

LOW-RANK MATRIX COMPLETION BY RIEMANNIAN OPTIMIZATION*

BART VANDEREYCKEN†

Abstract. The matrix completion problem consists of finding or approximating a low-rank matrix based on a few samples of this matrix. We propose a new algorithm for matrix completion that minimizes the least-square distance on the sampling set over the Riemannian manifold of fixed-rank matrices. The algorithm is an adaptation of classical nonlinear conjugate gradients, developed within the framework of retraction-based optimization on manifolds. We describe all the necessary objects from differential geometry necessary to perform optimization over this low-rank matrix manifold, seen as a submanifold embedded in the space of matrices. In particular, we describe how metric projection can be used as retraction and how vector transport lets us obtain the conjugate search directions. Finally, we prove convergence of a regularized version of our algorithm under the assumption that the restricted isometry property holds for incoherent matrices throughout the iterations. The numerical experiments indicate that our approach scales very well for large-scale problems and compares favorably with the state-of-the-art, while outperforming most existing solvers.

Key words. matrix completion, low-rank matrices, optimization on manifolds, differential geometry, nonlinear conjugate gradients, Riemannian manifolds, Newton

AMS subject classifications. 15A83, 65K05, 53B21

DOI. 10.1137/110845768

1. Introduction. Let $A \in \mathbb{R}^{m \times n}$ be an $m \times n$ matrix that is only known on a subset Ω of the complete set of entries $\{1, \dots, m\} \times \{1, \dots, n\}$. The low-rank matrix completion problem [16] consists of finding the matrix with lowest rank that agrees with A on Ω :

$$(1.1) \quad \begin{aligned} & \underset{X}{\text{minimize}} && \text{rank}(X), \\ & \text{subject to} && X \in \mathbb{R}^{m \times n}, P_{\Omega}(X) = P_{\Omega}(A), \end{aligned}$$

where

$$(1.2) \quad P_{\Omega}: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}, X_{i,j} \mapsto \begin{cases} X_{i,j} & \text{if } (i, j) \in \Omega, \\ 0 & \text{if } (i, j) \notin \Omega \end{cases}$$

denotes projection onto Ω . Without loss of generality, we assume $m \leq n$.

Due to the presence of noise, it is advisable to relax the equality constraint in (1.1) to allow for misfit. Then, given a tolerance $\varepsilon \geq 0$, a more robust version of the rank minimization problem becomes

$$(1.3) \quad \begin{aligned} & \underset{X}{\text{minimize}} && \text{rank}(X), \\ & \text{subject to} && X \in \mathbb{R}^{m \times n}, \|P_{\Omega}(X) - P_{\Omega}(A)\|_F \leq \varepsilon, \end{aligned}$$

where $\|X\|_F$ denotes the Frobenius norm of X .

*Received by the editors August 25, 2011; accepted for publication (in revised form) March 12, 2013; published electronically June 20, 2013.

<http://www.siam.org/journals/siopt/23-2/84576.html>

†Chair of Numerical Algorithms and HPC, MATHICSE, École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland (bartv@math.princeton.edu).

Matrix completion has a number of interesting applications such as collaborative filtering, system identification, and global positioning, but unfortunately it is NP hard. Recently, there has been a considerable body of work devoted to the identification of large classes of matrices for which (1.1) has a unique solution that can be recovered in polynomial time. In [15], for example, the authors show that when Ω is sampled uniformly at random, the nuclear norm relaxation

$$(1.4) \quad \begin{aligned} & \underset{X}{\text{minimize}} && \|X\|_*, \\ & \text{subject to} && X \in \mathbb{R}^{m \times n}, P_\Omega(X) = P_\Omega(A), \end{aligned}$$

can recover with high probability any matrix A of rank k that has so-called low incoherence, provided that the number of samples is large enough, $|\Omega| > Cnk \text{ polylog}(n)$. Related work has been done by [14, 16, 27, 28], which in particular also establish similar recovery results for the robust formulation (1.3).

Many of the potential applications for matrix completion involve very large data sets; the Netflix matrix, for example, has more than 10^8 entries [6]. It is therefore crucial to develop algorithms that can cope with such a large-scale setting, but, unfortunately, solving (1.4) by off-the-shelf methods for convex optimization scales very badly in the matrix dimension. This has spurred a considerable amount of algorithms that aim to solve the nuclear norm relaxation by specifically designed methods that try to exploit the low-rank structure of the solution; see, e.g., [13, 21, 35, 36, 37, 38, 50]. Other approaches include optimization on the Grassmann manifold [5, 18, 27]; atomic decompositions [33]; and nonlinear SOR [53].

1.1. The proposed method: Optimization on manifolds. We present a new method for low-rank matrix completion based on a direct optimization over the set of all fixed-rank matrices. By prescribing the rank of the global minimizer of (1.3), say k , the robust matrix completion problem is equivalent to

$$(1.5) \quad \begin{aligned} & \underset{X}{\text{minimize}} && f(X) := \frac{1}{2} \|P_\Omega(X - A)\|_F^2, \\ & \text{subject to} && X \in \mathcal{M}_k := \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) = k\}. \end{aligned}$$

It is well known that \mathcal{M}_k is a smooth (C^∞) manifold; it can, for example, be identified as the smooth part of the determinantal variety of matrices of rank at most k [10, Proposition 1.1]. Since the objective function f is also smooth, problem (1.5) is a smooth optimization problem that can be solved by methods from Riemannian optimization as introduced, among others, by [2, 4, 20]. Simply put, Riemannian optimization is the generalization of standard unconstrained optimization, where the search space is \mathbb{R}^n , to optimization of a smooth objective function on a Riemannian manifold.

For solving the optimization problem (1.5), in principle any method from optimization on Riemannian manifolds could be used. In this paper, we use a generalization of classical nonlinear conjugate gradients (CGs) on Euclidean space to perform optimization on manifolds; see, e.g., [2, 20, 49]. The reason for this choice compared to, say, Newton’s method was that CG performed best in our numerical experiments. The skeleton of the proposed method, LRGeomCG, is listed in Algorithm 1.

Algorithm 1 is derived using concepts from differential geometry, yet it closely resembles a typical nonlinear CG algorithm with Armijo line-search for unconstrained optimization. It is schematically visualized in Figure 1.1 for iteration number i and relies on the following crucial ingredients, which will be explained in more detail for (1.5) in section 2.

 ALGORITHM 1. LRGEOMCG: GEOMETRIC CG FOR (1.5).

Require: initial iterate $X_1 \in \mathcal{M}_k$, tolerance $\tau > 0$, tangent vector $\eta_0 = 0$

- 1: **for** $i = 1, 2, \dots$ **do**
 - 2: Compute the gradient
 $\xi_i := \text{grad } f(X_i)$ # see Algorithm 2
 - 3: Check convergence
 if $\|\xi_i\| \leq \tau$, then break
 - 4: Compute a conjugate direction by PR+
 $\eta_i := -\xi_i + \beta_i \mathcal{T}_{X_{i-1} \rightarrow X_i}(\eta_{i-1})$ # see Algorithm 4
 - 5: Compute an initial step t_i in closed-form from a straight-line search
 $t_i = \arg \min_t f(X_i + t \eta_i)$ # see Algorithm 5
 - 6: Perform Armijo backtracking to find the smallest integer $m \geq 0$ such that
 $f(X_i) - f(R_{X_i}(0.5^m t_i \eta_i)) \geq -0.0001 \times 0.5^m t_i \langle \xi_i, \eta_i \rangle$
 and obtain the new iterate
 $X_{i+1} := R_{X_i}(0.5^m t_i \eta_i)$ # see Algorithm 6
 - 7: **end for**
-




FIG. 1.1. Visualization of Algorithm 1: nonlinear CG on a Riemannian manifold.

1. The Riemannian gradient, denoted $\text{grad } f(X_i)$, is a specific tangent vector ξ_i which corresponds to the direction of steepest ascent of $f(X_i)$ but is restricted to only directions in the tangent space $T_{X_i} \mathcal{M}_k$.
2. The search direction $\eta_i \in T_{X_i} \mathcal{M}_k$ is conjugate to the gradient and is computed by a variant of the classical Polak–Ribière updating rule in nonlinear CG. This requires taking a linear combination of the Riemannian gradient with the previous search direction η_{i-1} . Since η_{i-1} does not lie in $T_{X_i} \mathcal{M}_k$, it needs to be transported to $T_{X_i} \mathcal{M}_k$. This is done by a mapping $\mathcal{T}_{X_{i-1} \rightarrow X_i} : T_{X_{i-1}} \mathcal{M}_k \rightarrow T_{X_i} \mathcal{M}_k$, the so-called vector transport.
3. As a tangent vector only gives a direction but not the line search itself on the manifold, a smooth mapping $R_{X_i} : T_{X_i} \mathcal{M}_k \rightarrow \mathcal{M}_k$, called the *retraction*, is needed to map tangent vectors to the manifold. Using the conjugate direction η_i , a line search can then be performed along the curve $t \mapsto R_{X_i}(t \eta_i)$. Step 6 uses a standard backtracking procedure where we have chosen to fix the constants. More judiciously chosen constants can sometimes improve the line search, but our numerical experiments indicate that this is not necessary in the setting under consideration.

1.2. Relation to existing manifold-related methods. At the time of submitting the present work, a large number of other matrix completion solvers based on Riemannian optimization have been proposed in [8, 9, 19, 27, 28, 40, 41, 43, 48]. Like the current paper, all of these algorithms use the concept of retraction-based optimization on the manifold of fixed-rank matrices, but they differ in their specific choice of Riemannian manifold structure and metric. It remains a topic of further investigation to assess the performance of these different geometries with respect to each other and to nonmanifold based solvers.

A first attempt at such a comparison has been very recently done in [44], where the previous geometries were tested to complete one large matrix. The authors reach the conclusion that for gradient-based algorithms there exists a set of geometries—including ours—that perform remarkably more or less the same with respect to the total computational time, and that overall most geometries outperformed the state-of-the-art. In addition, Newton-based algorithms performed well when high precision was required and now the algorithm based on our embedded submanifold geometry was clearly faster.

1.3. Outline of the paper. The plan of the paper is as follows. In the next section, the necessary concepts of differential geometry are explained to turn Algorithm 1 into a concrete method. The implementation of each step is explained in section 3. We also prove convergence of a slightly modified version of this method in section 4 under some assumptions which are reasonable for matrix completion. Numerical experiments and comparisons to the state-of-the-art are carried out in section 5. The last section is devoted to conclusions.

2. Differential geometry for low-rank matrix manifolds. In this section, we explain the differential geometry concepts used in Algorithm 1 applied to our particular matrix completion problem (1.5).

2.1. The Riemannian manifold. Let

$$\mathcal{M}_k = \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) = k\}$$

denote the manifold of fixed-rank matrices. Using the SVD, one has the equivalent characterization

$$(2.1) \quad \mathcal{M}_k = \{U\Sigma V^T : U \in \text{St}_k^m, V \in \text{St}_k^n, \Sigma = \text{diag}(\sigma_i), \sigma_1 \geq \dots \geq \sigma_k > 0\},$$

where St_k^m is the Stiefel manifold of $m \times k$ real, orthonormal matrices, and $\text{diag}(\sigma_i)$ denotes a diagonal matrix with σ_i on the diagonal. Whenever we use the notation $U\Sigma V^T$ in the rest of the paper, we mean matrices that satisfy (2.1). Furthermore, the constants k, m, n are always used to denote dimensions of matrices, and, for simplicity, we assume $1 \leq k < m \leq n$.

The following proposition shows that \mathcal{M}_k is indeed a smooth manifold. While the existence of such a smooth manifold structure, together with its tangent space, is well known (see, e.g., [25, 29] for applications of gradient flows on \mathcal{M}_k), more advanced concepts like retraction-based optimization on \mathcal{M}_k have only very recently been investigated; see [47, 48]. In contrast, the case of optimization on *symmetric* fixed-rank matrices has been studied in more detail in [26, 39, 45, 51].

PROPOSITION 2.1. *The set \mathcal{M}_k is a smooth submanifold of dimension $(m+n-k)k$ embedded in $\mathbb{R}^{m \times n}$. Its tangent space $T_X \mathcal{M}_k$ at $X = U \Sigma V^T \in \mathcal{M}_k$ is given by*

$$(2.2) \quad T_X \mathcal{M}_k = \left\{ [U \quad U_\perp] \begin{bmatrix} \mathbb{R}^{k \times k} & \mathbb{R}^{k \times (n-k)} \\ \mathbb{R}^{(m-k) \times k} & \mathbf{0}_{(m-k) \times (n-k)} \end{bmatrix} [V \quad V_\perp]^T \right\}$$

$$(2.3) \quad = \{UMV^T + U_p V^T + UV_p^T : M \in \mathbb{R}^{k \times k}, \\ U_p \in \mathbb{R}^{m \times k}, U_p^T U = 0, V_p \in \mathbb{R}^{n \times k}, V_p^T V = 0\}.$$

Proof. See [32, Example 8.14] for a proof that uses only elementary differential geometry based on local submersions. The tangent space is obtained by the first-order perturbation of the SVD and counting dimensions. \square

The *tangent bundle* is defined as the disjoint union of all tangent spaces,

$$T\mathcal{M}_k := \bigcup_{X \in \mathcal{M}_k} \{X\} \times T_X \mathcal{M}_k = \{(X, \xi) \in \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n} : X \in \mathcal{M}_k, \xi \in T_X \mathcal{M}_k\}.$$

By restricting the Euclidean inner product on $\mathbb{R}^{m \times n}$,

$$\langle A, B \rangle = \text{tr}(A^T B) \quad \text{with } A, B \in \mathbb{R}^{m \times n},$$

to the tangent bundle, we turn \mathcal{M}_k into a Riemannian manifold with Riemannian metric

$$(2.4) \quad g_X(\xi, \eta) := \langle \xi, \eta \rangle = \text{tr}(\xi^T \eta) \quad \text{with } X \in \mathcal{M}_k \text{ and } \xi, \eta \in T_X \mathcal{M}_k,$$

where the tangent vectors ξ, η are seen as matrices in $\mathbb{R}^{m \times n}$.

Once the metric is fixed, the notion of the gradient of an objective function can be introduced. For a Riemannian manifold, the Riemannian gradient of a smooth function $f : \mathcal{M}_k \rightarrow \mathbb{R}$ at $X \in \mathcal{M}_k$ is defined as the unique tangent vector $\text{grad } f(X)$ in $T_X \mathcal{M}_k$ such that

$$\langle \text{grad } f(X), \xi \rangle = Df(X)[\xi] \quad \text{for all } \xi \in T_X \mathcal{M}_k,$$

where we denoted directional derivatives by Df . Since \mathcal{M}_k is embedded in $\mathbb{R}^{m \times n}$, the Riemannian gradient is given as the orthogonal projection onto the tangent space of the gradient of f seen as a function on $\mathbb{R}^{m \times n}$; see, e.g., [2, equation (3.37)]. Defining $P_U := UU^T$ and $P_U^\perp := I - P_U$ for any $U \in \text{St}_k^m$, we denote the orthogonal projection onto the tangent space at X as

$$(2.5) \quad P_{T_X \mathcal{M}_k} : \mathbb{R}^{m \times n} \rightarrow T_X \mathcal{M}_k, Z \mapsto P_U Z P_V + P_U^\perp Z P_V + P_U Z P_V^\perp.$$

Then, using $P_\Omega(X - A)$ as the (Euclidean) gradient of $f(X) = \|P_\Omega(X - A)\|_F^2/2$, we obtain

$$(2.6) \quad \text{grad } f(X) := P_{T_X \mathcal{M}_k}(P_\Omega(X - A)).$$

2.2. Metric projection as retraction. As explained in section 1.1, we need a so-called retraction mapping to go back from an element in the tangent space to the manifold. Retractions are essentially first-order approximations of the exponential map of the manifold; see, e.g., [3, Definition 1]. In our setting, we have chosen metric projection as a retraction:

$$(2.7) \quad R_X : \mathcal{U}_X \rightarrow \mathcal{M}_k, \xi \mapsto P_{\mathcal{M}_k}(X + \xi) := \arg \min_{Z \in \mathcal{M}_k} \|X + \xi - Z\|_F,$$

where $\mathcal{U}_X \subset T_X \mathcal{M}_k$ is a suitable neighborhood around zero and $P_{\mathcal{M}_k}$ is the orthogonal projection onto \mathcal{M}_k . Based on [34, Lemma 2.1], it can be easily shown that R_X indeed satisfies the conditions of a retraction, which allows us to invoke the convergence proofs of [2].

Mapping R_X can be computed in closed-form by the SVD: Let $X \in \mathcal{M}_k$ be given, and then for sufficiently small $\xi \in \mathcal{U}_X \subset T_X \mathcal{M}_k$, we have

$$(2.8) \quad R_X(\xi) = P_{\mathcal{M}_k}(X + \xi) = \sum_{i=1}^k \sigma_i u_i v_i^T,$$

where σ_i, u_i, v_i are the (ordered) singular values and vectors of the SVD of $X + \xi$. It is obvious that when $\sigma_k = 0$, there is no best rank- k solution for (2.7), and additionally, when $\sigma_{k-1} = \sigma_k$ the minimization in (2.7) does not have a unique solution. In fact, since the manifold is nonconvex, metric projections can never be well defined on the whole tangent bundle. Fortunately a retraction has to be defined only locally because this is sufficient to establish convergence of the Riemannian algorithms.

2.3. Riemannian Newton on \mathcal{M}_k . Although we do not exploit second-order information in Algorithm 1, Newton or its variants may be preferable for some applications. In [44], for example, it was shown that the second-order Riemannian trust-region algorithm of [1] performs very well in combination with our embedded submanifold geometry. We will therefore give the following theorem with an explicit expression of the Riemannian Hessian of f . Its proof is a straightforward but technical analogue of that in [51] for symmetric fixed-rank matrices. It can be found in Appendix A of the extended technical report [52] of this paper.

PROPOSITION 2.2. *For any $X = U\Sigma V^T \in \mathcal{M}_k$ satisfying (2.1) and $\xi \in T_X \mathcal{M}_k$ satisfying (2.3), the Riemannian Hessian of f at X in the direction of ξ satisfies*

$$\begin{aligned} \text{Hess } f(X)[\xi] &= P_U P_\Omega(\xi) P_V + P_U^\perp [P_\Omega(\xi) + P_\Omega(X - A) V_p \Sigma^{-1} V^T] P_V \\ &\quad + P_U (P_\Omega(\xi) + U \Sigma^{-1} U_p^T P_\Omega(X - A)) P_V^\perp. \end{aligned}$$

2.4. Vector transport. Vector transport was introduced in [2] and [46] as a means to transport tangent vectors from one tangent space to another. In a similar way as retractions are approximations of the exponential mapping, vector transport is the first-order approximation of parallel transport, another important concept in differential geometry. See Figure 2.1 for an illustration of the retraction $\mathcal{T}_{X \rightarrow Y} : T_X \mathcal{M}_k \rightarrow T_Y \mathcal{M}_k$, and we refer to [46, Definition 2] for an exact definition.




FIG. 2.1. Vector transport $\mathcal{T}_{X \rightarrow Y}$ on a Riemannian manifold.

Since \mathcal{M}_k is an embedded submanifold of $\mathbb{R}^{m \times n}$, orthogonally projecting the translated tangent vector in $\mathbb{R}^{m \times n}$ onto the new tangent space constitutes a vector transport; see [2, section 8.1.2]. In other words, we define

$$(2.9) \quad \mathcal{T}_{X \rightarrow Y} : T_X \mathcal{M}_k \rightarrow T_Y \mathcal{M}_k, \xi \mapsto P_{T_Y \mathcal{M}_k}(\xi)$$

with $P_{T_Y \mathcal{M}_k}$ as defined in (2.5).

As explained in section 1.1, the conjugate search direction η_i in Algorithm 1 is computed as a linear combination of the gradient and the previous direction:

$$\eta_i = -\text{grad } f(X_i) + \beta_i \mathcal{T}_{X_{i-1} \rightarrow X_i}(\eta_{i-1}),$$

where we transported η_{i-1} . For β_i , we have chosen the geometrical variant of Polak–Ribière (PR+), as introduced in [2, Chapter 8.2]. Again using vector transport, this becomes

$$(2.10) \quad \beta_i = \frac{\langle \text{grad } f(X_i), \text{grad } f(X_i) - \mathcal{T}_{X_{i-1} \rightarrow X_i}(\text{grad } f(X_{i-1})) \rangle}{\langle \text{grad } f(X_{i-1}), \text{grad } f(X_{i-1}) \rangle}.$$

In order to prove convergence, we also enforce that the search direction η_i is sufficiently gradient-related in the sense that its angle with the gradient is never too small; see, e.g., [7, Chapter 1.2].

3. Implementation details. This section is devoted to the implementation of Algorithm 1. For most operations, we will also provide a flop count for the low-rank regime, i.e., $k \ll m \leq n$.

Low-rank matrices. Since every element $X \in \mathcal{M}_k$ is a rank k matrix, we store it as the result of a compact SVD:

$$X = U \Sigma V^T$$

with orthonormal matrices $U \in \text{St}_k^m$ and $V \in \text{St}_k^n$ and a diagonal matrix $\Sigma \in \mathbb{R}^{k \times k}$ with decreasing positive entries on the diagonal. For the computation of the objective function and the gradient, we also precompute the sparse matrix $X_\Omega := P_\Omega(X)$. As explained in the next paragraph, computing X_Ω costs $(m + 2|\Omega|)k$ flops.

Projection operator P_Ω . During the course of Algorithm 1, we require the application of P_Ω , as defined in (1.2), to certain low-rank matrices. By exploiting the low-rank form, this can be done efficiently as follows: Let Z be a rank- k_Z matrix with factorization $Z := Y_1 Y_2^T$. (k_Z is possibly different from k , the rank of the matrices in \mathcal{M}_k .) Define $Z_\Omega := P_\Omega(Z)$. Then element (i, j) of Z_Ω is given by

$$(3.1) \quad (Z_\Omega)_{i,j} = \begin{cases} \sum_{l=1}^{k_Z} Y_1(i, l) Y_2(j, l) & \text{if } (i, j) \in \Omega, \\ 0 & \text{if } (i, j) \notin \Omega. \end{cases}$$

The cost for computing Z_Ω is $2|\Omega|k_Z$ flops.

Tangent vectors. A tangent vector $\eta \in T_X \mathcal{M}_k$ at $X = U \Sigma V^T \in \mathcal{M}_k$ will be represented as

$$(3.2) \quad \eta = U M V^T + U_p V^T + U V_p^T,$$

where $M \in \mathbb{R}^{k \times k}$, $U_p \in \mathbb{R}^{m \times k}$ with $U^T U_p = 0$, and $V_p \in \mathbb{R}^{n \times k}$ with $V^T V_p = 0$. After vectorizing this representation, the inner product $\langle \eta, \nu \rangle$ can be computed in $2(m+n)k + 2k^2$ flops.

Since X is available as an SVD, the orthogonal projection onto the tangent space $T_X \mathcal{M}_k$ becomes

$$P_{T_X \mathcal{M}_k}(Z) := P_U Z P_V + (Z - P_U Z) P_V + P_U (Z^T - P_V Z^T)^T,$$

where $P_U := UU^T$ and $P_V := VV^T$. If Z can be efficiently applied to a vector, evaluating and storing the result of $P_{T_X \mathcal{M}_k}(Z)$ as a tangent vector in the format (3.2) can also be performed efficiently.

Riemannian gradient. From (2.6), the Riemannian gradient at $X = U\Sigma V^T$ is the orthogonal projection of $P_\Omega(X - A)$ onto the tangent space at X . Since the sparse matrix $A_\Omega = P_\Omega(A)$ is given and $X_\Omega = P_\Omega(X)$ was already precomputed, the gradient can be computed by Algorithm 2. The total cost is $2(n + 2m)k^2 + 4|\Omega|k$ flops.

ALGORITHM 2. CALCULATE GRADIENT $\text{grad} f(X)$.

Require: matrix $X = U\Sigma V^T \in \mathcal{M}_k$, sparse matrix $R = X_\Omega - A_\Omega \in \mathbb{R}^{m \times n}$

Ensure: $\text{grad} f(X) = U M V^T + U_p V^T + U V_p^T \in T_X \mathcal{M}_k$

- 1: $R_u \leftarrow R^T U, \quad R_v \leftarrow R V$ # $4|\Omega|k$ flops
 - 2: $M \leftarrow U^T R_v$ # $2mk^2$ flops
 - 3: $U_p \leftarrow R_v - U M, \quad V_p \leftarrow R_u - V M^T$ # $2(m + n)k^2$ flops
-

Vector transport. Let $\nu = U M V^T + U_p V^T + U V_p^T$ be a tangent vector at $X = U\Sigma V^T \in \mathcal{M}_k$. By (2.9), the vector

$$\nu_+ := \mathcal{T}_{X \rightarrow X_+}(\nu) = P_{T_{X_+} \mathcal{M}_k}(\nu)$$

is the transport of ν to the tangent space at some $X_+ = U_+ \Sigma_+ V_+^T$. It is computed by Algorithm 3 with a total cost of approximately $14(m + n)k^2 + 10k^3$ flops.

ALGORITHM 3. CALCULATE VECTOR TRANSPORT $\mathcal{T}_{X \rightarrow X_+}(\nu)$.

Require: matrices $X = U\Sigma V^T \in \mathcal{M}_k$ and $X_+ = U_+ \Sigma_+ V_+^T \in \mathcal{M}_k$,

tangent vector $\nu = U M V^T + U_p V^T + U V_p^T$

Ensure: $\mathcal{T}_{X \rightarrow X_+}(\nu) = U_+ M_+ V_+^T + U_{p_+} V_+^T + U_+ V_{p_+}^T \in T_{X_+} \mathcal{M}_k$

- 1: $A_v \leftarrow V^T V_+, \quad A_u \leftarrow U^T U_+$ # $2(m + n)k^2$
 - 2: $B_v \leftarrow V_p^T V_+, \quad B_u \leftarrow U_p^T U_+$ # $2(m + n)k^2$
 - 3: $M_+^{(1)} \leftarrow A_u^T M A_v, \quad U_+^{(1)} \leftarrow U(M A_v), \quad V_+^{(1)} \leftarrow V(M^T A_u)$ # $6k^3 + 2(m + n)k^2$
 - 4: $M_+^{(2)} \leftarrow B_u^T A_v, \quad U_+^{(2)} \leftarrow U_p A_v, \quad V_+^{(2)} \leftarrow V B_u$ # $2k^3 + 2(m + n)k^2$
 - 5: $M_+^{(3)} \leftarrow A_u^T B_v, \quad U_+^{(3)} \leftarrow U B_v, \quad V_+^{(3)} \leftarrow V_p A_u$ # $2k^3 + 2(m + n)k^2$
 - 6: $M_+ \leftarrow M_+^{(1)} + M_+^{(2)} + M_+^{(3)}$ # $2k^2$
 - 7: $U_{p_+} \leftarrow U_+^{(1)} + U_+^{(2)} + U_+^{(3)}, \quad U_{p_+} \leftarrow U_{p_+} - U_+(U_+^T U_{p_+})$ # $4mk^2$
 - 8: $V_{p_+} \leftarrow V_+^{(1)} + V_+^{(2)} + V_+^{(3)}, \quad V_{p_+} \leftarrow V_{p_+} - V_+(V_+^T V_{p_+})$ # $4nk^2$
-

Nonlinear CG. The geometric variant of Polak–Ribière is implemented as Algorithm 4 costing about $28(m + n)k^2 + 20k^3$ flops. In order to improve robustness, we use the PR+ variant and restart when the conjugate direction is almost orthogonal to the gradient.

Initial guess for line search. Although Algorithm 1 uses line search, a good initial guess can greatly enhance performance. We observed in our numerical experiments that an exact minimization on the tangent space alone (so, neglecting the retraction),

$$(3.3) \quad \min_t f(X + t\eta) = \frac{1}{2} \min_t \|P_\Omega(X) + tP_\Omega(\eta) - P_\Omega(A)\|_F^2,$$

ALGORITHM 4. COMPUTE THE CONJUGATE DIRECTION BY PR+.

Require: previous iterate X_{i-1} , previous gradient ξ_{i-1} , previous direction η_{i-1}
 current iterate X_i , current gradient ξ_i

Ensure: conjugate direction $\eta_i \in T_{X_i}\mathcal{M}_k$

- 1: Transport previous gradient and direction to current tangent space:
 $\bar{\xi}_i \leftarrow \mathcal{T}_{X_{i-1} \rightarrow X_i}(\xi_{i-1})$ # apply Algorithm 3
 $\bar{\eta}_i \leftarrow \mathcal{T}_{X_{i-1} \rightarrow X_i}(\eta_{i-1})$ # apply Algorithm 3
 - 2: Compute conjugate direction:
 $\delta_i \leftarrow \xi_i - \bar{\xi}_i$
 $\beta \leftarrow \max(0, \langle \delta_i, \xi_i \rangle / \langle \xi_{i-1}, \xi_{i-1} \rangle)$
 $\eta_i \leftarrow -\xi_i + \beta \bar{\eta}_i$
 - 3: Compute angle between conjugate direction and gradient:
 $\alpha \leftarrow \langle \eta_i, \xi_i \rangle / \sqrt{\langle \eta_i, \eta_i \rangle \langle \xi_i, \xi_i \rangle}$
 - 4: Reset to gradient if desired:
 if $\alpha \leq 0.1$, then $\eta_i \leftarrow \xi_i$
-

performed extremely well as an initial guess in the sense that backtracking is almost never necessary.

Equation (3.3) is a one-dimensional least-square fit for t on Ω . The closed-form solution of the minimizer t_* satisfies

$$t_* = \langle P_\Omega(\eta), P_\Omega(A - X) \rangle / \langle P_\Omega(\eta), P_\Omega(\eta) \rangle$$

and is unique when $\eta \neq 0$. As long as Algorithm 1 has not converged, η will always be nonzero since it is the direction of search. The solution to (3.3) is performed by Algorithm 5. In the actual implementation, the sparse matrices N and B are stored as the nonzero entries on Ω . Hence, the total cost is about $2mk^2 + 4|\Omega|(k + 1)$ flops.

ALGORITHM 5. COMPUTE THE INITIAL GUESS FOR LINE SEARCH $t_* = \arg \min_t f(X + t\eta)$.

Require: iterate $X = U\Sigma V^T$ and projection X_Ω , tangent vector $\eta = U M V^T + U_p V_p^T$, sparse matrix $R = A_\Omega - X_\Omega \in \mathbb{R}^{m \times n}$

Ensure: step length t_*

- 1: $N \leftarrow P_\Omega([U M + U_p \ U] [V \ V_p]^T)$ # $2mk^2 + 4|\Omega|k$ flops
 - 2: $t_* \leftarrow \text{tr}(N^T R) / \text{tr}(N^T N)$ # $4|\Omega|$ flops
-

Retraction. As shown in (2.8), the retraction (2.7) can be directly computed by the SVD of $X + \xi$. A full SVD of $X + \xi$ is, however, prohibitively expensive since it costs $O(n^3)$ flops. Fortunately, the matrix to retract has the particular form

$$X + \xi = [U \ U_p] \begin{bmatrix} \Sigma + M & I_k \\ I_k & 0 \end{bmatrix} [V \ V_p]^T$$

with $X = U\Sigma V^T \in \mathcal{M}_k$ and $\xi = U M V^T + U_p V_p^T \in T_X \mathcal{M}_k$. Algorithm 6 now performs a compact QR and an SVD of a small $2k$ -by- $2k$ matrix to reduce the flop count to $14(m + n)k^2 + C_{\text{SVD}}k^3$ when $k \ll \min(m, n)$.

Observe that the listing of Algorithm 6 uses MATLAB notation to denote common matrix operations. In addition, the flop count of computing the SVD in step 3 is given as $C_{\text{SVD}}k^3$ since it is difficult to estimate beforehand in general. In practice, the constant C_{SVD} is modest, say, less than 200. Furthermore, in step 4, we have added $\varepsilon_{\text{mach}}$ to Σ_s so that in the unlucky event of zero singular values, the retracted matrix is perturbed to a rank- k matrix in \mathcal{M}_k .

ALGORITHM 6. COMPUTE THE RETRACTION BY METRIC PROJECTION.

Require: iterate $X = U\Sigma V^T$, tangent vector $\xi = U_M V^T + U_p V^T + U V_p^T$

Ensure: retraction $R_X(\xi) = P_{\mathcal{M}_k}(X + \xi) = U_+ \Sigma_+ V_+^T$

1: $(Q_u, R_u) \leftarrow \text{qr}(U_p, 0), \quad (Q_v, R_v) \leftarrow \text{qr}(V_p, 0) \quad \# 10(m+n)k^2 \text{ flops}$

2: $S \leftarrow \begin{bmatrix} \Sigma + M & R_v^T \\ R_u & 0 \end{bmatrix}$

3: $(U_s, \Sigma_s, V_s) \leftarrow \text{svd}(S) \quad \# C_{\text{SVD}}k^3 \text{ flops}$

4: $\Sigma_+ \leftarrow \Sigma_s(1:k, 1:k) + \varepsilon_{\text{mach}}$

5: $U_+ \leftarrow [U \quad Q_u] U_s(:, 1:k), \quad V_+ \leftarrow [V \quad Q_v] V_s(:, 1:k) \quad \# 4(m+n)k^2 \text{ flops}$

Computational cost. Summing all the flop counts for Algorithm 1, we arrive at a cost per iteration of approximately

$$(48m + 44n)k^2 + 10|\Omega|k + \text{nb}_{\text{Armijo}}(8(m+n)k^2 + 2|\Omega|k),$$

where $\text{nb}_{\text{Armijo}}$ denotes the average number of Armijo backtracking steps. It is remarkable that in all experiments below we have observed that $\text{nb}_{\text{Armijo}} = 0$, that is, backtracking was never needed.

For our typical problems in section 5, the size of the sampling set satisfies $|\Omega| = \text{OS}(m+n-k)k$ with $\text{OS} > 2$ the oversampling factor (OS). When $m = n$ and $\text{nb}_{\text{Armijo}} = 0$, this brings the total flops per iteration to

$$(3.4) \quad t_{\text{theoretical}} = (92 + 20 \text{OS})nk^2.$$

Comparing this theoretical flop count with experimental results, we observe that sparse matrix vector multiplications and applying P_Ω require significantly more time than predicted by (3.4). This can be explained due to the lack of data locality for these sparse matrix operations, whereas the majority of the remaining time is spent by dense linear algebra, rich in BLAS3. After some experimentation, we estimated that for our MATLAB environment these sparse operations are penalized by a factor of about $C_{\text{sparse}} \simeq 5$. For this reason, we normalize the theoretical estimate (3.4) to obtain the following practical estimate:

$$(3.5) \quad t_{\text{practical}}^{\text{LRGeomCG}} = (92 + C_{\text{sparse}} 20 \text{OS})nk^2, \quad C_{\text{sparse}} \simeq 5.$$

For a sensible value of $\text{OS} = 3$, we can expect that about 75% of the computational time will be spent on operations with sparse matrices.

Comparison to existing methods. The vast majority of specialized solvers for matrix completion require computing the dominant singular vectors of a sparse matrix in each step of the algorithm. This is, for example, the case for approaches using soft- and hard-thresholding, like in [8, 13, 21, 35, 36, 37, 38, 43, 50]. Typically, PROPACK from [31] is used for computing such a truncated SVD. The use of sparse matrix-vector products and the potential convergence issues frequently make this the computationally most expensive part of all these algorithms.

On the other hand, our algorithm first projects any sparse matrix onto the tangent space and then computes the dominant singular vectors for a small $2k$ -by- $2k$ matrix. Apart from the application of P_Ω and a few sparse matrix-vector products, the rest of the algorithm consists solely of dense linear algebra operations. This is more robust and significantly faster than computing the SVD of a sparse matrix in each step. To the best of our knowledge, LMAFit [53] and other manifold-related algorithms

like [5, 18, 27, 28, 39] are the only competitive solvers where mostly dense linear algebra together with a few sparse matrix-vector products are sufficient.

Due to the difficulty of estimating practical flop counts for algorithms based on computing a sparse SVD, we will only compare the computational complexity of our method with that of LMAFit. An estimate similar to the one of above reveals that LMAFit has a computational cost of

$$(3.6) \quad t_{\text{practical}}^{\text{LMAFit}} = (18 + C_{\text{sparse}} 12 \text{ OS})nk^2, \quad C_{\text{sparse}} \simeq 5,$$

where the sparse manipulations involve $2k$ sparse matrix vector multiplications and one application of P_Ω to a rank k matrix. Comparing this estimate to (3.5), LMAFit should be about two times faster per iteration when $\text{OS} = 3$. As we will see later (Tables 5.1 and 5.2) this is almost exactly what we observe experimentally. Since LRGeomCG is more costly per iteration, it will have to converge much faster in order to compete with LMAFit. This is studied in detail in section 5.

4. Convergence for a modified cost function. In this section, we show the convergence of Algorithm 1 applied to a *modified* cost function.

4.1. Reconstruction on the tangent space. As the first step, we apply the general convergence theory for Riemannian optimization from [2].

PROPOSITION 4.1. *Let $\{X_i\}$ be an infinite sequence of iterates generated by Algorithm 1. Then, every accumulation point X_* of $\{X_i\}$ satisfies $P_{T_{X_*}\mathcal{M}_k} P_\Omega(X_*) = P_{T_{X_*}\mathcal{M}_k} P_\Omega(A)$.*

Proof. Since R_X is a smooth retraction, the Armijo-type line search together with the gradient-related search directions, allows us to conclude that any limit point of Algorithm 1 is a critical point by Theorem 4.3.1 in [2]. These critical points X_* are determined by $\text{grad } f(X) = 0$. By (2.6) this gives $P_{T_{X_*}\mathcal{M}_k} P_\Omega(X_* - A) = 0$. \square

The next step is establishing that there exist limit points. However, since the closure of \mathcal{M}_k are the matrices of rank *bounded* by k , a sequence in \mathcal{M}_k may have a limit point that is a matrix of rank lower than k , which is not in \mathcal{M}_k . In addition, the objective function $f(X)$ is not coercive because P_Ω has a nontrivial kernel.

These two observations, motivate our choice to prove convergence when Algorithm 1 is applied not to f but a regularized version, namely,

$$g: \mathcal{M} \rightarrow \mathbb{R}, \quad X \mapsto f(X) + \mu^2(\|X^\dagger\|_F^2 + \|X\|_F^2), \quad \mu > 0.$$

In g , the term $\|X^\dagger\|_F^2$ will act as barrier to lower-rank matrices, while $\|X\|_F^2$ makes it coercive. We will show below that as a consequence, all iterates $\{X_i\}$ stay in a compact subset of \mathcal{M}_k . In addition, $\text{grad } g(X)$ can be computed very cheaply after modifying step 2 in Algorithm 2: Since every $X \in \mathcal{M}_k$ is of rank k , the pseudoinverse is smooth on \mathcal{M}_k and hence $g(X)$ is also smooth. Using [23, Theorem 4.3], one can then show that the gradient of $X \mapsto \|X^\dagger\|_F^2 + \|X\|_F^2$ equals $2U(\Sigma - \Sigma^{-3})V^T$ for $X = U\Sigma V^T$.

Before continuing, we hasten to point out that although optimizing g using Algorithm 1 can be done at virtually no extra cost, there is actually no need for g in a practical algorithm for matrix completion; see Remark 4.3. As such our convergence result is artificial, but it suits our need without formally proving convergence of this more practical scheme, which would be quite technical.

PROPOSITION 4.2. *Let $\{X_i\}$ be an infinite sequence of iterates generated by Algorithm 1 but with the objective function $g(X) = f(X) + \mu^2(\|X^\dagger\|_F^2 + \|X\|_F^2)$ for some $0 < \mu < 1$. Then, $\lim_{i \rightarrow \infty} P_{T_{X_i}\mathcal{M}_k} P_\Omega(X_i - A) = \lim_{i \rightarrow \infty} 2\mu^2(X_i^\dagger + X_i)$.*

Proof. We will first show that the iterates stay in a closed and bounded subset of \mathcal{M}_k . Let $L = \{X \in \mathcal{M}_k : g(X) \leq g(X_0)\}$ be the level set at X_0 . By construction of the line search, all elements of $\{X_i\}$ stay inside L and we get

$$\frac{1}{2} \|P_\Omega(X_i - A)\|_F^2 + \mu^2 \|X_i^\dagger\|_F^2 + \mu^2 \|X_i\|_F^2 \leq C_0^2, \quad i > 0,$$

where $C_0^2 := g(X_0)$. This implies

$$\mu^2 \|X_i\|_F^2 \leq C_0^2 - \frac{1}{2} \|P_\Omega(X_i - A)\|_F^2 - \mu^2 \|X_i^\dagger\|_F^2 \leq C_0^2,$$

and we obtain an upper bound on the largest singular value:

$$\sigma_1(X_i) \leq \|X_i\|_F \leq C_0/\mu =: C^\sigma, \quad i > 0.$$

Similarly, we get a lower bound on the smallest singular value:

$$\mu^2 \|X_i^\dagger\|_F^2 = \sum_{j=1}^k \frac{\mu^2}{\sigma_j^2(X_i)} \leq C_0^2 - \frac{1}{2} \|P_\Omega(X_i - A)\|_F^2 - \mu^2 \|X_i\|_F^2 \leq C_0^2, \quad i > 0,$$

which implies that

$$\sigma_k(X_i) \geq \mu/C_0 =: C_\sigma, \quad i > 0.$$

It is clear that all elements of $\{X_i\}$ stay inside the set

$$B = \{X \in \mathcal{M}_k : \sigma_1(X) \leq C^\sigma, \sigma_k(X) \geq C_\sigma\}.$$

This set is closed and bounded, hence compact.

Next, suppose that $\lim_{i \rightarrow \infty} \|\text{grad } g(X_i)\|_F \neq 0$. Then there is a subsequence in $\{X_i\}_{i \in \mathcal{K}}$ such that $\|\text{grad } g(X)\|_F \geq \varepsilon > 0$ for all $i \in \mathcal{K}$. Since $X_i \in B$, this subsequence $\{X_i\}_{i \in \mathcal{K}}$ has a limit point X_* in B . By continuity of $\text{grad } g$, this implies that $\|\text{grad } g(X_*)\|_F \geq \varepsilon$ which contradicts Theorem 4.3.1 in [2] that every accumulation point is a critical point of g . Hence, $\lim_{i \rightarrow \infty} \|\text{grad } g(X_i)\|_F = 0$. \square

Remark 4.3. Proposition 4.2 is theoretical and μ^2 can be chosen arbitrarily small, for example, as small as $\varepsilon_{\text{mach}} \simeq 10^{-16}$. This means that as long as $\sigma_1(X_i) = O(1)$ and $\sigma_k(X_i) \gg \varepsilon_{\text{mach}}$, the regularization terms in g are negligible and one might as well have optimized the original objective function f , that is, g with $\mu = 0$. This is what we observed in all numerical experiments when $\text{rank}(A) \geq k$. In case $\text{rank}(A) < k$, it is obvious that now $\sigma_k(X_i) \rightarrow 0$ as $i \rightarrow \infty$. Theoretically one can still optimize g , thereby forcing the optimization problem to stay on \mathcal{M}_k , but in practice it makes more sense to transfer the original optimization problem to the manifold of rank $k - 1$ matrices. This can be accomplished by monitoring $\sigma_k(X_i)$ throughout the iteration. Since these rank drops do not occur in our experiments, proving convergence of such a method is beyond the scope of this paper.

A similar argument appears in the analysis of other methods for rank-constrained optimization too. For example, the proof of convergence in [12] for SDPLR [11] uses a regularization by $\mu \det(\Sigma)$ with Σ containing the nonzero singular values of X ; yet in practice, ones optimizes without this regularization using the observation that μ can be made arbitrarily small, and in practice one always observes convergence when $\mu = 0$. Analogous modifications also appear in [8, 9, 28].

4.2. Reconstruction on the whole space. Based on Proposition 4.2, we have that for a suitable choice of μ , any limit point satisfies

$$\|P_{T_{X_*} \mathcal{M}_k} P_{\Omega}(X_* - A)\|_F \leq \varepsilon, \quad \varepsilon > 0.$$

In other words, we will have approximately fitted X to the data A on the image of $P_{T_{X_*} \mathcal{M}_k} P_{\Omega}$, and, hopefully, this is sufficient to have $X_* = A$ on the whole space $\mathbb{R}^{m \times n}$. Without further assumptions on X_* and A , however, it is not reasonable to expect that X_* equals A since P_{Ω} usually has a very large null space.

Let us split the error $E := X - A$ into the three quantities,

$$E_1 = P_{T_X \mathcal{M}_k} P_{\Omega}(X - A), \quad E_2 = P_{N_X \mathcal{M}_k} P_{\Omega}(X - A), \quad \text{and} \quad E_3 = P_{\Omega}^{\perp}(X - A),$$

where $N_X \mathcal{M}_k$ is the normal space of X , i.e., the subspace of all matrices $Z \in \mathbb{R}^{m \times n}$ orthogonal to $T_X \mathcal{M}_k$.

Upon convergence of Algorithm 1, $\|E_1\|_F \rightarrow \varepsilon$, but how big can E_2 be? In general, E_2 can be arbitrarily large since it is the Lagrange multiplier of the fixed-rank constraint of X . In fact, when the rank of X is smaller than the exact rank of A , E_2 typically never vanishes, even though E_1 can converge to zero. On the other hand, when the ranks of X and A are the same, one typically observes that $\|E_1\|_F \rightarrow 0$ implies $\|E_2\|_F \rightarrow 0$. This can also be easily checked during the course of the algorithm, since $P_{\Omega}(X - A) = E_1 + E_2$ is computed explicitly for the computation of the gradient in Algorithm 2.

Suppose now that $\|E_1\|_F \leq \varepsilon$ and $\|E_2\|_F \leq \tau$, which means that X is in good agreement with A on Ω . The hope is that X and A agree on the complement of Ω too. This is of course not always the case, but it is exactly the crucial assumption of low-rank matrix completion that the observations on Ω are sufficient to complete A . To quantify this assumption, one can make use of standard theory in matrix completion literature; see, e.g., [15, 16, 27]. Since it is not the purpose of this paper to analyze the convergence properties of Algorithm 1 for various, rather theoretical random models, and we do not rely on this theory in the numerical experiments later on, we omit this discussion.

5. Numerical experiments. The implementation of LRGeomCG, i.e., Algorithm 1, was done in MATLAB R2012a on a desktop computer with a 3.10 GHz CPU and 8 GB of memory. All reported times are wall-clock time that include the setup phase of the solvers (if needed) but exclude setting up the (random) problem.

As is mostly done in the literature, we focus only on square matrices. Note however that the performance of most matrix completion solvers behaves very differently for highly nonsquare matrices, as observed in [8]. In our numerical experiments below we only compare with LMAFit of [53]. Extensive comparison with the other approaches from [13, 21, 35, 36, 37, 38, 50] revealed that LMAFit is overall the fastest method for square matrices. We refer to [42] for these results.

The implementation of LMAFit was the one provided by the authors.¹ All options were kept the same, except rank adaptivity was turned off. In sections 5.4 and 5.5 we also changed the stopping criteria as explained there. In order to have a fair comparison between LMAFit and LRGeomCG, the operation P_{Ω} was performed by the same MATLAB function `part_XY.m` in both solvers. In addition, the initial guesses were taken as random rank k matrices (more on this below).

¹The version used was downloaded August 14, 2012 from <http://lmafit.blogs.rice.edu/>.

In all experiments, we will complete random low-rank matrices which are generated as proposed in [13]. First, construct two random matrices $A_L, A_R \in \mathbb{R}^{n \times k}$ with i.i.d. standard Gaussian entries; then, assemble $A := A_L A_R^T$; finally, the set of observations Ω is sampled uniformly at random among all sets of cardinality $|\Omega|$. The resulting observed matrix to complete is now $A_\Omega := P_\Omega(A)$. Standard random matrix theory asserts that $\|A\|_F \simeq n\sqrt{k}$ and $\|A_\Omega\|_F \simeq \sqrt{|\Omega|k}$. In the later experiments, the test matrices are different and their construction will be explained there.

When we report on the relative error of an approximation X , we mean the error

$$(5.1) \quad \text{relative error} = \|X - A\|_F / \|A\|_F$$

computed on all the entries of A . On the other hand, the algorithms will compute a relative residual on Ω only:

$$(5.2) \quad \text{relative residual} = \|P_\Omega(X - A)\|_F / \|P_\Omega(A)\|_F.$$

Unless stated otherwise, we set as tolerance for the solvers a relative residual of 10^{-12} . Such a high precision is necessary to study the asymptotic convergence rate of the solvers, but where appropriate we will draw conclusions for moderate precisions too.

As initial guess X_1 for Algorithm 1, we construct a random rank- k matrix following the same procedure as above for M . The OS for a rank- k matrix is defined as the ratio of the number of samples to the degrees of freedom in a nonsymmetric matrix of rank k ,

$$(5.3) \quad \text{OS} = |\Omega| / (k(2n - k)).$$

Obviously one needs at least $\text{OS} \geq 1$ to uniquely complete any rank- k matrix. In the experiments below, we observe that regardless of the method, $\text{OS} > 2$ is needed to reliably recover an incoherent matrix of rank k after uniform sampling. In addition, each row and each column needs to be sampled at least once. By the coupon collector’s problem, this requires that $|\Omega| > Cn \log(n)$, but in all experiments below, Ω is always sufficiently large such that each row and column is sampled at least once. Hence, we will only focus on the quantity OS.

5.1. Influence of size and rank for fixed oversampling. As the first test, we complete random matrices A of exact rank k with LRGeomCG and LMAFit and we explicitly exploit the knowledge of the rank in the algorithms. The purpose of this test is to investigate the dependence of the algorithms on the size and the rank of the matrices. In section 5.3, we will study the influence of OS, but for now we fix the oversampling to $\text{OS} = 3$. In Table 5.1, we report on the mean run time and the number of iterations taking over 10 random instances. The run time and the iteration count of all these instances are visualized in Figures 5.1 and 5.2, respectively.

Overall, LRGeomCG needs less iterations than LMAFit for all problems. In Figure 5.1, we see that this is because LRGeomCG converges faster asymptotically although the convergence of LRGeomCG exhibits a slower transient behavior in the beginning. This phase is, however, only limited to the first 20 iterations. Since the cost per iteration for LMAFit is cheaper than for LRGeomCG (see the estimates (3.5) and (3.6)), this transient phase leads to a trade-off between runtime and precision. This is clearly visible in the timings of Figure 5.2: When sufficiently high accuracy is requested, LRGeomCG is always faster, whereas for low precision, LMAFit is faster.

Let us now check in more detail the dependence on the size n . It is clear from the left sides of Table 5.1 and Figure 5.1 that the iteration counts of LMAFit and

TABLE 5.1

The mean of the computational results for Figures 5.1–5.2 for solving with a tolerance of 10^{-12} .

n	Fixed rank $k = 40$				Fixed size $n = 8000$				
	LMAFit		LRGeomCG		LMAFit		LRGeomCG		
	Time(s.)	#its.	Time(s.)	#its.	k	Time(s.)	#its.	Time(s.)	#its.
1000	3.11	133	2.74	54.5	10	17.6	518	8.81	121
2000	9.13	175	6.43	61.3	20	23.9	297	14.6	86.5
4000	21.2	191	14.8	66.7	30	40.9	232	25.8	76.1
8000	53.3	211	36.5	71.7	40	52.7	211	35.8	71.7
16000	150	222	99.1	75.4	50	76.1	194	51.2	67.7
32000	383	233	254	79.1	60	96.6	186	67.0	66.1

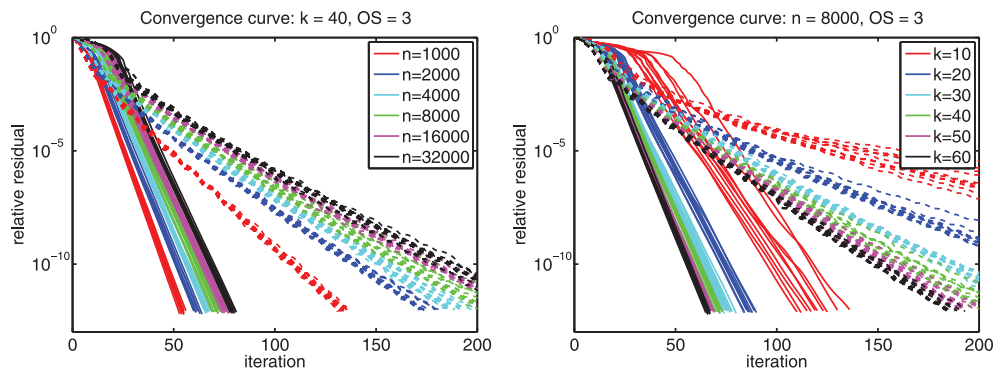


FIG. 5.1. Convergence curves for LMAFit (dashed line) and LRGeomCG (full line) for fixed oversampling of $OS = 3$. Left: variable size n and fixed rank $k = 40$; right: variable ranks k and fixed size $n = 8000$.

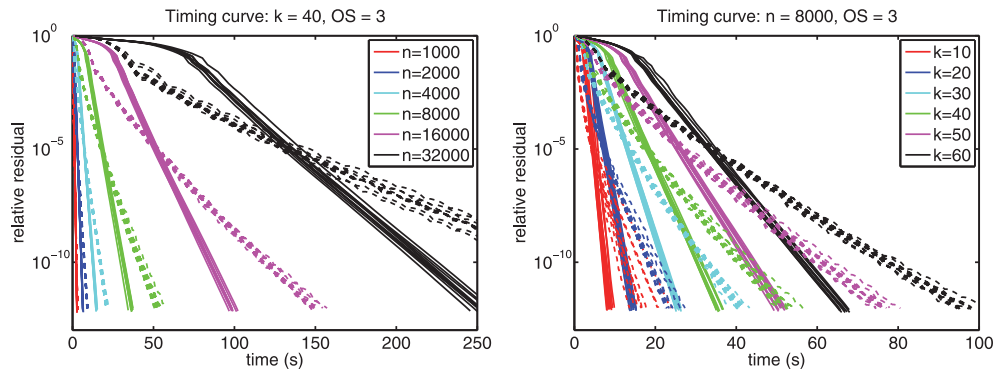


FIG. 5.2. Timing curves for LMAFit (dashed line) and LRGeomCG (full line) for fixed oversampling of $OS = 3$. Left: variable size n and fixed rank $k = 40$; right: variable ranks k and fixed size $n = 8000$.

LRGeomCG grow with n but seem to stagnate as $n \rightarrow \infty$. In addition, LMAFit needs about three times more iterations than LRGeomCG. Hence, even though each iteration of LRGeomCG is about two times more expensive than one iteration of LMAFit, LRGeomCG will eventually always be faster than LMAFit for sufficiently

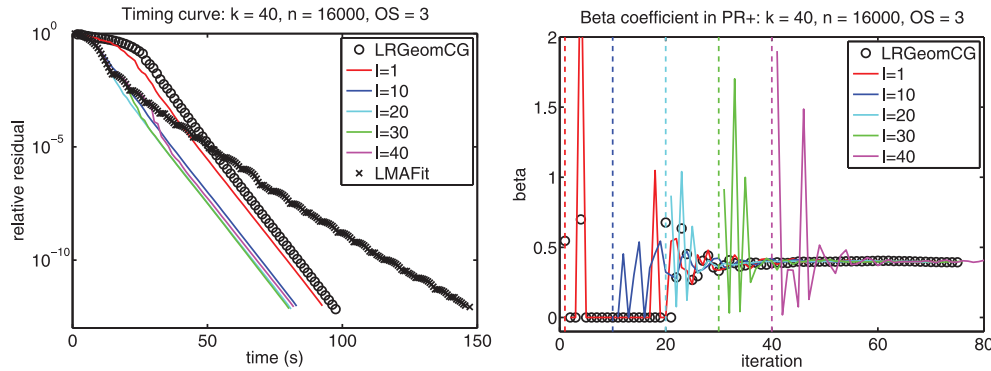


FIG. 5.3. Timing curve and β coefficient of the hybrid strategy for different values of I .

high precision. In Figure 5.2 on the left, one can, for example, see that this trade-off point happens at a precision of about 10^{-5} . For growing rank k , the conclusion of the same analysis is very much the same.

The previous conclusion depends critically on the convergence factor and, as we will see later on, this factor is determined by the amount of oversampling (OS). But for difficult problems (OS = 3 is near the lower limit of 2), we can already observe that LRGeomCG is about 50% faster than LMAFit.

5.2. Hybrid strategy. The experimental results from above clearly show that the transient behavior of LRGeomCG’s convergence is detrimental if only modest accuracy is required. The reason for this slow phase is the lack of nonlinear CG acceleration (β in Algorithm 4 is almost always zero) which essentially reduces LRGeomCG to a steepest descent algorithm. On the other hand, LMAFit is much less affected by slow convergence during the initial phase.

A simple heuristic to overcome this bad phase is a hybrid solver: For the first I iterations, we use LMAFit; after that, we hope that the nonlinear CG acceleration kicks in immediately and we can efficiently solve with LRGeomCG. In Figure 5.3, we have tested this strategy for different values of I on a problem of the previous section with $n = 16000$ and $k = 40$.

From the left panel of Figure 5.3, we get that the run time to solve for a tolerance of 10^{-12} is reduced from 97 sec (LRGeomCG) and 147 sec (LMAFit) to about 80 sec for the choices $I = 10, 20, 30, 40$. Performing only one iteration of LMAFit is less effective. For $I = 20$, the hybrid strategy has almost no transient behavior anymore and it is always faster than LRGeomCG or LMAFit alone. Furthermore, the choice of I is not very sensitive to the performance as long as it is between 10 and 40, and already for $I = 10$, there is a significant speedup of LRGeomCG noticeable for all tolerances.

The quantity β of PR+ in Algorithm 4 is plotted in the right panel of Figure 5.3. Until the 20th iteration, plain LRGeomCG shows no meaningful CG acceleration. For the hybrid strategies with $I > 10$, the acceleration kicks in almost immediately. Observe that all approaches converge to $\beta \simeq 0.4$.

While this experiment shows that there is potential to speed up LRGeomCG in the early phase, we do not wish to claim that the present strategy is very robust or directly applicable. The main point that we wish to make, however, is that the slow phase can be avoided in a relatively straightforward way and that LRGeomCG can be

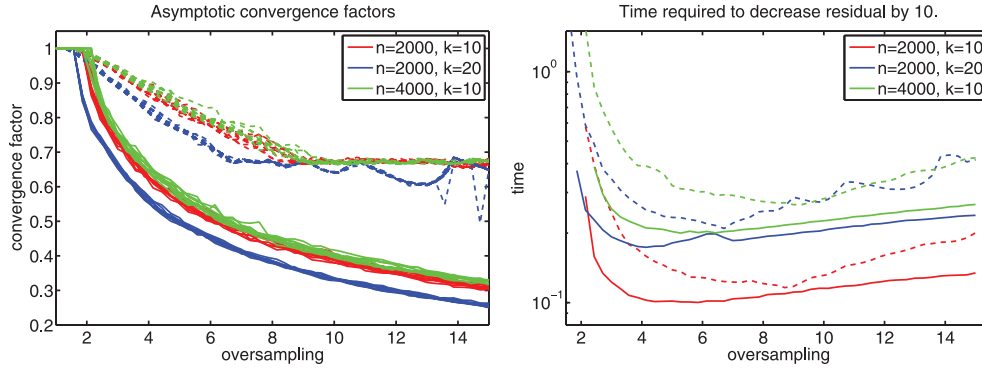


FIG. 5.4. Asymptotic convergence factor ρ (left) and mean time to decrease the error by a factor of 10 (right) for LMAFit (dashed line) and LRGeomCG (full line) in function of the OS.

warm started by any other method. In particular, this shows that LRGeomCG can be efficiently employed once the correct rank of the solution X is identified by some other method. As explained in the introduction, there are numerous other methods for low-rank matrix completion that can reliably identify this rank but have a slow asymptotic convergence. Combining them with LRGeomCG seems like a promising way forward.

5.3. Influence of oversampling. Next, we investigate the influence of oversampling on the convergence speed. We take 10 random problems with fixed rank and size, but vary the OS using 50 values between $1, \dots, 15$. In Figure 5.4, on the left, the asymptotic convergence factor ρ is visible for LMAFit and LRGeomCG for three different combinations of the size and rank. Since the convergence is linear and we want to filter out any transient behavior, this factor was computed as

$$\rho = \left(\frac{\|P_{\Omega}(X_{i_{\text{end}}} - A)\|_F}{\|P_{\Omega}(X_{10} - A)\|_F} \right)^{1/(i_{\text{end}} - 10)},$$

where i_{end} indicates the last iteration. A factor of one indicates failure to converge within 4000 steps. One can observe that all methods become slower as $\text{OS} \rightarrow 2$. Further, the convergence factor of LMAFit seems to stagnate for large values of OS while LRGeomCG's actually becomes better. Since for growing OS the completion problems become easier as more entries are available, only LRGeomCG shows the expected behavior. In contrast to our parameter-free choice of nonlinear CG, LMAFit needs to determine and adjust a certain acceleration factor dynamically based on the performance of the iteration. We believe the stagnation in Figure 5.4 on the left is due to a suboptimal choice of this factor in LMAFit.

In the right panel of Figure 5.4, the mean time to decrease the relative residual by a factor of 10 is displayed. These timings were again determined by neglecting the first 10 iterations and then interpolating the time needed for a reduction of the residual by 10. Similarly as before, we can observe that since the cost per iteration is cheaper for LMAFit, there is a range for OS around $7, \dots, 9$ where LMAFit is only slightly slower than LRGeomCG. However, for smaller and larger values of OS, LRGeomCG is always faster by a significant margin (observe the logarithmic scale).

5.4. Influence of noise. Next, we investigate the influence of noise by adding random perturbations to the rank- k matrix A . We define the noisy matrix $A^{(\varepsilon)}$ with

noise level ε as

$$A^{(\varepsilon)} := A + \varepsilon \frac{\|A_\Omega\|_F}{\|N_\Omega\|_F} N,$$

where N is a standard Gaussian matrix and Ω is the usual random sampling set. (See also [50, 53] for a similar setup.) The reason for defining $A^{(\varepsilon)}$ in this way is that we have

$$\|P_\Omega(A - A^{(\varepsilon)})\|_F = \varepsilon \frac{\|A_\Omega\|_F}{\|N_\Omega\|_F} \|P_\Omega(N)\|_F \simeq \varepsilon \sqrt{|\Omega|k}.$$

So, the best relative residual we may expect from an approximation $X^{\text{opt}} \simeq A$ is

$$(5.4) \quad \|P_\Omega(X^{\text{opt}} - A^{(\varepsilon)})\|_F / \|P_\Omega(A)\|_F \simeq \|P_\Omega(A - A^{(\varepsilon)})\|_F / \|P_\Omega(A)\|_F \simeq \varepsilon.$$

Similarly, the best relative error of X^{opt} should be on the order of the noise ratio too since

$$(5.5) \quad \|A - A^{(\varepsilon)}\|_F / \|A\|_F \simeq \varepsilon.$$

Due to the presence of noise, the relative residual (5.1) cannot go to zero but the Riemannian gradient of our optimization problem can. In principle, this suffices to detect convergence of LRGeomCG. In practice, however, we have noticed that this wastes a lot of iterations in case the iteration stagnates. A simply remedy is to monitor

$$(5.6) \quad \text{relative change at step } i = \left| 1 - \sqrt{f(X_i)/f(X_{i-1})} \right|$$

and stop the iteration when this value drops below a certain threshold. After some experimentation, we fixed this threshold to 10^{-3} . Such a stagnation detection is applied by most other low-rank completion solvers [50, 53], although its specific form varies. Our choice coincides with the one from [53], but we have lowered the threshold.

After equipping LMAFit and LRGeomCG with this stagnation detection, we can display the experimental results in Figure 5.5 and Table 5.2. We have only reported on one choice for the rank, size, and oversampling since the same conclusions can be reached for any other choice. Based on Figure 5.5, it is clear that the stagnation is effectively detected by both methods, except for the very noisy case of $\varepsilon = 1$. (Recall that we changed the original stagnation detection procedure in LMAFit to ours of (5.6) to be able to draw a fair comparison.)

Comparing the results with those of the previous section, the only difference is that the iteration stagnates when the error reaches the noise level. In particular, the iterations are undisturbed by the presence of the noise up until the very last iterations. In Table 5.2, one can observe that both methods solve all problems up to the noise level, which is the best that can be expected from (5.4) and (5.5). Again, LRGeomCG is faster when a higher accuracy is required, which, in this case, corresponds to small noise levels. Since the iteration is unaffected by the noise until the very end, the hybrid strategy of the previous section should be effective here too.

5.5. Exponentially decaying singular values. In the previous problems, the rank of the matrices could be unambiguously defined. Even in the noisy case, there was still a sufficiently large gap between the original nonzero singular values and the




FIG. 5.5. Timing curves for LMAFit (crosses) and LRGeomCG (circles) for different noise levels with size $n = 8000$, rank $k = 20$, and fixed oversampling of $OS = 3$. The noise levels ε are indicated by dashed lines in different color.

TABLE 5.2
Computational results for Figure 5.5.

ε	LMAFit				LRGeomCG			
	Time(s.)	#its.	Error	Residual	Time(s.)	#its.	Error	Residual
10^{-0}	2.6	31	1.04ε	0.59ε	6.8	42	1.52ε	0.63ε
10^{-2}	3.4	42	0.77ε	0.83ε	5.6	36	0.72ε	0.82ε
10^{-4}	6.4	79	0.83ε	0.83ε	7.5	46	0.72ε	0.82ε
10^{-6}	11	133	0.74ε	0.82ε	9.2	58	0.72ε	0.82ε
10^{-8}	15	179	0.96ε	0.85ε	11	70	0.72ε	0.82ε
10^{-10}	19	235	1.09ε	0.87ε	13	81	0.72ε	0.82ε

ones affected by noise. In this section, we will complete a matrix for which the singular values decay but do not become exactly zero.

Although a different setting than the previous problems, this type of matrix occurs frequently in the context of approximation of discretized functions or solutions of PDEs on tensorized domains; see, e.g., [30]. In this case, the problem of approximating discretized functions by low-rank matrices is primarily motivated by reducing the amount of data to store. On a continuous level, low-rank approximation corresponds in this setting to approximation by sums of separable functions. As such, one is typically interested in approximating the matrix up to a certain tolerance using the lowest rank possible.

The (exact) rank of a matrix is now better replaced by the numerical rank, or ε -rank [22, Chapter 2.5.5]. Depending on the required accuracy ε , the ε -rank is the quantity

$$(5.7) \quad k_\varepsilon(A) = \min_{\|A-B\|_2 \leq \varepsilon} \text{rank}(B).$$

Obviously, when $\varepsilon = 0$, one recovers the (exact) rank of a matrix. It is well known that the ε -rank of $A \in \mathbb{R}^{m \times n}$ can be determined from the singular values σ_i of A as follows

$$\sigma_1 \geq \sigma_{k_\varepsilon} > \varepsilon \geq \sigma_{k_\varepsilon+1} \geq \dots \geq \sigma_p, \quad p = \min(m, n).$$

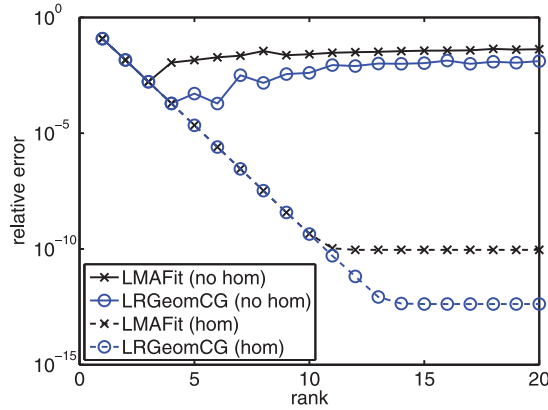


FIG. 5.6. Relative error of LMAFit and LRGeomCG after completion of a discretization of (5.8) for OS = 8 with and without a homotopy strategy.

In the context of the approximation of bivariate functions, the relation (5.7) is very useful in the following way. Let $f(x, y)$ be a function defined on a tensorized domain $(x, y) \in \mathcal{X} \times \mathcal{Y}$. After discretization of x and y , we arrive at a matrix A . If we allow for an absolute error in the spectral norm of the size ε , then (5.7) tells us that we can approximate M by a rank k_ε matrix. For example, we take the following simple bivariate function to complete:

$$(5.8) \quad f(x, y) = \frac{1}{1 + \|x - y\|_2^2}, \quad (x, y) \in [0, 1]^2.$$

Such functions occur as two-point correlations related to second-order elliptic problems with stochastic source terms; see [24].

We ran LRGeomCG and LMAFit after uniformly discretizing x and y in (5.8) resulting in a matrix of size $n = 8000$ and then sampling this matrix with OS = 8 (calculated for rank 20). Since a stopping condition on the residual has no meaning for this example, we only used a tolerance of 10^{-3} on the relative change (5.6) or a maximum of 500 iterations. In Figure 5.6, the final accuracy for different choices of the rank is visible for these two problem sets.

The results labeled “no hom” use the standard strategy of starting with random initial guesses for each value of the rank k . Clearly, this results in a very unsatisfactory performance of LRGeomCG and LMAFit since the error of the completion does not decrease with increasing rank. Remark that we have measured the relative error as

$$\|P_\Gamma(X_* - A)\|_F / \|P_\Gamma(A)\|_F,$$

where Γ is a random sampling set different from Ω but equally large.

The previous behavior is a clear example that the local optimizers of LMAFit and LRGeomCG are very far away from the global ones. However, there is a straightforward solution to this problem: Instead of taking a random initial guess for each k , we only take a random initial guess for $k = 1$. For all other $k > 1$, we use $X = U\Sigma V^T$, where

$$U = \begin{bmatrix} \tilde{U} & u \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \tilde{\Sigma} & 0 \\ 0 & \tilde{\Sigma}_{k-1, k-1} \end{bmatrix}, \quad V = \begin{bmatrix} \tilde{V} & v \end{bmatrix},$$

$\tilde{X} = \tilde{U}\tilde{\Sigma}\tilde{V}^T$ is the local optimizer for rank $k - 1$, and u and v are random unit-norm vectors orthogonal to U and V , respectively. We call this the *homotopy strategy* which is also used in one of LMAFit’s rank adaptivity strategies (but with zero vectors u and v). The error using this homotopy strategy is labeled “hom” in Figure 5.6 and it clearly performs much more favorably. In addition, LRGeomCG is able to obtain a smaller error.

6. Conclusions. The matrix completion consists of recovering a low-rank matrix based on a very sparse set of entries of this matrix. In this paper, we have presented a new method based on optimization on manifolds to solve large-scale matrix completion problems. Compared to most other existing methods in the literature, our approach consists of directly minimizing the least-square error of the fit directly over \mathcal{M}_k , the set of matrices of rank k .

The main contribution of this paper is to show that the lack of vector space structure of \mathcal{M}_k does not need be an issue when optimizing over this set and to illustrate this for low-rank matrix completion. Using the framework of retraction-based optimization, the necessary expressions were derived in order to minimize any smooth objective function over the Riemannian manifold \mathcal{M}_k . The geometry chosen for \mathcal{M}_k , namely, a submanifold embedded in the space of matrices, allowed for an efficient implementation of nonlinear CG for matrix completion. Indeed, the numerical experiments illustrated that this approach outperforms state-of-the-art solvers for matrix completion. In particular, the method achieved very fast asymptotic convergence factors without any tuning of parameters thanks to a simple computation of the initial guess for the line search.

A drawback of the proposed approach is, however, that the rank of the manifold is fixed. Although there are several problems where choosing the rank is straightforward, an integrated manifold-based solution, like in [17, 43], is desirable. In addition, our convergence proof relied on several safeguards and modifications in the algorithm that seem completely unnecessary in practice. Closing these gaps are currently topics of further research.

Acknowledgments. Parts of this paper have been prepared while the author was affiliated with the Seminar of Applied Mathematics, ETH Zürich. In addition, he gratefully acknowledges the helpful discussions with Daniel Kressner regarding low-rank matrix completion.

REFERENCES

- [1] P.-A. ABSIL, C. G. BAKER, AND K. A. GALLIVAN, *Trust-region methods on Riemannian manifolds*, *Found. Comput. Math.*, 7 (2007), pp. 303–330.
- [2] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, NJ, 2008.
- [3] P.-A. ABSIL AND J. MALICK, *Projection-like retractions on matrix manifolds*, *SIAM J. Optim.*, 22 (2012), pp. 135–158.
- [4] R. L. ADLER, J.-P. DEDIEU, J. Y. MARGULIES, M. MARTENS, AND M. SHUB, *Newton’s method on Riemannian manifolds and a geometric model for the human spine*, *IMA J. Numer. Anal.*, 22 (2002), pp. 359–390.
- [5] L. BALZANO, R. NOWAK, AND B. RECHT, *Online identification and tracking of subspaces from highly incomplete information*, in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Momticeo, IL, 2010.
- [6] J. BENNETT AND S. LANNING, *The Netflix Prize*, in *Proceedings of the KDD-Cup and Workshop at the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Jose, CA, 2007.
- [7] D. P. BERTSEKAS, *Nonlinear Programming*, Athena Scientific, Nashua, NH, 1999.

- [8] N. BOUMAL AND P.-A. ABSIL, *RTRMC: A Riemannian trust-region method for low-rank matrix completion*, in Proceedings of the Neural Information Processing Systems Conference (NIPS), Granada, Spain, 2011.
- [9] N. BOUMAL AND P.-A. ABSIL, *Low-Rank Matrix Completion via Trust-Regions on the Grassmann Manifold*, Technical report 2012.07, Université catholique de Louvain, INMA, Louvain-La-Neuve, Belgium, 2012.
- [10] W. BRUNS AND U. VETTER, *Determinantal rings*, Lecture Notes in Math. 1327, Springer-Verlag, Berlin, 1988.
- [11] S. BURER AND R. D. C. MONTEIRO, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Program., 95 (2003), pp. 329–357.
- [12] S. BURER AND R. D. C. MONTEIRO, *Local minima and convergence in low-rank semidefinite programming*, Math. Program., 103 (2005), pp. 427–444.
- [13] J.-F. CAI, E. J. CANDÈS, AND Z. SHEN, *A singular value thresholding algorithm for matrix completion*, SIAM J. Optim., 20 (2010), pp. 1956–1982.
- [14] E. J. CANDÈS AND Y. PLAN, *Matrix completion with noise*, Proceedings of the IEEE, 98 (2010), pp. 925–936.
- [15] E. J. CANDÈS AND T. TAO, *The power of convex relaxation: Near-optimal matrix completion*, IEEE Trans. Inform. Theory, 56 (2009), pp. 2053–2080.
- [16] E. CANDÈS AND B. RECHT, *Exact matrix completion via convex optimization*, Found. Comput. Math., 9 (2009), pp. 717–772.
- [17] T. P. CASON, P.-A. ABSIL, AND P. VAN DOOREN, *Iterative methods for low rank approximation of graph similarity matrices*, Linear Algebra Appl., 438 (2013), pp. 1863–1882.
- [18] W. DAI, O. MILENKOVIC, AND E. KERMAN, *Subspace evolution and transfer (SET) for low-rank matrix completion*, IEEE Trans. Signal Process., 59 (2011), pp. 3120–3132.
- [19] W. DAI AND O. MILENKOVIC, *SET: An algorithm for consistent matrix completion*, in International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Dallas, TX, 2010, pp. 3646–3649.
- [20] A. EDELMAN, T. A. ARIAS, AND S. T. SMITH, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 303–353.
- [21] D. GOLDFARB AND S. MA, *Convergence of fixed point continuation algorithms for matrix rank minimization*, Found. Comput. Math., 11 (2011), pp. 183–210.
- [22] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins Stud. Math. Sci., Baltimore, MD, 1996.
- [23] G. H. GOLUB AND V. PEREYRA, *The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate*, SIAM J. Numer. Anal., 10 (1973), pp. 413–432.
- [24] H. HARBRECHT, M. PETERS, AND R. SCHNEIDER, *On the low-rank approximation by the pivoted Cholesky decomposition*, Appl. Numer. Math., 62 (2012), pp. 428–440.
- [25] U. HELMKE AND J. B. MOORE, *Optimization and Dynamical Systems*, Springer-Verlag, London, 1994.
- [26] M. JOURNÉE, F. BACH, P.-A. ABSIL, AND R. SEPULCHRE, *Low-rank optimization on the cone of positive semidefinite matrices*, SIAM J. Optim., 20 (2010), pp. 2327–2351.
- [27] R. H. KESHAVAN, A. MONTANARI, AND S. OH, *Matrix completion from a few entries*, IEEE Trans. Inform. Theory, 56 (2010), pp. 2980–2998.
- [28] R. KESHAVAN, A. MONTANARI, AND S. OH, *Matrix completion from noisy entries*, J. Mach. Learn. Res., 11 (2010), pp. 2057–2078.
- [29] O. KOCH AND C. LUBICH, *Dynamical low-rank approximation*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 434–454.
- [30] D. KRESSNER AND C. TOBLER, *Krylov subspace methods for linear systems with tensor product structure*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 1688–1714.
- [31] R. M. LARSEN, *PROPACK—Software for Large and Sparse SVD Calculations*. <http://soi.stanford.edu/~rmunk/PROPACK> (2004).
- [32] J. M. LEE, *Introduction to smooth manifolds*, Grad. Texts in Math. 218, Springer-Verlag, New York, 2003.
- [33] K. LEE AND Y. BRESLER, *ADMiRA: Atomic decomposition for minimum rank approximation*, IEEE Trans. Inform. Theory, 56 (2010), pp. 4402–4416.
- [34] A. S. LEWIS AND J. MALICK, *Alternating projections on manifolds*, Math. Oper. Res., 33 (2008), pp. 216–234.
- [35] Z. LIN, M. CHEN, L. WU, AND YI MA, *The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices*, Technical report UILU-ENG-09-2215, University of Illinois, Urbana, IL, 2009.
- [36] Y.-J. LIU, D. SUN, AND K.-C. TOH, *An implementable proximal point algorithmic framework for nuclear norm minimization*, Math. Program., 133 (2012), pp. 399–436.

- [37] S. MA, D. GOLDFARB, AND L. CHEN, *Fixed point and Bregman iterative methods for matrix rank minimization*, Math. Program., 128 (2011), pp. 321–353.
- [38] R. MEKA, P. JAIN, AND I. S. DHILLON, *Guaranteed rank minimization via singular value projection*, in Proceedings of the Neural Information Processing Systems Conference (NIPS), 2010.
- [39] G. MEYER, S. BONNABEL, AND R. SEPULCHRE, *Linear regression under fixed-rank constraints: A Riemannian approach*, in Proceedings of the 28th International Conference on Machine Learning (ICML2011), Bellevue, WA, 2011.
- [40] G. MEYER, S. BONNABEL, AND R. SEPULCHRE, *Regression on fixed-rank positive semidefinite matrices: A Riemannian approach*, J. Mach. Learn. Res., 12 (2011), pp. 593–625.
- [41] G. MEYER, *Geometric Optimization Algorithms for Linear Regression on Fixed-Rank Matrices*, Ph.D. thesis, University of Liège, Liège, Belgium, 2011.
- [42] M. MICHENKOVÁ, *Numerical algorithms for low-rank matrix completion problems*, <http://www.math.ethz.ch/~kressner/students/michenkova.pdf> (2011).
- [43] B. MISHRA, G. MEYER, F. BACH, AND R. SEPULCHRE, *Low-Rank Optimization with Trace Norm Penalty*, preprint, 2011, <http://arxiv.org/abs/1112.2318>.
- [44] B. MISHRA, G. MEYER, S. BONNABEL, AND R. SEPULCHRE, *Fixed-Rank Matrix Factorizations and Riemannian Low-Rank Optimization*, preprint, 2012, <http://arxiv.org/abs/1209.0430>.
- [45] R. ORSI, U. HELMKE, AND J. B. MOORE, *A Newton-like method for solving rank constrained linear matrix inequalities*, Automatica, 42 (2006), pp. 1875–1882.
- [46] C. QI, K. A. GALLIVAN, AND P.-A. ABSIL, *Riemannian BFGS algorithm with applications*, in Recent Advances in Optimization and its Applications in Engineering, Springer, to appear.
- [47] U. SHALIT, D. WEINSHALL, AND G. CHECHIK, *Online learning in the manifold of low-rank matrices*, in Advances in Neural Inf. Process. Syst. 23, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, eds., 2010, pp. 2128–2136.
- [48] U. SHALIT, D. WEINSHALL, AND G. CHECHIK, *Online learning in the embedded manifold of low-rank matrices*, J. Mach. Learn. Res., 13 (2012), pp. 429–458.
- [49] S. T. SMITH, *Optimization techniques on Riemannian manifold*, in Hamiltonian and Gradient Flows, Algorithms and Control, vol. 3, AMS, Providence, RI, 1994, pp. 113–136.
- [50] K. TOH AND S. YUN, *An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems*, Pacific J. Optimization, 6 (2010), pp. 615–640.
- [51] B. VANDEREYCKEN AND S. VANDEWALLE, *A Riemannian optimization approach for computing low-rank solutions of Lyapunov equations*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2553–2579.
- [52] B. VANDEREYCKEN, *Low-rank matrix completion by Riemannian optimization (extended version)*, preprint, <http://arxiv.org/abs/1209.3834>.
- [53] Z. WEN, W. YIN, AND YIN ZHANG, *Solving a low-rank factorization model for matrix completion by a non-linear successive over-relaxation algorithm*, Math. Program. Comput., 4 (2012), pp. 333–361.