Left vs right representations for solving weighted low-rank approximation problems

Ivan Markovsky [a], Sabine Van Huffel [b],*

[a] *School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, United Kingdom*
[b] *Katholieke Universiteit Leuven, Departement Elektrotechniek, ESAT-SCD, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium*

Abstract

The weighted low-rank approximation problem in general has no analytical solution in terms of the singular value decomposition and is solved numerically using optimization methods. Four representations of the rank constraint that turn the abstract problem formulation into parameter optimization problems are presented. The parameter optimization problem is partially solved analytically, which results in an equivalent quadratically constrained problem. A commonly used re-parameterization avoids the quadratic constraint and makes the equivalent problem a nonlinear least squares problem, however, it might be necessary to change this re-parameterization during the iteration process. It is shown how the cost function can be computed efficiently in two special cases: row-wise and column-wise weighting.
© 2006 Elsevier Inc. All rights reserved.

1. Introduction

We consider the following matrix approximation problem.

---

* Corresponding author. Tel.: +32 16 32 17 03; fax: +32 16 32 19 70.
 E-mail addresses: ivan.markovsky@esat.kuleuven.be (I. Markovsky), sabine.vanhuffel@esat.kuleuven.be (S. Van Huffel).

Problem 1 (*Weighted low-rank approximation* (*WLRA*)). Given a matrix $D \in \mathbb{R}^{m \times n}$, a positive definite matrix $W \in \mathbb{R}^{mn \times mn}$, and an integer $r < \min(m, n)$, find an optimal W-weighted, rank-r approximation of D, defined as

$$\widehat{D}^* = \arg \min_{\widehat{D}} \mathrm{vec}^\top (D - \widehat{D}) W \mathrm{vec}(D - \widehat{D}) \quad \text{subject to } \mathrm{rank}(\widehat{D}) \leqslant r. \tag{WLRA}$$

For $W = I$, the cost function of (WLRA) is $\|D - \widehat{D}\|_F^2$ and the problem has an analytical solution in terms of the singular value decomposition (SVD) of D.

Theorem 1 (Eckart–Young–Mirsky [2]). *Let $D = U \Sigma V^\top$ be the SVD of D and partition U, $\Sigma =: \mathrm{diag}(\sigma_1, \dots, \sigma_q)$, $q := \min(m, n)$, and V as follows:*

$$U =: \begin{bmatrix} \overset{r}{U_1} & \overset{q-r}{U_2} \end{bmatrix} m, \quad \Sigma =: \begin{bmatrix} \overset{r}{\Sigma_1} & \overset{q-r}{0} \\ 0 & \Sigma_2 \end{bmatrix} \begin{matrix} r \\ q-r \end{matrix} \quad and \quad V =: \begin{bmatrix} \overset{r}{V_1} & \overset{q-r}{V_2} \end{bmatrix} n,$$

Then the rank-r matrix

$$\widehat{D}^* = U_1 \Sigma_1 V_1^\top$$

is such that

$$\|D - \widehat{D}^*\|_F^2 = \min_{\mathrm{rank}(\widehat{D}) \leqslant r} \|D - \widehat{D}\|_F^2 = \sigma_{r+1}^2 + \cdots + \sigma_q^2.$$

The solution $\widehat{D}^*$ is unique if and only if $\sigma_{r+1} \neq \sigma_r$.

Surprisingly, however, a similar analytical solution is not known for a general positive definite weight matrix W. Presently the most general WLRA problems that can be solved *exactly* using SVDs are those with a weight matrix of the form $W = W_r \otimes W_\ell$, where $W_r \in \mathbb{R}^{n \times n}$ and $W_\ell \in \mathbb{R}^{m \times m}$ are positive definite matrices and $\otimes$ is the Kronecker product. The solution procedure is based on the following theorem.

Theorem 2 (Two-sided WLRA). *Define the modified data matrix*

$$D_m := \sqrt{W_\ell} D \sqrt{W_r},$$

and let $\widehat{D}_m^*$ be the optimal (unweighted) low-rank approximation of D_m. Then

$$\widehat{D}^* := \left(\sqrt{W_\ell} \right)^{-1} \widehat{D}_m^* \left(\sqrt{W_r} \right)^{-1},$$

is a solution of the following two-sided WLRA problem

$$\min_{\widehat{D}} \| \sqrt{W_\ell} (D - \widehat{D}) \sqrt{W_r} \|_F^2 \quad \text{subject to } \mathrm{rank}(\widehat{D}) \leqslant r.$$

A solution always exists. It is unique if and only if $\widehat{D}_m^*$ is unique.

The link between two-sided WLRA and (WLRA) is

$$\left\| \sqrt{W_\ell} (D - \widehat{D}) \sqrt{W_r} \right\|_F^2 = \left\| \mathrm{vec} \left(\sqrt{W_\ell} (D - \widehat{D}) \sqrt{W_r} \right) \right\|^2$$
$$= \left\| \left(\sqrt{W_r} \otimes \sqrt{W_\ell} \right) \mathrm{vec}(D - \widehat{D}) \right\|^2$$
$$= \mathrm{vec}^\top (D - \widehat{D}) \left(W_r \otimes W_\ell \right) \mathrm{vec}(D - \widehat{D}).$$

Here we used the identities

$$\text{vec}(AXB) = (B^\top \otimes A)\text{vec}(X) \quad \text{and} \quad (A_1 \otimes B_1)(A_2 \otimes B_2) = (A_1 A_2) \otimes (B_1 B_2).$$

Of course, $W \in \mathbb{R}^{mn \times mn}$, $m, n > 1$, being representable as $W_r \otimes W_\ell$ with $W_r \in \mathbb{R}^{n \times n}$ and $W_\ell \in \mathbb{R}^{m \times m}$ is a rather restrictive condition, i.e., representing a given positive definite weight matrix W in the form $W_r \otimes W_\ell$, in general, requires approximation. It is shown in [10] that the optimal approximation (in Frobenius norm sense), i.e., the solution of the problem

$$\min_{W_r \in \mathbb{R}^{n \times n}, W_\ell \in \mathbb{R}^{m \times m}} \|W - W_r \otimes W_\ell\|_F \quad \text{subject to } W_r \text{ and } W_\ell \text{ are positive definite},$$

can be computed in analytical form in terms of the SVD of a matrix constructed from W. This result can be used for the computation of a suboptimal solution of the general WLRA problem as follows:

1. approximate the weight matrix W by the Kronecker product $W_r \otimes W_\ell$ with some positive definite matrices $W_r \in \mathbb{R}^{n \times n}$ and $W_\ell \in \mathbb{R}^{m \times m}$, using the result of [10], and
2. find an optimal WLRA of the data matrix D with the weight matrix $W_r \otimes W_\ell$, using Theorem 2.

The computed WLRA with the weight matrix $W_r \otimes W_\ell$ serves as a suboptimal solution for the desired WLRA with the weight matrix W. It is intuitively clear (and there is empirical evidence) that the better W is approximated, the closer the suboptimal solution to the optimal one is.

The suboptimal solution can be improved by numerical optimization methods. Since the problem is nonconvex there are no efficient optimization methods that are guaranteed to find a globally optimal solution. What is aimed at instead is a locally optimal solution nearby the given initial suboptimal approximation. Optimization methods for solving Problem 1 have been considered in the literature under different names:

- criss-cross multiple regression [3],
- Riemannian singular value decomposition [1],
- maximum likelihood principal component analysis [13],
- weighted low-rank approximation [4], and
- weighted total least squares [7,5].

Gabriel and Zamir [3] consider an element-wise weighted low-rank approximation problem, i.e., problem (WLRA) with diagonal weight matrix W, and propose an iterative solution method. Their method, however, does not necessarily converge to a minimum point, see the discussion in [3, Section 6, p. 491].

The Riemannian singular value decomposition framework of De Moor [1] includes the WLRA problem with rank specification $r = \min(m, n) - 1$ and a diagonal weight matrix W as a special case. In [1], an algorithm resembling the inverse power iteration algorithm is proposed. The method, however, has no proven convergence properties.

The maximum likelihood principal component analysis method of Wentzell et al. [13] is developed for applications in chemometrics, see also [9]. This method is an alternating least squares algorithm. It applies to the general WLRA problems and is globally convergent. The convergence rate, however, is linear and the method could be rather slow when the $r + 1$st and the rth singular values of the data matrix D are close to each other. In the unweighted case this situation corresponds to lack of uniqueness of the solution, cf., Theorem 1.

Manton et al. [4] treat the problem as an optimization over a Grassman manifold and propose steepest descent and Newton type algorithms. The least squares nature of the problem is not exploited in this work and the proposed algorithms are not globally convergent.

The method of [7,5] is a heuristic for solving the first order optimality condition of (WLRA). A solution of a nonlinear equation is sought instead of a minimum point of the original optimization problem. The method is locally convergent with super linear convergence rate. The method is not globally convergent and the region of convergence around a minimum point could be rather small in practice.

In [6, Chapter 3] it is shown that the above mentioned algorithms differ in

1. the way the rank constraint is represented and
2. the local optimization method applied to the resulting (from the particular representation used) parameter optimization problem.

Thus combining different rank representations with different local optimization methods, we obtain different solution methods for the WLRA problem.

An important remaining issue is the efficiency of the algorithm in special cases of the general WLRA approximation problem. Useful special cases are:

- *element-wise weighting*: $W = \mathrm{diag}(w_1, \ldots, w_{mn})$, where $w_i > 0$, for $i = 1, \ldots, mn$,
- *column-wise weighting*: $W = \mathrm{diag}(W_1, \ldots, W_n)$, where $W_i \in \mathbb{R}^{m \times m}$, $W_i > 0$, for $i = 1, \ldots, n$, and
- *row-wise weighting*: $T^\top W T = \mathrm{diag}(W_1, \ldots, W_m)$, where $W_i \in \mathbb{R}^{n \times n}$, $W_i > 0$, for $i = 1, \ldots, m$, and T is an $mn \times mn$ permutation matrix, such that $\mathrm{vec}(D) = T \mathrm{vec}(D^\top)$.

In this paper, we generalize the treatment of [6, Chapter 3], where left image and kernel representations and column-wise weighting, i.e., weight matrices $W = \mathrm{diag}(W_1, \ldots, W_n)$, are considered. Our renewed interest in the problem was triggered from the difficulty that if $m > n$ the left image and kernel representations lead to larger number of parameters. It is advantageous to use in such cases right image and kernel representations. Note that the problem can *not* be resolved by transposing the data matrix because then W is in general no longer block diagonal and this structure is exploited in the computations.

We start from a full weight matrix W and then specialize the results to column and row-wise weighting. Section 2 presents four representations of the rank constraint in a parametric form. For these representations equivalent optimization problems are derived in Section 3 by eliminating the variable $\widehat{D}$. Section 4 presents unique parameterizations of the rank constraint, which simplifies the equivalent optimization problems. The resulting problems are nonlinear least squares problems and are solved by standard optimization methods, such as the Levenberg–Marquardt method.

2. Low-rank representations

In order to solve (either analytically or numerically) the WLRA problem, we represent the rank constraint $\mathrm{rank}(\widehat{D}) \leqslant r$ in a parametric form. Four possibilities are what are called image and kernel representations with their left and right versions. The left image representation

$$\text{rank}(\widehat{D}) \leqslant r \Leftrightarrow \text{there is } P \in \mathbb{R}^{m \times r}, \quad P^\top P = I_r \text{ and } L \in \mathbb{R}^{r \times n}, \text{ such that } \widehat{D} = PL$$
$$(\ell\text{IM})$$

imposes an upper bound on the dimension of the column space of $\widehat{D}$. This representation has as a parameter the matrix P. The right image representation

$$\text{rank}(\widehat{D}) \leqslant r \Leftrightarrow \text{there is } P \in \mathbb{R}^{m \times r} \text{ and } L \in \mathbb{R}^{r \times n}, \quad LL^\top = I_r, \text{ such that } \widehat{D} = PL$$
$$(\text{rIM})$$

has as a parameter the matrix L and imposes an upper bound on the dimension of the row space of $\widehat{D}$.

The left kernel representation

$$\text{rank}(\widehat{D}) \leqslant r \Leftrightarrow \text{there is } R_\ell \in \mathbb{R}^{(m-r) \times m}, \quad R_\ell R_\ell^\top = I_{m-r}, \text{ such that } R_\ell \widehat{D} = 0$$
$$(\ell\text{KER})$$

has as a parameter the matrix R_ℓ and imposes a lower bound on the dimension of the left kernel of $\widehat{D}$. (The subscript "ℓ" in R_ℓ stands for "left".) The alternative (right) kernel representation

$$\text{rank}(\widehat{D}) \leqslant r \Leftrightarrow \text{there is} R_\text{r} \in \mathbb{R}^{n \times (n-r)}, \quad R_\text{r}^\top R_\text{r} = I_{n-r}, \text{ such that } \widehat{D} R_\text{r} = 0 \qquad (\text{rKER})$$

has as a parameter the matrix R_r and imposes a lower bound on the dimension of the right kernel of $\widehat{D}$.

The four representations (ℓIM), (rIM), (ℓKER), and (rKER), however, lead to different number of parameters to be optimized over in solving (WLRA). Fewer optimization variables result in less computations per iteration as well as typically fewer iteration steps. Thus in the numerical solution of the WLRA problem it is important to use a representation that leads to as few as possible optimization variables.

Looking at the dimensions of the parameters P, L, R_ℓ, and R_r, we see that

- (ℓKER) has the fewest free parameters if $m < n$ and $r < m/2$,
- (ℓIM) has the fewest free parameters if $m < n$ and $r > m/2$,
- (rKER) has the fewest free parameters if $m > n$ and $r < n/2$, and
- (rIM) has the fewest free parameters if $m > n$ and $r > n/2$.

Due to the constraints $P^\top P = I$, $LL^\top = I$, $R_\ell R_\ell^\top = I$, and $R_\text{r} R_\text{r}^\top = I$, however, not all elements of the parameters are free. Equivalently the parameters P, L, R_ℓ, and R_r are not unique. In Section 4, we derive so-called left and right input/output parameterizations of which the parameters $X_\ell \in \mathbb{R}^{r \times (m-r)}$ and $X_\text{r} \in \mathbb{R}^{r \times (n-r)}$ are unconstrained and unique.

3. Equivalent optimization problems

In what follows we use the notation

$$d := \text{vec}(D) \quad \text{and} \quad \hat{d} := \text{vec}(\widehat{D}).$$

Using the left image representation, (WLRA) yields the parameter optimization problem:

$$\min_{\substack{P \\ P^\top P = I}} \left(\underbrace{\min_{\hat{d},L}(d-\hat{d})^\top W(d-\hat{d}) \quad \text{subject to} \quad \widehat{D}=PL}_{f_{\ell\mathrm{IM}}(P)} \right). \tag{WLRA$_P$}$$

Using the right image representation, (WLRA) yields the parameter optimization problem:

$$\min_{\substack{L \\ LL^\top = I}} \left(\underbrace{\min_{\hat{d},P}(d-\hat{d})^\top W(d-\hat{d}) \quad \text{subject to} \quad \widehat{D}=PL}_{f_{\mathrm{rIM}}(L)} \right). \tag{WLRA$_L$}$$

Using the left kernel representation, (WLRA) yields the parameter optimization problem:

$$\min_{\substack{R_\ell \\ R_\ell R_\ell^\top = I}} \left(\underbrace{\min_{\hat{d}}(d-\hat{d})^\top W(d-\hat{d}) \quad \text{subject to} \quad R_\ell \widehat{D}=0}_{f_{\ell\mathrm{KER}}(R_\ell)} \right). \tag{WLRA$_{R_\ell}$}$$

Using the right kernel representation, (WLRA) yields the parameter optimization problem:

$$\min_{\substack{R_\mathrm{r} \\ R_\mathrm{r}^\top R_\mathrm{r} = I}} \left(\underbrace{\min_{\hat{d}}(d-\hat{d})^\top W(d-\hat{d}) \quad \text{subject to} \quad \widehat{D}R_\mathrm{r}=0}_{f_{\mathrm{rKER}}(R_\mathrm{r})} \right). \tag{WLRA$_{R_\mathrm{r}}$}$$

(WLRA$_P$), (WLRA$_L$), (WLRA$_{R_\ell}$), and (WLRA$_{R_\mathrm{r}}$) are double minimization problems with outer minimization over the parameters of the rank representation and inner minimization over the data approximation. The inner minimization problems are convex quadratic optimization problems and can be solved analytically, yielding explicit expressions for the cost functions to be minimized over the parameters of the rank representation.

Define for a symmetric positive definite matrix $W \in \mathbb{R}^{mn \times mn}$ and a full column rank matrix $K \in \mathbb{R}^{mn \times \bullet}$

$$f_{\mathrm{IM}}(K) := d^\top (W - WK(K^\top W K)^{-1}K^\top W)d \quad \text{and} \quad \hat{d}_{\mathrm{IM}}(K) := K(K^\top W K)^{-1}K^\top W d. \tag{IM SOL}$$

Theorem 3 (Equivalent problems to (WLRA$_P$) and (WLRA$_L$)). *The optimization problem* (WLRA$_P$) *is equivalent to*

$$\min_P f_{\ell\mathrm{IM}}(P) \quad \text{subject to} \quad P^\top P = I, \quad \text{where} \quad f_\ell\mathrm{IM}(P) = f_{\mathrm{IM}}(I_n \otimes P) \tag{WLRA$'_P$}$$

and the minimum is achieved at $\hat{d} = \hat{d}_{\mathrm{IM}}(I_n \otimes P)$. *Similarly,* (WLRA$_P$) *is equivalent to*

$$\min_L f_{\mathrm{rIM}}(L) \quad \text{subject to} \quad LL^\top = I, \quad \text{where} \quad f_{\mathrm{rIM}}(L) = f_{\mathrm{IM}}(L^\top \otimes I_m) \tag{WLRA$'_L$}$$

and the minimum is achieved at $\hat{d} = \hat{d}_{\mathrm{IM}}(L^\top \otimes I_m)$.

Proof. For $f_{\ell\mathrm{IM}}$, we have

$$f_{\ell\mathrm{IM}} = \min_{\hat{d},L}(d-\hat{d})^\top W(d-\hat{d}) \quad \text{subject to} \quad \hat{d} = \underbrace{(I_n \otimes P)}_{\mathbf{P}}\underbrace{\mathrm{vec}(L)}_{l}$$

$$= \min_l(d-\mathbf{P}l)^\top W(d-\mathbf{P}l).$$

This is a standard convex quadratic optimization problem, of which the solution is (IM SOL) with $K = \mathbf{P} \in \mathbb{R}^{mn \times rn}$.

For f_{rIM}, we have

$$f_{\mathrm{rIM}} = \min_{\hat{d}, P} (d - \hat{d})^\top W (d - \hat{d}) \quad \text{subject to } \hat{d} = \underbrace{(L^\top \otimes I_m)}_{\mathbf{L}} \underbrace{\mathrm{vec}(P)}_{p}$$

$$= \min_p (d - \mathbf{L}p)^\top W (d - \mathbf{L}p),$$

so that the solution is given by (IM SOL) with $K = \mathbf{L} \in \mathbb{R}^{mn \times mr}$. $\quad \square$

Define for a symmetric positive definite matrix $W \in \mathbb{R}^{mn \times mn}$ and a full row rank matrix $K \in \mathbb{R}^{mn \times \bullet}$

$$f_{\mathrm{KER}}(K) := d^\top K^\top (K W^{-1} K^\top)^{-1} K d, \quad \text{and}$$
$$\hat{d}_{\mathrm{KER}}(K) := \left(I - W^{-1} K^\top (K W^{-1} K^\top)^{-1} K \right) d. \qquad \text{(KER SOL)}$$

Theorem 4 (Equivalent problems to (WLRA$_{R_\ell}$) and (WLRA$_{R_r}$)). *The optimization problem* (WLRA$_{R_\ell}$) *is equivalent to*

$$\min_{R_\ell} f_{\ell\mathrm{KER}}(R_\ell) \quad \text{subject to } R_\ell R_\ell^\top = I, \quad \text{where } f_{\ell\mathrm{KER}(R_\ell)} = f_{\mathrm{KER}}(I_n \otimes R_\ell)$$

$$\text{(WLRA}'_{R_\ell})$$

and the minimum is achieved at $\hat{d} = \hat{d}_{\mathrm{KER}}(I_n \otimes R_\ell)$. *Similarly,* (WLRA$_{R_r}$) *is equivalent to*

$$\min_{R_r} f_{\mathrm{rKER}}(R_r) \quad \text{subject to } R_r^\top R_r = I, \quad \text{where } f_{\mathrm{rKER}}(R_r) = f_{\mathrm{KER}}(R_r^\top \otimes I_m)$$

$$\text{(WLRA}'_{R_r})$$

and the minimum is achieved at $\hat{d} = \hat{d}_{\mathrm{KER}}(R_r^\top \otimes I_m)$.

Proof. Let $\Delta D := D - \widehat{D}$ and $\Delta d := \mathrm{vec}(\Delta D)$. For $f_{\ell\mathrm{KER}}$, we have

$$f_{\ell\mathrm{KER}} = \min_{\Delta d} \Delta d^\top W \Delta d \quad \text{subject to } R_\ell(D - \Delta D) = 0$$

$$= \min_{\Delta d} \Delta d^\top W \Delta d \quad \text{subject to } (I_n \otimes R_\ell)\Delta d = \mathrm{vec}(R_\ell D),$$

which is a standard weighted least norm problem. The solution is given by (KER SOL) with $K = I_n \otimes R_\ell$. For f_{rKER}, we have

$$f_{\mathrm{rKER}} = \min_{\Delta d} \Delta d^\top W \Delta d \quad \text{subject to } (D - \Delta D) R_r = 0$$

$$= \min_{\Delta d} \Delta d^\top W \Delta d \quad \text{subject to } (R_r \otimes I_m)\Delta d = \mathrm{vec}(D R_r),$$

of which the solution is given by (KER SOL) with $K = R_r \otimes I_m$. $\quad \square$

4. Free parameters and input/output parameterizations

The outer minimization problems (WLRA$'_L$), (WLRA$'_P$), (WLRA$'_{R_\ell}$), and (WLRA$'_{R_r}$) are difficult constrained nonconvex optimization problems. In order to simplify them, we first elimi-

nate the quadratic constraints $P^\top P = I$, $LL^\top = I$, $R_\ell R_\ell^\top = I$, and $R_r^\top R_r = I$ by re-parameter-izing the problems. The parameters P, L, R_ℓ, and R_r are not unique and the nonuniqueness is up to a choice of basis for col span(P), row span(L), ker(R_ℓ), and leftker(R_r), respectively.

Theorem 5 (Equivalent parameters). *For any nonsingular matrices U, V, Q_ℓ, and Q_r of appropriate size,*

- $f_{\ell IM}(P) = f_{\ell IM}(PU)$,
- $f_{rIM}(L) = f_{rIM}(VL)$,
- $f_{\ell KER}(R_\ell) = f_{\ell KER}(Q_\ell R_\ell)$, *and*
- $f_{rKER}(R_r) = f_{rKER}(R_r Q_r)$.

In the above sense the parameters PU, VL, $Q_\ell R_\ell$, and $R_r Q_r$ are equivalent to P, L, R_ℓ, and R_r, respectively.

Proof. By the properties of the Kronecker product, we have

- $I_n \otimes (PU) = (I_n \otimes P)(I_n \otimes U)$,
- $(VL)^\top \otimes I_m = (L^\top \otimes I_m)(V^\top \otimes I_m)$,
- $I_n \otimes (Q_\ell R_\ell) = (I_n \otimes Q_\ell)(I_n \otimes R_\ell)$,
- $(R_r Q_r)^\top \otimes I_m = (Q_r^\top \otimes I_m)(R_r^\top \otimes I_m)$.

The matrices $I_n \otimes U$, $V^\top \otimes I_m$, $I_n \otimes R_\ell$, $R_r^\top \otimes I_m$ are nonsingular if and only if, respectively, the matrices U, V, Q_ℓ, Q_r are nonsingular.

Consider the solution (IM SOL) for the problems with image representations and let T be a nonsingular matrix of dimension coldim(K). We will show that $f_{IM}(K) = f_{IM}(KT)$. Indeed,

$$
\begin{aligned}
d_{IM}(KT) &:= KT(T^\top K^\top WKT)^{-1}T^\top K^\top Wd \\
&= KTT^{-1}(K^\top WK)^{-1}(T^\top)^{-1}T^\top K^\top Wd = d_{IM}(K). \qquad \square
\end{aligned}
$$

Consider then the solution (KER SOL) for the problems with kernel representations and let T be a nonsingular matrix of dimension row dim(K). We will show that $f_{KER}(K) = f_{KER}(TK)$. Indeed,

$$
\begin{aligned}
d_{KER}(TK) &:= (I - W^{-1}K^\top T^\top (TKW^{-1}K^\top T^\top)^{-1}TK)d \\
&= (I - W^{-1}K^\top T^\top (T^\top)^{-1}(KW^{-1}K^\top)T^{-1}TK)d = d_{KER}(K). \qquad \square
\end{aligned}
$$

Next we show special equivalent parameters that suggest a re-parameterization eliminating the nonuniqueness. Let $\Pi_\ell \in \mathbb{R}^{m \times m}$ and $\Pi_r \in \mathbb{R}^{n \times n}$ be permutation matrices and define the partitionings

$$
\Pi_\ell P =: \begin{bmatrix} P_1 \\ P_2 \end{bmatrix} \begin{matrix} r \\ m-r \end{matrix}, \quad
L\Pi_r =: \begin{bmatrix} L_1 & L_2 \end{bmatrix} r \begin{matrix} r & n-r \end{matrix}, \quad
R_\ell \Pi_\ell^{-1} =: \begin{bmatrix} R_{\ell,1} & R_{\ell,2} \end{bmatrix} \begin{matrix} r & m-r \\ m-r \end{matrix},
$$

$$
\Pi_r^{-1} R_r =: \begin{bmatrix} R_{r,1} \\ R_{r,2} \end{bmatrix} \begin{matrix} r \\ n-r \end{matrix} . \tag{I/O PRT}
$$

Since P and R_r are full column rank and L and R_ℓ are full row rank, then there are permutations Π_ℓ and Π_r, such that the square blocks P_1, L_1, $R_{\ell,2}$, and $R_{r,2}$ are nonsingular. In what follows, assume that Π_ℓ and Π_r are such permutations. Then equivalent parameters to $\begin{bmatrix} P_1 \\ P_2 \end{bmatrix}$, $\begin{bmatrix} L_1 & L_2 \end{bmatrix}$, $\begin{bmatrix} R_{\ell,1} & R_{\ell,2} \end{bmatrix}$, and $\begin{bmatrix} R_{r,1} \\ R_{r,2} \end{bmatrix}$ are

$$\widetilde{P} := \begin{bmatrix} I_r \\ P_2 P_1^{-1} \end{bmatrix}, \quad \widetilde{L} := \begin{bmatrix} I_r & L_1^{-1} L_2 \end{bmatrix}, \quad \widetilde{R}_\ell := \begin{bmatrix} -R_{\ell,2}^{-1} R_{\ell,1} & -I_{m-r} \end{bmatrix},$$

$$\widetilde{R}_r := \begin{bmatrix} -R_{r,1} R_{r,2}^{-1} \\ -I_{n-r} \end{bmatrix}.$$

Theorem 6 (Link between image and kernel parameters). *If col span $(P) = ker(R_\ell)$ and at least one of the blocks P_1 or $R_{\ell,2}$, defined in (I/O PRT), is nonsingular, then the other block is also nonsingular and*

$$P_2 P_1^{-1} = -R_{\ell,2}^{-1} R_{\ell,1} =: X_\ell \in \mathbb{R}^{(m-r) \times r}.$$

Similarly, if row span(L) = left ker(R_r) and at least one of the blocks L_1 or $R_{r,2}$, defined in (I/O PRT), is nonsingular, then the other block is also nonsingular and

$$L_1^{-1} L_2 = -R_{r,1} R_{r,2}^{-1} =: X_r. \in \mathbb{R}^{r \times (n-r)}.$$

Proof. For the first statement we have

$$\mathrm{col\,span}(P) = \ker(R_\ell) \Rightarrow \widetilde{R}_\ell \widetilde{P} = 0 \Rightarrow \begin{bmatrix} -R_{\ell,2}^{-1} R_{\ell,1} & -I_{m-r} \end{bmatrix} \begin{bmatrix} I_r \\ P_2 P_1^{-1} \end{bmatrix}$$

$$= 0 \Rightarrow P_2 P_1^{-1} = -R_{\ell,2}^{-1} R_{\ell,1}$$

and for the second one

$$\mathrm{row\,span}(L) = \mathrm{leftker}(R_r) \Rightarrow \widetilde{L} \widetilde{R}_r = 0 \Rightarrow \begin{bmatrix} I_r & L_1^{-1} L_2 \end{bmatrix} \begin{bmatrix} -R_{r,1} R_{r,2}^{-1} \\ -I_{n-r} \end{bmatrix}$$

$$\Rightarrow L_1^{-1} L_2 = -R_{r,1} R_{r,2}^{-1}. \quad \square$$

Define the two partitionings of the matrix $\widehat{D}$

$$\Pi_\ell \widehat{D} =: \begin{bmatrix} \widehat{A}_\ell \\ \widehat{B}_\ell \end{bmatrix} \begin{matrix} r \\ m-r \end{matrix} \quad \text{and} \quad \widehat{D} \Pi_r =: \begin{matrix} r \quad n-r \\ [\widehat{A}_r \quad \widehat{B}_r] \end{matrix}.$$

The representations

$$\text{there is } X_\ell \in \mathbb{R}^{(m-r) \times r}, \text{ such that } X_\ell \widehat{A}_\ell = \widehat{B}_\ell \Rightarrow \mathrm{rank}(\widehat{D}) \leqslant r \tag{lI/O}$$

and

$$\text{there is } X_r \in \mathbb{R}^{r \times (n-r)}, \text{ such that } \widehat{A}_r X_r = \widehat{B}_r \Rightarrow \mathrm{rank}(\widehat{D}) \leqslant r \tag{rI/O}$$

are called input/output ones. Their parameters X_ℓ and X_r are unconstrained. Using the left input/output parameterization (lI/O), the WLRA problem becomes

$$\min_{\widehat{A}_\ell, \widehat{B}_\ell, X_\ell} \mathrm{vec}^\top \left(D - \Pi_\ell^{-1} \begin{bmatrix} \widehat{A}_\ell \\ \widehat{B}_\ell \end{bmatrix} \right) W \mathrm{vec} \left(D - \Pi_\ell^{-1} \begin{bmatrix} \widehat{A}_\ell \\ \widehat{B}_\ell \end{bmatrix} \right) \quad \text{subject to } X_\ell \widehat{A}_\ell = \widehat{B}_\ell$$

$$\tag{lWTLS}$$

and using the right input/output parameterization (rI/O), the WLRA problem becomes

$$\min_{\widehat{A}_r, \widehat{B}_r, X_r} \text{vec}^\top \left(D - \begin{bmatrix} \widehat{A}_r & \widehat{B}_r \end{bmatrix} \Pi_r^{-1} \right) W \text{vec} \left(D - \begin{bmatrix} \widehat{A}_r & \widehat{B}_r \end{bmatrix} \Pi_r^{-1} \right)$$

$$\text{subject to } \widehat{A}_r X_r = \widehat{B}_r. \tag{rWTLS}$$

For fixed permutation matrices, problems (lWTLS) and (rWTLS) are known in the literature as weighted total least squares problems. Under the assumption that the permutations Π_ℓ and Π_r are properly chosen, (lWTLS) and (rWTLS) are equivalent to (WLRA).

Note 1. If the permutations Π_ℓ and Π_r are not properly chosen, the weighted total least squares problems might have no solution or be close to having no solution (which leads to ill conditioned computational problems). Such cases are called in the literature nongeneric total least squares problems and have been analyzed in [11,12,8]. In particular, the formulation of [8] defines a so-called core problem, corresponding to a given total least squares problem, that always has a unique solution. The solution of the core problem coincides with the classical total least squares solution whenever the latter exists and leads to the nongeneric total least squares solution, defined in [11,12], when the classical total least squares solution fails to exist. The problem of detecting when a given problem is nongeneric, however, involves checking singularity of a submatrix and is therefore not straightforward numerically. In addition, the core problem is currently formulated only for single right-hand-side problems (corresponding to low-rank approximation problems with rank reduction by one). Statistical methods for detecting nongeneric cases and generalization of the core problem formulation for the multiple-right-hand side total least squares problems are topics of current research.

We solve the unconstrained outer minimization problems

$$\min_{X_\ell} f_{\ell IM} \left(\Pi_\ell^{-1} \begin{bmatrix} I \\ X_\ell \end{bmatrix} \right), \quad \min_{X_r} f_{rIM} \left(\begin{bmatrix} I & X_r \end{bmatrix} \Pi_r^{-1} \right),$$

$$\min_{X_\ell} f_{\ell KER} \left(\begin{bmatrix} X_\ell & -I \end{bmatrix} \Pi_r \right), \quad \text{and} \quad \min_{X_r} f_{rKER} \left(\Pi_\ell \begin{bmatrix} X_r \\ -I \end{bmatrix} \right) \tag{WLRA''}$$

using local optimization methods for nonlinear least squares problems, e.g., the Levenberg–Marquardt method. Applying these methods, of prime importance is the number of optimization parameters and the efficiency of the cost function (and its derivatives) evaluation. From the point of view of minimal number of optimization variables the left (right) input/output parameterization should be chosen when $m < n$ ($m > n$), irrespective of the rank reduction r.

Next we verify on simulation examples that indeed the optimization problem using the "correct" (left or right) input/output representation is easier to solve in the sense that it requires fewer iterations steps under the same conditions (initial approximation and convergence tolerance).

Simulation example

The data matrix D is a sum of a random rank-r matrix and a random full rank matrix (errors-in-variables setup). The weight matrix W is a random positive definite matrix. In MATLAB D and W are generated as follows:

```
d0 = rand(m, r) * rand(r, n);
d0 = d0/norm(d0, 'fro');            %normalization
```

Table 1
Number of optimization variables and average iterations for the optimization algorithm solving (WLRA″) by left and right kernel representations

		m	2	3	4	5	6	7	8	9	10
		n	10	9	8	7	6	5	4	3	2
$r = 1$	Left		1/4	2/6	3/7	4/10	5/5	6/18	7/9	8/10	9/11
	Right		9/11	8/13	7/7	6/10	5/5	4/11	3/7	2/6	1/4
$r = 2$	Left		–	2/5	4/27	6/21	8/11	10/44	12/27	14/28	–
	Right		–	14/12	12/34	10/69	8/11	6/39	4/24	2/6	–
$r = 3$	Left		–	–	3/11	6/26	9/11	12/46	15/53	–	–
	Right		–	–	15/30	12/33	9/14	6/44	3/10	–	–
$r = 4$	Left		–	–	–	4/9	8/14	16/12	–	–	–
	Right		–	–	–	16/12	8/13	4/8	–	–	–

```
w = rand(m * n, 2 * m * n); w = w * w';
dt = reshape(sqrtm(w)\randn(m * n, 1), m, n);
dt = dt/norm(dt,' fro');          %normalization
 d = d0 + 0.001 * dt;
```

The matrix dimension m is varied from 2 to 10 and $n := 12 - m$. The desired rank r is varied from 1 to 4. For each combination of m and r, $N = 50$ independently generated data matrices are generated and for each one of them the problems (WLRA″) with left and right input/output representations are solved.

We use the Levenberg–Marquardt optimization algorithm from the Optimization toolbox of MATLAB (`lsqnonlin`) with cost function evaluations only. In both cases (left and right input/output representation), the algorithm is run from equivalent initial approximations obtained using the singular value decompositions (unweighted low-rank approximation) and the stopping criteria are set to the same tolerances. In this particular experiment, in all runs the same solution was found by using the two representations. As expected the average number of iterations, however, depends on the number of optimization parameters (elements of X_ℓ and X_r): the fewer the optimization variables, the fewer the average number of iterations for convergence. Numerical results are shown in Table 1 in the format "# of optimization variables/# of iterations".

5. Special cases of column and row-weighting

Up to now we have considered the WLRA problem with a full weight matrix W. Of interest, however, are the two special cases of element-wise weighting, i.e., W diagonal, column-wise weighting, i.e., W block diagonal with blocks of dimension $m \times m$, and row-wise weighting, i.e., $T^\top W T$ block diagonal with blocks of dimension $n \times n$, where T is the $mn \times mn$ permutation matrix, such that $\text{vec}(D) = T \text{vec}(D^\top)$. The special structure of W can be exploited for efficient cost function and first derivative evaluation.

Consider the column-wise weighting case, i.e., W is block-diagonal with blocks of dimension $m \times m$. It turns out that the cost function formulas using the left image and left kernel representations are easy to rewrite taking into account the special structure of W.

$$\hat{d}_{\ell\mathrm{IM}} := (I \otimes P)((I \otimes P)^\top W (I \otimes P))^{-1}(I \otimes P)^\top W d$$
$$\Leftrightarrow \hat{d}_{\ell\mathrm{IM},i} = P(P^\top W_i P)^{-1} P^\top W_i d_i,$$

where d_i is the ith column of D and $\hat{d}_i$ is the ith column of $\widehat{D}$. Similar decoupling into independent terms occurs for the left kernel representation:

$$\hat{d}_{\ell\mathrm{KER}} = \left(I_{mn} - W^{-1}(I \otimes R_\ell)^\top \left((I \otimes R_\ell) W (I \otimes R_\ell)^\top \right)^{-1} (I \otimes R_\ell) \right) d$$

$$\Leftrightarrow \hat{d}_{\ell\mathrm{KER},i} = \left(I_m - W_i^{-1} R_\ell^\top (R_\ell W_i R_\ell^\top)^{-1} R_\ell \right) d_i.$$

In the case of right image and right kernel representations the computational savings are achieved as follows. For the right image representation

$$\hat{d}_{\mathrm{rIM}} := (L^\top \otimes I_m) \left((L^\top \otimes I_m)^\top W (L^\top \otimes I_m) \right)^{-1} (L^\top \otimes I_m)^\top W d, \tag{*}$$

denoting the ith column of L by l_i, we have

$$(L^\top \otimes I_m)^\top W (L^\top \otimes I_m) = \begin{bmatrix} l_1 \otimes I_m & \cdots & l_n \otimes I_m \end{bmatrix}$$
$$\times \begin{bmatrix} W_1 & & \\ & \ddots & \\ & & W_n \end{bmatrix} \begin{bmatrix} l_1^\top \otimes I_m \\ \vdots \\ l_n^\top \otimes I_m \end{bmatrix} = \sum_{i=1}^n l_i l_i^\top \otimes W_i.$$

Let

$$p := \left(\sum_{i=1}^n l_i l_i^\top \otimes W_i \right)^{-1} \begin{bmatrix} l_1 \otimes W_1 & \cdots & l_n \otimes W_n \end{bmatrix} d \tag{**}$$

and $P := \begin{bmatrix} p_1 & \cdots & p_r \end{bmatrix}$, where $p^\top =: \begin{bmatrix} p_1^\top & \cdots & p_r^\top \end{bmatrix}$ and $p_i \in \mathbb{R}^m$. Note that $p = \mathrm{vec}(P)$. Finally,

$$\hat{d}_{\mathrm{rIM}} = (L^\top \otimes I_m) p \Rightarrow \widehat{D}_{\mathrm{rIM}} = PL.$$

Computing p using (**) and setting $\widehat{D}_{\mathrm{rIM}} = PL$ is a more efficient alternative to the direct computation of $\hat{d}_{\mathrm{rIM}}$ using (*).

6. Conclusions

We have presented four representations of the rank constraint that turn the abstract weighted low-rank problem formulation in concrete parameter optimization problems. The optimization problem was solved analytically for the approximating matrix $\widehat{D}$, which resulted in an equivalent optimization problem over the parameters of the rank representation only. We showed a unique parameterization that avoids the quadratic constraint and further transforms the optimization problem into a classical nonlinear least squares problem. Finally we showed how the cost function can be computed efficiently in two special cases: row-wise and column-wise weighting.

Acknowledgments

References

[1] B. De Moor, Structured total least squares and L_2 approximation problems, Linear Algebra Appl. 188–189 (1993) 163–207.

[2] G. Eckart, G. Young, The approximation of one matrix by another of lower rank, Psychometrika 1 (1936) 211–218.

[3] K. Gabriel, S. Zamir, Lower rank approximation of matrices by least squares with any choice of weights, Technometrics 21 (1979) 489–498.

[4] J. Manton, R. Mahony, Y. Hua, The geometry of weighted low-rank approximations, IEEE Trans. Signal Process. 51 (2) (2003) 500–514.

[5] I. Markovsky, M.-L. Rastello, A. Premoli, A. Kukush, S. Van Huffel, The element-wise weighted total least squares problem, Comput. Statist. Data Anal. 50 (1) (2005) 181–209.

[6] I. Markovsky, J.C. Willems, S. Van Huffel, B. De Moor, Exact and Approximate Modeling of Linear Systems: A Behavioral Approach, Monographs on Mathematical Modeling and Computation, vol. 11, SIAM, 2006.

[7] A. Premoli, M.-L. Rastello, The parametric quadratic form method for solving TLS problems with elementwise weighting, in: S. Van Huffel, P. Lemmerling (Eds.), Total Least Squares and Errors-in-Variables Modeling: Analysis, Algorithms and Applications, Kluwer, 2002, pp. 67–76.

[8] C. Paige, Z. Strakos, Core problems in linear algebraic systems, SIAM J. Matrix Anal. Appl. 27 (2005) 861–875.

[9] M. Schuermans, I. Markovsky, P. Wentzell, S. Van Huffel, On the equivalence between total least squares and maximum likelihood PCA, Anal. Chim. Acta 544 (2005) 254–267.

[10] C. Van Loan, N. Pitsianis, Approximation with Kronecker products, in: M. Moonen, G. Golub (Eds.), Linear Algebra for Large Scale and Real Time Applications, Kluwer Publications, 1993, pp. 293–314.

[11] S. Van Huffel, J. Vandewalle, Analysis and solution of the nongeneric total least squares problem, SIAM J. Matrix Anal. Appl. 9 (1988) 360–372.

[12] S. Van Huffel, J. Vandewalle, The Total Least Squares Problem: Computational Aspects and Analysis, SIAM, Philadelphia, 1991.

[13] P. Wentzell, D. Andrews, D. Hamilton, K. Faber, B. Kowalski, Maximum likelihood principle component analysis, J. Chemometrics 11 (1997) 339–366.