

Preconditioned Conjugate Gradient Algorithms for Graph Regularized Matrix Completion

Shuyu Dong¹, P.-A. Absil¹ and Kyle A. Gallivan² *

1- UCLouvain, ICTEAM Institute
Avenue Georges Lemaître 4 bte L4.05.01, 1348 Louvain-la-Neuve, Belgium

2- Florida State University, Department of Mathematics
208 Love Building, 1017 Academic Way, Tallahassee FL 32306-4510, USA

Abstract. Low-rank matrix completion is the problem of recovering the missing entries of a data matrix by using the assumption that a good low-rank approximation to the true matrix is possible. Much attention has been paid recently to exploiting correlations between the column/row entities through side information to improve the matrix completion quality. In this paper, we propose an efficient algorithm for solving the low-rank matrix completion with graph-based regularizers. Experiments on synthetic data show that our approach achieves significant speedup compared to the alternating minimization scheme.

1 Introduction

Low-rank matrix completion arises in applications such as recommender systems, forecasting and imputation of data. In certain situations, regularization methods, in addition to the low-rank assumption, that exploit other properties of the data matrix are needed since these (low-rank) data often come with other structures. Graph-based regularization methods [1, 2, 3, 4, 5] have gained much attention recently. These regularization methods model the pairwise similarities between the entries of the data matrix via undirected and weighted graphs and apply the graph-based regularizers to the low-rank matrix completion problem. Rao et al. [4] consider recovering a partially observed matrix by solving the following *graph-regularized* matrix factorization program

$$\underset{G \in \mathbb{R}^{m \times r}, H \in \mathbb{R}^{n \times r}}{\text{minimize}} \quad f(G, H) := \frac{1}{2} \|\mathcal{P}_\Omega(GH^T - M)\|_F^2 + \mathcal{R}_{L_r}(G) + \mathcal{R}_{L_c}(H), \quad (1)$$

where $\mathcal{P}_\Omega(\cdot)$ is the mask operator that reveals only the entries of $M \in \mathbb{R}^{m \times n}$ on an index set $\Omega \subset \llbracket m \rrbracket \times \llbracket n \rrbracket$ of size $|\Omega| \ll mn$ and the regularizers are defined according to the penalty function $\mathcal{R}_L(Q) := \frac{\alpha}{2} (\|Q\|_F^2 + \gamma \text{Tr}(Q^T L Q))$ for some given graph Laplacian matrix L . The graph Laplacian matrices L_r and L_c incorporate the pairwise correlations or similarities between the rows or columns

*This work was supported by (i) the Fonds de la Recherche Scientifique – FNRS and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project no 30468160, (ii) “Communauté française de Belgique - Actions de Recherche Concertées” (contract ARC 14/19-060). Part of this work was performed while the third author was a visiting professor at UCLouvain, funded by the Science and Technology Sector, and with additional support by the Netherlands Organization for Scientific Research.

of the data matrix M . Rao et al. [4] showed that problem (1) is equivalent to a weighted nuclear norm minimization problem and that the resulting low-rank approximation has a smaller upper bound for the recovery error than the standard nuclear norm minimization problem. In this previous work, an alternating minimization method is developed for solving the problem (1).

In this paper, we propose to solve problem (1) by using a nonlinear conjugate gradient method on the matrix product space $\mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}$, which is substantially different from alternating minimization algorithms for matrix factorization. While both schemes enjoy the benefit of a reduced computational complexity compared to nuclear norm minimization methods, our method on the matrix product space is particularly efficient, in comparison with alternating minimization, for solving the graph-regularized problem (1) when the true low-rank matrix has entries correlated in a manner close to the given graph information.

2 Preliminaries

Throughout this paper, we consider undirected graphs with nonnegative weights to model pairwise similarities between rows and columns of a data matrix. We denote by U the orthonormal matrix of eigenvectors and Λ the diagonal matrix of corresponding eigenvalues of a given graph Laplacian matrix L such that $L = U\Lambda U^T$. As in Rao et al. [4, §5.1], given the Laplacian matrices $L_r \in \mathbb{R}^{m \times m}$ and $L_c \in \mathbb{R}^{n \times n}$ of two graphs, we focus on a type of low-rank matrix exhibiting some row-wise and column-wise similarities that are prescribed by the graph information via L_r and L_c (or their eigenpairs) as follows

$$M^* = A_r Z^* A_c^T, \text{ with } Z^* = P^* Q^{*T}, \quad (2)$$

where $(P^*, Q^*) \in \mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}$ are random matrices of rank $r \ll \min(m, n)$ whose columns are i.i.d. Gaussian vectors. In (2), $A_r = U_r g(\Lambda_r)$, $A_c = U_c g(\Lambda_c)$ and g is a function acting elementwise on a diagonal matrix $\Lambda = \text{Diag}((\lambda_i))$ ¹ such that $g(\Lambda) = \text{Diag}((1/\sqrt{1 + \gamma\lambda_i}))$ for some hyperparameter $\gamma > 0$. We assume also that the low-rank matrix M^* is only observed on a (small) fraction of entries $\Omega \subset \llbracket m \rrbracket \times \llbracket n \rrbracket$. In order to exploit the given graph information, we add the graph Laplacian-based regularizers as used in related work (*e.g.* [4]) to the standard low-rank matrix completion objective. The model (2) is of particular interest for graph-regularized matrix completion not only because it models a wide range of real data matrices whose entries present pairwise similarities, but also because there is a relationship to the notion of a generalized nuclear norm of M^* . Rao et al. [4] showed that the regularizer $\mathcal{R}_{L_r}(G) + \mathcal{R}_{L_c}(H)$ is related to a generalized nuclear norm of $X = GH^T$, depending on the given graph information through A_r and A_c , and showed via empirical tests with several types of graphs that when such graph Laplacian matrices are known, solving the graph regularized problem (1) instead of the standard (graph-agnostic) matrix

¹The domain of definition of g is restricted to \mathbb{R}_+ since graph Laplacian matrices are known to be positive semi-definite.

completion problem yields better recovery results. The graph-based regularizers can be written as $\mathcal{R}_L(Q) = \frac{\alpha}{2}(\|Q\|_F^2 + \gamma \sum_{i \sim j} \sum_{k=1}^r W_{ij}(Q_{ik} - Q_{jk})^2)$, where W is the graph adjacency matrix such that $L = \text{Diag}(W\mathbf{1}) - W$. Intuitively, the second right hand-side term means that adding the graph Laplacian-based regularizers to the matrix completion objective promotes low-rank solutions that show pairwise similarities according to the given graphs.

3 Nonlinear Conjugate Gradient Descent on the Matrix Product Space

In this section, we propose to solve problem (1) with a nonlinear conjugate gradient descent algorithm by dealing directly with the pair of matrix factors $(G, H) \in \mathcal{M} := \mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}$. Our method produces iterates of a pair of low-rank matrix factors in the sense of a conjugate gradient over the matrix product space. This strategy, in a manner similar to Riemannian optimization on the manifold of fixed-rank matrices for graph-agnostic matrix completion [6, 7, 8], updates the low-rank matrix factors simultaneously.

Based on the Euclidean geometry of the product space \mathcal{M} , the gradient of the objective function of (1) at $X := (G, H) \in \mathcal{M}$ is as follows,

$$\text{grad}f(X) = (SH + \alpha_r(I_m + \gamma_r L_r)G, S^T G + \alpha_c(I_n + \gamma_c L_c)H), \quad (3)$$

where $S = P_\Omega(GH^T - M)$.

Adapting the approach of Mishra et al. [9], we use the diagonal blocks of the Hessian of f to construct a preconditioner for our gradient-based algorithms. The *Hessian* of f at $X \in \mathcal{M}$ acting on the tangent vector $\xi := (\xi_G, \xi_H) \in \mathcal{T}_X \mathcal{M} \equiv \mathcal{M}$ is

$$\nabla^2 f(X)[\xi] = \begin{pmatrix} P_\Omega(G\xi_H^T + \xi_G H^T)H + S\xi_H + \alpha_r(I_m + \gamma_r L_r)\xi_G, \\ P_\Omega(G\xi_H^T + \xi_G H^T)^T G + S^T \xi_G + \alpha_c(I_n + \gamma_c L_c)\xi_H \end{pmatrix}. \quad (4)$$

We define our preconditioner as an approximate inverse of $\xi \mapsto (\frac{\partial^2 f(X)}{\partial G^2}[\xi], \frac{\partial^2 f(X)}{\partial H^2}[\xi])$,

$$\text{Precon}(\xi) := (\xi_G(H^T H)^{-1}, \xi_H(G^T G)^{-1}). \quad (5)$$

For the Armijo-based line search (Algorithm 1, line 6), we compute the initial guess for the step size via the following exact line minimization:

$$t^* = \underset{t>0}{\text{argmin}} f(G + t\eta_G, H + t\eta_H) - f(G, H). \quad (6)$$

The solution to t^* is selected from the real positive roots of the derivative of this quartic function, which can be computed easily. We choose our initial rank r matrix via the r -SVD of the zero filled matrix $M_0 := M_{ij}$ for $(i, j) \in \Omega$ and 0 otherwise.

Based on the above elements, our algorithm is as follows,

Algorithm 1 (GRMC) Graph-Regularized Matrix Completion

Input: Subscripts $\Omega = \{(i_l, j_l) : l = 1, \dots, k\}$, observed matrix entries $P_\Omega(M)$ and rank value r . Parameters α, β .

Output: Matrix estimation \hat{X} .

- 1: *Initialization:* $X^0 := (U_0 \Sigma_0^{1/2}, V_0 \Sigma_0^{1/2})$, $(U_0, \Sigma_0, V_0) \leftarrow r$ -SVD(M_0).
 - 2: **repeat**
 - 3: *Compute the gradient* $\text{grad}f(X^t)$ using (3).
 - 4: *Stopping criterion check* with $\|\text{grad}f(X^t)\|$.
 - 5: *Descent direction:* compute η^t via the conjugate gradient rule (Hestenes-Stiefel+) using $\text{grad}f(X^t)$ and $\text{grad}f(X^{t-1})$.
 - 6: *Line search and update:* find $s_t > 0$ via line minimization and Armijo backtracking; $X^{t+1} = X^t + s_t \text{Precon}(\eta^t)$ using (5).
 - 7: $t \leftarrow t + 1$.
 - 8: **until** Stopping criterion satisfied.
-

4 Numerical Experiments

In this section, we compare our algorithm (with and without preconditioning) with the alternating minimization method (GRALS [4]) on synthetic data generated from the model (2). The $|\Omega|$ observed entries of M^* are uniformly sampled from $\llbracket m \rrbracket \times \llbracket n \rrbracket$ according to a given sample rate. For experiments in (B) and (C), the graph information is incorporated in M^* using A_r , for which a graph Laplacian matrix L_r (for simplicity, we let $L_c = 0$ such that $A_c = I_n$) is generated with the prototypical graph model `Community` using `GSPbox` [10].

(A) *Preliminary tests.* We consider completing synthetic low-rank matrices without any graph information. This corresponds to the case where $L_r = L_c = 0$ so that A_r and A_c reduce to identity matrices in the model (2).




Figure 1: Results per iteration under the setting $m = n = 300, r = 10, |\Omega|/mn = 13\%$. M^* is generated without graph information. (Left) RMSE (test entries); (Middle) Relative error; (Right) norm of the gradient $\|\text{grad}f(X)\|$.

Since the data matrices are graph-agnostic, we disable the graph-based regularizers for both algorithms by setting $\alpha_r = \alpha_c = 0$. Figure 1 shows the RMSE score (on test entries) per iteration given by our algorithm in comparison with

GRALS [4]; the relative error, measured by $\|X^t - M^*\|_F / \|M^*\|_F$, and the norm of $\text{grad}f(X^t)$ per iteration of our algorithm.

(B) *Graph-agnostic algorithms on low-rank matrices with graph information.* Given data matrices generated with the model (2), an “easy” scenario is when the sample rate is sufficiently large such that successful recovery is possible even without the graph regularizers. In such cases, we disable the regularizers by setting $\alpha_r = \alpha_c = 0$ for problem (1). Unlike the results (Figure 1) on graph-agnostic data matrices, Figure 2 shows that our algorithm is much faster than GRALS [4] when M^* is affected by pairwise correlations w.r.t. the model (2).





Figure 2: Results per iteration by MC algorithms ($\alpha_r = \alpha_c = 0$) under the setting $m = n = 300, r = 10, |\Omega|/mn = 30\%$. M^* is generated with non-trivial graph information (contained in A_r).

(C) *GRMC on low-rank matrices with graph information.* For experiments with low sampling rates, we test the graph-regularized algorithms for problem (1) with active parameters $\alpha_r, \gamma_r > 0$, which are selected by grid search. We can observe from Figure 3 that our algorithm is still faster than GRALS [4] in solving the graph-regularized problem (1).






Figure 3: From row 1 to 3: Results per iteration by GR-MC algorithms ($\alpha_r, \gamma_r > 0$) under the setting $m = n = 300$, $|\Omega|/mn = 10.3\%$, 17.0% and 23.4% . M^* is generated with non-trivial graph information (contained in A_r).

Conclusion From the experiments presented above, we observed that our non-linear conjugate gradient method with and without preconditioning on the matrix product space achieves significant speedup in solving the graph-regularized matrix completion problem (1) on synthetic data compared to the alternating minimization method (GRALS [4]).

References

- [1] T. Zhou, H. Shan, A. Banerjee, and G. Sapiro. Kernelized probabilistic matrix factorization: exploiting graphs and side information. *Icdm*, pages 403–414, 2012. *Cited on page 1*
- [2] V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst. Matrix Completion on Graphs. In *NIPS2014 - Robustness in High Dimension*, 2014. *Cited on page 1*
- [3] Z. Zhao, L. Zhang, X. He, and W. Ng. Expert Finding for Question Answering via Graph Regularized Matrix Completion. *IEEE Transactions on Knowledge and Data Engineering*, 27(4):993–1004, apr 2015. *Cited on page 1*
- [4] N. Rao, H.-F. Yu, P. Ravikumar, and I. S. Dhillon. Collaborative Filtering with Graph Information: Consistency and Scalable Methods. In *Advances in Neural Information Processing Systems 28*, pages 2107–2115. 2015. *5 citations: pages 1, 2, 4, 5, and 6*
- [5] H.-F. Yu, N. Rao, and I. S. Dhillon. Temporal Regularized Matrix Factorization for High-dimensional Time Series Prediction. In *Advances in Neural Information Processing Systems 29*, pages 847–855, 2016. *Cited on page 1*
- [6] B. Vandereycken. Low-Rank Matrix Completion by Riemannian Optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013. *Cited on page 3*
- [7] B. Mishra and R. Sepulchre. R3MC: A Riemannian three-factor algorithm for low-rank matrix completion. *Proceedings of the IEEE Conference on Decision and Control*, 2015-Febru(February):1137–1142, 2014. *Cited on page 3*
- [8] L. Cambier and P.-A. Absil. Robust Low-Rank Matrix Completion by Riemannian Optimization. *SISC, Siam Journal on Scientific Computing*, 38(5):1–25, 2015. *Cited on page 3*
- [9] B. Mishra, K. Adithya Apuroop, and R. Sepulchre. A Riemannian geometry for low-rank matrix completion *. Technical report, 2012. *Cited on page 3*
- [10] N. Perraudin, J. Paratte, D. Shuman, L. Martin, V. Kalofolias, P. Vandergheynst, and D. K. Hammond. GSPBOX: A toolbox for signal processing on graphs. *ArXiv e-prints*, aug 2014. *Cited on page 4*