



HAL
open science

Matrix Decomposition on Graphs: A Functional View

Abhishek Sharma, Maks Ovsjanikov

► **To cite this version:**

Abhishek Sharma, Maks Ovsjanikov. Matrix Decomposition on Graphs: A Functional View. 2020.
hal-02871840v2

HAL Id: hal-02871840

<https://hal.science/hal-02871840v2>

Preprint submitted on 16 Oct 2020 (v2), last revised 5 Feb 2021 (v3)

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 Functional View

Abhishek Sharma

LIX, Ecole Polytechnique, Institute Polytechnique de Paris

Maks Ovsjanikov

Abstract

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

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, graph is often used as a regularizer in various machine learning problems such as Dimensionality reduction [14], Hashing [19] or Matrix completion [16] to name a few. In this article, we focus on Dimensionality reduction and matrix completion and propose a principled framework that gives a unified solution to both these problems by modelling the extra geometric information available in terms of graphs.

Dimensionality reduction: Given a data matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ with n m -dimensional data vectors, most prior work related to PCA [2] can be broadly categorised in two themes: 1) matrix factorization approach of the classical PCA and its variants 2) matrix sub-

traction approach of robust PCA [7] and its variants. The former learns the projection $Q \in \mathbb{R}^{d \times n}$ of M in a d -dimensional linear space characterized by an orthonormal basis $U \in \mathbb{R}^{m \times d}$. Several followup works [14, 34, 15, 30] have shown that the clustering quality of PCA can be significantly improved by incorporating the data manifold information in the form of some underlying graph structure.

Instead of relying on matrix factorization, the second line of work directly estimates clean low rank approximation \mathbf{X} of data matrix \mathbf{M} by separating noise with a matrix additive model. Along these line, Fast Robust PCA on graphs (FRPCAG[27]) introduces a joint notion of low rankness for the rows and columns of a data matrix and proposes to jointly minimize the Dirichlet energy on the row and column graphs:

$$\min_{\mathbf{X}} \|\mathbf{M} - \mathbf{X}\|_1 + \gamma_1 \text{tr}(\mathbf{X}\mathbf{L}_1\mathbf{X}^\top) + \gamma_2 \text{tr}(\mathbf{X}^\top\mathbf{L}_2\mathbf{X}). \quad (1)$$

Here $\mathbf{L}_1, \mathbf{L}_2$ are Laplacian matrices of graphs built, respectively, from the rows and columns of the data matrix \mathbf{M} . Conceptually, minimizing the Dirichlet energy, $\text{tr}(\mathbf{X}\mathbf{L}_1\mathbf{X}^\top)$, promotes smoothness of \mathbf{X} by penalizing high frequency components of the signals on corresponding graphs. The authors of FRPCAG [27] demonstrate theoretically that under certain assumptions this minimization is connected with the spectrum of the underlying low rank matrix \mathbf{X} . Building on this idea, we instead directly constrain the low rank approximation by decomposing it using the first few eigenvectors of row and column graph Laplacians $\mathbf{X} = \Phi\mathbf{C}\Psi^\top$ and optimizing over the coupling matrix \mathbf{C} only. Our approach, similar in spirit to the matrix factorization approach [5, 3], leads to explicit control over the resulting rank of the matrix, and thereby, superior performance and significantly simpler optimization problems.

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 underlying matrix is small, the number of degrees of freedom decreases and thus, it is common to

find the lowest rank matrix that agrees with known measurements [8]. Under this low rank assumption, the problem is very similar to dimensionality reduction and can be rewritten as,

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

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 [16, 24], 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., [21] or non convex formulation with several hyperparameters [5] thereby making the overall optimization harder to optimize. Instead, our simple formulation based on the functional map representation [22], consisting of a single regularizer, mitigates the problems associated with [5].

Contributions. Our contributions are threefold. First, we propose a novel unified view of geometric matrix completion and graph regularized dimensionality reduction that is convex and smooth. Second, conceptually, our matrix decomposition formulation 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. Third, we propose a novel regularization inspired from the functional map literature that is shown to be competitive with a combination of several regularizers on various real world datasets.

2 Related work

Matrix completion and graph regularized PCA have been studied with many viewpoints and thus, exhaustive coverage of prior work is beyond the scope of this paper. In this section, we first briefly cover related work and then describe prior work that is directly related to our work. We refer to [27] for more details on PCA and related formulations.

Geometric matrix completion. A prominent relaxation of the rank operator in Eq. (2) 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 [4, 16, 24] such as graphs encoding relation between rows and columns. More recent work leverages deep learning on geometric domain [4, 21] to extract relevant information from geometric data such as graphs. As argued in [5], while these techniques achieve state-of-the-art results, their design is highly engineered and thus, non-intuitive.

Graph Regularized Dimensionality Reduction.

Jiang et. al. proposed Graph Laplacian PCA (GLPCA) [14] which imposes the graph regularization of principal components using the Dirichlet term for clustering in the low dimensional space. Similarly, the models proposed in [14, 34, 15, 30] leverage the graph structure to learn enhanced class structures. All these methods still suffer from non-convexity [14, 15, 30]. RPCAG [26] is convex but uses the nuclear norm relaxation that involves an expensive SVD step inhibiting its scalability to large datasets.

The idea of using two graph regularization terms has also been applied in co-clustering [11], Non negative matrix factorization [28] and more recently in the context of low-rank representation [33]. The co-clustering & NMF based models which use such a scheme [11], [28] suffer from non-convexity and the works of [33] use a nuclear-norm formulation making it difficult to scale. Note that there also exist methods that learn a union of low dimensional subspaces where each class belongs to a different subspace [9, 31] but they are not directly related to our approach.

Functional Maps. Our work is mainly inspired from the functional map framework [22], which is used ubiquitously in non-rigid shape correspondence, and has been extended to handle challenging partial matching cases, e.g. [18]. This framework has recently been adapted for geometric matrix completion in [5], 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. Moreover, their framework is tailored towards geometric matrix completion and does not target separability of data in some lower dimensional 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 \mathbf{L}

is given by $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where $\mathbf{D} = \text{diag}(\mathbf{W})$ is the *degree matrix*. \mathbf{L} is symmetric and positive semi-definite and therefore admits a spectral decomposition $\mathbf{L} = \mathbf{\Phi}\mathbf{\Lambda}\mathbf{\Phi}^\top$. It is well-known that spectrum of the Laplacian contains the structural information about the graph [29]. Let $\mathcal{G}_1 = (V_1, E_1, W_1)$, $\mathcal{G}_2 = (V_2, E_2, W_2)$ be two graphs, with $\mathbf{L}_1 = \mathbf{\Phi}\mathbf{\Lambda}_1\mathbf{\Phi}^\top$, $\mathbf{L}_2 = \mathbf{\Psi}\mathbf{\Lambda}_2\mathbf{\Psi}^\top$ being their corresponding graph Laplacians. The bases $\mathbf{\Phi}$, $\mathbf{\Psi}$ can be used to represent functions on these graphs. 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 $V_1 \times V_2$, on which two nodes (u, u') , (v, v') are adjacent if either $u = v$ and $(u', v') \in E_2$ or $u' = v'$ and $(u, v) \in E_1$.

Functional maps. Let \mathbf{X} be a function defined on $\mathcal{G}_1 \square \mathcal{G}_2$. It can be encoded as a matrix of size or $|V_1| \times |V_2|$. Then it can be represented using the bases $\mathbf{\Phi}$, $\mathbf{\Psi}$ of the individual graph Laplacians, $\mathbf{C} = \mathbf{\Phi}^\top \mathbf{X} \mathbf{\Psi}$. In the shape processing community, such \mathbf{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 \mathbf{C} rather than the matrix \mathbf{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 \mathcal{G}_1 and \mathcal{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. For example, given two functions, $\mathbf{x} = \mathbf{\Phi}\boldsymbol{\alpha}$ on \mathcal{G}_1 and $\mathbf{y} = \mathbf{\Psi}\boldsymbol{\beta}$ on \mathcal{G}_2 , one can use \mathbf{C} to map between their representations $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, i.e., $\boldsymbol{\alpha} = \mathbf{\Phi}^\top \mathbf{x} = \mathbf{C}\mathbf{\Psi}^\top \mathbf{y} = \mathbf{C}\boldsymbol{\beta}$.

4 Low Rank Matrix Decomposition on Graphs

We assume that we are given a set of samples in some matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$. In addition, we construct two graphs $\mathcal{G}_r, \mathcal{G}_c$, encoding relations between the rows and the columns, respectively. For matrix completion problem, this matrix \mathbf{M} is not completely known so we are also given a binary indicator mask \mathbf{S} that indicates 1 for measured samples and 0 for missing ones. We represent the Laplacians of these graphs and their spectral decompositions by $\mathbf{L}_r = \mathbf{\Phi}\mathbf{\Lambda}_r\mathbf{\Phi}^\top$, $\mathbf{L}_c = \mathbf{\Psi}\mathbf{\Lambda}_c\mathbf{\Psi}^\top$. We minimize the objective function of the following form:

$$\min_{\mathbf{X}} E_{\text{data}}(\mathbf{X}) + \mu E_{\text{reg}}(\mathbf{X}) \quad (3)$$

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

$$E_{\text{data}}(\mathbf{X}) = \|(\mathbf{X} - \mathbf{M}) \odot \mathbf{S}\|_F^2, \quad (4)$$

As observed in [5], we can decompose $\mathbf{X} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$. Remarkably, the data term itself, as we show in our

experiments, when expressed through the functional map i.e. $\mathbf{X} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$ already recovers low-rank matrices and outperforms the recent approach of [5] 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

Our first observation is that by using a reduced basis to represent a function \mathbf{X} on the product space $\mathcal{G}_1 \square \mathcal{G}_2$ already provides a strong regularization, which can be sufficient to recover a low rank matrix approximation even from a sparse signal.

Specifically, suppose that we constrain \mathbf{X} to be a matrix such that $\mathbf{X} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$ for some matrix \mathbf{C} . Note that if $\mathbf{\Phi}$ and $\mathbf{\Psi}$ have k columns each then \mathbf{C} must be a $k \times k$ matrix. We would like to argue that solving Eq. (4) under the constraint that $\mathbf{X} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$ 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 \mathbf{M} has rank r . Consider its singular value decomposition $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$. If \mathbf{M} has rank r , then $\mathbf{\Sigma}$ is a diagonal matrix with r non-zero entries. We will call \mathbf{M} *basis-consistent* with respect to $\mathbf{\Phi}$, $\mathbf{\Psi}$ if the first r left singular vectors U_r (i.e., those corresponding to non-zero singular values) lie in the span of $\mathbf{\Phi}$, and the first r right singular vectors V_r lie in the span of $\mathbf{\Psi}$. In other words, there exist some matrices \mathbf{R}, \mathbf{Q} s.t. $U_r = \mathbf{\Phi}\mathbf{R}$ (note that this implies $k \geq r$) and $V_r = \mathbf{\Psi}\mathbf{Q}$.

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 \mathbf{C} . In other words, given \mathbf{Y} that is basis-consistent, there is some functional map \mathbf{C} s.t. $\mathbf{Y} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$. Conversely any $\mathbf{X} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$ has rank at most k and must be basis-consistent by construction.

Second, suppose we are optimizing Eq (4) under the constraint $\mathbf{X} = \mathbf{\Phi}\mathbf{C}\mathbf{\Psi}^\top$ and that the optimum, i.e., the ground truth matrix \mathbf{M} , is basis-consistent. Then since the energy $E(\mathbf{C})$ is convex, given enough known samples to fully constrain the corresponding linear system, we are guaranteed to recover the optimum low-rank basis-consistent matrix. We note briefly that the argument above can also be made approximate, when the ground truth matrix is not exactly, but only approximately basis consistent, by putting appropriate error bounds.

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 \mathbf{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.

Connections and Differences to FRPCAG [27]

We do not target the Robust PCA problem [7] as done in FRPCAG. FRPCAG obtains a low rank approximation by minimizing Dirichlet energy on the two graphs and thereby, implicitly obtains a low rank approximation. In contrast, we explicitly factorize the data matrix. As shown in our experiments below, this explicit control over the resulting low rank of matrix, by optimizing over \mathbf{C} , yields superior clustering results over FRPCAG.

4.2 Functional Regularization

Our E_{reg} contains a single regularization term on the functional map induced between row space and column space described next.

Laplacian Commutativity as a Regularizer 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 [23]. Namely, in the context of surfaces, the functional map is often expected to *commute with the Laplace-Beltrami operator*:

$$E_{\text{reg}} = \|\mathbf{C}\mathbf{\Lambda}_r - \mathbf{\Lambda}_c\mathbf{C}\|^2, \quad (5)$$

where $\mathbf{\Lambda}_r$ and $\mathbf{\Lambda}_c$ are diagonal matrices of Laplacian eigenvalues of the source graph (row graph) and target graph (column graph).

For shape matching problems, this constraint helps to find better mappings because functional maps that commute with the Laplacian must arise from near isometric point-to-point correspondences [25, 22]. 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

Eq. (3) reduces to

$$\min_{\mathbf{C}} \|\mathbf{X} - \mathbf{M}\|_{\odot}^2 + \|\mathbf{C}\mathbf{\Lambda}_r - \mathbf{\Lambda}_c\mathbf{C}\|^2 \quad (6)$$

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

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 [18, 5]. In practice, we observe faster convergence if we replace \mathbf{C} with $\mathbf{P}\mathbf{C}\mathbf{Q}^\top$, and let all three \mathbf{P} , \mathbf{C} and \mathbf{Q} be free variables.

Differences from SGMC [5] In addition to Dirichlet energy on the two graphs, [5] also introduces two regularization on the transformation matrix \mathbf{P} , \mathbf{Q} . Moreover, the authors of [5] also propose a multi-resolution spectral loss named SGMC-Zoomout (SGMC-Z) [20] with its own hyperparameters (step size between different resolutions) besides several scalars to weigh different regularizations.

4.3 Implementation

The optimization is carried out using gradient descent in Tensorflow [1].

Graphs Construction Following [27], we use two types of graphs G_1 and G_2 in our proposed model. The graph G_1 is constructed between the data samples or the columns of the data matrix and the graph G_2 is constructed between the features or the rows of the data matrix. The graphs are undirected and built using a standard K-nearest neighbor strategy. We connect each x_i to its K nearest neighbors x_j where K is 10. This is followed by the graph weight matrix A computation as

$$A_{ij} = \begin{cases} \exp\left(-\frac{\|x_i - x_j\|_2^2}{\sigma^2}\right) & \text{if } x_j \text{ is connected to } x_i \\ 0 & \text{otherwise.} \end{cases}$$

Initialization Similar to [5], we initialize the \mathbf{P} and \mathbf{Q} with an identity matrix with size equal to that of underlying matrix \mathbf{M} corresponding to respective dataset and \mathbf{C} by projecting $\mathbf{X} \odot \mathbf{S}$ on the first eigen vector of \mathbf{L}_c and \mathbf{L}_r .

Hyperparameters Our formulation contains two hyperparameters namely the size of \mathbf{C} and the weighing scalar μ . 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. We set μ and learning rate to be .00001 for all the experiments. Size of \mathbf{C} is our only

Table 1: Clustering Results on Benchmark Datasets.

Dataset	Samples	PCA	LE	GLPCA	GRPCA	FGRPCA	Ours
ORL	400	57	56	68	74	77	79
COIL20	1404	67	56	66	65	68	71
MFEAT	400	82	90	71	80	85	90
BCI	400	52	52	52	53	52	53

tunable hyperparameter and we report it explicitly in each experiment below.

5 Results

This section is divided into subsections. The goal of the subsection 5.1 is to validate our dimension reduction framework for the task of clustering and classification. In the subsection 5.2, we evaluate our model performance for matrix completion problem on both synthetic and real world datasets.

5.1 Graph Regularized Dimensionality Reduction

5.1.1 Datasets

We use 4 well-known benchmarks and perform our clustering experiments on following databases: ORL, BCI, COIL20, and MFEAT. ORL¹ is a face database with small pose variations. COIL20² is a dataset of objects with significant pose changes. MFeat³ consists of features extracted from handwritten numerals whereas BCI database comprises of features extracted from a Brain Computer Interface setup⁴.

5.1.2 Baselines

We compare the clustering performance of our model with 5 other dimensionality reduction models. Apart from classical PCA, the rest of the models exploit graph information.

Models using graph structure: We compare 1) Graph Laplacian PCA (GLPCA)[14] 2) Laplacian Eigenmaps (LE) 3) Robust PCA on graphs RPCAG [26] 4) Fast Robust PCA on graphs FRPCAG [27] 5) Our proposed model. Note that RPCAG and FRPCAG are closest to our approach and known to outperform other graph regularized models such as Manifold Regularized Matrix Factorization (MMF) [34], Non-negative Matrix Factorization (NMF)[17], Graph Regularized Non-negative Matrix Factorization (GNMF) [6]. We

obtain FRPCAG and RPCAG results by running the open source implementation provided by the authors on the four datasets. Note that we run the clustering on all the samples of COIL20 and all the features of MFEAT whereas FRPCAG only use a subset of them in their paper. FRPCAG contains two hyperparameters, namely weighing scalars for Dirichlet energy. For these scalars, we pick the best value from the set (1,10,50,100) based on empirical performance. For PCA, we use the first 40 principal components. For our method, the only hyper-parameter is the dimensionality of matrix C . We pick the best value out of (50, 100). We pre-process the datasets to zero mean and unit standard deviation along the features.

5.1.3 Clustering Metric

We follow the standard evaluation protocol and use clustering purity to evaluate our method. To compute purity, each cluster is assigned to the class which is most frequent in the cluster, and then the accuracy of this assignment is measured by counting the number of correctly assigned and dividing by the total no. of samples. We report the maximum clustering error from 10 runs of k-means and summarize our findings in Table 1.

Table 2: Classification Results on Benchmark Datasets.

Dataset	PCA	LE	FGRPCA	Ours
ORL	63	56	66	68
COIL20	88	78	88	89
MFEAT	97	94	97	97
BCI	52	48	53	55

As shown in Table 1, our model obtains superior or competitive performance over all other baselines. We provide the runtime comparison in the supplement.

5.1.4 Classification

We further evaluate our framework on the classification task on the same 4 datasets. We perform classification with PCA, LE and our data representations using a KNN classifier. We randomly select 30% of labeled data, and use the rest to evaluate. We repeat this 5

¹cl.cam.ac.uk/research/dtg/attachive/facedatabase.html

²cs.columbia.edu/CAVE/software/softlib/coil-20.php

³archive.ics.uci.edu/ml/datasets/Multiple+Features

⁴olivier.chapelle.cc/ssl-book/benchmarks.html

Table 3: Comparative results to test the dependence on the rank of the underlying random matrix of size 150×200

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
20	3e-2	1e-2	5e-2

times and summarize the average classification accuracy in Table 2. Our method obtains competitive accuracy compared to other baselines. PCA representation with first 40 components already provides very competitive classification results on several datasets.

5.2 Geometric Matrix Completion experiments

This section is divided into two subsections. The goal of first subsection 5.2.1 is to extensively compare between our approach and Spectral Geometric Matrix Completion (SGMC)[5] on a synthetic example of a community structured graphs. In the second subsection 5.2.2, we compare with all approaches on various real world recommendation benchmarks. Note that we use SGMC and [5] interchangeably in this section.

5.2.1 Experiments on synthetic datasets

For a fair comparison with [5], we use graphs taken from the synthetic Netflix dataset. Synthetic Netflix is a small synthetic dataset constructed by [16] and [21], in which the user and item graphs have strong community structure. Similar to [5], 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.

Baselines

- **Ours-FM**: This baseline only optimizes for \mathcal{C} without any regularization. All results are obtained with \mathcal{C} of size 30×30 .
- **SGMC**: All results are obtained with their open source code with their optimal parameters.

Test Error. To evaluate the performance of the algorithms in this section, we report the *root mean squared*

Table 4: Comparative results to test 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 5: Comparative results to test the robustness in the presence of 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

error,

$$\text{RMSE}(\mathbf{X}, \mathbf{S}) = \sqrt{\frac{\|(\mathbf{X} - \mathbf{M}) \odot \mathbf{S}\|_F^2}{\sum_{i,j} \mathbf{S}_{i,j}}} \quad (7)$$

computed on the complement of the training set. Here \mathbf{X} is the recovered matrix and \mathbf{S} is the binary mask representing the support of the set on which the RMSE is computed.

We compare the two approaches on different constraints as follows:

Rank of the underlying matrix. We explore the effect of the rank of the underlying matrix, showing that as the rank increases up to 15 to 20, it becomes harder for both methods to recover the matrix. As the rank increases, the reconstruction error increases, but it increases slower for our method than for SGMC. For the training set we used 10% of the points chosen at random (same training set for all experiments summarized in Table 3). We remark that Ours-FM consistently outperforms SGMC for all rank values.

Sampling density. We investigate the effect of the number of samples on the reconstruction error. 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 4. 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 study the effect of noisy graphs on the performance. We follow the same experimental

Table 6: Test error on Flixster and MovieLens-100K

Model	Flixster	ML-100K
MC [8]	1.533	0.973
GMC [16]	–	0.996
GRALS [24]	1.245	0.945
RGCNN [21]	0.926	0.929
GC-MC [4]	0.917	0.910
Ours-FM	1.02	1.12
DMF[3]	1.06	0.922
SGMC	0.900	0.912
SGMC-Z	0.888	0.913
Ours	0.888	0.915

setup as [5] 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 5 demonstrates that our method is robust to substantial amounts of noise in graphs. Surprisingly, Ours-FM demonstrates even stronger resilience to noise.

5.2.2 Results on recommender systems datasets

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 [12] contains ratings of 1682 items by 943 users whereas Flixter [13] contains ratings of 3000 items by 3000 users. All baseline numbers, except Ours-FM, are taken from [21] and [5].

Baselines In addition to **SGMC** and **SGMC(Z)**, we also compare with **DMF**[3]. This is a matrix factorization approach that was adapted for matrix completion tasks by [5]. Note that this approach does not incorporate any geometric information. We explain some observations from Table 6: 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 [5].

6 Conclusion and Future Work

In this work, we provide a novel unified view of geometric matrix completion and graph regularized dimensionality reduction and 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, we propose a novel regularization and show, through extensive experimentation on real and synthetic datasets, that our single regularization is very competitive when compared to the combination of several different regularizations proposed before for geometric matrix completion problem.

Extracting geometric information from graph structured data is a core task in several domains from few shot learning [10], zero shot learning [32] in computer vision, machine learning to knowledge graph based problems in natural language processing since graphs appear everywhere. For future work, we plan to extend our framework to several such large scale problems and also test its robustness to noise and corruptions in input data.

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] H. Abdi and L. J. Williams. Principal component analysis. Wiley Interdisciplinary Reviews: Computational Statistics, 2010.
- [3] S. Arora, N. Cohen, W. Hu, and Y. Luo. Implicit regularization in deep matrix factorization. *arXiv preprint arXiv:1905.13655*, 2019.
- [4] R. v. d. Berg, T. N. Kipf, and M. Welling. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.
- [5] A. Boyarski, S. Vedula, and A. Bronstein. Deep matrix factorization with spectral geometric regularization. *arXiv preprint arXiv:1911.07255*, 2020.
- [6] D. Cai, X. He, J. Han, and T. Huang. Graph regularized nonnegative matrix factorization for

- data representation. *IEEE transactions on pattern analysis and machine intelligence*, 33 8:1548–60, 2011.
- [7] E. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58:11:1–11:37, 2011.
- [8] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.
- [9] E. Elhamifar and R. Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:2765–2781, 2013.
- [10] S. Gidaris and N. Komodakis. Generating classification weights with gnn denoising autoencoders for few-shot learning. *CVPR*, 2019.
- [11] Q. Gu and J. Zhou. Co-clustering on manifolds. In *KDD*, 2009.
- [12] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):19, 2016.
- [13] 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.
- [14] B. Jiang, C. Ding, B. Luo, and J. Tang. Graph-laplacian pca: Closed-form solution and robustness. In *CVPR*, pages 3492–3498, 2013.
- [15] T. Jin, J. Yu, J. You, K. Zeng, C. Li, and Z. Yu. Low-rank matrix factorization with multiple hypergraph regularizer. *Pattern Recognition.*, 48:1011–1022, 2015.
- [16] V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst. Matrix completion on graphs. *arXiv preprint arXiv:1408.1717*, 2014.
- [17] D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.
- [18] 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.
- [19] W. Liu, J. Wang, S. Kumar, and S. Chang. Hashing with graphs. In *ICML*, 2011.
- [20] S. Melzi, J. Ren, E. Rodola, M. Ovsjanikov, and P. Wonka. Zoomout: Spectral upsampling for efficient shape correspondence. *arXiv preprint arXiv:1904.07865*, 2019.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] S. Rosenberg. *The Laplacian on a Riemannian manifold: an introduction to analysis on manifolds*, volume 31. Cambridge University Press, 1997.
- [26] N. Shahid, V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst. Robust principal component analysis on graphs. In *ICCV*, pages 2812–2820, 2015.
- [27] N. Shahid, N. Perraudin, V. Kalofolias, G. Puy, and P. Vandergheynst. Fast robust pca on graphs. *IEEE Journal of Selected Topics in Signal Processing*, 10:740–756, 2016.
- [28] F. Shang, L. Jiao, and F. Wang. Graph dual regularization non-negative matrix factorization for co-clustering. *Pattern Recognition.*, 45:2237–2250, 2012.
- [29] D. Spielman. Spectral graph theory. *Lecture Notes, Yale University*, pages 740–0776, 2009.
- [30] L. Tao, H. Ip, Y. Wang, and X. Shu. Low rank approximation with sparse integration of multiple manifolds for data representation. *Applied Intelligence*, 42:430–446, 2014.
- [31] R. Vidal and P. Favaro. Low rank subspace clustering (lrsc). *Pattern Recognit. Lett.*, 43:47–61, 2014.
- [32] X. Wang, Y. Ye, and A. Gupta. Zero-shot recognition via semantic embeddings and knowledge graphs. *CVPR*, 2018.
- [33] M. Yin, J. Gao, Z. Lin, Q. Shi, and Y. Guo. Dual graph regularized latent low-rank representation for subspace clustering. *IEEE Transactions on Image Processing*, 24:4918–4933, 2015.
- [34] Z. Zhang and K. Zhao. Low-rank matrix approximation with manifold regularization. *PAMI*, 35:1717–1729, 2013.