

Rethinking Symmetric Matrix Factorization: A More General and Better Clustering Perspective

Mengyuan Zhang
Clemson University
Clemson, SC, USA
mengyuz@clemson.edu

Kai Liu
Clemson University
Clemson, SC, USA
kail@clemson.edu

Abstract—Nonnegative matrix factorization (NMF) is widely used for clustering with strong interpretability. Among general NMF problems, symmetric NMF is a special one that plays an important role in graph clustering where each element measures the similarity between data points. Most existing symmetric NMF algorithms require factor matrices to be nonnegative, and only focus on minimizing the gap between similarity matrix and its approximation for clustering, without giving a consideration to other potential regularization terms which can yield better clustering. In this paper, we explore factorizing a symmetric matrix that does not have to be nonnegative, presenting an efficient factorization algorithm with a regularization term to boost the clustering performance. Moreover, a more general framework is proposed to solve symmetric matrix factorization problems with different constraints on the factor matrices.

Index Terms—symmetric matrix factorization, generalized optimization framework, clustering

I. INTRODUCTION

Nonnegative matrix factorization (NMF) [15] problem is formulated as the following: given a data matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}_+^{m \times n}$ containing n observations, each observation denoted as \mathbf{x}_i is an m dimensional vector, where $\mathbb{R}_+^{m \times n}$ denotes the set of $m \times n$ element-wise nonnegative matrices. NMF aims to find a lower-rank matrix approximation represented by:

$$\mathbf{X} \approx \mathbf{U}\mathbf{V}^T, \quad (1)$$

where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k] \in \mathbb{R}_+^{m \times k}$, and $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k] \in \mathbb{R}_+^{n \times k}$. Typically squared Frobenius norm is used to measure the distance between \mathbf{X} and $\mathbf{U}\mathbf{V}^T$, so the objective in NMF is formulated as [10]:

$$\min_{\mathbf{U}, \mathbf{V} \geq 0} \|\mathbf{X} - \mathbf{U}\mathbf{V}^T\|_F^2, \quad (2)$$

where $\|\mathbf{Q}\|_F^2 = \sum_{i,j} \mathbf{Q}(i,j)^2$. Usually, k is assumed to be smaller than $\min\{m, n\}$, thus NMF can be regarded as a lower-rank approximation problem.

Apparently, the NMF paradigm described above conducts clustering based on input data directly and assumes data is well linearly separable. However, for data that lies in a specific manifold (say a certain sphere or two moons), it will yield a poor result. Therefore, graph clustering is introduced to overcome the difficulty based on a matrix that measures the similarity between each data point [12]. The factorization of similarity matrix $\mathbf{A} \in \mathbb{R}_+^{n \times n}$ will yield a lower-rank matrix

$\mathbf{H} \in \mathbb{R}_+^{n \times k}$ which plays a similar role as \mathbf{V} for cluster assignment [7], [12], [20], [22]. Specifically, symmetric NMF formulates the objective as:

$$\min_{\mathbf{H} \geq 0} \|\mathbf{A} - \mathbf{H}\mathbf{H}^T\|_F^2, \quad (3)$$

where k is the number of clusters. Compared to classical NMF, symmetric NMF is more flexible in terms of admitting any reasonable measurement with mixed signs such as cosine similarity [21], [32], [34].

Previous work on symmetric NMF mostly requires that the matrix \mathbf{H} is nonnegative. Therefore, even \mathbf{A} is not explicitly constrained to be non-negative, in practice, it is equivalent to setting the negative elements to be 0. When symmetric NMF is applied to graph clustering, the result is directly obtained from \mathbf{H} while ignoring some other techniques such as graph regularization to promote clustering performance. In this paper, we comprehensively study symmetric matrix factorization with its application in graph clustering. Our contribution is threefold:

- We first extend vanilla symmetric NMF and study a more general case where there is no non-negative constraint on \mathbf{H} and interpret it in a meaningful manner. We propose a very efficient updating algorithm that can be extended to the non-negative case.
- A regularization term is added to boost the clustering performance. Instead of merely focusing on minimizing the objective $\|\mathbf{A} - \mathbf{H}\mathbf{H}^T\|_F^2$, we impose graph regularization term to ensure that data points with higher similarity value share more similar cluster indicators, and vice versa, with \mathbf{H} admitting mixed signs.
- We propose a general framework that can deal with symmetric matrix factorization problems with various constraints by learning the stepsize adaptively.

This paper is organized as follows: Section II describes the motivation behind the problem. Section III provides the problem formulation and optimization methods, including one column-wise update method with fast speed, and one more general framework. Section IV shows the convergence rate of our method. Section V reports the experimental results on both image datasets and text datasets, followed by the Conclusion Section. Our code is available on GitHub.¹

¹<https://github.com/clair-lab/Symmetric-Matrix-Factorization>




Fig. 1: An illustration of symmetric matrix factorization on $\mathbf{A} \in \mathbb{S}^n$. The green block indicates the similarity is high, the corresponding rows (first two) in \mathbf{H} should also be similar, so they are filled with the same color. The red block indicates the similarity is low, where the corresponding rows (5th and 7th, respectively) in \mathbf{H} should be dissimilar, filled with different colors.

II. MOTIVATION

In symmetric NMF for clustering, the objective function (3) is to measure the gap between the original similarity matrix \mathbf{A} and $\mathbf{H}\mathbf{H}^T$, where \mathbf{H} is the clustering assignment matrix with nonnegative constraint. However, most algorithms only aim to minimize the gap $\|\mathbf{A} - \mathbf{H}\mathbf{H}^T\|_F^2$, while ignoring the potential over-fitting which may lower the clustering performance. Following the idea in graph regularization [4], [6]: data points that have high similarity (in \mathbf{A}) should have closer clustering indicators (rows of \mathbf{H}), and vice versa, which is demonstrated as Fig. 1. Accordingly, the regularization term is given by:

$$\min_{\mathbf{H} \geq 0} \sum_{i,j=1}^n \mathbf{A}_{ij} \|\mathbf{h}^i - \mathbf{h}^j\|_2^2, \quad (4)$$

where \mathbf{h}^i denotes i -th row in \mathbf{H} .

Though theoretically, \mathbf{A} admits mixed signs, however, due to the nonnegative constraint on \mathbf{H} , negative elements are treated as 0 after projection and play no role in learning. Therefore, we remove the nonnegative constraint on \mathbf{H} . Naturally, if \mathbf{A}_{ij} is negative, the indicators should be significantly different and $\mathbf{A}_{ij} \|\mathbf{h}^i - \mathbf{h}^j\|_2^2$ remains negatively small, which is in accordance with the spirit of graph regularization. Negative element h_{ij} denotes the **unlikelihood** of i -th data belonging to j -th cluster while positive represents the very **likelihood**. Therefore we formulate the objective with regularization as:

$$\min_{\mathbf{H}} \|\mathbf{A} - \mathbf{H}\mathbf{H}^T\|_F^2 + \lambda \sum_{i,j=1}^n \mathbf{A}_{ij} \|\mathbf{h}^i - \mathbf{h}^j\|_2^2, \quad (5)$$

where λ is the tuning regularization parameter.

To verify whether the regularization term can help to boost the clustering performance, a pilot experiment is conducted on COIL-20 data set [23]. Fig. 2 shows clustering accuracy [33] comparison obtained from the same initialization and constraint (admitting mixed signs on \mathbf{H}) with the only difference being the existence of the regularization term. We can see that the clustering performance is significantly improved by incorporating the regularization term into the objective, especially when the number of clusters grows larger.

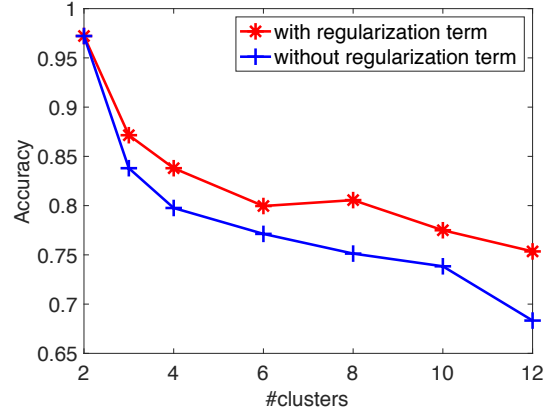


Fig. 2: Accuracy of clustering comparison on COIL20 data set with different cluster numbers.

Moreover, regarding \mathbf{H} , there can be different constraints on it. One can verify that if $\mathbf{H}^T \mathbf{H} = \mathbf{I}$, then it becomes spectral clustering [1], [24]; if all is zero except one element is 1 in each row of \mathbf{H} , then it is K -means; if each row is nonnegative and the sum is 1 [16], [26], then it is the probability distribution to each cluster. Instead of proposing ad-hoc algorithms with various constraints, we will systematically address the problem by providing a general framework that can obtain optimal solutions with updates efficiently.

III. FORMULATION AND ALGORITHM

In this section, we propose our algorithm to solve (5) which enjoys faster convergence than its existing counterparts.

A. Reformulation

By noting the regularization term is in row-wise form, we are to reformulate the objective in a more compact way as:

Theorem 1. (5) is equivalent to:

$$\min_{\mathbf{H}} \|\mathbf{A} - \lambda \mathbf{L} - \mathbf{H}\mathbf{H}^T\|_F^2, \quad (6)$$

where \mathbf{L} is the Laplacian matrix given by $\mathbf{L} := \mathbf{D} - \mathbf{A} \in \mathbb{R}^{n \times n}$ and \mathbf{D} is the degree matrix, which is diagonal defined as: $D(i, i) = \sum_j (A(i, j))$.

Proof of Theorem 1. Following [4], we have:

$$\sum_{i,j=1}^n \mathbf{A}_{ij} \|\mathbf{h}^i - \mathbf{h}^j\|_2^2 = 2\text{tr}(\mathbf{H}^T \mathbf{L} \mathbf{H}). \quad (7)$$

Therefore, (5) is equivalent to

$$\min_{\mathbf{H}} \|\mathbf{A} - \mathbf{H}\mathbf{H}^T\|_F^2 + 2\lambda \text{tr}(\mathbf{H}^T \mathbf{L} \mathbf{H}). \quad (8)$$

By expanding (6), we obtain:

$$\begin{aligned} & \min_{\mathbf{H}} \|\mathbf{A} - \lambda \mathbf{L} - \mathbf{H}\mathbf{H}^T\|_F^2 \\ &= \min_{\mathbf{H}} \text{tr}(\mathbf{A} - \lambda \mathbf{L} - \mathbf{H}\mathbf{H}^T)^T (\mathbf{A} - \lambda \mathbf{L} - \mathbf{H}\mathbf{H}^T) \\ &= \min_{\mathbf{H}} \|\mathbf{A} - \mathbf{H}\mathbf{H}^T\|_F^2 + 2\lambda \text{tr}(\mathbf{H}^T \mathbf{L} \mathbf{H}) \\ &+ \text{terms irrelevant to } \mathbf{H}. \end{aligned} \quad (9)$$

Thus we conclude that solution to (5) is the same as to (6) in terms of optimizing \mathbf{H} . \square

B. Column-Wise Fast Update

Given the new formulation in (6), we now turn to provide detailed updating rule for \mathbf{H} .

We first denote $\mathbf{A} - \lambda \mathbf{L}$ as \mathbf{M} , obviously \mathbf{M} is still symmetric, though not necessarily positive definite. The optimization problem thereafter can be reformulated as:

$$\min_{\mathbf{H}} f(\mathbf{H}) = \|\mathbf{M} - \mathbf{H}\mathbf{H}^T\|_F^2. \quad (10)$$

The above problem, theoretically speaking, has closed solutions, though not unique (unless \mathbf{M} is negative semi-definite). One can see that if \mathbf{H} is an optimal solution, then $\mathbf{H}\mathbf{R}$ admits the same objective as long as $\mathbf{R}\mathbf{R}^T = \mathbf{I}$ ($\mathbf{R} \in \mathbb{R}^{k \times k}$). Apparently, if \mathbf{M} is positive definite, from Eckart-Young-Mirsky theorem, $\mathbf{H}(:, i) = \pm \sqrt{\sigma_i} \mathbf{v}_i$ will obtain the minimal where σ_i is the top i -th eigenvalue and \mathbf{v}_i is the corresponding eigenvector of \mathbf{M} . When \mathbf{M} is negative definite, then $\mathbf{H}^* = 0$ is unique solution. If \mathbf{M} has mixed signs in the eigenvalues, then $\mathbf{H}(:, i) = \pm \sqrt{\max\{\sigma_i, 0\}} \mathbf{v}_i$. However, one significant disadvantage of the above method is when the size of \mathbf{M} is very large, conducting eigenvalue decomposition is extremely computationally demanding. As a contribution, in our paper, we seek an alternative that is more applicable in practice. In light of the non-convexity of (10), where optimizing \mathbf{H} is challenging due to its high order, we turn to optimize:

$$\min_{\mathbf{H}, \mathbf{P}} \|\mathbf{M} - \mathbf{H}\mathbf{P}^T\|_F^2 + \lambda \|\mathbf{H} - \mathbf{P}\|_F^2, \quad (11)$$

where as long as λ is sufficiently large, hopefully, we have $\mathbf{H}^* = \mathbf{P}^*$, which is not difficult to obtain optimal solutions by utilizing any practical method such as alternating minimization. The following theorem will give a more specific bound:

Theorem 2. *Let $(\mathbf{H}^*, \mathbf{P}^*)$ be critical points of (11), where $\sigma_n(\cdot)$ denotes the n -th largest eigenvalue and $\lambda > \frac{\|\mathbf{H}^* \mathbf{P}^{*T}\|_F - \sigma_n(\mathbf{M})}{2}$, then $\mathbf{H}^* = \mathbf{P}^*$ and \mathbf{H}^* is a critical point of original problem (10).*

Proof of Theorem 2. We first introduce the following lemma which is very useful for later proof [17], [35].

Lemma 1. *For any symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ and positive semi-definite matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$, we have:*

$$\sigma_n(\mathbf{A}) \text{trace}(\mathbf{B}) \leq \text{trace}(\mathbf{A}\mathbf{B}) \leq \sigma_1(\mathbf{A}) \text{trace}(\mathbf{B}),$$

where $\sigma_i(\mathbf{A})$ is the i -th largest eigenvalue of \mathbf{A} .

Now we turn to check the sub-differential of f at critical point $(\mathbf{H}^*, \mathbf{P}^*)$ which is:

$$\begin{aligned} \partial_{\mathbf{H}^*} f(\mathbf{H}^*, \mathbf{P}^*) &= 2[(\mathbf{H}^* \mathbf{P}^{*T} - \mathbf{M})\mathbf{P}^* + \lambda(\mathbf{H}^* - \mathbf{P}^*)] = 0, \\ \partial_{\mathbf{P}^*} f(\mathbf{H}^*, \mathbf{P}^*) &= 2[(\mathbf{P}^* \mathbf{H}^{*T} - \mathbf{M})\mathbf{H}^* - \lambda(\mathbf{H}^* - \mathbf{P}^*)] = 0. \end{aligned} \quad (12)$$

By subtracting the second line from the first, we have:

$$(2\lambda \mathbf{I} + \mathbf{M})(\mathbf{H}^* - \mathbf{P}^*) = \mathbf{P}^* \mathbf{H}^{*T} \mathbf{H}^* - \mathbf{H}^* \mathbf{P}^{*T} \mathbf{P}^*. \quad (13)$$

By taking the inner product $\mathbf{H}^* - \mathbf{P}^*$ on both sides:

$$\begin{aligned} &\langle 2\lambda \mathbf{I} + \mathbf{M}, (\mathbf{H}^* - \mathbf{P}^*)(\mathbf{H}^* - \mathbf{P}^*)^T \rangle \\ &= \langle \mathbf{P}^* \mathbf{H}^{*T} \mathbf{H}^* - \mathbf{H}^* \mathbf{P}^{*T} \mathbf{P}^*, \mathbf{H}^* - \mathbf{P}^* \rangle. \end{aligned} \quad (14)$$

Applying Lemma 1 on the LHS, we have:

$$\langle 2\lambda \mathbf{I} + \mathbf{M}, (\mathbf{H}^* - \mathbf{P}^*)(\mathbf{H}^* - \mathbf{P}^*)^T \rangle \geq (2\lambda + \sigma_n(\mathbf{M})) \|\mathbf{H}^* - \mathbf{P}^*\|_F^2, \quad (15)$$

while applying Lemma 1 on the other side we have:

$$\begin{aligned} &\langle \mathbf{P}^* \mathbf{H}^{*T} \mathbf{H}^* - \mathbf{H}^* \mathbf{P}^{*T} \mathbf{P}^*, \mathbf{H}^* - \mathbf{P}^* \rangle \\ &= \left\langle \frac{\mathbf{P}^* \mathbf{H}^{*T} + \mathbf{H}^* \mathbf{P}^{*T}}{2}, (\mathbf{H}^* - \mathbf{P}^*)(\mathbf{H}^* - \mathbf{P}^*)^T \right\rangle \\ &\quad - \frac{\|\mathbf{H}^* \mathbf{P}^{*T} - \mathbf{P}^* \mathbf{H}^{*T}\|_F^2}{2} \\ &\leq \left\langle \frac{\mathbf{P}^* \mathbf{H}^{*T} + \mathbf{H}^* \mathbf{P}^{*T}}{2}, (\mathbf{H}^* - \mathbf{P}^*)(\mathbf{H}^* - \mathbf{P}^*)^T \right\rangle \quad (16) \\ &\leq \sigma_1 \left(\frac{\mathbf{P}^* \mathbf{H}^{*T} + \mathbf{H}^* \mathbf{P}^{*T}}{2} \right) \|\mathbf{H}^* - \mathbf{P}^*\|_F^2 \\ &\leq \left\| \frac{\mathbf{P}^* \mathbf{H}^{*T} + \mathbf{H}^* \mathbf{P}^{*T}}{2} \right\|_F \|\mathbf{H}^* - \mathbf{P}^*\|_F^2 \\ &\leq \|\mathbf{H}^* \mathbf{P}^{*T}\|_F \|\mathbf{H}^* - \mathbf{P}^*\|_F^2. \end{aligned}$$

Combining the above two equations we have:

$$(2\lambda + \sigma_n(\mathbf{M})) \|\mathbf{H}^* - \mathbf{P}^*\|_F^2 \leq \|\mathbf{H}^* \mathbf{P}^{*T}\|_F \|\mathbf{H}^* - \mathbf{P}^*\|_F^2. \quad (17)$$

Thus, if $\lambda > \frac{\|\mathbf{H}^* \mathbf{P}^{*T}\|_F - \sigma_n(\mathbf{M})}{2}$, then $\mathbf{H}^* = \mathbf{P}^*$ and any critical points satisfying (12) are also those for (10). \square

The following lemma gives a bound for $\|\mathbf{H}^* \mathbf{P}^{*T}\|_F$.

Lemma 2. *For (11), suppose the objective decreases with initialization $\mathbf{P}_0 = \mathbf{H}_0$, then for any $k \geq 0$, the iterate $(\mathbf{H}_k, \mathbf{P}_k)$ generated by any algorithm satisfies:*

$$\|\mathbf{H}_k \mathbf{P}_k^T\|_F \leq \|\mathbf{M} - \mathbf{H}_0 \mathbf{P}_0^T\|_F + \|\mathbf{M}\|_F. \quad (18)$$

Proof of Lemma 2. By the assumption that the algorithm decreases the objective function, we have:

$$\begin{aligned} \|\mathbf{M} - \mathbf{H}_k \mathbf{P}_k^T\|_F^2 + \lambda \|\mathbf{H}_k - \mathbf{P}_k\|_F^2 &\leq \|\mathbf{M} - \mathbf{H}_0 \mathbf{H}_0^T\|_F^2 \\ &\implies \|\mathbf{M} - \mathbf{H}_k \mathbf{P}_k^T\|_F^2 \leq \|\mathbf{M} - \mathbf{H}_0 \mathbf{H}_0^T\|_F^2 \\ &\implies \|\mathbf{H}_k \mathbf{P}_k^T\|_F \leq \|\mathbf{M}\|_F + \|\mathbf{M} - \mathbf{H}_k \mathbf{P}_k^T\|_F \\ &\leq \|\mathbf{M}\|_F + \|\mathbf{M} - \mathbf{H}_0 \mathbf{P}_0^T\|_F. \end{aligned} \quad \square$$

Corollary 1. *If $\lambda > \frac{1}{2}(\|\mathbf{M}\|_F + \|\mathbf{M} - \mathbf{H}_0 \mathbf{P}_0^T\|_F - \sigma_n(\mathbf{M}))$ and $\mathbf{P}_0 = \mathbf{H}_0$, then any algorithm decreases the objective in (11) will result in $\mathbf{P}^* = \mathbf{H}^*$.*

Proof. This is established by Lemma 2 and Theorem 2. \square

Discussion 1. *In our case, \mathbf{M} is symmetric, not necessarily positive semi-definite (PSD). To compute $\sigma_n(\mathbf{M})$, which can*

Algorithm 1 Efficient update to optimize (11)

Initialization: $\mathbf{H}_0 = \mathbf{P}_0 \in \mathbb{R}^{n \times k}$.

while not converge **do**

for $i = 1 : k$ **do**

$$\bar{\mathbf{M}} = \mathbf{M} - \sum_{j \neq i} \mathbf{h}_j \mathbf{p}_j^T.$$

$$\mathbf{h}_i^+ = \frac{(\bar{\mathbf{M}} + \lambda \mathbf{I}) \mathbf{p}_i}{\|\mathbf{p}_i\|^2 + \lambda}, \quad \mathbf{p}_i^+ = \frac{(\bar{\mathbf{M}} + \lambda \mathbf{I}) \mathbf{h}_i^+}{\|\mathbf{h}_i^+\|^2 + \lambda}.$$

end for

end while

Output: $\mathbf{H}^* = \mathbf{P}^*$.

be negative, it is very time-consuming if eigenvalue decomposition is utilized given its complexity level being $\mathcal{O}(n^3)$ when n is large. Therefore, we can divide it into the following cases:

- when \mathbf{M} is PSD, then $\lambda > \frac{1}{2}(\|\mathbf{M}\|_F + \|\mathbf{M} - \mathbf{H}_0 \mathbf{P}_0^T\|_F)$ will naturally satisfy the requirement, admitting desired optimal solutions.
- when \mathbf{M} is not PSD, we can first compute the leading eigenvalue (t) by power iteration or Lanczos method [13], which gives $\mathcal{O}(\log(n)/k)$ and $\mathcal{O}((\log(n)/k)^2)$ convergence rate, respectively [31]. By setting $\lambda > \frac{1}{2}(\|\mathbf{M}\|_F + \|\mathbf{M} - \mathbf{H}_0 \mathbf{P}_0^T\|_F + |t|)$, one can verify that it will obtain desired solution with significantly reduced complexity than eigenvalue decomposition.

To end this proof part, we propose an efficient algorithm that will **decrease the objective monotonically**.

By noticing $\mathbf{H} \mathbf{P}^T = \mathbf{h}_i \mathbf{p}_i^T + \sum_{j \neq i} \mathbf{h}_j \mathbf{p}_j^T$, we can optimize \mathbf{H} and \mathbf{P} column by column. Different from the column-wise update process proposed in other studies [27], [29], there is no assumption needed in our update algorithm, the update formula in our method is much more straightforward and a detailed proof of sufficient decrease in objective function with our method is provided.

Denote $\bar{\mathbf{M}} = \mathbf{M} - \sum_{j \neq i} \mathbf{h}_j \mathbf{p}_j^T$, we have:

$$\begin{aligned} \mathbf{h}_i^+ &= \arg \min_{\mathbf{h}_i} \|\bar{\mathbf{M}} - \mathbf{h}_i \mathbf{p}_i^T\|^2 + \lambda \|\mathbf{h}_i - \mathbf{p}_i\|^2 \\ &= \arg \min_{\mathbf{h}_i} (\|\mathbf{p}_i\|^2 + \lambda) \mathbf{h}_i^2 - 2 \langle \mathbf{h}_i, \lambda \mathbf{p}_i + \bar{\mathbf{M}} \mathbf{p}_i \rangle \\ &= \frac{(\bar{\mathbf{M}} + \lambda \mathbf{I}) \mathbf{p}_i}{\|\mathbf{p}_i\|^2 + \lambda}. \end{aligned} \quad (19)$$

Moreover, by noticing the strongly convexity ($\alpha = 2(\|\mathbf{p}_i\|^2 + \lambda)$) of objective w.r.t. \mathbf{h} , one have:

$$\begin{aligned} &f(\mathbf{h}_i, \mathbf{p}_i) - f(\mathbf{h}_i^+, \mathbf{p}_i) \\ &\geq \langle \nabla_{\mathbf{h}_i} f(\mathbf{h}_i^+, \mathbf{p}_i), \mathbf{h}_i - \mathbf{h}_i^+ \rangle + \frac{\alpha}{2} \|\mathbf{h}_i - \mathbf{h}_i^+\|^2 \\ &\geq \lambda \|\mathbf{h}_i - \mathbf{h}_i^+\|^2, \end{aligned} \quad (20)$$

which indicates a sufficient decrease by updating \mathbf{h}_i . Similarly, one can have the same conclusion while updating \mathbf{p}_i .

Fig. 3 shows the convergence curves of Algorithm1, SymANLS [35] and alternating direction method of multipliers (ADMM) [3] when solving symmetric matrix factorization problem. From the figure, we see Algorithm 1 converges very fast which indicates its superiority.

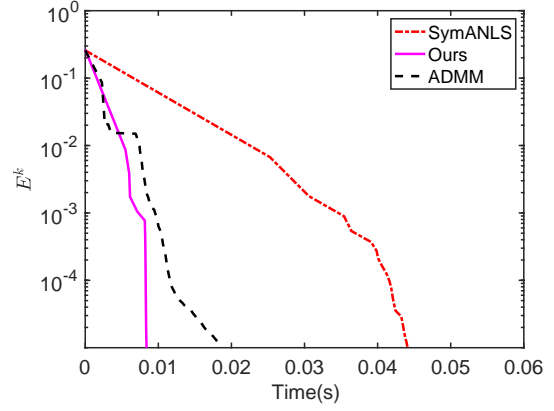


Fig. 3: Typical convergence curve which shows the superiority of our proposed method.

C. A More General Framework for Constrained Optimization

Though the above subsection describes a simple, yet very efficient algorithm to obtain an optimal solution, still it can't deal with most constraint problems such as $\|\mathbf{h}\|_2 = 1$, $\|\mathbf{h}\|_0 \leq s$, etc. [18], where different constraint indicate various meanings such as sparsity, probability distribution etc. In this subsection we propose a more general framework to solve (10).

Recall the relationship between second-order Taylor expansion and gradient descent:

$$f(\mathbf{y}) \approx f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{1}{2} \langle \mathbf{y} - \mathbf{x}, \nabla^2 f(\mathbf{x})(\mathbf{y} - \mathbf{x}) \rangle. \quad (21)$$

If we replace Hessian $\nabla^2 f(\mathbf{x})$ with $\frac{1}{t} \mathbf{I}$, then $\min f(\mathbf{y})$ is to minimize $\|\mathbf{y} - \mathbf{x} + t \nabla f(\mathbf{x})\|_2^2$, where t is the step size in gradient descent method. It can be verified that:

$$\nabla f(\mathbf{H}) = 4(\mathbf{H} \mathbf{H}^T \mathbf{H} - \mathbf{M} \mathbf{H}). \quad (22)$$

By invoking (21) to update \mathbf{H}^+ , we get:

$$\begin{aligned} \mathbf{H}^+ &= \arg \min_{\mathbf{H}' \in \mathbb{H}} f(\mathbf{H}') \\ &= \arg \min_{\mathbf{H}' \in \mathbb{H}} f(\mathbf{H}) + \langle \nabla f(\mathbf{H}), \mathbf{H}' - \mathbf{H} \rangle + \frac{1}{2t} \|\mathbf{H}' - \mathbf{H}\|_F^2 \\ &= \arg \min_{\mathbf{H}' \in \mathbb{H}} \frac{1}{2t} \|\mathbf{H}' - (\mathbf{H} - t \nabla f(\mathbf{H}))\|_F^2 \\ &= \mathcal{P}_{\mathbb{H}}(\mathbf{H} - t \nabla f(\mathbf{H})). \end{aligned} \quad (23)$$

\mathbb{H} denotes the feasible set satisfying the constraint. Stepsize t in the above update rule should be relatively small to avoid gradient explosion. However, if it is too small, the convergence becomes slow, which should be avoided as well. As a contribution, we propose a method with **adaptive stepsize** which will make the objective decrease monotonically.

Algorithm 2 provides a generalized framework to solve any symmetric matrix factorization with different constraints. Below we provide some concrete examples:

Algorithm 2 Optimize (10) where $\mathbf{H} \in \mathbb{H}$

Input: $\mathbf{M} = \mathbf{A} - \lambda \mathbf{L}$.

Initialization: $\mathbf{H}_0 \in \mathbb{H}$, $i = 0$.

while $i < K$ **do**

$$\nabla f(\mathbf{H}_i) = 4(\mathbf{H}_i \mathbf{H}_i^T \mathbf{H}_i - \mathbf{M} \mathbf{H}_i).$$

$$L_i = 4\sigma_{\max}(\mathbf{H}_i \mathbf{H}_i^T - \mathbf{M}) + 8\sigma_{\max}(\mathbf{H}_i^T \mathbf{H}_i).$$

$$\mathbf{H}_{i+1} = \mathcal{P}_{\mathbb{H}}(\mathbf{H}_i - t \nabla f(\mathbf{H}_i)), \text{ where } t = \frac{1}{2L_i}.$$

$$i = i + 1.$$

end while

Output: \mathbf{H}_K .

- *Example I:* Nonnegative constraint $\mathbf{H} \geq 0^2$:

$$\mathbf{H}^+ = \max\{\mathbf{H} - t \nabla f(\mathbf{H}), 0\}. \quad (24)$$

- *Example II:* Unit constraint $\|\mathbf{h}\|_2 = 1$:

$$\mathbf{h}^+ = \frac{\mathbf{h} - t \nabla f(\mathbf{h})}{\|\mathbf{h} - t \nabla f(\mathbf{h})\|_2}. \quad (25)$$

- *Example III:* Sparsity constraint $\|\mathbf{h}\|_0 \leq s$:
WLOG, assume the top s entry with maximum magnitude in \mathbf{h} is indexed as $[1, s]$, then [19]

$$\mathbf{h}^+ = \begin{cases} (\mathbf{h} - t \nabla f(\mathbf{h}))_j & \text{if } j \in [1, s], \\ 0 & \text{otherwise.} \end{cases} \quad (26)$$

- *Example IV:* Orthogonality Constraint $\mathbf{H}^T \mathbf{H} = \mathbf{I}$:

$$\mathbf{H}^+ = \mathbf{U}^T \mathbf{V}, \quad (27)$$

where $[\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}] = \text{svd}(\mathbf{H} - t \nabla f(\mathbf{H}))$.

- *Example V:* ℓ_1 -norm constraint on $\|\mathbf{h}\|_1 \leq \alpha$ [2]. By denoting $\mathcal{T}_{\lambda}(\mathbf{h}) = [\mathbf{h} - \lambda \mathbf{e}]_+ \odot \text{sgn}(\mathbf{h})$, we have:

$$\mathbf{h}^+ = \begin{cases} \mathbf{h} & \text{if } \|\mathbf{h}\|_1 \leq \alpha, \\ \mathcal{T}_{\lambda^*}(\mathbf{h}) & \text{otherwise,} \end{cases} \quad (28)$$

where λ^* is positive root for $\|\mathcal{T}_{\lambda}(\mathbf{h})\|_1 = \alpha$ which can be solved within $\mathcal{O}(n \log(n))$ [5].

IV. CONVERGENCE ANALYSIS

In the last section we mentioned the step size t should not be too large or small, and in this section, we will determine the best t in each update which guarantees the objective decreases monotonically by introducing the following lemma to begin:

Lemma 3. For a function f with a Lipschitz continuous gradient L , if $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2 \leq L\|\mathbf{x} - \mathbf{y}\|_2$ then $f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|_2^2$.

Proposition 1. For (10), in each i -th update, $L_i = 4\sigma_{\max}(\mathbf{H}_i \mathbf{H}_i^T - \mathbf{M}) + 8\sigma_{\max}(\mathbf{H}_i^T \mathbf{H}_i)$.

Proof of Proposition 1. For sake of simplicity, we denote \mathbf{H} as \mathbf{H}_i . With Lemma 3, it is equivalent to show $f(\mathbf{H}') -$

²One can verify that Algorithm 1 can also work if we set every column nonnegative, which can simply be obtained via $\mathbf{h}_i^+ = \max\{\frac{(\overline{\mathbf{M}} + \lambda \mathbf{I})\mathbf{p}_i}{\|\mathbf{p}_i\|^2 + \lambda}, 0\}$, $\mathbf{p}_i^+ = \max\{\frac{(\overline{\mathbf{M}} + \lambda \mathbf{I})\mathbf{h}_i^+}{\|\mathbf{h}_i^+\|^2 + \lambda}, 0\}$.

$f(\mathbf{H}') - \langle \nabla f(\mathbf{H}), \mathbf{H}' - \mathbf{H} \rangle \leq \frac{L}{2}\|\mathbf{H}' - \mathbf{H}\|_F^2$. By denoting \mathbf{H}' as $\mathbf{H} + \Delta \mathbf{H}$, we have:

$$\begin{aligned} & f(\mathbf{H} + \Delta \mathbf{H}) - f(\mathbf{H}) - \langle \Delta \mathbf{H}, \nabla f(\mathbf{H}) \rangle \\ &= 2\langle \mathbf{H} \mathbf{H}^T - \mathbf{M}, \Delta \mathbf{H} \Delta \mathbf{H}^T \rangle + \|\mathbf{H} \Delta \mathbf{H}^T + \Delta \mathbf{H} \mathbf{H}^T\|_F^2 \\ &\leq 2\langle \mathbf{H} \mathbf{H}^T - \mathbf{M}, \Delta \mathbf{H} \Delta \mathbf{H}^T \rangle + \|\mathbf{H} \Delta \mathbf{H}^T\|_F^2 + \|\Delta \mathbf{H} \mathbf{H}^T\|_F^2 \\ &= 2\text{tr}(\Delta \mathbf{H}^T (\mathbf{H} \mathbf{H}^T - \mathbf{M}) \Delta \mathbf{H}) + 4\text{tr}(\Delta \mathbf{H} \mathbf{H}^T \mathbf{H} \Delta \mathbf{H}^T) \\ &\leq [2\sigma_{\max}(\mathbf{H} \mathbf{H}^T - \mathbf{M}) + 4\sigma_{\max}(\mathbf{H}^T \mathbf{H})] \|\Delta \mathbf{H}\|_F^2 = \frac{L}{2} \|\Delta \mathbf{H}\|_F^2, \end{aligned} \quad (29)$$

where $\sigma_{\max}(\cdot)$ denotes the maximum singular value. \square

Below we show that the objective function $f(\mathbf{H})$ in Algorithm 2 has sufficient decrease in each update with step size $t = \frac{1}{2L}$ and the generated sequence is convergent.

Theorem 3. Let $g(\mathbf{H}_k) := f(\mathbf{H}_k) + \mathcal{C}(\mathbf{H}_k)$ be the objective function sequence generated by Algorithm 2 with constant step size $t_k = \frac{1}{2L_k}$. Then the sequence $g(\mathbf{H}_k)$ obeys sufficient decrease: $g(\mathbf{H}_{k-1}) - g(\mathbf{H}_k) \geq \frac{L_k}{2} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2$.

Proof of Theorem 3. From (23), noting that \mathbf{H}_k minimizes $J(\mathbf{H}) = \langle \nabla f(\mathbf{H}_{k-1}), \mathbf{H} - \mathbf{H}_{k-1} \rangle + \frac{1}{2t} \|\mathbf{H} - \mathbf{H}_{k-1}\|_F^2 + \mathcal{C}(\mathbf{H})$, where $\mathcal{C}(\mathbf{H})$ represents any constraint $\mathcal{C}(\mathbf{H}) = \begin{cases} 0, & \mathbf{H} \in \mathbb{H} \\ \infty, & \text{else} \end{cases}$, thus we naturally have $J(\mathbf{H}_k) \leq J(\mathbf{H}_{k-1})$, which implies:

$$\begin{aligned} & \langle \nabla f(\mathbf{H}_{k-1}), \mathbf{H}_k - \mathbf{H}_{k-1} \rangle + \frac{1}{2t} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2 \\ & \leq \mathcal{C}(\mathbf{H}_{k-1}) - \mathcal{C}(\mathbf{H}_k). \end{aligned} \quad (30)$$

According to the definition of Lipschitz continuous, Lemma 3, and when $t = \frac{1}{2L_k}$ in (30), we have:

$$\begin{aligned} & f(\mathbf{H}_k) + \mathcal{C}(\mathbf{H}_k) - f(\mathbf{H}_{k-1}) - \mathcal{C}(\mathbf{H}_{k-1}) \\ & \leq \langle \nabla f(\mathbf{H}_{k-1}), \mathbf{H}_k - \mathbf{H}_{k-1} \rangle + \frac{L_k}{2} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2 \\ & \quad + \mathcal{C}(\mathbf{H}_k) - \mathcal{C}(\mathbf{H}_{k-1}) \\ & \leq -L_k \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2 + \frac{L_k}{2} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2 \\ & = -\frac{L_k}{2} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2. \end{aligned} \quad (31)$$

\square

One can see that if $\mathbf{H}_k \in \mathbb{H}$, then (31) degenerates to $f(\mathbf{H}_{k-1}) - f(\mathbf{H}_k) \geq \frac{L_k}{2} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2$. Therefore, (10) decreases with update. Now repeating for all k will give $\frac{\bar{L}}{2} \sum_{k=1}^{\infty} \|\mathbf{H}_k - \mathbf{H}_{k-1}\|_F^2 \leq g(\mathbf{H}_0)$, where $\bar{L} := \min\{L_1, \dots, L_k\} > 0$ and establishes its convergence.

Theorem 4. In Algorithm 2, to ensure $\min \|\nabla f(\mathbf{H}) + \partial \mathcal{C}(\mathbf{H})\|_F^2 \leq \epsilon$, we need at most $T = \mathcal{O}(\frac{1}{\epsilon})$ iterations.

This indicates that the algorithm reaches a critical point at least a sub-linear convergence rate. Below is the proof:

Proof. First, by definition:

$$\mathbf{H}^+ = \arg \min_{\mathbf{H}'} \langle \nabla f(\mathbf{H}), \mathbf{H}' - \mathbf{H} \rangle + \frac{1}{2t} \|\mathbf{H}' - \mathbf{H}\|_F^2 + \mathcal{C}(\mathbf{H}'), \quad (32)$$




Fig. 4: Convergence comparison w.r.t. $E^k := \frac{\|M - H^k (H^k)^T\|_F^2}{\|M\|_F^2}$. **Left:** objective value versus iteration numbers. **Middle:** objective value versus time. **Right:** projected gradient descent with different step-size settings.




Fig. 5: Accuracy (Y-axis) comparison with λ changes (X-axis). **Left:** on a subset of 5 clusters from COIL-20. **Middle:** on a subset of 12 clusters from Reuters-21578. **Right:** on a subset of 10 clusters from TDT2.

which implies:

$$0 \in \nabla f(\mathbf{H}) + \frac{1}{t}(\mathbf{H}^+ - \mathbf{H}) + \partial\mathcal{C}(\mathbf{H}^+). \quad (33)$$

WLOG, now we define $\mathbf{A}_{\mathbf{H}^+} \in \nabla f(\mathbf{H}^+) + \partial\mathcal{C}(\mathbf{H}^+)$, apparently it indicates the gap to \mathbf{H}^* which satisfies $0 \in \nabla f(\mathbf{H}^*) + \partial\mathcal{C}(\mathbf{H}^*)$. Accordingly:

$$\mathbf{A}_{\mathbf{H}^+} \in \nabla f(\mathbf{H}^+) + \partial\mathcal{C}(\mathbf{H}^+) = \nabla f(\mathbf{H}^+) - \nabla f(\mathbf{H}) - \frac{1}{t}(\mathbf{H}^+ - \mathbf{H}). \quad (34)$$

As f is L -Lipschitz gradient continuous, then:

$$\begin{aligned} \|\mathbf{A}_{\mathbf{H}^+}\| &\leq \|\nabla f(\mathbf{H}^+) - \nabla f(\mathbf{H})\| + \frac{1}{t}\|\mathbf{H}^+ - \mathbf{H}\| \\ &\leq L\|\mathbf{H}^+ - \mathbf{H}\| + \frac{1}{t}\|\mathbf{H}^+ - \mathbf{H}\| \\ &= 3L\|\mathbf{H}^+ - \mathbf{H}\|, \end{aligned} \quad (35)$$

where the first line comes from subadditivity inequality, the second line is by definition of Lipschitz gradient continuity, while the last line is by definition of $t = \frac{1}{2L}$.

On the other hand, $f(\mathbf{H}) - f(\mathbf{H}^+) \geq \frac{L}{2}\|\mathbf{H}^+ - \mathbf{H}\|_F^2$, therefore:

$$\begin{aligned} \|\mathbf{A}_{\mathbf{H}^+}\|_F^2 &\leq 9L^2\|\mathbf{H}^+ - \mathbf{H}\|_F^2 \\ &\leq 9L^2 * \frac{2}{L}(f(\mathbf{H}) - f(\mathbf{H}^+)) \\ &= 18L(f(\mathbf{H}) - f(\mathbf{H}^+)). \end{aligned} \quad (36)$$

By repeating the above for all k :

$$\sum_i \|\mathbf{A}_i\|_F^2 \leq 18L(f(\mathbf{H}_0) - f(\mathbf{H}_k)) \leq 18Lf(\mathbf{H}_0). \quad (37)$$

Thus, $\min \|\mathbf{A}_i\|_F^2 \leq \frac{18Lf(\mathbf{H}_0)}{T}$, that is as long as $T = \frac{18Lf(\mathbf{H}_0)}{\epsilon} = \mathcal{O}(\frac{1}{\epsilon})$, then $\min \|\mathbf{A}_i\|_F = \epsilon$ which finishes the proof. Apparently, our algorithm has a at least sub-linear convergence rate. \square

We compare our proposed method with the nonnegative constraint on \mathbf{H} using projected gradient descent and the Newton-like method in SymNMF [12] in terms of convergence speed on synthetic data in Fig. 4. When run with the adaptive step-size setting, it converges way faster than others.

TABLE I: Datasets Information

Dataset	#Clusters	#Samples	Dimensionality
COIL-20	20	1440	1024
CIFAR-10	10	3000	1024
Reuters-21578	41	8213	18933
TDT2	30	9394	36771

TABLE II: Normalized mutual information (NMI) of different algorithms on four datasets with varying numbers of clusters

Method	COIL-20			CIFAR-10			Reuters-21578			TDT2		
	2	10	20	3	6	10	2	8	15	2	10	20
<i>K</i> -means	0.901	0.624	0.591	0.296	0.287	0.201	0.785	0.553	0.421	0.752	0.532	0.501
NMF	0.907	0.729	0.522	0.308	0.288	0.195	0.819	0.752	0.598	0.822	0.666	0.600
Spectral	0.877	0.701	0.677	0.309	0.298	0.201	0.828	0.611	0.499	0.829	0.607	0.592
SymHALS	0.911	0.688	0.652	0.308	0.255	0.198	0.855	0.631	0.552	0.822	0.611	0.588
SNMF	0.911	0.659	0.638	0.311	0.302	0.188	0.872	0.566	0.531	0.751	0.689	0.603
SymNMF	0.951	0.739	0.662	0.307	0.289	0.209	0.897	0.692	0.605	0.802	0.671	0.662
GNMF	0.951	0.701	0.652	0.319	0.301	0.217	0.852	0.772	0.618	0.798	0.699	0.652
DSC	0.949	0.752	0.701	0.318	0.302	0.297	0.876	0.779	0.682	0.851	0.691	0.664
Ours	0.958	0.797	0.725	0.341	0.319	0.302	0.901	0.798	0.755	0.861	0.729	0.705

TABLE III: Clustering accuracy (AC) of different algorithms on four datasets with varying numbers of clusters

Method	COIL-20			CIFAR-10			Reuters-21578			TDT2		
	2	10	20	3	6	10	2	8	15	2	10	20
<i>K</i> -means	0.921	0.674	0.631	0.316	0.297	0.221	0.815	0.563	0.503	0.800	0.581	0.533
NMF	0.923	0.765	0.586	0.330	0.306	0.198	0.900	0.777	0.616	0.839	0.693	0.629
Spectral	0.898	0.737	0.702	0.332	0.316	0.228	0.889	0.645	0.513	0.896	0.639	0.602
SymHALS	0.927	0.703	0.682	0.321	0.287	0.219	0.881	0.658	0.575	0.851	0.639	0.601
SNMF	0.932	0.686	0.653	0.328	0.312	0.202	0.890	0.586	0.552	0.781	0.703	0.630
SymNMF	0.972	0.772	0.695	0.321	0.305	0.231	0.911	0.721	0.627	0.823	0.703	0.689
GNMF	0.968	0.722	0.683	0.343	0.326	0.232	0.871	0.802	0.628	0.825	0.715	0.676
DSC	0.972	0.788	0.722	0.339	0.321	0.318	0.903	0.811	0.703	0.863	0.720	0.687
Ours	0.975	0.815	0.787	0.369	0.346	0.328	0.922	0.823	0.782	0.878	0.756	0.736

V. EXPERIMENTS

A. Datasets

Two image datasets and two text datasets are used in the experiment: COIL-20 [23], CIFAR-10 [11], Reuters-21578 [25], and TDT2 [28]. Detailed descriptions of the number of clusters, number of samples, and dimensionality of these datasets can be found in Table I.

B. Experimental Settings

Clustering performances of the following 9 algorithms are compared:

- 1) Standard *K*-means;
- 2) NMF using alternating nonnegative least squares algorithm [10]; The data matrix \mathbf{X} is transformed into its normalized-cut weighted version;
- 3) Spectral clustering (Spectral) [24], [30];
- 4) Hierarchical Alternating Least Squares (HALS) for symmetric NMF (SymHALS) [35];
- 5) Symmetric NMF using Procrustes rotations (SNMF) [8];
- 6) Symmetric NMF (SymNMF) [12];
- 7) Graph regularized nonnegative matrix factorization (GNMF) [4];
- 8) Deep subspace clustering (DSC) [9];
- 9) Our method. Algorithm 2 is used to solve the objective function (10).

In order to randomize the experiments, we conduct the evaluation using subsets of the whole datasets with different

cluster numbers. For each selected number of clusters K , 10 test runs are conducted on a randomly chosen subset with K clusters. When K is the total number of clusters in the complete data set, the test runs are repeated on the entire data set. The symmetric matrix \mathbf{A} can be obtained by utilizing any similarity measures, for simplicity we use the inner product similarity. Throughout the experiments, we use Matlab R2019a on a laptop with a 1.4 GHz QuadCore Intel Core i5 processor.

C. Results and Analysis

The clustering quality is measured by normalized mutual information (NMI) [14], a measurement of similarity from information theory, and clustering accuracy (AC), the percentage of items correctly clustered with the maximum bipartite matching [33]. AC is defined as follows:

$$AC = \frac{\sum_{i=1}^n \delta(r_i, \text{map}(l_i))}{n}, \quad (38)$$

where l_i is the obtained cluster label, r_i is the original provided label, n is the number of total samples, $\delta(x, y)$ equals 1 when $x = y$ and equals 0 otherwise, and $\text{map}(l_i)$ is the permutation function that maps each l_i to the equivalent cluster label provided via Hungarian algorithm.

Fig. 5 shows how the clustering accuracy of our method varies with different values of λ . The performance is not changing dramatically with respect to the parameter λ , and our method has consistently ideal performance if λ is within a reasonable range. It's reasonable to observe that the optimal

value of λ is slightly dependent on the data since the dimensionality and magnitude of the data can all have some effect on it. Experiment results of normalized mutual information and clustering accuracy on the four datasets are shown in Table II and Table III. We report the mean of NMI and AC for each given cluster number K over 10 test runs, the highest accuracy for each K is highlighted. We can see that for both image data and text data, our proposed method can always achieve the best clustering performance among all the methods, and the improvement is significant, both in NMI and AC. One potential reason that the performance on CIFAR-10 is not that good as on COIL-20 may be the images from CIFAR-10 have more complex and varying backgrounds than the images from COIL-20. Although as the number of clusters increases, all methods' clustering performance is getting worse, our method is relatively stable with an increasing number of clusters compared to other methods.

VI. CONCLUSION

In this paper, we study the symmetric matrix factorization problem with a regularization term. We propose an efficient column-wise update rule and provide a general framework that can be extended to solve symmetric matrix factorization problems with various constraints. We prove the convergence rate with theoretical analysis. The results of extensive experiments on real-world data sets validate the effectiveness of our algorithm and its superiority in data clustering tasks.

REFERENCES

- [1] Francis Bach and Michael Jordan. Learning spectral clustering. *Advances in neural information processing systems*, 16(2):305–312, 2004.
- [2] Amir Beck. *First-order methods in optimization*. SIAM, 2017.
- [3] Stephen Boyd, Neal Parikh, and Eric Chu. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- [4] Deng Cai, Xiaofei He, Jiawei Han, and Thomas S Huang. Graph regularized nonnegative matrix factorization for data representation. *IEEE transactions on pattern analysis and machine intelligence*, 33(8):1548–1560, 2010.
- [5] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l_1 -ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279, 2008.
- [6] Quanquan Gu and Jie Zhou. Co-clustering on manifolds. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 359–368, 2009.
- [7] Zhaoshui He, Shengli Xie, Rafal Zdunek, Guoxu Zhou, and Andrzej Cichocki. Symmetric nonnegative matrix factorization: Algorithms and applications to probabilistic clustering. *IEEE Transactions on Neural Networks*, 22(12):2117–2131, 2011.
- [8] Kejun Huang, Nicholas D Sidiropoulos, and Ananthram Swami. Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition. *IEEE Transactions on Signal Processing*, 62(1):211–224, 2013.
- [9] Pan Ji, Tong Zhang, Hongdong Li, Mathieu Salzmann, and Ian Reid. Deep subspace clustering networks. *arXiv preprint arXiv:1709.02508*, 2017.
- [10] Hyunsoo Kim and Haesun Park. Nonnegative matrix factorization based on alternating nonnegativity constrained least squares and active set method. *SIAM journal on matrix analysis and applications*, 30(2):713–730, 2008.
- [11] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [12] Da Kuang, Chris Ding, and Haesun Park. Symmetric nonnegative matrix factorization for graph clustering. In *Proceedings of the 2012 SIAM international conference on data mining*, pages 106–117. SIAM, 2012.
- [13] Jacek Kuczynski and Henryk Woźniakowski. Estimating the largest eigenvalue by the power and lanczos algorithms with a random start. *SIAM journal on matrix analysis and applications*, 13(4):1094–1122, 1992.
- [14] Andrea Lancichinetti, Santo Fortunato, and János Kertész. Detecting the overlapping and hierarchical community structure in complex networks. *New journal of physics*, 11(3):033015, 2009.
- [15] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- [16] Tao Li and Chris Ding. The relationships among various nonnegative matrix factorization methods for clustering. In *Sixth International Conference on Data Mining (ICDM'06)*, pages 362–371. IEEE, 2006.
- [17] Xiao Li, Zhihui Zhu, Qiuwei Li, and Kai Liu. A provable splitting approach for symmetric nonnegative matrix factorization. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [18] Kai Liu. Spherical matrix factorization. *arXiv preprint arXiv:2111.14536*, 2021.
- [19] Kai Liu, Qiuwei Li, Hua Wang, and Gongguo Tang. Spherical principal component analysis. *arXiv preprint arXiv:1903.06877*, 2019.
- [20] Kai Liu, Xiangyu Li, Zhihui Zhu, Lodewijk Brand, and Hua Wang. Factor-bounded nonnegative matrix factorization. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(6):1–18, 2021.
- [21] Kai Liu and Hua Wang. High-order co-clustering via strictly orthogonal and symmetric l_1 -norm nonnegative matrix tri-factorization. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 2018.
- [22] Kai Liu, Hua Wang, Shannon Risacher, Andrew Saykin, and Li Shen. Multiple incomplete views clustering via non-negative matrix factorization with its application in alzheimer's disease analysis. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 1402–1405. IEEE, 2018.
- [23] S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (coil-20). *Technical Report CUCS-005-96*, 1996.
- [24] Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*, pages 849–856, 2002.
- [25] Reuters-21578. <http://www.research.att.com/lewis>. 1987.
- [26] Mikkel N Schmidt, Ole Winther, and Lars Kai Hansen. Bayesian non-negative matrix factorization. In *International Conference on Independent Component Analysis and Signal Separation*, pages 540–547. Springer, 2009.
- [27] Qingjiang Shi, Haoran Sun, Songtao Lu, Mingyi Hong, and Meisam Razaviyayn. Inexact block coordinate descent methods for symmetric nonnegative matrix factorization. *IEEE Transactions on Signal Processing*, 65(22):5995–6008, 2017.
- [28] TDT2. <http://projects.ldc.upenn.edu/tdt2/>. 1998.
- [29] Arnaud Vandaele, Nicolas Gillis, Qi Lei, Kai Zhong, and Inderjit Dhillon. Efficient and non-convex coordinate descent for symmetric nonnegative matrix factorization. *IEEE Transactions on Signal Processing*, 64(21):5571–5584, 2016.
- [30] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [31] John Wright and Yi Ma. *High-dimensional data analysis with low-dimensional models: Principles, computation, and applications*. Cambridge University Press, 2022.
- [32] Wenhui Wu, Yuheng Jia, Sam Kwong, and Junhui Hou. Pairwise constraint propagation-induced symmetric nonnegative matrix factorization. *IEEE transactions on neural networks and learning systems*, 29(12):6348–6361, 2018.
- [33] Wei Xu, Xin Liu, and Yihong Gong. Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '03*, page 267–273, New York, NY, USA, 2003. Association for Computing Machinery.
- [34] Yan Zhao, Xing Chen, Jun Yin, and Jia Qu. Snmfsmma: using symmetric nonnegative matrix factorization and kronecker regularized least squares to predict potential small molecule-microrna association. *RNA biology*, 17(2):281–291, 2020.
- [35] Zhihui Zhu, Xiao Li, Kai Liu, and Qiuwei Li. Dropping symmetry for fast symmetric nonnegative matrix factorization. *arXiv preprint arXiv:1811.05642*, 2018.