

Structured low-rank approximation with missing data

Ivan Markovsky and Konstantin Usevich

School of Electronics and Computer Science
University of Southampton
{im,ku}@ecs.soton.ac.uk

Abstract

The approach of *SIAM J. Matrix Anal. Appl.*, 26(4):1083–1099 for solving structured total least squares problems is generalized to weighted structured low-rank approximation with missing data. The method proposed is based on elimination of the correction matrix and solution of the resulting nonlinear least squares problem by local optimization methods. The elimination step is a singular linear least-norm problem, which admits an analytic solution. Two approaches are proposed for the nonlinear least-squares minimization: minimization subject to equality constraints and unconstrained minimization with regularized cost function. The method is generalized to weighted low-rank approximation with singular weight matrix and is illustrated on matrix completion, system identification, and data-driven simulation problems. An extended version of the paper is a literate program, implementing the method and reproducing the presented results.

Keywords: low-rank approximation, structured total least squares, variable projections, missing data, system identification.

1 Introduction and notation

The paper describes a solution method for matrix structured low-rank approximation, *i.e.*, approximation of a given matrix by another matrix whose elements satisfy certain predefined relations (matrix structure) and whose rank is less than or equal to a predefined value. The combination of matrix and low-rank structure makes structured low-rank approximation a tool for data modeling. Low-rank property of a matrix is equivalent to existence of an exact low-complexity linear model for the data. Moreover, the rank of the matrix is related to the complexity of the model. The structure, imposed on the approximation, is related to properties of the model. For example, Hankel structure corresponds to time-invariance of a linear dynamical model for the data.

Structured low-rank approximation has been studied in the literature from different viewpoints: numerical algorithm for computing locally optimal or suboptimal solutions, statistical properties of the resulting estimators, and applications. The subject is closely related to the structured total least squares method. The similarities and differences between the low-rank approximation and total least squares paradigms are well documented in the literature [12, 6, 8] and will not be repeated here.

A novel feature of the low-rank approximation problem, considered in this paper, is that elements of the data matrix can be missing (not specified). Missing data may occur in practical applications due to malfunctioning of measurement device, communication channel, or storing device. In such cases, the best strategy is to collect a complete data record by repeating the data collection experiment. In other applications, however, the missing data problem is intrinsic and can not be avoided by repeated experiments. An example of such an application is the prediction of the user ratings of products (recommendation systems), where all users rate some, but rarely all, products and the task is to predict the missing ratings. Methods for solving unstructured low-rank approximation problems with missing data have been proposed in the literature [4, 5, 16, 1, 7], however, to the best of the authors knowledge none of these methods can deal with matrix structure. In this paper, we generalize the method of [13] to structured low-rank approximation with missing data.

Problem formulation

We denote missing data values by the symbol NaN (“not a number”). The considered low-rank approximation problem is

$$\begin{aligned} & \text{minimize} \quad \text{over } \hat{p} \in \mathbb{R}^{n_p} \quad \sum_{\{i \mid p_i \neq \text{NaN}\}} (p_i - \hat{p}_i)^2 \\ & \text{subject to} \quad \text{rank}(\mathcal{S}(\hat{p})) \leq r, \end{aligned} \quad (\text{SLRA})$$

where

$$\mathcal{S} : \mathbb{R}^{n_p} \rightarrow \mathbb{R}^{m \times n}, \quad \text{defined by} \quad \mathcal{S}(\hat{p}) = S_0 + \sum_{i=1}^{n_p} S_i \hat{p}_i, \quad (\mathcal{S})$$

is the matrix structure—an affine function from the structure parameter space \mathbb{R}^{n_p} to the set of matrices $\mathbb{R}^{m \times n}$. With \mathcal{G} denoting the vector of indices of the given values $\{i \mid p_i \neq \text{NaN}\}$ (in decreasing order) and $p_{\mathcal{G}}$ denoting the subvector of p with indices in \mathcal{G} , the approximation criterion can be written as

$$\sum_{i \in \mathcal{G}} (p_i - \hat{p}_i)^2 = \|p_{\mathcal{G}} - \hat{p}_{\mathcal{G}}\|_2^2.$$

Using the kernel representation of the rank constraint

$$\text{rank}(\mathcal{S}(\hat{p})) \leq r \iff \text{there is } R \in \mathbb{R}^{(m-r) \times m}, \text{ such that } R\mathcal{S}(\hat{p}) = 0 \text{ and } R \text{ has full row rank,}$$

the following equivalent problem to (SLRA) is obtained

$$\begin{aligned} & \text{minimize} \quad \text{over } \hat{p} \in \mathbb{R}^{n_p} \text{ and } R \in \mathbb{R}^{(m-r) \times m} \quad \|p_{\mathcal{G}} - \hat{p}_{\mathcal{G}}\|_2^2 \\ & \text{subject to} \quad R\mathcal{S}(\hat{p}) = 0 \quad \text{and} \quad R \text{ has full row rank.} \end{aligned} \quad (\text{SLRA}_R)$$

Problem (SLRA_R) is a double minimization over the parameters R and \hat{p}

$$\text{minimize} \quad \text{over } R \in \mathbb{R}^{(m-r) \times m} \quad M(R) \quad \text{subject to} \quad R \text{ has full row rank,} \quad (\text{SLRA}'_R)$$

where

$$M(R) := \min_{\hat{p}} \|p_{\mathcal{G}} - \hat{p}_{\mathcal{G}}\|_2^2 \quad \text{subject to} \quad R\mathcal{S}(\hat{p}) = 0. \quad (\text{INNER})$$

The evaluation of the cost function M , *i.e.*, solving (INNER) for a given value of R , is referred to as the *inner minimization problem*. This problem is solved analytically in Section 2. The remaining problem of minimizing M over R is referred to as the *outer minimization problem*. It is a nonlinear least-squares problem, which, in general, admits no analytic solution. General purpose local optimization methods are used in Section 4 for its numerical solution. In Section 3, the approach is generalized to weighted 2-norm approximation criteria with singular weight matrix. Numerical examples of solving approximation problems with missing data by the proposed methods are shown in Section 5.

Notation

In the rest of the paper, we use the following notation.

- $A_{\mathcal{I}, \mathcal{J}}$ is the submatrix of A with rows in \mathcal{I} and columns in \mathcal{J} . The row/column index can be replaced by the symbol “:”, in which case all rows/columns are selected.
- $\mathcal{M} / \mathcal{G}$ is the vector of indices of p (in decreasing order) that are missing / given.
- A^+ is the pseudo inverse of A and A^\perp is a matrix whose rows form a basis for the left null space of A .
- For given \mathcal{S} and $R \in \mathbb{R}^{(m-r) \times m}$, we define the matrix

$$G := [\text{vec}(RS_1) \quad \cdots \quad \text{vec}(RS_{n_p})] \in \mathbb{R}^{(m-r)n \times n_p}, \quad (G)$$

where $\text{vec}(\cdot)$ is the column-wise vectorization operator.

2 Analytical solution of the inner minimization problem

In this section, we consider the inner minimization problem (INNER).

Problem 1. Given linear structure \mathcal{S} , structure parameter vector $p \in \mathbb{R}^{n_p} \cup \text{NaN}$, and a kernel parameter $R \in \mathbb{R}^{(m-r) \times m}$, find the cost function $M(R)$, defined in (INNER), and a value of \hat{p} that attains the minimum.

Theorem 2. Under the following assumptions:

1. $G_{:, \mathcal{M}}$ is full column rank,
2. $1 \leq (m-r)n - n_m \leq n_p - n_m$, and
3. $\bar{G} := G_{:, \mathcal{M}}^\perp G_{:, \mathcal{G}}$ is full row rank,

Problem 1 has a unique minimum

$$M(R) = s^\top (\bar{G} \bar{G}^\top)^{-1} s, \quad \text{where } s := (\bar{G} p_{\mathcal{G}} - G_{:, \mathcal{M}}^\perp \text{vec}(RS_0)), \quad (\text{M})$$

attained by

$$\hat{p}_{\mathcal{G}} = p_{\mathcal{G}} - \bar{G}^\top (\bar{G} \bar{G}^\top)^{-1} s \quad \text{and} \quad \hat{p}_{\mathcal{M}} = -G_{:, \mathcal{M}}^+ G_{:, \mathcal{G}} \hat{p}_{\mathcal{G}}. \quad (\hat{p})$$

Proof. Defining

$$\Delta p_{\mathcal{G}} := p_{\mathcal{G}} - \hat{p}_{\mathcal{G}}$$

and using the identity

$$R \mathcal{S}(\hat{p}) = 0 \quad \Longleftrightarrow \quad G \hat{p} = -\text{vec}(RS_0),$$

we have

$$R \mathcal{S}(\hat{p}) = -\text{vec}(RS_0) \quad \Longleftrightarrow \quad \begin{bmatrix} G_{:, \mathcal{G}} & G_{:, \mathcal{M}} \end{bmatrix} \begin{bmatrix} p_{\mathcal{G}} - \Delta p_{\mathcal{G}} \\ \hat{p}_{\mathcal{M}} \end{bmatrix} = -\text{vec}(RS_0).$$

Therefore, (INNER) is equivalent to

$$M(R) := \min_{\Delta p_{\mathcal{G}} \in \mathbb{R}^{n_p - n_m}, \hat{p}_{\mathcal{M}} \in \mathbb{R}^{n_m}} \|\Delta p_{\mathcal{G}}\|_2^2 \quad \text{subject to} \quad \begin{bmatrix} G_{:, \mathcal{G}} & G_{:, \mathcal{M}} \end{bmatrix} \begin{bmatrix} \Delta p_{\mathcal{G}} \\ -\hat{p}_{\mathcal{M}} \end{bmatrix} = G_{:, \mathcal{G}} p_{\mathcal{G}} - \text{vec}(RS_0),$$

which is a generalized linear least norm problem. The solution follows from Lemma 3. \square

Generalized least norm problem

Lemma 3. Consider the generalized linear least norm problem

$$f = \min_{x, y} \|x\|_2^2 \quad \text{subject to} \quad Ax + By = c, \quad (\text{GLN})$$

with $A \in \mathbb{R}^{m \times n_x}$, $B \in \mathbb{R}^{m \times n_y}$, and $c \in \mathbb{R}^m$. Under the following assumptions:

1. B is full column rank,
2. $1 \leq m - n_y \leq n_x$, and
3. $\bar{A} := B^\perp A$ is full row rank,

problem (GLN) has a unique solution

$$\begin{aligned} f &= c^\top (B^\perp)^\top (\bar{A} \bar{A}^\top)^{-1} B^\perp c, \\ x &= A^\top (B^\perp)^\top (\bar{A} \bar{A}^\top)^{-1} B^\perp c \quad \text{and} \quad y = B^+(c - Ax). \end{aligned} \quad (\text{SOL})$$

Proof. Under assumption 1, B has a nontrivial left kernel of dimension $m - n_y$. Therefore for the nonsingular matrix $T = \begin{bmatrix} B^+ \\ B^\perp \end{bmatrix} \in \mathbb{R}^{m \times m}$

$$TB = \begin{bmatrix} B^+ \\ B^\perp \end{bmatrix} B = \begin{bmatrix} T^+ B \\ T^\perp B \end{bmatrix} = \begin{bmatrix} I_{n_y} \\ 0 \end{bmatrix}.$$

Pre-multiplying both sides of the constraint of (GLN) by T , we have the following equivalent constraint

$$\begin{bmatrix} B^+ Ax \\ B^\perp Ax \end{bmatrix} + \begin{bmatrix} y \\ 0 \end{bmatrix} = \begin{bmatrix} B^+ c \\ B^\perp c \end{bmatrix}.$$

The first equation

$$y = B^+(c - Ax)$$

uniquely determines y , given x . The second equation

$$B^\perp Ax = B^\perp c \quad (*)$$

defines a linear constraint for x only. By assumption 2, it is an underdetermined system of linear equations. Therefore, (GLN) is equivalent to the following standard least norm problem

$$f = \min_x \|x\|_2^2 \quad \text{subject to} \quad B^\perp Ax = B^\perp c. \quad (\text{GLN}')$$

By assumption 3 the solution is unique and is given by (SOL). \square

Note 4 (About assumptions 1–3). Assumption 1 is a necessary condition for uniqueness of the solution. Relaxing assumptions 1 implies that any vector in the affine space

$$\mathcal{Y} = B^+(c - Ax) + \text{null}(B)$$

is a solution to (GLN). Assumption 2 ensures that the problem is a least norm problem and has a nontrivial solution. In the case $m = n_y$, the problem has a trivial solution $f = 0$. In the case $m - n_y > n_x$, the problem generically has no solution because the constraint (*) is an overdetermined system of equations. Assumption 3 is also required for uniqueness of the solution. It can also be relaxed, making y nonunique.

Note 5 (Link to weighted least norm problems with singular weight matrix). Consider the weighted least norm problem

$$\min_z z^\top W z \quad \text{subject to} \quad Dz = c,$$

with singular positive semidefinite weight matrix W . Using a change of variables $\bar{z} = T^{-1}z$, where T is a nonsingular matrix, we obtain the equivalent problem

$$\min_{\bar{z}} \bar{z}^\top T^\top W T \bar{z} \quad \text{subject to} \quad D T \bar{z} = c.$$

There exists a nonsingular matrix T , such that

$$T^\top W T = \begin{bmatrix} I_{n_x} & \\ & 0 \end{bmatrix}.$$

Partitioning \bar{z} and $\bar{D} := D T^{-1}$ conformably as

$$\bar{z} = \begin{bmatrix} x \\ y \end{bmatrix} \quad \text{and} \quad \bar{D} = \begin{bmatrix} A & B \end{bmatrix}$$

we obtained problem (GLN).

3 Weighted approximation

Problem (SLRA) is generalized in this section to the weighted structured low-rank approximation problem

$$\begin{aligned} & \text{minimize} \quad \text{over } \hat{p} \in \mathbb{R}^{n_p} \quad (p_{\mathcal{G}} - \hat{p}_{\mathcal{G}})^\top W_g (p_{\mathcal{G}} - \hat{p}_{\mathcal{G}}) \\ & \text{subject to} \quad \text{rank}(\mathcal{S}(\hat{p})) \leq r, \end{aligned} \quad (\text{WSLRA})$$

where W_g is a positive definite matrix. The change of variables

$$p'_{\mathcal{G}} = \sqrt{W_g} p_{\mathcal{G}} \quad \text{and} \quad \hat{p}'_{\mathcal{G}} = \sqrt{W_g} \hat{p}_{\mathcal{G}} \quad (p \mapsto p')$$

reduces Problem (WSLRA) to an equivalent unweighted problem (SLRA). We have

$$\mathcal{S}(\hat{p}) = S_0 + \text{vec}^{-1}(\mathbf{S}\hat{p}), \quad \text{where} \quad \mathbf{S} := [\text{vec}(S_1) \quad \cdots \quad \text{vec}(S_{n_p})] \in \mathbb{R}^{mn \times n_p}. \quad (\mathbf{S})$$

The structure \mathcal{S}' of the equivalent problem is defined by the matrices S_0 and $\mathbf{S}' = [\text{vec}(S'_1) \quad \cdots \quad \text{vec}(S'_{n_p})]$, where

$$\mathbf{S}'_{:, \mathcal{G}} = \mathbf{S}_{:, \mathcal{G}} \sqrt{W_g^{-1}} \quad \text{and} \quad \mathbf{S}'_{:, \mathcal{M}} = \mathbf{S}_{:, \mathcal{M}}. \quad (\mathcal{S} \mapsto \mathcal{S}')$$

We showed that problem (WSLRA) is solved by:

1. preprocessing the data p and the structure \mathcal{S} , as in $(p \mapsto p')$ and $(\mathcal{S} \mapsto \mathcal{S}')$,
2. solving the equivalent unweighted problem with structure parameter vector p' , structure specification \mathcal{S}' , and rank specification r , and
3. postprocessing the solution \hat{p}' , obtained in step 2, in order to obtain the solution $\hat{p}_{\mathcal{G}} = \sqrt{W_g^{-1}} \hat{p}'_{\mathcal{G}}$ of the original problem.

Using the transformation $(p \mapsto p')$, $(\mathcal{S} \mapsto \mathcal{S}')$ and the solution (M) of (SLRA), we obtain the following explicit expression for the cost function of (WSLRA)

$$M(R) = (\bar{G}p_{\mathcal{G}} - G_{:, \mathcal{M}}^\perp \text{vec}(RS_0))^\top W_g^{-1} \bar{G}^\top (\bar{G}W_g^{-1} \bar{G}^\top)^{-1} \bar{G}W_g^{-1} (\bar{G}p_{\mathcal{G}} - G_{:, \mathcal{M}}^\perp \text{vec}(RS_0)), \quad (\text{M}_W)$$

where $\bar{G} = G_{:, \mathcal{M}}^\perp G_{:, \mathcal{G}}$ and G is defined in (G).

Note 6 (Weighted structured low-rank approximation with a singular weight matrix). A more general formulation of problem (WSLRA) is

$$\begin{aligned} & \text{minimize} \quad \text{over } \hat{p} \in \mathbb{R}^{n_p} \quad (p - \hat{p})^\top W (p - \hat{p}) \\ & \text{subject to} \quad \text{rank}(\mathcal{S}(\hat{p})) \leq r, \end{aligned} \quad (\text{WSLRA}')$$

with positive semidefinite weight matrix W . Problem (WSLRA') can be reduced to an equivalent unweighted low-rank approximation problem with missing data (SLRA). There exists a nonsingular matrix T , such that

$$T^\top W T = \begin{bmatrix} I_{n_g} & \\ & 0 \end{bmatrix}.$$

Defining structure \mathcal{S}' , specified by the matrices S_0 and $\mathbf{S}' = \mathbf{S}T$, and changing the variables $p' = T^{-1}p$, $\hat{p}' = T^{-1}\hat{p}$, we obtain an equivalent problem in the form (SLRA) with last $n_p - \text{rank}(W)$ values missing.

Note 7 (Solving (SLRA) as weighted unstructured problem). Consider an instance of problem (SLRA), refer to as problem P1, with structure $\mathcal{S} = \mathcal{S}_1$ and an instance of problem (WSLRA), refer to as problem P2, with unstructured correction and weight matrix

$$W = \mathbf{S}_1 \mathbf{S}_1^\top. \quad (\mathcal{S} \mapsto W)$$

It can be verified by inspection that the cost functions (M) and (M_W) of problems P1 and P2, respectively, coincide. The weight matrix $W \in \mathbb{R}^{mn \times mn}$, defined in $(\mathcal{S} \mapsto W)$, however is singular ($\text{rank}(W)$ is equal to the number of structure parameters of problem P1, which is less than mn). In the derivation of the cost function (M_W) it is assumed that W_g is positive definite, so that minimization of (M_W) is not equivalent to problem P2. Using pseudo-inverse instead of inverse in $(p \mapsto p')$ and $(\mathcal{S} \mapsto \mathcal{S}')$ and observing that $(W^+)^+ = W$, minimization of (M_W) for problem P2 yields an equivalent problem to problem P1.

4 Outer minimization problem

The outer minimization problem (SLRA_R) is a nonlinear least-squares problem, which we solve by general purpose local optimization methods. In order to apply standard optimization methods, however, we need first to replace the rank constraint with equivalent equality or inequality constraints.

The full row rank constraint on R is equivalent to and can be enforced in the parameter optimization method by the equality constraint

$$RR^\top = I_{m-r}. \quad (\text{f.r.r. } R)$$

Then the outer minimization problem becomes a constrained nonlinear least squares problems

$$\text{minimize over } R \in \mathbb{R}^{(m-r) \times m} \quad M(R) \quad \text{subject to} \quad RR^\top - I_{m-r} = 0, \quad (\text{SLRA}'_R)$$

which can be solved by general purpose constrained optimization methods [14]. Another approach of solving the outer minimization problem is to reformulate it as a regularized unconstrained nonlinear least squares problem by adding the regularization term $\gamma \|RR^\top - I_{m-r}\|_F^2$ to the cost function, i.e.,

$$\text{minimize over } R \in \mathbb{R}^{(m-r) \times m} \quad M(R) + \gamma \|RR^\top - I_{m-r}\|_F^2. \quad (\text{SLRA}''_R)$$

The parameter γ should be chosen “large enough” in order to enforce the constraint (f.r.r. R). A corollary of the following theorem shows that $\gamma = \|p_{\mathcal{G}}\|_2^2$ is sufficiently large.

Theorem 8. *Let $M : \mathbb{R}^{(m-r) \times m} \rightarrow \mathbb{R}_+$ be a homogeneous function, i.e., $M(R) = M(TR)$, for any R and a nonsingular $m \times m$ matrix T . The optimal solutions of problem (SLRA''_R) with $\gamma = \max_R M(R)$ coincide with the optimal solutions of (SLRA'_R).*

Proof. Let R be an solution to (SLRA''_R). We will show that

$$\|RR^\top - I_{m-r}\|_F^2 = m - r - \text{rank}(R). \quad (*)$$

There exists an orthogonal matrix U diagonalizing RR^\top . We have

$$\begin{aligned} \|RR^\top - I_{m-r}\|_F^2 &= \|URR^\top U^\top - I_{m-r}\|_F^2 \\ &= \|\text{diag}(a_1, \dots, a_{\text{rank}(R)}, 0, \dots, 0) - I_{m-r}\|_F^2, \quad \text{where } a_i > 0 \\ &= \sum_{i=1}^{\text{rank}(R)} (a_i - 1)^2 + m - r - \text{rank}(R). \end{aligned}$$

Suppose that $a_i \neq 1$ for some i . The matrix

$$R' = \text{diag}(1, \dots, 1, 1/\sqrt{a_i}, 1, \dots, 1)R$$

has the same kernel and rank as R , so that by the homogeneity property of M , $M(R) = M(R')$. However, we have

$$\|RR^\top - I_{m-r}\|_F^2 > \|R'R'^\top - I_{m-r}\|_F^2,$$

so that R' achieves smaller value of the cost function of (SLRA''_R) than R . This is a contradiction. Therefore, $a_i = 1$ for all i . This concludes the proof of (*).

So far we showed that the cost function of (SLRA''_R) is

$$M(R) + \gamma(m - r - \text{rank}(R)). \quad (M'')$$

Denote by M_r^* the optimal value of (SLRA'_R) (the index in the subscript is the upper bound for the rank) and note that the optimal value of (SLRA''_R) is equal to M_r^* provided that the solution R of (SLRA''_R) is full row rank. Therefore, in order to prove the theorem it is sufficient to show that the cost function of (SLRA''_R) achieves its minimum for a full rank R , i.e.,

$$M_r^* < M_{r-i}^* + \gamma i \quad \Longleftrightarrow \quad \gamma > \frac{1}{i}(M_r^* - M_{r-i}^*), \quad \text{for } i = 1, 2, \dots, m - r. \quad (**)$$

Since $M_r^* > 0$, $\gamma = \max_R M(R)$ is a sufficient condition for (**). \square

Note 9 (Initial approximation). Solving the outer minimization problem by either constrained or regularized local minimization requires an initial approximation for the parameter R , *i.e.*, a suboptimal solution of the structured low-rank approximation problem. Such a solution can be computed from a heuristic that ignores the data matrix structure \mathcal{S} and fills in zeros for the missing values. The resulting unstructured low-rank approximation problem can then be solved analytically in terms of the singular value decomposition.

Note 10 (Efficient computation and software implementation). Efficient evaluation of the cost function and its derivatives in the special case of mosaic-Hankel matrix structure is presented in a companion paper [17]. The method, presented in this paper (general linear structure) and the efficient methods of [17] are implemented in Matlab (using Optimization Toolbox) and C++, respectively. Description of the software and overview of its applications is given in [10].

5 Applications

As an illustration of how the developed methods can be used in practice and as a verification of their effectiveness, we present in this section three sample applications:

- unstructured noisy matrix completion,
- scalar autonomous system identification with missing data, and
- data-driven simulation.

Numerical examples comparing the methods developed in the paper with alternative methods, specifically developed for these applications, are shown. All simulations are done in Matlab and are reproducible in the sense of [2]. An extended version [11] of this paper is a literate program (in noweb format [15]), implementing the methods in the paper and generating the presented numerical results. The necessary m-files can be downloaded from

<http://eprints.soton.ac.uk/340718>.

5.1 Unstructured matrix with missing data

In the case of unstructured data matrix, the results obtained by the methods in the paper are compared with the results of alternative methods

- the alternating projections method of [7] and
- the singular value thresholding method of [3].

The alternating projections method for weighted low-rank approximation uses an image representation PL , where P is $m \times r$ and L is $r \times n$, of the $m \times n$ rank- r matrix $\mathcal{S}(\hat{p})$. The algorithm iteratively minimizes the cost function over P with fixed L from the previous iteration step and over L with fixed P to its previously computed values. Both problems—minimization over P and minimization over L —are weighted linear least-squares problems, so that they can be solved globally and reliably. The cost function value is monotonically non-increasing over the iterations of the alternating projections method. The method is adapted in [7] to the case of missing data and is effective in solving large scale noisy matrix completion problems. A Matlab implementation is available from <http://eprints.ecs.soton.ac.uk/18296/>.

Singular value thresholding is a method for low-rank matrix completion, *i.e.*, a low-rank approximation with missing and exact values only. Although singular value thresholding is initially designed for the exact data case, it is shown to handle noisy data as well. Therefore, it solve low-rank approximation problems with missing data. The method is based on convex relaxation of the rank constraint and does not require an initial approximation. A Matlab implementation is available from <http://svt.caltech.edu/>

The results of a numerical example with a data matrix with the following pattern of missing values

$$\begin{bmatrix} \text{NaN} & \times & \times & \text{NaN} & \times & \times & \text{NaN} & \times & \times & \times \\ \times & \text{NaN} & \times & \times & \text{NaN} & \times & \times & \text{NaN} & \times & \times \\ \times & \times & \text{NaN} & \times & \times & \text{NaN} & \times & \times & \text{NaN} & \times \end{bmatrix}$$

and rank one specification are shown in Tables 1 and 2. Tables 1 shows the approximation errors $M_r^* = \|p_{\mathcal{M}} - \hat{p}_{\mathcal{M}}\|$ achieved by the algorithms upon convergence and Tables 2 shows the number of iterations performed by the algorithms.

# of missing values	1	2	3	4	5	6	7	8	9
(SLRA' _R)	0.8859	0.8642	0.8598	0.7900	0.7570	0.7568	0.6659	0.6031	0.6022
(SLRA'' _R)	0.8859	0.8639	0.8593	0.7900	0.7570	0.7568	0.6658	0.6031	0.6022
alternating projections	0.8859	0.8639	0.8593	0.7900	0.7570	0.7568	0.6658	0.6031	0.6022
singular value thresholding	0.8892	0.8668	1.0539	1.0393	1.0467	1.0468	1.0310	0.9890	3.8635

Table 1: Approximation error $M_1^* \times 10^{-3}$ for the compared methods on problems with 1, ..., 9 missing values.

# of missing values	1	2	3	4	5	6	7	8	9
(SLRA' _R)	4	3	4	3	4	4	4	3	10
(SLRA'' _R)	3	4	5	3	5	5	5	6	14
alternating projections	3	3	4	4	4	4	4	4	7

Table 2: Number of iterations for convergence of the compared methods on the problems in Table 1.

The results show that the proposed methods based on quadratic equality constraint (SLRA'_R) and regularization (SLRA''_R) achieve the same approximation error as the alternating projections method and require similar number of iterations (starting from the same suboptimal initial approximation, see Note 9). The approximation error achieved by the singular value thresholding method increases with the increase of the number of missing values.

5.2 System identification with missing data

Identification of a linear time-invariant system from a noisy trajectory of the system is a mosaic-Hankel structured low-rank approximation problem [17]. In the simplest case of a scalar autonomous system, the structure is Hankel

$$\mathcal{H}_{m,n}(p) := \begin{bmatrix} p_1 & p_2 & p_3 & \cdots & p_n \\ p_2 & p_3 & \ddots & & p_{n+1} \\ p_3 & \ddots & & & \vdots \\ \vdots & & & & \\ p_m & p_{m+1} & \cdots & & p_{m+n-1} \end{bmatrix} \in \mathbb{R}^{m \times n}$$

and the identification problem is

$$\begin{aligned} & \text{minimize} && \text{over } \hat{y} \in \mathbb{R}^T && \|\mathbf{y}_d - \hat{\mathbf{y}}\|_2^2 \\ & \text{subject to} && \text{rank}(\mathcal{H}_{\ell+1, T-\ell}(\hat{\mathbf{y}})) \leq \ell, \end{aligned}$$

where ℓ is the system's lag (assumed known) and

$$\mathbf{y}_d = (y_d(1), \dots, y_d(T)) \in \mathbb{R}^T$$

is the given trajectory. As an application of the algorithms developed in the paper, we consider autonomous system identification when samples of y_d are missing at arbitrary locations.

Figure 1 shows the results of a simulation example with the system defined by the difference equation

$$y(t) = 1.456y(t-1) - 0.81y(t-2).$$

The data is a trajectory of the system perturbed by additive noise with $T = 50$ samples. (For details on the simulation setup, see [11]) The missing values are distributed periodically with a period of 3 samples, so that standard system identification methods are not applicable. In this experiment the default initial approximation results in convergence to a poor local minimum and is replaced by $[0.8 \quad -1.5 \quad 1]$.




Figure 1: System identification with periodically missing data (crosses on the t -axis). Noisy samples (circles), optimal approximation (dashed blue), and true trajectory (solid red).

5.3 Data-driven simulation

Data-driven simulation problems [9] are special cases of missing data low-rank approximation. The to-be-simulated system is assumed to be linear time-invariant with a known upper bound ℓ of the lag. The system is implicitly specified by a trajectory $w_d^{(1)} = (u_d^{(2)}, y_d^{(2)}) \in (\mathbb{R}^q)^{T_1}$. The to-be-simulated trajectory $w_d^{(2)} = (u_d^{(2)}, y_d^{(2)}) \in (\mathbb{R}^q)^{T_2}$ is specified by the initial conditions $w_{\text{ini}} = (w_d^{(2)}(1), \dots, w_d^{(2)}(\ell))$ and the input $(u_d^{(2)}(\ell+1), \dots, u_d^{(2)}(T_2))$. The data-driven simulation problem is a mosaic-Hankel structured low-rank approximation problem

$$\begin{aligned} & \text{minimize} \quad \text{over } \hat{w} \quad \|w_d - \hat{w}\|_2^2 \\ & \text{subject to} \quad \text{rank} \left(\begin{bmatrix} \mathcal{H}_{\ell+1}(u_d^{(1)}) & \mathcal{H}_{\ell+1}(u_d^{(2)}) \\ \mathcal{H}_{\ell+1}(y_d^{(1)}) & \mathcal{H}_{\ell+1}(y_d^{(2)}) \end{bmatrix} \right) \leq 2\ell + 1, \end{aligned}$$

with missing data being the to-be-simulated response $(y_d^{(2)}(\ell+1), \dots, y_d^{(2)}(T_2))$.

Consider a simulation example with the second order single-input single-output system, defined by the difference equation

$$y(t) = 1.456y(t-1) - 0.81y(t-2) + u(t) - u(t-1).$$

The data $w_d^{(1)}$ is a trajectory of the system generated from random input and additive noise. The to-be-simulated trajectory $w_d^{(2)}$ is the impulse response \bar{h} of the system, *i.e.*, the response under zero initial conditions and pulse input:

$$u_d^{(2)} = (\underbrace{0, \dots, 0}_{\ell}, \underbrace{1, 0, \dots, 0}_{\text{pulse input}}), \quad \text{and} \quad y_d^{(2)} = (\underbrace{0, \dots, 0}_{\ell}, \underbrace{\hat{h}(0), \hat{h}(1), \dots, \hat{h}(T_2 - \ell - 1)}_{\text{impulse response}}).$$

Figure 2 shows the true and estimated by the proposed low-rank approximation method impulse responses.

6 Conclusions

A variable-projection-like approach for structured low-rank approximation with missing data was developed. The approach was furthermore generalized to weighted structured low-rank approximation with singular weight matrix. Two optimization strategies were proposed for the nonlinear least-squares optimization: optimization subject to quadratic equality constraints and regularized unconstrained optimization. The problem and solution methods developed have applications in matrix completion (unstructured problems), system identification with missing data, and data-driven simulation and control (mosaic-Hankel structured problems). The performance of the methods was illustrated on simulation examples from these applications and was compared with the performance of problem specific methods.




Figure 2: Data-driven simulation of impulse response: true (red solid line), optimal approximation (dashed blue).

Acknowledgements

The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement number 258581 “Structured low-rank approximation: Theory, algorithms, and applications”.

References

- [1] M. Brand. Fast low-rank modifications of the thin singular value decomposition. *Lin. Algebra and its Appl.*, 415(1):20–30, 2006.
- [2] J. Buckheit and D. Donoho. *Wavelets and statistics*, chapter "Wavelab and reproducible research". Springer-Verlag, Berlin, New York, 1995.
- [3] J.-F. Cai, E. Candés, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, January 2010.
- [4] K. Gabriel and S. Zamir. Lower rank approximation of matrices by least squares with any choice of weights. *Technometrics*, 21:489–498, 1979.
- [5] B. Grung and R. Manne. Missing values in principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 42:125–139, 1998.
- [6] I. Markovsky. Structured low-rank approximation and its applications. *Automatica*, 44(4):891–909, 2008.
- [7] I. Markovsky. *Algorithms and iterate programs for weighted low-rank approximation with missing data*, volume 3 of *Springer Proc. Mathematics*, pages 255–273. Springer, 2011.
- [8] I. Markovsky. *Low Rank Approximation: Algorithms, Implementation, Applications*. Springer, 2012.
- [9] I. Markovsky and P. Rapisarda. Data-driven simulation and control. *Int. J. Control*, 81(12):1946–1959, 2008.
- [10] I. Markovsky and K. Usevich. Software for weighted structured low-rank approximation. Technical Report 339974, Univ. of Southampton, <http://eprints.soton.ac.uk/339974>, 2012.
- [11] I. Markovsky and K. Usevich. Structured low-rank approximation with missing values. Technical Report 340718, Univ. of Southampton, <http://eprints.soton.ac.uk/340718>, 2012.
- [12] I. Markovsky and S. Van Huffel. Overview of total least squares methods. *Signal Proc.*, 87:2283–2302, 2007.
- [13] I. Markovsky, S. Van Huffel, and R. Pintelon. Block-Toeplitz/Hankel structured total least squares. *SIAM J. Matrix Anal. Appl.*, 26(4):1083–1099, 2005.
- [14] J. Nocedal and S. Wright. *Numerical optimization*. Springer-Verlag, 1999.
- [15] N. Ramsey. Iterate programming simplified. *IEEE Software*, 11:97–105, 1994.
- [16] N. Srebro. *Learning with matrix factorizations*. PhD thesis, MIT, 2004.
- [17] K. Usevich and I. Markovsky. Efficient algorithms for mosaic-Hankel weighted low-rank approximation. Technical report, Univ. of Southampton, <http://eprints.soton.ac.uk/>, 2012.