))$ with

$$\boldsymbol{b}_2 = -(\boldsymbol{S}')^{\mathsf{T}} \operatorname{vec}(\boldsymbol{B}_1) \tag{25}$$

which completes the proof.

5.3 The proof of the main theorem

Proof of Theorem 3.3. Now let us denote by $M^*(\varepsilon)$ the candidate dual certificate (Definition 4.2) for the scaled case. Thanks to Proposition 5.5, we have rank $\{A(\varepsilon)\} = m$ for small ε , therefore, by (13), we have that

$$\|\boldsymbol{M}^*(\varepsilon)\|_F^2 = \boldsymbol{b}^{\mathsf{H}}(\varepsilon)(\boldsymbol{A}(\varepsilon)\boldsymbol{A}^{\mathsf{H}}(\varepsilon))^{-1}\boldsymbol{b}(\varepsilon).$$

Note that from (19) we have

$$\boldsymbol{A}(\varepsilon)\boldsymbol{A}^{\mathsf{H}}(\varepsilon) = \varepsilon^{2}((\boldsymbol{S}')^{\mathsf{T}}(\boldsymbol{I}\otimes\boldsymbol{Q}_{1})(\boldsymbol{S}') + o(\varepsilon)),$$

which is real analytic in a neighborhood of 0, with $((S')^{\mathsf{T}}(I \otimes Q_1)(S')$ nonsingular by Proposition 5.5. Therefore, we get that, from (22):

nerefore, we get that, from (22):

$$\|\boldsymbol{M}^{*}(\varepsilon)\|_{F}^{2} = \varepsilon^{2}(\boldsymbol{b}_{2} + o(\varepsilon))^{\mathsf{H}} \left(\boldsymbol{A}(\varepsilon)\boldsymbol{A}^{\mathsf{H}}(\varepsilon)\right)^{-1} \varepsilon^{2}(\boldsymbol{b}_{2} + o(\varepsilon))$$
$$= \varepsilon^{2} \left(\boldsymbol{b}_{2}^{\mathsf{H}} \left((\boldsymbol{S}')^{\mathsf{T}}(\boldsymbol{I} \otimes \boldsymbol{Q}_{1})(\boldsymbol{S}')\right)^{-1} \boldsymbol{b}_{2} + o(1)\right),$$

which is analytic in a neighborhood of 0. This implies that there exists a neighborhood of 0 such that

$$\|\boldsymbol{M}^*(\varepsilon)\|_2^2 \le \|\boldsymbol{M}^*(\varepsilon)\|_F^2 < 1,$$

which concludes the proof.

Finally, we give the proof of Proposition 4.3.

Proposition 5.7 (Proposition 4.3, reformulated). Under Assumptions 3.1 to 3.3 let $\hat{p} = \hat{p}^{(RMIN)}$, and let $M^*(\varepsilon)$ denote the candidate dual certificate for the scaled problem (NNMIN- ε). Then, for small $\varepsilon > 0$, the candidate dual certificate exists, is unique, and its squared Frobenius norm (see (13)) has the expansion

$$\|\boldsymbol{M}^*(\varepsilon)\|_F^2 = m_0\varepsilon^2 + o(\varepsilon^2),$$

where m_0 is a constant defined as follows:

$$m_0 = \boldsymbol{b}_2^{\mathsf{H}}((\boldsymbol{S}')^{\mathsf{T}}(\boldsymbol{I}\otimes\boldsymbol{Q}_1)\boldsymbol{S}')^{-1}\boldsymbol{b}_2.$$

The term \mathbf{b}_2 has the expansion $\mathbf{b}_2 = -(\mathbf{S}')^{\mathsf{T}} \operatorname{vec}(\mathbf{B}_1)$ (see (25), the coefficient corresponding to ε^2 from (22)), where \mathbf{B}_1 is the first-order term in the expansion (24), which can be found as

$$\boldsymbol{B}_1 = \boldsymbol{U}_0 \boldsymbol{\Sigma}_0^{-1} \boldsymbol{U}_0^{\mathsf{H}} \boldsymbol{\Delta}; \tag{26}$$

where in (26), U_0 and Σ_0 are the terms of the SVD of $S_0(\widehat{p}) = U_0 \Sigma_0 V_0^{\mathsf{H}}$ and

$$oldsymbol{\Delta} := egin{pmatrix} oldsymbol{0}_{K-\kappa} & oldsymbol{0} \\ oldsymbol{0} & oldsymbol{I}_\kappa \end{bmatrix} \end{pmatrix}$$

is the notation that we introduce for the unscaled block of $S_{\varepsilon}(\hat{p})$ (so that $S_{\varepsilon}(\hat{p}) = S_0(\hat{p}) + \varepsilon \Delta$).

Proof. To show that B_1 has the form (26), we denote $Z(\varepsilon)$ as in (23), with a particular choice of $W = U_0$. Recall that $Z(\varepsilon)$ is a perturbation of a full-row-rank matrix Z_0

$$\boldsymbol{Z}(\varepsilon) = \boldsymbol{Z}_0 + \varepsilon \boldsymbol{Z}_1,$$

with

$$oldsymbol{Z}_0 = oldsymbol{U}_0^{\mathsf{H}} \mathcal{S}_0(\widehat{oldsymbol{p}}), \quad oldsymbol{Z}_1 = oldsymbol{U}_0^{\mathsf{H}} oldsymbol{\Delta}_2$$

where $\langle \mathbf{Z}_0, \mathbf{Z}_1 \rangle_F = 0$. Then from [GL17, Lemma 3.2, eq. (28)], we get that the polar factor of $\mathbf{Z}(\varepsilon)$ has expansion

$$\mathcal{P}(\boldsymbol{Z}(\varepsilon)) = \mathcal{P}(\boldsymbol{Z}_0) + \varepsilon(\boldsymbol{Z}_0 \boldsymbol{Z}_0^{\mathsf{H}})^{-\frac{1}{2}} \boldsymbol{Z}_1 + o(\varepsilon).$$

Bringing it all together and using $Z_0 = U_0^{\mathsf{H}} \mathcal{S}_0(\widehat{p}) = \Sigma_0 V_0^{\mathsf{H}}$, we get

$$B_1 = U_0 (Z_0 Z_0^{\mathsf{H}})^{-\frac{1}{2}} Z_1 = U_0 (\Sigma_0^2)^{-\frac{1}{2}} Z_1 = U_0 (\Sigma_0)^{-1} U_0^{\mathsf{H}} \Delta$$

which completes the proof.

6 Numerical examples

In this section we describe some expository numerical examples describing how the scaling proposed enables exact recovery. We later focus on forecasting time series and data-driven simulation to offer just two application areas. Exact recovery guarantees obtained via nuclear norm minimization for structured matrices can be extended beyond these examples based on Hankel-type structures to other formats that naturally appear in diverse applications. For example, Toeplitz matrices—which often model convolution operators in imaging, communications, and system identification—arise when linear systems exhibit time-invariant or shift-invariant properties. Circulant matrices, which are a special case of Toeplitz matrices where each row is a cyclic shift of the previous one, are common in the analysis of digital filters and in channel estimation problems because they diagonalize under the discrete Fourier transform. Such structures also emerge in array processing and radar signal detection where block-Toeplitz or block-circulant matrices characterize multi-channel or spatial correlations. In these contexts, exact recovery results can be used to robustly interpolate missing measurements or reconstruct degraded signals while preserving the inherent matrix structure. This not only leads to more efficient computational algorithms

(by exploiting fast Fourier transform techniques in the circulant case) but also reinforces the reliability of system identification and deconvolution methods in applications ranging from advanced medical imaging to network communications. Several application areas for which the result in this paper has significant consequence are also described in [Mar19] and the references therein.

In our current implementation, we leverage CVX [GB14, GB08], a software for convex optimization, which provides a rapid and versatile framework for our computations. This setup enables us to efficiently handle a variety of problem instances without extensive tailoring to specific applications. Due to the non specific optimization routines we have used in this paper, we may experience some loss in efficiency since it is not specifically optimized for any particular problem or application. This generality means that while our approach can be applied to a wide range of scenarios, it might not be as finely tuned or perform as well as methods designed for specific tasks. However, if needed, our method can serve as an effective initial step for non-convex optimization routines.

Efficient algorithms for matrix completion leverage a variety of advanced computational techniques to improve performance and scalability. One primary approach involves first-order optimization methods, such as those described in [Bec17]. Randomized linear algebra also plays a crucial role in efficient matrix completion [HMT11]. By using random sampling and projections, these methods can approximate the leading singular values and vectors with high accuracy while significantly reducing the computational load compared to traditional deterministic algorithms. The efficient computation of SVD is particularly important in matrix completion. Structured matrices further enhance the efficiency of SVD computations (e.g., [LQ00]). Algorithms that exploit these structures can perform SVD more rapidly.

6.1 The rank-1 case

In this example, we revisit the numerical study of [UC16], and take $p_k = \lambda^k$, $\lambda \in \mathbb{R}$, with structure as described in (1). The aim is to investigate for which range of λ and ε do the solutions of (RMIN) and (NNMIN- ε) coincide. It is shown in [UC16, Thm. 6] that for the case with no matrix scaling (i.e. $\varepsilon = 1$) if $|\lambda| < 1$ then the solution of (NNMIN) is unique and coincides with the solution of (RMIN), namely $p_{n+k} = \lambda^{n+k}$. Otherwise, the solutions do not coincide. In this numerical exercise we show that the scaling (2) and consequent solving of the scaled optimization problem (NNMIN- ε) indeed increases the range of $|\lambda|$ for which we get exact (perfect) recovery of missing values. We use (12) to compute candidate dual certificates to evidence this.



Figure 1: Plot of the Frobenius norm of the candidate dual certificate as given in (13) against λ for the form of p as described, with n = 10, L = 4, $\kappa = m$.

Figure 1 is a plot of the Frobenius norm of the candidate dual certificate as given in (13) against λ for the form of p as described, with n = 10, L = 4, $\kappa = m$, for different numbers of unknown observations m and scaling ε , as stated within the figure. The horizontal 'dashed' lines denote the value when the Frobenius norm of the candidate dual certificate is one. By Lemma 4.1, when this norm is strictly less than one then the solution of (NNMIN- ε) is unique, coincides with that of (RMIN), and thus

we have exact recovery of the *m* missing values. The range of λ for which we have exact recovery is a function of the scaling parameter ε , the range increases as more scaling (in the sense of taking ε small) is applied. For *p* with exponential growth, the more observations that are unknown, then the more scaling is required for exact recovery. Note that the point at which the solution with no scaling ($\varepsilon = 1$) intersects the horizontal 'success' line, is the value of λ as given in [GU18, Corollary 4.3], which gives bounds on λ for which the solution of (NNMIN) coincides with that of (RMIN).

6.2 The rank-r case

We extend the example of the previous section for a more challenging problem, and instead take $p_k = \sum_{j=1}^r \lambda_j^k$, $\lambda_j \in \mathbb{R}$, with again Hankel structure as described in (1). We take m = 1, L = 10, n = 20, $\kappa = 1$ and each λ_j is independently sampled uniformly in the range (0, u), for some selected u > 0. We explore the influence of u and r and how they interact with the necessary amount of scaling (dictated by the parameter ε) to yield the Frobenius norm of the candidate dual certificate as given in (13) to be smaller than 1.



Figure 2: Plots (3-d and contour) of the proportion of times the Frobenius norm of the candidate dual certificate as given in (13) exceeds one, with ε (lhs) against r = 1, 2, ..., 8 (bottom), for different u. Parameter settings are $m = 1, L = 10, n = 20, \kappa = 1$.

Figure 2 contains plots (3-d and contour) of the proportion of times this Frobenius norm of the candidate dual certificate (13) exceeds one, with ε plotted against r = 1, 2, ..., 8, for different u. We take 10,000 Monte-Carlo simulations. The amount of scaling required for exact recovery of the missing

values increases with r and u. This is consistent with the intuition as written in [UC16] and [GU18], that the observed vector needs to be sufficiently damped in order for (NNMIN) to yield an identical solution to (RMIN).

In the examples considered so far, we have opted to compute the Frobenius norm of the candidate dual certificate, $||M^*||_F$, as given in (13), instead of its spectral norm $||M^*||_2$ by virtue of Theorem 3.3, thus obtaining a weaker sufficient condition. We now investigate the gap between these norms. We take the same signal as described at the opening of this section, with L = 10, n = 20, r = 3 and $\kappa = m$, and inspect the ratio $||M^*||_F/||M^*||_2$ for different m and u.

Figures 3 and 4 contain boxplots of the ratio $||M^*||_F/||M^*||_2$ for different m and u (as stated in the caption), $\varepsilon = 0.5$ and $\varepsilon = 0.1$ respectively, taken over 10,000 Monte-Carlo simulations. The ratio appears more variable as the number of missing values m increases, and for small u the ratio seems to be a piece-wise linear function of m, increasing until beginning to (approximately) plateau after a particular number of missing observations. For small ε the ratio is smaller.



Figure 3: Boxplots of the ratio $\|M^*\|_F / \|M^*\|_2$ for different m and u (as stated in the caption), $\varepsilon = 0.5$.



Figure 4: Boxplots of the ratio $\|M^*\|_F / \|M^*\|_2$ for different m and u (as stated in the caption), $\varepsilon = 0.1$.

We now discuss the result of Proposition 5.7. Using the notation as given in this proposition, Figure 5 is a plot of the ratio $\|M^*\|_F / \varepsilon \sqrt{m_0}$ against ε and r for the example considered with $\kappa = 1$, for different u, again taken over 10,000 Monte-Carlo simulations. The leading term approximation of the Frobenius norm of the candidate dual certificate M^* can be seen to be an excellent approximation, with deviation only for more 'complex' time series for large r or u, but this deviation can be mitigated with a selection of a smaller ε .

6.3 Forecasting time series

In this section we provide figures directly showing the impact of scaling upon forecasting a given signal/time series. Take $p_k = \exp(\lambda k) \sin(2\pi\omega k)$, $\lambda = \omega = 0.1$, with Hankel structure as described in (1). We take m = 5, L = 4, n = 10, $\kappa = 5$. Practically one may view the problem considered here



Figure 5: Ratio $\|M^*\|_F / \varepsilon \sqrt{m_0}$ against ε and r for the example considered with $\kappa = 1$ and different u.

as forecasting ahead five observations having already recorded ten exact observations of a signal/time series.

Figure 6 contains plots of the signal p with the solutions obtained from (NNMIN- ε) and (NNMIN) (i.e., nuclear norm minimization with and without scaling, respectively). The failure of (NNMIN) to exactly recover the time series is clearly evident, and the introduction of scaling gives exact recovery of the m = 5 missing values once ε is taken sufficiently small. The final value of $\varepsilon \approx 0.0694$ is that found from (8), showing its value in selecting the scaling parameter ε in practice.



Figure 6: Plots of the signal p (blue curve) with the solutions obtained from (NNMIN- ε) (black cross, with ε as given) and (NNMIN) (orange circle).

6.4 Data-driven simulation

In this subsection, we describe an example of structured matrix completion that appears in modeling (linear time-invariant) dynamical systems (with p inputs and m outputs) that transform an input sequence $u(t) \in \mathbb{R}^p$ to the output sequence $y(t) \in \mathbb{R}^m$. The simulation problem [Mar17] aims at computing the output time series for a given dynamical system, input and initial conditions; the system parameters often have to be estimated prior to simulation. The data-driven approach to simulation [Mar19, DM19] avoids the step of estimating system parameters and is rather based on previously recorded data of inputs and outputs produced by the system. As shown in [Mar19, DM19], data-driven simulation can be reformulated as structured matrix completion. More details on the data-driven simulation problem are provided in Appendices 7.3 and 7.4, and we give below just necessary details to formulate the matrix completion problem.

Let $u_d \in \mathbb{R}^{m \times T_d}$ and $y_d \in \mathbb{R}^{p \times T_d}$ be some known input-output pair of time series ("data")

$$u_d = (u_d(1), \dots, u_d(T_d)), \quad y_d = (y_d(1), \dots, y_d(T_d)),$$

and a new input time series $m{u}_s \in \mathbb{R}^{m imes T_s}$ the goal is to find the corresponding output $m{y}_s \in \mathbb{R}^{p imes T_s}$

$$\boldsymbol{u}_s = (\boldsymbol{u}_s(1), \dots, \boldsymbol{u}_s(T_s)), \quad \boldsymbol{y}_s = \left(\boldsymbol{y}_s(1), \dots, \boldsymbol{y}_s(\ell), \underbrace{\boldsymbol{y}_s(\ell+1), \dots, \boldsymbol{y}_s(T_s)}_{\text{unknown}}\right)$$

assuming the knowledge of its first $\ell \ll T_s$ values (initial conditions).

By choosing the parameter $L \ge \ell + 1$ and denoting $K_d = T_d - L + 1$ and $K_s = T_s - L + 1$, the approach of [DM19] consists in minimizing the rank of the following $L(m + p) \times (K_d + K_s)$ matrix

	$oldsymbol{u}_d(1)$	$oldsymbol{u}_d(2)$		$oldsymbol{u}_d(K_d)$	$oldsymbol{u}_s(1)$	$oldsymbol{u}_s(2)$	·	···		$\boldsymbol{u}_s(K_s)$	
$\mathcal{M}_L =$	$oldsymbol{u}_d(2)$	$oldsymbol{u}_d(3)$		÷	$oldsymbol{u}_s(2)$	$oldsymbol{u}_s(3)$	· · ·	·		÷	
	:	.·*		÷	÷	·	.·*	···		÷	
	$\boldsymbol{u}_d(L)$	· · ·		$\boldsymbol{u}_d(T_d)$	$oldsymbol{u}_s(L)$		· · ·			$oldsymbol{u}_s(T_s)$	
	$oldsymbol{y}_d(1)$	$oldsymbol{y}_d(2)$	• • •	$oldsymbol{y}_d(K_d)$	$oldsymbol{y}_s(1)$	•••	$oldsymbol{y}_s(\ell)$	$\boldsymbol{y}_s(\ell+1)$	• • •	$\boldsymbol{y}_s(K_s)$	
	$oldsymbol{y}_d(2)$	$oldsymbol{y}_d(3)$		÷	•					÷	•
	•	· · [·]		÷	$oldsymbol{y}_s(\ell)$. · ·	· · ·		÷	
	÷	.·*		÷	$\boldsymbol{y}_s(\ell+1)$		[.]	.·*		÷	
	÷	· · ·	•••	÷			[.]	[.]	•••	÷	
	$oldsymbol{y}_d(L)$			$oldsymbol{y}_d(T_d)$	ig igyinup igyinupig		.·*	.· [.]		$\boldsymbol{y}_s(T_s)$	
										((27)

We consider a randomly generated single-input single-output (m = p = 1) linear time-invariant system of order $\ell = 6$ (generated by MATLAB's function DRSS). The input is a standard normal zeromean white Gaussian noise sequence u_d of length $T_d = 40$ and we denote the corresponding (noise-free) output as y_d . The simulation input u_s of length $T_s = 20 + \ell$ is generated in the same way, and the (noisefree) simulation output is denoted by \tilde{y}_s . All simulations assume to start from zero initial conditions. We set the number of block rows L of the Hankel matrix to $L = \ell + 1$.

To study the effect of the data scaling, we replace the data $w_s = (u_s, y_s)$ by εw_s and investigate the results for values of ε in the range $[10^{-3}, 10^3]$. The nuclear norm minimization using the proposed block-column scaling (NNMIN- ε) returns the (scaled) simulation output \hat{y}_s , which is then compared to the true simulation output \tilde{y}_s after rescaling the result by multiplying it with ε^{-1} . Figure 7 contains the result of this analysis, showing the relative simulation error $\|\hat{y}_s - \tilde{y}_s\|/\|\tilde{y}_s\|$ with \tilde{y}_s the true simulation output. At large values $\varepsilon \gg 1$ the simulation signal is close to a zero signal (and the relative error is one). As ε decreases, the reconstruction error gradually decreases and for a sufficiently small value $\varepsilon \ll 1$ the method returns the true simulation output \tilde{y}_s .

7 Conclusion

We have introduced a technique to achieve exact recovery of missing entries of structured matrices using the commonly used nuclear norm relaxation of low-rank matrix completion where the observations are recorded without noise. To do this we use block-column scaling of the given matrix. As well as describing the theoretical basis of this approach we have offered several examples showing the usefulness and merit of our proposed technique. We provide a result suggesting a value for our scaling parameter which will be helpful for practical use of our results. There is significant potential for synergizing our theoretical advancements with several practical applications in fields such as time-series analysis, datadriven simulation, and more. Future work will consider the hard problem of when observations are noisy and not exact, and to derive conditions for when Assumption 3.1 is verifiable for general affine matrix structures.



Figure 7: The relative reconstruction error in the data-driven simulation problem exhibits a transition to a zero error for sufficiently small values of the scaling parameter ε . This effect is shown on the simulated signal where large values of $\varepsilon \gg 1$ result in a zero simulation output (relative error of 1), while small values of $\varepsilon \ll 1$ lead to a perfect reconstruction of the signal in the noiseless case, indicated by the bold line connecting the circled points. We additionally show simulation results for intermediate values of ε in grey lines that have a nonzero relative error — their corresponding relative errors are indicated on the left panel by circles.

Appendix

7.1 Forecasting for Hankel matrices

Proof of Lemma 3.5. Note that $S_0(p)$ depends only on $p'_0 = (p_0)_{1:n+m-\kappa}$

$$\mathcal{S}_0(oldsymbol{p}) := egin{bmatrix} \mathcal{H}_L(oldsymbol{p}_0') & oldsymbol{0}. \end{bmatrix}$$

where $\mathcal{H}_L(p'_0)$ is a subblock of $\mathcal{H}_L(p_0)$. Under the assumption [GNZ01, §5.3], p'_0 has a unique minimal rank completion for any number of forecasted values.

7.2 Subdifferential of the nuclear norm

Definition 7.1 (Subdifferential of a function f at x). For a convex function $f : \mathbb{R}^N \to \mathbb{R}$, the subdifferential [HUL01, p. 167] of f at x, denoted by $\partial f(x)$, is defined as the set

$$\partial f(\boldsymbol{x}) \stackrel{\mathsf{def}}{=} \left\{ \boldsymbol{z} \in \mathbb{R}^N : f(\boldsymbol{y}) \ge f(\boldsymbol{x}) + \langle \boldsymbol{z}, \boldsymbol{y} - \boldsymbol{x} \rangle, \text{ for all } \boldsymbol{y} \in \mathbb{R}^N \right\}.$$

Notice that if f is differentiable at x, then the subdifferential contains a single element, namely the gradient, i.e., $\partial f(x) = \nabla f(x)$.

The subdifferential of the nuclear norm of a matrix X, is related to the SVD of X as follows. Consider a rank-*r* matrix $X \in \mathbb{R}^{L \times K}$, and its compact SVD as

$$X = U\Sigma V^{\top},$$

with orthonormal $U \in \mathbb{R}^{L \times r}$ and $V \in \mathbb{R}^{K \times r}$, and $\Sigma \in \mathbb{R}^{r \times r}$ a diagonal matrix with the nonzero singular values. Let also $U_{\perp} \in \mathbb{R}^{L \times (L-r)}$ and $V_{\perp} \in \mathbb{R}^{K \times (K-r)}$ be the matrices containing the remaining left and right singular vectors (spanning the left and right nullspaces of X, respectively). Then the subdifferential of the nuclear norm [Wat92, p.41] at X is the following set of matrices:

$$\partial \|\boldsymbol{X}\|_* = \left\{ \begin{bmatrix} \boldsymbol{U} & \boldsymbol{U}_{\perp} \end{bmatrix} \begin{bmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{Z} \end{bmatrix} \begin{bmatrix} \boldsymbol{V} & \boldsymbol{V}_{\perp} \end{bmatrix}^{\mathsf{T}} : \boldsymbol{Z} \in \mathbb{R}^{(L-r) \times (K-r)}, \|\boldsymbol{Z}\|_2 \leq 1 \right\},$$

In what follows, we use a slightly more compact notation. Denote by $B \stackrel{\text{def}}{=} UV^{\mathsf{T}}$ the *polar factor* of X and by Q_1 and Q_2 the orthogonal projectors on the left and right nullspace of X respectively:

$$\boldsymbol{Q}_1 = \boldsymbol{U}_\perp \boldsymbol{U}_\perp^\mathsf{T} = \boldsymbol{I}_L - \boldsymbol{U}\boldsymbol{U}^\mathsf{T}, \quad \boldsymbol{Q}_2 = \boldsymbol{V}_\perp \boldsymbol{V}_\perp^\mathsf{T} = \boldsymbol{I}_K - \boldsymbol{V}\boldsymbol{V}^\mathsf{T}.$$

Then the subdifferential of the nuclear norm at X can be expressed as

$$\partial \|\boldsymbol{X}\|_{*} = \left\{ \boldsymbol{B} + \boldsymbol{Q}_{1} \boldsymbol{M} \boldsymbol{Q}_{2} : \boldsymbol{M} \in \mathbb{R}^{L \times K} \text{ with } \|\boldsymbol{M}\|_{2} \leq 1 \right\}.$$
(28)

Let $f(p) = ||S(p)||_*$ denote the function that maps coefficients p to the nuclear norm of the associated structured matrix S(p). By the chain rule, the subdifferential of f at p is (see [UC16])

$$\partial f(\boldsymbol{p}) = \left\{ \left[\begin{array}{ccc} \langle \boldsymbol{S}_1, \boldsymbol{H} \rangle_F & \cdots & \langle \boldsymbol{S}_N, \boldsymbol{H} \rangle_F \end{array} \right]^\top : \boldsymbol{H} \in \partial \left\| \boldsymbol{X} \right\|_* \text{ and } \boldsymbol{X} = \mathcal{S}(\boldsymbol{p}) \right\},$$
(29)

where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product. Combining (28) with (29) and the standard necessary and sufficient conditions for convex optimization problems, we can easily obtain conditions for the structured nuclear norm minimization case.

7.3 Linear dynamical systems and block-Hankel matrices

Willems' behavioral system theory [PW98] defines a system \mathcal{B} as the set of its admissible trajectories w. A trajectory w of a discrete-time q-variate system \mathcal{B} (in shorthand notation ' $w \in \mathcal{B}$ ') is a sequence $w = (w(1), \ldots, w(T))$, with $w(t) \in \mathbb{R}^q$ for $t = 1, \ldots, T$. In this paper we consider the class of q-variate linear time-invariant (LTI) systems \mathcal{L}^q . An LTI system $\mathcal{B} \in \mathcal{L}^q$ (with m inputs and p outputs, such that q = p + m) has a kernel representation

$$\mathcal{B} = \left\{ \boldsymbol{w} \mid \boldsymbol{R}_0 \boldsymbol{w}(t) + \dots + \boldsymbol{R}_\ell \boldsymbol{w}(t+\ell) = \boldsymbol{0}, \text{ for } t \ge 1 \right\},\tag{30}$$

where $\mathbf{R} = \begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \dots & \mathbf{R}_\ell \end{bmatrix}$, in which $\mathbf{R}_i \in \mathbb{R}^{p \times q}$ is a kernel parameter that specifies the system. This formulation can be viewed as a difference equation describing the admissible trajectories $\mathbf{w} \in \mathcal{B}$. The minimal value for ℓ for which (30) holds, is an invariant of the system \mathcal{B} and is called the *lag*.

For a trajectory w, the block-Hankel matrix is defined similarly to (1) as

$$\mathcal{H}_{L}(\boldsymbol{w}) \stackrel{\text{def}}{=} \begin{bmatrix} \boldsymbol{w}(1) & \boldsymbol{w}(2) & \cdots & \boldsymbol{w}(T-L+1) \\ \boldsymbol{w}(2) & \boldsymbol{w}(3) & \cdots & \vdots \\ \vdots & \vdots & & \vdots \\ \boldsymbol{w}(L) & \boldsymbol{w}(L+1) & \cdots & \boldsymbol{w}(T) \end{bmatrix} \in \mathbb{R}^{qL \times (T-L+1)}.$$
(31)

The block-Hankel matrix captures into the language of linear algebra the linear time-invariance of \mathcal{B} , and is closely connected to the kernel representation of $\mathcal{B} \in \mathcal{L}^q$. For a trajectory $w \in \mathcal{B}$, the block-Hankel matrix $H_L(w)$ with $L \ge \ell + 1$ is (row) rank-deficient, since $RH_{\ell+1}(w) = 0$. More precisely, the rank of the block-Hankel matrix $H_L(w)$ is bounded [Mar19]

$$\operatorname{rank} \boldsymbol{H}_L(\boldsymbol{w}) \le mL + p\ell, \tag{32}$$

for $T \gg L$, and under the persistency of excitation condition.

A trajectory $\boldsymbol{w} \in \mathcal{B}$ admits a partitoning into inputs and outputs: i.e., there exists two time series $\boldsymbol{u}(t) \in \mathbb{R}^m$ (input) and $\boldsymbol{y}(t) \in \mathbb{R}^p$ (output), so that the vector $\boldsymbol{w}(t)$ is a permutation of the elements of the vector $\begin{bmatrix} \boldsymbol{u}(t) \\ \boldsymbol{y}(t) \end{bmatrix}$; we will write $\boldsymbol{w} = (\boldsymbol{u}, \boldsymbol{y})$ for short and assume that the permutation is known.

Remark 7.2 (SISO LTI systems). The case m = p = 1 is referred to as single-input-single-output (SISO); in this case the lag ℓ is equal to the system order. Moreover, the block-Hankel matrix has rank rank $\mathbf{H}_{\ell+1}(\mathbf{w}) \leq 2\ell + 1$ (rank deficiency 1), with generically equality.

7.4 Data-driven simulation as matrix completion

The simulation of a system for a given input signal can be stated as follows. Given a system $\mathcal{B} \in \mathcal{L}^q$, an input u_s , and initial conditions w_i

$$\boldsymbol{w}_{i} = (\boldsymbol{w}_{i}(-\ell+1), \boldsymbol{w}_{i}(-\ell+2), \dots, \boldsymbol{w}_{i}(-1), \boldsymbol{w}_{i}(0)), \qquad (33)$$

find the output y_s such that $w_i \wedge w_s = w_i \wedge (u_s, y_s) \in \mathcal{B}$. This is a basic problem in system theory and is studied in various formulations. In *data-driven simulation*, the system \mathcal{B} is defined implicitly by a given trajectory $w_d = (w_d(1), \ldots, w_d(T_d)) \in \mathcal{B}$, where 'd' stands for data. In this context, the simulation of the output $y_s = (y_s(1), \ldots, y_s(T_s))$ for a given input $u_s = (u_s(1), \ldots, u_s(T_s))$, where 's' stands for simulation, is done without first identifying the system \mathcal{B} . We assume that the input u_d of the given trajectory $w_d = (u_d, y_d)$ is persistently exciting, so w_d completely specifies the system \mathcal{B} [WRMD05]. The data-driven simulation problem can then formally be stated as follows.

Problem 7.3 (Exact data-driven simulation).

Given a trajectory $w_d = (u_d, y_d) \in \mathcal{B} \in \mathcal{L}^q$, an input u_s , and initial conditions w_i , find the output y_s , such that $w_i \wedge w_s = w_i \wedge (u_s, y_s) \in \mathcal{B}$.

Remark 7.4. If we assume that the initial conditions are zero, i.e., $w_i \equiv 0$, the specification of initial conditions is done by prepending the trajectory w with ℓ zeros.

The rank-deficiency of the Hankel matrix $H_L(w)$, for $L \ge \ell + 1$, is closely related to the kernel representation of the system \mathcal{B} and its interpretation in terms of difference equations. The mosaic Hankel matrix $\begin{bmatrix} H_L(w') & H_L(w'') \end{bmatrix}$ built from two trajectories $w', w'' \in \mathcal{B}$ has the same rank as $H_L(w)$, since w' and w'' satisfy the same difference equations, so rank $\begin{bmatrix} H_L(w') & H_L(w'') \end{bmatrix} \le mL + p\ell$.

This observation is crucial for formulating the data-driven simulation problem as a matrix completion problem. Putting $w_d = (u_d, y_d)$ and $w_s = (u_s, y_s)$ (where w_s is only partially known), the unknown output trajectory y_s should be determined such that

$$\operatorname{rank} \begin{bmatrix} \boldsymbol{H}_L(\boldsymbol{w}_d) & \boldsymbol{H}_L(\boldsymbol{w}_s) \end{bmatrix} \le mL + p\ell, \text{ for } L \ge \ell + 1,$$
(34)

implying that both trajectories belong to the same system \mathcal{B} . Remark that again we need to specify initial conditions to uniquely determine the simulation output \boldsymbol{y}_s . If we denote by $\boldsymbol{w}'_s = \boldsymbol{w}_i \wedge \boldsymbol{w}_s$ the simulation trajectory prepended with initial conditions, then data-driven simulation problem can be phrased as the following block-Hankel matrix completion problem.

Problem 7.5 (Data-driven simulation via Hankel completion). *Given a trajectory* $w_d = (u_d, y_d) \in \mathcal{B} \in \mathcal{L}^q$, an input u_s , and initial conditions w_i , find the output y_s from the following minimization problem.

$$\min_{\boldsymbol{y}_s} \operatorname{rank} \left[\begin{array}{c} \boldsymbol{H}_L(\boldsymbol{w}_d) & \boldsymbol{H}_L(\boldsymbol{w}'_s) \end{array} \right],$$

$$(35)$$

where $\boldsymbol{w}_s = (\boldsymbol{u}_s, \boldsymbol{y}_s)$.

By rearranging the rows of the mosaic Hankel matrix in (35), it is easy to see that Problem 7.5 is equivalent to rank minimization of the matrix in (27) (where the fixed element in the right blocks in (27) are due to initial conditions).

References

- [Bec17] Amir Beck. First-order methods in optimization. SIAM, 2017.
- [BG17] Holly Butcher and Jonathan Gillard. Simple nuclear norm based algorithms for imputing missing data and forecasting in time series. *Statistics and Its Interface*, 10(3):19–25, 2017.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

- [CBSW14] Yudong Chen, Srinadh Bhojanapalli, Sujay Sanghavi, and Rachel Ward. Coherent matrix completion. In *International Conference on Machine Learning*, pages 674–682. PMLR, 2014.
- [CGO22] Marcus Carlsson, Daniele Gerosa, and Carl Olsson. An unbiased approach to low rank recovery. *SIAM Journal on Optimization*, 32(4):2969–2996, 2022.
- [CP10] Emmanuel J. Candès and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [CR12] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119, 2012.
- [CT10] Emmanuel J. Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [DM19] Philippe Dreesen and Ivan Markovsky. Data-driven simulation using the nuclear norm heuristic. In *Proc. of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2019)*, pages 8207–8211, Brighton, UK, 2019.
- [DP15] Liang Dai and Kristiaan Pelckmans. On the nuclear norm heuristic for a Hankel matrix completion problem. *Automatica*, 51:268 272, 2015.
- [FNP⁺21] Simon Foucart, Deanna Needell, Reese Pathak, Yaniv Plan, and Mary Wootters. Weighted matrix completion from non-random, non-uniform sampling patterns. *IEEE Transactions on Information Theory*, 67(2):1264–1290, 2021.
- [GB08] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- [GB14] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. https://cvxr.com/cvx, March 2014.
- [GL17] Evan S. Gawlik and Melvin Leok. Iterative computation of the Fréchet derivative of the polar decomposition. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1354–1379, 2017.
- [GNZ01] Nina Golyandina, Vladimir Nekrutkin, and Anatoly A. Zhigljavsky. *Analysis of time series structure: SSA and related techniques.* CRC press, 2001.
- [GU18] Jonathan Gillard and Konstantin Usevich. Structured low-rank matrix completion for forecasting in time series analysis. *International Journal of Forecasting*, 34(4):582–597, 2018.
- [GU23] Jonathan Gillard and Konstantin Usevich. Hankel low-rank approximation and completion in time series analysis and forecasting: a brief review. *Statistics and Its Interface*, 2023. in press.
- [HMT11] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [HR84] Georg Heinig and Karla Rost. *Algebraic methods for Toeplitz-like matrices and operators*. Birkhäuser, Boston, 1984.
- [HUL01] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of Convex Analysis*. Springer Verlag, Heidelberg, 2001.

- [JONS13] Prateek Jain, Sewoong Oh, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 665–674, 2013.
- [Kat95] Tosio Kato. *Perturbation theory for linear operators*. Springer-Verlag, 2nd corrected edition, 1995.
- [Laz16] Damiana Lazzaro. A nonconvex approach to low-rank matrix completion using convex optimization. *Numerical Linear Algebra with Applications*, 23(5):801–824, 2016.
- [LQ00] Franklin T. Luk and Sanzheng Qiao. A fast eigenvalue algorithm for hankel matrices. *Linear Algebra and Its Applications*, 316(1-3):171–182, 2000.
- [Mar08] Ivan Markovsky. Structured low-rank approximation and its applications. *Automatica*, 44(4):891–909, 2008.
- [Mar12] Ivan Markovsky. How effective is the nuclear norm heuristic in solving data approximation problems? In *Proc. of the 16th IFAC Symposium on System Identification*, pages 316–321, Brussels, 2012.
- [Mar17] Ivan Markovsky. A missing data approach to data-driven filtering and control. *IEEE Trans. Automat. Contr.*, 62:1972–1978, 2017.
- [Mar19] Ivan Markovsky. Low-Rank Approximation. Springer, New York, 2019.
- [MD22] Ivan Markovsky and Florian Dörfler. Data-driven dynamic interpolation and approximation. *Automatica*, 135:110008, 2022.
- [PW98] Jan Willem Polderman and Jan C. Willems. *Introduction to mathematical systems theory: a behavioral approach*. Springer, New York, 1998.
- [SS10] Nathan Srebro and Russ R. Salakhutdinov. Collaborative filtering in a non-uniform world: Learning with the weighted trace norm. *Advances in neural information processing systems*, 23, 2010.
- [TL21] Abiy Tasissa and Rongjie Lai. Low-rank matrix completion in a general non-orthogonal basis. *Linear Algebra and its Applications*, 625:81–112, 2021.
- [TW16] Jacob Tanner and Ke Wei. Low rank matrix completion by alternating steepest descent methods. *Applied and Computational Harmonic Analysis*, 40:417–429, 2016.
- [UC16] Konstantin Usevich and Pierre Comon. Hankel low-rank matrix completion: Performance of the nuclear norm relaxation. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):637–646, 2016.
- [Van13] Benoît Vandereycken. Low-rank matrix completion by riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013.
- [Wat92] G. Alistair Watson. Characterization of the subdifferential of some matrix norms. *Linear Algebra Appl*, 170(1):33–45, 1992.
- [WRMD05] Jan C. Willems, Paolo Rapisarda, Ivan Markovsky, and Bart L. M. De Moor. A note on persistency of excitation. *Syst. Control Lett.*, 54(4):325–329, 2005.
- [XJ12] Yun-Hai Xiao and Zheng-Fen Jin. An alternating direction method for linear-constrained matrix nuclear norm minimization. *Numerical Linear Algebra with Applications*, 19(3):541–554, 2012.