

Graph-Based Matrix Completion Applied to Weather Data

Benoît Loucheur

ICTEAM Institute

UCLouvain

Louvain-la-Neuve, Belgium

benoit.loucheur@uclouvain.be

P.-A. Absil

ICTEAM Institute

UCLouvain

Louvain-la-Neuve, Belgium

pa.absil@uclouvain.be

Michel Journée

Dept. Climatology

Royal Meteo. Inst. of Belgium

Brussels, Belgium

michel.journee@meteo.be

Abstract—Low-rank matrix completion is the task of recovering unknown entries of a matrix by assuming that the true matrix admits a good low-rank approximation. Sometimes additional information about the variables is known, and incorporating this information into a matrix completion model can lead to a better completion quality. We consider the situation where information between the column/row entities of the matrix is available as a weighted graph. In this framework, we address the problem of completing missing entries in air temperature data recorded by weather stations. We construct test sets by holding back data at locations that mimic real-life gaps in weather data. On such test sets, we show that adequate spatial and temporal graphs can significantly improve the accuracy of the completion obtained by graph-regularized low-rank matrix completion methods.

Index Terms—Matrix Completion, Low Rank, Graph, Regularization, Time Series, Air Temperature

I. INTRODUCTION

Matrix completion seeks to find a low-rank matrix when only a fraction of its entries are available. This branch of research has application in numerous research projects such as collaborative filtering (Netflix challenge [1]) [2], traffic sensing [3], [4], image inpainting [5], and gene expression imputation [6].

Sometimes it is known that there is a particular underlying structure that connects the data. Take for example the most famous case of collaborative filtering, the Netflix problem. In this case, each row corresponds to a user and each column to a movie. It has been shown that this rating matrix [7] admits good low-rank approximations. The low-rank structure can be understood from the fact that movies can be organized in a few *genres*, which are rated similarly by the same kinds of users. However, this low rank property is not the only form of structure; there can also be a graphical structure that represents connections between the row/column entities. In such a representation, the rows, resp. columns, are modeled as the nodes of a weighted undirected graph where the weights represent the similarity between the nodes. For example, in the case of the Netflix problem, one can create a graph that connects movies with a weight equal to the number of common actors between those movies.

Missing data in weather time series can appear for various reasons, usually related to sensors deficiencies or communication intermittency. It was recently shown in [8] that matrix

completion methods are able to outperform the state of the art for the completion of missing data in daily extreme temperature series from a network of weather stations.

In this paper, we are interested in the application of graph-regularized matrix completion on air temperature data recorded with a 10-minute time resolution provided by the Royal Meteorological Institute (RMI) of Belgium. We show that adding a regularization by spatial and temporal graphs considerably improves the accuracy of low-rank matrix completion. We also investigate how the pattern of missing data influences the outcomes. A comparison of performance with the state of the art is also performed.

Notation: Lowercase and uppercase boldface letters stand for vectors and matrices respectively. $\|\mathbf{A}\|_F$ is the Frobenius norm of the matrix \mathbf{A} . $\|\mathbf{x}\|_*$ denotes the ℓ_1 norms of vector \mathbf{x} . Finally, \mathbf{I} is the identity matrix and $\mathbf{0}$ is the zero matrix.

II. BACKGROUND AND RELATED WORKS

Matrix completion is the task of recovering all the entries of a matrix \mathbf{M} by observing only a subset Ω of them. Formally, given a matrix \mathbf{M} of size $m \times n$, we have only access to $|\Omega| \ll m \cdot n$ entries and the goal is to predict the remaining unobserved ones in $\bar{\Omega}$.

A. Convex Models

A classical version of the matrix completion problem consists in finding a minimum rank matrix \mathbf{X} which is equal to the observation matrix \mathbf{M} in the set Ω :

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \quad & \text{rank}(\mathbf{X}), \\ \text{s.t.} \quad & \mathbf{M}_{ij} = \mathbf{X}_{ij}, \forall i, j \in \Omega. \end{aligned} \quad (1)$$

Unfortunately, solving such problem is hard [9], and moreover this problem formulation is inadequate when the available data matrix is not exactly low rank due, e.g., to measurement noise. In order to reduce the complexity of this problem, [9] proposed to deal with the tightest possible convex relaxation of the rank operator, which is called the nuclear norm. The relaxed formulation is as follows:

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \quad & \|\mathbf{X}\|_* \\ \text{s.t.} \quad & \mathbf{M}_{ij} = \mathbf{X}_{ij}, \forall i, j \in \Omega. \end{aligned} \quad (2)$$

The model can then be generalized to take into account the measurement noise. The constraint is then relaxed and the optimization problem becomes:

$$\min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \frac{1}{2} \|\mathcal{P}_\Omega(\mathbf{M}) - \mathcal{P}_\Omega(\mathbf{X})\|_F^2 + \lambda \|\mathbf{X}\|_*, \quad (3)$$

where $\lambda > 0$ is a regularization parameter.

B. Factorized Models

When the sparse matrix \mathbf{X} is large, which is usually the case for recommendation systems, it is more efficient to impose a low-rank structure on \mathbf{X} by representing it in factorized form, $\mathbf{X} = \mathbf{A}\mathbf{B}^T$, where \mathbf{A} and \mathbf{B} have few columns.

$$\min_{\substack{\mathbf{A} \in \mathbb{R}^{m \times r} \\ \mathbf{B} \in \mathbb{R}^{n \times r}}} \frac{1}{2} \|\mathcal{P}_\Omega(\mathbf{M}) - \mathcal{P}_\Omega(\mathbf{A}\mathbf{B}^T)\|_F^2 + \frac{\lambda_a}{2} \|\mathbf{A}\|_F + \frac{\lambda_b}{2} \|\mathbf{B}\|_F, \quad (4)$$

where $\lambda_a, \lambda_b > 0$ are two regularization parameters and r is the target rank of the matrix.

Formulation (3) and (4) are closely related in view of the following property of the nuclear norm [10], [11]:

$$\|\mathbf{X}\|_* = \min_{\substack{\mathbf{A} \in \mathbb{R}^{m \times r} \\ \mathbf{B} \in \mathbb{R}^{n \times r}}} \|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2 \quad (5) \\ \text{s.t. } \mathbf{X} = \mathbf{A}\mathbf{B}^T.$$

C. Graph-based Regularization

For simplicity, we present the case of a graph on the rows of \mathbf{M} ; since the columns of \mathbf{M} are the rows of \mathbf{M}^T , the extension to a graph on the columns of \mathbf{M} is direct. Hence assume the availability of a graph $\mathcal{G}^a = (\mathbf{V}^a, \mathbf{E}^a, \mathbf{W}^a)$ on the rows with vertices $\mathbf{V}^a = \{1, \dots, m\}$, edges $\mathbf{E}^a \subseteq \mathbf{V}^a \times \mathbf{V}^a$ and non-negative weights on the edges represented by the symmetric $m \times m$ matrix \mathbf{W}^a . If there is an edge between the nodes i and j , then $\mathbf{W}_{ij}^a = \mathbf{W}_{ji}^a \neq 0$. This information in the form of a graph can then be incorporated into the minimization (6) using the following term [12]:

$$\frac{1}{2} \sum_{i,j} \mathbf{W}_{ij}^a \|\mathbf{a}_i - \mathbf{a}_j\|_2^2 = \text{Tr}(\mathbf{A}^T \mathbf{Lap}(\mathbf{W}^a) \mathbf{A}), \quad (6)$$

where $\mathbf{Lap}(\mathbf{W}^a) = \mathbf{D}^a - \mathbf{W}^a$ is the graph Laplacian of \mathbf{W}^a , \mathbf{D}^a is the diagonal degree matrix, with $\mathbf{D}_{i,i}^a = \sum_{j \neq i} \mathbf{W}_{i,j}^a$. The left-hand side term of (6) favors a similarity between the rows \mathbf{a}_i and \mathbf{a}_j of \mathbf{A} whenever \mathbf{W}_{ij}^a is large. The reasoning is the same for the graph on the columns \mathcal{G}^b .

In view of the above, the problem of matrix completion over graphs can be formulated as follows [13]:

$$\min_{\substack{\mathbf{A} \in \mathbb{R}^{m \times r} \\ \mathbf{B} \in \mathbb{R}^{n \times r}}} \frac{1}{2} \|\mathcal{P}_\Omega(\mathbf{M}) - \mathcal{P}_\Omega(\mathbf{A}\mathbf{B}^T)\|_F^2 + \frac{\lambda_a}{2} \left\{ \text{Tr}(\mathbf{A}^T \mathbf{Lap}(\mathbf{W}^a) \mathbf{A}) + \text{Tr}(\mathbf{B}^T \mathbf{Lap}(\mathbf{W}^b) \mathbf{B}) \right\} + \frac{\lambda_a}{2} \|\mathbf{A}\|_F^2 + \frac{\lambda_b}{2} \|\mathbf{B}\|_F^2. \quad (7)$$

This formulation having two terms of regularization by graphs, it is more commonly called Graph-Regularized Matrix Completion (GRMC). A method called GRALS, to address (7) by applying a conjugate gradient method alternatively to \mathbf{A} and \mathbf{B} , is provided in [14].

D. State of the Art

In [8], the following methods were compared on weather data completion tasks.

IDW (Inverse Distance Weighting) is a widely-used method for estimating missing data in meteorological and other time series datasets. The approach uses a weighted average of neighboring observed data points, with weights computed by the inverse distance between the missing data points and the neighbors. IDW is simple, computationally efficient, and can produce accurate estimation when the underlying data has a smooth spatial structure. However, its performance is affected by outliers or uneven data distributions.

PCA [15] (Principal Component Analysis) is another method for the estimation of missing data. The approach involves decomposing the data into its principal components, which capture the most important patterns or variation in the data. Then, the missing data can be estimated using a linear combination of the principal components based on the observed values in the corresponding time periods.

SoftImpute [16] is a matrix completion algorithm that uses the Singular Value Decomposition (SVD) to estimate the low-rank structure and then applies a soft-thresholding operator to shrink small singular values. The formulation of SoftImpute is as follow:

$$\min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \|\mathbf{X}\|_*, \quad (8) \\ \text{s.t. } \mathcal{P}_\Omega(\mathbf{M}) - \mathcal{P}_\Omega(\mathbf{X}) \leq \delta.$$

Finally, RTRMC [17] addresses the matrix completion problem by applying a trust-region methods to the following optimization problem on the Grassmann manifold $\mathcal{G}^{m \times r}$:

$$\min_{\mathbf{A} \in \mathcal{G}^{m \times r}} \min_{\mathbf{B} \in \mathbb{R}^{n \times r}} \frac{1}{2} \|\mathcal{P}_\Omega(\mathbf{M}) - \mathcal{P}_\Omega(\mathbf{A}\mathbf{B}^T)\|_F^2 + \frac{\lambda^2}{2} \|\mathcal{P}_{\bar{\Omega}}(\mathbf{A}\mathbf{B}^T)\|_F^2, \quad (9)$$

where $\bar{\Omega}$ is the complement of the set Ω .

III. GRAPH GENERATION

In addition to the estimation of the classical hyperparameters of a matrix completion method, we also need to generate two graphs, one for the rows of the matrix (i.e., each timestamp) and one for the columns (i.e., each station). In this subsection, different ways to create our spatial and temporal graphs are described.

A. Spatial Graph

The graph of the matrix columns is in our case the spatial graph. Each node represents a weather station and we are free to create a weighted edge or not between two stations. We have compared different graph possibilities with various levels of complexity.

The most basic approach is the K-Nearest Neighbors (KNN), in which case each station is connected by an edge to its K nearest neighbors using the geographic distance. (see Fig. 1a). It is then possible to add constraints on the generation of the KNN graph by adding as weights for the edges the inverse of the distance between the two stations. Thus, the closer the stations are, the more strongly they are connected (see Fig. 1b). The visualization of the weights in the graph is done via the transparency and thickness of the edges. The more important the weight is, the darker and thicker the edge will be. A third version of the spatial graph consists in adding a constraint of maximum altitude difference between two stations. If the altitude difference between two stations is greater than a chosen threshold value, then this connection is ignored. This avoids the rather special case where two stations are very close spatially but have a large difference in altitude, e.g., more than 100 meters (see Fig. 1c). Such cases can appear in areas with complex orography and can negatively impact the completion of missing data. By comparing the Fig. 1b and Fig. 1c, no changes are observed for the edges in the north of Belgium. This is due to the relatively flat terrain in this geographical area which will then never trigger the altitude limit condition. In contrast, the southern part of Belgium has a much more complex orography which causes changes in the edges.

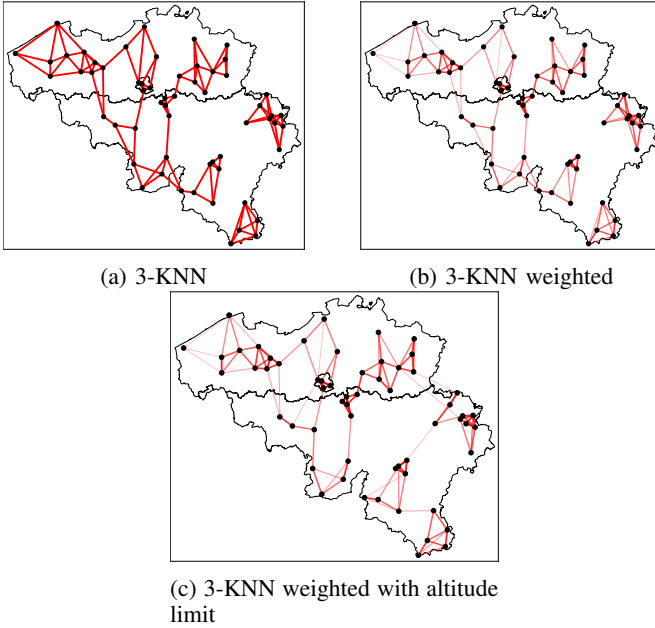


Fig. 1: Spatial Graph examples

B. Temporal Graph

The graph of the matrix rows is in our case the temporal graph. Each node represents a date and to link them we use the notion of lag set \mathcal{L} which is the (repeating) dependency pattern of the graph.

Fig. 2a shows the simplest lag set when $\mathcal{L} = [1]$. This means that the temperature measured at time t will be connected to the temperature at time $t - 10$ minutes and $t + 10$ minutes.

Fig. 2b shows the lag set $\mathcal{L} = [1, 2, 3]$, which creates a link between the current measurement and the three previous and following ones.

The first lag set only tries to capture the immediate relationship or dependency between two consecutive measures, whereas the second lag set tries to capture a slightly longer-term relationship.




Fig. 2: Temporal Graph with different lag sets \mathcal{L}

Once the temporal graph is created, we have the possibility to choose weights for its edges. In our experiments, we chose two different rules for the attribution of weights. The first trivial one is when $w_i = 1$ which allows us to have in a non-weighted graph. The other possibility considered is to use the rule $w_i = \frac{1}{i}$ to emphasize the importance of short-term relationships in the data, as shorter lags often reflect more immediate or direct dependencies, while longer lags may be more influenced by noise or other indirect factors.

IV. MACHINE LEARNING MODEL

All our experiments were performed with data provided by the RMI, consisting of air temperature measurements from 50 automatic weather stations located in Belgium, which provide data with a temporal resolution of 10-min. We consider in this study the data for the period from January 1, 2020 to March 1, 2022. The source code is available at <https://github.com/bloucheur/GRMCWeatherData>. Implementations of the existing algorithms are also publicly available, however the dataset we use is confidential. An open source dataset containing similar data is available.

A. Generation of Missing Weather Data

To evaluate the performance of our model, we have to synthetically create missing entries in the matrix \mathbf{M} . These synthetic missing data will be used to evaluate the error committed by our model once the completion of these missing data is calculated. However, as we will show in our results in Section V, the pattern of missing data significantly impacts the quality of the completion. In our case, two types of patterns are considered. The first one, represented in Fig. 3a, considers that the missing data are in the form of long duration

blocks (between one and 3 days). While the second scenario, represented in Fig. 3b, considers synthetic holes of small duration (10 minutes to 2 hours). It is important to note that in total the number of missing data generated for both scenarios is the same. Furthermore, the number of artificially generated missing data is 10% of the maximum number of entries in the matrix, more formally: $|\Omega_{\text{Block}}| = |\Omega_{\text{Spread}}| = 0.1 * mn$.




Fig. 3: The two scenarios considered, the hatched areas represent the missing data.

Note that this condition is not applied in Fig. 3a and Fig. 3b as these represent an artist's view of our two scenarios.

B. Hyperparameter Tuning

The graph regularized matrix completion model we study, named GRALS [14], has hyperparameters to choose from. In particular, the value of the rank (r) for the factorization, the three constants of regularization (λ_L , λ_a and λ_b) as well as variables of our design for the generation of our graphs which are explained above. We also want to simulate our missing data according to particular patterns as shown in Fig. 3. The machine learning model that corresponds to our expectations is the Monte Carlo Cross-Validation [18]. The advantage of this evaluation technique is that it allows us to simulate missing data according to a pattern that we can define.

We first split our datasets in two, the data from January 1, 2020 to August 31, 2021 are part of the training set, the rest (i.e. September 1, 2021 to March 1, 2022) are part of the testing set. For each experiment, we consider one week of data for all stations, i.e., the observation matrix $\mathbf{M} \in \mathbb{R}^{1009 \times 50}$. In order to perform the hyperparameter tuning, we select 10 weeks without gaps and that do not intersect in the training set. For each of these 10 weeks, we generate 5 patterns of missing data, to create our sparse matrix \mathbf{M} , according to the scenario imposed for the experiment. The optimal set of hyperparameters is determined as the one that results into the lowest RMSE when averaged over the 5×10 folds. It is important to note that the weeks selected for the two scenarios are the same, only the way the missing data are generated differs. For the final evaluation on the test set, the same missing data generation methodology is applied with 5 different weeks and 3 generations of missing data patterns.

V. EXPERIMENTAL RESULTS

A. Comparison of the methods

TABLE I shows the values of RMSE for every method on the test set.

Methods	Block	Spread
IDW	0.61	0.57
PCA	0.78	0.71
SoftImpute	0.50	0.41
RTRMC	0.48	0.50
GRALS	0.51	0.45

TABLE I: Average RMSE (in °C) on the test set for every methods considered. Results for the two types of missing data generation with different constraints applied on the model.

The results indicate that the performance of GRALS, under our experimental conditions, is comparable to the other low-rank matrix completion methods (SoftImpute and RTRMC).

In the remainder of this experimental section, we will conduct ablation studies on GRALS to evaluate how much the achieved RMSE depends on its spatial and temporal graph regularization.

B. Graph Part

TABLE II represents all the hyperparameters of the graph regularized matrix completion model that we have. For each hyperparameter, we define a set of possible and relevant values. Due to the high number of hyperparameters in our model, we generate all possible combinations of hyperparameters, and a sampling without replacement is performed. This approach is based on the RandomizedSearchCV model from Scikit-learn [19].

Hp	Values tested	Block	Spread
r	$\{2, \dots, 20\}$	10	11
λ_L	$\{0.1, 0.01, 0.001, 0.005\}$	0.001	0.001
λ_a	$\{0.1, 0.01, 0.001, 0.005\}$	0.005	0.001
λ_b	$\{0.1, 0.01, 0.001, 0.005\}$	0.005	0.005
K	$\{1, 2, \dots, 5\}$	4	3
weighted	$\{\text{True}, \text{False}\}$	True	True
altitude_limit	$\{\text{True}, \text{False}\}$	True	True
\mathcal{L}	$\{[1], [1, 2], [1, 2, 3], [1, 2, 3, 4, 5]\}$	[1]	[1, 2, 3]

TABLE II: Set of hyperparameters (Hp) with their optimal values for the Block and Spread missing data patterns on the training set.

The last two columns of this table represent the optimal hyperparameters obtained during the training phase for the two scenarios considered.

Whether it is the block or spread scenario, the spatial graph requires a large number of K neighbours and it is important to note that both require a weighted spatial graph coupled with the altitude boundary condition.

TABLE III represents different RMSE values on the test set by imposing or not a priori conditions on some hyperparameters. The Case #1 does not impose any constraints, and we consider these results as a reference.

	GRALS Conditions	Block	Spread
Case #1	No constraints	0.51	0.45
Case #2	$\lambda_L = \lambda_a = \lambda_b = 0$	0.95	0.93
Case #3	$\lambda_a = \lambda_b = 0$	0.90	0.91
Case #4	$\lambda_L = 0$	0.76	0.67
Case #5	$\text{Lap}(\mathbf{W}^b) = \mathbf{0}$	0.71	0.61
Case #6	$\text{Lap}(\mathbf{W}^a) = \mathbf{0}$	0.55	0.6

TABLE III: Average RMSE (in °C) on the test set. Results for the two types of missing data generation with different constraints applied on the model.

Setting all regularization terms to zero ($\lambda_L = \lambda_a = \lambda_b = 0$) drastically decreases the performance of the completion regardless of the missing data pattern scenario.

Case #5 and Case #6 show the impact of considering only a temporal graph ($\text{Lap}(\mathbf{W}^b) = \mathbf{0}$) or only a spatial graph ($\text{Lap}(\mathbf{W}^a) = \mathbf{0}$). A comparison between Case #1 and Case #6 shows that the temporal graph contributes little to the resulting completion performance in the case of the Block scenario. On the other hand, in the Spread scenario, removing the temporal graph significantly impacts the completion with an average RMSE that increases by 35%.

Then, by comparing Case #1 and Case #5, we see that the spatial graph is important for both scenarios. This is a rather expected result: the layout of the weather stations as well as the continuous nature of the temperature in space, reinforces the idea that linking stations close in space could improve the results.

VI. CONCLUSION

It is known [8] that the low-rank matrix completion methods SoftImpute and RTRMC are very effective for completing missing meteorological data. Our experiments (Section V) have shown that a *graph-regularized* low-rank matrix completion method, termed GRALS, is also very effective, and that its graph regularization is essential to obtain competitive results. This opens two avenues of research: (i) enhance RTRMC and SoftImpute with graph regularization and (ii) improve the construction of the graphs.

Currently, our graphs are built using only the metadata at our disposal (i.e., geographic coordinates of the stations as well as their altitudes). In our future research, we will attempt to build adaptive graphs based on the measurements. By defining the notion of distance between two time series, it is possible to perform clustering and thus generate a spatial graph. For the temporal graph, it is possible to fit an autoregressive model on the known data, thus giving us weight values for each weight of the lag set.

REFERENCES

- [1] James Bennett, Stan Lanning, and Netflix, “The Netflix prize,” in *In KDD Cup and Workshop in conjunction with KDD*, 2007.
- [2] Jasson D. M. Rennie and Nathan Srebro, “Fast maximum margin matrix factorization for collaborative prediction,” in *Proceedings of the 22nd International Conference on Machine Learning*, New York, NY, USA, 2005, ICML ’05, p. 713–719, Association for Computing Machinery.
- [3] Rong Du, Cailian Chen, Bo Yang, Ning Lu, Xinping Guan, and Xuemin Shen, “Effective urban traffic monitoring by vehicular sensor networks,” *IEEE Transactions on Vehicular Technology*, vol. 64, no. 1, pp. 273–286, 2015.
- [4] Rong Du, Cailian Chen, Bo Yang, and Xinping Guan, “Vanet based traffic estimation: A matrix completion approach,” in *2013 IEEE Global Communications Conference (GLOBECOM)*, 2013, pp. 30–35.
- [5] Hongyang Xue, Shengming Zhang, and Deng Cai, “Depth image inpainting: Improving low rank matrix completion with low gradient regularization,” *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4311–4320, sep 2017.
- [6] Arnav Kapur, Kshitij Marwah, and Gil Alterovitz, “Gene expression prediction using low-rank matrix completion,” *BMC Bioinformatics*, vol. 17, 06 2016.
- [7] Robert Bell and Yehuda Koren, “Lessons from the Netflix prize challenge,” *SIGKDD Explorations*, vol. 9, pp. 75–79, 12 2007.
- [8] Benoît Loucheur, P.-A. Absil, and Michel Journée, “Gap filling in air temperature series by matrix completion methods,” in *Proceedings of the 30th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN2022)*, 2022, pp. 369–374, i6doc.
- [9] Emmanuel J. Candès and Benjamin Recht, “Exact matrix completion via convex optimization,” *Found. Comput. Math.*, vol. 9, no. 6, pp. 717–772, dec 2009.
- [10] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM Review*, vol. 52, no. 3, pp. 471–501, jan 2010.
- [11] Trevor Hastie, Rahul Mazumder, Jason D. Lee, and Reza Zadeh, “Matrix completion and low-rank SVD via fast alternating least squares,” *Journal of Machine Learning Research*, vol. 16, no. 104, pp. 3367–3402, 2015.
- [12] Mikhail Belkin and Partha Niyogi, “Laplacian eigenmaps for dimensionality reduction and data representation,” *Neural Computation*, vol. 15, no. 6, pp. 1373–1396, 2003.
- [13] Vassilis Kalofolias, Xavier Bresson, Michael M. Bronstein, and Pierre Vandergheynst, “Matrix completion on graphs,” *CoRR*, vol. abs/1408.1717, 2014.
- [14] Nikhil Rao, Hsiang-Fu Yu, Pradeep K. Ravikumar, and Inderjit S. Dhillon, “Collaborative filtering with graph information: Consistency and scalable methods,” in *Advances in Neural Information Processing Systems* 27, 2015.
- [15] Hege Hisdal and Ole Tveito, “Extension of runoff series using empirical orthogonal functions,” *Hydrological Sciences Journal*, vol. 38, pp. 33–49, 02 1993.
- [16] Jian-Feng Cai, Emmanuel J. Candès, and Zuowei Shen, “A singular value thresholding algorithm for matrix completion,” *SIAM Journal on Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [17] N. Boumal and P.-A. Absil, “RTRMC: A Riemannian trust-region method for low-rank matrix completion,” in *Advances in Neural Information Processing Systems 24 (NIPS)*, pp. 406–414, 2011.
- [18] Richard R. Picard and R. Dennis Cook, “Cross-validation of regression models,” *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 575–583, 1984.
- [19] James Bergstra and Yoshua Bengio, “Random search for hyperparameter optimization,” *J. Mach. Learn. Res.*, vol. 13, pp. 281–305, feb 2012.