
Fast Rank-1 NMF for Missing Data with KL Divergence

Kazu Ghalamkari

Mahito Sugiyama

National Institute of Informatics

The Graduate University for Advanced Studies, SOKENDAI

Abstract

We propose a fast non-gradient-based method of rank-1 non-negative matrix factorization (NMF) for missing data, called A1GM, that minimizes the KL divergence from an input matrix to the reconstructed rank-1 matrix. Our method is based on our new finding of an analytical closed-formula of the best rank-1 non-negative multiple matrix factorization (NMMF), a variety of NMF. NMMF is known to exactly solve NMF for missing data if positions of missing values satisfy a certain condition, and A1GM transforms a given matrix so that the analytical solution to NMMF can be applied. We empirically show that A1GM is more efficient than a gradient method with competitive reconstruction errors.

1 INTRODUCTION

The tabular form is one of the most common data types across different fields, and includes purchasing data, processed sensor data, and images. Tabular datasets are often treated as matrices for analysis. To date, many methods have been developed to extract essential information from matrices (Thompson, 1984; Wold et al., 1987; Comon, 1994; Blei et al., 2003). Non-negative matrix factorization (NMF) is one of the most popular techniques (Lee and Seung, 1999). NMF extracts factors in a dataset by decomposing a given non-negative matrix $\mathbf{X} \in \mathbb{R}_{\geq 0}^{I \times J}$ into a product \mathbf{AB} of two matrices $\mathbf{A} \in \mathbb{R}_{\geq 0}^{I \times r}$ and $\mathbf{B} \in \mathbb{R}_{\geq 0}^{r \times J}$ so that the predetermined cost becomes smaller. The matrix rank of the product \mathbf{AB} is less than or equal to the hyperparameter $r \in \mathbb{N}$.

Standard NMF, which uses the least squares error

$\|\mathbf{X} - \mathbf{AB}\|_F$ as the cost function, is widely used in various applications, including face recognition (Rajapakse et al., 2004), recommender systems (Takács et al., 2008), and text analysis (Xu et al., 2003). Although it is an NP-hard problem to find the best decomposition that minimizes the cost function exactly for any $r > 1$ (Vavasis, 2010), in an exceptional case of $r = 1$, the best decomposition is obtained in polynomial time (Gonzalez, 2007). Therefore it is possible to obtain the best decomposition in a reasonable time if $r = 1$ and a number of NMF algorithms are developed based on rank-1 NMF (Biggs et al., 2008; Liao and Zhang, 2013). Rank-1 NMF approximates an input matrix \mathbf{X} by the Kronecker product $\mathbf{a} \otimes \mathbf{b}$ of two non-negative vectors \mathbf{a} and \mathbf{b} , called dominant factors. Since the vectors \mathbf{a} and \mathbf{b} correspond to the largest principal components restricted to the first quadrant in eigenvalue decomposition, they are representative features that roughly describe the input matrix.

However, it is problematic if a matrix includes missing values, which often occurs in real-world datasets in practice. When the cost function is the least squares error, NMF for missing data — that is, a matrix with missing values, which we call *missing NMF* (Kim and Choi, 2009) — is known to become an NP-hard problem, even for $r = 1$ (Gillis and Glineur, 2011). Although this is a crucial drawback of missing NMF, missing NMF with other cost functions has not been well studied to date. In this paper, we show that there are certain cases in which rank-1 missing NMF can be exactly solved in polynomial time when the cost function is defined as the KL divergence.

Our key idea is that, instead of directly solving missing NMF, we focus on non-negative multiple matrix factorization (NMMF) (Takeuchi et al., 2013a), a variant of NMF. We derive a closed formula that globally minimizes the cost function of NMMF when the target rank $r = 1$, which is our main theoretical contribution. NMMF conducts simultaneous factor-sharing decomposition of multiple matrices, which has been used in purchase forecast systems (Kohjima et al., 2016) and recommender systems (Zhang et al., 2016). Interest-

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

ingly, if the cost function is given as the KL divergence, NMMF is equivalent to missing NMF when missing values are clustered in the lower right corner of the input matrix (Takeuchi et al., 2013a). Using this relationship between missing NMF and NMMF, we can efficiently compute the exact solution of rank-1 missing NMF with the KL divergence by our solution to NMMF when we can locate missing values in a rectangular region by permuting rows and columns.

Moreover, to treat matrices in which we cannot locate missing values in a rectangular region by permutations of rows and columns, we provide a method of finding an approximate solution of rank-1 missing NMF by adding more missing values so that we can use the closed formula of the best rank-1 NMMF. We call our novel method *A1GM* (Analytical solution for rank-1 NMF with Grid-based Missing values). We empirically show that A1GM is more efficient than an existing gradient-based method for rank-1 missing NMF with the competitive reconstruction error.

We summarize our contribution as follows:

- We derive a closed formula of the best rank-1 NMMF, which extracts the most dominant factors faster than the existing gradient method.
- We develop a novel efficient method to solve rank-1 missing NMF, called A1GM, and prove that A1GM globally minimizes a cost function under an assumption about the position of missing values.
- We empirically show that A1GM is more efficient than an existing method for missing NMF with competitive reconstruction error.

1.1 Notations and problem setup

Matrices are denoted by bold capital letters like \mathbf{X} and \mathbf{Y} , and vectors are denoted by lower-case bold alphabets like \mathbf{a} and \mathbf{b} . The total sum of a matrix or a vector is represented as $S(\cdot)$. The i th component of a vector \mathbf{a} is written in its non-bold letter as a_i . The Kronecker product of two vectors \mathbf{a} and \mathbf{b} is denoted by $(\mathbf{a} \otimes \mathbf{b})$, which is a rank-1 matrix, and each element is defined as $(\mathbf{a} \otimes \mathbf{b})_{ij} = a_i b_j$. The $I \times J$ all-one and all-zero matrices are denoted by $\mathbf{1}_{I \times J}$ and $\mathbf{0}_{I \times J}$, respectively. The identity matrix is denoted by \mathbf{I} . The transpose of a matrix \mathbf{X} is denoted by \mathbf{X}^T . The element-wise product of two matrices \mathbf{A} and \mathbf{B} is denoted by $\mathbf{A} \circ \mathbf{B}$. For a pair of natural numbers n and m ($\geq n$), we denote by $[n, m] = \{n, n+1, \dots, m-1, m\}$. We abbreviate $[1, m]$ as $[m]$. The set difference of B and A is denoted by $B \setminus A$. When we use the Kullback–Leibler (KL) divergence $D(\mathbf{X}, \mathbf{Y})$ for matrices \mathbf{X} and \mathbf{Y} , it is

defined as follows: (Lee and Seung, 2001)

$$D(\mathbf{X}, \mathbf{Y}) = \sum_{i,j} \left\{ \mathbf{X}_{ij} \log \frac{\mathbf{X}_{ij}}{\mathbf{Y}_{ij}} - \mathbf{X}_{ij} + \mathbf{Y}_{ij} \right\}.$$

In this paper, we treat two tasks, rank-1 NMMF and rank-1 missing NMF. We consistently assume that these cost functions are defined by the above KL divergence throughout the paper.

Rank-1 NMMF Rank-1 NMMF simultaneously decomposes three matrices $\mathbf{X} \in \mathbb{R}_{>0}^{I \times J}$, $\mathbf{Y} \in \mathbb{R}_{>0}^{N \times J}$ and $\mathbf{Z} \in \mathbb{R}_{>0}^{I \times M}$ into three rank-1 matrices $\mathbf{w} \otimes \mathbf{h}$, $\mathbf{a} \otimes \mathbf{h}$, and $\mathbf{w} \otimes \mathbf{b}$, respectively, using non-negative vectors $\mathbf{w} \in \mathbb{R}_{>0}^I$, $\mathbf{h} \in \mathbb{R}_{>0}^J$, $\mathbf{a} \in \mathbb{R}_{>0}^N$, and $\mathbf{b} \in \mathbb{R}_{>0}^M$. The cost function of NMMF is defined as

$$D(\mathbf{X}, \mathbf{w} \otimes \mathbf{h}) + \alpha D(\mathbf{Y}, \mathbf{a} \otimes \mathbf{h}) + \beta D(\mathbf{Z}, \mathbf{w} \otimes \mathbf{b}), \quad (1)$$

and the task of rank-1 NMMF is to find vectors \mathbf{w} , \mathbf{h} , \mathbf{a} , and \mathbf{b} that minimize the above cost. We assume that the scaling parameters α and β are non-negative

Rank-1 missing NMF Rank-1 NMF for a matrix with missing values (rank-1 missing NMF) is the task of finding two non-negative vectors $\mathbf{w} \in \mathbb{R}_{\geq 0}^I$ and $\mathbf{h} \in \mathbb{R}_{\geq 0}^J$ that minimize a weighted cost function $D_{\Phi}(\mathbf{X}, \mathbf{w} \otimes \mathbf{h})$ defined as

$$D_{\Phi}(\mathbf{X}, \mathbf{w} \otimes \mathbf{h}) = D(\Phi \circ \mathbf{X}, \Phi \circ (\mathbf{w} \otimes \mathbf{h})) \quad (2)$$

for a non-negative input matrix $\mathbf{X} \in \mathbb{R}_{\geq 0}^{I \times J}$ and a binary weight matrix $\Phi \in \{0, 1\}^{I \times J}$. The weight matrix indicates the position of missing values; that is, $\Phi_{ij} = 0$ if the entry \mathbf{X}_{ij} is missing, $\Phi_{ij} = 1$ otherwise. Note that the above cost function (2) is always convex in \mathbf{w} and \mathbf{h} .

2 THE A1GM METHOD

In this section, we present our proposed method for rank-1 missing NMF, called A1GM. We derive a closed formula of the best rank-1 NMMF in Section 2.1, followed by introducing A1GM using the closed formula in Section 2.2.

2.1 A closed formula of the best rank-1 NMMF

We give the following closed formula of the best rank-1 NMMF that exactly minimizes the cost function in Equation (1), which is our main theoretical contribution. This formula efficiently extracts only the most dominant factors from three input matrices.

Theorem 1 (the best rank-1 NMMF). *For any given three positive matrices $\mathbf{X} \in \mathbb{R}_{>0}^{I \times J}$, $\mathbf{Y} \in \mathbb{R}_{>0}^{N \times J}$, and*

$\mathbf{Z} \in \mathbb{R}_{>0}^{I \times M}$ and two parameters $\alpha, \beta \geq 0$, four non-negative vectors $\mathbf{w} \in \mathbb{R}_{\geq 0}^I, \mathbf{h} \in \mathbb{R}_{\geq 0}^J, \mathbf{a} \in \mathbb{R}_{\geq 0}^N$, and $\mathbf{b} \in \mathbb{R}_{\geq 0}^M$ that minimize the cost function in Equation (1) is given as

$$\begin{aligned} w_i &= \frac{\sqrt{S(\mathbf{X})}}{S(\mathbf{X}) + \beta S(\mathbf{Z})} \left\{ \sum_{j=1}^J \mathbf{X}_{ij} + \sum_{m=1}^M \beta \mathbf{Z}_{im} \right\}, \\ h_j &= \frac{\sqrt{S(\mathbf{X})}}{S(\mathbf{X}) + \alpha S(\mathbf{Y})} \left\{ \sum_{i=1}^I \mathbf{X}_{ij} + \sum_{n=1}^N \alpha \mathbf{Y}_{nj} \right\}, \\ a_n &= \frac{\sum_{j=1}^J \mathbf{Y}_{nj}}{\sqrt{S(\mathbf{X})}}, \quad b_m = \frac{\sum_{i=1}^I \mathbf{Z}_{im}}{\sqrt{S(\mathbf{X})}}. \end{aligned}$$

The time complexity to obtain the best rank-1 NMMF is $O(IJ + NJ + IM)$. Note that, for $N = M = 0$, our result in Theorem 1 coincides with the best rank-1 NMF minimizing the KL divergence from an input matrix \mathbf{X} to the reconstructed rank-1 matrix shown in Ho and Van Dooren (2008).

In subsections 2.1.1 – 2.1.3, we introduce information geometric formulation of NMMF using the log-linear model and derive the closed-form solution given in Theorem 1. Although they can be viewed as a proof of Theorem 1, we provide it in the main text as it includes various interesting properties between NMMF and information geometry, particularly characterization of rank-1 decomposition via parameters of the exponential family, which we believe is valuable for further development of this line of research.

2.1.1 Projection theory in information geometry

To prove Theorem 1, we first prepare the general projection theory in information geometry (Amari, 2016).

The exponential family is a set of probability distributions that can be represented as $\log P_{\boldsymbol{\theta}}(x) = C(x) + \sum_{i=1}^N \theta_i F_i(x) - \psi(\boldsymbol{\theta})$ using *natural parameters* $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N) \in \mathbb{R}^N$. It is known that *expectation parameters* $\eta_i = \mathbb{E}_{P_{\boldsymbol{\theta}}}[F_i]$ can be uniquely determined by the normalization factor $\psi(\boldsymbol{\theta})$ and natural parameters $\boldsymbol{\theta}$. Since both natural parameters $\boldsymbol{\theta}$ and expectation parameters $\boldsymbol{\eta} = (\eta_1, \dots, \eta_N)$ can identify a distribution, we can use them as coordinate systems for distributions, which is a typical approach in information geometry.

Let \mathcal{Q} be the set of distributions that satisfy a condition $f(\boldsymbol{\theta}_{1:n}) = 0$ for a linear function $f(\cdot)$ on a part of the natural parameters $\boldsymbol{\theta}_{1:n} = (\theta_1, \dots, \theta_n)$, where we assume that this part is from 1 to n without loss of generality. This set \mathcal{Q} is a subspace characterized by the $\boldsymbol{\theta}$ -coordinate system. Let us find a distribution $Q \in \mathcal{Q}$ that is nearest to a given distribution P in terms of the




Figure 1: A partial order structure for NMMF for three input matrices $\mathbf{X} \in \mathbb{R}_{>0}^{I \times J}, \mathbf{Y} \in \mathbb{R}_{>0}^{N \times J}$, and $\mathbf{Z} \in \mathbb{R}_{>0}^{I \times M}$. Only natural parameters on gray-colored nodes can have non-zero values if and only if $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable.

KL divergence; that is, solve the optimization problem

$$Q = \operatorname{argmin}_{Q' \in \mathcal{Q}} D(P, Q').$$

This constrained optimization is called the *m-projection* of P onto \mathcal{Q} . When the condition of \mathcal{Q} is linear to $\boldsymbol{\theta}$, the projection destination Q always uniquely exists and this optimization becomes a convex problem. Moreover, this *m-projection* does not change the rest of the part of expectation parameters $\boldsymbol{\eta}_{n+1:N} = (\eta_{n+1}, \dots, \eta_N)$ (Amari, 2008). In this paper, we call this property *expectation conservation law in m-projection*. It has already been reported that we can conduct non-negative tensor Tucker rank reduction efficiently using this conservation law (Ghalamkari and Sugiyama, 2021).

2.1.2 Modeling

The input of NMMF is a triple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, where $\mathbf{X} \in \mathbb{R}_{>0}^{I \times J}, \mathbf{Y} \in \mathbb{R}_{>0}^{N \times J}$, and $\mathbf{Z} \in \mathbb{R}_{>0}^{I \times M}$. For simplicity, we normalize them beforehand so that their sum is 1; that is, $S(\mathbf{X}) + S(\mathbf{Y}) + S(\mathbf{Z}) = 1$. It is straightforward to eliminate this assumption using the property of KL divergence, $\alpha D(\mathbf{X}, \mathbf{Y}) = D(\alpha \mathbf{X}, \alpha \mathbf{Y})$, for any non-negative number α . We model these three matrices using a single discrete distribution equipped with a structured sample space.

Let p be a probability mass function with the sample

space Ω given as

$$\begin{aligned} \Omega &= \Omega_{\mathbf{X}} \cup \Omega_{\mathbf{Y}} \cup \Omega_{\mathbf{Z}}, \text{ where} \\ \Omega_{\mathbf{X}} &= [I] \times [J], \quad \Omega_{\mathbf{Y}} = [I+1, I+N] \times [J], \\ \Omega_{\mathbf{Z}} &= [I] \times [J+1, J+M]. \end{aligned}$$

For each $(k, l) \in \Omega$, probability $p(k, l)$ is defined as

$$p(k, l) = \exp \left(\sum_{(s,t) \leq (k,l)} \theta_{st} \right), \quad (3)$$

where we use a partial order “ \leq ” over elements of Ω , defined as $(k, l) \leq (s, t)$ if and only if $k \leq s$ and $l \leq t$. It is clear that this distribution in Equation (3) belongs to the exponential family and $\boldsymbol{\theta} = (\theta_{ij})_{(i,j) \in \Omega \setminus \{(1,1)\}}$ are natural parameters of the distribution, hence there is a one-to-one correspondence between $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ and $\boldsymbol{\theta}$. The natural parameter θ_{11} , in which $(1, 1)$ is the smallest element in the partially ordered space Ω , corresponds to the normalization factor $\psi(\boldsymbol{\theta})$. The natural parameters are identified so that they satisfy

$$\mathbf{X}_{ij} = p(i, j), \quad \mathbf{Y}_{nj} = p(I+n, j), \quad \mathbf{Z}_{im} = p(i, J+m),$$

where the subspace $\Omega_{\mathbf{X}}$ corresponds to \mathbf{X} , $\Omega_{\mathbf{Y}}$ to \mathbf{Y} , and $\Omega_{\mathbf{Z}}$ to \mathbf{Z} . This model is a special case of a log-linear model on a poset (Sugiyama et al., 2017). Figure 1 illustrates the partial order for the input triple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$. There are other possible ways to model $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ as a probability distribution using different partial order structure. However, the solution formula that we obtain in Theorem 1 does not depend on the modeling.

Since we use the exponential function in our modeling in Equation (3), our model cannot treat 0. Therefore, our mathematical statements are guaranteed with only positive inputs.

Our model belongs to the exponential family, which makes it possible to analyze NMMF from the viewpoint of information geometry (Amari, 2016). The generalization of this model, the log-linear model on posets, is flexible to handle discrete structures, so our model is able to handle the partial order structure of the NMMF. As described in the previous subsection, we can also identify an exponential distribution by the expectation parameters $\boldsymbol{\eta}$ instead of the natural parameters $\boldsymbol{\theta}$. Using results provided in (Sugiyama et al., 2017) based on the incidence algebra between natural parameters $\boldsymbol{\theta}$ and expectation parameters $\boldsymbol{\eta}$, we can also obtain the expectation parameters $\boldsymbol{\eta}$ using the following formula:

$$\eta_{kl} = \sum_{(k,l) \leq (s,t)} p(s, t).$$

Therefore we have a one-to-one mapping between an input tuple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, $\boldsymbol{\theta}$, and $\boldsymbol{\eta}$.




Figure 2: Information geometric view of rank-1 NMMF. Rank-1 NMMF is m -projection onto simultaneous rank-1 subspace from a triple of input three matrices, where one-body expectation parameters do not change.

2.1.3 Derivation of the exact solution of rank-1 NMMF

Let $\mathbf{w} \in \mathbb{R}_{>0}^I$, $\mathbf{h} \in \mathbb{R}_{>0}^J$, $\mathbf{a} \in \mathbb{R}_{>0}^N$, and $\mathbf{b} \in \mathbb{R}_{>0}^M$. If three positive matrices $\mathbf{X} \in \mathbb{R}_{>0}^{I \times J}$, $\mathbf{Y} \in \mathbb{R}_{>0}^{N \times J}$, and $\mathbf{Z} \in \mathbb{R}_{>0}^{I \times M}$ can be decomposed into a form $\mathbf{w} \otimes \mathbf{h}$, $\mathbf{a} \otimes \mathbf{h}$, and $\mathbf{w} \otimes \mathbf{b}$, we say that $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is *simultaneously rank-1 decomposable*. Note that the subspace \mathcal{Q} in Section 2.1.1 corresponds to the set of simultaneous rank-1 decomposable tuples.

To describe the necessary and sufficient conditions for $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ to be simultaneously rank-1 decomposable, we define *one-body* parameters and *two-body* parameters, which were also used by Ghalamkari and Sugiyama (2021). A parameter θ_{ij} or η_{ij} is a one-body parameter if one of the indices is 1; that is, $i = 1$ or $j = 1$. Parameters other than one-body parameters are two-body parameters; that is, θ_{ij}, η_{ij} are two-body parameters if $i \neq 1$ and $j \neq 1$. Using these parameters, we obtain the following propositions.

Proposition 1 (simultaneous rank-1 θ -condition). *A triple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable if and only if its all two-body natural parameters are 0; that is, for any $(i, j) \in \Omega$,*

$$\theta_{ij} = 0 \quad \text{if } i \neq 1 \text{ and } j \neq 1.$$

In Figure 1, only gray-colored elements have non-zero values of natural parameters $\boldsymbol{\theta}$ if and only if $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable. Moreover, we provide the simultaneously rank-1 decomposable condition using $\boldsymbol{\eta}$ -parameters as well as on $\boldsymbol{\theta}$ -parameters.

Proposition 2 (simultaneous rank-1 η -condition). *A triple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable if and only if its all two-body expectation parameters are factorizable as a product of two one-body parameters:*

for any $(i, j) \in \Omega$,

$$\eta_{ij} = \eta_{i1}\eta_{1j}.$$

We call a subspace that satisfies simultaneous rank-1 condition *simultaneous rank-1 subspace*. From the viewpoint of information geometry, we can understand the best rank-1 NMMF as follows (shown in Figure 2). The input of NMMF $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ corresponds to a point in the space described by the θ -coordinate system. The best rank-1 NMMF is an m -projection onto the simultaneous rank-1 subspace from the input point.

Since the m -projection is a convex optimization, we can get the projection destination by a gradient method. However, it requires appropriate settings for initial values, stopping criterion, and learning rates. In addition, the gradient method is often computationally time-consuming because it involves iterative procedures.

Our closed analytical formula of the projection destination in Theorem 1 solves all the drawbacks of the gradient-based optimization. According to the expectation conservation law in this m -projection onto simultaneous rank-1 subspace, two-body η -parameters do not change in the m -projection. That is, for any $(i, j) \in \Omega$,

$$\eta_{1j} = \bar{\eta}_{1j}, \quad \eta_{i1} = \bar{\eta}_{i1},$$

where η is the expectation parameter of input, and $\bar{\eta}$ is the expectation parameter after the m -projection onto simultaneous rank-1 subspace. By the definition of expectation parameters, we obtain

$$\begin{cases} \eta_{1j}^{\mathbf{X}} - \eta_{1,j+1}^{\mathbf{X}} = (S(\mathbf{w}) + S(\mathbf{a})) h_j \\ \eta_{i1}^{\mathbf{X}} - \eta_{i+1,1}^{\mathbf{X}} = w_i (S(\mathbf{h}) + S(\mathbf{b})) \\ \eta_{1j}^{\mathbf{Y}} - \eta_{1,j+1}^{\mathbf{Y}} = S(\mathbf{w}) b_j \\ \eta_{i1}^{\mathbf{Z}} - \eta_{i+1,1}^{\mathbf{Z}} = a_i S(\mathbf{h}). \end{cases}$$

where $\eta_{ij}^{\mathbf{X}} = \eta_{ij}$ for $(i, j) \in \Omega_{\mathbf{X}}$, $\eta_{i-I,j}^{\mathbf{Y}} = \eta_{i,j}$ for $(i, j) \in \Omega_{\mathbf{Y}}$, and $\eta_{i,j-J}^{\mathbf{Z}} = \eta_{i,j}$ for $(i, j) \in \Omega_{\mathbf{Z}}$.

The expectation conservation law guarantees that the values of the left-hand sides do not change before the m -projection, and they also do not change after the m -projection. Since the sum of each matrix $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ is represented by one-body η -parameters, the sum of each matrix does not change either in m -projection. Using these facts and multiplying these equations together, we can derive Theorem 1. A complete proof of Theorem 1 is in Supplementary Material.

2.2 Rank-1 missing NMF using Theorem 1

In missing NMF, NMF for missing data, when a binary weight matrix $\Phi \in \{0, 1\}^{I+N, M+J}$ and an input matrix




Figure 3: Examples of matrices with non-grid-like missing values (left) and grid-like missing values (right). Meshed blocks are missing values. We can create grid-like missing values by increase missing values.

$\mathbf{T} \in \mathbb{R}_{\geq 0}^{I+N, M+J}$ are given in the form of

$$\Phi = \begin{bmatrix} \mathbf{1}_{IJ} & \mathbf{1}_{IM} \\ \mathbf{1}_{NJ} & \mathbf{0}_{NM} \end{bmatrix}, \quad \mathbf{T} = \begin{bmatrix} \mathbf{X} & \mathbf{Z} \\ \mathbf{Y} & \mathbf{U} \end{bmatrix}, \quad (4)$$

the cost function of missing NMF is equivalent to that of NMMF (Takeuchi et al., 2013a):

$$\begin{aligned} & \operatorname{argmin}_{\mathbf{T}_1: \operatorname{rank}(\mathbf{T}_1)=1} D_{\Phi}(\mathbf{T}, \mathbf{T}_1) \\ & = \operatorname{argmin}_{\mathbf{w}, \mathbf{h}, \mathbf{a}, \mathbf{b}} D(\mathbf{X}, \mathbf{w} \otimes \mathbf{h}) + D(\mathbf{Y}, \mathbf{a} \otimes \mathbf{h}) + D(\mathbf{Z}, \mathbf{w} \otimes \mathbf{b}). \end{aligned}$$

Therefore, using the closed formula of the best rank-1 NMMF in Theorem 1, we can solve the rank-1 missing NMF when the binary weight matrix Φ is given in the form in Equation (4). In what follows, we always assume that positions of zero entries of Φ coincide with those of missing values of a given input matrix \mathbf{T} , and we identify them in our discussion.

2.2.1 A1GM with grid-like missing values

Using the fact that rank-1 missing NMF is homogeneous for row and column permutations, we can develop a method for rank-1 missing NMF, called A1GM.

Proposition 3 (Homogeneity of rank-1 missing NMF). *Let $\operatorname{NMF}_1(\Phi, \mathbf{X})$ be the best rank-1 matrix $\mathbf{w} \otimes \mathbf{h}$, which minimizes the cost function in Equation (2). For any permutation matrices \mathbf{G} and \mathbf{H} , it holds that*

$$\operatorname{NMF}_1(\mathbf{G}\Phi\mathbf{H}, \mathbf{G}\mathbf{T}\mathbf{H}) = \mathbf{G}^T \operatorname{NMF}_1(\Phi, \mathbf{T}) \mathbf{H}^T.$$

We can then guarantee that A1GM gives the best solution that globally minimizes Equation (1) if an input binary weight is grid-like, defined as follows.

Definition 1 (grid-like binary weight matrix). *Let $\Phi \in \{0, 1\}^{I+N, J+M}$ be a binary weight matrix. If there exist two sets $S^{(1)} \subset [I+N]$ and $S^{(2)} \subset [J+M]$ such that*

$$\Phi_{ij} = \begin{cases} 0 & \text{if } i \in S^{(1)} \text{ and } j \in S^{(2)}, \\ 1 & \text{otherwise,} \end{cases}$$




Figure 4: Sketch of the algorithm of A1GM. Meshed blocks are missing values. In Step 1, we increase missing values so that they become grid-like. In Step 2, we gather missing values in the block at the bottom right by low and column permutations. In Step 3, we use the closed formula of the best rank-1 NMMF in Theorem 1. In this example, we get $\mathbf{w} = (1.5, 1.3, 1.9)$, $\mathbf{a} = (1.9, 1.1)$, $\mathbf{h} = (1.8, 1.6, 1.3)$, $\mathbf{a} = (0.85, 3.4)$. Finally, we get two vectors as the output by the re permutation. We use two significant digits in this figure.

Φ is called *grid-like*.

Real-world tabular datasets tend to have missing values on only certain rows or columns. Therefore, the binary weight matrix Φ often becomes grid-like in practice (we show examples in Section 3). Figure 3 illustrates examples of matrices with grid-like missing values.

When $\Phi \in \{0, 1\}^{I+N, J+M}$ is grid-like, we can transform it in the form given in Equation (4) using only row and column permutations. Let $S^{(1)} \subset [I+N]$ and $S^{(2)} \subset [J+M]$ with $|S^{(1)}| = K$ and $|S^{(2)}| = L$ be the row and column index sets for zero entries in Φ . For the block at the bottom right of Ω whose row and column indices are specified as

$$\begin{aligned} B^{(1)} &= [I+N-K+1, I+N], \\ B^{(2)} &= [J+M-L+1, J+M], \end{aligned}$$

we can collect all the zero entries of Φ in the rectangular region $B^{(1)} \times B^{(2)}$ using only row and column permutations. Formally, for a grid-like binary weight matrix Φ , there are row $\mathcal{G} : [I+N] \rightarrow [I+N]$ and column $\mathcal{H} : [J+M] \rightarrow [J+M]$ permutations satisfying

$$(\mathbf{G}\Phi\mathbf{H})_{ij} = \begin{cases} 0 & \text{if } i \in B^{(1)} \text{ and } j \in B^{(2)} \\ 1 & \text{otherwise} \end{cases}$$

where \mathbf{G} and \mathbf{H} are corresponding permutation matrices to \mathcal{G} and \mathcal{H} , respectively.

We can obtain \mathbf{G} and \mathbf{H} as follows. First, we focus on row permutation \mathcal{G} . We want to include each row $j \in S^{(1)} \cap B^{(1)c}$, where $B^{(1)c} = [I+N] \setminus B^{(1)}$, in $B^{(1)}$ by row permutation \mathcal{G} , which can be achieved by any one-to-one mapping from $S^{(1)} \cap B^{(1)c}$ to $S^{(1)c} \cap B^{(1)}$, where $S^{(1)c} = [I+N] \setminus S^{(1)}$. Note that $|S^{(1)} \cap B^{(1)c}| = |S^{(1)c} \cap B^{(1)}|$ always holds. The corresponding permutation matrix is given as

$$\mathbf{G} = \prod_{k \in S^{(1)} \cap B^{(1)c}} \mathbf{R}^{k \leftrightarrow \mathcal{G}(k)}$$

where $\mathbf{R}^{k \leftrightarrow l}$ is a permutation matrix, which switches the k -th row and the l -th row; that is,

$$\mathbf{R}_{ij}^{k \leftrightarrow l} = \begin{cases} 0 & \text{if } (i, j) = (k, k) \text{ or } (l, l), \\ 1 & \text{if } (i, j) = (k, l) \text{ or } (l, k), \\ \mathbf{I}_{ij} & \text{otherwise.} \end{cases}$$

Since $S^{(1)} \cap B^{(1)c}$ and $S^{(1)c} \cap B^{(1)}$ are disjoint, it holds that $\mathbf{G} = \mathbf{G}^T$.

In the same way, any one-to-one mapping from $S^{(2)} \cap B^{(2)c}$ to $S^{(2)c} \cap B^{(2)}$ can be \mathcal{H} and the corresponding permutation matrix is given as

$$\mathbf{H} = \prod_{k \in S^{(2)} \cap B^{(2)c}} \mathbf{R}^{k \leftrightarrow \mathcal{H}(k)},$$

which is also a symmetric matrix.

The above discussion leads to the following procedure of the best rank-1 missing NMF for an input matrix \mathbf{T} if a binary weight matrix Φ is grid-like. The first step is to find proper permutations \mathbf{G} and \mathbf{H} to collect the missing values in the lower right corner. In the next step, we obtained $\text{NMF}_1(\mathbf{G}\Phi\mathbf{H}, \mathbf{G}\mathbf{T}\mathbf{H})$ using the closed formula of the best rank-1 NMMF. In the final step, we operated the inverse permutations of \mathbf{G} and \mathbf{H} to the result of the previous step; that is, $\mathbf{G}^{-1} \text{NMF}_1(\mathbf{G}\Phi\mathbf{H}, \mathbf{G}\mathbf{T}\mathbf{H}) \mathbf{H}^{-1}$. Note that $\mathbf{G}^{-1} = \mathbf{G}^T = \mathbf{G}$ and $\mathbf{H}^{-1} = \mathbf{H}^T = \mathbf{H}$ always holds since these permutation matrices are orthogonal and symmetrical.

2.2.2 A1GM for general case

If a binary weight matrix Φ is not grid-like, the above procedure cannot be directly applied. To treat any matrices with missing values, we *increase* missing values so that Φ becomes grid-like. Although this strategy is counter-intuitive because we lose some information, which may cause a larger reconstruction error, we gain the efficiency instead using our closed-form solution in

Theorem 1, and, as we empirically show in the next section, the error increase is not significant in many datasets. Examples of this step are demonstrated in Figure 3.

In the worst case, the number of missing values after this step becomes k^2 for k missing values. If every row or column has at least one missing value, all indices are missing after this step, for which our algorithm does not work. Thus, our method is not suitable if there are too many missing values in a matrix.

We illustrate an example of the overall procedure of A1GM in Figure 4 and show its algorithm in Algorithm 1. Since the time complexity of each process of A1GM is at most linear with respect to the number of entries of an input matrix, the time complexity is $O((I + N)(J + M))$.

Relation to *em*-algorithm As a method to solve rank-1 NMF with missing values, the *em*-algorithm repeats the following *e*- and *m*-steps after filling missing values of an input matrix \mathbf{T} with arbitrary values (Zhang et al., 2006).

***m*-step** : Get the rank-1 approximation of the input matrix \mathbf{T} that minimizes the KL divergence from \mathbf{T} .

***e*-step** : Overwrite missing values of \mathbf{T} by the obtained matrix in the *m*-step with keeping other values.

This algorithm also minimizes the cost function (2) indirectly (Amari, 2016). For grid-like data, A1GM directly finds the convergence point of the *em*-algorithm without performing the above iterations.

3 EXPERIMENTS

In this section, we empirically investigate the efficiency and effectiveness of our proposed method, A1GM, on synthetic and real data. We use three types of data: (i) synthetic data with missing values at the lower right corner, (ii) synthetic data with random grid-like missing values, and (iii) real data with grid-like and non-grid-like missing values. It is guaranteed that A1GM always finds the best solution for any data of (i) and (ii). Thus, we only investigate efficiency in our experiments for (i) and (ii). By contrast, for data (iii), the reconstruction error can be worse than the existing methods due to increased missing values in A1GM. Therefore, in the experiment for data (iii), we investigate both efficiency (running time) and effectiveness (reconstruction error).

We use KL-WNMF as a comparison method (Blondel et al., 2007). KL-WNMF is a commonly used gradient method that reduces the KL-based cost in Equation (2) by multiplicative updates. Although faster NMF methods, such as ALS (Kim et al., 2014) and ADMM (Hajinezhad et al., 2016), have been developed,

Algorithm 1: A1GM

input : A binary weight matrix $\Phi \in \{0, 1\}^{I \times J}$,
a matrix $\mathbf{X} \in \mathbb{R}_{\geq 0}^{I \times J}$
output : Dominant factors $\mathbf{c} \in \mathbb{R}_{\geq 0}^I$ and
 $\mathbf{d} \in \mathbb{R}_{\geq 0}^J$
A1GM(Φ, \mathbf{X})
 $S^{(1)} \leftarrow \emptyset; S^{(2)} \leftarrow \emptyset$
for $(i, j) \in [I] \times [J]$ **do**
 if $\Phi_{ij} = 0$ **then**
 $S^{(1)} \leftarrow S^{(1)} \cup \{i\}; S^{(2)} \leftarrow S^{(2)} \cup \{j\}$
 $B^{(1)} \leftarrow \{I - |S^{(1)}| + 1, I - |S^{(1)}| + 2, \dots, I\}$
 $B^{(2)} \leftarrow \{J - |S^{(2)}| + 1, J - |S^{(2)}| + 2, \dots, J\}$
 $\text{perm1} \leftarrow (1, \dots, I)$
 for $k \in \{1, 2, \dots, |S^{(1)} \cap B^{(1)c}|\}$ **do**
 $i \leftarrow k$ th smallest element of $S^{(1)} \cap B^{(1)c}$
 $j \leftarrow k$ th smallest element of $S^{(1)c} \cap B^{(1)}$
 swap($\text{perm1}[i], \text{perm1}[j]$)
 $\text{perm2} \leftarrow (1, \dots, J)$
 for $k \in \{1, 2, \dots, |S^{(2)} \cap B^{(2)c}|\}$ **do**
 $i \leftarrow k$ th smallest element of $S^{(2)} \cap B^{(2)c}$
 $j \leftarrow k$ th smallest element of $S^{(2)c} \cap B^{(2)}$
 swap($\text{perm2}[i], \text{perm2}[j]$)
 $\mathbf{X} \leftarrow \mathbf{X}[\text{perm1}, \text{perm2}]$
 $\mathbf{w}, \mathbf{h}, \mathbf{a}, \mathbf{b} \leftarrow$ the best rank-1 NMMF of \mathbf{X}
 $\mathbf{c} \leftarrow$ concatenate \mathbf{w} and \mathbf{a} ; $\mathbf{d} \leftarrow$ concatenate \mathbf{h} and \mathbf{b}
 $\mathbf{c} \leftarrow \mathbf{c}[\text{perm1}]; \mathbf{d} \leftarrow \mathbf{d}[\text{perm2}]$
return \mathbf{c}, \mathbf{d}

they are just as fast as a multiplicative update when the target rank is small (Song et al., 2014). Moreover, as we will show in Sections 3.1 and 3.2, KL-WNMF converges within only 2–4 iterations in our experiments. In addition, the *em*-algorithm needs more iterations since it minimizes the cost function indirectly. Thus, these techniques are considered to be ineffective for speeding up rank-1 KL-WNMF. This is why we only compared A1GM with simple KL-WNMF.

We implemented KL-WNMF by referring to the original paper (Blondel et al., 2007). The stopping criterion of KL-WNMF follows the implementation of the standard NMF in scikit-learn (Pedregosa et al., 2011). The initial values of KL-WNMF are determined by sampling from a uniform distribution from 0 to 1.

All methods are implemented in **Julia 1.6**. We used **BenchmarkTools** to measure the running time (Chen and Revels, 2016). Experiments were conducted on Ubuntu 20.04.1 with a single core of 2.1GHz Intel Xeon CPU Gold 5218 and 128GB of memory.

Table 1: Performance of A1GM compared to KL-WNMF on 17 real datasets.

DataSet	size	# missing values	increase rate	relative error	relative runtime
Autompg	(398, 8)	6	1	1	0.12957
DailySunSpot	(73718, 9)	3247	1	1	0.12845
CaliforniaHousing	(20640, 9)	207	1	1	0.11821
MTSLibrary	(1533078, 4)	1247722	1	1	0.18327
BigMartSaleForecas	(8522, 5)	1463	1	1	0.12699
BoardGameGeekData	(101375, 17)	21	1	1	0.14625
CreditCardApproval	(590, 7)	25	1.92	1.0018	0.12212
HumanResourceAnaly	(14999, 7)	519	1.96146	1.0168	0.11858
heartdisease	(303, 14)	6	2	1	0.12259
lungcancer	(32, 57)	5	2	1.0001	0.13803
PerthHousePrice	(33656, 14)	16585	2.61345	1.0004	0.15382
SleepData	(62, 8)	12	2.75	1.0211	0.18208
arrhythmia	(452, 280)	408	4.70588	1.0148	0.11387
Bostonhousing	(506, 14)	120	5.6	1.003	0.1097
LifeExpectancyData	(2938, 19)	2563	7.04097	5.7983	0.095773
HCCSurvivalDataSet	(165, 50)	826	8.3632	3.2898	0.07113
wiki4HE	(913, 53)	1995	18.10175	1.2363	0.066256

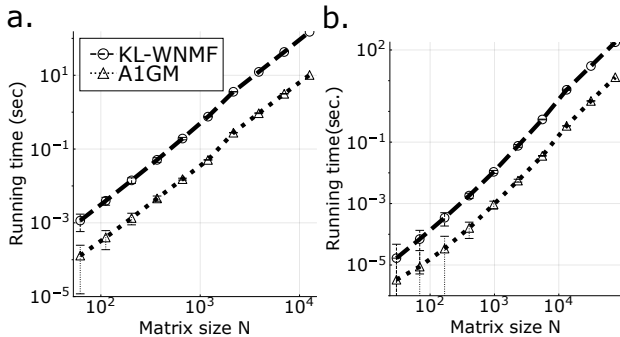


Figure 5: Running time comparison of the proposed method A1GM (triangle, dots line) and KL-WNMF (circle, dashed line) with respect to the matrix size N . (a) Missing values are at the lower right corner. (b) Missing value positions are grid-like. We plot mean \pm S.D. of five trials.

3.1 Synthetic datasets

Missing values in the lower right corner We prepared synthetic matrices $\mathbf{X} \in \mathbb{R}^{N \times N}$ and their weights $\Phi \in \{0, 1\}^{N \times N}$. We assumed that each input weight Φ is in the form of Equation (4). We measured the running time to obtain rank-1 decomposition of \mathbf{X} with varying the matrix size N . As Figure 5(a) shows, A1GM is an order of magnitude faster than the existing gradient method. The number of iterations of the existing method until convergence was between 2 and 4. A1GM just applies the closed formula in Theorem 1 to parts of input matrices.

Random grid-like missing values We also prepared synthetic matrices and its binary weight matrices $\Phi \in \{0, 1\}^{N \times N}$. We assumed that every input weight

matrix Φ is grid-like and we set the ratio of missing values to be 5 percent. We measured the running time of A1GM to complete the best rank-1 missing NMF compared with KL-WNMF by varying the matrix size N . As Figure 5(b) shows, our method is always faster than the gradient method. The number of iterations of the existing method required for convergence was between 2 and 4. Note that in these datasets A1GM does not need to increase missing values.

3.2 Real datasets

We used 17 real datasets. We downloaded tabular datasets that have missing values from the Kaggle databank¹ or UCI dataset.² If a dataset contains negative values, we converted them to their absolute values. Zero values in a matrix were replaced with the average value of the matrix to make them all positive. See Supplementary Material for the sources of these datasets. We evaluate the relative error as

$$D_{\Phi}(\mathbf{X}, \text{A1GM}(\mathbf{X})) / D_{\Phi}(\mathbf{X}, \text{WNMF}(\mathbf{X})),$$

where $\text{WNMF}(\mathbf{X})$ and $\text{A1GM}(\mathbf{X})$ are the rank-1 reconstructed matrices by KL-WNMF and A1GM, respectively, and the binary weight matrix Φ indicates locations of missing values of \mathbf{X} . We also compared the relative running time of A1GM to KL-WNMF.

The results are summarized in Table 1. In the table, the column **increase rate** means the ratio of the number of missing values after addition in A1GM to the original number of missing values. If **increase rate** is 1, it means that the location of missing values of the dataset is originally grid-like. For such datasets,

¹<https://www.kaggle.com/datasets>

²<https://archive.ics.uci.edu/ml/>

it is theoretically guaranteed that our method A1GM always provides the best rank-1 missing NMF, which minimizes the KL divergence in Equation (2). It is reasonable that the reconstructed matrix by KL-WNMF and that by A1GM are the same since the cost function (2) is convex. The number of iterations of the existing method required for convergence was between 2 and 4 for real datasets.

We can see that A1GM is much faster than KL-WNMF for all the datasets. Moreover, the relative error remains low even if missing values of datasets are not grid-like for most of datasets. In some real data, missing values are likely to be biased towards a particular row or column. As a result, they become grid-like by just adding small number of missing values. In these cases, our proposed method can conduct rank-1 missing NMF rapidly with competitive errors to KL-WNMF. By contrast, a large amount of information is lost after the increasing missing value step for some datasets (large **increase rates**). As a result, our method is not suitable for obtaining an accurately reconstructed rank-1 matrix, even though it is much faster than the existing method.

4 CONCLUSION

In this paper, we have derived the closed analytical formula of the best rank-1 NMMF. To obtain this formula, we have used the conservation law in m -projection in information geometry by modeling matrices as a log-linear model on a poset. Using the formula, we have developed a novel method of rank-1 NMF for missing data, called A1GM. We have shown that A1GM obtains the best rank-1 NMF when missing values are located in a grid-like manner. When the location of missing values is not grid-like, we increase the number of missing values so that they become grid-like, to which again we can use the closed formula of the best rank-1 NMMF. We empirically show that A1GM, which is not based on the gradient descent method, is more efficient than the existing gradient method for rank-1 missing NMF.

As noted, our method has two main limitations. First, because our modeling uses a log-linear model, we cannot handle zero values in a matrix. Second, the performance of A1GM is not expected to be convincing if there are a huge number of missing values. NMMF and NMF for missing data have been extended to tensors as NMFT (Takeuchi et al., 2013b) and WNFT (Ozerov et al., 2013), respectively. Generalization of our study to tensors is an interesting area for future work.

Acknowledgement

This work was supported by JSPS KAKENHI Grant Number JP20J23179 (KG), JP21H03503 (MS), and JST, PRESTO Grant Number JPMJPR1855, Japan (MS).

References

- Amari, S. (2008). Information geometry and its applications: Convex function and dually flat manifold. In *LIX Fall Colloquium on Emerging Trends in Visual Computing*, pages 75–102. Springer.
- Amari, S. (2016). *Information Geometry and Its Applications*. Springer.
- Biggs, M., Ghodsi, A., and Vavasis, S. (2008). Nonnegative matrix factorization via rank-one downdate. In *Proceedings of the 25th International Conference on Machine Learning*, pages 64–71.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3:993–1022.
- Blondel, V. D., Ho, N.-D., and Van Dooren, P. (2007). Weighted nonnegative matrix factorization and face feature extraction. *Preprint*.
- Chen, J. and Revels, J. (2016). Robust benchmarking in noisy environments. *arXiv:1608.04295*.
- Comon, P. (1994). Independent component analysis, a new concept? *Signal Processing*, 36(3):287–314.
- Ghalamkari, K. and Sugiyama, M. (2021). Fast Tucker rank reduction for non-negative tensors using mean-field approximation. In *Advances in Neural Information Processing Systems 34*.
- Gillis, N. and Glineur, F. (2011). Low-rank matrix approximation with weights or missing data is NP-hard. *SIAM Journal on Matrix Analysis and Applications*, 32(4):1149–1165.
- Gonzalez, E. F. (2007). *Efficient Alternating Gradient-type Algorithms for the Approximate Non-negative Matrix Factorization Problem*. PhD thesis, Rice University.
- Hajinezhad, D., Chang, T.-H., Wang, X., Shi, Q., and Hong, M. (2016). Nonnegative matrix factorization using admm: Algorithm and convergence analysis. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4742–4746. IEEE.
- Ho, N.-D. and Van Dooren, P. (2008). Non-negative matrix factorization with fixed row and column sums. *Linear Algebra and its Applications*, 429(5–6):1020–1025.

- Kim, J., He, Y., and Park, H. (2014). Algorithms for nonnegative matrix and tensor factorizations: A unified view based on block coordinate descent framework. *Journal of Global Optimization*, 58(2):285–319.
- Kim, Y.-D. and Choi, S. (2009). Weighted nonnegative matrix factorization. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1541–1544. IEEE.
- Kohjima, M., Matsubayashi, T., and Sawada, H. (2016). Non-negative multiple matrix factorization with Euclidean and Kullback-Leibler mixed divergences. In *2016 23rd International Conference on Pattern Recognition*, pages 2515–2520. IEEE.
- Lee, D. and Seung, H. S. (2001). Algorithms for non-negative matrix factorization. In Leen, T., Dietterich, T., and Tresp, V., editors, *Advances in Neural Information Processing Systems*, volume 13, pages 556–562. MIT Press.
- Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791.
- Liao, Q. and Zhang, Q. (2013). Efficient rank-one residue approximation method for graph regularized non-negative matrix factorization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 242–255. Springer.
- Ozerov, A., Duong, N., and Chevallier, L. (2013). Weighted nonnegative tensor factorization: on monotonicity of multiplicative update rules and application to user-guided audio source separation.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Rajapakse, M., Tan, J., and Rajapakse, J. (2004). Color channel encoding with NMF for face recognition. In *2004 International Conference on Image Processing, 2004. ICIP'04.*, volume 3, pages 2007–2010. IEEE.
- Song, D., Meyer, D. A., and Min, M. R. (2014). Fast nonnegative matrix factorization with rank-one admm. In *NIPS 2014 Workshop on Optimization for Machine Learning (OPT2014)*.
- Sugiyama, M., Nakahara, H., and Tsuda, K. (2017). Tensor balancing on statistical manifold. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3270–3279.
- Takács, G., Pilászy, I., Németh, B., and Tikk, D. (2008). Investigation of various matrix factorization methods for large recommender systems. In *2008 IEEE International Conference on Data Mining Workshops*, pages 553–562. IEEE.
- Takeuchi, K., Ishiguro, K., Kimura, A., and Sawada, H. (2013a). Non-negative multiple matrix factorization. In *Twenty-Third International Joint Conference on Artificial Intelligence*, pages 1713–1720.
- Takeuchi, K., Tomioka, R., Ishiguro, K., Kimura, A., and Sawada, H. (2013b). Non-negative multiple tensor factorization. In *2013 IEEE 13th International Conference on Data Mining*, pages 1199–1204. IEEE.
- Thompson, B. (1984). *Canonical correlation analysis: Uses and interpretation*. Number 47. Sage.
- Vavasis, S. A. (2010). On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377.
- Wold, S., Esbensen, K., and Geladi, P. (1987). Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 2(1–3):37–52.
- Xu, W., Liu, X., and Gong, Y. (2003). 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*, pages 267–273.
- Zhang, G., He, M., Wu, H., Cai, G., and Ge, J. (2016). Non-negative multiple matrix factorization with social similarity for recommender systems. In *Proceedings of the 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*, pages 280–286.
- Zhang, S., Wang, W., Ford, J., and Makedon, F. (2006). Learning from incomplete ratings using non-negative matrix factorization. In *Proceedings of the 2006 SIAM International Conference on Data Mining*, pages 549–553. SIAM.

A Proofs

Proposition 1 (simultaneous rank-1 θ -condition). *A triple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable if and only if its all two-body natural parameters are 0, that is, for any $(i, j) \in \Omega$,*

$$\theta_{ij} = 0 \quad \text{if } i \neq 1 \text{ and } j \neq 1.$$

Proof. It is clear from the definition that $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable if and only if we can represent the corresponding distribution p as $p(i, j) = s(i)t(j)$ for any index $(i, j) \in \Omega$.

First, we show that we can represent $p(i, j) = s(i)t(j)$ for any index $(i, j) \in \Omega$ if all two-body natural parameters are 0. If all two-body natural parameters are 0, we get a distribution corresponding to $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ as

$$p(i, j) = \exp(\theta_{11}) \exp\left(\sum_{i'=2}^i \theta_{i'1}\right) \exp\left(\sum_{j'=2}^j \theta_{1j'}\right).$$

The normalization condition $\sum_{ij} p(i, j) = 1$ leads to

$$\exp(\theta_{11}) = \frac{1}{\sum_{i=1}^{I+N} \exp\left(\sum_{i'=2}^i \theta_{i'1}\right)} \frac{1}{\sum_{j=1}^{J+M} \exp\left(\sum_{j'=2}^j \theta_{j'1}\right)}.$$

The distribution can then be written as

$$\begin{aligned} p(i, j) &= \frac{\exp\left(\sum_{i'=2}^i \theta_{i'1}\right)}{\sum_{i=1}^{I+N} \exp\left(\sum_{i'=2}^i \theta_{i'1}\right)} \frac{\exp\left(\sum_{j'=2}^j \theta_{1j'}\right)}{\sum_{j=1}^{J+M} \exp\left(\sum_{j'=2}^j \theta_{j'1}\right)} \\ &\equiv s(i)t(j). \end{aligned}$$

Thus, $p(i, j) = s(i)t(j)$ holds.

Next, we show that all two-body natural parameters are 0 if we can represent $p(i, j) = s(i)t(j)$ for any index $(i, j) \in \Omega$. We consider a contraposition of the statement; that is, there is a index $(i, j) \in \Omega$ that does not hold $p(i, j) = s(i)t(j)$ if a two-body natural parameter is not 0. When $p(i, j) = s(i)t(j)$, it holds that

$$\frac{p(i, j)}{p(i, j+1)} = \frac{p(i', j)}{p(i', j+1)}. \quad (5)$$

for any index $(i, j) \in \Omega$. However, for example, if $\theta_{22} \neq 0$, we can immediately confirm Equation (5) is not correct for $i = 1, j = 1$, and $i' = 2$. Therefore, the proposition has been proved. \square

Proposition 2 (simultaneous rank-1 η -condition). *A triple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable if and only if all of its two-body expectation parameters are factorizable as a product of two one-body parameters: for any $(i, j) \in \Omega$,*

$$\eta_{ij} = \eta_{i1}\eta_{1j}.$$

Proof. As we can see in Proposition 1, $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is simultaneously rank-1 decomposable if and only if the distribution becomes

$$p(i, j) = s(i)t(j)$$

Then,

$$\begin{aligned} \eta_{ij} &= \sum_{(i', j') \geq (i, j)} p(i', j') = \sum_{(i', j') \geq (i, j)} s(i')t(j') = \sum_{i' \geq i} s(i') \sum_{(i', j') \in \Omega} s(i')t(j') \sum_{j' \geq j} t(j') \\ &= \sum_{i' \geq i} s(i') \sum_{j' \geq 1} t(j') \sum_{i' \geq 1} s(i') \sum_{j' \geq j} t(j') \\ &= \eta_{i1}\eta_{1j}. \end{aligned}$$

We used the normalization condition $\sum_{ij} p(i, j) = 1$. \square

Theorem 1 (the best rank-1 NMMF). *For any given three positive matrices $\mathbf{X} \in \mathbb{R}_{>0}^{I \times J}$, $\mathbf{Y} \in \mathbb{R}_{>0}^{N \times J}$, and $\mathbf{Z} \in \mathbb{R}_{>0}^{I \times M}$ and two parameters $\alpha, \beta \geq 0$, four non-negative vectors $\mathbf{w} \in \mathbb{R}_{\geq 0}^I$, $\mathbf{h} \in \mathbb{R}_{\geq 0}^J$, $\mathbf{a} \in \mathbb{R}_{\geq 0}^N$, and $\mathbf{b} \in \mathbb{R}_{\geq 0}^M$ that minimize the cost function in Equation (1) is given as*

$$\begin{aligned} w_i &= \frac{\sqrt{S(\mathbf{X})}}{S(\mathbf{X}) + \beta S(\mathbf{Z})} \left\{ \sum_{j=1}^J \mathbf{X}_{ij} + \sum_{m=1}^M \beta \mathbf{Z}_{im} \right\}, \\ h_j &= \frac{\sqrt{S(\mathbf{X})}}{S(\mathbf{X}) + \alpha S(\mathbf{Y})} \left\{ \sum_{i=1}^I \mathbf{X}_{ij} + \sum_{n=1}^N \alpha \mathbf{Y}_{nj} \right\}, \\ a_n &= \frac{\sum_{j=1}^J \mathbf{Y}_{nj}}{\sqrt{S(\mathbf{X})}}, \quad b_m = \frac{\sum_{i=1}^I \mathbf{Z}_{im}}{\sqrt{S(\mathbf{X})}}. \end{aligned}$$

Proof. First, we show this theorem with $\alpha = \beta = 1$, followed by generalizing the result for any non-negative α and β . Hereinafter, we use overline for quantities on the simultaneous rank-1 subspace; for example, $(\overline{\mathbf{X}}, \overline{\mathbf{Y}}, \overline{\mathbf{Z}})$ as rank-1 matrices sharing factors, $\overline{\eta}$ as expectation parameters for distributions on simultaneous rank-1 subspace. We decompose input matrices \mathbf{X} , \mathbf{Y} and \mathbf{Z} as $\overline{\mathbf{X}} = \mathbf{w} \otimes \mathbf{h}$, $\overline{\mathbf{Y}} = \mathbf{a} \otimes \mathbf{h}$, and $\overline{\mathbf{Z}} = \mathbf{w} \otimes \mathbf{b}$, respectively, so that they minimize the cost function

$$D(\mathbf{X}, \mathbf{w} \otimes \mathbf{h}) + \alpha D(\mathbf{Y}, \mathbf{a} \otimes \mathbf{h}) + \beta D(\mathbf{Z}, \mathbf{w} \otimes \mathbf{b}). \quad (6)$$

To simplify, we define as follows:

$$\begin{aligned} \eta_{ij}^{\mathbf{X}} &= \eta_{ij} \quad \text{for } (i, j) \in \Omega_{\mathbf{X}}, \\ \eta_{i-I, j}^{\mathbf{Y}} &= \eta_{ij} \quad \text{for } (i, j) \in \Omega_{\mathbf{Y}}, \\ \eta_{i, j-J}^{\mathbf{Z}} &= \eta_{ij} \quad \text{for } (i, j) \in \Omega_{\mathbf{Z}}. \end{aligned}$$

According to the expectation conservation law in this m -projection, it holds that

$$\begin{aligned} \eta_{ij}^{\mathbf{X}} &= \overline{\eta}_{ij}^{\mathbf{X}} \quad \text{for } i = 1 \text{ or } j = 1, \\ \eta_{nj}^{\mathbf{Y}} &= \overline{\eta}_{nj}^{\mathbf{Y}} \quad \text{for } n = 1, \\ \eta_{im}^{\mathbf{Z}} &= \overline{\eta}_{im}^{\mathbf{Z}} \quad \text{for } m = 1, \end{aligned}$$

where $(\eta^{\mathbf{X}}, \eta^{\mathbf{Y}}, \eta^{\mathbf{Z}})$ is the expectation parameter of input, and $(\overline{\eta}^{\mathbf{X}}, \overline{\eta}^{\mathbf{Y}}, \overline{\eta}^{\mathbf{Z}})$ is the expectation parameter after the m -projection. By the definition of expectation parameters and the conservation law, we obtain

$$\begin{cases} \eta_{1j}^{\mathbf{X}} - \eta_{1,j+1}^{\mathbf{X}} = (S(\mathbf{w}) + S(\mathbf{a})) h_j, \\ \eta_{i1}^{\mathbf{X}} - \eta_{i+1,1}^{\mathbf{X}} = w_i (S(\mathbf{h}) + S(\mathbf{b})), \\ \eta_{1j}^{\mathbf{Y}} - \eta_{1,j+1}^{\mathbf{Y}} = S(\mathbf{w}) b_j, \\ \eta_{i1}^{\mathbf{Z}} - \eta_{i+1,1}^{\mathbf{Z}} = a_i S(\mathbf{h}). \end{cases}$$

We multiply these equations together and simplify them, resulting in

$$\begin{aligned} \overline{\mathbf{X}}_{ij} &= h_i w_j = \frac{(\eta_{i1}^{\mathbf{X}} - \eta_{i+1,1}^{\mathbf{X}}) (\eta_{1j}^{\mathbf{X}} - \eta_{1,j+1}^{\mathbf{X}})}{(S(\mathbf{w}) + S(\mathbf{a})) (S(\mathbf{h}) + S(\mathbf{b}))}, \\ \overline{\mathbf{Y}}_{ij} &= a_i h_j = \frac{(\eta_{i1}^{\mathbf{Y}} - \eta_{i+1,1}^{\mathbf{Y}}) (\eta_{1j}^{\mathbf{X}} - \eta_{1,j+1}^{\mathbf{X}})}{(S(\mathbf{w}) + S(\mathbf{a})) S(\mathbf{h})}, \\ \overline{\mathbf{Z}}_{ij} &= w_i b_j = \frac{(\eta_{i1}^{\mathbf{Z}} - \eta_{i+1,1}^{\mathbf{Z}}) (\eta_{1j}^{\mathbf{Z}} - \eta_{1,j+1}^{\mathbf{Z}})}{S(\mathbf{w}) (S(\mathbf{h}) + S(\mathbf{b}))}. \end{aligned}$$

Since the sum of each matrix \mathbf{X} , \mathbf{Y} and \mathbf{Z} are represented by conserved quantities, $S(\mathbf{X}) = \eta_{11}^{\mathbf{X}} - \eta_{11}^{\mathbf{Y}} - \eta_{11}^{\mathbf{Z}}$, $S(\mathbf{Y}) = \eta_{11}^{\mathbf{Y}}$, and $S(\mathbf{Z}) = \eta_{11}^{\mathbf{Z}}$, the sum of each matrix also do not change in this projection, that is, $S(\mathbf{X}) =$

$S(\bar{\mathbf{X}})$, $S(\mathbf{Y}) = S(\bar{\mathbf{Y}})$, and $S(\mathbf{Z}) = S(\bar{\mathbf{Z}})$. As a result, it follows that

$$\begin{aligned} S(\mathbf{X}) &= S(\bar{\mathbf{X}}) = S(\mathbf{w} \otimes \mathbf{h}) = S(\mathbf{w})S(\mathbf{h}), \\ S(\mathbf{Y}) &= S(\bar{\mathbf{Y}}) = S(\mathbf{a} \otimes \mathbf{h}) = S(\mathbf{a})S(\mathbf{h}), \\ S(\mathbf{Z}) &= S(\bar{\mathbf{Z}}) = S(\mathbf{w} \otimes \mathbf{b}) = S(\mathbf{w})S(\mathbf{b}), \end{aligned}$$

and we obtain

$$\left\{ \begin{aligned} \bar{\mathbf{X}}_{ij} &= \frac{S(\mathbf{X})}{(S(\mathbf{X}) + S(\mathbf{Y}))(S(\mathbf{X}) + S(\mathbf{Z}))} \left(\sum_{j'} \mathbf{X}_{ij'} + \sum_{j'} \mathbf{Z}_{ij'} \right) \left(\sum_{i'} \mathbf{X}_{i'j} + \sum_{i'} \mathbf{Y}_{i'j} \right), \\ \bar{\mathbf{Y}}_{ij} &= \frac{1}{S(\mathbf{X}) + S(\mathbf{Y})} \left(\sum_{i'} \mathbf{X}_{i'j} + \sum_{i'} \mathbf{Y}_{i'j} \right) \left(\sum_{j'} \mathbf{Y}_{ij'} \right), \\ \bar{\mathbf{Z}}_{ij} &= \frac{1}{S(\mathbf{X}) + S(\mathbf{Z})} \left(\sum_{j'} \mathbf{X}_{ij'} + \sum_{j'} \mathbf{Z}_{ij'} \right) \left(\sum_{i'} \mathbf{Z}_{ij'} \right). \end{aligned} \right.$$

Using the general property of the KL divergence, $\lambda D(\mathbf{P}, \mathbf{Q}) = D(\lambda \mathbf{P}, \lambda \mathbf{Q})$ for any matrices \mathbf{P}, \mathbf{Q} and non-negative number λ , the above result with general α and β is obtained as

$$\left\{ \begin{aligned} \bar{\mathbf{X}}_{ij} &= \frac{S(\mathbf{X})}{(S(\mathbf{X}) + \alpha S(\mathbf{Y}))(S(\mathbf{X}) + \beta S(\mathbf{Z}))} \left(\sum_{j'} \mathbf{X}_{ij'} + \sum_{j'} \beta \mathbf{Z}_{ij'} \right) \left(\sum_{i'} \mathbf{X}_{i'j} + \sum_{i'} \alpha \mathbf{Y}_{i'j} \right), \\ \bar{\mathbf{Y}}_{ij} &= \frac{1}{\alpha (S(\mathbf{X}) + \alpha S(\mathbf{Y}))} \left(\sum_{i'} \mathbf{X}_{i'j} + \sum_{i'} \alpha \mathbf{Y}_{i'j} \right) \left(\sum_{j'} \alpha \mathbf{Y}_{ij'} \right), \\ \bar{\mathbf{Z}}_{ij} &= \frac{1}{\beta (S(\mathbf{X}) + \beta S(\mathbf{Z}))} \left(\sum_{j'} \mathbf{X}_{ij'} + \sum_{j'} \beta \mathbf{Z}_{ij'} \right) \left(\sum_{i'} \beta \mathbf{Z}_{ij'} \right). \end{aligned} \right.$$

Thus, the theorem was proved. \square

Proposition 3 (Homogeneity of rank-1 missing NMF). *Let $\text{NMF}_1(\Phi, \mathbf{X})$ be the best rank-1 matrix $\mathbf{w} \otimes \mathbf{h}$ which minimizes the cost function in Equation (2). For any permutation matrices \mathbf{G} and \mathbf{H} , it holds that*

$$\text{NMF}_1(\mathbf{G}\Phi\mathbf{H}, \mathbf{G}\mathbf{T}\mathbf{H}) = \mathbf{G}^T \text{NMF}_1(\Phi, \mathbf{T}) \mathbf{H}^T.$$

Proof. We assume $\mathbf{T} \in \mathbb{R}_{\geq 0}^{n \times m}$ and $\Phi \in \{0, 1\}^{n \times m}$. Let $\mathbf{G} \in \{0, 1\}^{n \times n}$ and $\mathbf{H} \in \{0, 1\}^{m \times m}$ be the permutation matrices corresponding to the mappings \mathcal{G} and \mathcal{H} , respectively. This means that, for a given matrix \mathbf{X} and its row and column permutation $\mathbf{X}' = \mathbf{G}\mathbf{X}\mathbf{H}$, we have $\mathbf{X}'_{g(i), h(j)} = \mathbf{X}_{i,j}$. We can also apply the permutation matrices \mathbf{G} and \mathbf{H} to vectors $\mathbf{w} \in \mathbb{R}_{\geq 0}^n$ and $\mathbf{h} \in \mathbb{R}_{\geq 0}^m$. For $\mathbf{w}' = \mathbf{G}\mathbf{w}$ and $\mathbf{h}' = \mathbf{H}\mathbf{h}$, it holds that $w'_{g(i)} = w_i$ and $h'_{h(j)} = h_j$. We define \mathbf{w}^* and \mathbf{h}^* as

$$\begin{aligned} \mathbf{w}^*, \mathbf{h}^* &\equiv \underset{\mathbf{w}, \mathbf{h}}{\text{argmin}} D_{\mathbf{G}\Phi\mathbf{H}}(\mathbf{G}\mathbf{X}\mathbf{H}, \mathbf{w} \otimes \mathbf{h}) \\ &= \underset{\mathbf{w}, \mathbf{h}}{\text{argmin}} \sum_{i=1}^n \sum_{j=1}^m (\Phi_{g(i)h(j)} a_i b_j - \Phi_{g(i)h(j)} \mathbf{X}_{g(i)h(j)} \log a_i b_j). \end{aligned}$$

Table 2: Dataset details.

dataset	# zeros	# negatives	URL
Autompg	0	0	https://www.kaggle.com/uciml/autompg-dataset
DailySunSpot	16994	3247	https://www.kaggle.com/abhinand05/daily-sun-spot-data-1818-to-2019
CaliforniaHousing	0	20640	https://www.kaggle.com/harrywang/housing?select=housing.csv
MTSLibrary	200932	0	https://www.kaggle.com/sharzh23/mts-library?select=interactions.csv
BigMartSaleForecas	526	0	https://www.kaggle.com/arashnic/big-mart-sale-forecast?select=train.csv
BoardGameGeekData	520624	21	https://www.kaggle.com/mandshaw/games-0918
CreditCardApproval	797	0	https://www.kaggle.com/redwue/credit-card-approval?select=train.csv
HumanResourceAnaly	27510	0	https://www.kaggle.com/cezarschroeder/human-resource-analytics-dataset
heartdisease	1149	0	https://archive.ics.uci.edu/ml/datasets/Heart+Disease
lungcancer	107	0	https://archive.ics.uci.edu/ml/datasets/Lung+Cancer
PerthHousePrice	0	33656	https://www.kaggle.com/syuzai/perth-house-prices
SleepData	0	0	https://www.kaggle.com/mathurinache/sleep-dataset
arrhythmia	67256	14250	https://archive.ics.uci.edu/ml/datasets/Arrhythmia
Bostonhousing	812	0	https://www.kaggle.com/altavish/boston-housing-dataset
LifeExpectancyData	3385	0	https://www.kaggle.com/kumarajarshi/life-expectancy-who
HCCSurvivalDataSet	2416	0	https://archive.ics.uci.edu/ml/datasets/HCC+Survival
wiki4HE	1801	0	https://archive.ics.uci.edu/ml/datasets/wiki4HE

We replace \mathbf{w} with $\mathbf{G}\mathbf{w}$ and \mathbf{h} with $\mathbf{H}\mathbf{h}$ and we get

$$\begin{aligned}
 \mathbf{G}\mathbf{w}^*, \mathbf{H}\mathbf{h}^* &= \operatorname{argmin}_{\mathbf{w}, \mathbf{h}} \sum_{i=1}^n \sum_{j=1}^m (\Phi_{g(i)h(j)} a_{g(i)} b_{h(j)} - \Phi_{g(i)h(j)} \mathbf{X}_{g(i)h(j)} \log a_{g(i)} b_{h(j)}) \\
 &= \operatorname{argmin}_{\mathbf{w}, \mathbf{h}} \sum_{g(i)=1}^n \sum_{h(j)=1}^m (\Phi_{g(i)h(j)} a_{g(i)} b_{h(j)} - \Phi_{g(i)h(j)} \mathbf{X}_{g(i)h(j)} \log a_{g(i)} b_{h(j)}) \\
 &= \operatorname{argmin}_{\mathbf{w}, \mathbf{h}} \sum_{i=1}^n \sum_{j=1}^m (\Phi_{ij} a_i b_j - \Phi_{ij} \mathbf{X}_{ij} \log a_i b_j) \\
 &= D_{\Phi}(\mathbf{X}, \mathbf{w} \otimes \mathbf{h}).
 \end{aligned}$$

Thus, it holds that $\mathbf{w}^* \otimes \mathbf{h}^* = \operatorname{NMF}_1(\mathbf{G}\Phi\mathbf{H}, \mathbf{G}\mathbf{X}\mathbf{H})$ and $\mathbf{G}(\mathbf{w}^* \otimes \mathbf{h}^*)\mathbf{H} = \operatorname{NMF}_1(\Phi, \mathbf{X})$. Therefore, we have

$$\operatorname{NMF}_1(\mathbf{G}\Phi\mathbf{H}, \mathbf{G}\mathbf{T}\mathbf{H}) = \mathbf{G}^T \operatorname{NMF}_1(\Phi, \mathbf{T}) \mathbf{H}^T.$$

We use the fact that permutation matrix is always orthogonal; that is, $\mathbf{G}^{-1} = \mathbf{G}^T$ and $\mathbf{H}^{-1} = \mathbf{H}^T$. \square

B Details of datasets

We provide the list of the source of the real datasets in Table 2.