



Binary Orthogonal Non-negative Matrix Factorization

Sajad Fathi Hafshejani^(✉), Daya Gaur, Shahadat Hossain,
and Robert Benkoczi

Department of Math and Computer Science, University of Lethbridge,
Lethbridge, AB, Canada

{sajad.fathihafshejan, daya.gaur, shahadat.hossain,
robert.benkoczi}@uleth.ca

Abstract. We propose a method for computing binary orthogonal non-negative matrix factorization (BONMF) for clustering and classification. The method is tested on several representative real-world data sets. The numerical results confirm that the method has improved accuracy compared to the related techniques. The proposed method is fast for training and classification and space efficient.

Keywords: Binary orthogonal non-negative matrix factorization · Non-convex optimization problem · Classification

1 Introduction

For a data matrix \mathbf{X} of size $m \times n$, $\mathbf{WH} \approx \mathbf{X}$ (where \mathbf{W} is of size $m \times k$, \mathbf{H} is of size $k \times n$) is considered as a low rank approximation ($k \ll n$) of the data matrix \mathbf{X} . Low rank approximations are essential in machine learning applications and especially in natural language processing and topic modelling where the data matrix is constructed over a collection of words from a vocabulary and a usually large collection of documents [10, 21, 24, 34].

Singular value decomposition (SVD) [28] is an early approach for computing such a low rank approximation of data. SVD minimizes the Frobenius norm and the spectral norm simultaneously; not only that, the columns of \mathbf{W} are orthogonal, and the rows of \mathbf{H} are also orthogonal. However, the entries in \mathbf{W} , \mathbf{H} may be negative, which reduces the utility of SVD for data matrix \mathbf{X} in which the entries are positive as the factors in \mathbf{W} do not have an intuitive explanation. Non-negative matrix factorization (NMF), $\mathbf{WH} \approx \mathbf{X}$ and $\mathbf{X}, \mathbf{W}, \mathbf{H} \geq \mathbf{0}$, was introduced by Paatero and Tapper [20] to overcome this difficulty of interpretation of the factors. NMF was shown to be NP-complete by Vavasis [27]. NMF does not require the columns of \mathbf{W} to be orthogonal, and this is considered a severe drawback in some applications as the columns (factors) of \mathbf{W} are not separable by a large angle. Keeping this limitation in mind Ding et al. [34] introduced orthogonality constraints in NMF, $\mathbf{X} \approx \mathbf{WH}$ and $\mathbf{X}, \mathbf{W}, \mathbf{H} \geq \mathbf{0}$, the rows of \mathbf{H} are orthogonal and demonstrated that is an effective approach for clustering of documents.

We consider the following problem: given a $m \times n$ data matrix \mathbf{X} , we wish to represent \mathbf{X} as a product of two matrices \mathbf{W}, \mathbf{H} with dimensions $m \times k$ and $k \times n$ respectively with the following restrictions: entries in \mathbf{W} are positive, the entries in \mathbf{H} are either 0 or 1, $\mathbf{H}\mathbf{H}^T = I$ and the norm $\|\mathbf{X} - \mathbf{W}\mathbf{H}\|^2$ is minimized. Additionally, we want k to be small compared to n, m . The columns of the data matrix \mathbf{X} can be thought of as the n samples. Low rank \mathbf{W} represents the latent features. We call this problem the binary orthogonal nonnegative matrix factorization problem (BONMF).

1.1 Contributions

This paper gives a new method (Algorithm 1) for computing a binary orthogonal NMF using the two-phase iterative approach. In the first phase, we use a known update rule [16] to compute the factor \mathbf{W} . In the second phase, we use the observation that the binary constraints on \mathbf{H} have a geometric interpretation. This gives an efficient rule to update \mathbf{H} in each iteration (Eq. (12)). The entries in \mathbf{H} are binary, and they are computed column-wise. If all the entries in \mathbf{H} are non-zero, then $O(nk)$ space is needed. However, \mathbf{H} is binary, and the rows of \mathbf{H} are orthogonal. Therefore, only $O(n)$ space is needed. If we compute the entries of \mathbf{H} columns-wise, intermediate states also need $O(n)$ space. The computation for each column of \mathbf{H} takes $O(n^2k)$ steps. Therefore, the method is space efficient.

We evaluate the method’s performance (in Sect. 4) for training and testing on reference data sets from the ML repository. The experiments demonstrate that the training and the classification phase are efficient (Table 2). The method is accurate and outperforms the state of art methods (Table 2). This method uses k dot products of m element vectors to update each column of the coefficient matrix \mathbf{H} where k is the number of classes in the data set. This is a significant reduction in the computation needed compared to the algorithms of [15, 32, 34] in the classification phase. The method is also space efficient as \mathbf{H} is sparse.

2 Related Work

We begin with NMF and the related background needed to describe our algorithm. Given a non-negative matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, a non-negative matrix factorization of \mathbf{X} finds two non-negative matrices $\mathbf{W} \in \mathbb{R}^{m \times k}$ and $\mathbf{H} \in \mathbb{R}^{k \times n}$ with $k \ll \min(m, n)$ such that:

$$\mathbf{X} \approx \mathbf{W}\mathbf{H},$$

and the entries in \mathbf{W}, \mathbf{H} are positive. The factorization has a natural interpretation [15] and can be computed using various unsupervised machine learning methods. Due to its intuitive interpretation, NMF has found numerous applications such as data consolidation [8], image clustering [6], topic modelling [2], community detection [29], recommender systems [22], and gene expression profiling [33].

BONMF is different from NMF. The entries in \mathbf{H} are restricted to binary. If the columns of \mathbf{H} are orthogonal, then the columns can be used to cluster the data. Therefore, BONMF factorization has several exciting applications [31].

Orthogonal NMF (ONMF) in which $\mathbf{X} \approx \mathbf{W}\mathbf{H}$ and $\mathbf{W}, \mathbf{H} \geq \mathbf{0}$ and $\mathbf{H}\mathbf{H}^T = \mathbf{I}$ was defined by Ding et al. [34] who gave an algorithm based on solving the Lagrangian relaxation. The entries in \mathbf{H} in ONMF are not required to be binary. ONMF use for data clustering was popularized by Seung and Lee [23]. One of the first notable applications of orthogonal NMF to document clustering is in [30] who gave improved algorithms and showed that ONMF performed better at document clustering than NMF. K-means [19] is one of the most widely used algorithms for unsupervised learning. Bauckhage [3] showed that the objective function of K-means can be rewritten as ONMF if the entries in \mathbf{H} are binary, and the following condition holds:

$$\sum_i \mathbf{H}_{ij} = 1 \quad \forall j \quad (1)$$

Therefore, BONMF is equivalent to K-means clustering. BONMF was also studied by Zdunek [31] and differs from the well-studied non-negative matrix factorization (NMF). Lee et al. [16] studied BONMF without the condition (1) on \mathbf{H} and gave an algorithm for determining such a factorization. However, applications to classification are not many. In this paper, we study BONMF for its use in prediction and classifying data, including clustering.

2.1 NMF

NMF can be formulated as the following optimization problem that minimizes the square of the Frobenius norm:¹

$$F(\mathbf{W}, \mathbf{H}) = \min_{\mathbf{W}, \mathbf{H} \geq \mathbf{0}} \frac{1}{2} \|\mathbf{X} - \mathbf{W}\mathbf{H}\|_F^2. \quad (2)$$

Most of the methods for computing NMF are based on iterative update rules. A popular set of update rules given below is due to Lee and Seung [15], the iteration number is in superscript.

$$\mathbf{W}_{ia}^{t+1} = \mathbf{W}_{ia}^t \frac{(\mathbf{X}\mathbf{H}^{tT})_{ia}}{(\mathbf{W}^t\mathbf{H}^t\mathbf{H}^{tT})_{ia}}, \quad \forall i, a; \quad (3)$$

$$\mathbf{H}_{bj}^{t+1} = \mathbf{H}_{bj}^t \frac{(\mathbf{W}^{t+1T}\mathbf{X})_{bj}}{(\mathbf{W}^{t+1T}\mathbf{W}^{t+1}\mathbf{H}^t)_{bj}}, \quad \forall b, j. \quad (4)$$

For many more variations on such update rules, see [11]. Optimization approaches such as block-coordinate descent, projected gradient descent, and alternating non-negative least squares (ANLS) [18] have also been used for NMF. ANLS transforms the problem in (2) into two convex optimization problems:

¹ $\|\mathbf{A}\|_F = \sqrt{\text{tr}(\mathbf{A}^T \times \mathbf{A})} = \sum_{i,j} |a_{ij}|^2$.

$$\mathbf{W}_{t+1} = \min_{\mathbf{W} \geq 0} f(\mathbf{W}, \mathbf{H}_t) = \min_{\mathbf{W} \geq 0} \frac{1}{2} \|\mathbf{X} - \mathbf{W}\mathbf{H}_t\|_F^2, \quad (5)$$

$$\mathbf{H}_{t+1} = \min_{\mathbf{H} \geq 0} f(\mathbf{W}_{t+1}, \mathbf{H}) = \min_{\mathbf{H} \geq 0} \frac{1}{2} \|\mathbf{X} - \mathbf{W}_{t+1}\mathbf{H}\|_F^2. \quad (6)$$

We can solve the optimization problems given by (5) and (6) in a few ways. [12, 13] gave the Rank-one Residue Iteration (RRI) algorithm for computing NMF. This algorithm was also independently proposed by Cichocki et al. [5], which is called the Hierarchical Alternating Least Squares (HALS) algorithm. The solution to (5) and (6) in HALS/RRI is given by explicit formulas, which make for easy implementation. Kim et al. [14] used Newton and quasi-Newton methods to solve (5), (6) and showed that their method has faster convergence. However, these methods require determining a suitable active set of the constraints in each iteration [4]. Two efficient algorithms for approximately orthogonal NMF were given by Li et al. [17]. Asymmetric NMF with Beta-divergences approach was studied by Lee et al. [16].

NMF is a quadratic boolean optimization problem, so it can also be solved using the Quantum Simulated Annealing (QSA) approach of Farhi et al. [7]. Recently, Golden and O'Malley [9] used a combination of forward and reverse annealing in the quantum annealing to obtain improved performance of QSA for NMF.

2.2 Binary Orthogonal NMF

Given a non-negative matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, a BONMF of \mathbf{X} finds the non-negative matrix $\mathbf{W} \in \mathbb{R}^{m \times k}$ and a binary $\mathbf{H} \in \{0, 1\}^{k \times n}$ with $k \ll \min(m, n)$. The BONMF can be written as the following optimization problem:

$$F(\mathbf{W}, \mathbf{H}) = \min_{\mathbf{W} \in \mathbb{R}^{m \times k}, \mathbf{H} \in \{0, 1\}^{k \times n}} \frac{1}{2} \|\mathbf{X} - \mathbf{W}\mathbf{H}\|_F^2. \quad (7)$$

Using the ANLS approach [18] we can transform (7) into the following sub-problems:

$$\mathbf{W}_{t+1} = \min_{\mathbf{W} \geq 0} \frac{1}{2} \|\mathbf{X} - \mathbf{W}\mathbf{H}_t\|_F^2, \quad (8)$$

$$\mathbf{H}_{t+1} = \min_{\mathbf{H} \in \{0, 1\}} \frac{1}{2} \|\mathbf{X} - \mathbf{W}_{t+1}\mathbf{H}\|_F^2. \quad (9)$$

The problem (8) can be solved using the update rule (3) of [15]. Sub-problem (9) is solved in two different ways in the following papers. Zhang et al. [32] update each row of the matrix \mathbf{H} using the following strategy:

$$\mathbf{h} = \text{sgn} \left(\mathbf{X}^T \mathbf{z} - \frac{1}{2} \mathbf{I} \mathbf{z}^T \mathbf{z} - \mathbf{H}^T \mathbf{W}^T \mathbf{z} \right), \quad (10)$$

where

$$\text{sgn}(x) = \begin{cases} 1, & \text{if } x > 0 \\ 0, & \text{otherwise,} \end{cases}$$

and \mathbf{z} is the k -th column of \mathbf{W} , and \mathbf{W}' is the matrix of \mathbf{W} excluding \mathbf{z} ; \mathbf{h}^T is the k -th row of \mathbf{H} and \mathbf{H}' is the matrix of \mathbf{H} excluding \mathbf{h}^T . In addition, $\mathbf{I} \in \mathbb{R}^n$ is a vector whose entries are all one. Zdunek [31] presented another method for updating \mathbf{H} under the assumption that \mathbf{H} is orthogonal, which uses simulated annealing. Since they use a different approach, we don't describe it in detail.

3 The Algorithm

This section describes our approach. The method solves the optimization problems given by (8) and (9). To solve (8), we use the update rule given by equation (3) [15] where \mathbf{W} is computed using

$$\mathbf{W}_{ia}^{t+1} \leftarrow \mathbf{W}_{ia}^t \frac{(\mathbf{X}\mathbf{H}^{tT})_{ia}}{(\mathbf{W}^t\mathbf{H}^t\mathbf{H}^{tT})_{ia}}, \quad \forall i, a.$$

Given \mathbf{X}, \mathbf{H} , to solve (9) we write the problem as:

$$F(\mathbf{H}) = \min_{\mathbf{H} \in \{0,1\}^{k \times n}} \|\mathbf{X} - \mathbf{W}\mathbf{H}\|_F^2. \quad (11)$$

Each column of the matrix \mathbf{H} is computed in two steps as follows:

- In the first step, we calculate the angular distance between column i of \mathbf{X} and column j of matrix \mathbf{W} to obtain $\mathbf{H}_{j,i}$.

$$\mathbf{H}_{j,i} = \frac{\langle \mathbf{X}_{:,i}, \mathbf{W}_{:,j} \rangle}{\|\mathbf{X}_{:,i}\| \|\mathbf{W}_{:,j}\|}, \quad (12)$$

where $\mathbf{X}_{:,i}$ denotes the i -th column of matrix \mathbf{X} and $\langle \cdot, \cdot \rangle$ is the inner product.

- In the second step, the maximum value (any) in each matrix column \mathbf{H} is changed to 1, and other values are changed to 0. The process can be summarized as follows:

$$\mathbf{H}_{j,i} = \begin{cases} 1, & \text{if } \mathbf{H}_{j,i} = \max \mathbf{H}_{:,i} \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

The pseudo-code for the method is in Algorithm 1. These steps are executed column by column for \mathbf{H} .

Algorithm 1. BONMF

-
- 1: INPUT: Matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, and T
 - 2: OUTPUT: Matrices $\mathbf{W} \in \mathbb{R}^{m \times k}$ and $\mathbf{H} \in \{0, 1\}^{k \times n}$
 - 3: Initialize matrices \mathbf{W} and \mathbf{H}
 - 4: WHILE *iterations* < *max* & \neg *convergence*
 - 5: $\mathbf{W}_{ia}^{t+1} \leftarrow \mathbf{W}_{ia}^t \frac{(\mathbf{X}\mathbf{H}^{tT})_{ia}}{(\mathbf{W}^t\mathbf{H}^t\mathbf{H}^{tT})_{ia}}, \quad \forall i, a;$ ▷ Update \mathbf{W} using (3).
 - 6: $\mathbf{H}_{j,i}^{t+1} \leftarrow \frac{\langle \mathbf{x}_{:,i}, \mathbf{w}_{:,j}^{t+1} \rangle}{\|\mathbf{x}_{:,i}\| \|\mathbf{w}_{:,j}^{t+1}\|} \quad \forall j, i$ ▷ Update \mathbf{H} using (12)
 - 7: $\mathbf{H}_{j,i}^{t+1} \leftarrow \begin{cases} 1, & \text{if } \mathbf{H}_{j,i}^{t+1} = \max \mathbf{H}_{:,i}^{t+1} \\ 0, & \text{otherwise} \end{cases} \quad \forall j, i$ ▷ Update \mathbf{H} using (13)
 - 8: END
 - 9: return \mathbf{W} and \mathbf{H}
-

4 Empirical Evaluation

This section examines three characteristics of Algorithm 1. We study the time needed for classification, the accuracy, and the time required for computing the factorization (training time) for the data sets shown in Table 1. The data sets are representative of the varying complexity of machine learning; some are easy (digits), some are hard (diabetes), and some have a significant number of features (ORL). These are popular datasets from the OpenML repository [26]. These data sets have multiple single label classes and serve as a nice testbed for evaluating unsupervised learning algorithms, even in deep learning.

Table 1. Data Sets

Name	# samples	# features	# classes
ORL	400	4,096	40
Optdigits	5,619	65	10
Phishing	11,055	68	2
Monkey	471	6	2
Pendigits	10,992	17	2
Diabetes	7,67	8	2
W8a	49,748	300	2
Banking	8237	13	3
Svmguide	3,087	5	2

We compare the performance of Algorithm 1 with the algorithms for orthogonal matrix factorization [34], non-negative matrix factorization [15], and semi-

binary non negative matrix factorization [32]. We examine the relative performance of these algorithms for accuracy and classification time. We use two methods for ONMF to classify a new data point j (column vector of \mathbf{X}). Typically, ONMF uses the index i of the maximum entry in column $\mathbf{H}_{:,j}$ for classification, which gives the cluster to which data j belongs. The data points in cluster i may have different labels, and the label of j is the label of the point in cluster i that is closest (distance-wise). We refer to this default scheme for determining the label as ONMF in Table 2. The second scheme we use to determine the label uses the label on i' , which is the point in cluster i that forms the smallest angle data point (vector for j); then, the label of i' is used to classify point j . The cluster to which data point j belongs is again computed based on the angles to the columns of \mathbf{W} , the closest column of \mathbf{W} determines the cluster, and the closest point in the cluster (angle-wise) determines the label. The second scheme is ONMF-cos in Table 2. The other two algorithms that we used to compare are i) the popular and the foundational algorithm of [15] for NMF, labelled ‘‘Lee and Seung’’ and ii) the algorithm of [32] for NMF with the constraint that the entries in \mathbf{H} are binary (labelled ‘‘Zhang et al.’’). We use only the matrix-based method factorization algorithms closest to the K-means for evaluation. As part of a future study, it would be interesting to see how these algorithms perform against a highly optimized implementation of K-means.

We report on experiments that were run on a laptop (i5-7200U, 12GB of RAM). The algorithms [15,32] were coded in Python 3.1. We used the number of classes as the rank in factorization. Eighty percent of the data was used for training, and the remaining was used for testing the accuracy. We use the python library (ionmf.factorization.onmf) for ONMF [25]. Initialization of \mathbf{W} , \mathbf{H} is done using the following scheme: we sort the columns of the matrix \mathbf{X} based on its norm. To determine the i^{th} column of \mathbf{W} , use the average of ten randomly chosen columns from the first thirty columns of \mathbf{X} as in [1]. Initial matrix \mathbf{H}^0 is computed using $\mathbf{H}^0 = (\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T \mathbf{X}$. Since the initial values of \mathbf{W} are random, we run the algorithm thirty times and report the averages in Table 2. The first thing to note is that in Algorithm 1 extra computation is needed to convert \mathbf{H} to binary in each iteration. This computation increases the time needed for factorization relative to ONMF and NMF and is linear in the size of \mathbf{H} . However, given the factorization, the classification phase is more efficient, and \mathbf{H} is sparse.

4.1 Classification

In the basic NMF approach given by update rules (3) and (4) (as in [15]), the number of steps needed for the classification of new data (column vector of \mathbf{X}) is proportional to the number of columns in the factorization \mathbf{W} , \mathbf{H} . We need to calculate the angle between the coefficient vector for the new data and all the columns of \mathbf{H} (as many as the columns in \mathbf{X}) to determine the label for the data. Algorithm 1 does not share this disadvantage. We can compute the angle of the sample to every column of \mathbf{W} (a low-rank matrix) and use the closest column to determine the label. This observation is reflected in data in the row labelled ‘‘TT (s)’’ in Table 2.

4.2 Accuracy

The accuracy of the five methods is presented in Table 2. The entries in bold font indicate that a particular method had the most accuracy). Six of eight data sets (except pendigits and banking) have improved accuracy for classification when cos angles are used to measure similarity, and the utility of using the angle (12) is evident. The method presented here has the best accuracy on six of the eight data sets (except optdigits, phishing). Regarding classification time, it performs best on seven of the eight datasets. Note that the other method ONMF+cos which is as competitive as Algorithm 1 uses $O(nk)$ space to store \mathbf{H} whereas Algorithm 1 uses $O(n)$ space even in the intermediate stages of the calculations.

Table 2. Numerical Results

Name		ONMF Strazar et al.	Lee and Seung	Algorithm 1	ONMF+ cos Strazar et al.	Zhang et al.
ORL	TT (s)	1.62	1.63	2.28	1.63	9.48
	CT (s)	0.24	0.178	0.09	0.09	0.20
	AC (%)	85.00	85.00	89.99	89.99	85.00
Optdigits	TT (s)	0.38	0.62	2.60	0.36	3.72
	CT (s)	45.24	38.97	13.03	13.23	37.91
	AC (%)	53.45	62.27	80.78	88.96	48.12
Phishing	TT (s)	0.87	0.80	3.42	1.05	1.96
	CT (s)	156.24	142.92	52.22	53.14	146.40
	AC (%)	54.76	54.76	91.85	92.14	54.76
Monkey	TT (s)	0.01	0.01	0.10	0.01	0.06
	CT (s)	0.45	0.34	0.14	0.12	0.31
	AC (%)	48.80	48.80	80.95	80.95	53.12
Diabetes	TT (s)	0.02	0.01	1.016	0.01	0.59
	CT (s)	0.80	0.85	0.24	0.25	0.67
	AC (%)	51.72	51.72	68.96	68.96	68.95
Banking	TT (s)	0.05	0.06	2.01	0.07	1.56
	CT (s)	77.07	70.87	25.25	27.42	73.74
	AC (%)	87.05	76.09	87.05	87.05	87.05
Svmguide	TT (s)	0.01	0.01	0.82	0.02	0.40
	CT (s)	13.90	12.05	4.93	45.27	13.51
	AC (%)	73.70	65.73	80.17	80.17	79.23
Pendigits	TT (s)	2.33	2.72	5.08	2.22	4.86
	CT (s)	161.30	161.50	53.19	54.39	170.83
	AC (%)	90.47	82.29	90.47	90.47	90.47
W8a	TT (s)	29.80	27.25	37.96	26.30	32.93
	CT (s)	2851.52	2721.82	999.72	1006.71	2659.24
	AC (%)	95.25	95.25	97.14	97.14	97.14

1. TT is the training time in seconds.
2. CT is the classification time in seconds.
3. AC is the accuracy in %.

4.3 Running Time and Space

Table 2 shows the running time for the training and classification phase of the five algorithms. The entries in bold signify that the running time is the smallest. The proposed Algorithm 1 is as fast and accurate as the other best method in the table, which is the modified version of ONMF in which angles are used for classification. However, we cannot directly compare the factorization returned by the two algorithms, as ONMF returns \mathbf{H} with orthogonal rows in which entries are real. In contrast, the method proposed here returns a sparse \mathbf{H} , which is binary and has orthogonal rows. Algorithm 1 is space efficient compared to the ONMF+cos method, which is an essential consideration for large data sets.

Based on the data in Table 2, we can conclude that Algorithm 1 has competitive accuracy and leads to better clustering with a natural interpretation and a sparse representation. It also classifies new data faster.

5 Conclusion

This paper gives a new geometric approach for binary orthogonal non-negative matrix factorization. The proposed method is space efficient. We also compared the proposed Algorithm with three other methods [15, 32, 34] for accuracy and time on representative datasets in machine learning. Our experiments show that the method is fast and accurate on the data sets tested.

Acknowledgements. The authors would like to thank Chirag Wadhwa for discussion and comments.

References

1. Albright, R., et al.: Algorithms, initializations, and convergence for the nonnegative matrix factorization. Technical report. 919. NCSU Technical Report Math 81706 (2006). <http://meyer.math.ncsu.edu/Meyer/Abstracts/Publications.html>
2. Arora, S., et al.: A practical algorithm for topic modeling with provable guarantees. In: International Conference on Machine Learning, pp. 280–288. PMLR (2013)
3. Bauckhage, C.: K-means clustering is matrix factorization. arXiv preprint [arXiv:1512.07548](https://arxiv.org/abs/1512.07548) (2015)
4. Bertsekas, D.P.: Projected Newton methods for optimization problems with simple constraints. *SIAM J. Control. Optim.* **20**(2), 221–246 (1982)
5. Cichocki, A., Zdunek, R., Amari, S.: Hierarchical ALS algorithms for nonnegative matrix and 3D tensor factorization. In: Davies, M.E., James, C.J., Abdallah, S.A., Plumbley, M.D. (eds.) ICA 2007. LNCS, vol. 4666, pp. 169–176. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74494-8_22
6. Fu, X., et al.: Nonnegative matrix factorization for signal and data analytics: identifiability, algorithms, and applications. *IEEE Signal Process. Mag.* **36**(2), 59–80 (2019)
7. Farhi, E., et al.: Quantum computation by adiabatic evolution. arXiv preprint [quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106) (2000)

8. Gao, M., et al.: Feature fusion and non-negative matrix factorization based active contours for texture segmentation. *Signal Process.* **159**, 104–118 (2019)
9. Golden, J., O'Malley, D.: Reverse annealing for nonnegative/binary matrix factorization. *PLoS ONE* **16**(1), e0244026 (2021)
10. Haddock, J., et al.: Semi-supervised nonnegative matrix factorization for document classification. In: 2021 55th Asilomar Conference on Signals, Systems, and Computers, pp. 1355–1360. IEEE (2021)
11. Han, L., Neumann, M., Prasad, U.: Alternating projected Barzilai-Borwein methods for nonnegative matrix factorization. *Electron. Trans. Numer. Anal.* **36**(6), 54–82 (2009)
12. Ho, N.-D.: Nonnegative matrix factorization algorithms and applications. Ph.D. thesis. Citeseer (2008)
13. Ho, N.-D., Van Dooren, P., Blondel, V.D.: Descent methods for nonnegative matrix factorization. In: Van Dooren, P., Bhattacharyya, S., Chan, R., Olshevsky, V., Routray, A. (eds.) *Numerical Linear Algebra in Signals, Systems and Control. Lecture Notes in Electrical Engineering*, vol. 80, pp. 251–293 (2011). https://doi.org/10.1007/978-94-007-0602-6_13
14. Kim, D., Sra, S., Dhillon, I.S. Fast Newton-type methods for the least squares nonnegative matrix approximation problem. In: *Proceedings of the 2007 SIAM International Conference on Data Mining*, pp. 343–354. SIAM (2007)
15. Lee, D.D., Seung, H.S.: Learning the parts of objects by non-negative matrix factorization. *Nature* **401**(6755), 788–791 (1999)
16. Lee, H., Yoo, J., Choi, S.: Semi-supervised nonnegative matrix factorization. *IEEE Signal Process. Lett.* **17**(1), 4–7 (2009)
17. Li, B., Zhou, G., Cichocki, A.: Two efficient algorithms for approximately orthogonal nonnegative matrix factorization. *IEEE Signal Process. Lett.* **22**(7), 843–846 (2014)
18. Lin, C.-J.: Projected gradient methods for nonnegative matrix factorization. *Neural Comput.* **19**(10), 2756–2779 (2007)
19. MacQueen, J.: Classification and analysis of multivariate observations. In: 5th Berkeley Symp. Math. Statist. Probability, pp. 281–297 (1967)
20. Paatero, P., Tapper, U.: Positive matrix factorization: a nonnegative factor model with optimal utilization of error estimates of data values. *Environmetrics* **5**(2), 111–126 (1994)
21. Pauca, V.P., et al.: Text mining using non-negative matrix factorizations. In: *Proceedings of the 2004 SIAM International Conference on Data Mining*, pp. 452–456. SIAM (2004)
22. Ran, X., et al.: A differentially private nonnegative matrix factorization for recommender system. *Inf. Sci.* **592**, 21–35 (2022)
23. Seung, D., Lee, L.: Algorithms for non-negative matrix factorization. *Adv. Neural. Inf. Process. Syst.* **13**, 556–562 (2001)
24. Shahnaz, F., et al.: Document clustering using nonnegative matrix factorization. *Inf. Process. Manage.* **42**(2), 373–386 (2006)
25. Stražar, M., et al.: Orthogonal matrix factorization enables integrative analysis of multiple RNA binding proteins. *Bioinformatics* **32**(10), 1527–1535 (2016)
26. Vanschoren, J., et al.: OpenML: networked science in machine learning. *ACM SIGKDD Explor. Newsl.* **15**(2), 49–60 (2014)
27. Vavasis, S.A.: On the complexity of nonnegative matrix factorization. *SIAM J. Optim.* **20**(3), 1364–1377 (2010)
28. Wold, S., Esbensen, K., Geladi, P.: Principal component analysis. *Chemometr. Intell. Lab. Syst.* **2**(1–3), 37–52 (1987)

29. Ye, Z., et al.: CDCN: a new NMF-based community detection method with community structures and node attributes. *Wireless Commun. Mobile Comput.* **2021**, 1–12 (2021)
30. Yoo, J.H., Choi, S.J.: Nonnegative matrix factorization with orthogonality constraints. *J. Comput. Sci. Eng.* **4**(2), 97–109 (2010)
31. Zdunek, R.: Data clustering with semi-binary nonnegative matrix factorization. In: Rutkowski, L., Tadeusiewicz, R., Zadeh, L.A., Zurada, J.M. (eds.) *ICAISC 2008. LNCS (LNAI)*, vol. 5097, pp. 705–716. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-69731-2_68
32. Zhang, M., et al.: Non-negative matrix factorization for binary space learning. In: *2021 13th International Conference on Advanced Computational Intelligence (ICACI)*, pp. 215–219. IEEE (2021)
33. Zhu, Y.-L., et al.: Ensemble adaptive total variation graph regularized NMF for single cell RNA-seq data analysis. *Curr. Bioinform.* **16**(8), 1014–1023 (2021)
34. Ding, C., et al.: Orthogonal nonnegative matrix t-factorizations for clustering. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 126–135 (2006)