Regularisation for PCA- and SVD-type matrix factorisations

Abdolrahman Khoshrou[a,b,1,*], Eric J. Pauwels[a,2]

[a]Centrum Wiskunde & Informatica, Science Park 123, 1098 XG, Amsterdam, The Netherlands
[b]Department of Mathematics and Computer Science, Delft University of Technology, The Netherlands

Abstract

Singular Value Decomposition (SVD) and its close relative, Principal Component Analysis (PCA), are well-known linear matrix decomposition techniques that are widely used in applications such as dimension reduction and clustering. However, an important limitation of SVD/PCA is its sensitivity to noise in the input data. In this paper, we take another look at the problem of regularisation and show that different formulations of the minimisation problem lead to qualitatively different solutions.

Keywords: Singular value decomposition (SVD), Principal component analysis (PCA), matrix factorisation, regularisation, dimensionality reduction, graph Laplacian, feature manifold.

1. Introduction and Motivation

1.1. Introduction and Related Work

Singular Value Decomposition (SVD) and its close relative, Principal Component Analysis (PCA), are well-known linear matrix factorisation techniques that are widely used in applications as varied as dimension reduction and clustering, matrix completion [1] (e.g. for recommender systems), dictionary learning [2] and time series analysis [3]. In a surprising turn of events, (deep) matrix factorisation also plays a role in the implicit regularisation that enables acceptable generalisation in deep learning [4].

In their abstract version, SVD and PCA amount to two different but related types of matrix factorisation. More precisely, given a general (data) matrix A, the aim is to approximate it as a product of simpler (i.e. lower-rank) matrices. Specifically:

- PCA-type decomposition: $A \approx PQ^T$ where the columns of Q are orthonormal, i.e. $Q^TQ = I$;
- SVD-type decomposition: $A \approx PBQ^T$ where B is diagonal, while $P^TP = I,\ Q^TQ = I$.

The approximation in the above equations is measured in terms of the Frobenius (matrix) norm which for an arbitrary matrix $X \in \mathbb{R}^{n\times m}$ is defined as:

$$|X|_F^2 = \sum_{i=1}^{n}\sum_{j=1}^{m} x_{ij}^2 = \mathrm{Tr}(XX^T) = \mathrm{Tr}(X^TX) = |X^T|_F^2. \tag{1}$$

(In the remainder of the paper, we will drop the subscript F).

Although these factorisation techniques are both conceptually simple and effective, it is well-known that they are sensitive to noise and outliers in the input data. As a consequence, some modifications of the original algorithms have been proposed to alleviate the effect of these disturbances [5, 6]. Candes et al. [7] introduce

*Corresponding author: *a.khoshrou@cwi.nl*
[1]Abdolrahman Khoshrou is a (guest) PhD student at TU Delft. He is also with the Intelligent and Autonomous Systems Group at National Dutch Mathematics and Informatics Center (CWI).
[2]Eric Pauwels is a senior researcher and the leader of the Intelligent and Autonomous Systems Group at CWI.

Robust PCA (RPCA) which aims to separate signal from outliers by decomposing any given matrix into the sum of a low-rank approximation and a sparse matrix of outliers. An extension of this work for inexact recovery of the data is presented in [8]. Another example of sparse PCA using low rank approximation is proposed in [9].

Adding a regularisation term is another versatile way to tackle the problem of noisy input. For instance, Dumitrescu et al. [10] show how a regularized version of K-SVD algorithm can be adapted to the Dictionary Learning (DL) problem. However, the presence of noise in the input is not the only reason to invoke regularisation. Recent research [11] shows that in many real data sets, it is not only the observed data that lie on a (non-)linear low dimensional manifold, but this also applies to the features. He et al. [12] point out that if the columns of the matrix A are interpreted as data points, then the rows are features. The neighbourhood structure of both the data points and the features give rise to distinct graphs (the so-called data and the feature graph) and hence, to corresponding graph Laplacians (L_d and L_f respectively). The resulting regularised PCA is referred to as the *graph-dual Laplacian PCA* (gDLPCA) and for a given data matrix A, is obtained by minimising the functional:

$$J(V,Y) = |A - VY|^2 + \alpha \,\mathrm{Tr}(V^T L_d V) + \beta \,\mathrm{Tr}(Y L_f Y^T) \qquad \text{subject to } V^T V = I \tag{2}$$

The ability of the graph dual regularization technique to incorporate both data and feature structure has deservedly attracted considerable attention in dimensionality reduction applications [13, 14, 12].

In the present paper, we take the functional. (2) as a starting point and investigate the two factorisation approaches mentioned above (invoking eq. (1) to recast the trace as a norm):

- PCA-type decomposition ($A \approx PQ^T$) by minimising the regularisation functional:

$$|A - PQ^T|^2 + \lambda |DP|^2 + \mu |GQ|^2 \tag{3}$$

- SVD-type decomposition ($A \approx PBQ^T$) by minimising the regularisation functional:

$$|A - PBQ^T|^2 + \lambda |DP|^2 + \mu |GQ|^2 \tag{4}$$

The minimisation of the functional (3) was discussed in [12], but the proposed solution contains an error which we correct in this paper. In addition, we also provide an algorithm to solve functional (4), which somewhat surprisingly is quite different from the one for (3).

The remainder of this paper is organised as follows: We finalise this section by recapitulating some important facts facts about SVD. In section 2 and 3 we derive an algorithm for minimisation of the regularised version of PCA-type and SVD-type factorisation, respectively. In section 4 how gradient descent can be implemented by drawing on some elementary facts from Lie-group theory. Finally, we conclude by giving some pointers to potential extensions.

1.2. Brief recap of Singular Value Decomposition (SVD)

For the sake of completeness, we first recall the well-known SVD result; for more details we refer to standard textbooks such as [15][16].

Theorem 1 (Singular Value Decomposition, SVD). *Any real-valued $n \times m$ matrix A can be factorized into the product of three matrices:*

$$A = USV^T \qquad \text{where} \quad U \in O(n) \quad \text{and} \quad V \in O(m) \quad \text{are orthonormal,} \tag{5}$$

and S is an $n \times m$ diagonal matrix where the elements on the main "diagonal" (so-called singular values) are non-negative (i.e. $\sigma_i := S_{ii} \geq 0$ for $1 \leq i \leq \min(n,m)$).
Assuming that the rank $rk(A) = r \leq \min(n,m)$, we can sort the singular values such that

$$\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_r > 0 = \sigma_{r+1} = \ldots = \sigma_{\min(n,m)}$$

and recast eq. (5) *as*

$$A = \sum_{i=1}^{r} \sigma_i U_i V_i^T \qquad \textit{where } U_i, V_i \textit{ are the } i\textit{-th columns of } U \textit{ and } V\textit{, respectively.} \tag{6}$$

For the singular values sorted as above, we introduce the short-hand notation $U_{(1:k)}$ *and* $V_{(1:k)}$ *to denote the matrix comprising the first* k *columns of* U *and* V*, respectively:*

$$U_{(1:k)} := [U_1, U_2, \ldots, U_k] \qquad \textit{and} \qquad V_{(1:k)} := [V_1, V_2, \ldots, V_k].$$

In this notation, eq. (6) *can be expressed concisely as:*

$$A = U_{(1:r)}\, diag(\sigma_1, \ldots, \sigma_r)\, V_{(1:r)}^T. \tag{7}$$

■

To appreciate the significance of Theorem 1, it is helpful to highlight its geometric interpretation. Recall that any $n \times m$ matrix A gives rise to a corresponding linear transformation $A : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ that maps the standard basis in $\mathbb{R}^m$ into the columns of A:

$$A\mathbf{e}_k = A_k \qquad \text{where} \quad \mathbf{e}_k = (0, 0, \ldots, 0, 1, 0, \ldots, 0)^T.$$

Roughly speaking, the SVD theorem therefore tells us that it is always possible to select an *orthonormal* basis in $\mathbb{R}^m$ (columns of V) that is mapped (up to non-negative scaling factors, i.e. the singular values) into an *orthonormal* basis in $\mathbb{R}^n$ (columns of U). This is immediately obvious from eq. (6):

$$AV_\ell = \sum_{k=1}^{r} \sigma_k U_k V_k^T V_\ell = \sum_{k=1}^{r} \sigma_k U_k \delta_{k\ell} = \sigma_\ell U_\ell.$$

where $\delta_{k\ell}$ is a Kronecker delta function. It is worth noting that insisting on the orthogonality of V ($V^T V = I$) is not restrictive. Indeed, a linear transformation is completely and uniquely determined by specifying its effect on any basis, and there is no loss of generality by insisting on the orthonormality of this basis. However, the non-trivial message of this theorem is this orthonormal basis (V) can be chosen in such a way that its image U under A is also orthonormal (again, up to non-negative scalings). Furthermore, in a generic case (where all singular values are different) the singular value decomposition is unique, up to an arbitrary relabeling of the *basis-vectors* and a simultaneous sign-flip of corresponding columns in U and V, i.e. $(U_\ell, V_\ell) \to (-U_\ell, -V_\ell)$ for any number of columns.

The importance of the SVD result, and the starting point for this paper, is the following well-known minimisation result (more details can be found in [17, 18]).

Theorem 2 (Eckart-Young-Mirsky Theorem: Optimal low rank approximation). *Let us consider an* $n \times m$ *matrix* A *with rank* $rk(A) = r \leq \min(n, m)$*. For* $k < r$*, finding the rank-*k *matrix* A_k *that is closest to* A *in (Frobenius) norm gives rise to the following constrained minimisation problem:*

$$\min_{A_k} |A - A_k|^2 \qquad \textit{subject to} \quad rk(A_k) \leq k.$$

The solution to this problem is obtained by truncating the SVD expansion eq. (6) *after the* k*-th largest singular value:*

$$A_k = \sum_{i=1}^{k} \sigma_i U_i V_i^T = U_{(1:k)}\, diag(\sigma_1, \ldots, \sigma_k)\, V_{(1:k)}^T. \tag{8}$$

■

Recall that a rank-k matrix of size $n \times m$ can always be written as a product $A_k = PQ^T$ where $P \in \mathbb{R}^{n\times k}$ and $Q \in \mathbb{R}^{m\times k}$ are matrices of full rank k. Again, in this factorisation, there is no loss of generality in requiring $Q^TQ = I_k$. In fact, it is necessary to remove indeterminacy due to arbitrary but trivial rescalings such as $P \longmapsto rP$ while $Q \longmapsto (1/r)Q$ (with $r \neq 0$), and the like. Hence, one can reformulate Theorem 2 as the factorisation result in Theorem 3.

Theorem 3 (PCA-type factorisation). *Assume that the $n \times m$ matrix A has rank $rk(A) = r \leq \min(n, m)$. We now define the functional $G(P, Q)$ as follow:*

$$G(P, Q) = |A - PQ^T|^2 \tag{9}$$

and the corresponding constrained optimisation problem:

$$\min_{P,Q} G(P, Q) \qquad \text{subject to} \quad rk(P) = rk(Q) = k \quad \text{and} \quad Q^TQ = I_k \tag{10}$$

where $k < r$. A solution to the above constrained minimisation problem (in $P \in \mathbb{R}^{n\times k}$ and $Q \in \mathbb{R}^{m\times k}$) is given by (using the SVD notation given above):

$$Q = V_{(1:k)} \qquad \text{and} \qquad P = U_{(1:k)}\, diag(\sigma_1, \ldots, \sigma_k) \tag{11}$$

hence:

$$PQ^T = \sum_{i=1}^{k} \sigma_i U_i V_i^T. \tag{12}$$

From (11) *this it also follows that P^TP is diagonal, but not necessarily equal to the identity.* ■

Note that If we drop the insistence on the diagonal form for P^TP (i.e. P need no longer be an orthogonal frame), then the solution is no longer unique. Indeed, by taking any $k \times k$ orthogonal matrix R with $R^TR = I_k = RR^T$, it is clear that $P' = PR$ and $Q' = QR$ are also solutions. In this case: $Q'^TQ' = R^TQ^TQR = I_k$ but $P'^TP' = R^TP^TPR = R^T(SS^T)R$ is in general a positive definite symmetric matrix.

2. Regularisation for PCA-type factorisation

2.1. Regularised PCA

The following theorem outlines an obvious generalisation to the regularised version of the minimisation problem.

Theorem 4 (Regularised PCA). *Let A be an $n \times m$ matrix of rank $r \leq \min(n, m)$. For $k \leq r$, let $P \in \mathbb{R}^{n\times k}$ and $Q \in \mathbb{R}^{m\times k}$ full rank matrices (i.e. of rank k). Furthermore, for arbitrary (non-zero) integers d and g we introduce regularisation matrices $D \in \mathbb{R}^{d\times n}$ and $G \in \mathbb{R}^{g\times m}$, as well as weights $\lambda, \mu \geq 0$.*

We now define the following functional F in the variables P and Q:

$$F(P, Q) = |A - PQ^T|^2 + \lambda\,|DP|^2 + \mu\,|GQ|^2 \tag{13}$$

and pose the corresponding constrained optimisation problem:

$$\min_{P,Q} F(P, Q) \qquad \text{subject to} \qquad Q^TQ = I_k. \tag{14}$$

Introducing short-hand notation $L := D^TD \in \mathbb{R}^{n\times n}$ and $M := G^TG \in \mathbb{R}^{m\times m}$ (both symmetric and positive semi-definite), the solution of the constrained optimisation problem (14) *is constructed as follows:*

- *The k columns of the $m \times k$ matrix Q are the eigenvectors of the $m \times m$ matrix:*

$$K := A^T(I_n + \lambda L)^{-1}A - \mu\, M$$

corresponding to the k largest eigenvalues;

- *Furthermore:* $P = (I_n + \lambda L)^{-1} AQ$

For the sake of completeness, we reiterate that the condition $Q^T Q = I_k$ is not restrictive but necessary to eliminate arbitrary rescalings. In passing, we point out that result above corrects an error in [12] where it is incorrectly stated that $P = AQ$.

Proof. Since the variable P in the functional (13) in unconstrained, we can identify the optimum in P (for fixed Q) by computing the gradient:

$$\frac{1}{2}\nabla_P F = (PQ^T - A)Q + \lambda D^T DP \tag{15}$$

and solving for P:

$$\nabla_P F = 0 \quad \Rightarrow \quad P\underbrace{Q^T Q}_{I_k} - AQ + \lambda LP = 0 \quad \Rightarrow \quad (I_k + \lambda L)P = AQ. \tag{16}$$

This condition needs to hold at the solution point. By first re-writing $F(P, Q)$ formula as the trace of matrices and then plugging in (16), we have:

$$\begin{aligned} F(P,Q) &= \mathrm{Tr}\left[(A - PQ^T)(A^T - QP^T)\right] + \lambda\,\mathrm{Tr}(P^T LP) + \mu\,\mathrm{Tr}(Q^T MQ) \\ &= \mathrm{Tr}\left[AA^T - AQP^T - PQ^T A^T + PQ^T QP^T\right] + \lambda\,\mathrm{Tr}(P^T LP) + \mu\,\mathrm{Tr}(Q^T MQ) \end{aligned}$$

Considering the fact that the trace operator is invariant under transposition as well as cyclic permutation, and plugging in eq. (16) we arrive at:

$$\begin{aligned} F(P,Q) &= \mathrm{Tr}\left[AA^T - 2(I_n + \lambda L)PP^T + PP^T\right] + \lambda\,\mathrm{Tr}(P^T LP) + \mu\,\mathrm{Tr}(Q^T MQ) \\ &= \mathrm{Tr}\left(AA^T - PP^T - 2\lambda LPP^T\right) + \lambda\,\mathrm{Tr}(P^T LP) + \mu\,\mathrm{Tr}(Q^T MQ) \\ &= \mathrm{Tr}(AA^T) - \mathrm{Tr}(PP^T) - 2\lambda\,\mathrm{Tr}(LPP^T) + \lambda\,\mathrm{Tr}(P^T LP) + \mu\,\mathrm{Tr}(Q^T MQ) \\ &= \mathrm{Tr}(AA^T) - \mathrm{Tr}(P^T P) - \lambda\,\mathrm{Tr}(P^T LP) + \mu\,\mathrm{Tr}(Q^T MQ) \\ &= \mathrm{Tr}(AA^T) - \mathrm{Tr}\left[P^T \underbrace{(I_n + \lambda L)P}_{AQ}\right] + \mu\,\mathrm{Tr}(Q^T MQ). \end{aligned} \tag{17}$$

Extracting P and its transpose from eq. (16):

$$P = (I_n + \lambda L)^{-1} AQ \quad \Rightarrow \quad P^T = Q^T A^T (I_n + \lambda L)^{-1} \qquad \text{as } L \text{ is symmetric} \tag{18}$$

we arrive at:

$$F(P,Q) = \mathrm{Tr}(AA^T) - \mathrm{Tr}\left[Q^T \left(A^T (I_n + \lambda L)^{-1} A - \mu M\right) Q\right]. \tag{19}$$

Therefore, in order to minimize F, one must maximize the right-most term as $\mathrm{Tr}(AA^T)$ is a constant. This is achieved by selecting for Q, eigenvectors corresponding to the k largest eigenvalues of $(A^T (I_n + \lambda L)^{-1} A - \mu M)$. Once Q is determined, P is obtained via eq. (18).

As a concluding remark, we point out that the matrix $I_n + \lambda L$ is always invertible. Indeed, since $L = D^T D$ is positive semi-definite and symmetric, it has a complete set of eigenvectors with corresponding non-negative eigenvalues, i.e., $L = W\Lambda W^T$, where W is orthogonal (i.e. $W^T W = WW^T = I_n$) and $\Lambda \geq 0$. Hence, the matrix $(I_n + \lambda L)$ has strictly positive diagonal elements, and is indeed invertible. ■

Some illustrative numerical experiments can be found [19].

2.2. Some special cases

- $\boxed{\lambda = 0 \text{ and } \mu = 0}$: In that case, Q comprises the first k eigenvectors of $K = A^T A$ and $P = AQ$, which means that we end up with the standard SVD, as expected. Some numerical experiments can be found [20].

- $\boxed{D = I_n \text{ and } \mu = 0}$: The following section provides an overview of the results in [10] where a regularized K-SVD problem is addressed. In the aforementioned paper, the authors consider a special case, where $\mu = 0$ and $D = I_n$. Since this implies that $L = D^T D = I_n$ and $\mu M = 0$, the matrix K simplifies to

$$K = \frac{1}{1+\lambda} A^T A$$

The eigenvectors of K are therefore the right singular vectors of A (i.e. the eigenvectors of $A^T A$). Hence $Q = V_{(1:k)}$, and as a result:

$$P = \frac{1}{1+\lambda} AQ \quad \textit{and} \quad AQ = U_{(1:k)}\, diag(\sigma_1, \ldots, \sigma_k).$$

In particular, for $k = 1$ (the rank-1 reconstruction), we obtain:

$$Q = \mathbf{v}_1 \qquad \text{and} \qquad P = \frac{\sigma_1}{1+\lambda} \mathbf{u}_1$$

which is the result that can be found in [10]. The experiments are available in [21].

3. Regularisation for SVD-type factorisation

We now turn our attention to the SVD-type factorisation which looks for an approximation of the form:

$$A \approx PBQ^T \qquad \text{subject to:} \qquad Q^T Q = I_k, \quad |P_i| = 1 \ \forall i \in \{1, 2, \ldots, k\}, \quad \text{and } B \text{ diagonal.}$$

Loosely speaking, since the columns of P and Q are of unit length, they only pins down the structure of A, whereas the diagonal matrix $B = diag(\beta_1, \beta_2, \ldots, \beta_k)$ captures the *amplitude* of the corresponding structures. Similar to before, the columns of Q are orthonormal, i.e., we again insist on $Q^T Q = I_k$. However, unlike before, the columns of P are now only required to have unit length.

In light of the aforementioned SVD-type matrix factorisation technique, Theorems 5 and 6 provide an alternative solution to the lower-rank matrix approximation problem. For notational convenience, Theorem 5 first addresses the simplified case for $\mu = 0$. Finally, in Theorem 6 we return to the general case.

Theorem 5 (Regularised SVD). *Let A be an $n \times m$ matrix of rank $r \leq \min(n, m)$. For $k \leq r$, let $P \in \mathbb{R}^{n \times k}$ and $Q \in \mathbb{R}^{m \times k}$ of rank k, while $B \in \mathbb{R}^{k \times k}$ diagonal (i.e. $B = diag(\beta_1, \beta_2, \ldots, \beta_k)$). Furthermore, for arbitrary non-zero integer d we introduce regularisation matrix $D \in \mathbb{R}^{d \times n}$, as well as weight $\lambda \geq 0$. Finally, we introduce the short-hand notation $L := D^T D \in \mathbb{R}^{n \times n}$ (symmetric and positive-definite). We are now in a position to define the following functional F in the variables P, Q and B:*

$$F(P, Q, B) = |A - PBQ^T|^2 + \lambda\, |DP|^2, \tag{20}$$

and the corresponding constrained optimisation problem:

$$\min_{P,Q,B} F(P, Q, B) \qquad \textit{subject to:} \qquad Q^T Q = I_k, \quad |P_i| = 1 \ \forall i \in \{1, 2, \ldots, k\}, \quad \textit{and } B \textit{ diagonal.} \tag{21}$$

This problem is solved by the solution Algorithm 1 specified below.

Algorithm 1: Proposed RSVD method ($\mu = 0$)

Input: A, k, λ, D
Output: P, B, Q
Initialization
while *no convergence* **do**

1. *Determine the $m \times k$ matrix $Q = [\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k]$ (with orthonormal columns: $Q^T Q = I_k$) such that the sum of the smallest eigenvalue of each of the k symmetric matrices $S(\mathbf{q}_i)$ is minimal, i.e.:*

$$\min_Q \psi(Q) = \min_Q \sum_{i=1}^{k} \lambda_1(\mathbf{q}_i) \qquad \text{such that } Q^T Q = I_k$$

   where $\lambda_1(\mathbf{q}_i) = \min(eig(S(\mathbf{q}_i))$. To this end we use gradient descent (see Section 4).
2. *For each $\mathbf{q}_i$ as determined above, take $\mathbf{p}_i$ to be the eigenvector $W_1(\mathbf{q}_i)$ corresponding to the smallest eigenvector $\lambda_1(\mathbf{q}_i)$. Construct the $n \times k$ matrix $P = [\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k]$.*
3. *Finally, set $B = diag(\beta_1, \ldots, \beta_n)$ where $\beta_i = (P^T A Q)_{ii}$.*

end

Proof. Since B is unconstrained, we can determine its optimal value by computing the derivative with respect to B and equating it to zero:

$$\nabla_B F(P, Q, B) = \nabla_B |A - PBQ^T|^2. \tag{22}$$

Expanding the norm in terms of a trace (cf. eq. (1)), and using the invariance of a trace under transposition, we arrive at (recall $Q^T Q = I_k$):

$$\begin{aligned}
|A - PBQ^T|^2 &= \mathrm{Tr}\left[(A - PBQ^T)(A^T - QBP^T)\right] \\
&= \mathrm{Tr}(AA^T) - 2\mathrm{Tr}(AQBP^T) + \mathrm{Tr}(PB^2P^T) \\
&= |A|^2 - 2\,\mathrm{Tr}(P^T AQB) + \mathrm{Tr}(B^2 P^T P) \\
&= |A|^2 - 2\sum_{i=1}^{k}(P^T AQ)_{ii}\beta_i + \sum_{i=1}^{k}(P^T P)_{ii}\beta_i^2 \qquad (B \text{ is diagonal}) \qquad (23) \\
&= |A|^2 - 2\sum_{i=1}^{k}(P^T AQ)_{ii}\beta_i + \sum_{i=1}^{k}\beta_i^2 \qquad (|P_i| = 1 \Rightarrow (P^T P)_{ii} = 1). \qquad (24)
\end{aligned}$$

We therefore calculate the gradient of the functional F with respect to B as follow:

$$\frac{\partial}{\partial \beta_i} |A - PBQ^T|^2 = 2\,(\beta_i - (P^T AQ)_{ii}).$$

For given P and Q, we find the optimal B by insisting that the resulting gradient vanishes, which yields:

$$\beta_i = (P^T AQ)_{ii} \qquad \forall i \in \{1, 2, \ldots, k\}. \tag{25}$$

Plugging this optimal choice back into eq. (24) the functional (20) simplifies to

$$|A - PBQ^T|^2 = |A|^2 - \sum_{i=1}^{k}\beta_i^2 \tag{26}$$

To recast eq. (26) in terms of P and Q (in order to eliminate B), we observe that for an arbitrary matrix H we have $H_{ij} = \mathbf{e}_i^T H \mathbf{e}_j$, where $\mathbf{e}_i = (0, 0, \ldots, 1, \ldots, 0)^T$ are the standard basis vectors. Hence, using the fact that the

diagonal of a matrix is unchanged under transposition, we conclude that

$$\beta_i = \begin{cases} (P^T A Q)_{ii} & = & \mathbf{e}_i^T P^T A Q \, \mathbf{e}_i = \mathbf{p}_i^T A \mathbf{q}_i \\ (Q^T A^T P)_{ii} & = & \mathbf{e}_i^T Q^T A^T P \, \mathbf{e}_i = \mathbf{q}_i^T A^T \mathbf{p}_i \end{cases}$$

where $\mathbf{p}_i, \mathbf{q}_i$ are the i-th columns of P and Q, respectively. i.e. $P = [\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k]$ and $Q = [\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k]$. As a consequence,

$$\sum_{i=1}^{k} \beta_i^2 = \sum_{i=1}^{k} \mathbf{p}_i^T A \mathbf{q}_i \, \mathbf{q}_i^T A^T \mathbf{p}_i. \tag{27}$$

As a final step, we introduce the notation $L = D^T D$ to recast the regularisation term as:

$$|DP|^2 = \mathrm{Tr}(P^T L P) = \sum_{i=1}^{k} \mathbf{e}_i^T P^T L P \, \mathbf{e}_i = \sum_{i=1}^{k} \mathbf{p}_i^T L \, \mathbf{p}_i. \tag{28}$$

Plugging eqs. (27) and (28) into eq. (20), we obtain the following simplified form for the functional F (assuming that we eliminate B by using its optimal value):

$$F(P, Q) = |A|^2 + F_1(P, Q), \qquad \text{where} \qquad F_1(P, Q) = \sum_{i=1}^{k} \mathbf{p}_i^T (\lambda L - A \mathbf{q}_i \mathbf{q}_i^T A^T) \mathbf{p}_i. \tag{29}$$

Introducing the notation $S(\mathbf{q}) := \lambda L - A \mathbf{q} \mathbf{q}^T A^T$, we conclude that

$$F_1(P, Q) = \sum_{i=1}^{k} \mathbf{p}_i^T S(\mathbf{q}_i) \mathbf{p}_i.$$

Since each $S(\mathbf{q})$ is a symmetric matrix, it can be diagonalised with respect to an orthonormal basis, i.e. there is an orthogonal $n \times n$ matrix W (with $W^T W = W W^T = I_n$) and a diagonal matrix $\Lambda = diag(\lambda_1, \ldots, \lambda_n)$ (ordered $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n$), both depending on $\mathbf{q}$ such that

$$S(\mathbf{q}) = W(\mathbf{q}) \Lambda(\mathbf{q}) W(\mathbf{q})^T,$$

i.e. the columns of W are the eigenvectors of $S(\mathbf{q})$, with the corresponding eigenvalues on the diagonal of Λ. By introducing the notation $\lambda_1(S(\mathbf{q}))$ to denote the smallest eigenvalue of $\Lambda(\mathbf{q})$, we obtain the minimal value $\mathbf{p}_i^T S(\mathbf{q}_i) \mathbf{p}_i = \lambda_1(\mathbf{q}_i)$ when choosing $\mathbf{p}_i$ to be the (unit) eigenvector $(W_1(\mathbf{q}_i))$ corresponding to the smallest eigenvalue. As a consequence, the solution strategy boils down to steps in Algorithm 1.

This choice of P, Q and B solves the constrained minimisation problem (21). Notice that due to the fact that P and B matrices are determined after finding Q, this optimisation problem can essentially be translated into a search in the space of Q matrices. Some illustrative numerical experiments are available at [22].

■

We conclude this section by giving a slightly more general version ($\mu \neq 0$) of the previous theorem, thus re-establishing the symmetry between P and Q.

Theorem 6 (Regularised SVD, symmetric version). *Let A be an $n \times m$ matrix of rank $r \leq \min(n, m)$. For $k \leq r$, let $P \in \mathbb{R}^{n \times k}$ and $Q \in \mathbb{R}^{m \times k}$ of rank k, while $B \in \mathbb{R}^{k \times k}$ diagonal (i.e. $B = diag(\beta_1, \beta_2, \ldots, \beta_k)$). Furthermore, for arbitrary non-zero integers d and g we introduce regularisation matrices $D \in \mathbb{R}^{d \times n}$, and $G \in \mathbb{R}^{g \times m}$, as well as weights $\lambda, \mu \geq 0$. Finally, we introduce the short-hand notation $L := D^T D \in \mathbb{R}^{n \times n}$ and $M := G^T G \in \mathbb{R}^{m \times m}$ symmetric and positive-definite). We are now in a position to define the following functional F in the variables P, Q and B:*

$$F(P, Q, B) = |A - P B Q^T|^2 + \lambda |DP|^2 + \mu |GQ|^2 \tag{30}$$

and the corresponding constrained optimisation problem:

$$\min_{P,Q,B} F(P, Q, B) \qquad \textit{subject to:} \qquad Q^T Q = I_k, \quad |P_i| = 1, \ \forall i \in \{1, 2, \ldots, k\} \quad \textit{and } B \textit{ diagonal.} \tag{31}$$

This problem is solved by the solution specified in Algorithm 2.

Proof. Using the notation introduced above and in Theorem 5, we see that

$$|GQ|^2 = \mathrm{Tr}(Q^T M Q) = \sum_{i=1}^{k} \mathbf{q}_i^T M \mathbf{q}_i.$$

Hence, the functional (30) can be recast as:

$$F(P,Q) = |A|^2 + F_2(P,Q), \qquad \text{where} \qquad F_2(P,Q) = \sum_{i=1}^{k} \mathbf{p}_i^T (\lambda L - A\mathbf{q}_i\mathbf{q}_i^T A^T)\mathbf{p}_i + \mu \sum_{i=1}^{k} \mathbf{q}_i^T M \mathbf{q}_i. \tag{32}$$

The minimum of each term in the first summation in F_2 is equal to the smallest eigenvalue $\lambda_1(S(\mathbf{q}_i))$. Finding the minimum for the constrained optimisation problem (31) therefore amounts to finding the minimum of the functional:

$$\psi(Q) := \sum_{i=1}^{k} \left(\lambda_1(S(\mathbf{q}_i)) + \mu\, \mathbf{q}_i^T M \mathbf{q}_i\right) \tag{33}$$

subject to the constraint $Q^T Q = I_k$. Therefore, the minimisation problem again calls for a minimisation in Q space, as the optimal choice for P (corresponding eigen-vectors) follows automatically. We therefore arrive at the following Algorithm 2. Some illustrative numerical examples are available in [22].

■

Algorithm 2: Proposed RSVD method ($\mu \neq 0$)

Input: A, k, μ, λ, D, G
Output: P, B, Q
Initialization
while *no convergence* **do**

1. Recall that for any unit vector $\mathbf{q} \in \mathbb{R}^m$ we define $S(\mathbf{q}) = \lambda L - Aqq^T A^T$. Since this is a symmetric $n \times n$ matrix, it has a complete set of eigenvectors and corresponding eigenvalues. Denote the smallest eigenvalue of each $S(\mathbf{q}_i)$ as $\lambda_1(S(\mathbf{q}_i))$.
2. For a given $m \times k$ matrix $Q = [\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k]$ (with orthonormal columns: $Q^T Q = I_k$) compute the functional:
$$\psi(Q) := \sum_{i=1}^{k} \left(\lambda_1(S(\mathbf{q}_i)) + \mu\, \mathbf{q}_i^T M \mathbf{q}_i\right)$$
and use gradient descent (on the compact *torus domain*, see section 4) to find the minimum.
3. For each $\mathbf{q}_i$ as determined above, take $\mathbf{p}_i$ to be the eigenvector $W_1(\mathbf{q}_i)$ corresponding to the smallest eigenvector $\lambda_1(S(\mathbf{q}_i))$. Construct the $n \times k$ matrix $P = [\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_k]$.
4. Finally, set $B = diag(\beta_1, \ldots, \beta_n)$ where $\beta_i = (P^T A Q)_{ii}$.

end

4. Computational Aspects

4.1. Gradient Descent on the Unitary Domain

From Algorithm 2 it becomes clear that the full regularisation problem can be reduced to the simpler constrained minimisation problem detailed in eq. (33). Since the ψ-functional is smooth on a compact domain, this minimum is guaranteed to exist and one can use gradient descent to locate it. However, gradient descent needs to respect the constraint $Q^T Q = I_k$. This can be achieved by applying orthogonal transformations to the current Q matrix, as this will preserve orthonormality. Specifically, recall all orthogonal $m \times m$ matrices with determinant equal to 1 (rather than -1), constitute a multiplicative group denoted as $\mathcal{SO}(m)$ and formally defined as:

$$\mathcal{SO}(m) = \left\{ R \in \mathbb{R}^{m\times m} \mid RR^T = I_m = R^T R, \quad \text{and} \quad \det(R) = 1 \right\}$$

It is then straightforward to check that for any $R \in \mathcal{SO}(m)$, it holds that if $\bar{Q} = RQ$, the condition $Q^T Q = I_k$ implies that $\bar{Q}^T \bar{Q} = I_k$. It therefore follows that we can generate the "*infinitesimal variations*" needed to compute the gradient $\nabla_Q \psi(Q)$ by applying "sufficiently small" orthogonal matrices to the current value of Q. More precisely, we draw on the fact that $\mathcal{SO}(m)$ is actually a Lie-group [23] and that therefore each $R \in \mathcal{SO}(m)$ can be generated by exponentiating an element from its Lie-algebra $so(m) = \left\{ K \in \mathbb{R}^{m\times m} \mid K^T = -K \right\}$ (the skew-symmetric matrices):

$$R = \exp(tK) \equiv I_m + tK + \frac{1}{2!}t^2K^2 + \ldots + \frac{1}{n!}t^nK^n + \ldots \qquad (\text{with } K^T = -K)$$

By choosing t sufficiently small, one obtains an orthogonal transformation that is close to the identity I_m. Furthermore, it suffices to restrict the variations to orthogonal transformations that result from exponentiating a basis for the space of skew-symmetric matrices. Such a basis is provided by the $m(m-1)/2$ skew-symmetric matrices K_{ij} (where $1 \le i < j \le m$) for which the matrix element k, ℓ is given by:

$$K_{ij}(k,\ell) = \begin{cases} 1 & \text{if} \quad k = i,\ \ell = j \\ -1 & \text{if} \quad k = j,\ \ell = i \\ 0 & \text{otherwise} \end{cases}$$

Given the current value Q_0, we construct nearby values for Q by looping over $K_{12}, K_{13}, K_{23}, \ldots etc$ and constructing the corresponding orthogonal matrices $R_{12}(t) = \exp(tK_{12}), \ldots, etc$. Denoting these "*infinitesimal*" rotation matrices as R_α (where $\alpha = 1, \ldots, m(m-1)/2$), we see that the partial derivatives with respect to these rotations can be estimated as:

$$\frac{\partial \psi(Q)}{\partial R_\alpha} \approx \frac{\psi(R_\alpha(t)Q_0) - \psi(Q_0)}{t} \qquad (\text{for } t \text{ sufficiently small}).$$

From these results we can select the infinitesimal rotation that results in the steepest descent.

Since computing ψ is computationally expensive (it requires determining eigenvalues) a viable alternative to computing the gradient, is random descent: generate random rotations (by exponentiating random skew matrices) and check whether they result in a lower ψ-value. As soon as one is found, proceed in that direction, and repeat the process.

4.2. Illustrative example: Smoothing a noisy matrix

As common in the literature e.g., [11, 12, 24], we start from the assumption that the $n \times m$ data matrix A has a relatively smooth underlying structure that is corrupted by noise:

$$A = \mathbf{u}\mathbf{v}^T + \tau Z,$$

where the $n \times m$ matrix Z has independent standard normal entries, and τ controls the size of the noise.

To recover the underlying "signals" $\mathbf{u}$ and $\mathbf{v}$, we minimise the SVD-type regularisation functional (30) where the smoothness of the result is enforced by using regularisation matrices D and F that extract the second derivative, i.e.

$$D = F = \begin{bmatrix} -1 & 1 & 0 & & \cdots & & 0 \\ 1 & -2 & 1 & 0 & \cdots & & 0 \\ 0 & 1 & -2 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & -2 & 1 & \cdots & 0 \\ 0 & & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & & \cdots & 0 & 1 & -2 & 1 \\ 0 & & \cdots & & 0 & 1 & -1 \end{bmatrix}$$

A typical result for a rank-1 ($k = 1$) approximation is depicted in Figure 1, and compared to the standard SVD solution. This illustrative example is available in [25].

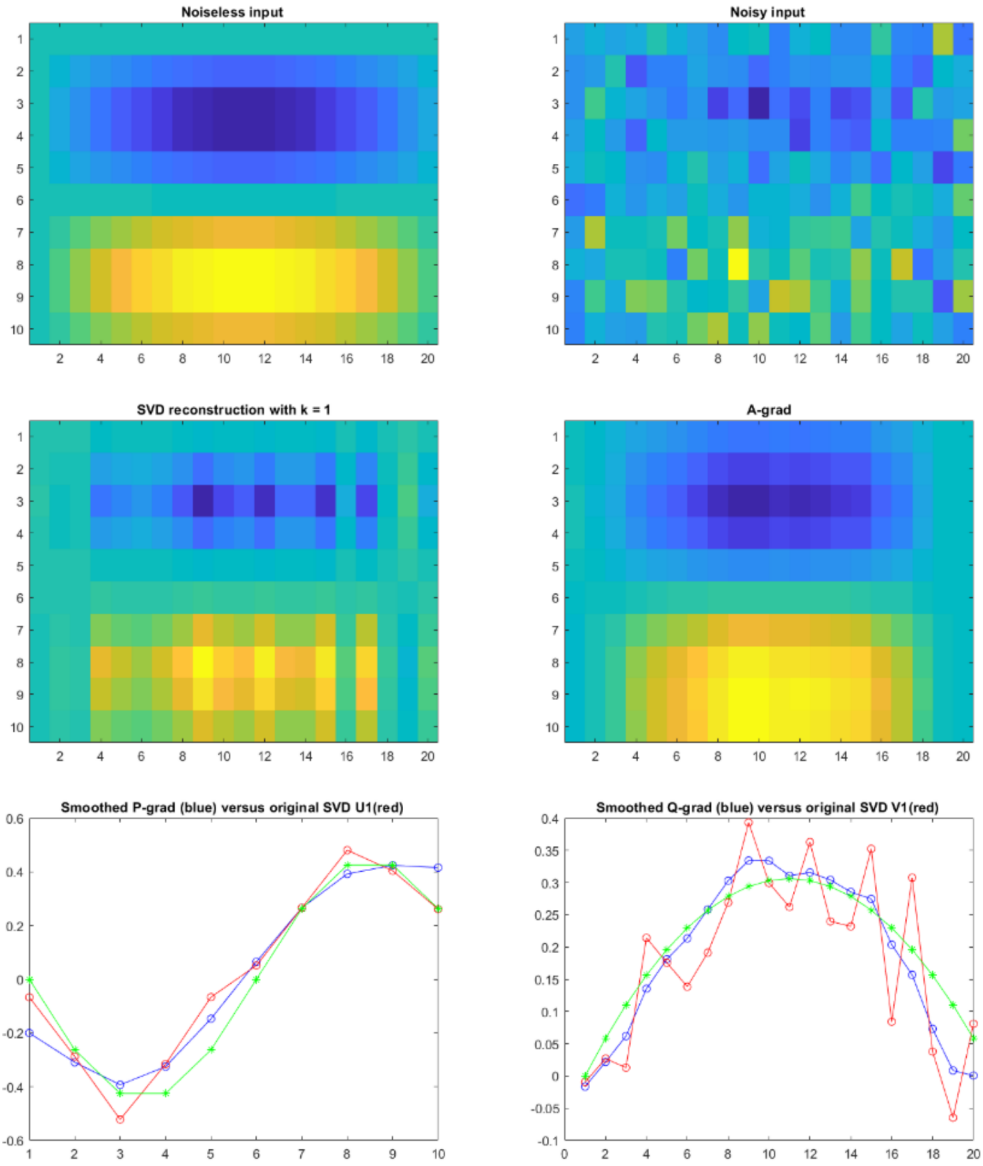

Figure 1: Reconstruction of noisy matrix based on RSVD. Top left: noise-less rank-1 matrix $\mathbf{u}\mathbf{v}^T$, (image) , top right: noisy input image $\mathbf{u}\mathbf{v}^T + \tau Z$ (high noise level), Middle left: standard rank-1 SVD reconstruction, middle right: RSVD reconstruction (D and F are 2nd deriv matrices. weight parameters $\lambda = \mu = 1.5$). Bottom: comparison of standard SVD $U(:,1)$ (red) versus P (blue), and $V(:,1)$ (red) (left) versus Q (blue) (right). The actual $\mathbf{u}$ and $\mathbf{v}$ for the noiseless input signal are drawn in green.

5. Conclusions and Future Research

Singular Value Decomposition (SVD) and Principal Component Analysis (PCA) are important matrix factorisation techniques that underpin numerous applications. However, it is well-known that disturbances in the input (noise, outliers or missing values) have a significant effect on the outcome. For that reason we investigate regularisation in two different but related versions of the factorisation, and detail the solution algorithms.

An important topic for further research would be to find ways in which the gradient descent procedure in Algorithms 1 and 2 can be accelerated by taking advantage of the fact that the functional is very smooth and locally approximately quadratic. It would also be useful to derive some estimates for appropriate values for the weights λ and μ in terms of noise characteristics corrupting the underlying signal. Finally, although the P matrix in algorithm 2 has unit-length columns, we were not able to prove that these columns are also orthogonal ($P^T P = I$) as is the case in standard SVD. In fact, numerical experiments seem to indicate that such a constraint is not compatible with minimisation of the functional. This requires further theoretical elucidation.

Acknowledgment

The authors gratefully acknowledge partial support by the Dutch NWO ESI-Bida project NEAT (647.003.002).

References

[1] M. A. Davenport, J. Romberg, An overview of low-rank matrix recovery from incomplete observations, IEEE Journal of Selected Topics in Signal Processing 10 (4) (2016) 608–622. doi:10.1109/jstsp.2016.2539100.
URL http://dx.doi.org/10.1109/JSTSP.2016.2539100

[2] I. Tošić, P. Frossard, Dictionary learning, IEEE Signal Processing Magazine 28 (2) (2011) 27–38.

[3] A. Khoshrou, E. J. Pauwels, Data-driven pattern identification and outlier detection in time series, in: Science and Information Conference, Springer, 2018, pp. 471–484.

[4] S. Gunasekar, B. Woodworth, S. Bhojanapalli, B. Neyshabur, N. Srebro, Implicit regularization in matrix factorization (2017). arXiv:1705.09280.

[5] J. P. Brooks, J. H. Dulá, E. L. Boone, A pure l1-norm principal component analysis, Computational statistics & data analysis 61 (2013) 83–98.

[6] N. Kwak, Principal component analysis by $L\{p\}$-norm maximization, IEEE Transactions on Cybernetics 44 (5) (2013) 594–609.

[7] E. J. Candès, X. Li, Y. Ma, J. Wright, Robust principal component analysis?, Journal of the ACM (JACM) 58 (3) (2011) 1–37.

[8] Z. Zhou, X. Li, J. Wright, E. Candes, Y. Ma, Stable principal component pursuit, in: 2010 IEEE international symposium on information theory, IEEE, 2010, pp. 1518–1522.

[9] H. Shen, J. Z. Huang, Sparse principal component analysis via regularized low rank matrix approximation, Journal of multivariate analysis 99 (6) (2008) 1015–1034.

[10] B. Dumitrescu, P. Irofti, Regularized k-svd, IEEE Signal Processing Letters 24 (3) (2017) 309–313.

[11] T. Jin, J. Yu, J. You, K. Zeng, C. Li, Z. Yu, Low-rank matrix factorization with multiple hypergraph regularizer, Pattern Recognition 48 (3) (2015) 1011–1022.

[12] J. He, Y. Bi, B. Liu, Z. Zeng, Graph-dual laplacian principal component analysis, Journal of Ambient Intelligence and Humanized Computing 10 (8) (2019) 3249–3262.

[13] M. Yin, J. Gao, Z. Lin, Q. Shi, Y. Guo, Dual graph regularized latent low-rank representation for subspace clustering, IEEE Transactions on Image Processing 24 (12) (2015) 4918–4933.

[14] N. Shahid, N. Perraudin, V. Kalofolias, G. Puy, P. Vandergheynst, Fast robust pca on graphs, IEEE Journal of Selected Topics in Signal Processing 10 (4) (2016) 740–756.

[15] G. Strang, Introduction to linear algebra, Vol. 3, Wellesley-Cambridge Press Wellesley, MA, 1993.

[16] R. Horn, C. Johnson, Matrix Analysis, Cambridge University Press, 1985.

[17] G. H. Golub, C. F. Van Loan, Matrix computations, Vol. 3, JHU press, 2013.

[18] C. Eckart, G. Young, The approximation of one matrix by another of lower rank, Psychometrika 1 (3) (1936) 211–218.

[19] code for theorem 4, https://www.dropbox.com/s/fjtz7fh1hyf9cdm/theorem_4.m?dl=0, created: 2021, June.

[20] code for special case:$\mu = 0$ and $\lambda = 0$, https://www.dropbox.com/s/ngjksurfepn8dml/special_case_mu_0_lambda_0.m?dl=0, created: 2021, June.

[21] code for special case: $\mu = 0$ and $d = i_n$, https://www.dropbox.com/s/ab1rfiquiyuzuvz/special_case_mu_0_D_In.m?dl=0, created: 2021, June.

[22] code for factorisation svd-type theorems, https://www.dropbox.com/sh/f257tzsuttbp1ro/AABaJc1IVXZFQFVQKnpIGjr7a?dl=0, created: 2021, June.

[23] A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, A. Zanna, Lie-group methods, Acta numerica 9 (2000) 215–365.

[24] M. Gavish, D. L. Donoho, The optimal hard threshold for singular values is $4/\sqrt{3}$, IEEE Transactions on Information Theory 60 (8) (2014) 5040–5053.

[25] code for the numerical experiments section, https://www.dropbox.com/sh/tcl7lag80cimibw/AAD3QNx8FSTOX-c3-wtAh-UFa?dl=0, created: 2021, June.