

Recovering Joint Probability of Discrete Random Variables from Pairwise Marginals

Shahana Ibrahim and Xiao Fu

Abstract—Learning the joint probability of random variables (RVs) is the cornerstone of statistical signal processing and machine learning. However, direct nonparametric estimation for high-dimensional joint probability is in general impossible, due to the curse of dimensionality. Recent work has proposed to recover the joint probability mass function (PMF) of an arbitrary number of RVs from three-dimensional marginals, leveraging the algebraic properties of low-rank tensor decomposition and the (unknown) dependence among the RVs. Nonetheless, accurately estimating three-dimensional marginals can still be costly in terms of sample complexity, affecting the performance of this line of work in practice in the sample-starved regime. Using three-dimensional marginals also involves challenging tensor decomposition problems whose tractability is unclear. This work puts forth a new framework for learning the joint PMF using only pairwise marginals, which naturally enjoys a lower sample complexity relative to the third-order ones. A coupled nonnegative matrix factorization (CNMF) framework is developed, and its joint PMF recovery guarantees under various conditions are analyzed. Our method also features a Gram–Schmidt (GS)-like algorithm that exhibits competitive runtime performance. The algorithm is shown to provably recover the joint PMF up to bounded error in finite iterations, under reasonable conditions. It is also shown that a recently proposed economical expectation maximization (EM) algorithm guarantees to improve upon the GS-like algorithm’s output, thereby further lifting up the accuracy and efficiency. Real-data experiments are employed to showcase the effectiveness.

Index Terms—joint probability learning, nonnegative matrix factorization, probability tensors, two-dimensional marginals

I. INTRODUCTION

Estimating the joint probability of random variables (RVs) from noisy, incomplete and limited observations lies at the heart of statistical signal processing and machine learning. Having the joint probability allows us to construct statistically optimal estimators and detectors under certain metrics. For example, the *maximum a posteriori* (MAP) and the *minimum mean squared error* (MMSE) principles can be easily carried out, if the joint probability of the RVs involved has been estimated [1]. These principles are widely used in tasks such as speech processing, filter design, symbol detection, compressive sensing, and remote sensing; see, e.g., [2]–[5]. In addition, joint probability is the key ingredient for estimating information-theoretic measures, e.g., mutual information and entropy [6], which are often integrated into the solutions

X. Fu and S. Ibrahim are with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331, United States. email (xiao.fu,ibrahi)@oregonstate.edu

This work is supported by the National Science Foundation under NSF-ECCS 1608961, NSF-ECCS 1808159, and NSF IIS-1910118, and the Army Research Office under ARO W911NF-19-1-0247 and W911NF-19-1-0407.

of challenging signal processing problems; see examples in blind source separation [7], image processing [8] and radar waveform design [9]. Joint probability estimation is also the linchpin of data mining, information retrieval, and machine intelligence [10]. Driven by its critical role across multiple domains, joint probability learning using both parametric and nonparametric methods has been a long-term interest in the signal processing community for decades; see, e.g., [11]–[15].

Joint probability estimation poses a variety of challenges in both theory and methods. In particular, in the high-dimensional regime, directly estimating the joint probability via ‘structure-free’ methods such as sample averaging is considered infeasible [16]. Suppose that there are N RVs, where each has an I -dimensional alphabet. To estimate the joint probability reliably, we generally need $S \gg \Omega(I^N)$ ‘diverse’ N -dimensional examples (without any entries missing). This is because the probability of encountering most N -tuples is very low. Therefore, only a small portion of the empirical distribution will be non-zero given a reasonable amount of data samples—this makes the approach considerably inaccurate in most cases.

Many workarounds have been proposed to circumvent this challenge in the literature. For example, linear approximations have been widely used, e.g., the linear MMSE (LMMSE) estimator [17]. Logistic regression, kernel methods, and neural networks can be understood as nonlinear function approximation-based counterparts [10], [18]. These are effective surrogates, but do not directly address the fundamental challenge in estimating high-dimensional joint probability from limited samples. Another commonly adopted route in statistical learning is to make use of problem-specific structural assumptions such as graphical models and prior distributions to help reduce the model complexity [19]–[21]—but making these assumptions restricts the methods to specific applications.

Very recently, Kargas *et al.* proposed a new framework for *blindly* estimating the *joint probability mass function* (PMF) of N discrete finite-alphabet RVs [12]. Building upon a link between high-dimensional joint PMFs and low-rank tensors, the work in [12] shows that any N -dimensional joint PMF admits a naive Bayes *representation*. In addition, if the RVs are ‘reasonably dependent’, the joint PMF can be recovered via jointly decomposing the three-dimensional marginal PMFs (which are third-order tensors).

The work in [12] has shown promising results. However, many challenges remain. First, accurately estimating three-dimensional marginal PMFs is not a trivial task, which requires enumerating the co-occurrences of three RVs. This is particularly hard if the alphabets of the RVs are large and the data is sparse. Second, tensor decomposition is a challenging

optimization problem [22]. Factoring a large number of latent factor-coupled tensors as in [12] also gives rise to scalability issues. A natural question is: can we use pairwise marginals, i.e., joint PMFs of *two* RVs, to achieve the same goal of joint PMF recovery? Pairwise marginals are much easier to estimate relative to triple ones. In addition, the pairwise marginals give rise to probability mass matrices other than tensors—which may circumvent tensor decomposition in algorithm design and thus lead to more lightweight solutions.

Notably, a recent work in [11] offers an *expectation maximization* (EM) algorithm that directly estimates the latent factors of the N th-order probabilistic tensor from the ‘raw data’, instead of working with a large number of three-dimensional marginals. The EM algorithm is well-motivated, since it is associated with the *maximum likelihood estimator* (MLE) under the naive Bayes model in [12]. In addition, it admits simple and economical updates, and thus is much more scalable relative to the coupled tensor decomposition approach. However, performance characterizations (e.g., estimation accuracy under finite samples and convergence properties) of the EM algorithm have been elusive.

Contributions. In this work, we propose a new framework that utilizes pairwise marginals to recover the joint PMF of an arbitrary number of discrete finite-alphabet RVs. We address both the theoretical aspects (e.g., recoverability) and the algorithmic aspects. Our contributions are as follows:

- **A Pairwise Marginal-based Framework.** On the theory side, we offer in-depth analyses for the recoverability of the joint PMF from pairwise marginals. We show that the joint PMF can be recovered via jointly factoring the pairwise marginals under nonnegativity constraints in a latent factor-coupled manner. Unlike [12] which leverages the uniqueness of tensor decomposition to establish recoverability, our analysis utilizes the model identifiability of *nonnegative matrix factorization* (NMF)—leading to recoverability conditions that have a different flavor compared to those in [12]. More importantly, accurately estimating pairwise marginals is much easier relative to the three-dimensional counterparts, in terms of the amount of data samples needed.

- **Provable Algorithms.** On the algorithm side, we propose to employ a block coordinate descent (BCD)-based algorithm for the formulated coupled NMF problem. More importantly, we propose a pragmatic and easy-to-implement initialization approach, which is based on performing a Gram–Schmidt (GS)-like structured algorithm on a carefully constructed ‘virtual NMF’ model. We show that this approach works provably well even when there is modeling noise—e.g., noise induced by finite-sample estimation of the pairwise marginals. This is in contrast to the method in [12], whose recoverability guarantees are based on the assumption that the third-order marginals are perfectly accessible. We also show that the EM algorithm from [11] can provably improve upon proper initialization (e.g., that given by the GS-like algorithm) towards the target latent model parameters that are needed to reconstruct the joint PMF.

- **Extensive Validation.** We test the proposed approach on a large variety of synthetic and real-world datasets. In particular, we validate our theory via performing classification on four different UCI datasets and three movie recommendation tasks.

Part of the paper was presented at the Asilomar signal processing conference 2020 [23]. The journal version additionally includes new theoretical results, e.g., the recoverability analysis for the coupled NMF formulation, and a BCD-based CNMF algorithm, the detailed proofs for the Gram–Schmidt-like procedure, and the optimality analysis for the EM algorithm. More real data experiments are also presented.

Notation. We use $x, \mathbf{x}, \mathbf{X}, \underline{\mathbf{X}}$ to denote a scalar, vector, matrix and tensor, respectively. \mathbb{D}_{KL} denotes the KL divergence. $\kappa(\mathbf{X})$ denotes the condition number of the matrix \mathbf{X} and is given by $\kappa(\mathbf{X}) = \frac{\sigma_{\max}(\mathbf{X})}{\sigma_{\min}(\mathbf{X})}$ where σ_{\max} and σ_{\min} are the largest and smallest singular values of \mathbf{X} , respectively. $\|\mathbf{x}\|_0$ denotes the ‘zero norm’ of the vector \mathbf{x} , i.e., the number of nonzero elements in \mathbf{x} . $\mathbf{X} \geq \mathbf{0}$ implies that all the elements of \mathbf{X} are nonnegative. $\text{cone}(\mathbf{X})$ represents the conic hull formed by the columns of \mathbf{X} and $\text{conv}(\mathbf{X})$ denotes the convex hull formed by the columns of \mathbf{X} . \emptyset denotes the empty set. $\|\mathbf{X}\|_2$ represents the 2-norm of matrix \mathbf{X} and $\|\mathbf{X}\|_2 = \sigma_{\max}(\mathbf{X})$. $\|\mathbf{X}\|_{\text{F}}$ denotes the Frobenius norm of \mathbf{X} . $\|\mathbf{x}\|_2$ and $\|\mathbf{x}\|_1$ denote ℓ_2 and ℓ_1 norm of vector \mathbf{x} , respectively. $\text{vec}(\mathbf{X})$ denotes the vectorization operator that concatenates the columns of \mathbf{X} . \top and \dagger denote transpose and pseudo-inverse, respectively. $|\mathcal{C}|$ denotes the cardinality of set \mathcal{C} . $\text{Diag}(\mathbf{x})$ is a diagonal matrix that holds the entries of \mathbf{x} as its diagonal elements. ‘ \odot ’ denotes the Khatri–Rao product. $[T]$ denotes $\{1, \dots, T\}$, where T is an integer. \mathbf{e}_n denotes the unit vector with n th element being one. $f(x) = O(g(x))$ (where $f(\cdot)$ and $g(\cdot)$ are nonnegative) means that there exist a positive constant C and x_0 such that $f(x) \leq Cg(x), \forall x \geq x_0$. Similarly, $f(x) = \Omega(g(x))$ means that there exist a positive constant c and x_0 such that $f(x) \geq cg(x), \forall x \geq x_0$.

II. BACKGROUND

Consider a set of discrete and finite-alphabet RVs, i.e., Z_1, \dots, Z_N . We will use $\Pr(i_1, \dots, i_N)$ as the shorthand notation to represent

$$\Pr\left(Z_1 = z_1^{(i_1)}, \dots, Z_N = z_N^{(i_N)}\right)$$

in the sequel, where $\{z_n^{(1)}, \dots, z_n^{(I_n)}\}$ denotes the alphabet of Z_n . Suppose that we do not have access to the joint PMF $\Pr(i_1, \dots, i_N)$. Instead, we have access to the lower-dimensional marginal distributions, e.g., $\Pr(i_j, i_k, i_\ell)$ and $\Pr(i_j, i_k)$ for different j, k, ℓ . Kargas *et al.* asked the following research question [12]:

Can the joint distribution of Z_1, \dots, Z_N be recovered from the low-dimensional marginals without knowing any structural information about the underlying probabilistic model?

Note that estimating a high-dimensional joint distribution from lower-dimensional marginals is challenging, but not entirely impossible—if some structural information between the RVs (e.g., tree or Markov chain) is known; see discussions in [12]. However, learning the structural information from data *per se* is often a hard problem [19].

A. Tensor Algebra

The work in [12] utilizes a connection between joint PMFs and low-rank tensors under the *canonical polyadic decomposition* (CPD) model [24], [25]. If an N -way tensor $\underline{\mathbf{X}} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ has CP rank F , it can be written as:

$$\underline{\mathbf{X}}(i_1, i_2, \dots, i_N) = \sum_{f=1}^F \boldsymbol{\lambda}(f) \prod_{n=1}^N \mathbf{A}_n(i_n, f), \quad (1)$$

where $\mathbf{A}_n \in \mathbb{R}^{I_n \times F}$ is called the mode- n latent factor. In the above, $\boldsymbol{\lambda} = [\boldsymbol{\lambda}(1), \dots, \boldsymbol{\lambda}(F)]^\top$ with $\|\boldsymbol{\lambda}\|_0 = F$ is employed to ‘absorb’ the norms of columns (so that the columns of \mathbf{A}_n can be normalized w.r.t. a certain ℓ_q norm for $q \geq 1$). We use the shorthand notation

$$\underline{\mathbf{X}} = [\![\boldsymbol{\lambda}, \mathbf{A}_1, \dots, \mathbf{A}_N]\!]$$

to denote the CPD in (1). Note that $F \gg I_n$ for $n = 1, \dots, N$ could happen. In fact, the tensor rank can reach $O(I^{N-1})$ if $I_1 = \dots = I_N = I$ [25].

B. Joint PMF Recovery: A Tensor Perspective

A key takeaway in [12] is that *any* joint PMF admits a naive Bayes (NB) model *representation*; i.e., any joint PMF can be generated from a latent variable model with just one hidden variable. It follows that the joint PMF of $\{Z_n\}_{n=1}^N$ can always be decomposed as

$$\Pr(i_1, i_2, \dots, i_N) = \sum_{f=1}^F \Pr(f) \prod_{n=1}^N \Pr(i_n|f), \quad (2)$$

where $\Pr(f) := \Pr(H = f)$ is the prior distribution of a latent variable H and $\Pr(i_n|f) := \Pr(Z_n = z_n^{(i_n)}|H = f)$ are the conditional distributions. Remarkably, such representation is *universal* if one allows the hidden variable to have a sufficiently rich finite alphabet:

Theorem 1 [12] *The maximum F needed in order to describe an arbitrary PMF using a naive Bayes model is bounded by $F \leq \min_k \prod_{\substack{n=1 \\ n \neq k}}^N I_n$.*

To see this, one can represent any joint PMF as an N th-order tensor by letting $\underline{\mathbf{X}}(i_1, \dots, i_N) = \Pr(i_1, \dots, i_N)$ and $\mathbf{A}_n(i_n, f) = \Pr(i_n|f)$, $\boldsymbol{\lambda}(f) = \Pr(f)$. Since any nonnegative tensor admits a nonnegative polyadic decomposition with nonnegative latent factors, any joint PMF admits a naive Bayes representation with finite latent alphabet.

A relevant question is: what does ‘low tensor rank’ mean for the joint PMF tensors? In [12], the authors argue that reasonably dependent RVs lead to a joint PMF tensor with relatively small F —which is normally the case in statistical learning and inference. For example, classification is concerned with inferring Z_N (label) from Z_1, \dots, Z_{N-1} (features), and the dependence amongs the RVs are leveraged for building up effective classifiers.

Another key observation in [12] is that the marginal distribution of any subset of RVs is an induced CPD model. From

the law of total probability, if one marginalize the joint PMF down to $\Pr(i_j, i_k, i_\ell) \forall j, k, \ell \in [N]$, the expression is

$$\Pr(i_j, i_k, i_\ell) = \sum_{f=1}^F \Pr(f) \Pr(i_j|f) \Pr(i_k|f) \Pr(i_\ell|f).$$

Let $\underline{\mathbf{X}}_{jkl}(i_j, i_k, i_\ell) = \Pr(i_j, i_k, i_\ell)$. Then, we have $\underline{\mathbf{X}}_{jkl} = [\![\boldsymbol{\lambda}, \mathbf{A}_j, \mathbf{A}_k, \mathbf{A}_\ell]\!]$, where $\{\mathbf{A}_n\}_{n=1}^N$ and $\boldsymbol{\lambda}$ are defined as before. Another important observation here is that *the marginal PMFs $\underline{\mathbf{X}}_{jkl}$ ’s and the joint PMF $\underline{\mathbf{X}}$ share the same \mathbf{A}_j , \mathbf{A}_k , \mathbf{A}_ℓ and $\boldsymbol{\lambda}$* . It is readily seen that if the $\underline{\mathbf{X}}_{jkl}$ ’s admit essentially unique CPD, then \mathbf{A}_n ’s and $\boldsymbol{\lambda}$ can be identified from the marginals—and thus $\Pr(i_1, \dots, i_N)$ can be recovered via identifying the latent factors.

The tensor-based approach in [12] has shown that recovering high-dimensional joint PMF from low-dimensional marginals is possible. However, a couple of major hurdles exist for applying it in practice. First, estimating three-dimensional marginals $\Pr(i_j, i_k, i_\ell)$ is still not easy, since one needs many co-realizations of three RVs. For sparse datasets, this is particularly hard. Second, tensor decomposition is a hard computation problem [22]. The coupled tensor decomposition (CTD) algorithm in [12] that involves many tensors is even more challenging in practice.

C. Maximum Likelihood Estimation and EM

A recent work in [11] offers an alternative for estimating $\Pr(f)$ and $\Pr(i_n|f)$. Instead of working with third-order statistics, the model estimation problem is formulated as a maximum likelihood (ML) estimation problem in [11], and an EM algorithm is employed to handle the ML estimator (MLE). The upshot of this approach is that the EM updates are much more scalable relative to the CTD updates in [12]. However, there are two challenges of applying EM. First, the recoverability guarantees of the joint PMF using the MLE framework is unclear under finite number of observations. Second, as the authors of [11] have noticed, the EM algorithm sometimes does not converge to a good solution if randomly initialized. This is perhaps due to the nonconvex nature of the ML estimation problem.

III. PROPOSED APPROACH

To advance the task of joint PMF recovery from marginal distributions, we propose a pairwise marginal-based approach. Note that the pairwise marginal is given by

$$\begin{aligned} \Pr(i_j, i_k) &= \sum_{f=1}^F \Pr(f) \Pr(i_j|f) \Pr(i_k|f) \\ &\iff \mathbf{X}_{jk} = \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top, \quad \mathbf{X}_{jk}(i_j, i_k) = \Pr(i_j, i_k), \end{aligned}$$

where $\mathbf{D}(\boldsymbol{\lambda}) = \text{Diag}(\boldsymbol{\lambda})$.

In practice, the pairwise marginals \mathbf{X}_{jk} ’s are estimated from realizations of the joint PMF. Consider a set of realizations (data samples) of $\Pr(Z_1, \dots, Z_N)$, denoted as $\{\mathbf{d}_s \in \mathbb{R}^N\}_{s=1}^S$. The following sample averaging estimator can be employed:

$$\widehat{\mathbf{X}}_{jk}(i_j, i_k) = \frac{1}{|\mathcal{S}_{jk}|} \sum_{s \in \mathcal{S}_{jk}} \mathbb{I} \left[\mathbf{d}_s(j) = z_j^{(i_j)}, \mathbf{d}_s(k) = z_k^{(i_k)} \right],$$

where $\mathbf{d}_s(n)$ denotes the realization of Z_n in the s -th sample, $\mathbb{I}[E] = 1$ if the event E happens and $\mathbb{I}[E] = 0$ otherwise, S is the number of co-realizations of Z_1, \dots, Z_N (i.e., the number of ‘samples’), and $\mathcal{S}_{jk} \subseteq [S]$ denotes the set of indices of the data samples in which the realizations of both Z_j and Z_k are observed.

Notably, with the same amount of data (i.e., with a fixed S), the second-order statistics can be estimated to a much higher accuracy, compared to the third-order ones [26]. This is particularly articulated when the dataset is sparse. To see this, assume that every RV has an alphabet size I and is observed with probability $p \in (0, 1]$ (which means that every entry of \mathbf{d}_s is observed with probability p —and \mathbf{d}_s is very sparse when p is small). Then, on average Sp^2 and Sp^3 co-realizations of every pair and triple of RVs, respectively, can be observed in the samples. Under such cases, Theorem 1 in [26] shows that the average estimation error (over all possible joint distributions) for $\underline{\mathbf{X}}_{jkl}$ in terms of ℓ_1 distance is $O(\sqrt{I^3/Sp^3})$, while that for \mathbf{X}_{jk} is $O(\sqrt{I^2/Sp^2})$. One can see that the difference is particularly significant when p is small—i.e., when the dataset is considerably sparse.

A. Challenges

Despite the sample complexity appeal of using pairwise marginals, the *recoverability* of the joint PMF under such settings is nontrivial to establish. The main analytical tool for proving recoverability in [12] is the essential uniqueness of tensor decomposition—i.e., the latent factors of low-rank tensors are identifiable up to resolvable ambiguities under mild conditions. However, pairwise distributions such as $\mathbf{X}_{jk} = \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top$ are matrices, and low-rank matrix decomposition is in general *nonunique*. A natural thought to handle the identifiability problem would be employing certain NMF tools [27], [28], since the latent factors are all nonnegative, per their physical interpretations. However, the identifiability of NMF models holds only if $F \leq \min\{I_j, I_k\}$ (and preferably $F \ll \min\{I_j, I_k\}$). The pairs $\mathbf{X}_{jk} = \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top \in \mathbb{R}^{I_j \times I_k}$ inherit the inner dimension F (i.e., the column dimension of \mathbf{A}_j) from the joint PMF of all the variables, which is the tensor rank of an N th-order tensor. As mentioned, the tensor rank F could be much larger than the I_j ’s. Hence, one may not directly use the available NMF uniqueness results on individual \mathbf{X}_{jk} ’s to argue for joint PMF recoverability. Nonetheless, as we will show shortly, NMF identifiability can be applied onto carefully constructed *coupled* NMF models to establish recoverability.

B. Preliminaries: NMF and Model Identifiability

To see how we approach these challenges, let us first briefly introduce some pertinent preliminaries of NMF. Consider a nonnegative data matrix $\mathbf{X} \in \mathbb{R}^{L \times K}$. Assume that the matrix is generated by $\mathbf{X} = \mathbf{W} \mathbf{H}^\top$, where $\mathbf{W} \in \mathbb{R}^{L \times F}$ and $\mathbf{H} \in \mathbb{R}^{K \times F}$. Also assume that $\mathbf{W} \geq \mathbf{0}$ and $\mathbf{H} \geq \mathbf{0}$ and $F \leq \min\{L, K\}$. NMF aims to factor \mathbf{X} for identifying \mathbf{W} and \mathbf{H} up to certain trivial ambiguities.

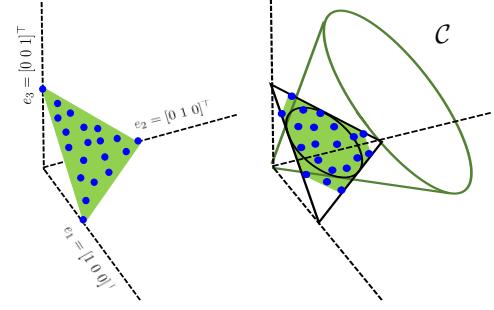


Fig. 1. Left: an $\mathbf{H} \in \mathbb{R}^{20 \times 3}$ that satisfies the separability condition. Right: an $\mathbf{H} \in \mathbb{R}^{20 \times 3}$ that satisfies the sufficiently scattered condition. The triangle corresponds to probability simplex Δ_3 ; and the dots correspond to rows of \mathbf{H} rescaled to have unit ℓ_1 -norm. The inner circle on Δ_3 corresponds to the intersection of the second-order cone \mathcal{C} and Δ_3 ; and the shaded region is the intersection of cone $\{\mathbf{H}^\top\}$ and Δ_3 .

A number of conditions on \mathbf{W} and/or \mathbf{H} allow us to establish identifiability of the latent factors. The first one is the so-called separability condition:

Definition 1 (Separability) [29] If $\mathbf{H} \geq \mathbf{0}$, and $\boldsymbol{\Lambda} = \{l_1, \dots, l_F\}$ such that $\mathbf{H}(\boldsymbol{\Lambda}, :) = \boldsymbol{\Sigma}$ holds, where $\boldsymbol{\Sigma} = \text{Diag}(\alpha_1, \dots, \alpha_F)$ and $\alpha_f > 0$, then, \mathbf{H} satisfies the separability condition. When $\boldsymbol{\Lambda} = \{l_1, \dots, l_F\}$ satisfies $\|\mathbf{H}(l_f, :) - \mathbf{e}_f\|_2 \leq \varepsilon$ for $f = 1, \dots, F$, \mathbf{H} is called ε -separable.

If one of \mathbf{W} and \mathbf{H} satisfies the separability condition and the other has full column rank, there exist polynomial-time algorithms that can provably identify \mathbf{W} and \mathbf{H} up to scaling and permutation ambiguities. Among these algorithms, a number of Gram–Schmidt-like lightweight algorithms exist; see, e.g., [30]–[34]. A more relaxed condition is as follows:

Definition 2 (Sufficiently Scattered) [35]–[37] Assume that $\mathbf{H} \geq \mathbf{0}$ and $\mathcal{C} \subseteq \text{cone}\{\mathbf{H}^\top\}$ where $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^F \mid \mathbf{x}^\top \mathbf{1} \geq \sqrt{F-1}\|\mathbf{x}\|_2\}$. In addition, assume that $\text{cone}\{\mathbf{H}^\top\} \not\subseteq \text{cone}\{\mathbf{Q}\}$ for any orthonormal $\mathbf{Q} \in \mathbb{R}^{F \times F}$ except for the permutation matrices. Then, \mathbf{H} is called sufficiently scattered.

The sufficiently scattered condition subsumes the separability condition as a special case. An important result is that if \mathbf{W} and \mathbf{H} are both sufficiently scattered, then the model $\mathbf{X} = \mathbf{W} \mathbf{H}^\top$ is unique up to scaling and permutation ambiguities [35]. The two conditions are illustrated in Fig. 1. More detailed discussions and illustrations can be seen in [27], [38].

In principle, the ‘taller’ \mathbf{H} is, the chance that \mathbf{H} (approximately) attains the separability or sufficiently scattered condition is higher, if \mathbf{H} is generated following a certain continuous distribution. This intuition was formalized in [39], under the assumption that $\mathbf{H}(l, :)$ ’s are drawn from the probability simplex uniformly at random and then positively scaled. The work in [39] also shows that the sufficiently scattered condition is easier to satisfy relative to separability, under the same generative model; see Lemma 6 in Appendix E in the supplementary materials.

C. Recoverability Analysis

Our goal is to identify \mathbf{A}_n and $\boldsymbol{\lambda}$ from the available pairwise marginals $\mathbf{X}_{jk} = \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top$'s, so that the joint PMF can be reconstructed following (2). Note that applying existing NMF techniques onto \mathbf{X}_{jk} is unlikely to work, since when $F > \min\{I_j, I_k\}$, 'fat' $\mathbf{A}_j \in \mathbb{R}^{I_j \times F}$'s can never be separable or sufficiently scattered. To circumvent this, our idea is as follows. Consider a splitting of the indices of the N variables, i.e., $\mathcal{S}_1 = \{\ell_1, \dots, \ell_M\}$ and $\mathcal{S}_2 = \{\ell_{M+1}, \dots, \ell_N\}$ such that $\mathcal{S}_1 \cup \mathcal{S}_2 = \{1, \dots, N\}$ and $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$. Then, we construct the following matrix:

$$\begin{aligned} \widetilde{\mathbf{X}} &= \begin{bmatrix} \mathbf{X}_{\ell_1 \ell_{M+1}} & \dots & \mathbf{X}_{\ell_1 \ell_N} \\ \vdots & \vdots & \vdots \\ \mathbf{X}_{\ell_M \ell_{M+1}} & \dots & \mathbf{X}_{\ell_M \ell_N} \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} \mathbf{A}_{\ell_1} \\ \vdots \\ \mathbf{A}_{\ell_M} \end{bmatrix}}_{\mathbf{W}} \underbrace{\mathbf{D}(\boldsymbol{\lambda})[\mathbf{A}_{\ell_{M+1}}^\top, \dots, \mathbf{A}_{\ell_N}^\top]}_{\mathbf{H}^\top}. \end{aligned} \quad (3)$$

Note that in this 'virtual NMF' model, \mathbf{W} and \mathbf{H} matrices have sizes of $MI \times F$ and $(N - M)I \times F$, respectively, if $I_1 = \dots = I_N = I$. The idea is to construct $\widetilde{\mathbf{X}}$ such that $F \leq \min\{MI, (N - M)I\}$ so that \mathbf{W} and \mathbf{H} may satisfy the separability or sufficiently scattered condition—and thus $\mathbf{A}_1, \dots, \mathbf{A}_N$ can be identified via factoring $\widetilde{\mathbf{X}}$. This is a straightforward application of the NMF identifiability results introduced in the previous subsection.

The challenge lies in finding a suitable splitting of $\mathcal{S}_1, \mathcal{S}_2$ such that \mathbf{W} and \mathbf{H} are sufficiently scattered and is highly nontrivial. Note that even if \mathbf{W} and \mathbf{H} are given, checking if these matrices are sufficiently scattered is hard [35]. It is also impractical to try every possible index splitting. To address this challenge, we consider the following factorization problem:

$$\text{find } \mathbf{A}_n, \quad n = 1, \dots, N, \quad \boldsymbol{\lambda} \quad (4a)$$

$$\text{subject to } \mathbf{X}_{jk} = \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top, \quad \forall j, k \in \Omega \quad (4b)$$

$$\mathbf{1}^\top \mathbf{A}_j = \mathbf{1}^\top, \quad \mathbf{A}_j \geq \mathbf{0} \quad (4c)$$

$$\mathbf{1}^\top \boldsymbol{\lambda} = 1, \quad \boldsymbol{\lambda} \geq \mathbf{0}, \quad (4d)$$

where Ω contains the index set of (j, k) 's such that $j < k$ and the joint PMF $\Pr(i_j, i_k)$ is accessible. Note that the above can be understood as a latent factor-coupled NMF problem. Regarding the identifiability of the conditional PMFs and the latent prior, we have the following result:

Theorem 2 (Recoverability) *Assume that that $\Pr(i_j, i_k)$'s for $j, k \in \Omega$ are available and that $\Pr(f) \neq 0$ for $f = 1, \dots, F$. Suppose that there exists $\mathcal{S}_1 = \{\ell_1, \dots, \ell_M\}$ and $\mathcal{S}_2 = \{\ell_{M+1}, \dots, \ell_Q\}$ such that $Q \leq N$ and $\mathcal{S}_1 \cup \mathcal{S}_2 \subseteq \{1, \dots, N\}$, $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$. Also assume the following conditions hold:*

- (i) *the matrices $[\mathbf{A}_{\ell_1}^\top, \dots, \mathbf{A}_{\ell_M}^\top]^\top$ and $[\mathbf{A}_{\ell_{M+1}}^\top, \dots, \mathbf{A}_{\ell_Q}^\top]^\top$ are sufficiently scattered;*
- (ii) *all pairwise marginal distributions $\Pr(i_j, i_k)$'s for $j \in \mathcal{S}_1$ and $k \in \mathcal{S}_2$ are available;*
- (iii) *every T -concatenation of \mathbf{A}_n 's, i.e., $[\mathbf{A}_{n_1}^\top, \dots, \mathbf{A}_{n_T}^\top]^\top$, is*

a full column rank matrix, if $I_{n_1} + \dots + I_{n_T} \geq F$;

(iv) for every $j \notin \mathcal{S}_1 \cup \mathcal{S}_2$ there exists a set of $r_t \in \mathcal{S}_1 \cup \mathcal{S}_2$ for $t = 1, \dots, T$ such that $\Pr(i_j, i_{r_t})$ or $\Pr(i_{r_t}, i_j)$ are available.

Then, solving Problem (4) recovers $\Pr(i_j|f)$ and $\Pr(f)$ for $j = 1, \dots, N$, $f = 1, \dots, F$, thereby the joint PMF $\Pr(i_1, \dots, i_N)$.

The proof is relegated to Appendix A, which is reminiscent of identifiability of a similar coupled NMF problem that arises in crowdsourced data labeling [39], with proper modifications to handle the cases where $F > I_n$ (since the crowdsourcing model always has $F = I_n$, which is a simpler case). In particular, conditions (iii) and (iv) are employed to accommodate the more challenging cases. The criterion spares one the effort for first finding \mathcal{S}_1 and \mathcal{S}_2 and then constructing the matrix $\widetilde{\mathbf{X}}$. Instead, under Theorem 2, the goal is simply finding an approximate solution of (4)—which is much more approachable in practice.

Remark 1 *Under our pairwise marginal based approach, a key enabler for recovering the joint PMF is the identifiability of \mathbf{A}_n 's up to an identical column permutation ambiguity. At first glance, the identifiability of the latent factors seems to be irrelevant since what we truly care about is the ambient joint PMF tensor $\underline{\mathbf{X}} \in \mathbb{R}^{I_1 \times \dots \times I_N}$. However, since the tensor is heavily marginalized (i.e., one only observes pairwise marginals), the identifiability of \mathbf{A}_n 's is leveraged on to assemble the joint PMF following the model in (2) under our framework.*

IV. ALGORITHMS AND PERFORMANCE ANALYSES

In this section, we develop algorithms under the coupled NMF framework and analyze performance under realistic conditions, e.g., cases with finite samples and missing values.

A. BCD for Problem (4)

To handle Problem (4), we employ the standard procedure based on *block coordinate descent* (BCD) for handling coupled tensor/matrix decomposition algorithms in the literature [12], [39], [40]. To be specific, we cyclically minimize the constrained optimization problem w.r.t. \mathbf{A}_k , when fixing \mathbf{A}_j for all $j \neq k$ and $\boldsymbol{\lambda}$. Consequently, the subproblem is as follows:

$$\underset{\mathbf{A}_k}{\text{minimize}} \sum_{j \in \Omega_k} \mathbb{D}(\mathbf{X}_{jk}, \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top) \quad (5a)$$

$$\text{subject to } \mathbf{1}^\top \mathbf{A}_k = \mathbf{1}^\top, \quad \mathbf{A}_k \geq \mathbf{0}, \quad (5b)$$

where Ω_k is the index set of j such that $\Pr(i_j, i_k)$ is available and $\mathbb{D}(\cdot, \cdot)$ is a certain 'distance' measure. Under a number of $\mathbb{D}(\cdot, \cdot)$ -functions, e.g., the Euclidean distance or the Kullback–Leibler (KL) divergence, the subproblems are convex and thus can be solved via off-the-shelf algorithms. In this work, we employ the KL divergence as the fitting measure since it is

natural for measuring the similarity of PMFs. Consequently, the subproblems are

$$\mathbf{A}_k \leftarrow \arg \min_{\substack{\mathbf{1}^\top \mathbf{A}_k = \mathbf{1}^\top \\ \mathbf{A}_k \geq \mathbf{0}}} \sum_{j \in \Omega_k} \mathbb{D}_{\text{KL}}(\mathbf{X}_{jk}, \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top), \quad (6a)$$

$$\boldsymbol{\lambda} \leftarrow \arg \min_{\substack{\mathbf{1}^\top \boldsymbol{\lambda} = \mathbf{1} \\ \boldsymbol{\lambda} \geq \mathbf{0}}} \sum_{j, k \in \Omega} \mathbb{D}_{\text{KL}}(\mathbf{X}_{jk}, \mathbf{A}_j \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_k^\top). \quad (6b)$$

We employ the *mirror descent* algorithm to solve the two subproblems. Mirror descent admits closed-form updates for handling KL-divergence based linear model fitting under simplex constraints (which is also known as exponential gradient in the literature); see, e.g., [32]. The algorithm CNMF-OPT is described in (1).

Algorithm 1: CNMF-OPT

```

input : data samples  $\{\mathbf{d}_s\}_{s=1}^S$ 
1 estimate second order statistics  $\widehat{\mathbf{X}}_{jk}$ ;
2 get initial estimates  $\{\widehat{\mathbf{A}}_n\}_{n=1}^N$  and  $\widehat{\boldsymbol{\lambda}}$  (e.g., using CNMF-SPA); repeat
3   for  $n = 1$  to  $N$  do
4     | update  $\widehat{\mathbf{A}}_n \leftarrow$  (6a) using mirror descent;
5   end
6   update  $\widehat{\boldsymbol{\lambda}} \leftarrow$  (6b) using mirror descent;
7 until some stopping criterion is reached;
output: estimates  $\{\widehat{\mathbf{A}}_n\}_{n=1}^N$ ,  $\widehat{\boldsymbol{\lambda}}$ 

```

Notably, the work in [12] also has some experiments (without recoverability analysis) using pairwise marginals to recover the joint PMF using a similar formulation (using Euclidean distance). Our rationale for reaching the formulation in (4) is very different, which is motivated by the NMF theory. Also, the result using pairwise marginals in [12] seems not promising. In this work, our analysis in Theorems 2 has suggested that the reason why pairwise marginals do not work in [12] is likely due to optimization pitfalls, rather than lack of recoverability. In the next subsection, we show that with a carefully designed initialization scheme, accurately recovering joint PMFs from pairwise marginals is viable.

B. Key Step: Gram–Schmidt-like Initialization

Directly applying the BCD algorithm to the joint PMF recovery problem has a number of challenges. First, the CNMF problem is nonconvex, and thus global optimality is not ensured. Second, the computational complexity grows quickly with N and I , giving rise to potential scalability issues. In this subsection, we offer a fast algorithm that can provably identify $\Pr(i_n|f)$ and $\Pr(f)$ up to a certain accuracy—under more stringent conditions relative to those in Theorem 2. The algorithm is reminiscent of a Gram–Schmidt-like algorithm for NMF with careful construction of a ‘virtual NMF model’.

Let us assume that there is a splitting $\mathcal{S}_1 = \{\ell_1, \dots, \ell_M\}$ and $\mathcal{S}_2 = \{\ell_{M+1}, \dots, \ell_N\}$ such that the $\mathbf{X} = \mathbf{W}\mathbf{H}^\top$ in (3) admits a full rank \mathbf{W} and a separable \mathbf{H} . Under the separability condition, we have $\mathbf{H}(\mathbf{A}, :) = \boldsymbol{\Sigma} = \text{Diag}(\alpha_1, \dots, \alpha_F)$ and $\mathbf{W}\boldsymbol{\Sigma} = \widehat{\mathbf{X}}(:, \mathbf{A})$. Hence, the coupled NMF task boils down to identifying the index set \mathbf{A} . The so-called *successive projection algorithm* (SPA) from the NMF literature [30]–[32], [34] can be employed for this purpose. A remark is that SPA is

a Gram–Schmidt-like algorithm, which only consists of norm comparison and orthogonal projection.

Once \mathbf{W} is identified, one can recover $\mathbf{A}_{\ell_n} \in \mathbb{R}^{I_{\ell_n} \times F}$ for $\ell_n \in \mathcal{S}_1$ up to identical column permutations, by extracting the corresponding rows of \mathbf{W} (cf. line 6 in Algorithm 2). Unlike general NMF models, since every column of \mathbf{A}_n is a conditional PMF, there is no scaling ambiguity. The \mathbf{H} matrix can be estimated using (constrained) least squares, and \mathbf{A}_{ℓ_n} for $\ell_n \in \mathcal{S}_2$ can then be extracted in a similar way. Denote (3) as $\widehat{\mathbf{X}} = \mathbf{W}\mathbf{D}(\boldsymbol{\lambda})\widehat{\mathbf{H}}^\top$, where $\widehat{\mathbf{H}} = [\mathbf{A}_{\ell_{M+1}}^\top \dots \mathbf{A}_{\ell_N}^\top]^\top$. Then, the PMF of the latent variable can be estimated via $\boldsymbol{\lambda} = (\widehat{\mathbf{H}} \odot \mathbf{W})^\dagger \text{vec}(\widehat{\mathbf{X}})$, where we have used the fact that the Khatri–Rao product $\widehat{\mathbf{H}} \odot \mathbf{W}$ has full column rank since both \mathbf{W} and $\widehat{\mathbf{H}}$ have full column rank.

A remark is that the SPA-estimated $\widehat{\mathbf{A}}_{\ell_n}$ is at best the column-permuted version of \mathbf{A}_{ℓ_n} [30]. However, since the permutation ambiguity across all the \mathbf{A}_n ’s and the $\boldsymbol{\lambda}$ are identical per the above procedure, the existence of column permutations does not affect the ‘assembling’ of $\Pr(i_n|f)$ and $\Pr(f)$ to recover $\Pr(i_1, \dots, i_N)$. We refer to this procedure as *coupled NMF via SPA* (CNMF-SPA); see Algorithm 2.

From Algorithm 2, one can see that the procedure is rather simple since SPA is a Gram–Schmidt-like procedure and the other steps are either convex quadratic programming or least squares. Hence, CNMF-SPA can be employed for initializing computationally more intensive algorithms, e.g., CNMF-OPT.

Algorithm 2: CNMF-SPA

```

input : data samples  $\{\mathbf{d}_s\}_{s=1}^S$  and  $M$ 
1 estimate second order statistics  $\widehat{\mathbf{X}}_{jk}$ ;
2 split  $\{1, \dots, N\}$  into  $\mathcal{S}_1 = \{1, \dots, M\}$  and  $\mathcal{S}_2 = \{M+1, \dots, N\}$ ;
3 Construct  $\widehat{\mathbf{X}}$ ;
4 Estimate  $\widehat{\mathbf{W}}$  using the SPA algorithm [30] to select  $\mathbf{A}$ ;
5  $b \leftarrow 0$ ;
6 for  $n = 1$  to  $M$  do
7   |  $\widehat{\mathbf{A}}_n \leftarrow \widehat{\mathbf{W}}(b + [1, \dots, I_n], :)$ ;
8   | normalize columns of  $\widehat{\mathbf{A}}_n$  with respect to  $\ell_1$  norm;
9   |  $b \leftarrow b + I_n$ ;
10 end
11  $\widehat{\mathbf{H}} \leftarrow \arg \min_{\mathbf{H} \geq 0} \|\widehat{\mathbf{X}} - \widehat{\mathbf{W}}\mathbf{H}^\top\|_F^2$ ;
12  $b \leftarrow 0$ ;
13 for  $n = M+1$  to  $N$  do
14   |  $\widehat{\mathbf{A}}_n \leftarrow \widehat{\mathbf{H}}(b + [1, \dots, I_n], :)$ ;
15   | normalize columns of  $\widehat{\mathbf{A}}_n$  with respect to  $\ell_1$  norm;
16   |  $b \leftarrow b + I_n$ ;
17 end
18  $\widehat{\mathbf{W}}^\top \leftarrow [\widehat{\mathbf{A}}_1^\top, \dots, \widehat{\mathbf{A}}_M^\top]^\top$ ;
19  $\widehat{\mathbf{H}}^\top \leftarrow [\widehat{\mathbf{A}}_{M+1}^\top \dots \widehat{\mathbf{A}}_N^\top]^\top$ ;
20  $\widehat{\boldsymbol{\lambda}} \leftarrow (\widehat{\mathbf{H}} \odot \widehat{\mathbf{W}})^\dagger \text{vec}(\widehat{\mathbf{X}})$ ;
output: estimates  $\{\widehat{\mathbf{A}}_n\}_{n=1}^N$ ,  $\widehat{\boldsymbol{\lambda}}$ 

```

C. Performance Analysis of CNMF-SPA

In principle, to make CNMF-SPA work, one needs to identify a splitting scheme \mathcal{S}_1 and \mathcal{S}_2 such that \mathbf{H} in (3) satisfies the separability condition. It is impossible to know such a splitting *a priori*. Testing all combinations of \mathcal{S}_1 and \mathcal{S}_2 gives rise to a hard combinatorial problem. In addition, the pairwise marginals \mathbf{X}_{jk} are not perfectly estimated in practice.

In this subsection, we offer performance analysis of CNMF-SPA by taking the above aspects into consideration.

From our extensive experiments, we observed that using the ‘naive’ construction $\mathcal{S}_1 = \{1, \dots, M\}$ and $\mathcal{S}_2 = \{M + 1, \dots, N\}$ seems to work reasonably well for various tasks. To establish theoretical understanding to such effectiveness, recall that we use S and p to denote the number of available realizations of $\Pr(Z_1, \dots, Z_N)$ and the probability of observing each entry of data sample \mathbf{d}_s , respectively. For simplicity, we assume that $I_n = I$ for all n . We impose the following generative model for \mathbf{A}_m :

Assumption 1 Assume that every row of $\mathbf{A}_n \in \mathbb{R}^{I_m \times F}$ for all n is generated from the $(F - 1)$ -probability simplex uniformly at random and then positively scaled by a scalar, so that $\mathbf{1}^\top \mathbf{A}_n = \mathbf{1}^\top$ is respected¹.

Under the above settings, we show that the following holds:

Theorem 3 Assume that $\|\widehat{\mathbf{X}}_{jk}(:, q)\|_1 \geq \eta > 0$ for any q, j, k . Also, assume that $M \geq F/I$, $p \geq (\frac{8}{S} \log(4/\delta))^{1/2}$,

$$S = \Omega\left(\frac{M^2 I \log(1/\delta)}{\sigma_{\max}^2(\mathbf{W}) \eta^2 \varepsilon^2 p^2}\right),$$

$$N = M + \Omega\left(\frac{\varepsilon^{-2(F-1)}}{IF} \log\left(\frac{F}{\delta}\right)\right),$$

where $0 < \varepsilon \leq \frac{M \min\left(\frac{1}{2\sqrt{F-1}}, \frac{1}{4}\right)}{2\kappa(\mathbf{W})(1+80\kappa^2(\mathbf{W}))}$. Then, under the defined S, p and Assumption 1, CNMF-SPA outputs $\widehat{\mathbf{A}}_m$ ’s such that

$$\min_{\mathbf{\Pi}: \text{permutation}} \|\widehat{\mathbf{A}}_m \mathbf{\Pi} - \mathbf{A}_m\|_2 = O\left(\kappa^2(\mathbf{W}) \sqrt{F} \zeta\right) \quad (7)$$

for $m \in \mathcal{S}_1$ with a probability greater than or equal to $1 - \delta$, where $\zeta = \max(\sigma_{\max}(\mathbf{W}) \varepsilon, M \sqrt{I \log(1/\delta)} / \eta p \sqrt{S})$.

The proof is relegated to Appendix B. Theorem 3 asserts that using CNMF-SPA on the constructed $\widetilde{\mathbf{X}}$ in (3) approximately recovers the joint PMF—if the N, I, S are large enough. Note that if \mathbf{A}_m for all $m \in \mathcal{S}_1$ can be accurately estimated, the estimation accuracy of \mathbf{A}_n for all $n \in \mathcal{S}_2$ and $\boldsymbol{\lambda}$ can also be guaranteed and quantified, following the standard sensitivity analyses of least squares; see, e.g., [42].

Remark 2 Theorem 3 is not entirely surprising. The insight behind the proof is to recast the constructed $\widetilde{\mathbf{X}}$ as an equivalent noisy NMF model with the separability condition holding exactly on the right latent factor. This way, the noise robustness analysis for SPA [30] can be applied. The key is to quantify the noise bound in this equivalent model. Such ‘virtual noise’ is contributed by the finite sample-induced error for estimating the pairwise marginals and the violation to the separability condition. In particular, classical concentration theorems are leveraged to quantify the former, and Assumption 1 for the latter. Note that Assumption 1 is more of a working assumption, rather than an exact model for capturing reality—it helps formalize the fact that \mathbf{H} in (3) more likely attains the ε -separability condition when $|\mathcal{S}_2|$ grows [39]. In principle, if

¹Such scaling exists under mild conditions; see [41, Proposition 1].

\mathbf{A}_m is drawn from any joint absolutely continuous distribution, a similar conclusion can be reached—but using the uniform distribution simplifies the analysis.

D. More Discussions: EM Meets CNMF

As mentioned, a recent work in [11] proposed an EM algorithm for solving the maximum likelihood estimation problem under (2). To see the idea, let $f_s \in \{1, \dots, F\}$ be the realization of the ‘latent variable’ H in the s th realization of $\Pr(Z_1, \dots, Z_N)$. The joint log-likelihood of the observed data $\{\mathbf{d}_s\}_{s=1}^S$ and the corresponding latent variable $\{f_s\}_{s=1}^S$ as a function of the latent factor $\boldsymbol{\theta} = [\text{vec}(\mathbf{A}_1)^\top, \dots, \text{vec}(\mathbf{A}_N)^\top, \boldsymbol{\lambda}^\top]^\top$ can be expressed as

$$\begin{aligned} \mathcal{L}(\{\mathbf{d}_s, f_s\}_{s=1}^S; \boldsymbol{\theta}) \\ = \log \left(\prod_{s=1}^S \boldsymbol{\lambda}(f_s) \prod_{n=1}^N \prod_{i=1}^{I_n} \mathbf{A}_n(i, f_s) \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i)}) \right). \end{aligned} \quad (8)$$

The EM algorithm estimates $\boldsymbol{\theta}$ via attempting to maximize the log-likelihood. The resulting updates are simple and economical. However, the optimality of the algorithm is unclear. It was also observed in [11] that the algorithm sometimes converges to some undesired solutions if not carefully initialized.

Naturally, our idea is using CNMF-SPA to initialize the EM algorithm in [11]. As one will see, this combination oftentimes offers competitive performance. This combination also features an appealing runtime performance, since both algorithms consist of lightweight updates. We refer to this procedure as CNMF-SPA-initialized EM (CNMF-SPA-EM).

In this work, we also present performance characterizations for the EM algorithm. To proceed, we make the following assumption:

Assumption 2 Define $\overline{D}_1 = \min_{f \neq f'} \frac{1}{N} \sum_{n=1}^N p \mathbb{D}_{\text{KL}}(\mathbf{A}_n(:, f), \mathbf{A}_n(:, f'))$, $\overline{D}_2 = \frac{2}{N} \min_{f \neq f'} \log(\boldsymbol{\lambda}(f) / \boldsymbol{\lambda}(f'))$ and $\overline{D} = (\overline{D}_1 + \overline{D}_2)/2$. Assume that $\mathbf{A}_n(i, f) \geq \rho_1$ and $\boldsymbol{\lambda}(f) \geq \rho_2$ for all n, i, f , and that $\mathbf{A}_n, \boldsymbol{\lambda}$ and the initial estimates $\widehat{\mathbf{A}}_n^0, \widehat{\boldsymbol{\lambda}}^0$ satisfy

$$|\widehat{\mathbf{A}}_n^0(i, f) - \mathbf{A}_n(i, f)| \leq \delta_1 := \frac{4}{\rho_1(4 + \overline{D})}$$

$$|\widehat{\boldsymbol{\lambda}}^0(f) - \boldsymbol{\lambda}(f)| \leq \delta_2 := \frac{4}{\rho_2(4 + N\overline{D})}.$$

Here, \overline{D}_1 can be understood as a measure for the ‘conditioning’ of \mathbf{A}_n under the KL divergence—a larger \overline{D}_1 implies more diverse columns of \mathbf{A}_n , and thus a better ‘condition number’. In addition, $|\overline{D}_2|$ measures how far $\boldsymbol{\lambda}$ is away from the uniform distribution. Using Assumption 2, we show that the EM algorithm can provably improve upon reasonable initial guesses (e.g., those given by CNMF-SPA), under some conditions. To be specific, we have the following theorem:

Theorem 4 Let $\delta_{\min} = \min(\delta_1, \delta_2)$. Assume that the following hold:

$$N \geq \max\left(\frac{33 \log(3SF/\mu)}{\rho_1 \overline{D}_1}, \frac{4 \log(4SF^2/(3p\rho_2\mu))}{\overline{D}}\right),$$

$$S \geq \frac{192F^2 \log(12NFI/\mu)}{p^2 \rho_2^2 \delta_{\min}^2}, \quad \overline{D} \geq \max\left\{\frac{8 - 4\rho_1^2}{\rho_1^2}, \frac{8 - 4\rho_2^2}{N\rho_2^2}\right\}.$$

Then, under Assumption 2, the EM algorithm in [11] outputs $\hat{\mathbf{A}}_n(i, f), \hat{\boldsymbol{\lambda}}(f)$ that satisfy the following with a probability greater than or equal to $1 - \mu$:

$$|\hat{\mathbf{A}}_n(i, f) - \mathbf{A}_n(i, f)|^2 \leq \frac{48 \log(12NFI/\mu)}{Sp\lambda(f)} \leq \delta_1^2,$$

$$|\hat{\boldsymbol{\lambda}}(f) - \boldsymbol{\lambda}(f)|^2 \leq \frac{192F^2\lambda(f) \log(12NFI/\mu)}{S} \leq \delta_2^2.$$

The proof is relegated to Appendix C in the supplementary materials. Our proof extends the analysis of a different EM algorithm proposed in [43] that is designed for learning the Dawid-Skene model in crowdsourcing. The EM algorithm there effectively learns a naive Bayes model when the latent variable admits a uniform distribution, i.e., $\boldsymbol{\lambda}(f) = 1/F$ for all f . The algorithm in [43] and the EM algorithm in [11] are closely related, but the analysis in [43] does not cover the latter—since the EM in [11] does not restrict $\boldsymbol{\lambda}$ to be a uniform PMF. Extending the analysis in [43] turns out to be nontrivial, which requires careful characterization for the ‘difficulty’ introduced by $\boldsymbol{\lambda}$ ’s deviation from the uniform distribution. Our definition of \bar{D}_2 and Theorem 4 fill the gap.

V. EXPERIMENTS

In this section, we use synthetic and real data experiments to showcase the effectiveness of the proposed algorithms.

A. Baselines

We use the CTD based method in [12] and the MLE-based method in [11] as the major benchmarks, since they are general joint PMF estimation approaches as the proposed methods. Since the MLE based EM algorithm in [11] may be sensitive to initialization, an iterative optimization algorithm designed for the same MLE formulation (cf. the alternating directions (AD) algorithm in [11]) is used to initialize the EM in [11]. This baseline is denoted as the MLE-AD-EM method in our experiments.

For the real data experiments, we test the proposed approaches on two core tasks in machine learning, namely, classification and recommender systems. For the classification task, we use 4 different classifiers from the MATLAB *Statistics and Machine Learning Toolbox*, i.e., linear SVM, kernel SVM with radial basis function (RBF), multinomial logistic regression and naive Bayes classifier. For the user-item recommender systems, we use the celebrated *biased matrix factorization* (BMF) method [44] as a baseline. We also compare with results obtained by *global average* of the ratings, the *user average*, and the *item average* for predicting the missing entries in the user-item matrix.

All the algorithms are implemented in MATLAB 2018b and run on a desktop machine equipped with i7 3.40 GHZ CPU. The iterative algorithms used in the experiments are stopped when the relative change in the objective function value is less than 10^{-6} .

TABLE I
MSE & MRE PERFORMANCE ON SYNTHETIC DATA WITH
 $N = 5, F = 5, I = 10, p = 0.5, \varepsilon = 0.1$.

| Algorithms | Metric | $S = 10^3$ | $S = 10^4$ | $S = 10^5$ | $S = 10^6$ |
|------------------------------|--------|---------------|---------------|---------------|---------------|
| CNMF-SPA [Proposed] | MSE | 0.0669 | 0.0270 | 0.0210 | 0.0204 |
| CNMF-OPT [Proposed] | MSE | 0.0519 | 0.0242 | 0.0210 | 0.0204 |
| CNMF-SPA-EM [Proposed] | MSE | 0.0553 | 0.0247 | 0.0208 | 0.0204 |
| RAND-EM | MSE | 0.0734 | 0.0317 | 0.0289 | 0.0336 |
| CTD [Kargas et al.] | MSE | 0.1523 | 0.0256 | 0.0212 | 0.0205 |
| MLE-AD-EM [Yeredor & Haardt] | MSE | 0.0909 | 0.0368 | 0.0266 | 0.0296 |
| CNMF-SPA [Proposed] | MRE | 0.7965 | 0.3321 | 0.1052 | 0.0346 |
| CNMF-OPT [Proposed] | MRE | 0.6746 | 0.2327 | 0.0721 | 0.0239 |
| CNMF-SPA-EM [Proposed] | MRE | 0.6788 | 0.2341 | 0.0664 | 0.0219 |
| RAND-EM | MRE | 0.8078 | 0.2820 | 0.1845 | 0.2610 |
| CTD [Kargas et al.] | MRE | 0.9076 | 0.2920 | 0.0952 | 0.0329 |
| MLE-AD-EM [Yeredor & Haardt] | MRE | 0.8363 | 0.3552 | 0.1384 | 0.1593 |

TABLE II
MSE & MRE PERFORMANCE ON SYNTHETIC DATA WITH
 $N = 5, F = 5, I = 10, p = 0.5, \varepsilon = 0.3$.

| Algorithms | Metric | $S = 10^3$ | $S = 10^4$ | $S = 10^5$ | $S = 10^6$ |
|------------------------------|--------|---------------|---------------|---------------|---------------|
| CNMF-SPA [Proposed] | MSE | 0.0880 | 0.0288 | 0.0238 | 0.0237 |
| CNMF-OPT [Proposed] | MSE | 0.0810 | 0.0247 | 0.0204 | 0.0210 |
| CNMF-SPA-EM [Proposed] | MSE | 0.0813 | 0.0249 | 0.0206 | 0.0217 |
| RAND-EM | MSE | 0.0883 | 0.0292 | 0.0216 | 0.0276 |
| CTD [Kargas et al.] | MSE | 0.1534 | 0.0266 | 0.0204 | 0.0206 |
| MLE-AD-EM [Yeredor & Haardt] | MSE | 0.0911 | 0.0312 | 0.0313 | 0.0233 |
| CNMF-SPA [Proposed] | MRE | 0.8933 | 0.4395 | 0.3201 | 0.3128 |
| CNMF-OPT [Proposed] | MRE | 0.8010 | 0.2675 | 0.1021 | 0.0784 |
| CNMF-SPA-EM [Proposed] | MRE | 0.8133 | 0.2658 | 0.1132 | 0.1151 |
| RAND-EM | MRE | 0.8214 | 0.3080 | 0.1390 | 0.2239 |
| CTD [Kargas et al.] | MRE | 0.9101 | 0.3298 | 0.1017 | 0.0346 |
| MLE-AD-EM [Yeredor & Haardt] | MRE | 0.8752 | 0.3593 | 0.2666 | 0.1413 |

B. Synthetic-data Experiments

We consider $N = 5$ RV’s where each variable takes $I = 10$ discrete values. The entries of the conditional PMF matrices (factor matrices) $\mathbf{A}_n \in \mathbb{R}^{I_n \times F}$ and the prior probability vector $\boldsymbol{\lambda} \in \mathbb{R}^F$ are randomly generated from a uniform distribution between 0 and 1 and the columns are rescaled to have unit ℓ_1 norm with rank $F = 5$. We generate S realizations of the joint PMF by randomly hiding each variable realization with observation probability $p = 0.5$. In particular, to test the performance of CNMF-SPA under various ε -separability conditions, we fix $M = 3$ and manually ensure the ε -separability condition on \mathbf{H} in (3) by fixing $\varepsilon \in \{0.1, 0.3\}$. To run the proposed algorithms, the low-order marginals are then estimated from the realizations via sample averaging. The *mean square error* (MSE) of the factors (see [35] for definition) and the *mean relative error* (MRE) of the recovered joint PMFs (see [12]) are evaluated. We should mention that, ideally, MRE is more preferred for evaluation, but it is hard to compute (due to memory issues) for large N . MSE is an intermediate metric since it does not directly measure the recovery performance. Low MSEs imply good joint PMF recovery performance, but the converse is not necessarily true. The results are averaged from 20 random trials.

Table I-II present the MSE and MRE results for various values of ε and S , where we manually control ε when generating the \mathbf{A}_n ’s. One can see that the proposed CNMF-SPA method works reasonably well for both ε ’s under test. It performs better when ε is small, corroborating our analysis in Sec IV-C. In addition, CNMF-SPA is quite effective for initializing the EM algorithm whereas randomly initialized EM (RAND-EM) and the MLE-AD-EM method from [11] sometimes

TABLE III
MSE PERFORMANCE ON SYNTHETIC DATA WITH
 $N = 5, F = 5, I = 10, S = 100000, \varepsilon = 0.1$.

| Algorithms | Metric | $p = 0.1$ | $p = 0.05$ | $p = 0.01$ | $p = 0.005$ |
|------------------------------|--------|---------------|---------------|---------------|---------------|
| CNMF-SPA [Proposed] | MSE | 0.0370 | 0.1132 | 0.4674 | 0.4948 |
| CNMF-OPT [Proposed] | MSE | 0.0257 | 0.0564 | 0.3992 | 0.7481 |
| CNMF-SPA-EM [Proposed] | MSE | 0.0311 | 0.0816 | 0.2908 | 0.2046 |
| RAND-EM | MSE | 0.0703 | 0.0923 | 0.0978 | 0.1014 |
| CTD [Kargas <i>et al.</i>] | MSE | 0.0606 | 0.2759 | 0.4459 | 0.6922 |
| MLE-AD-EM [Yeredor & Haardt] | MSE | 0.0697 | 0.1094 | 0.0824 | 0.0862 |

struggle to attain good performance. Notably, the proposed pairwise marginals based methods consistently outperform the three-dimensional marginals based method, i.e., CTD from [12], especially when S is small. This highlights the sample complexity benefits of the proposed methods. Another point is that when CNMF-SPA does not result in good accuracy, then the follow-up EM stage (i.e., CNMF-SPA-EM) does not perform well (e.g., when $S = 10^3$). However, the proposed coupled factorization based approach, i.e., CNMF-OPT, still works well in this ‘sample-starved regime’, showing the power of recoverability guaranteed problem formulation.

It should be noted that marginal distribution based methods’ performance relies on the estimation accuracy of the marginal PMFs. Table III presents the MSEs of the algorithms under various *very* small p ’s, with $N = 5, F = 5, I = 10, S = 100,000$ and $\varepsilon = 0.1$. One can see that when $p \leq 0.05$, the triple marginal based method (i.e., the CTD method) does not produce promising results. This is because triples are rarely observed when $p = 0.05$ and thus $\widehat{\mathbf{X}}_{jkl}$ ’s are fairly noisy. The proposed method works well even when $p = 0.05$, again showing the sample complexity advantages of using pairwise marginals. When $p \leq 0.01$, both CTD and the proposed approach cannot accurately estimate the marginal distributions. In such cases, the EM algorithm that does not use higher-order co-occurrences shows its advantages.

Table IV-V show the MSE results of the algorithms tested on cases where $N = 15, I = 10$ and $F = 10$ for various values of p and S —recall that a small p means the data has many missing observations. The MRE is not presented since instantiating 15th-order tensors is prohibitive in terms of memory. The results for MLE-AD-EM is not presented since the AD stage consumes a large amount of time under this high-dimensional PMF setting. We do not manually impose ε -separability here, with the hope that the \mathbf{H} matrix is likely ε -separable under large N and I , as argued in Theorem 3 and Remark 2. One can see that the proposed CNMF-based methods exhibit better MSE performance compared to other baselines. Also, the advantage of CNMF-OPT is more articulated under challenging settings (e.g., small p and small S), where one can see that CNMF-OPT almost always outperforms the other algorithms.

Table VI shows the average runtime of the algorithms under different problem sizes. One can see that for the smaller case (cf. the left column), all the methods exhibit good runtime performance except for MLE-AD-EM. The elongated runtime of MLE-AD-EM is due to the computationally expensive AD iterations used for initialization. For the high-dimensional

TABLE IV
MSE PERFORMANCE ON SYNTHETIC DATA WITH
 $N = 15, F = 10, I = 10, p = 0.5$.

| Algorithms | Metric | $S = 10^3$ | $S = 10^4$ | $S = 10^5$ | $S = 10^6$ |
|-----------------------------|--------|---------------|---------------|---------------|---------------|
| CNMF-SPA [Proposed] | MSE | 0.1183 | 0.1030 | 0.1063 | 0.1041 |
| CNMF-OPT [Proposed] | MSE | 0.0218 | 0.0042 | 0.0022 | 0.0020 |
| CNMF-SPA-EM [Proposed] | MSE | 0.0894 | 0.0110 | 0.0056 | 0.0018 |
| RAND-EM | MSE | 0.0376 | 0.0112 | 0.0149 | 0.0069 |
| CTD [Kargas <i>et al.</i>] | MSE | 0.0329 | 0.0359 | 0.0404 | 0.0355 |

TABLE V
MSE PERFORMANCE ON SYNTHETIC DATA FOR
 $N = 15, F = 10, I = 10, p = 0.2$.

| Algorithms | Metric | $S = 10^3$ | $S = 10^4$ | $S = 10^5$ | $S = 10^6$ |
|-----------------------------|--------|---------------|---------------|---------------|---------------|
| CNMF-SPA [Proposed] | MSE | 0.2884 | 0.1181 | 0.0987 | 0.0996 |
| CNMF-OPT [Proposed] | MSE | 0.2277 | 0.0423 | 0.0062 | 0.0021 |
| CNMF-SPA-EM [Proposed] | MSE | 0.2274 | 0.0672 | 0.0165 | 0.0070 |
| RAND-EM | MSE | 0.2333 | 0.0782 | 0.0349 | 0.0313 |
| CTD [Kargas <i>et al.</i>] | MSE | 0.4015 | 0.0862 | 0.0081 | 0.0260 |

TABLE VI
AVERAGE RUNTIME (SEC.) OF THE ALGORITHMS ON SYNTHETIC
DATA $S = 10^5$ UNDER VARIOUS $[N, F, I]$ ’S.

| Algorithms | $[5, 5, 10]$ | $[15, 10, 10]$ |
|------------------------------|--------------|-----------------------------------|
| CNMF-SPA [Proposed] | 0.139 | 0.865 |
| CNMF-OPT [Proposed] | 1.633 | 94.931 |
| CNMF-SPA-EM [Proposed] | 0.490 | 4.091 |
| RAND-EM | 0.353 | 3.217 |
| CTD [Kargas <i>et al.</i>] | 0.741 | 50.559 |
| MLE-AD-EM [Yeredor & Haardt] | 124.449 | 12.949×10^3 [†] |

[†] This number is reported from a single trial due to the elongated runtime of the algorithm under this high-dimensional setting.

case where $N = 15, F = I = 10$, CNMF-SPA and EM are much faster due to their computationally economical updates. Considering the MSE and MRE results in Tables I-V, CNMF-SPA and CNMF-SPA-EM seem to provide a preferred good balance between speed and accuracy. Furthermore, the proposed CNMF-OPT offers the best accuracy for sample-starved cases—at the expense of higher computational costs.

C. Real-data Experiments: Classification Tasks

We conduct classification tasks on a number of UCI datasets (see <https://archive.ics.uci.edu>) using the joint PMF recovery algorithms and dedicated data classifiers as we detailed in Sec. V-A. The details of the datasets are given in Table VII. Note that in the datasets, some of the features are continuous-valued. To make the joint PMF recovery methods applicable, we follow the setup in [12] to discretize those continuous features. We use $I_{\text{ave}} = (1/N) \sum_{n=1}^N I_n$ to denote the averaged alphabet size of the discretized features for each dataset in Table VII ². We split each dataset into training, validation and testing sets with a ratio of 50 : 20 : 30. For our approach, we estimate the joint PMF of the features and the label using the training set, and then predict the labels on the testing data by constructing an MAP predictor following [12]. The rank F for the joint PMF recovery methods and the number of iterations needed for the methods involving EM are chosen using the validation set. We set $M = 5$ for all datasets. We perform 20 trials for each dataset with randomly

²The detailed setup and source code are available at <https://github.com/shahanaibrahimosu/joint-probability-estimation>.

TABLE VII
DETAILS OF UCI DATASETS.

| Dataset | # Samples | # Features | I_{avg} | # Classes |
|----------|-----------|------------|------------------|-----------|
| Votes | 435 | 17 | 2 | 2 |
| Car | 1728 | 7 | 4 | 4 |
| Nursery | 12960 | 8 | 4 | 4 |
| Mushroom | 8124 | 22 | 6 | 2 |

TABLE VIII
CLASSIFICATION RESULTS ON UCI DATASET ‘VOTES’.

| Algorithm | Accuracy (%) | Time (sec.) |
|------------------------------|----------------------------------|-------------|
| CNMF-SPA [Proposed] | 90.07 \pm 1.30 | 0.005 |
| CNMF-OPT [Proposed] | 94.94\pm2.13 | 9.720 |
| CNMF-SPA-EM [Proposed] | 92.82 \pm 2.17 | 0.018 |
| CTD [Kargas <i>et al.</i>] | 92.37 \pm 2.88 | 6.413 |
| MLE-AD-EM [Yeredor & Haardt] | 90.99 \pm 4.18 | 10.556 |
| SVM | 94.34 \pm 1.68 | 0.057 |
| Logistic Regression | 94.05 \pm 1.24 | 0.301 |
| SVM-RBF | 92.21 \pm 2.66 | 0.012 |
| Naive Bayes | 90.46 \pm 1.98 | 0.032 |

TABLE IX
CLASSIFICATION RESULTS ON UCI DATASET ‘CAR’.

| Algorithm | Accuracy (%) | Time (sec.) |
|------------------------------|----------------------------------|-------------|
| CNMF-SPA [Proposed] | 70.31 \pm 2.17 | 0.007 |
| CNMF-OPT [Proposed] | 85.00 \pm 1.80 | 5.550 |
| CNMF-SPA-EM [Proposed] | 87.42\pm1.12 | 0.017 |
| CTD [Kargas <i>et al.</i>] | 84.21 \pm 1.27 | 3.601 |
| MLE-AD-EM [Yeredor & Haardt] | 87.25 \pm 1.41 | 6.321 |
| SVM | 84.59 \pm 3.21 | 0.149 |
| Logistic Regression | 83.09 \pm 2.42 | 1.528 |
| SVM-RBF | 77.13 \pm 4.29 | 0.876 |
| Naive Bayes | 83.32 \pm 2.23 | 0.023 |

partitioned training/testing/validation sets, and report the mean and standard deviation of the classification accuracy.

Tables VIII-XI show the results for the UCI datasets ‘Votes’, ‘Car’, ‘Nursery’ and ‘Mushroom’, respectively. In all cases, one can see that the proposed approaches CNMF-SPA, CNMF-OPT and CNMF-SPA-EM are promising in terms of offering competitive performance for data classification. We should remark that our method is not designed for data classification, but estimating the joint PMF of the features and the label. The good performance on classification clearly support our theoretical analysis and the usefulness of the proposed algorithms.

A key observation is that the CPD method in [12] does not perform as well compared to the proposed pairwise marginals based methods, especially when the number of data samples is small (e.g., the ‘Votes’ dataset in Table VIII). This echos our motivation for this work—using pairwise marginals is advantageous over using third-order ones when fewer samples are available. In terms of runtime, the proposed CNMF-SPA and CNMF-SPA-EM methods are faster than other methods by several orders of magnitude.

D. Real-data Experiments: Recommender Systems

We test the approaches using the MovieLens 20M dataset [45]. Again, following the evaluation strategy in [12], we first round the ratings to the closest integers so that every movie’s rating resides in $\{1, 2, \dots, 5\}$. We choose different movie genres, namely, action, animation and romance, and select a subset of movies from each genre. Each subset contains

TABLE X
CLASSIFICATION RESULTS ON UCI DATASET ‘NURSERY’.

| Algorithm | Accuracy (%) | Time (sec.) |
|------------------------------|----------------------------------|-------------|
| CNMF-SPA [Proposed] | 97.48 \pm 0.20 | 0.008 |
| CNMF-OPT [Proposed] | 98.16\pm0.22 | 6.719 |
| CNMF-SPA-EM [Proposed] | 98.04 \pm 0.35 | 0.041 |
| CTD [Kargas <i>et al.</i>] | 97.70 \pm 0.22 | 4.318 |
| MLE-AD-EM [Yeredor & Haardt] | 98.04 \pm 0.33 | 8.920 |
| SVM | 97.42 \pm 0.79 | 0.706 |
| Logistic Regression | 98.02 \pm 0.15 | 5.019 |
| SVM-RBF | 70.92 \pm 4.70 | 0.738 |
| Naive Bayes | 98.03 \pm 0.20 | 0.046 |

TABLE XI
CLASSIFICATION RESULTS ON UCI DATASET ‘MUSHROOM’.

| Algorithm | Accuracy (%) | Time (sec.) |
|-----------------------------|----------------------------------|-------------|
| CNMF-SPA [Proposed] | 91.86 \pm 6.33 | 0.025 |
| CNMF-OPT [Proposed] | 96.70 \pm 0.82 | 39.349 |
| CNMF-SPA-EM [Proposed] | 99.47\pm0.68 | 0.210 |
| CTD [Kargas <i>et al.</i>] | 96.09 \pm 0.46 | 24.116 |
| SVM | 97.60 \pm 0.28 | 34.443 |
| Logistic Regression | 96.75 \pm 0.62 | 2.875 |
| SVM-RBF | 97.16 \pm 0.41 | 1.523 |
| Naive Bayes | 94.60 \pm 0.56 | 0.063 |

TABLE XII
MOVIE RECOMMENDATION RESULTS ON THE ACTION SET.

| Algorithm | RMSE | MAE | Time (s) |
|-----------------------------|-------------------------------------|-------------------------------------|----------|
| CNMF-SPA [Proposed] | 0.8497 \pm 0.0114 | 0.6663 \pm 0.0059 | 0.031 |
| CNMF-OPT [Proposed] | 0.8167 \pm 0.0035 | 0.6321 \pm 0.0040 | 70.018 |
| CNMF-SPA-EM [Proposed] | 0.7840\pm0.0025 | 0.5991\pm0.0031 | 2.424 |
| CTD [Kargas <i>et al.</i>] | 0.8770 \pm 0.0088 | 0.6649 \pm 0.0076 | 52.253 |
| BMF | 0.8011 \pm 0.0012 | 0.6260 \pm 0.0013 | 46.637 |
| Global Average | 0.9468 \pm 0.0018 | 0.6956 \pm 0.0017 | – |
| User Average | 0.8950 \pm 0.0010 | 0.6825 \pm 0.0010 | – |
| Movie Average | 0.8847 \pm 0.0018 | 0.6982 \pm 0.0012 | – |

30 movies. Hence, for every subset, $N = 30$, Z_i represents the rating of movie i , and Z_i ’s alphabet is $\{1, 2, \dots, 5\}$. We predict the rating for a movie, say, movie N , by user k via computing $\mathbb{E}[i_N | r_k(1), \dots, r_k(N-1)]$, where $r_k(i)$ denotes the rating of movie i by user k (i.e., using the MMSE estimator). This can be done via estimating $\Pr(i_1, \dots, i_N)$. Note that we leave out the MLE-AD-EM algorithm for the MovieLens datasets due to its scalability challenges.

We create the validation and testing sets by randomly hiding 20% and 30% of the dataset for each trial. The remaining 50% is used for training (i.e., learning the joint PMF in our approach). The results are taken from 20 random trials as before. In addition to the root MSE (RMSE) performance, we also report the *mean absolute error* (MAE) of the predicted ratings, which is more meaningful in the presence of outlying trials; see the definitions of RMSE and MAE in [12].

Table XII-XIV present the results for three different subsets, respectively. One can see that the proposed methods are promising—their predictions are either comparable or better than the BMF approach. Note BMF is specialized for recommender systems, while the proposed approaches are for generic joint PMF recovery. The fact that our methods perform better suggests that the underlying joint PMF is well captured by the proposed CNMF approach.

TABLE XIII
MOVIE RECOMMENDATION RESULTS ON THE ANIMATION SET.

| Algorithm | RMSE | MAE | Time (s) |
|-----------------------------|-------------------------------------|-------------------------------------|----------|
| CNMF-SPA [Proposed] | 0.8705 \pm 0.0095 | 0.6798 \pm 0.0060 | 0.028 |
| CNMF-OPT [Proposed] | 0.8124\pm0.0031 | 0.6241\pm0.0041 | 61.018 |
| CNMF-SPA-EM [Proposed] | 0.8170 \pm 0.0075 | 0.6317 \pm 0.0086 | 2.424 |
| CTD [Kargas <i>et al.</i>] | 0.8300 \pm 0.0053 | 0.6335 \pm 0.0029 | 48.253 |
| BMF | 0.8408 \pm 0.0023 | 0.6553 \pm 0.0015 | 46.637 |
| Global Average | 0.9371 \pm 0.0021 | 0.7042 \pm 0.0014 | — |
| User Average | 0.8850 \pm 0.0009 | 0.6632 \pm 0.0011 | — |
| Movie Average | 0.9027 \pm 0.0019 | 0.6900 \pm 0.0013 | — |

TABLE XIV
MOVIE RECOMMENDATION RESULTS ON THE ROMANCE SET.

| Algorithm | RMSE | MAE | Time (s) |
|-----------------------------|-------------------------------------|-------------------------------------|----------|
| CNMF-SPA [Proposed] | 0.9280 \pm 0.0066 | 0.7376 \pm 0.0076 | 0.032 |
| CNMF-OPT [Proposed] | 0.9076 \pm 0.0014 | 0.7123 \pm 0.0029 | 60.762 |
| CNMF-SPA-EM [Proposed] | 0.9057\pm0.0052 | 0.7106\pm0.0049 | 1.881 |
| CTD [Kargas <i>et al.</i>] | 0.9498 \pm 0.0085 | 0.7416 \pm 0.0054 | 47.010 |
| BMF | 0.9337 \pm 0.0007 | 0.7463 \pm 0.0009 | 31.823 |
| Global Average | 1.0019 \pm 0.0007 | 0.8078 \pm 0.0008 | — |
| User Average | 1.0195 \pm 0.0007 | 0.7862 \pm 0.0008 | — |
| Movie Average | 0.9482 \pm 0.0007 | 0.7599 \pm 0.0007 | — |

VI. CONCLUSION

We proposed a new framework for recovering joint PMF of a finite number of discrete RVs from marginal distributions. Unlike a recent approach that relies on three-dimensional marginals, our approach only uses two-dimensional marginals, which naturally has reduced-sample complexity and a lighter computational burden. We showed that under certain conditions, the recoverability of the joint PMF can be guaranteed—with only noisy pairwise marginals accessible. We proposed a coupled NMF formulation as the optimization surrogate for this task, and proposed to employ a Gram–Schmidt-like scalable algorithm as its initialization. We showed that the initialization method is effective even under the finite-sample case (whereas the existing tensor-based method does not offer sample-complexity characterization). We also showed that the Gram–Schmidt-like algorithm can provably enhance the performance of an EM algorithm that is proposed from an ML estimation perspective for the joint PMF recovery problem. We tested the proposed approach on classification tasks and recommender systems, and the results corroborate our analyses.

APPENDIX A PROOF OF THEOREM 2

The first part of the proof is reminiscent of the identifiability proof of a similar problem in [39]. The difference lies in how to handle ‘‘fat’’ A_n ’s. In [39], the A_n ’s are all square nonsingular matrices. In our case, A_n ’s can be ‘‘fat’’ matrices. This creates some more challenges.

Step 1): Under our assumption, there exists a splitting denoted by $\mathcal{S}_1 = \{\ell_1, \dots, \ell_M\}$, $\mathcal{S}_2 = \{\ell_{M+1}, \dots, \ell_Q\}$ $\mathcal{S}_1 \cup \mathcal{S}_1 \subseteq [N]$, $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$ such that

$$\widetilde{\mathbf{X}} = \begin{bmatrix} \mathbf{X}_{\ell_1, \ell_{M+1}} & \dots & \mathbf{X}_{\ell_1, \ell_Q} \\ \vdots & \vdots & \vdots \\ \mathbf{X}_{\ell_M, \ell_{M+1}} & \dots & \mathbf{X}_{\ell_M, \ell_Q} \end{bmatrix} = \mathbf{W} \mathbf{H}^\top, \quad (9)$$

where $\mathbf{W} = [\mathbf{A}_{\ell_1}^\top, \dots, \mathbf{A}_{\ell_M}^\top]^\top$ and $\mathbf{H} = [\mathbf{A}_{\ell_{M+1}}^\top, \dots, \mathbf{A}_{\ell_Q}^\top]^\top \mathbf{D}(\boldsymbol{\lambda})$. Note that $Q < N$ is allowed.

Since we have assumed that \mathbf{W} and \mathbf{H} are sufficiently scattered (since $\|\boldsymbol{\lambda}\|_0 = F$ by the assumption that $\Pr(f) \neq 0$ for all f). By Theorem 4 in [35], one can see that by solving (4), we always have $\widehat{\mathbf{W}} = \mathbf{W} \mathbf{D} \boldsymbol{\Pi}$ and $\widehat{\mathbf{H}} = \mathbf{H} \boldsymbol{\Pi} \mathbf{D}^{-1}$, where $\widehat{\mathbf{W}} = [\widehat{\mathbf{A}}_{\ell_1}^\top, \dots, \widehat{\mathbf{A}}_{\ell_M}^\top]^\top$ and $\widehat{\mathbf{H}} = [\widehat{\mathbf{A}}_{\ell_{M+1}}^\top, \dots, \widehat{\mathbf{A}}_{\ell_Q}^\top]^\top$. Since the column norms of A_n ’s are known, by column normalization of each $\widehat{\mathbf{A}}_n$ with respect to ℓ_1 -norm, we have

$$\widehat{\mathbf{A}}_n = \mathbf{A}_n \boldsymbol{\Pi}, \quad \forall n \in \mathcal{S}_1 \cup \mathcal{S}_2. \quad (10)$$

Step 2): Now we show that A_j for $j \notin \mathcal{S}_1 \cup \mathcal{S}_2$ can also be identified up to the same permutation ambiguity; i.e., we hope to show that any $\widehat{\mathbf{A}}_j$ that is a solution of (4) satisfies $\widehat{\mathbf{A}}_j = \mathbf{A}_j \boldsymbol{\Pi}$, with the same $\boldsymbol{\Pi}$ as in (10).

Let us denote $\widehat{\mathbf{A}}_n$ and $\widehat{\boldsymbol{\lambda}}$ as any optimal solution of Problem (4). Since there exists a construction (9), and $\widehat{\mathbf{A}}_{\ell_q} = \mathbf{A}_{\ell_q} \boldsymbol{\Pi}$, $q \in \{1, \dots, Q\}$ can be identified by our construction, it is easy to see that $\widehat{\boldsymbol{\lambda}} = \boldsymbol{\Pi}^\top \boldsymbol{\lambda}$. Indeed, from (9), one can see that any optimal solution of Problem (4) satisfies

$$\widetilde{\mathbf{X}} = \begin{bmatrix} \widehat{\mathbf{A}}_{\ell_1} \\ \vdots \\ \widehat{\mathbf{A}}_{\ell_M} \end{bmatrix} \mathbf{D}(\widehat{\boldsymbol{\lambda}}) \underbrace{[\widehat{\mathbf{A}}_{\ell_{M+1}}^\top, \dots, \widehat{\mathbf{A}}_{\ell_Q}^\top]}_{\mathbf{Z}_2^\top}. \quad (11)$$

Again, since \mathbf{Z}_1 and \mathbf{Z}_2 are sufficient scattered, we have $\text{rank}(\mathbf{Z}_1) = \text{rank}(\mathbf{Z}_2) = F$. Consequently, we have $\widehat{\boldsymbol{\lambda}} = (\mathbf{Z}_2 \odot \mathbf{Z}_1)^\dagger \text{vec}(\widetilde{\mathbf{X}}) = \boldsymbol{\Pi}^\top \boldsymbol{\lambda}$, due to (9) and the fact that $\text{rank}(\mathbf{Z}_2 \odot \mathbf{Z}_1) = F$, which is a result of $\text{rank}(\mathbf{Z}_1) = \text{rank}(\mathbf{Z}_2) = F$ [25]. In the above, ‘‘ \odot ’’ denotes the Khatri–Rao product.

Note that by condition iii), we have $r_1, \dots, r_T \in \mathcal{S}_1 \cup \mathcal{S}_2$ such that $\mathbf{X}_{r_t j} = \widehat{\mathbf{A}}_{r_t} \mathbf{D}(\widehat{\boldsymbol{\lambda}}) \widehat{\mathbf{A}}_j^\top$, for $t = 1, \dots, T$. These equalities can be re-expressed as follows

$$\begin{bmatrix} \mathbf{X}_{r_1 j} \\ \vdots \\ \mathbf{X}_{r_T j} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{r_1} \\ \vdots \\ \mathbf{A}_{r_T} \end{bmatrix} \mathbf{D}(\boldsymbol{\lambda}) \mathbf{A}_j^\top = \begin{bmatrix} \widehat{\mathbf{A}}_{r_1} \\ \vdots \\ \widehat{\mathbf{A}}_{r_T} \end{bmatrix} \mathbf{D}(\widehat{\boldsymbol{\lambda}}) \widehat{\mathbf{A}}_j^\top \quad (12)$$

Note that $\mathbf{D}(\widehat{\boldsymbol{\lambda}}) = \boldsymbol{\Pi}^\top \mathbf{D}(\boldsymbol{\lambda}) \boldsymbol{\Pi}^\top$, since $\widehat{\boldsymbol{\lambda}} = \boldsymbol{\Pi}^\top \boldsymbol{\lambda}$. By the assumption that any T -concatenation of A_n ’s have full column rank and $\widehat{\mathbf{A}}_{r_t} = \mathbf{A}_{r_t} \boldsymbol{\Pi}$, one can see that $\widehat{\mathbf{A}}_j = \mathbf{A}_j \boldsymbol{\Pi}$, $j \notin \mathcal{S}_1 \cup \mathcal{S}_2$. Since we have assumed that for every $j \notin \mathcal{S}_1 \cup \mathcal{S}_2$, there exists a set of r_1, \dots, r_T such that condition iii) is satisfied.

Once A_n ’s and $\boldsymbol{\lambda}$ are estimated, the joint probability $\Pr(i_1, i_2, \dots, i_N), i_n \in \{1, \dots, I_n\}$ can be estimated by

$$\begin{aligned} \Pr(i_1, i_2, \dots, i_N) &= \sum_{f=1}^F \widehat{\boldsymbol{\lambda}}(f) \prod_{n=1}^N \widehat{\mathbf{A}}_n(i_n, f) \\ &= \sum_{f=1}^F \boldsymbol{\lambda}(f) \prod_{n=1}^N \mathbf{A}_n(i_n, f). \end{aligned} \quad (13)$$

The last equality holds because the unified permutation ambiguity across A_n ’s and $\boldsymbol{\lambda}$ does not affect the reconstruction of $\Pr(i_1, i_2, \dots, i_N)$.

APPENDIX B PROOF OF THEOREM 3

Consider a matrix factorization model as below:

$$\widetilde{\mathbf{X}} = \mathbf{W}\mathbf{H}^\top, \quad (14)$$

where $\mathbf{W} \in \mathbb{R}^{L \times F}$, $\mathbf{H} \in \mathbb{R}^{K \times F}$, $\mathbf{W} \geq \mathbf{0}$ and $\mathbf{H} \geq \mathbf{0}$.

The SPA algorithm for factoring $\widetilde{\mathbf{X}}$ into \mathbf{W} and \mathbf{H} consists of two key steps:

- 1) Normalize the columns of $\widetilde{\mathbf{X}}$ w.r.t. their ℓ_1 norms.
- 2) Apply an F step Gram-Schmidt-like procedure to pick up $\Lambda = \{l_1, \dots, l_F\}$.

Note that $\widetilde{\mathbf{X}}(:, \Lambda) = \mathbf{W}\Sigma$; i.e., \mathbf{W} can be identified up to a column scaling ambiguity through this simple procedure. Gillis and Vavasis [30] have shown that under the model in (14), SPA is provably robust to noise in estimating the factor matrix \mathbf{W} (see Lemma 4 in Appendix E from [30]). To proceed, first characterize the “noise” in our virtual NMF model.

Consider the pairwise marginals \mathbf{X}_{jk} ’s which are used to construct the matrix $\widetilde{\mathbf{X}}$ in (3). \mathbf{X}_{jk} ’s are estimated by sample averaging of a finite number of realizations and thus the estimated \mathbf{X}_{jk} (denoted as $\widehat{\mathbf{X}}_{jk}$) is always noisy; i.e., we have

$$\widehat{\mathbf{X}}_{jk} = \mathbf{X}_{jk} + \mathbf{N}_{jk}, \quad (15)$$

where the noise matrix $\mathbf{N}_{jk} \in \mathbb{R}^{I \times I}$, assuming $I_n = I$ for all $n \in \{1, \dots, N\}$. We have the following proposition:

Proposition 1 Let $p \in (0, 1]$ be the probability that an RV is observed. Let S be the number of available realizations of N RVs. Assume that $p \geq (\frac{8}{S} \log(2/\delta))^{1/2}$. Then, with probability at least $1 - \delta$, $\|\mathbf{X}_{jk} - \widehat{\mathbf{X}}_{jk}\|_F = \|\mathbf{N}_{jk}\|_F \leq \phi$, holds for any distinct j, k where $\phi = \frac{\sqrt{2}(1 + \sqrt{\log(2/\delta)})}{(p\sqrt{S})}$.

The proof of Proposition 1 is given in Sec. B-A.

By the definition of the Frobenius norm, we have $\sum_{c=1}^I \|\mathbf{N}_{ij}(:, c)\|_2^2 = \|\mathbf{N}_{ij}\|_F^2 \leq \phi^2$. Applying norm equivalence $\frac{\|\mathbf{N}_{ij}(:, c)\|_1}{\sqrt{I}} \leq \|\mathbf{N}_{ij}(:, c)\|_2$, we get

$$\|\mathbf{N}_{ij}(:, c)\|_1^2 \leq I\phi^2 \implies \|\mathbf{N}_{ij}(:, c)\|_1 \leq \sqrt{I}\phi. \quad (16)$$

By using the estimates $\widehat{\mathbf{X}}_{jk}$, the model given by (3) can be represented as

$$\widehat{\mathbf{X}} = \widetilde{\mathbf{X}} + \widetilde{\mathbf{N}}, \quad (17)$$

where the (i, j) th block of $\widetilde{\mathbf{N}}$ is \mathbf{N}_{ij} . Note that $\widehat{\mathbf{X}}, \widetilde{\mathbf{X}}$ and $\widetilde{\mathbf{N}}$ all have the same size of $L \times K$. Assuming $I_n = I$ for all $n \in \{1, \dots, N\}$, we have $L = MI$ and $K = (N - M)I$. Also note that \mathbf{W} has a size of $L \times F$ and \mathbf{H} has a size of $K \times F$.

Since any column of $\widetilde{\mathbf{N}}$ is formed from the columns of M number of \mathbf{N}_{ij} ’s, we have

$$\|\widetilde{\mathbf{N}}(:, q)\|_1 \leq M\sqrt{I}\phi, \quad (18)$$

by the triangle inequality and (16).

As mentioned, before performing SPA, the columns of $\widetilde{\mathbf{X}}$ are normalized with respect to the ℓ_1 -norm. In the noiseless case, let us denote $\overline{\mathbf{X}}(:, q) = \frac{\widetilde{\mathbf{X}}(:, q)}{\|\widetilde{\mathbf{X}}(:, q)\|_1}$. Then, we have

$$\overline{\mathbf{X}}(:, q) = \sum_{f=1}^F \frac{\mathbf{W}(:, f)}{\|\mathbf{W}(:, f)\|_1} \frac{\|\mathbf{W}(:, f)\|_1 \mathbf{H}(q, f)}{\|\sum_{f=1}^F \mathbf{W}(:, f) \mathbf{H}(q, f)\|_1},$$

or equivalently $\overline{\mathbf{X}} = \overline{\mathbf{W}}\overline{\mathbf{H}}^\top$, where $\overline{\mathbf{W}}(:, f) = \frac{\mathbf{W}(:, f)}{\|\mathbf{W}(:, f)\|_1}$ and $\overline{\mathbf{H}}(q, r) = \frac{\|\mathbf{W}(:, f)\|_1 \mathbf{H}(q, f)}{\|\sum_{f=1}^F \mathbf{W}(:, f) \mathbf{H}(q, f)\|_1}$. One can verify that $\overline{\mathbf{H}}\mathbf{1} = \mathbf{1}^\top$, which is critical for applying SPA. When the data is noisy, i.e., $\widehat{\mathbf{X}} = \widetilde{\mathbf{X}} + \widetilde{\mathbf{N}}$, we hope to show that the normalized data can be represented as follows:

$$\overline{\mathbf{X}} = \overline{\mathbf{W}}\overline{\mathbf{H}}^\top + \overline{\mathbf{N}}, \quad (19)$$

where $\overline{\mathbf{X}}$ is the column normalized version (with respect to the ℓ_1 norm) of $\widetilde{\mathbf{X}}$, and $\overline{\mathbf{W}}$ and $\overline{\mathbf{H}}$ are defined as above.

From the assumption $\|\widehat{\mathbf{X}}_{ij}(:, c)\|_1 \geq \eta$ for any $i \neq j$ and $c \in \{1, \dots, I\}$, we get $\|\widetilde{\mathbf{X}}(:, q)\|_1 \geq M\eta$ for any q . Combining Lemma 5 and Eq. (18), we get

$$\|\overline{\mathbf{N}}(:, q)\|_1 \leq \frac{2\sqrt{I}\phi}{\eta}. \quad (20)$$

Applying norm equivalence, we further have $\|\overline{\mathbf{N}}(:, q)\|_2 \leq \|\overline{\mathbf{N}}(:, q)\|_1$ and hence we get

$$\|\overline{\mathbf{N}}(:, q)\|_2 \leq \frac{2\sqrt{I}\phi}{\eta}. \quad (21)$$

Fact 1 Assume that $\|\overline{\mathbf{N}}(:, q)\|_2 \leq \varphi$ for any q and that $\overline{\mathbf{H}}$ satisfies ε -separability assumption in the model (19). Suppose

$$(\sigma_{\max}(\overline{\mathbf{W}})\varepsilon + \varphi) \leq \sigma_{\min}(\overline{\mathbf{W}})\varphi\tilde{\kappa}^{-1},$$

where $\tilde{\kappa} = (1 + 80\kappa^2(\overline{\mathbf{W}}))$ and $\varrho = \min\left(\frac{1}{2\sqrt{F-1}}, \frac{1}{4}\right)$. Then, SPA identifies an index set $\widehat{\Lambda} = \{\widehat{l}_1, \dots, \widehat{l}_F\}$ such that

$$\max_{1 \leq f \leq F} \min_{\widehat{l}_f \in \widehat{\Lambda}} \left\| \overline{\mathbf{W}}(:, f) - \overline{\mathbf{X}}(:, \widehat{l}_f) \right\|_2 \leq (\sigma_{\max}(\overline{\mathbf{W}})\varepsilon + \varphi)\tilde{\kappa}, \quad (22)$$

Proof: From the assumption that $\overline{\mathbf{H}}$ satisfies ε -separability, there exists a set of indices $\Lambda = \{l_1, \dots, l_F\}$ such that $\overline{\mathbf{H}}(\Lambda, :) = \mathbf{I}_F + \mathbf{E}$, $\mathbf{E} \in \mathbb{R}^{F \times F}$ is the error matrix with $\|\mathbf{E}(l, :)\|_2 \leq \varepsilon$. and \mathbf{I}_F is the identity matrix of size $F \times F$. Without loss of generality, we assume $\Lambda = \{1, \dots, F\}$. Hence, one can write the normalized model given in (19) as

$$\begin{aligned} \overline{\mathbf{X}} &= \overline{\mathbf{W}}\overline{\mathbf{H}}^\top + \overline{\mathbf{N}} = \overline{\mathbf{W}}[\mathbf{I}_F + \mathbf{E}^\top, (\mathbf{H}^*)^\top] + \overline{\mathbf{N}} \\ &= \overline{\mathbf{W}}[\mathbf{I}_F, (\mathbf{H}^*)^\top] + [\overline{\mathbf{W}}\mathbf{E}^\top, \mathbf{0}] + \overline{\mathbf{N}}, \end{aligned}$$

where the zero matrix $\mathbf{0}$ has the same dimension as that of \mathbf{H}^* . By defining the noise matrix $\mathbf{N} \in \mathbb{R}^{L \times K}$ such that $\mathbf{N} := [\overline{\mathbf{W}}\mathbf{E}^\top, \mathbf{0}] + \overline{\mathbf{N}}$, we have $\overline{\mathbf{X}} = \overline{\mathbf{W}}[\mathbf{I}_F, (\mathbf{H}^*)^\top] + \mathbf{N}$. Then, for any $q \in \{1, \dots, K\}$, the following inequality holds:

$$\begin{aligned} \|\mathbf{N}(:, q)\|_2 &\leq \|\overline{\mathbf{W}}\|_2 \|\mathbf{E}(q, :)\|_2 + \|\overline{\mathbf{N}}(:, q)\|_2 \\ &\leq \sigma_{\max}(\overline{\mathbf{W}})\varepsilon + \varphi. \end{aligned} \quad (23)$$

Combining (23) and Lemma 4 (see Appendix E), we obtain (22) if $(\sigma_{\max}(\bar{\mathbf{W}})\varepsilon + \varphi) \leq \sigma_{\min}(\bar{\mathbf{W}})\varrho\tilde{\kappa}^{-1}$. \blacksquare

The square of the left hand side of (22) can be written as

$$\begin{aligned} & \frac{1}{M^2} \sum_{m=1}^M \max_{1 \leq f \leq F} \min_{\hat{\mathbf{A}} \in \hat{\mathbf{A}}} \left\| \mathbf{A}_{\ell_m}(:, f) - \hat{\mathbf{A}}_{\ell_m}(:, f) \right\|_2^2 \\ & \geq \frac{1}{M^2} \max_{1 \leq f \leq F} \min_{\hat{\mathbf{A}} \in \hat{\mathbf{A}}} \left\| \mathbf{A}_{\ell_m}(:, f) - \hat{\mathbf{A}}_{\ell_m}(:, f) \right\|_2^2, \end{aligned} \quad (24)$$

for any $m \in \{1, \dots, M\}$, where the first equality is due to $\bar{\mathbf{W}} = [\mathbf{A}_{\ell_1}^\top, \dots, \mathbf{A}_{\ell_M}^\top]^\top/M$, in which $\hat{\mathbf{A}}_{\ell_m}$ denotes the corresponding estimate of \mathbf{A}_{ℓ_m} .

Since $\|\mathbf{W}(:, f)\|_1 = M$ for any f , we have $\bar{\mathbf{W}} = \mathbf{W}/M$. Therefore, we have $\sigma_{\max}(\bar{\mathbf{W}}) = \sigma_{\max}(\mathbf{W})/M$, $\sigma_{\min}(\bar{\mathbf{W}}) = \sigma_{\min}(\mathbf{W})/M$,

$$\kappa(\bar{\mathbf{W}}) = \frac{\sigma_{\max}(\bar{\mathbf{W}})}{\sigma_{\min}(\bar{\mathbf{W}})} = \kappa(\mathbf{W}). \quad (25)$$

Therefore, by combining (21), (24), (25) and Fact 1, SPA estimates \mathbf{A}_{ℓ_m} for any $m \in \{1, \dots, M\}$ such that

$$\begin{aligned} & \max_{1 \leq f \leq F} \min_{\hat{\mathbf{A}} \in \hat{\mathbf{A}}} \left\| \mathbf{A}_{\ell_m}(:, f) - \hat{\mathbf{A}}_{\ell_m}(:, f) \right\|_2 \\ & \leq \left(\sigma_{\max}(\mathbf{W})\varepsilon + \frac{2M\sqrt{I}\phi}{\eta} \right) \tilde{\kappa}, \end{aligned} \quad (26)$$

if the below condition holds:

$$\frac{\sigma_{\max}(\mathbf{W})}{M}\varepsilon + \frac{2\sqrt{I}\phi}{\eta} \leq \sigma_{\min}(\mathbf{W})\varrho\tilde{\kappa}^{-1}. \quad (27)$$

Letting $\varepsilon = \frac{M\varrho}{2\kappa(\bar{\mathbf{W}})\tilde{\kappa}}$, from (27), we get the condition on ϕ as follows:

$$\phi \leq \frac{\eta\sigma_{\max}(\mathbf{W})\varepsilon}{4M\sqrt{I}}. \quad (28)$$

From Proposition 1, we have $\phi \leq \frac{\sqrt{2}(1+\sqrt{\log(2/\delta)})}{p\sqrt{S}}$ with probability greater than or equal to $1 - \delta$. Combining with (28), we get the number of realizations S required to get the estimation error bound (26) as below:

$$S \geq \frac{32M^2I(1 + \sqrt{\log(2/\delta)})^2}{\sigma_{\max}^2(\mathbf{W})\eta^2\varepsilon^2p^2}.$$

Note that Lemma 1 holds if $\bar{\mathbf{H}}$ satisfies ε -separability condition. By combining Assumption 1 and Lemma 6, we get that $\bar{\mathbf{H}}$ satisfies ε -separability assumption with probability greater than $1 - \rho$, if

$$(N - M)I = \Omega\left(\frac{\varepsilon^{-2(F-1)}}{F} \log\left(\frac{F}{\rho}\right)\right). \quad (29)$$

By substituting ϕ in (26) and using the fact that for any matrix $\mathbf{A} \in \mathbb{R}^{I \times F}$, the matrix 2-norm $\|\mathbf{A}\|_2 \leq \sqrt{F} \max_{1 \leq f \leq F} \|\mathbf{A}(:, f)\|_2$, we get the result (7) in the theorem. Letting $\rho = \frac{\delta}{2}$ in (29), the proof is completed.

A. Proof Proposition 1

Recall $\mathbf{d}_s \in \mathbb{R}^N$ denotes the s th realization of the joint PMF $\Pr(Z_1, \dots, Z_N)$. For simplicity, we assume that 0 does not belong to the alphabets of Z_1, \dots, Z_N , and we use the notation $\mathbf{d}_s(j) = 0$ to represent that ' Z_j is not observed in the s th realization'.

For S realizations of the joint PMF, i.e., $\{\mathbf{d}_s\}_{s=1}^S$, the sample averaging expressions for estimating \mathbf{X}_{jk} is defined as follows:

$$\widehat{\mathbf{X}}_{jk}(i_j, i_k) = \frac{1}{|\mathcal{S}_{jk}|} \sum_{s \in \mathcal{S}_{jk}} \mathbb{I}[\mathbf{d}_s(j) = z_j^{(i_j)}, \mathbf{d}_s(k) = z_k^{(i_k)}],$$

where $\mathcal{S}_{jk} = \{s \mid \mathbb{I}[\mathbf{d}_s(j) \neq 0, \mathbf{d}_s(k) \neq 0]\}$.

Let us construct a random variable $V_{j,s}$, where $V_{j,s} = 1$ if Z_j is observed in \mathbf{d}_s ; otherwise $V_{j,s} = 0$. With this definition, we can construct a derived RV S_{jk} that is sum of S i.i.d. Bernoulli random variables such that $S_{jk} = \sum_{s=1}^S \mathbb{I}[V_{j,s} = 1]$ and $V_{k,s} = 1$, where we have $\mathbb{E}[S_{jk}] = Sp^2$.

In order to characterize the random variable S_{jk} , we can use Chernoff lower tail bound such that for $0 < t < 1$,

$$\Pr[S_{jk} \geq (1 - t)Sp^2] \geq 1 - e^{-Sp^2t^2/2}. \quad (30)$$

Combining Lemma 19 in [43] and (30), we have

$$\begin{aligned} & \Pr\left[\|\widehat{\mathbf{X}}_{jk} - \mathbf{X}_{jk}\|_{\text{F}} \leq \frac{(1 + \sqrt{\log(1/\delta)})}{\sqrt{(1-t)Sp^2}}\right] \\ & = \Pr\left[\|\widehat{\mathbf{X}}_{jk} - \mathbf{X}_{jk}\|_{\text{F}} \leq \frac{(1 + \sqrt{\log(1/\delta)})}{\sqrt{S_{jk}}}, S_{jk} \geq (1-t)Sp^2\right] \\ & \geq 1 - \delta - e^{-Sp^2t^2/2}, \end{aligned} \quad (31)$$

where we have applied the De Morgan's law and the union bound to obtain the last inequality.

Since (31) holds for any $t \in (0, 1)$, we set $t = 1/2$ for expression simplicity. Then, we have

$$\begin{aligned} & \Pr\left[\|\widehat{\mathbf{X}}_{jk} - \mathbf{X}_{jk}\|_{\text{F}} \leq \frac{\sqrt{2}(1 + \sqrt{\log(1/\delta)})}{p\sqrt{S}}\right] \\ & \geq 1 - \delta - e^{-Sp^2/8}. \end{aligned} \quad (32)$$

It follows that if $p^2 \geq \frac{8}{S} \log(1/\delta)$, the right hand side of (32) is greater than $1 - 2\delta$.

APPENDIX C PROOF OF THEOREM 4

We first introduce the EM algorithm in [11]. Then, we will show that the EM algorithm improves upon good initializations (e.g., those given by the CNMF-SPA).

A. An EM Algorithm for Joint PMF Learning [11]

The EM algorithm proposed in [11] for handling maximizing the log-likelihood in (8) has the following E-step and M-step:

E-step: The posterior of the latent variable $q_{s,f} = \Pr(f_s = f \mid \mathbf{d}_1, \dots, \mathbf{d}_S, \hat{\boldsymbol{\theta}})$ is estimated via $\hat{q}_{s,f} = \frac{\exp(\log(\hat{\lambda}(f)) + \sum_{n=1}^N \sum_{i=1}^{I_n} \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i)}) \log(\hat{A}_n(i, f)))}{\sum_{f'=1}^F \exp(\log(\hat{\lambda}(f')) + \sum_{n=1}^N \sum_{i=1}^{I_n} \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i)}) \log(\hat{A}_n(i, f')))}$.

M-step: Using the estimated $\hat{q}_{s,f}$, $\hat{\mathbf{A}}_n$ and $\hat{\boldsymbol{\lambda}}$ that maximize the likelihood function are computed using the following:

$$\hat{\mathbf{A}}_n(i, f) \leftarrow \frac{\sum_{s=1}^S \hat{q}_{s,f} \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i)})}{\sum_{i'=1}^{I_n} \sum_{s=1}^S \hat{q}_{s,f} \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i')}), \quad (33a)$$

$$\hat{\boldsymbol{\lambda}}(f) \leftarrow \frac{\sum_{s=1}^S \hat{q}_{s,f}}{\sum_{f'=1}^F \sum_{s=1}^S \hat{q}_{s,f'}}. \quad (33b)$$

B. Performance Analysis for EM

We employ the paradigm in [43] for proving the optimality of an EM algorithm. There, the EM algorithm learns a naive Bayes model as in (2) but assuming a uniform latent distribution, i.e., $\boldsymbol{\lambda}(f) = 1/F$. In our case, we do not assume uniform prior for $\boldsymbol{\lambda}$ and thus our proof covers more general cases. To begin with, we follow the idea in [43] to define a number of events as below:

$$\begin{aligned} \mathcal{E}_1 : & \sum_{n=1}^N \sum_{i=1}^{I_n} \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i)}) \log \left(\frac{\mathbf{A}_n(i, f_s)}{\mathbf{A}_n(i, f)} \right) \geq N \bar{D}_1 / 2, \\ \mathcal{E}_2 : & \left| \sum_{s=1}^S \mathbb{I}(f_s = f) \mathbb{I}(\mathbf{d}_s(n) = z_n^{(i)}) - S \boldsymbol{\lambda}(f) p \mathbf{A}_n(i, f) \right| \leq S t_{nif}, \\ \mathcal{E}_3 : & \left| \sum_{s=1}^S \mathbb{I}(f_s = f) \mathbb{I}(\mathbf{d}_s(n) \neq 0) - S \boldsymbol{\lambda}(f) p \right| \leq S t_{nif}, \\ \mathcal{E}_4 : & \left| \sum_{s=1}^S \mathbb{I}(f_s = f) - S \boldsymbol{\lambda}(f) \right| \leq S c_f, \end{aligned}$$

where the events are defined for all f, n, i, f_s , $\mathbf{d}_s(n) \neq 0$ represents that n -th RV is observed with any value from its alphabet set in the s -th sample, and $t_{nif} > 0, c_f > 0$ are scalars. Note that \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 are the same as those defined in [43], while \mathcal{E}_4 is introduced in our proof to accommodate the general $\boldsymbol{\lambda}$.

First, we consider the E-step. The parameter $\hat{q}_{s,f}$ can be bounded using the following lemma:

Lemma 1 Assume that the event \mathcal{E}_1 happens and also assume that $\mathbf{A}_n, \boldsymbol{\lambda}$ and the initial estimates satisfy $|\hat{\mathbf{A}}_n(i, f) - \mathbf{A}_n(i, f)| \leq \delta_1$, $\mathbf{A}_n(i, f) \geq \rho_1$, $|\hat{\boldsymbol{\lambda}}(f) - \boldsymbol{\lambda}(f)| \leq \delta_2$ and $\boldsymbol{\lambda}(f) \geq \rho_2$ for all n, i, f . Also assume that $\delta_1 < \rho_1$ and $\delta_2 < \rho_2$. Then, $\hat{q}_{s,f}$ satisfies the following:

$$|\hat{q}_{s,f} - \mathbb{I}(f_s = f)| \leq v, \quad \forall f, s. \quad (34)$$

where $v = \exp(-(N \bar{D}_1 + 1 - \frac{1}{\rho_2(\rho_2 - \delta_2)} + N(1 - \frac{1}{\rho_1(\rho_1 - \delta_1)})) + \log(F))$.

The proof of Lemma 1 is given in Appendix D-A. The next lemma shows that the subsequent M-step estimates the \mathbf{A}_n 's and $\boldsymbol{\lambda}$ up to bounded errors:

Lemma 2 Assume that $\mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4$ holds. Suppose $\hat{q}_{s,f}$ satisfies the following:

$$|\hat{q}_{s,f} - \mathbb{I}(f_s = f)| \leq \beta, \quad \forall f, s, \quad (35)$$

where $\beta > 0$ is a scalar. Then $\hat{\mathbf{A}}_n$ and $\hat{\boldsymbol{\lambda}}$ updated by (33) are bounded by:

$$|\hat{\mathbf{A}}_n(i, f) - \mathbf{A}_n(i, f)| \leq \frac{2t_{nif} + 2S\beta}{\lambda(f)p - t_{nif} - \beta}, \quad (36a)$$

$$|\hat{\boldsymbol{\lambda}}(f) - \boldsymbol{\lambda}(f)| \leq \frac{c_f + \beta + F\beta}{1 - F\beta}. \quad (36b)$$

The proof of Lemma 2 is given in Appendix D-B. To proceed, we have the following lemma:

Lemma 3 Assume that $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4$ happens. Also assume that $|\hat{\mathbf{A}}_n^0(i, f) - \mathbf{A}_n(i, f)| \leq \delta_1 := \frac{4}{\rho_1(4 + \bar{D})}$, $\mathbf{A}_n(i, f) \geq \rho_1$, $|\hat{\boldsymbol{\lambda}}^0(f) - \boldsymbol{\lambda}(f)| \leq \delta_2 := \frac{4}{\rho_2(4 + \bar{N}\bar{D})}$, $\boldsymbol{\lambda}(f) \geq \rho_2$ for all n, i, f and $\bar{D} \geq \max \left\{ \frac{8-4\rho_1^2}{\rho_1^2}, \frac{8-4\rho_2^2}{N\rho_2^2} \right\}$. Suppose that the following holds $\forall g \in \{t_{nif}\}_{n,i,f}, \{c_f\}_f$:

$$2 \exp(-\frac{N\bar{D}}{2} + \log(F)) \leq g \leq \frac{p\rho_2}{8F} \xi, \quad (37)$$

where $\xi = \min \left(\frac{4}{\rho_1(4 + \bar{D})}, \frac{4}{\rho_2(4 + \bar{N}\bar{D})} \right)$. Then, by updating the parameters using EM at least once (i.e., after running the EM algorithm for at least one iteration), we have the following:

$$|\hat{\mathbf{A}}_n(i, f) - \mathbf{A}_n(i, f)| \leq \frac{4t_{nif}}{\lambda(f)p} \leq \delta_1, \quad (38a)$$

$$|\hat{\boldsymbol{\lambda}}(f) - \boldsymbol{\lambda}(f)| \leq 8Fc_f \leq \delta_2. \quad (38b)$$

The proof of Lemma 3 is given in Appendix D-C.

Let us characterize the probability of the intersection $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4$ happening. Theorem 4 in [43] characterizes the probabilities of \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 happening. Specifically, we get the following results from [43]:

$$\Pr(\mathcal{E}_1) \geq 1 - SF \exp \left(-\frac{N\bar{D}_1}{33 \log(1/\rho_1)} \right), \quad (39a)$$

$$\Pr(\mathcal{E}_2) \geq 1 - \sum_{n=1}^N \sum_{f=1}^F \sum_{i=1}^{I_n} 2 \exp \left(-\frac{St_{nif}^2}{3p\lambda(f)} \right), \quad (39b)$$

$$\Pr(\mathcal{E}_3) \geq 1 - \sum_{n=1}^N \sum_{f=1}^F \sum_{i=1}^{I_n} 2 \exp \left(-\frac{St_{nif}^2}{3p\lambda(f)} \right). \quad (39c)$$

In order to characterize $\Pr(\mathcal{E}_4)$, we observe that $\sum_{s=1}^S \mathbb{I}(f_s = f)$ is the sum of i.i.d. Bernoulli random variables. Using the Chernoff bound, for a particular f , we have

$$\Pr \left[\left| \sum_{s=1}^S \mathbb{I}(f_s = f) - S\lambda(f) \right| \geq Sc_f \right] \leq 2e^{-\frac{Sc_f^2}{3\lambda(f)}}. \quad (40)$$

By taking the union bound over all $f \in \{1, \dots, F\}$, we obtain

$$\Pr(\mathcal{E}_4) \geq 1 - \sum_{f=1}^F 2 \exp(-Sc_f^2/(3\lambda(f))). \quad (41)$$

Applying the union bound and the De Morgan's law, one can see that $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4$ happens with a probability $\geq 1 - SFe^{-\frac{N\bar{D}_1}{33 \log(1/\rho_1)}} - \sum_{n=1}^N \sum_{f=1}^F \sum_{i=1}^{I_n} 4e^{-\frac{St_{nif}^2}{3p\lambda(f)}} - \sum_{f=1}^F 2e^{-\frac{Sc_f^2}{3\lambda(f)}}$.

To ensure that the estimation error bounds for \mathbf{A}_n and $\boldsymbol{\lambda}$ given by (38) hold with probability greater than $1 - \epsilon$, a sufficient condition is that the following being satisfied simultaneously:

$$N \geq \frac{33 \log(1/\rho_1) \log(3SF/\epsilon)}{\bar{D}_1} \quad (42)$$

$$S \geq \frac{3p\boldsymbol{\lambda}(f) \log(12NFI/\epsilon)}{t_{nif}^2} \quad (43)$$

$$S \geq \frac{3\boldsymbol{\lambda}(f) \log(6F/\epsilon)}{c_f^2}. \quad (44)$$

We can assign specific values to t_{nif} and c_f such that the above conditions are satisfied. Let

$$t_{nif} := \sqrt{\frac{3p\boldsymbol{\lambda}(f) \log(12NFI/\epsilon)}{S}}, \quad (45a)$$

$$c_f := \sqrt{\frac{3\boldsymbol{\lambda}(f) \log(12NFI/\epsilon)}{S}}. \quad (45b)$$

By this selection of t_{nif} and c_f , the conditions in (43) and (44) hold. To enforce the condition (37), the following equalities have to hold:

$$\begin{aligned} \sqrt{\frac{3p\boldsymbol{\lambda}(f) \log(12NFI/\epsilon)}{S}} &\geq 2 \exp\left(-\frac{\bar{N}\bar{D}}{2} + \log(F)\right) \\ \sqrt{\frac{3\boldsymbol{\lambda}(f) \log(12NFI/\epsilon)}{S}} &\leq \frac{p\rho_2}{8F} \xi = \frac{p\rho_2\delta_{\min}}{8F}, \end{aligned}$$

where $\delta_{\min} = \min(\delta_1, \delta_2)$. The above can be implied by the following:

$$N \geq \frac{4 \log(2SF^2/(3p\rho_2 \log(12NFI/\epsilon)))}{\bar{D}} \quad (46)$$

$$S \geq \frac{192F^2 \log(12NFI/\epsilon)}{p^2 \rho_2^2 \delta_{\min}^2}, \quad (47)$$

where we have used $1 \geq \boldsymbol{\lambda}(f) \geq \rho_2$.

Using the inequality $\log x > 1 - \frac{1}{x}$, $x > 0$, we can express the condition (46) as

$$N \geq \frac{4 \log(2SF^2/(3p\rho_2(1 - \epsilon/(12NFI))))}{\bar{D}}$$

Since $\epsilon \leq 1$ and the product $NFI > 1$, we always have $1 - \epsilon/(12NFI) > \epsilon/2$, which makes the above condition as follows:

$$N \geq \frac{4 \log(4SF^2/(3p\rho_2\epsilon))}{\bar{D}}. \quad (48)$$

Combing the two conditions (42) and (48), we have

$$N \geq \max\left(\frac{33 \log(3SF/\epsilon)}{\rho_1 \bar{D}_1}, \frac{4 \log(4SF^2/(3p\rho_2\epsilon))}{\bar{D}}\right), \quad (49)$$

where we have used the fact that $\log(1/\rho_1) \leq (1/\rho_1) - 1 < 1/\rho_1$.

To summarize, if (47) and (49) hold and t_{nif} and c_f are chosen to be as in (45), then, with probability at least $1 - \epsilon$, the following inequalities hold by Lemma 3:

$$|\hat{\mathbf{A}}_n(i, f) - \mathbf{A}_n(i, f)|^2 \leq \frac{16t_{nif}^2}{p^2 \boldsymbol{\lambda}(f)^2} \leq \frac{48 \log(12NFI/\epsilon)}{Sp\boldsymbol{\lambda}(f)}$$

$$|\hat{\boldsymbol{\lambda}}(f) - \boldsymbol{\lambda}(f)|^2 \leq 64F^2 c_f^2 \leq \frac{192F^2 \boldsymbol{\lambda}(f) \log(12NFI/\epsilon)}{S}.$$

This completes the proof.

REFERENCES

- [1] H. L. Van Trees, K. L. Bell, and Z. Tian, *Detection Estimation and Modulation Theory, Part I: Detection, Estimation, and Filtering Theory, 2nd Edition*. John Wiley and Sons, New York, NY, 2013.
- [2] Y. Ephraim, "A Bayesian estimation approach for speech enhancement using hidden Markov models," *IEEE Trans. Signal Process.*, vol. 40, no. 4, pp. 725–735, 1992.
- [3] G. Colavolpe and A. Barbieri, "On MAP symbol detection for ISI channels using the Ungerboeck observation model," *IEEE Commun. Lett.*, vol. 9, no. 8, pp. 720–722, 2005.
- [4] S. Rangan, A. K. Fletcher, and V. K. Goyal, "Asymptotic analysis of MAP estimation via the replica method and applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1902–1923, 2012.
- [5] A. A. Farag, R. M. Mohamed, and A. El-Baz, "A unified framework for MAP estimation in remote sensing image segmentation," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 7, pp. 1617–1634, 2005.
- [6] T. M. Cover, *Elements of information theory*. John Wiley & Sons, 1999.
- [7] A. Taleb and C. Jutten, "Source separation in post-nonlinear mixtures," *IEEE Trans. Signal Process.*, vol. 47, no. 10, pp. 2807–2820, 1999.
- [8] P. Thevenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Trans. Image Process.*, vol. 9, no. 12, pp. 2083–2099, 2000.
- [9] Y. Yang and R. S. Blum, "MIMO radar waveform design based on mutual information and minimum mean-square error estimation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 1, pp. 330–343, 2007.
- [10] K. P. Murphy, *Machine learning: A probabilistic perspective*. MIT press, 2012.
- [11] A. Yeredor and M. Haardt, "Maximum likelihood estimation of a low-rank probability mass tensor from partial observations," *IEEE Signal Process. Lett.*, vol. 26, no. 10, pp. 1551–1555, Oct 2019.
- [12] N. Kargas, N. D. Sidiropoulos, and X. Fu, "Tensors, learning, and 'Kolmogorov extension' for finite-alphabet random vectors," *IEEE Trans. Signal Process.*, vol. 66, no. 18, pp. 4854–4868, 2018.
- [13] O. Dabeer, "Joint probability mass function estimation from asynchronous samples," *IEEE Trans. Signal Process.*, vol. 61, no. 2, pp. 355–364, 2013.
- [14] Hunsop Hong and D. Schonfeld, "A new approach to constrained expectation-maximization for density estimation," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2008, pp. 3689–3692.
- [15] Jenq-Neng Hwang, Shyh-Rong Lay, and A. Lippman, "Nonparametric multivariate density estimation: a comparative study," *IEEE Trans. Signal Process.*, vol. 42, no. 10, pp. 2795–2810, 1994.
- [16] M. J. Wainwright, *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019, vol. 48.
- [17] S. M. Kay, *Fundamentals of statistical signal processing*. Prentice Hall PTR, 1993.
- [18] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [19] M. I. Jordan, *Learning in graphical models*. Springer Science & Business Media, 1998, vol. 89.
- [20] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, Sep. 2001.
- [21] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.
- [22] C. J. Hillar and L.-H. Lim, "Most tensor problems are NP-hard," *Journal of the ACM (JACM)*, vol. 60, no. 6, p. 45, 2013.
- [23] S. Ibrahim and X. Fu, "Recovering joint PMF from pairwise marginals," in *54th Asilomar Conference on Signals, Systems and Computers*, Oct. 2020.
- [24] R. Harshman, "Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multi-modal factor analysis," *UCLA Working Papers in Phonetics*, vol. 16, 1970.
- [25] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, "Tensor decomposition for signal processing and machine learning," *IEEE Trans. Signal Process.*, vol. 65, no. 13, pp. 3551–3582, 2017.
- [26] Y. Han, J. Jiao, and T. Weissman, "Minimax estimation of discrete distributions under ℓ_1 loss," *IEEE Trans. Info. Theory*, vol. 61, no. 11, pp. 6434–6354, 2015.
- [27] X. Fu, K. Huang, N. D. Sidiropoulos, and W.-K. Ma, "Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications," *IEEE Signal Process. Mag.*, vol. 36, no. 2, pp. 59–80, March 2019.

- [28] N. Gillis, "The why and how of nonnegative matrix factorization," *Regularization, Optimization, Kernels, and Support Vector Machines*, vol. 12, p. 257, 2014.
- [29] D. Donoho and V. Stodden, "When does non-negative matrix factorization give a correct decomposition into parts?" in *Advances in Neural Information Processing Systems*, 2003, pp. 1141–1148.
- [30] N. Gillis and S. Vavasis, "Fast and robust recursive algorithms for separable nonnegative matrix factorization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 4, pp. 698–714, April 2014.
- [31] X. Fu, W.-K. Ma, T.-H. Chan, and J. M. Bioucas-Dias, "Self-dictionary sparse regression for hyperspectral unmixing: Greedy pursuit and pure pixel search are related," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 6, pp. 1128–1141, 2015.
- [32] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, and M. Zhu, "A practical algorithm for topic modeling with provable guarantees," in *Proceedings of International Conference on Machine Learning*, 2013.
- [33] T.-H. Chan, W.-K. Ma, A. Ambikapathi, and C.-Y. Chi, "A simplex volume maximization framework for hyperspectral endmember extraction," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 11, pp. 4177–4193, Nov. 2011.
- [34] U. Araújo, B. Saldanha, R. Galvão, T. Yoneyama, H. Chame, and V. Visani, "The successive projections algorithm for variable selection in spectroscopic multicomponent analysis," *Chemometrics and Intelligent Laboratory Systems*, vol. 57, no. 2, pp. 65–73, 2001.
- [35] K. Huang, N. Sidiropoulos, and A. Swami, "Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition," *IEEE Trans. Signal Process.*, vol. 62, no. 1, pp. 211–224, 2014.
- [36] X. Fu, W.-K. Ma, K. Huang, and N. D. Sidiropoulos, "Blind separation of quasi-stationary sources: Exploiting convex geometry in covariance domain," *IEEE Trans. Signal Process.*, vol. 63, no. 9, pp. 2306–2320, May 2015.
- [37] X. Fu, K. Huang, B. Yang, W. Ma, and N. D. Sidiropoulos, "Robust volume minimization-based matrix factorization for remote sensing and document clustering," *IEEE Trans. Signal Process.*, vol. 64, no. 23, pp. 6254–6268, Dec 2016.
- [38] X. Fu, K. Huang, and N. D. Sidiropoulos, "On identifiability of nonnegative matrix factorization," *IEEE Signal Process. Lett.*, vol. 25, no. 3, pp. 328–332, 2018.
- [39] S. Ibrahim, X. Fu, N. Kargas, and K. Huang, "Crowdsourcing via pairwise co-occurrences: Identifiability and algorithms," in *Advances in Neural Information Processing Systems*, 2019, pp. 7845–7855.
- [40] N. Kargas and N. D. Sidiropoulos, "Learning mixtures of smooth product distributions: Identifiability and algorithm," in *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019, pp. 388–396.
- [41] B. Yang, X. Fu, N. D. Sidiropoulos, and K. Huang, "Learning nonlinear mixtures: Identifiability and algorithm," *IEEE Trans. Signal Process.*, vol. 68, pp. 2857–2869, 2020.
- [42] P.-Å. Wedin, "Perturbation theory for pseudo-inverses," *BIT Numerical Mathematics*, vol. 13, no. 2, pp. 217–232, Jun 1973.
- [43] Y. Zhang, X. Chen, D. Zhou, and M. I. Jordan, "Spectral methods meet EM: A provably optimal algorithm for crowdsourcing," *J. Mach. Learn. Res.*, vol. 17, no. 102, pp. 1–44, 2016.
- [44] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, Aug. 2009.
- [45] F. M. Harper and J. A. Konstan, "The MovieLens datasets: History and context," *ACM Trans. Interact. Intell. Syst.*, vol. 5, no. 4, pp. 19:1–19:19, Dec. 2015.



Xiao Fu (Senior Member, IEEE) received the B.Eng. and MSc. degrees from the University of Electronic Science and Technology of China (UESTC), Chengdu, China, in 2005 and 2010, respectively. He received the Ph.D. degree in Electronic Engineering from The Chinese University of Hong Kong (CUHK), Shatin, N.T., Hong Kong, in 2014. He was a Postdoctoral Associate with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA, from 2014 to 2017. Since 2017, he has been an Assistant Professor with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA. His research interests include the broad area of signal processing and machine learning.

Dr. Fu received a Best Student Paper Award at ICASSP 2014, and was a recipient of the Outstanding Postdoctoral Scholar Award at University of Minnesota in 2016. His coauthored papers that received Best Student Paper Awards from IEEE CAMSAP 2015 and IEEE MLSP 2019, respectively. He serves as a member of the Sensor Array and Multichannel Technical Committee (SAM-TC) of the IEEE Signal Processing Society (SPS). He is also a member of the Signal Processing for Multisensor Systems Technical Area Committee (SPMuS-TAC) of EURASIP. He is the Treasurer of the IEEE SPS Oregon Chapter. He serves as an Editor of SIGNAL PROCESSING. He was a tutorial speaker at ICASSP 2017 and SIAM Conference on Applied Linear Algebra 2021.



Shahana Ibrahim received the B.Tech. degree in electronics and communication engineering from National Institute of Technology, Calicut, India, in 2012, and the M.S. degree in electrical engineering in 2019 from Oregon State University, Corvallis, OR, USA, where she is currently working toward the Ph.D. degree. She had been working as a System Validation Engineer with Texas Instruments, Bengaluru, India, from 2012 to 2017. Her research interests are in the broad area of statistical machine learning and signal processing.