



HAL
open science

Matrix Decomposition on Graphs: A Simplified Functional View

Abhishek Sharma, Maks Ovsjanikov

► **To cite this version:**

Abhishek Sharma, Maks Ovsjanikov. Matrix Decomposition on Graphs: A Simplified Functional View. ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), May 2022, Singapore, France. pp.3358-3362, 10.1109/ICASSP43922.2022.9746243 . hal-04479857

HAL Id: hal-04479857

<https://hal.science/hal-04479857>

Submitted on 27 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MATRIX DECOMPOSITION ON GRAPHS: A SIMPLIFIED FUNCTIONAL VIEW

Abhishek Sharma, Maks Ovsjanikov

LIX, Ecole Polytechnique, IP Paris

ABSTRACT

We propose a simplified functional view of matrix decomposition problems on graphs such as geometric matrix completion. Our unifying framework is based on the key idea that using a reduced basis to represent functions on the product space is sufficient to recover a low rank matrix approximation even from a sparse signal. We validate our framework on several real and synthetic benchmarks where it either outperforms very competitive baselines or achieves competitive results at a fraction of the computational effort of prior work.

Index Terms— Geometric Matrix Completion, Functional Maps, Low Rank Estimators

1. INTRODUCTION

The assumption that high-dimensional data samples lie on or close to a smooth low-dimensional manifold is exploited as a regularizer or prior in many machine learning algorithms. Often, the low-dimensional manifold information is exploited via a graph structure between the data samples. As a result, graphs are often used as a regularizers in various machine learning problems such as dimensionality reduction [10], hashing [16] or matrix completion [11] to name a few. In this article, we focus on geometric matrix completion.

Matrix completion deals with the recovery of missing values of a matrix of which we have only measured a subset of the entries. In general, without any constraints, this problem is ill-posed. However if the rank of the underlying matrix is small, it is common to find the lowest rank matrix that agrees with known measurements [5]. Under this low rank assumption, the problem is very similar to dimensionality reduction and can be rewritten as,

$$\min_{\mathbf{X}} \text{rank}(\mathbf{X}) + \frac{\mu}{2} \|(\mathbf{X} - \mathbf{M}) \odot \mathbf{S}\|_F^2. \quad (1)$$

Here \mathbf{X} stands for the unknown matrix, $\mathbf{M} \in \mathbb{R}^{m \times n}$ for the ground truth matrix, \mathbf{S} is a binary mask representing the input support, and \odot denotes the Hadamard product. Various problems in collaborative filtering can be posed as a *matrix completion* problem [11, 21], where for example the columns and rows represent users and items, respectively, and matrix values represent a score determining the preference of user for

that item. Often, additional structural information is available in the form of column and row graphs representing similarity of users and items, respectively. Most of the prior work that incorporates geometric structure into matrix completion problems is either based on highly engineered frameworks, e.g., [18] or non convex formulation with several hyperparameters [4] thereby making the overall optimization harder to optimize. Instead, our simple formulation based on the functional map representation [19], with a single regularizer, mitigates the problems associated with [4].

Contributions. We propose a convex and smooth matrix decomposition formulation that explicitly imposes and optimizes for a low rank approximation and, as we demonstrate below, is empirically more accurate in recovering a low rank matrix approximation than competitive baselines. Our simplified functional framework also proposes a novel regularization that is shown to be competitive with a combination of several regularizers on various real world datasets. Moreover, we also outline a condition under which a functional map based framework can recover the low rank matrix.

2. RELATED WORK

Matrix completion has been studied with many viewpoints. In this section, we first briefly cover related work on geometric matrix completion and then describe prior work that is directly related to our work.

Geometric matrix completion. A prominent relaxation of the rank operator in equation (1) is to constrain the space of solutions to be smooth w.r.t. some geometric structure of the matrix rows and columns. There exist several prior works on geometric matrix completion that exploit geometric information [3, 11, 21] such as graphs encoding relations between rows and columns. More recent works leverage deep learning on geometric domains [3, 18] to extract relevant information from geometric data such as graphs. While these techniques achieve state-of-the-art results, their design is highly engineered and thus, non-intuitive and often lacks a proper theoretical foundation.

Low Rank Estimators. In classical matrix completion or estimation literature, there is large body of work that assumes the underlying signal matrix \mathbf{M} to be low rank and then tries to estimate it using truncated SVD methods [14, 12, 7, 6, 13]

Thanks to ERC grant EXPROTEA for funding.

as it is the best approximation of a given rank r to the data in the least squares sense. Most of these work estimate this unknown rank and provide bounds on optimality of hard thresholded SVD in an asymptotic framework. Our method is not directly related to these work and we explain it in more detail in the methodology section 4.

Functional Maps. Our work is mainly inspired from the functional map framework [19], which is used ubiquitously in non-rigid shape correspondence, and has been extended to handle challenging partial matching cases, e.g. [15]. This framework has recently been adapted for geometric matrix completion in [4], where the authors propose to build a functional map between graphs of rows and columns. However, they 1) impose several non convex regularization terms each with a scaling hyperparameter and some even with different initialization 2) explore a huge range of hyperparameter space.

3. PRELIMINARIES

In this section, we cover some preliminaries about product graphs and functional maps.

Product graphs Let $\mathcal{G} = (V, E, W)$ be a (weighted) graph with its vertex set V and edge set E and adjacency matrix denoted by W . Graph Laplacian L is given by $L = D - W$, where $D = \text{diag}(W)$ is the *degree matrix*. L is symmetric and positive semi-definite and therefore admits a spectral decomposition $L = \Phi \Lambda \Phi^\top$. Let $\mathcal{G}_1 = (V_1, E_1, W_1)$, $\mathcal{G}_2 = (V_2, E_2, W_2)$ be two graphs, with $L_1 = \Phi \Lambda_1 \Phi^\top$, $L_2 = \Psi \Lambda_2 \Psi^\top$ being their corresponding graph Laplacians. We define the Cartesian product of \mathcal{G}_1 and \mathcal{G}_2 , denoted by $\mathcal{G}_1 \square \mathcal{G}_2$, as the graph with vertex set $V_1 \times V_2$.

Functional maps. Let X be a function defined on $\mathcal{G}_1 \square \mathcal{G}_2$. It can be encoded as a matrix of size $|V_1| \times |V_2|$. Then it can be represented using the bases Φ, Ψ of the individual graph Laplacians, $C = \Phi^\top X \Psi$. In the shape processing community, such C is called a *functional map*, as it is used to map between the functional spaces of \mathcal{G}_1 and \mathcal{G}_2 . One of the advantages of working with the functional map representation C rather than the matrix X is that its size is typically much smaller, and is only controlled by the size of the basis, independent of the number of nodes in G_1 and G_2 , resulting in simpler optimization problems. Moreover, the projection onto a basis also provides a strong regularization, which can itself be beneficial for both shape matching, and, as we show below, matrix completion.

4. LOW RANK MATRIX DECOMPOSITION

We assume that we are given a set of samples in some matrix $M \in \mathbb{R}^{m \times n}$. In addition, we construct two graphs $\mathcal{G}_r, \mathcal{G}_c$, encoding relations between the rows and the columns, respectively. We represent the Laplacians of these graphs and their spectral decompositions by $L_r = \Phi \Lambda_r \Phi^\top$, $L_c = \Psi \Lambda_c \Psi^\top$.

For the matrix completion problem, the matrix M is not completely known so we are also given a binary indicator mask S that indicates 1 for measured samples and 0 for missing ones. We minimize the objective function of the following form:

$$\min_{\mathbf{X}} E_{\text{data}}(\mathbf{X}) + \mu E_{\text{reg}}(\mathbf{X}) \quad (2)$$

with E_{data} denoting a data term of the form

$$E_{\text{data}}(\mathbf{X}) = \|(\mathbf{X} - \mathbf{M}) \odot \mathbf{S}\|_F^2, \quad (3)$$

As observed in [4], we can decompose $\mathbf{X} = \Phi \mathbf{C} \Psi^\top$. Remarkably, the data term itself, as we show in our experiments, when expressed through the functional map i.e. $\mathbf{X} = \Phi \mathbf{C} \Psi^\top$ already recovers low-rank matrices and outperforms the recent approach of [4] on synthetic geometric experiments for matrix completion and obtains competitive results on dimensionality reduction tasks. Before we explain the choice and motivation of our regularizer E_{reg} , we explain next why the data term itself already works remarkably well on rank constrained geometric problems.

4.1. Motivation and Analysis

Suppose that we constrain \mathbf{X} to be a matrix such that $\mathbf{X} = \Phi \mathbf{C} \Psi^\top$ for some matrix \mathbf{C} . Note that if Φ and Ψ have k columns each then \mathbf{C} must be a $k \times k$ matrix. We would like to argue that solving equation (3) under the constraint that $\mathbf{X} = \Phi \mathbf{C} \Psi^\top$ will recover the underlying ground truth signal if it is low rank and satisfies an additional condition that we call basis consistency.

For this suppose that the ground truth hidden signal M has rank r . Consider its singular value decomposition $M = U \Sigma V^\top$. As M has rank r , Σ is a diagonal matrix with r non-zero entries. We will call M *basis-consistent* with respect to Φ, Ψ if the first r left singular vectors U_r (i.e., those corresponding to non-zero singular values) lie in the span of Φ , and the first r right singular vectors V_r lie in the span of Ψ . In other words, there exist some matrices R, Q s.t. $U_r = \Phi R$ (note that this implies $k \geq r$) and $V_r = \Psi Q$. Given this definition, it is easy to see that all basis-consistent matrices with rank $r \leq k$ can be represented by some functional map C . In other words, given Y that is basis-consistent, there is some functional map C s.t. $Y = \Phi C \Psi^\top$. Conversely any $X = \Phi C \Psi^\top$ has rank at most k and must be basis-consistent by construction.

Second, suppose we are optimizing equation (3) under the constraint $\mathbf{X} = \Phi \mathbf{C} \Psi^\top$ and that the optimum, i.e., the ground truth matrix M , is basis-consistent. Then since the energy $E_{\text{data}}(C)$ is convex, given k^2 known samples to fully constrain the corresponding linear system, we are guaranteed to recover the optimum low-rank basis-consistent matrix.

This simple observation suggests that by restricting $\mathbf{X} = \Phi \mathbf{C} \Psi^\top$ and optimizing over the matrices \mathbf{C} instead of \mathbf{X} already provides a strong regularization that can help recover

appropriate low-rank signals even without any other regularization. Further, it avoids solving complex optimization problems involving iterative SVD, since C becomes the only free variable, which can be optimized directly. For problems such as geometric matrix completion, we observe that a weak additional regularization is often sufficient to obtain state-of-the-art results. More formally, we state our result as follows

Proposition 1 *We recover an optimal low rank matrix with high probability as long as the underlying latent matrix X is basis consistent.*

Proof: The proof is based on the main result (Theorem 1 in [5]) in low rank exact matrix recovery method. [5] prove that there is a unique rank k matrix that agrees with the sampled values with high probability and thus, recovers this underlying hidden signal matrix. Our method also recovers a rank k matrix by construction. Since our problem is convex, our method will recover the best rank k matrix that is within the span of the eigenfunctions. If the underlying matrix is basis consistent, then our method will recover the same exact matrix as a low rank exact recovery method (by definition of basis consistency).

Note that we verified on the Synthetic Netflix dataset that our basis consistency condition is indeed satisfied.

4.2. Laplacian Commutativity as a Regularizer

Our E_{reg} contains a single regularization term on the functional map induced between row space and column space described next. We propose to use the simplest possible regularizer, which furthermore leads to a convex optimization problem and can achieve state-of-the-art results. For this we borrow a condition that is prominent in the functional map literature [20]. Namely, in the context of surfaces, the functional map is often expected to *commute with the Laplace-Beltrami operator*: $E_{\text{reg}} = \|C\Lambda_r - \Lambda_c C\|^2$ where Λ_r and Λ_c are diagonal matrices of Laplacian eigenvalues of the source graph (row graph) and target graph (column graph). More broadly, commutativity with the Laplacian imposes a diagonal structure of the functional map, which intuitively promotes preservation of low frequency eigenfunctions used in the basis. In the context of matrix completion this can be interpreted simply as approximate preservation of global low frequency signals defined on the two graphs.

Given these above definitions, our problem defined in equation (2) reduces to

$$\min_C \|(\mathbf{X} - \mathbf{M}) \odot \mathbf{S}\|_F^2 + \mu * \|C\Lambda_r - \Lambda_c C\|^2 \quad (4)$$

where $\mathbf{X} = \Phi C \Psi^\top$

As noted in several works, isometry between two spaces is a key to functional map representation. Assuming isometry between real world graphs is however over optimistic. Thus, one way to work under relaxed isometry condition is to instead

align the eigen basis with additional transformation matrix to achieve diagonal functional map matrix [15, 4]. In practice, we observe faster convergence if we replace C with PCQ^\top , and let all three P , C and Q be free variables.

Differences from SGMC [4] Even though both methods, ours and SGMC build on the functional map framework, there is a fundamental difference between the two. SGMC focus is on high complexity functional map based model (large values of C , multiple resolutions of C, P, Q) and thus, requires a variety of (non-convex) regularizers. In contrast, our core idea is on the low rank matrix recovery based on the functional map based decomposition alone $\mathbf{X} = \Phi C \Psi^\top$ (See 'Ours-FM' baseline in experiments Section 5.2).

To outline the differences more precisely, in addition to Dirichlet energy on the two graphs, [4] also introduces two regularizations on the transformation matrix P, Q . Additionally, [4] also uses a multi-resolution spectral loss named SGMC-Zoomout (SGMC-Z) [17] with its own hyperparameters (step size between different resolutions) besides several scalars to weigh different regularizations.

Hyperparameters The optimization is carried out using gradient descent in Tensorflow [1]. For all experiments, we set μ and the learning rate to be .00001 for all the experiments. We report the size of C explicitly in each experiment below. For geometric matrix completion, we divide the number of available entries in the matrix randomly into training and validation set in a 95 to 5 ratio respectively.

5. EXPERIMENTS

In the first half of this section, we extensively compare between our approach and Spectral Geometric Matrix Completion (SGMC)[4] on a synthetic example of a community structured graphs. In the latter half, we compare with all approaches on various real world recommendation benchmarks. For a fair comparison with [4], we use graphs taken from the synthetic Netflix dataset. Synthetic Netflix is a small synthetic dataset constructed by [11] and [18], in which the user and item graphs have strong community structure. Similar to [4], we use a randomly generated low rank matrix on the product graph $\mathcal{G}_c \square \mathcal{G}_r$ to test the matrix completion accuracy. Synthetic Netflix is useful in conducting controlled experiments to understand the behavior of geometry-exploiting algorithms. We consider the following two baselines:

Ours-FM: This baseline only optimizes for C without any regularization. All results are obtained with C of size 30×30 . This value was chosen after a cross validation from a set of 20, 30, 40. **SGMC:** All results are obtained with their open source code with their optimal parameters.

To evaluate the performance of the algorithms in this section, we report the *root mean squared error* computed on the complement of the training set.

We compare the two approaches on different constraints

Table 1. Testing the dependence on the density of the sampling set for a random rank 10 matrix of size 150×200 .

Density in %	Ours	Ours-FM	SGMC
1	2e-2	2e-2	1e-1
5	8e-7	1e-3	5e-4
10	2e-7	5e-5	2e-4
20	1e-7	2e-5	1e-4

Table 2. Testing the robustness to noisy graphs.

Noise	Ours	Ours-FM	SGMC
5	1e-3	2e-3	5e-3
10	4e-3	3e-3	1e-2
20	6e-3	6e-3	1e-2

ranging from rank of the underlying matrix to the sampling density. Note that optimality bounds for classical matrix completion algorithms also depend on constraints such as sampling density, noise variance etc.

Rank of the underlying matrix. We observe that as the rank increases up to 15, it becomes harder for both methods to recover the matrix. We remark that Ours-FM alone recovers the low rank very effectively. However, on real data, we find the additional regularizer in Ours to be more effective than Ours-FM. We also remark that Ours-FM consistently outperforms SGMC for all rank values. For the training set we used 10% of the points chosen at random (same training set for all experiments summarized in Table 3).

Sampling density. We demonstrate that in the data-poor regime, our regularization is strong enough to recover the matrix, compared to performance achieved by incorporating geometric regularization through SGMC. These experiments are summarized in Table 1. Note that gap between us and SGMC remains high even when the sample density increases to 20%. Even when using 1% of the samples, we perform better than SGMC.

Noisy graphs. We follow the same experimental setup as [4] and perturb the edges of the graphs by adding random Gaussian noise with zero mean and tunable standard deviation to the adjacency matrix. We discard the edges that became negative as a result of the noise, and symmetrize the adjacency matrix. Table 2 demonstrates that our method is

Table 3. Perturbation in the rank of the underlying matrix.

Rank	Ours	Ours-FM	SGMC
5	1e-7	2e-5	1e-4
10	2e-7	2e-5	2e-4
12	5e-7	4e-5	9e-4
15	6e-3	1e-3	1e-2

Table 4. Test error on Flixster and Movielens-100K

Model	Flixster	ML-100K
MC [5]	1.533	0.973
GMC [11]	–	0.996
GRALS [21]	1.245	0.945
RGCNN [18]	0.926	0.929
GC-MC [3]	0.917	0.910
Ours-FM	1.02	1.12
DMF[2]	1.06	0.922
SGMC	0.900	0.912
SGMC-Z	0.888	0.913
Ours	0.888	0.915

robust to substantial amounts of noise in graphs. Surprisingly, Ours-FM demonstrates even stronger resilience to noise.

Runtime Comparison. Our method runs at least 20 times faster than SGMC when compared on synthetic experiments described above. This is not surprising as SGMC involves optimizing various regularizers and with high values of P, C, Q .

In addition to synthetic Netflix, we also validate our method on two more recommender systems datasets for which row and column graphs are available. Movielens-100K [8] contains ratings of 1682 items by 943 users whereas Flixter [9] contains ratings of 3000 items by 3000 users. All baseline numbers, except Ours-FM, are taken from [18] and [4]. In addition to **SGMC** and **SGMC(Z)**, we also compare with **DMF[2]**. This is a matrix factorization approach that was adapted for matrix completion tasks by [4]. Note that this approach does not incorporate any geometric information. We explain some observations from Table 4: First, our baseline, Ours-FM, obtains surprisingly good performance across datasets. This underscores the regularization brought in by the Laplacian eigen-basis of row and column graphs. Second, non geometric model such as DMF shows competitive performance with all the other methods on ML-100K. This suggests that the geometric information is not very useful for this dataset. Third, our proposed algorithm is competitive with the other methods while being simple and interpretable. Lastly, these experimental results validate the effectiveness of our single regularization when compared to the combination of several non-convex regularizations introduced in [4].

Conclusion We establish empirically and theoretically that using a reduced basis to represent a function on the product space of two graphs already provides a strong regularization, that is sufficient to recover a low rank matrix approximation. Moreover, by extensive experiments, we show that our functional map based framework is very competitive when compared to some complex baselines proposed before.

6. REFERENCES

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] S. Arora, N. Cohen, W. Hu, and Y. Luo. Implicit regularization in deep matrix factorization. In *Neurips*, 2019.
- [3] R. v. d. Berg, T. N. Kipf, and M. Welling. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.
- [4] A. Boyarski, S. Vedula, and A. Bronstein. Deep matrix factorization with spectral geometric regularization. *arXiv preprint arXiv:1911.07255*, 2020.
- [5] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.
- [6] S. Chatterjee. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43, 12 2012.
- [7] D. Donoho and M. Gavish. The optimal hard threshold for singular values is. *IEEE Transactions on Information Theory*, 60, 05 2013.
- [8] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):19, 2016.
- [9] M. Jamali and M. Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 135–142. ACM, 2010.
- [10] B. Jiang, C. Ding, B. Luo, and J. Tang. Graph-laplacian pca: Closed-form solution and robustness. In *CVPR*, pages 3492–3498, 2013.
- [11] V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst. Matrix completion on graphs. *arXiv preprint arXiv:1408.1717*, 2014.
- [12] O. Klopp. Rank penalized estimators for high-dimensional matrices. *Electron. J. Stat.*, 5:1161–1183, 2011.
- [13] O. Klopp. Matrix completion by singular value thresholding: sharp bounds. *Electron. J. Stat.*, 9(2):2348–2369, 2015.
- [14] V. Koltchinskii, K. Lounici, and A. B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Annals of Statistics*, 39(5):2302–2329, 10 2011.
- [15] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully spectral partial shape matching. In *Computer Graphics Forum*, volume 36, pages 247–258. Wiley Online Library, 2017.
- [16] W. Liu, J. Wang, S. Kumar, and S. Chang. Hashing with graphs. In *ICML*, 2011.
- [17] S. Melzi, J. Ren, E. Rodola, M. Ovsjanikov, and P. Wonka. Zoomout: Spectral upsampling for efficient shape correspondence. In *Siggraph-Asia*, 2019.
- [18] F. Monti, M. Bronstein, and X. Bresson. Geometric matrix completion with recurrent multi-graph neural networks. In *Advances in Neural Information Processing Systems*, pages 3697–3707, 2017.
- [19] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012.
- [20] M. Ovsjanikov, E. Corman, M. Bronstein, E. Rodolà, M. Ben-Chen, L. Guibas, F. Chazal, and A. Bronstein. Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*, page 9. ACM, 2016.
- [21] N. Rao, H.-F. Yu, P. K. Ravikumar, and I. S. Dhillon. Collaborative filtering with graph information: Consistency and scalable methods. In *Advances in neural information processing systems*, pages 2107–2115, 2015.