Privacy-Preserving Distributed Nonnegative Matrix Factorization

Ehsan Lari¹, Reza Arablouei², Stefan Werner¹

¹Department of Electronic Systems, Norwegian University of Science and Technology, Trondheim, Norway ²CSIRO's Data61, Pullenvale QLD 4069, Australia

Abstract-Nonnegative matrix factorization (NMF) is an effective data representation tool with numerous applications in signal processing and machine learning. However, deploying NMF in a decentralized manner over ad-hoc networks introduces privacy concerns due to the conventional approach of sharing raw data among network agents. To address this, we propose a privacypreserving algorithm for fully-distributed NMF that decomposes a distributed large data matrix into left and right matrix factors while safeguarding each agent's local data privacy. It facilitates collaborative estimation of the left matrix factor among agents and enables them to estimate their respective right factors without exposing raw data. To ensure data privacy, we secure information exchanges between neighboring agents utilizing the Paillier cryptosystem, a probabilistic asymmetric algorithm for publickey cryptography that allows computations on encrypted data without decryption. Simulation results conducted on synthetic and real-world datasets demonstrate the effectiveness of the proposed algorithm in achieving privacy-preserving distributed NMF over ad-hoc networks.

I. INTRODUCTION

Nonnegative matrix factorization (NMF) [1]-[5] is a specific case of constrained low-rank matrix approximation [6] and a linear dimensionality reduction (LDR) technique aimed at representing nonnegative data more compactly through nonnegative factors. NMF, originally introduced as positive matrix factorization in [7], has gained significant research interest, particularly after being popularized by [8]. It has found widespread applications in various fields such as signal and image processing, data mining and analytics, machine learning, and federated learning. Examples include air emission control [7], visual object recognition [9], video backgroundforeground separation [10], spectral unmixing [11], text mining [12], blind source separation [13], clustering [14], collaborative filtering [15], computational biology [16], music analysis [17], molecular pattern discovery [3], efficient implementation of deep neural networks [18], and detecting malware activities [19]. Its popularity stems from its utility in identifying and extracting meaningful features from data in addition to serving as a powerful LDR technique.

Distributed optimization and estimation algorithms have garnered significant attention due to the ubiquity of data dispersed across multiple agents within network environments. The existing distributed NMF algorithms align with this trend, addressing scenarios where data is distributed among a network of agents. However, the conventional approach of sharing raw data among neighboring agents poses inherent security risks and compromises the privacy of sensitive information [20]–[22]. Consequently, there is a pressing need for privacypreserving distributed NMF algorithms that ensure the security and confidentiality of each agent's local data.

The Paillier cryptosystem [23] is a fundamental tool for enhancing privacy in distributed algorithms. As a probabilistic asymmetric algorithm for public-key cryptography, it is specifically designed to provide secure homomorphic encryption. Its primary advantage lies in its homomorphic properties, which enable computations to be performed on encrypted data without decryption. The Paillier cryptosystem has proven to be effective in improving privacy in various applications, including smart grids [24]–[26], machine learning [27], smart homes [28], and federated learning [29], [30]. However, the potential advantages of its utilization in addressing the distributed NMF problem have not been explored in the literature.

In this paper, we introduce a privacy-preserving distributed NMF (PPDNMF) algorithm tailored for scenarios where the data matrix to be factorized is distributed among agents within an ad-hoc network. Each agent holds a subset of the columns of the data matrix. Our goal is to perform NMF of the entire data dispersed over the network in a fully distributed and secure manner. Specifically, agents participate in a distributed and collaborative process to estimate both the left and right factors, exchanging information exclusively with their immediate neighbors over secure communication links. While taking part in this collaborative process, agents maintain the privacy of their local data and the corresponding right factor estimates. We utilize the block coordinate-descent (BCD) algorithm and the alternating direction method of multipliers (ADMM) to develop our distributed NMF algorithm. Furthermore, to ensure privacy preservation, we integrate the Paillier cryptosystem into our algorithm. We evaluate the performance of the proposed algorithm through simulations using synthetic and real data. demonstrating its efficacy in achieving results comparable to those obtained by the centralized alternative.

II. DISTRIBUTED NMF

The objective of NMF is to approximate a data matrix $\mathbf{Z} \in \mathbb{R}^{L \times M}$ consisting of nonnegative entries using the product of left and right factor matrices, both with nonnegative entries. That is, $\mathbf{Z} = \mathbf{X}\mathbf{Y}$ where $\mathbf{X} \in \mathbb{R}^{L \times K}$ and $\mathbf{Y} \in \mathbb{R}^{K \times M}$, typically with $K < \min(L, M)$. This approximation represents the

This work was partially supported by the Research Council of Norway and the Research Council of Finland (Grant 354523).

L-dimensional datapoints (columns of the data matrix) within a *K*-dimensional linear subspace spanned by the columns of the left factor, whose coordinates are given by the columns of the right factor. The nonnegativity constraint on the factors induces sparsity, further enhancing the compactness of the representation. Moreover, in many applications, the factors' nonnegativity is essential to their physical plausibility and intuitive interpretability.

We utilize the least-squares criterion, which is appropriate when the perturbation in the data matrix \mathbf{Z} can be modeled as a Gaussian process. Therefore, the NMF problem can be formulated as

$$\begin{split} \min_{\mathbf{X},\mathbf{Y}} & \frac{1}{2} \|\mathbf{Z} - \mathbf{X}\mathbf{Y}\|_{\mathsf{F}}^2 \\ \text{s. t. } & \mathbf{X} \ge 0, \mathbf{Y} \ge 0, \end{split} \tag{1}$$

where $\|\cdot\|_{\mathsf{F}}$ denotes the Frobenius norm. We consider the scenario where \mathbf{Z} is distributed over a network with N agents such that we have $\mathbf{Z} = [\mathbf{Z}_1, \dots, \mathbf{Z}_N]$ and consequently $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_N]$ with $\mathbf{Y}_i \in \mathbb{R}^{K \times M_i}$ and $\sum_{i=1}^N M_i = M$. Therefore, we rewrite (1) as

$$\min_{\mathbf{X}, \{\mathbf{Y}_i\}} \frac{1}{2} \sum_{i=1}^{N} \|\mathbf{Z}_i - \mathbf{X}\mathbf{Y}_i\|_{\mathsf{F}}^2 \tag{2}$$
s. t. $\mathbf{X} \ge 0, \mathbf{Y}_i \ge 0.$

In a fully distributed approach, every agent, indexed by i, aims to estimate **X** and its own \mathbf{Y}_i using its local data \mathbf{Z}_i and by exchanging information solely with its immediate neighbors through single-hop communication. To this end, we utilize the BCD algorithm and iteratively solve two optimization subproblems for **X** and **Y**. That is, we repeat the following alternating minimizations until convergence is achieved:

$$\mathbf{X}^{(n)} = \min_{\mathbf{X}} \frac{1}{2} \sum_{i=1}^{N} \left\| \mathbf{Z}_{i} - \mathbf{X} \mathbf{Y}_{i}^{(n-1)} \right\|_{\mathsf{F}}^{2}$$
(3)

$$\mathbf{Y}_{i}^{(n)} = \min_{\mathbf{Y}_{i}} \frac{1}{2} \left\| \mathbf{Z}_{i} - \mathbf{X}^{(n)} \mathbf{Y}_{i} \right\|_{\mathsf{F}}^{2}, \forall i \in \{1, \dots, N\} \quad (4)$$

s.t. $\mathbf{Y}_{i} \ge 0.$

The superscript (n) denotes the estimate of its respective parameter at the *n*th BCD iteration.

The solution of (4) can be localized straightforwardly, provided that each agent has access to the estimate $\mathbf{X}^{(n)}$. To solve (3) in a fully distributed manner, we introduce the variable \mathbf{X}_i at each agent *i* as a local copy of \mathbf{X} and enforce it to be equal to those of the agents within the immediate neighborhood of agent *i*, thereby achieving consensus across the network. Thus, we reformulate (3) into the following equivalent form

$$\mathbf{X}_{i}^{(n)} = \min_{\mathbf{X}_{i}} \frac{1}{2} \left\| \mathbf{Z}_{i} - \mathbf{X}_{i} \mathbf{Y}_{i}^{(n-1)} \right\|_{\mathsf{F}}^{2} + \imath(\mathbf{X}_{i})$$
(5)
s.t. $\mathbf{X}_{i} = \mathbf{X}_{j} \quad \forall j \in \mathcal{N}_{i}, \ \forall i \in \{1, \dots, N\}$

where $i(\cdot)$ denotes the indicator function accounting for the nonegativity constraint and \mathcal{N}_i denotes the set of neighbors

of agent *i* with cardinality $d_i = |\mathcal{N}_i|$. Subsequently, we decompose and decouple the optimization problems at the agents by introducing the auxiliary variables $\mathbf{U}_i, \mathbf{S}_{i,j} \in \mathbb{R}^{L \times K}$ and rewriting the optimization in (5) as

$$\min_{\mathbf{X}_{i},\mathbf{U}_{i},\mathbf{S}_{i,j}} \frac{1}{2} \left\| \mathbf{Z}_{i} - \mathbf{U}_{i} \mathbf{Y}_{i}^{(n-1)} \right\|_{\mathsf{F}}^{2} + \imath(\mathbf{X}_{i}) \quad (6)$$

$$\mathbf{U}_{i} = \mathbf{X}_{i}$$
s. t.
$$\mathbf{S}_{i,j} = \mathbf{U}_{i} \quad \forall j \in \mathcal{N}_{i}, \; \forall i \in \{1, \dots, N\}$$

$$\mathbf{S}_{j,i} = \mathbf{S}_{i,j}.$$

We can express the corresponding aggregate augmented Lagrangian function as

$$\mathcal{L} (\{\mathbf{X}_{i}\}, \{\mathbf{U}_{i}\}, \{\mathbf{S}_{i,j}\}, \{\mathbf{P}_{i}\}, \{\mathbf{Q}_{i,j}\}) = \frac{1}{2} \sum_{i=1}^{N} \left\| \mathbf{Z}_{i} - \mathbf{U}_{i} \mathbf{Y}_{i}^{(n-1)} \right\|_{\mathsf{F}}^{2} + \sum_{i=1}^{N} \imath(\mathbf{X}_{i}) + \frac{\mu}{2} \sum_{i=1}^{N} \left\| \mathbf{X}_{i} - \mathbf{U}_{i} - \mathbf{P}_{i} \right\|_{\mathsf{F}}^{2} + \frac{1}{2} \sum_{i=1}^{N} \sum_{j \in \mathcal{N}_{i}} \rho_{i,j} \left\| \mathbf{U}_{i} - \mathbf{S}_{i,j} - \mathbf{Q}_{i,j} \right\|_{\mathsf{F}}^{2},$$
(7)

where μ and $\rho_{i,j}$ are penalty parameters and $\mathbf{P}_i, \mathbf{Q}_{i,j} \in \mathbb{R}^{L \times K}$ are scaled Lagrange multipliers. We maintain μ consistent across all agents and iterations. However, we allow each $\rho_{i,j}$, unique to the edge connecting agents *i* and *j*, to vary over iterations [31]. Additionally, we consider $\rho_{i,j} = \rho_{j,i} \forall i, j$.

A. Estimating the Left Factor

Minimizing (7) using ADMM, which leads to the elimination of the auxiliary variables $\{S_{i,j}\}$, yields the following iterations at each agent *i* [32]:

$$\begin{aligned} \mathbf{X}_{i}^{(n,m)} &= \Pi_{\geq 0} \left\{ \mathbf{U}_{i}^{(m-1)} + \mathbf{P}_{i}^{(m-1)} \right\} \end{aligned} \tag{8} \\ \mathbf{U}_{i}^{(m)} &= \left[\mathbf{Z}_{i} \mathbf{Y}_{i}^{(n-1)\mathsf{T}} + \mu \left(\mathbf{X}_{i}^{(n,m)} - \mathbf{P}_{i}^{(m-1)} \right) \right. \\ &+ \rho_{i}^{(m)} \left(\mathbf{U}_{i}^{(m-1)} + 2\mathbf{Q}_{i}^{(m-1)} - \mathbf{Q}_{i}^{(m-2)} \right) \right] \\ &\times \left[\mathbf{Y}_{i}^{(n-1)} \mathbf{Y}_{i}^{(n-1)\mathsf{T}} + \left(\mu + \rho_{i}^{(m)} \right) \mathbf{I} \right]^{-1} \end{aligned} \tag{9}$$

$$\mathbf{P}_{i}^{(m)} = \mathbf{P}_{i}^{(m-1)} - \left(\mathbf{X}_{i}^{(n,m)} - \mathbf{U}_{i}^{(m)}\right)$$
(10)

$$\mathbf{Q}_{i}^{(m)} = \mathbf{Q}_{i}^{(m-1)} + \sum_{j \in \mathcal{N}_{i}} \rho_{i,j}^{(m)} \left(\mathbf{U}_{j}^{(m)} - \mathbf{U}_{i}^{(m)} \right).$$
(11)

Here, (m) denotes the ADMM iteration index, $\Pi_{\geq 0}$ represents the projection onto the nonnegative orthant, and $(\cdot)^{\mathsf{T}}$ stands for matrix transpose. In addition, we define $\rho_i^{(m)} = \sum_{j \in \mathcal{N}_i} \rho_{i,j}^{(m)}$ and $\mathbf{Q}_i^{(m)} = \sum_{j \in \mathcal{N}_i} \rho_{i,j}^{(m)} \mathbf{Q}_{i,j}^{(m)}$. These ADMM iterations can be executed in a fully distributed manner, relying solely on locally available information and single-hop communications. Upon convergences of the algorithm, we utilize the latest estimates $\mathbf{X}_i^{(n,m)}$ for optimizing \mathbf{Y}_i in the subsequent BCD iteration, i.e., $\mathbf{X}_i^{(n)} \leftarrow \mathbf{X}_i^{(n,m)}$. Note that we enforce the nonnegativity constraint and consensus simultaneously.

B. Estimating the Right Factor

Similarly, we can employ ADMM to iteratively solve (4) as follows:

$$\mathbf{Y}_{i}^{(n,k)} = \Pi_{\geq 0} \left\{ \mathbf{V}_{i}^{(k-1)} + \mathbf{R}_{i}^{(k-1)} \right\}$$
(12)

$$\mathbf{V}_{i}^{(k)} = \left(\mathbf{X}_{i}^{(n)\mathsf{T}}\mathbf{X}_{i}^{(n)} + \eta\mathbf{I}\right)^{-1}$$

$$\times \left[\mathbf{X}_{i}^{(n)\mathsf{T}} \mathbf{Z}_{i} + \eta \left(\mathbf{Y}_{i}^{(n,k)} - \mathbf{R}_{i}^{(k-1)} \right) \right]$$
(13)

$$\mathbf{R}_{i}^{(k)} = \mathbf{R}_{i}^{(k-1)} - \left(\mathbf{Y}_{i}^{(n,k)} - \mathbf{V}_{i}^{(k)}\right).$$
(14)

Here, (k) represents the ADMM iteration index and η is the penalty parameter. Once convergence is attained, we utilize the latest estimates $\mathbf{Y}_i^{(n,k)}$ to update \mathbf{X}_i estimates in the subsequent BCD iteration, i.e., $\mathbf{Y}_i^{(n)} \leftarrow \mathbf{Y}_i^{(n,k)}$.

Note that, we employ warm start in both ADMM algorithms for estimating the left and right factors. At the onset of each BCD iteration, we initialize both ADMM inner iterations using the most recent estimates from the preceding iterations.

C. Synchronization and Stopping

Our algorithm does not require waiting for all agents to converge in either BCD or ADMM iterations. During X_i updates, the first agent to converge or reach a predetermined maximum number of ADMM iterations stops updating and raises a flag, signaling its neighbors to stop as well. This message is then propagated through the network until all agents stop updating. After completing the collaborative X_i update iterations, each agent can immediately start updating its Y_i estimate independently. When an agent stops \mathbf{Y}_i update due to convergence or reaching the corresponding iteration limit, it begins the first iteration of \mathbf{X}_i update and shares its \mathbf{U}_i with its neighbors. Using warm start, in the first iteration of \mathbf{X}_i update, each agent utilizes the neighbor \mathbf{U}_i values from the previous BCD iteration. Afterwards, if an agent has not received all required U_i updates from its neighbors, subsequent iterations are postponed until they are all acquired, ensuring synchronization of X_i updates among all agents. To decide when to terminate the BCD iterations, one can employ the same strategy as described for X_i updates. That is, the initial agent to discern convergence or reach a predefined maximum number of BCD iterations halts its updates and notifies its neighbors. This notification propagates across the network until all agents cease updating.

D. Convergence Analysis

We can express the optimization problem (6) as

$$\min_{\boldsymbol{\mathcal{X}},\boldsymbol{\mathcal{U}}} f(\boldsymbol{\mathcal{X}}) + g(\boldsymbol{\mathcal{U}})$$
(15)

s. t.
$$\mathbf{A}\boldsymbol{\mathcal{X}} + \mathbf{B}\boldsymbol{\mathcal{U}} = \mathbf{0},$$
 (16)

where

$$f(\boldsymbol{\mathcal{X}}) = \sum_{i=1}^{N} i(\mathbf{X}_{i}), \ g(\boldsymbol{\mathcal{U}}) = \frac{1}{2} \sum_{i=1}^{N} \left\| \mathbf{Z}_{i} - \mathbf{U}_{i} \mathbf{Y}_{i}^{(n-1)} \right\|_{\mathsf{F}}^{2},$$
$$\boldsymbol{\mathcal{X}} = \begin{bmatrix} \mathbf{X}_{1}, \cdots, \mathbf{X}_{N} \end{bmatrix}^{\mathsf{T}}, \ \boldsymbol{\mathcal{U}} = \begin{bmatrix} \mathbf{U}_{1}, \cdots, \mathbf{U}_{N} \left| \mathbf{S}_{1,1}, \cdots, \mathbf{S}_{N,N} \right]^{\mathsf{T}},$$

Algorithm 1: The DDNME algorithm as agent

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \ \mathbf{B} = \begin{bmatrix} -\mathbf{I} & \mathbf{0} \\ \hline \mathcal{D} & -\mathcal{A} \\ \hline \mathbf{0} & \mathcal{I} \end{bmatrix}.$$

and $\mathcal{D} = \text{bdiag}(d_1 \mathbf{I}_{L \times K}, \dots, d_N \mathbf{I}_{L \times K})$. In addition, $\mathcal{A} \in \mathbb{R}^{NL \times N^2K}$ and $\mathcal{I} \in \mathbb{R}^{NL \times N^2K}$ represent modified versions of the adjacency and edge-node incidence matrices of the network, respectively [33]. Consequently, the convergence of the ADMM iterations (8)-(11) can be proven using the approach proposed in [34]–[36]. The convergence of (12)-(14) can also be verified in a similar manner.

III. PRIVACY-PRESERVING DISTRIBUTED NMF

In this section, we provide a brief overview of the Paillier cryptosystem, which we employ to enhance the privacy of the distributed NMF algorithm developed in section II. Subsequently, we introduce our proposed privacy-preserving distributed NMF (PPDNMF) algorithm.

A. Paillier Cryptosystem

In the Paillier cryptosystem, a public key and a private key are utilized. The public key is broadcast publicly, allowing other users to encrypt messages. However, decrypting the messages is only possible with the private key, which remains unknown to other users. This cryptosystem exhibits additive homomorphism [31], [37], [38], i.e., $\mathcal{E}(m_3(m_1 + m_2)) = (\mathcal{E}(m_1)\mathcal{E}(m_2))^{m_3}$.



Fig. 1. The normalized mean-square error (NMSE) values of PPDNMF and centralized algorithm versus the BCD iteration index for different values of N_{max} on (a) synthetic data and (b) MIT-CBCL database. (c) The original and reconstructed faces #1 and #2429 from the MIT-CBCL database.

B. Privacy Preservation

The only update equation among (8)-(14) that relies on information received from neighboring agents is (11). To protect the privacy of agents at this step, we adopt a similar approach to [31] and enable the agents to encrypt all messages communicated with their neighbors. To this end, we decompose each edge-specific penalty parameter as $\rho_{i,j}^{(m)} = g_{i\to j}^{(m)}g_{j\to i}^{(m)}$ where $g_{i\to j}^{(m)}$ and $g_{j\to i}^{(m)}$ are exclusively known to agents *i* and *j*, respectively. In addition, we implement a secure data exchange procedure as outlined in lines 10-15 of Algorithm 1, which provides a summary of the proposed PPDNMF algorithm. Consequently, note the following:

- Data exchanged between agents *i* and *j* is encrypted, rendering it inaccessible to other agents or eavesdroppers, even if intercepted.
- The parameter $g_{i \rightarrow j}^{(m)}$ is unique to each edge and iteration. Therefore, an agent cannot infer the private information $\mathbf{U}_{j}^{(m)}$ of any of its neighbors by decrypting the messages it receives from them, as each neighbor j uses its unique $g_{j \rightarrow i}^{(m)}$ in its encrypted message to agent i.
- The Paillier cryptosystem is intended for encrypting scalar unsigned integers. To encrypt the entries of $\mathbf{U}_i^{(m)}$, which are typically floating-point values, we initially quantize them. This involves multiplying each entry by a positive integer N_{\max} , which determines the quantization resolution, and then rounding the result to the nearest integer. To undo the quantization, we divide the decrypted values by N_{\max} .
- To guarantee convergence, we ensure that the parameters $g_{i \rightarrow j}^{(m)}$ increase monotonically over iterations without becoming unbounded [31]. Thus, we select each parameter uniformly from the interval $\left(g_{i \rightarrow j}^{(m-1)}, g_i\right)$, where g_i is a predefined positive constant, known only to agent *i*, and $g_{i \rightarrow j}^{(0)} = 0$.

IV. SIMULATION RESULTS

In this section, we conduct a series of numerical experiments to evaluate the performance of our PPDNMF algorithm. We consider a network consisting of N = 10 agents, interconnected arbitrarily, with each agent having three neighbors on average. We test our algorithm on two datasets, namely, a synthetic dataset and the MIT-CBCL face database [39].

The agents collaboratively factorize a data matrix $\mathbf{Z} \in \mathbb{R}^{L \times M}$ to left and right factor matrices $\mathbf{X} \in \mathbb{R}^{L \times K}$ and $\mathbf{Y} \in \mathbb{R}^{K \times M}$, where $L \in \{30, 361\}, M \in \{200, 2429\}$, and $K \in \{5, 49\}$ in our two experiments. We set the number of BCD iterations to 100 and the number of ADMM iterations to 30. In our implementation of the Paillier cryptosystem, we use 128-bit public and private keys. To handle the encryption of negative quantized values (note line 10 in Algorithm 1), we convert them to positive integers by adding the public key to them [37].

We evaluate the performance of PPDNMF in comparison with the centralized algorithm, i.e., where all data is available at a central hub. To quantify the performance, we utilize the normalized mean-square error (NMSE) at each BCD iteration, defined as $\frac{1}{N}\sum_{i=1}^{N} \|\mathbf{Z}_{i} - \mathbf{X}_{i}^{(n)}\mathbf{Y}_{i}^{(n)}\|_{\mathsf{F}} / \|\mathbf{Z}_{i}\|_{\mathsf{F}}$. In addition, we average the presented results over 100 independent trials.

In our first experiment utilizing synthetic data, we draw the entries of the nonnegative factor matrices $\mathbf{X} \in \mathbb{R}^{30 \times 5}$ and $\mathbf{Y} \in \mathbb{R}^{5 imes 200}$ independently from exponential distributions with parameter values 0.033 and 0.8, respectively. We calculate the data matrix as $\mathbf{Z} = \mathbf{X}\mathbf{Y} + \mathbf{\Gamma}$, where we draw the entries of Γ independently from a Gaussian distribution with zero mean and variance 3.6×10^{-4} , resulting in an SNR of approximately 20dB. We set $\mu = 0.1$, $\eta = 1$, and $q_i = 0.033$ for all agents. We consider \mathbf{Z} to be distributed among the agents such that each agent has a varying number of columns between four and 40. We present the NMSE learning curves of PPDNMF for different values of $N_{\rm max}$ alongside that of the corresponding centralized algorithm in Fig. 1(a). We observe from Fig. 1(a) that the proposed PPDNMF algorithm closely approximates the performance of the centralized algorithm in terms of both convergence rate and steady-state NMSE. Additionally, it is also evident that a higher value of $N_{\rm max}$ leads to a lower steadystate NMSE.

Our second experiment involves the MIT-CBCL face

database, which comprises 2429 monochromic face images in its training set. We distribute the associated data matrix among the agents such that each agent has between 224 and 245 columns. For this experiment, we set $\mu = 2$, $\eta = 2$, and $g_i = 0.05$ for all agents. The results presented in Fig. 1(b) underscore the effectiveness of PPDNMF. Notably, PPDNMF exhibits robust performance even with $N_{\text{max}} = 10$. Furthermore, we compare the original faces #1 and #2429 and their reconstructed versions by PPDNMF using $N_{\text{max}} = 10^6$ in Fig. 1(c). The reconstructed faces closely resemble their original counterparts.

V. CONCLUSION

We introduced a novel privacy-preserving distributed nonnegative matrix factorization algorithm that employs the Paillier cryptosystem to enable secure collaboration among agents, thereby safeguarding their privacy and mitigating the risk of sensitive data leakage over ad-hoc networks. Our simulation results, based on both synthetic and real data, confirmed the efficacy of the proposed algorithm. In future work, we plan to conduct a comprehensive theoretical privacy analysis of the proposed algorithm, exploring its resilience across various attack scenarios.

REFERENCES

- N. Gillis, "The why and how of nonnegative matrix factorization," Connections, vol. 12, no. 2, 2014.
- [2] A. Cichocki, R. Zdunek, A. H. Phan, and S.-i. Amari, Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation. John Wiley & Sons, 2009.
- [3] M. W. Berry, M. Browne, A. N. Langville, V. P. Pauca, and R. J. Plemmons, "Algorithms and applications for approximate nonnegative matrix factorization," *Comput. Stat. Data Anal.*, vol. 52, no. 1, pp. 155– 173, 2007.
- [4] Y.-X. Wang and Y.-J. Zhang, "Nonnegative matrix factorization: A comprehensive review," *IEEE Trans. Knowl. Data Eng.*, vol. 25, no. 6, pp. 1336–1353, 2013.
- [5] D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Proc. Adv. Neural Inf. Process. Syst.*, T. Leen, T. Dietterich, and V. Tresp, Eds., vol. 13. MIT Press, 2000.
- [6] M. Udell, C. Horn, R. Zadeh, S. Boyd et al., "Generalized low rank models," Found. Trends Mach. Learn., vol. 9, no. 1, pp. 1–118, 2016.
- [7] P. Paatero and U. Tapper, "Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values," *Environmetrics*, vol. 5, no. 2, pp. 111–126, 1994.
- [8] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [9] M. W. Spratling and P. Dayan, "Learning image components for object recognition." J. Mach. Learn. Res., vol. 7, no. 5, 2006.
- [10] A. Kumar and V. Sindhwani, "Near-separable non-negative matrix factorization with ℓ₁ and Bregman loss functions," in *Proc. SIAM Int. Conf. Data Min.*, 2015, pp. 343–351.
- [11] W.-K. Ma, J. M. Bioucas-Dias, T.-H. Chan, N. Gillis, P. Gader, A. J. Plaza, A. Ambikapathi, and C.-Y. Chi, "A signal processing perspective on hyperspectral unmixing: Insights from remote sensing," *IEEE Signal Process. Mag.*, vol. 31, no. 1, pp. 67–81, 2013.
- [12] D. Godfrey, C. Johns, C. Meyer, S. Race, and C. Sadek, "A case study in text mining: Interpreting twitter data from world cup tweets," *arXiv* preprint arXiv:1408.5427, 2014.
- [13] T.-H. Chan, W.-K. Ma, C.-Y. Chi, and Y. Wang, "A convex analysis framework for blind separation of non-negative sources," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 5120–5134, 2008.
- [14] A. C. Türkmen, "A review of nonnegative matrix factorization methods for clustering," arXiv preprint arXiv:1507.03194, 2015.
- [15] P. Melville and V. Sindhwani, "Recommender systems." Encycl. Mach. Learn., vol. 1, pp. 829–838, 2010.

- [16] K. Devarajan, "Nonnegative matrix factorization: an analytical and interpretive tool in computational biology," *PLoS Comput. Biol.*, vol. 4, no. 7, p. e1000029, 2008.
- [17] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the itakura-saito divergence: With application to music analysis," *Neural Comput.*, vol. 21, no. 3, pp. 793–830, 2009.
- [18] S. Bhattacharya and N. D. Lane, "Sparsification and separation of deep learning layers for constrained resource inference on wearables," in *Proc. ACM Conf. Embedded Netw. Sensor Syst.*, 2016, pp. 176–189.
- [19] Y.-W. Chang, H.-Y. Chen, C. Han, T. Morikawa, T. Takahashi, and T.-N. Lin, "FINISH: Efficient and scalable NMF-based federated learning for detecting malware activities," *IEEE Trans. Emerg. Top. Comput.*, vol. 11, no. 4, pp. 934–949, 2023.
- [20] Y. Qian, C. Tan, D. Ding, H. Li, and N. Mamoulis, "Fast and secure distributed nonnegative matrix factorization," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 2, pp. 653–666, 2022.
- [21] P. Mai and Y. Pang, "Privacy-preserving multiview matrix factorization for recommender systems," *IEEE Trans. Artif. Intell.*, vol. 5, no. 1, pp. 267–277, 2024.
- [22] N. K. D. Venkategowda and S. Werner, "Privacy-preserving distributed maximum consensus," *IEEE Signal Process. Lett.*, vol. 27, pp. 1839– 1843, 2020.
- [23] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proc. Int. Conf. Theory Appl. Cryptogr. Techn.* Springer, 1999, pp. 223–238.
- [24] Y. Yan, Z. Chen, V. Varadharajan, M. J. Hossain, and G. E. Town, "Distributed consensus-based economic dispatch in power grids using the paillier cryptosystem," *IEEE Trans. Smart Grid*, vol. 12, no. 4, pp. 3493–3502, 2021.
- [25] R. Lu, X. Liang, X. Li, X. Lin, and X. Shen, "Eppa: An efficient and privacy-preserving aggregation scheme for secure smart grid communications," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, pp. 1621–1631, 2012.
- [26] H. Shen, M. Zhang, and J. Shen, "Efficient privacy-preserving cube-data aggregation scheme for smart grids," *IEEE Trans. Inf. Forensics Secur.*, vol. 12, no. 6, pp. 1369–1381, 2017.
- [27] M. Shen, X. Tang, L. Zhu, X. Du, and M. Guizani, "Privacy-preserving support vector machine training over blockchain-based encrypted IoT data in smart cities," *IEEE Internet Things J.*, vol. 6, pp. 7702–7712, 2019.
- [28] S. M. Errapotu, J. Wang, Y. Gong, J.-H. Cho, M. Pan, and Z. Han, "Safe: Secure appliance scheduling for flexible and efficient energy consumption for smart home IoT," *IEEE Internet Things J.*, vol. 5, no. 6, pp. 4380– 4391, 2018.
- [29] B. Li, Y. Wu, J. Song, R. Lu, T. Li, and L. Zhao, "DeepFed: Federated deep learning for intrusion detection in industrial cyber-physical systems," *IEEE Trans. Ind. Inform.*, vol. 17, no. 8, pp. 5615–5624, 2021.
- [30] Q. Xu, Y. Lan, Z. Su, D. Fang, and H. Zhang, "Verifiable and privacypreserving cooperative federated learning in uav-assisted vehicular networks," in *Proc. IEEE Int. Conf. Commun.*, 2023, pp. 2288–2293.
- [31] C. Zhang, M. Ahmad, and Y. Wang, "ADMM based privacy-preserving decentralized optimization," *IEEE Trans. Inf. Forensics Secur.*, vol. 14, no. 3, pp. 565–580, 2019.
- [32] G. B. Giannakis, Q. Ling, G. Mateos, I. D. Schizas, and H. Zhu, Decentralized learning for wireless communications and networking. Springer, 2017.
- [33] E. Wei and A. Ozdaglar, "Distributed alternating direction method of multipliers," in *Proc. IEEE Conf. Decis. Control*, 2012, pp. 5445–5450.
- [34] W. Deng and W. Yin, "On the global and linear convergence of the generalized alternating direction method of multipliers," J. Sci. Comput., vol. 66, pp. 889–916, 2016.
- [35] Y. Wang, W. Yin, and J. Zeng, "Global convergence of admm in nonconvex nonsmooth optimization," J. Sci. Comput., vol. 78, pp. 29– 63, 2019.
- [36] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, 2011.
- [37] K. Kogiso and T. Fujita, "Cyber-security enhancement of networked control systems using homomorphic encryption," in *Proc. IEEE Conf. Decis. Control*, 2015, pp. 6836–6843.
- [38] M. Ruan, H. Gao, and Y. Wang, "Secure and privacy-preserving consensus," *IEEE Trans. Automat. Control*, vol. 64, pp. 4035–4049, 2019.
- MIT-[39] R. Skelley, and Fischer. J. B. Heisele. "The database." CBCL facial [Online]. expression Available: http://cbcl.mit.edu/software-datasets/FaceData2.html