
Crowdsourcing via Pairwise Co-occurrences: Identifiability and Algorithms

Shahana Ibrahim

School of Elec. Eng. & Computer Sci.
Oregon State University
Corvallis, OR 97331
ibrahish@oregonstate.edu

Xiao Fu*

School of Elec. Eng. & Computer Sci.
Oregon State University
Corvallis, OR 97331
xiao.fu@oregonstate.edu

Nikos Kargas

Department of Elec. & Computer Eng.
University of Minnesota
Minneapolis, MN 55455
kaga005@umn.edu

Kejun Huang

Department of Computing & Info. Sci. & Eng.
University of Florida
Gainesville, FL 32611
kejun.huang@ufl.edu

Abstract

The data deluge comes with high demands for data labeling. Crowdsourcing (or, more generally, ensemble learning) techniques aim to produce accurate labels via integrating noisy, non-expert labeling from annotators. The classic Dawid-Skene estimator and its accompanying expectation maximization (EM) algorithm have been widely used, but the theoretical properties are not fully understood. Tensor methods were proposed to guarantee identification of the Dawid-Skene model, but the sample complexity is a hurdle for applying such approaches—since the tensor methods hinge on the availability of third-order statistics that are hard to reliably estimate given limited data. In this paper, we propose a framework using pairwise co-occurrences of the annotator responses, which naturally admits lower sample complexity. We show that the approach can identify the Dawid-Skene model under realistic conditions. We propose an algebraic algorithm reminiscent of convex geometry-based structured matrix factorization to solve the model identification problem efficiently, and an identifiability-enhanced algorithm for handling more challenging and critical scenarios. Experiments show that the proposed algorithms outperform the state-of-art algorithms under a variety of scenarios.

1 Introduction

Background. The drastically increasing availability of data has successfully enabled many timely applications in machine learning and artificial intelligence. At the same time, most supervised learning tasks, e.g., the core tasks in computer vision, natural language processing, and speech processing, heavily rely on labeled data. However, labeling data is not a trivial task—it requires educated and knowledgeable annotators (which could be human workers or machine classifiers), to work under a reliable way. More importantly, it needs an effective mechanism to integrate the possibly different labeling from multiple annotators. Techniques addressing this problem in machine learning are called *crowdsourcing* [25] or more generally, *ensemble learning* [9].

*The work is supported in part by the National Science Foundation under projects ECCS 1808159 and NSF ECCS 1608961, and by the Army Research Office (ARO) under projects ARO W911NF-19-1-0247 and ARO W911NF-19-1-0407.

Crowdsourcing has a long history in machine learning, which can be traced back to the 1970s [7]. Many models and methods have appeared since then [23, 24, 22, 35, 39, 29, 38]. Intuitively, if a number of reliable annotators label the same data samples, then *majority voting* among the annotators is expected to work well. However, in practice, not all the annotators are equally reliable—e.g., different annotators could be specialized for recognizing different classes. In addition, not all the annotators are labeling all the data samples, since data samples are often dispatched to different groups of annotators in a certain way. Under such circumstances, majority voting is not very promising.

A more sophisticated way is to treat the crowdsourcing problem as a model identification problem. The arguably most popular generative model in crowdsourcing is the Dawid-Skene model [7], where every annotator is assigned with a ‘confusion matrix’ that decides the probability of an annotator giving class label ℓ when the ground-truth label is g . If such confusion matrices and the probability mass function (PMF) of the ground-truth label can be identified, then a maximum likelihood (ML) or a maximum *a posteriori* (MAP) estimator for the true label of any given sample can be constructed. The Dawid-Skene model is quite simple and succinct, and some of the model assumptions (e.g., the conditional independence of the annotator responses) are actually debatable. Nonetheless, this model has been proven very useful in practice [31, 38, 15, 24, 29, 40].

Theoretical aspects for the Dawid-Skene model, however, are less well understood. In particular, it had been unclear if the model could be identified via the accompanying *expectation maximization* (EM) algorithm proposed in the same paper [7], until some recent works addressing certain special cases [24]. The works in [38, 40] put forth tensor methods for learning the Dawid-Skene model. These methods admit model identifiability, and also can be used to effectively initialize the classic EM algorithm provably [40]. The challenge is that tensor methods utilize third-order statistics of the data samples, which are rather hard to estimate reliably in practice given limited data [20].

Contributions. In this work, we propose an alternative for identifying the Dawid-Skene model, without using third-order statistics. Our approach is based on utilizing the pairwise co-occurrences of annotators’ responses to data samples—which are second-order statistics and thus are naturally much easier to estimate compared to the third-order ones. We show that, by judiciously combining the co-occurrences between different annotator pairs, the confusion matrices and the ground-truth label’s prior PMF can be provably identified, under realistic conditions (e.g., when there exists a relatively well-trained annotator among all annotators). This is reminiscent of nonnegative matrix theory and convex geometry [14, 16]. Our approach is also naturally robust to spammers as well as scenarios where every annotator only labels partial data. We offer two algorithms under the same framework. The first algorithm is algebraic, and thus is efficient and suitable for handling very large-scale crowdsourcing problems. The second algorithm offers enhanced identifiability guarantees, and is able to deal with more critical cases (e.g., when no highly reliable annotators exist), with the price of using a computationally more involved iterative optimization algorithm. Experiments show that both approaches outperform a number of competitive baselines.

2 Background

The Dawid-Skene Model. Let us consider a dataset $\{\mathbf{f}_n\}_{n=1}^N$, where $\mathbf{f}_n \in \mathbb{R}^d$ is a data sample (or, feature vector) and N is the number of samples. Each \mathbf{f}_n belongs to one of K classes. Let y_n be the ground-truth label of the data sample \mathbf{f}_n . Suppose that there are M annotators who work on the dataset $\{\mathbf{f}_n\}_{n=1}^N$ and provide labels. Let $X_m(\mathbf{f}_n)$ represent the response of the annotator m to \mathbf{f}_n . Hence, X_m can be understood as a discrete random variable whose alphabet is $\{1, \dots, K\}$. In crowdsourcing or ensemble learning, our goal is to estimate the true label corresponding to each item \mathbf{f}_n from the M annotator responses. Note that in a realistic scenario, an annotator will likely only work on part of the dataset, since having all annotators work on all the samples is much more costly.

In 1979, Dawid and Skene proposed an intuitively pleasing model for estimating the ‘true response’ of the patients from recorded answers [7], which is essentially a crowdsourcing/ensemble learning problem. This model has sparked a lot of interest in the machine learning community [31, 38, 15, 24, 29, 40]. The Dawid-Skene model in essence is a *naive Bayesian model* [33]. In this model, the ground-truth label of a data sample is a latent discrete random variable, Y , whose values are different class indices. The ambient variables are the responses given by different annotators, denoted as X_1, \dots, X_M , where M is the number of annotators. The key assumption in the Dawid-Skene model is that given the ground-truth label, the responses of the annotators are conditionally independent. Of course, the Dawid-Skene model is a simplified version of reality, but has been proven very useful—and it has been a workhorse for crowdsourcing since its proposal.

Under the Dawid-Skene model, one can see that

$$\Pr(X_1 = k_1, \dots, X_M = k_M) = \sum_{k=1}^K \prod_{m=1}^M \Pr(X_m = k_m | Y = k) \Pr(Y = k), \quad (1)$$

where $k \in \{1, \dots, K\}$ denotes the index of a given class, and k_m denotes the response of the m -th annotator. If one defines a series of matrices $\mathbf{A}_m \in \mathbb{R}^{K \times K}$ and let

$$\mathbf{A}(k_m, k) := \Pr(X_m = k_m | Y = k), \quad (2)$$

then $\mathbf{A}_m \in \mathbb{R}^{K \times K}$ can be understood as the ‘confusion matrix’ of annotator m : It contains all the conditional probabilities of annotator m labeling a given data sample as from class k_m while the ground-truth label is k . Also define a vector $\mathbf{d} \in \mathbb{R}^K$ such that $\mathbf{d}(k) := \Pr(Y = k)$; i.e., the prior PMF of the ground-truth label Y . Then the crowdsourcing problem boils down to estimating \mathbf{A}_m for $m = 1, \dots, M$ and \mathbf{d} .

Prior Art. In the seminal paper [7], Dawid and Skene proposed an EM-based algorithm to estimate $\Pr(X_m = k_m | Y = k)$ and $\Pr(Y = k)$. Their formulation is well-motivated from an ML viewpoint, but also has some challenges. First, it is unknown if the model is identifiable, especially when there is a large number of unrecorded responses (i.e., missing values)—but model identification plays an essential role in such estimation problems [14]. Second, since the ML estimator is a nonconvex optimization criterion, the solution quality of the EM algorithm is not easy to characterize in general. More recently, tensor methods were proposed to identify the Dawid-Skene model [40, 38]. Take the most recent work in [38] as an example. The approach considers estimating the joint probability $\Pr(X_i = k_i, X_j = k_j, X_\ell = k_\ell)$ for different triples i, j, ℓ . Such joint PMFs can be regarded as third-order tensors, and the confusion matrices and the prior \mathbf{d} are latent factors of these tensors. The upshot is that identifiability of \mathbf{A}_m and \mathbf{d} can be elegantly established leveraging tensor algebra [34, 26]. The challenge, however, is that reliably estimating $\Pr(X_i = k_i, X_j = k_j, X_\ell = k_\ell)$ is quite hard, since it normally needs a large number of annotator responses. Another tensor method in [40] judiciously partitions the data and works with group statistics between three groups, which is reminiscent of the graph statistics proposed in [1]. The method is computationally more tractable, leveraging orthogonal tensor decomposition. Nevertheless, the challenge again lies in sample complexity: the group/graph statistics are still third-order statistics.

3 Proposed Approach

In this section, we propose a model identification approach that only uses second-order statistics, in particular, pairwise co-occurrences $\Pr(X_i = k_i, X_j = k_j)$.

Problem Formulation. Let us consider the following pairwise joint PMF:

$$\Pr(X_m = k_m, X_\ell = k_\ell) = \sum_{k=1}^K \Pr(Y = k) \Pr(X_m = k_m | Y = k) \Pr(X_\ell = k_\ell | Y = k).$$

Letting $\mathbf{R}_{m,\ell}(k_m, k_\ell) = \Pr(X_m = k_m, X_\ell = k_\ell)$, and using the matrix notations that we defined, we have $\mathbf{R}_{m,\ell}(k_m, k_\ell) = \sum_{k=1}^K \Pr(Y = k) \Pr(X_m = k_m | Y = k) \Pr(X_\ell = k_\ell | Y = k)$ —or, in a more compact form:

$$\mathbf{R}_{m,\ell}(k_m, k_\ell) = \sum_{k=1}^K \mathbf{d}(k) \mathbf{A}_m(k_m, k) \mathbf{A}_\ell(k_\ell, k) \iff \mathbf{R}_{m,\ell} := \mathbf{A}_m \mathbf{D} \mathbf{A}_\ell^\top,$$

where we have $\mathbf{D} = \text{Diag}(\mathbf{d})$, which is a diagonal matrix. Note that \mathbf{A}_m is a confusion matrix, i.e., its columns are respectable probability measures. In addition, \mathbf{d} is a prior PMF. Hence, we have

$$\mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top, \quad \mathbf{A}_m \geq \mathbf{0}, \quad \forall m, \quad \mathbf{1}^\top \mathbf{d} = 1, \quad \mathbf{d} \geq \mathbf{0}. \quad (3)$$

In practice, $\mathbf{R}_{m,\ell}$ ’s are not available but can be estimated via sample averaging. Specifically, if we are given the annotator responses $X_m(\mathbf{f}_n)$, then

$$\widehat{\mathbf{R}}_{m,\ell}(k_m, k_\ell) = \frac{1}{|\mathcal{S}_{m,\ell}|} \sum_{n \in \mathcal{S}_{m,\ell}} I[X_m(\mathbf{f}_n) = k_m, X_\ell(\mathbf{f}_n) = k_\ell],$$

where $\mathcal{S}_{m,\ell}$ is the index set of samples which both annotators m and ℓ have worked on. Here, $I[\cdot]$ is an indicator function: If the event E happens, then $I[E] = 1$, and $I[E^c] = 0$ otherwise. It is readily seen that

$$\mathbb{E}[I(X_m(\mathbf{f}_n) = k_m, X_\ell(\mathbf{f}_n) = k_\ell)] = \mathbf{R}_{m,\ell}(k_m, k_\ell), \quad (4)$$

where the expectation is taken over data samples. Note that the sample complexity for reliably estimating $\mathbf{R}_{m,\ell}$ is much lower relative to that of estimating $\mathbf{R}_{m,n,\ell}$ [40, 1], and the latter is needed in tensor based methods, e.g., [38]. To be specific, to achieve $|\mathbf{R}_{m,\ell}(k_m, k_\ell) - \widehat{\mathbf{R}}_{m,\ell}(k_m, k_\ell)| \leq \epsilon$ with a probability greater than $1 - \delta$, $\mathcal{O}(\epsilon^{-2}(\log \frac{1}{\delta}))$ joint responses from annotators m and ℓ are needed. However, in order to attain the same accuracy for $\widehat{\mathbf{R}}_{m,n,\ell}(k_m, k_n, k_\ell)$, the number of joint responses from annotators m,n and ℓ is required to be atleast $\mathcal{O}(K\epsilon^{-2}(\log \frac{K}{\delta}))$, where K is the number of classes (also see supplementary materials Sec. J for a short discussion).

An Algebraic Algorithm. Assume that we have obtained $\mathbf{R}_{m,\ell}$'s for different pairs of m, ℓ . We now show how to identify \mathbf{A}_m 's and \mathbf{d} from such second-order statistics. Let us take the estimation of \mathbf{A}_m as an illustrative example. First, we construct a matrix \mathbf{Z}_m as follows:

$$\mathbf{Z}_m = [\mathbf{R}_{m,m_1}, \mathbf{R}_{m,m_2}, \dots, \mathbf{R}_{m,m_{T(m)}}], \quad (5)$$

where $m_t \neq m$ for $t = 1, \dots, T(m)$ denote the indices of annotators who have co-labeled data samples with annotator m , and the integer $T(m)$ denotes the number of such annotators. Due to the underlying model of $\mathbf{R}_{m,\ell}$ in (3), we have

$$\mathbf{Z}_m = [\mathbf{A}_m \mathbf{D} \mathbf{A}_{m_1}^\top, \dots, \mathbf{A}_m \mathbf{D} \mathbf{A}_{T(m)}^\top] = \mathbf{A}_m [\mathbf{D} \mathbf{A}_{m_1}^\top, \dots, \mathbf{D} \mathbf{A}_{T(m)}^\top] \in \mathbb{R}^{K \times KT(m)}.$$

Let us define $\mathbf{H}_m^\top = [\mathbf{D} \mathbf{A}_{m_1}^\top, \dots, \mathbf{D} \mathbf{A}_{T(m)}^\top] \in \mathbb{R}^{K \times KT(m)}$. This leads to the model $\mathbf{Z}_m = \mathbf{A}_m \mathbf{H}_m^\top$. We propose to identify \mathbf{A}_m from \mathbf{Z}_m . The key enabling postulate is that, among all annotators, some \mathbf{A}_ℓ 's should be *diagonally dominant*—if there exist annotators who are reasonably trained. In other words, for a reasonable annotator ℓ , $\Pr(X_\ell = j | Y = j)$ should be greater than $\Pr(X_\ell = j | Y = k)$ and $\Pr(X_\ell = j | Y = i)$ for $k, i \neq j$. To see the intuition of the algorithm, consider an ideal case where for each class k , there exists an annotator $m_{t(k)} \in \{m_1, \dots, m_{T(m)}\}$ such that

$$\Pr(X_{m_{t(k)}} = k | Y = k) = 1, \quad \Pr(X_{m_{t(k)}} = k | Y = j) = 0, \quad j \neq k. \quad (6)$$

This physically means that annotator $m_{t(k)}$ is very good at recognizing class k and never confuses other classes with class k . Under such circumstances, one can use the following procedure to identify \mathbf{A}_m . First, let us normalize the columns of \mathbf{Z}_m via $\overline{\mathbf{Z}}_m(:, q) = \mathbf{Z}_m(:, q) / \|\mathbf{Z}_m(:, q)\|_1$ for $q = \{1, \dots, KT(m)\}$. This way, we have a normalized model $\overline{\mathbf{Z}}_m = \overline{\mathbf{A}}_m \overline{\mathbf{H}}_m^\top$, where

$$\overline{\mathbf{A}}_m(:, k) = \frac{\mathbf{A}_m(:, k)}{\|\mathbf{A}_m(:, k)\|_1} = \mathbf{A}_m(:, k), \quad \overline{\mathbf{H}}_m(q, :) = \frac{\mathbf{H}_m(q, :) \|\mathbf{A}_m(:, k)\|_1}{\|\mathbf{Z}_m(:, q)\|_1}. \quad (7)$$

where the second equality above is because $\|\mathbf{A}_m(:, k)\|_1 = 1$ [cf. Eq. (3)]. After normalization, it can be verified that

$$\overline{\mathbf{H}}_m \mathbf{1} = \mathbf{1}, \quad \overline{\mathbf{H}}_m \geq \mathbf{0}, \quad (8)$$

i.e., all the rows of $\overline{\mathbf{H}}_m$ reside in the $(K - 1)$ -probability simplex. In addition, by the assumption in (6), it is readily seen that there exists $A_q = \{q_1, \dots, q_K\} \subset \{1, \dots, L_m\}$ where $L_m = KT(m)$ such that

$$\overline{\mathbf{H}}_m(A_q, :) = \mathbf{I}_K, \quad (9)$$

i.e., an identity matrix is a submatrix of $\overline{\mathbf{H}}_m$ (after proper row permutations). Consequently, we have $\mathbf{A}_m = \overline{\mathbf{Z}}_m(:, A_q)$ —i.e., \mathbf{A}_m can be identified from $\overline{\mathbf{Z}}_m$ up to column permutations. The task also boils down to identifying A_q . This turns out to be a well-studied task in the context of *separable nonnegative matrix factorization* [17, 16, 14], and an algebraic algorithm exists:

$$\widehat{q}_k = \arg \max_{q \in \{1, \dots, L_m\}} \left\| \mathbf{P}_{\mathbf{A}_m(:, 1:k-1)}^\perp \overline{\mathbf{Z}}_m(:, q) \right\|_2^2, \quad \forall k. \quad (10)$$

where $\widehat{\mathbf{A}}_m(:, 1:k-1) = [\overline{\mathbf{Z}}_m(:, \widehat{q}_1), \dots, \overline{\mathbf{Z}}_m(:, \widehat{q}_{k-1})]$ and $\mathbf{P}_{\widehat{\mathbf{A}}_m(:, 1:k-1)}^\perp$ is a projector onto the orthogonal complement of $\text{range}(\widehat{\mathbf{A}}_m(:, 1:k-1))$ and we let $\mathbf{P}_{\widehat{\mathbf{A}}_m(:, 1:0)}^\perp := \mathbf{I}$.

It has been shown in [17, 2] that the so-called *successive projection algorithm* (SPA) in Eq. (10) identifies A_q in K steps. This is a very plausible result, since the procedure admits Gram-Schmitt-like lightweight steps and thus is quite scalable. See more details in Sec. F.1.

Each of the A_m 's can be estimated from the corresponding \bar{Z}_m by repeatedly applying SPA, and we call this simple procedure *multiple SPA* (MultiSPA) as we elaborate in Algorithm 1.

Algorithm 1 MultiSPA

Input: Annotator Responses $\{X_m(\mathbf{f}_n)\}$.

Output: \hat{A}_m for $m = 1, \dots, M$, $\hat{\mathbf{d}}$.

estimate second order statistics $\hat{\mathbf{R}}_{m,\ell}$;

for $m = 1$ **to** M **do**

 construct $\hat{\mathbf{Z}}_m$ and normalize columns to unit ℓ_1 norm;

 estimate \hat{A}_m using Eq. (10);

end for

fix permutation mismatch between \hat{A}_m and \hat{A}_ℓ for all $m \neq \ell$;

estimate $\hat{\mathbf{D}} = \hat{A}_m^{-1} \mathbf{R}_{m,\ell} (\hat{A}_\ell^\top)^{-1}$ (and take average over all pairs (m, ℓ) if needed).;

extract the prior $\hat{\mathbf{d}} = \text{diag}(\hat{\mathbf{D}})$.

Of course, assuming that (6) or (9) holds perfectly may be too ideal. It is more likely that there exist some annotators who are good at recognizing certain classes, but still have some possibilities of being confused. It is of interest to analyze how SPA can do under such conditions. Another challenge is that one may not have $\mathbf{R}_{m,\ell}$ perfectly estimated, since only limited number of samples are available. It is desirable to understand the sample complexity of applying SPA to Dawid-Skene identification. We answer these two key technical questions in the following theorem:

Theorem 1. Assume that annotators m and t co-label at least S samples $\forall t \in \{m_1, \dots, m_{T(m)}\}$, and that $\hat{\mathbf{Z}}_m$ is constructed using $\hat{\mathbf{R}}_{m,m_{T(m)}}$'s according to Eq. (5). Also assume that the constructed $\hat{\mathbf{Z}}_m$ satisfies $\|\hat{\mathbf{Z}}_m(:, l)\|_1 \geq \eta, \forall l \in \{1, \dots, KT(m)\}$, where $\eta \in (0, 1]$. Suppose that $\text{rank}(\mathbf{A}_m) = \text{rank}(\mathbf{D}) = K$ for $m = 1, \dots, M$, and that for every class index $k \in \{1, \dots, K\}$, there exists an annotator $m_{t(k)} \in \{m_1, \dots, m_{T(m)}\}$ such that

$$\Pr(X_{m_{t(k)}} = k | Y = k) \geq (1 - \epsilon) \sum_{j=1}^K \Pr(X_{m_{t(k)}} = k | Y = j), \quad (11)$$

where $\epsilon \in [0, 1]$. Then, if $\epsilon \leq \mathcal{O}\left(\max\left(K^{-1}\kappa^{-3}(\mathbf{A}_m), \sqrt{\ln(1/\delta)}(\sigma_{\max}(\mathbf{A}_m)\sqrt{S}\eta)^{-1}\right)\right)$, with probability greater than $1 - \delta$, the SPA algorithm in (10) can estimate an \hat{A}_m such that

$$\left(\min_{\mathbf{\Pi}} \|\hat{A}_m \mathbf{\Pi} - \mathbf{A}_m\|_{2,\infty}\right) \leq \mathcal{O}\left(\sqrt{K}\kappa^2(\mathbf{A}_m) \max\left(\sigma_{\max}(\mathbf{A}_m)\epsilon, \sqrt{\ln(1/\delta)}(\sqrt{S}\eta)^{-1}\right)\right) \quad (12)$$

where $\mathbf{\Pi} \in \mathbb{R}^{K \times K}$ is a permutation matrix, $\|\mathbf{Y}\|_{2,\infty} = \max_{\ell} \|\mathbf{Y}(:, \ell)\|_2$, $\sigma_{\max}(\mathbf{A}_m)$ is the largest singular value of \mathbf{A}_m , and $\kappa(\mathbf{A}_m)$ is the condition number of \mathbf{A}_m .

In the above Theorem, the assumption $\|\hat{\mathbf{Z}}_m(:, l)\|_1 \geq \eta$ means that the proposed algorithm favors cases where more co-occurrences are observed, since $\hat{\mathbf{Z}}_m$'s elements are averaged number of co-occurrences—which makes a lot of sense. In addition, Eq. (11) relaxes the ideal assumption in (6), allowing the ‘good annotator’ $m_{t(k)}$ to confuse class $j \neq k$ with class k up to a certain probability, thereby being more realistic. The proof of Theorem 1 is reminiscent of the noise robustness of the SPA algorithm [17, 2]; see the supplementary materials (Sec. F.1). A direct corollary is as follows:

Corollary 1. Assume that the conditions in Theorem 1 hold for $\hat{\mathbf{Z}}_m$ and $\mathbf{A}_m, \forall m \in \{1, \dots, M\}$. Then, the estimation error bound in (12) holds for every MultiSPA-output $\hat{A}_m, \forall m \in \{1, \dots, M\}$.

Theorem 1 and Corollary 1 are not entirely surprising due to the extensive research on SPA-like algorithms [2, 17, 11, 30, 5]. The implication for crowdsourcing, however, is quite intriguing. First, one can see that if an annotator m does not label all the data samples, it does not necessarily hurt

the model identifiability—as long as annotator m has co-labeled some samples with a number of other annotators, identification of \mathbf{A}_m is possible. Second, assume that there exists a well-trained annotator m^* whose confusion matrix is diagonally dominant, then for every annotator m who has co-labeled samples with annotator m^* , the matrix $\overline{\mathbf{H}}_m$ can easily satisfy (11) by letting $m_{t(k)} = m^*$ for all k . In practice, one would not know who is m^* —otherwise the crowdsourcing problem would be trivial. However, one can design a dispatch strategy such that every pair of annotators m and ℓ co-label a certain amount of data. This way, it guarantees that \mathbf{A}_{m^*} appears in everyone else’s $\overline{\mathbf{H}}_m$ and thus ensures identifiability of all \mathbf{A}_m ’s for $m \neq m^*$. This insight may shed some light on how to effectively dispatch data to annotators.

Another interesting question to ask is *does having more annotators help?* Intuitively, having more annotators should help: If one has more rows in $\overline{\mathbf{H}}_m$, then it is more likely that some rows approach the vertices of the probability simplex—which can then enable SPA. We use the following simplified generative model and theorem to formalize the intuition:

Theorem 2. *Let $\rho > 0, \varepsilon > 0$, and assume that the rows of $\overline{\mathbf{H}}_m$ are generated within the $(K - 1)$ -probability simplex uniformly at random. If the number of annotators satisfies $M \geq \mathcal{O}\left(\frac{\varepsilon^{-2(K-1)}}{K} \log\left(\frac{K}{\rho}\right)\right)$, then, with probability greater than or equal to $1 - \rho$, there exist rows of $\overline{\mathbf{H}}_m$ indexed by q_1, \dots, q_K such that $\|\overline{\mathbf{H}}_m(q_k, :) - \mathbf{e}_k^\top\|_2 \leq \varepsilon$, $k = 1, \dots, K$.*

Note that Theorem 2 implies (11) under proper ε and ρ —and thus having more annotators indeed helps identify the model. The above can be shown by utilizing the Chernoff-Hoeffding inequality, and the detailed proof can be found in the supplementary materials (Sec. G).

After obtaining $\widehat{\mathbf{A}}_m$ ’s, \mathbf{d} can be estimated via various ways—see the supplementary materials in Sec. D. Using $\widehat{\mathbf{d}}$ and $\widehat{\mathbf{A}}_m$ ’s together, ML and MAP estimators for the true labels can be built up [38].

4 Identifiability-enhanced Algorithm

The MultiSPA algorithm is intuitive and lightweight, and is effective as we will show in the experiments. One concern is that perhaps the assumption in (11) may be violated in some cases. In this section, we propose another model identification algorithm that is potentially more robust to critical scenarios. Specifically, we consider the following feasibility problem:

$\text{find } \{\mathbf{A}_m\}_{m=1}^M, \mathbf{D} \tag{13a}$
$\text{subject to } \mathbf{R}_{m,\ell} = \mathbf{A}_m \mathbf{D} \mathbf{A}_\ell^\top, \forall m, \ell \in \{1, \dots, M\} \tag{13b}$
$\mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top, \mathbf{A}_m \geq \mathbf{0}, \forall m, \mathbf{1}^\top \mathbf{d} = 1, \mathbf{d} \geq \mathbf{0}. \tag{13c}$

The criterion in (13) seeks confusion matrices and a prior PMF that fit the available second-order statistics. The constraints in (13c) reflect the fact that the columns of \mathbf{A}_m ’s are conditional PMFs and the prior \mathbf{d} is also a PMF.

To proceed, let us first introduce the following notion from convex geometry [14, 28]:

Definition 1. (*Sufficiently Scattered*) *A nonnegative matrix $\mathbf{H} \in \mathbb{R}^{L \times K}$ is sufficiently scattered if 1) $\text{cone}\{\mathbf{H}^\top\} \supseteq \mathcal{C}$, and 2) $\text{cone}\{\mathbf{H}^\top\}^* \cap \text{bd}\mathcal{C}^* = \{\lambda \mathbf{e}_k \mid \lambda \geq 0, k = 1, \dots, K\}$. Here, $\mathcal{C} = \{\mathbf{x} \mid \mathbf{x}^\top \mathbf{1} \geq \sqrt{K-1} \|\mathbf{x}\|_2\}$, $\mathcal{C}^* = \{\mathbf{x} \mid \mathbf{x}^\top \mathbf{1} \geq \|\mathbf{x}\|_2\}$. In addition, $\text{cone}\{\mathbf{H}^\top\} = \{\mathbf{x} \mid \mathbf{x} = \mathbf{H}^\top \boldsymbol{\theta}, \forall \boldsymbol{\theta} \geq \mathbf{0}\}$ and $\text{cone}\{\mathbf{H}^\top\}^* = \{\mathbf{y} \mid \mathbf{x}^\top \mathbf{y} \geq 0, \forall \mathbf{x} \in \text{cone}\{\mathbf{H}^\top\}\}$ are the conic hull of \mathbf{H}^\top and its dual cone, respectively, and bd is the boundary of a closed set.*

The sufficiently scattered condition has recently emerged in convex geometry-based matrix factorization [28, 13]. This condition models how the rows of \mathbf{H} are spread in the nonnegative orthant. In principle, the sufficiently scattered condition is much easier to be satisfied relative to the condition as in (9), or, the so-called *separability condition* under the context of nonnegative matrix factorization [10, 17]. \mathbf{H} satisfying the separability condition is the extreme case, meaning that $\text{cone}\{\mathbf{H}^\top\} = \mathbb{R}_+^K$. However, the sufficiently scattered condition only requires $\mathcal{C} \subseteq \text{cone}\{\mathbf{H}^\top\}$ —which is naturally much more relaxed; also see [14] and the supplementary materials for detailed illustrations (Sec. E).

Regarding identifiability of $\mathbf{A}_1, \dots, \mathbf{A}_M$ and \mathbf{d} , we have the following result:

Theorem 3. Assume that $\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{A}_m) = K$ for all $m = 1, \dots, M$, and that there exist two subsets of the annotators, indexed by \mathcal{P}_1 and \mathcal{P}_2 , where $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$ and $\mathcal{P}_1 \cup \mathcal{P}_2 \subseteq \{1, \dots, M\}$. Suppose that from \mathcal{P}_1 and \mathcal{P}_2 the following two matrices can be constructed: $\mathbf{H}^{(1)} = [\mathbf{A}_{m_1}^\top, \dots, \mathbf{A}_{m_{|\mathcal{P}_1|}}^\top]^\top$, $\mathbf{H}^{(2)} = [\mathbf{A}_{\ell_1}^\top, \dots, \mathbf{A}_{\ell_{|\mathcal{P}_2|}}^\top]^\top$, where $m_t \in \mathcal{P}_1$ and $\ell_j \in \mathcal{P}_2$. Furthermore, assume that i) both $\mathbf{H}^{(1)}$ and $\mathbf{H}^{(2)}$ are sufficiently scattered; ii) all \mathbf{R}_{m_t, ℓ_j} 's for $m_t \in \mathcal{P}_1$ and $\ell_j \in \mathcal{P}_2$ are available; and iii) for every $m \notin \mathcal{P}_1 \cup \mathcal{P}_2$ there exists a $\mathbf{R}_{m, r}$ available, where $r \in \mathcal{P}_1 \cup \mathcal{P}_2$. Then, solving Problem (13) recovers \mathbf{A}_m for $m = 1, \dots, M$ and $\mathbf{D} = \text{Diag}(\mathbf{d})$ up to identical column permutation.

The proof of Theorem 3 is relegated to the supplementary results (Sec. H). Note that the theorem holds under the the existence of \mathcal{P}_1 and \mathcal{P}_2 , but there is no need to know the sets *a priori*. Generally speaking, a ‘taller’ matrix $\mathbf{H}^{(i)}$ would have a better chance to have its rows sufficiently spread in the nonnegative orthant under the same intuition of Theorem 2. Thus, having more annotators also helps to attain the sufficiently scattered condition. Nevertheless, formally showing the relationship between the number of annotators and $\mathbf{H}^{(i)}$ for $i = 1, 2$ being sufficiently scattered is more challenging than the case in Theorem 2, since the sufficiently scattered condition is a bit more abstract relative to the separability condition—the latter specifically assumes e_k 's exist as rows of $\mathbf{H}^{(i)}$ while the former depends on the ‘shape’ of the conic hull of $(\mathbf{H}^{(i)})^\top$, which contains an infinite number of cases. Towards this end, let us first define the following notion:

Definition 2. Assume that there exist $\widetilde{\mathbf{H}} \in \mathbb{R}^{L \times K}$ such that $\widetilde{\mathbf{H}}$ is sufficiently scattered. Also assume \mathcal{V} is the row index set of $\widetilde{\mathbf{H}}$ such that $\widetilde{\mathbf{H}}(\mathcal{V}, :)$ collects the extreme rays of $\text{cone}\{\widetilde{\mathbf{H}}^\top\}$. If there exist row indices $\ell_v \in \{1, \dots, L\}$ for all $v \in \mathcal{V}$, such that $\|\widetilde{\mathbf{H}}(v, :) - \mathbf{H}(\ell_v, :)\|_2 \leq \varepsilon$, then $\mathbf{H} \in \mathbb{R}^{L \times K}$ is called ε -sufficiently scattered.

One can see that an ε -sufficiently scattered matrix is sufficiently scattered when $\varepsilon \rightarrow 0$. With this definition, we show the following theorem:

Theorem 4. Let $\rho > 0$, $\frac{\rho}{2} > \varepsilon > 0$, and assume that the rows of $\mathbf{H}^{(1)}$ and $\mathbf{H}^{(2)}$ are generated from \mathbb{R}^K uniformly at random. If the number of annotators satisfies $M \geq \mathcal{O}\left(\frac{(K-1)^2}{K\alpha^{2(K-2)}\varepsilon^2} \log\left(\frac{K(K-1)}{\rho}\right)\right)$, where $\alpha = 1$ for $K = 2$, $\alpha = 2/3$ for $K = 3$ and $\alpha = 1/2$ for $K > 3$, then with probability greater than or equal to $1 - \rho$, $\mathbf{H}^{(1)}$ and $\mathbf{H}^{(2)}$ are ε -sufficiently scattered.

The proof of Theorem 4 is relegated to the supplementary materials (Sec. I). One can see that to satisfy ε -sufficiently scattered condition, M is smaller than that in Theorem 2. Conditions i)-iii) in Theorem 3 and Theorem 4 together imply that if we have enough annotators, and if many pairs co-label a certain number of data, then it is quite possible that one can identify the Dawid-Skene model via simply finding a feasible solution to (13). This feasibility problem is nonconvex, but can be effectively approximated; see the supplementary materials (Sec. C). In a nutshell, we reformulate the problem as a Kullback-Leibler (KL) divergence-based constrained fitting problem and handle it using alternating optimization. Since nonconvex optimization relies on initialization heavily, we use MultiSPA to initialize the fitting stage—which we will refer to as the MultiSPA-KL algorithm.

5 Experiments

Baselines. The performance of the proposed approach is compared with a number of competitive baselines, namely, Spectral-D&S [40], TensorADMM [38], and KOS [23], EigRatio [6], GhoshSVD [15] and MinmaxEntropy [41]. The performance of the Majority Voting scheme and the Majority Voting initialized Dawid-Skene (MV-D&S) estimator [7] are also presented. We also use MultiSPA to initialize EM algorithm (named as MultiSPA-D&S). Note that KOS, EigRatio and MinmaxEntropy work with more complex models relative to the Dawid-Skene model, but are considered as good baselines for the crowdsourcing/ensemble learning tasks. After identifying the model parameters, we construct a MAP predictor following [38] and observe the result. The algorithms are coded in Matlab.

Synthetic-data Simulations. Due to page limitations, synthetic data experiments demonstrating model identifiability of the proposed algorithms are presented in the supplementary materials (Sec. A).

Integrating Machine Classifiers. We employ different UCI datasets (<https://archive.ics.uci.edu/ml/datasets.html>; details in Sec. B). For each of the datasets under test, we use a collection of different classification algorithms to annotate the data samples. Different classification algorithms from the MATLAB machine learning toolbox (<https://www.mathworks.com/>

Table 1: Classification Error (%) on UCI Datasets; see runtime tabulated in Sec. B.

Algorithms	Nursery			Mushroom			Adult		
	$p = 1$	$p = 0.5$	$p = 0.2$	$p = 1$	$p = 0.5$	$p = 0.2$	$p = 1$	$p = 0.5$	$p = 0.2$
MultiSPA	2.83	4.54	17.96	0.02	0.293	6.35	15.71	16.05	17.66
MultiSPA-KL	2.72	4.26	13.06	0.00	0.152	5.89	15.66	15.98	17.63
MultiSPA-D&S	2.82	4.44	13.39	0.00	0.194	6.17	15.74	16.29	23.88
Spectral-D&S	3.14	37.2	44.29	0.00	0.198	6.17	15.72	16.31	23.97
TensorADMM	17.97	7.26	19.78	0.06	0.237	6.18	15.72	16.05	25.08
MV-D&S	2.92	66.48	66.61	0.00	47.99	48.63	15.76	75.21	75.13
Minmax-entropy	3.63	26.31	11.09	0.00	0.163	8.14	16.11	16.92	15.64
EigenRatio	N/A	N/A	N/A	0.06	0.329	5.97	15.84	16.28	17.69
KOS	4.21	6.07	13.48	0.06	0.576	6.42	17.19	24.97	38.29
Ghosh-SVD	N/A	N/A	N/A	0.06	0.329	5.97	15.84	16.28	17.71
Majority Voting	2.94	4.83	19.75	0.14	0.566	6.57	15.75	16.21	20.57
Single Best	3.94	N/A	N/A	0.00	N/A	N/A	16.23	N/A	N/A
Single Worst	15.65	N/A	N/A	7.22	N/A	N/A	19.27	N/A	N/A

products/statistics.html) such as various k -nearest neighbour classifiers, support vector machine classifiers, and decision tree classifiers are employed to serve as our machine annotators. In order to train the annotators, we use 20% of the samples to act as training data. After the data samples are trained, we use the annotators to label the unseen data samples. In practice, not all samples are labeled by an annotator due to several factors such as annotator capacity, difficulty of the task, economical issues and so on. To simulate such a scenario, each of the trained algorithms is allowed to label a data sample with probability $p \in (0, 1]$. We test the performance of all the algorithms under different p 's—and a smaller p means a more challenging scenario. All the results are averaged from 10 random trials.

Table 1 shows the classification error of the algorithms under test. Since GhoshSVD and EigenRatio works only on binary tasks, they are not evaluated for the Nursery dataset where $K = 4$. The ‘single best’ and ‘single worst’ rows correspond to the results of using the classifiers individually when $p = 1$, as references. The best and second-best performing algorithms are highlighted in the table. One can see that the proposed methods are quite promising for this experiment. Both algorithms largely outperform the tensor based methods TensorADMM and Spectral-D&S in this case, perhaps because the limited number of available samples makes the third-order statistics hard to estimate. It is also observed that the proposed algorithms enjoy favorable runtime; see supplementary materials (cf. Table 8 in Sec. B). Using the MultiSPA to initialize EM (i.e. MultiSPA-D&S) also works well, which offers another viable option that strikes a good balance between runtime and accuracy.

Amazon Mechanical Turk Crowdsourcing Data. In this section, the performance of the proposed algorithms are evaluated using the Amazon Mechanical Turk (AMT) data (<https://www.mturk.com>) in which human annotators label various classification tasks. Data description is given in the supplementary materials Sec. B. Table 2 shows the classification error and the runtime performance of the algorithms under test. One can see that MultiSPA has a very favorable execution time, because it is a Gram-Schmitt-like algorithm. MultiSPA-KL uses more time, because it is an iterative optimization method—with better accuracy paid off. Since TensorADMM algorithm does not scale well, the results are not reported for very large datasets (i.e., TREC and RTE). Similar as before, since Web and Dog are multi-class datasets, EigenRatio and GhoshSVD are not applicable. From the results, it can be seen that the proposed algorithms outperform many existing crowdsourcing algorithms in both classification accuracy and runtime. In particular, one can see that the algebraic algorithm MultiSPA gives very similar results compared to the computationally much more involved algorithms. This shows the potential for its application in big data crowdsourcing.

Table 2: Classification Error (%) and Run-time (sec) : AMT Datasets

Algorithms	TREC		Bluebird		RTE		Web		Dog	
	(%) Error	(sec) Time	(%) Error	(sec) Time	(%) Error	(sec) Time	(%) Error	(sec) Time	(%) Error	(sec) Time
MultiSPA	31.47	50.68	13.88	0.07	8.75	0.28	15.22	0.54	17.09	0.07
MultiSPA-KL	29.23	536.89	11.11	1.94	7.12	17.06	14.58	12.34	15.48	15.88
MultiSPA-D&S	29.84	53.14	12.03	0.09	7.12	0.32	15.11	0.84	16.11	0.12
Spectral-D&S	29.58	919.98	12.03	1.97	7.12	6.40	16.88	179.92	17.84	51.16
TensorADMM	N/A	N/A	12.03	2.74	N/A	N/A	N/A	N/A	17.96	603.93
MV-D&S	30.02	3.20	12.03	0.02	7.25	0.07	16.02	0.28	15.86	0.04
Minmax-entropy	91.61	352.36	8.33	3.43	7.50	9.10	11.51	26.61	16.23	7.22
EigenRatio	43.95	1.48	27.77	0.02	9.01	0.03	N/A	N/A	N/A	N/A
KOS	51.95	9.98	11.11	0.01	39.75	0.03	42.93	0.31	31.84	0.13
GhoshSVD	43.03	11.62	27.77	0.01	49.12	0.03	N/A	N/A	N/A	N/A
Majority Voting	34.85	N/A	21.29	N/A	10.31	N/A	26.93	N/A	17.91	N/A

6 Conclusion

In this work, we have revisited the classic Dawid-Skene model for multi-class crowdsourcing. We have proposed a second-order statistics-based approach that guarantees identifiability of the model parameters, i.e., the confusion matrices of the annotators and the label prior. The proposed method naturally admits lower sample complexity relative to existing methods that utilize tensor algebra to ensure model identifiability. The proposed approach also has an array of favorable features. In particular, our framework enables a lightweight algebraic algorithm, which is reminiscent of the Gram-Schmitt-like SPA algorithm for nonnegative matrix factorization. We have also proposed a coupled and constrained matrix factorization criterion that enjoys enhanced-identifiability, as well as an alternating optimization algorithm for handling the identification problem. Real-data experiments show that our proposed algorithms are quite promising for integrating crowdsourced labeling.

References

- [1] Anandkumar, A., Ge, R., Hsu, D., and Kakade, S. M. A tensor approach to learning mixed membership community models. *The Journal of Machine Learning Research*, 15(1):2239–2312, 2014.
- [2] Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., Wu, Y., and Zhu, M. A practical algorithm for topic modeling with provable guarantees. In *Proceedings of ICML*, 2013.
- [3] Baraniuk, R., Davenport, M., DeVore, R., and Wakin, M. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [4] Bertsekas, D. P. *Nonlinear programming*. Athena Scientific, 1999.
- [5] Chan, T.-H., Ma, W.-K., Ambikapathi, A., and Chi, C.-Y. A simplex volume maximization framework for hyperspectral endmember extraction. *IEEE Trans. Geosci. Remote Sens.*, 49(11): 4177–4193, Nov. 2011.
- [6] Dalvi, N., Dasgupta, A., Kumar, R., and Rastogi, V. Aggregating crowdsourced binary ratings. In *Proceedings of the 22Nd International Conference on World Wide Web*, pp. 285–294, New York, NY, USA, 2013. ACM.
- [7] Dawid, A. P. and Skene, A. M. Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied statistics*, pp. 20–28, 1979.
- [8] Deng, J., Dong, W., Socher, R., Li, L., and and. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, June 2009.
- [9] Dietterich, T. G. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pp. 1–15. Springer, 2000.
- [10] Donoho, D. and Stodden, V. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in neural information processing systems*, volume 16, 2003.

- [11] Fu, X., Ma, W.-K., Chan, T.-H., and Bioucas-Dias, J. M. Self-dictionary sparse regression for hyperspectral unmixing: Greedy pursuit and pure pixel search are related. *IEEE J. Sel. Topics Signal Process.*, 9(6):1128–1141, 2015.
- [12] Fu, X., Huang, K., Yang, B., Ma, W.-K., and Sidiropoulos, N. D. Robust volume minimization-based matrix factorization for remote sensing and document clustering. *IEEE Trans. Signal Process.*, 64(23):6254–6268, 2016.
- [13] Fu, X., Huang, K., and Sidiropoulos, N. D. On identifiability of nonnegative matrix factorization. *IEEE Signal Process. Lett.*, 25(3):328–332, 2018.
- [14] Fu, X., Huang, K., Sidiropoulos, N. D., and Ma, W.-K. Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications. *arXiv preprint arXiv:1803.01257*, 2018.
- [15] Ghosh, A., Kale, S., and McAfee, P. Who moderates the moderators?: crowdsourcing abuse detection in user-generated content. In *Proceedings of the 12th ACM conference on Electronic commerce*, pp. 167–176. ACM, 2011.
- [16] Gillis, N. The why and how of nonnegative matrix factorization. *Regularization, Optimization, Kernels, and Support Vector Machines*, 12:257, 2014.
- [17] Gillis, N. and Vavasis, S. Fast and robust recursive algorithms for separable nonnegative matrix factorization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(4):698–714, April 2014.
- [18] Huang, K., Sidiropoulos, N., and Swami, A. Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition. *IEEE Trans. Signal Process.*, 62(1): 211–224, 2014.
- [19] Huang, K., Sidiropoulos, N. D., and Liavas, A. P. A flexible and efficient algorithmic framework for constrained matrix and tensor factorization. *IEEE Trans. Signal Process.*, 64(19):5052–5065, 2016.
- [20] Huang, K., Fu, X., and Sidiropoulos, N. D. Learning hidden markov models from pairwise co-occurrences with applications to topic modeling. In *Proceedings of ICML 2018*, 2018.
- [21] Jonker, R. and Volgenant, T. Improving the hungarian assignment algorithm. *Operations Research Letters*, 5(4):171–175, 1986.
- [22] Karger, D. R., Oh, S., and Shah, D. Budget-optimal crowdsourcing using low-rank matrix approximations. 2011.
- [23] Karger, D. R., Oh, S., and Shah, D. Efficient crowdsourcing for multi-class labeling. *ACM SIGMETRICS Performance Evaluation Review*, 41(1):81–92, 2013.
- [24] Karger, D. R., Oh, S., and Shah, D. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 62(1):1–24, 2014.
- [25] Kittur, A., Chi, E. H., and Suh, B. Crowdsourcing user studies with mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 453–456. ACM, 2008.
- [26] Kolda, T. G. and Bader, B. W. Tensor decompositions and applications. *SIAM review*, 51(3): 455–500, 2009.
- [27] Lease, M. and Kazai., G. Overview of the trec 2011 crowdsourcing track. 2011.
- [28] Lin, C.-H., Ma, W.-K., Li, W.-C., Chi, C.-Y., and Ambikapathi, A. Identifiability of the simplex volume minimization criterion for blind hyperspectral unmixing: The no-pure-pixel case. *IEEE Trans. Geosci. Remote Sens.*, 53(10):5530–5546, Oct 2015.
- [29] Liu, Q., Peng, J., and Ihler, A. T. Variational inference for crowdsourcing. In *Advances in neural information processing systems*, pp. 692–700, 2012.

- [30] Nascimento, J. and Bioucas-Dias, J. Vertex component analysis: A fast algorithm to unmix hyperspectral data. *IEEE Trans. Geosci. Remote Sens.*, 43(4):898–910, 2005.
- [31] Raykar, V. C., Yu, S., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., and Moy, L. Learning from crowds. *Journal of Machine Learning Research*, 11(Apr):1297–1322, 2010.
- [32] Razaviyayn, M., Hong, M., and Luo, Z.-Q. A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization*, 23(2): 1126–1153, 2013.
- [33] Robert, C. Machine learning, a probabilistic perspective, 2014.
- [34] Sidiropoulos, N. D., De Lathauwer, L., Fu, X., Huang, K., Papalexakis, E. E., and Faloutsos, C. Tensor decomposition for signal processing and machine learning. *IEEE Trans. Signal Process.*, 65(13):3551–3582, 2017.
- [35] Snow, R., O’Connor, B., Jurafsky, D., and Ng, A. Y. Cheap and fast—but is it good?: evaluating non-expert annotations for natural language tasks. In *Proceedings of the conference on empirical methods in natural language processing*, pp. 254–263. Association for Computational Linguistics, 2008.
- [36] Stein, P. A Note on the Volume of a Simplex. *The American Mathematical Monthly*, 73(3), 1966. doi: 10.2307/2315353.
- [37] Stephane Boucheron, Gabor Lugosi, O. B. Concentration Inequalities, 2004. URL: http://www.econ.upf.edu/~lugosi/mlss_conc.pdf.
- [38] Traganitis, P. A., Pages-Zamora, A., and Giannakis, G. B. Blind multiclass ensemble classification. *IEEE Trans. Signal Process.*, 66(18):4737–4752, 2018.
- [39] Welinder, P., Branson, S., Perona, P., and Belongie, S. J. The multidimensional wisdom of crowds. In *Advances in neural information processing systems*, pp. 2424–2432, 2010.
- [40] Zhang, Y., Chen, X., Zhou, D., and Jordan, M. I. Spectral methods meet em: A provably optimal algorithm for crowdsourcing. In *Advances in neural information processing systems*, pp. 1260–1268, 2014.
- [41] Zhou, D., Liu, Q., Platt, J., and Meek, C. Aggregating ordinal labels from crowds by minimax conditional entropy. In *Proceedings of ICML*, volume 32, pp. 262–270, 2014.

Supplementary Materials for “Crowdsourcing via Pairwise Co-occurrences: Identifiability and Algorithms”

A Synthetic Data Experiments

In the first experiment, we consider that $M = 25$ annotators are available to annotate $N = 10,000$ items, each belonging to one of $K = 3$ classes. The true label for each item is sampled uniformly from $\{1, \dots, K\}$, i.e., the prior probability vector \mathbf{d} is fixed to be $\mathbf{d} = [1/3, 1/3, 1/3]^\top$. For generating the confusion matrices, two different cases are considered

- **Case 1:** an annotator is chosen uniformly at random and is assigned an ideal confusion matrix, i.e., an identity matrix \mathbf{I}_3 . This ensures the assumption as given by Eq.(9) (or Eq. (6)).
- **Case 2:** an annotator m is chosen uniformly at random and its confusion matrix is made diagonally dominant such that $\mathbf{A}_m(k, k) > \mathbf{A}_m(k', k)$, for $k', k \in \{1, \dots, K\}, k \neq k'$. To achieve this, the elements of each column of \mathbf{A}_m is drawn from a uniform distribution between 0 and 1. The columns are then normalized using their respective ℓ_1 -norms. After that, for each column, the elements are re-organized such that the corresponding diagonal entry is dominant in that column and then normalized with respect to ℓ_1 -norm. In this way, Eq. (11) in Theorem 1 may be (approximately) satisfied.

In both the cases, for the remaining annotators, the confusion matrices \mathbf{A}_m are randomly generated; the elements are first drawn following the uniform distribution between 0 and 1, and then the columns are normalized with respect to the ℓ_1 -norm. Once \mathbf{A}_m 's are generated, the responses from each annotator m for the items with true labels $g \in \{1, \dots, K\}$ are randomly chosen from $\{1, \dots, K\}$ using the probability distribution $\mathbf{A}_m(:, g)$. An annotator response for each item is retained for the estimation of \mathbf{A}_m with probability $p \in (0, 1]$. In other words, with probability $1 - p$, each response is made 0. In this way, our simulated scenario is expected to mimic realistic situations where we have a combination of reliable and unreliable annotators, each labeling parts of the items. Using the generated responses, we construct $\widehat{\mathbf{R}}_{m,\ell}$'s and then follow the proposed approach to identify the confusion matrices and the prior \mathbf{d} .

The accuracy of the estimation is measured using *mean squared error* (MSE) defined as

$$\text{MSE} = \min_{\pi(k) \in \{1, \dots, K\}} \frac{1}{K} \sum_{k=1}^K \|\mathbf{A}_m(:, \pi(k)) - \widehat{\mathbf{A}}_m(:, k)\|_2^2 \quad (14)$$

where $\widehat{\mathbf{A}}_m$ is the estimate of \mathbf{A}_m and $\pi(k)$'s are used to fix the column permutation.

The average (MSE) of the confusion matrices for various values of p under the above mentioned cases are shown in Table 3 and Table 4 where the proposed methods, `MultiSPA` and `MultiSPA-KL` are compared with the baselines `Spectral-E&M`, `TensorADMM` and `MV-D&S` since these methods are also Dawid-Skene model identification approaches. As `MV-D&S` becomes numerically unstable for smaller values of p , those results are not reported in the table. All the results are averaged from 10 trials.

From the two tables, one can see that `MultiSPA` works reasonably well for both cases. As expected, it exhibits lower MSEs for case 1, since the condition in (6) is perfectly enforced. Nevertheless, in both cases, using `MultiSPA` to initialize the KL algorithm identifies the confusion matrices to a very high accuracy. It is observed that `MultiSPA-KL` outperforms the baselines in terms of the estimation accuracy—which may be a result of using second order statistics.

Under the same settings as in case 2, the true labels are estimated using the MAP/ML predictor as in [38] (in this case, ML and MAP are the same since the prior PMF is a uniform distribution). The classification error and the runtime of the crowdsourcing algorithms are computed and shown in Table 5.

In the next experiment with case 2, the true labels are sampled with unequal probability. Specifically, \mathbf{d} is set to be $[\frac{1}{6}, \frac{2}{3}, \frac{1}{6}]^\top$ with all other parameters and conditions same as in the first experiment. Using the MAP predictor, the true labels are estimated for the proposed algorithms for various values of p and the results are shown in Table 6. It can be inferred from the results that both the proposed

Table 3: Average MSE of the confusion matrices \mathbf{A}_m for case 1.

Algorithms	$p = 0.2$	$p = 0.3$	$p = 0.5$	$p = 1$
MutLiSPA	0.0184	0.0083	0.0063	0.0034
MultiSPA-KL	0.0019	0.0009	0.0004	1.73E-04
Spectral D&S	0.0320	0.0112	0.0448	1.74E-04
TensorADMM	0.0026	0.0011	0.0005	1.88E-04
MV-D&S	–	–	0.0173	1.84E-04

Table 4: Average MSE of the confusion matrices \mathbf{A}_m for case 2.

Algorithms	$p = 0.2$	$p = 0.3$	$p = 0.5$	$p = 1$
MutLiSPA	0.0229	0.0188	0.0115	0.0102
MultiSPA-KL	0.0029	0.0014	0.0005	1.67E-04
Spectral D&S	0.0348	0.0265	0.0391	1.67E-04
TensorADMM	0.0031	0.0016	0.0006	1.93E-04
MV-D&S	–	–	0.0028	5.88E-04

algorithms MultiSPA and MultiSPA-KL grantee better classification accuracy when the true label distribution of the items is not balanced.

In the next experiment, the effect of the number of annotators (M) in the estimation accuracy of the confusion matrices is investigated. According to Theorem 2 and 4, the proposed methods will benefit from the availability of more annotators (i.e., a larger M). For $N = 10,000$, $K = 3$, $\mathbf{d} = [\frac{1}{6}, \frac{2}{3}, \frac{1}{6}]^T$, $p = 0.5$ and the true confusion matrices \mathbf{A}_m being generated as in case 2, the MSEs under various values of M are plotted in Figure 1. One can see that MultiSPA-KL achieves better accuracy relative to MultiSPA under the same M 's, which corroborates our results in Theorem 4.

B More Details on UCI and AMT Dataset Experiments

UCI data. The details of the UCI datasets employed in the real data experiments is given in Table 7. To be more specific, the Adult dataset predicts the income of a person into $K = 2$ classes based on 14 attributes. The Mushroom dataset has 22 attributes of certain variations of mushrooms and the task there predicts either ‘edible’ or ‘poisonous’. The Nursery dataset predicts applications to one of the 4 categories based on 8 attributes of the financial and social status of the parents.

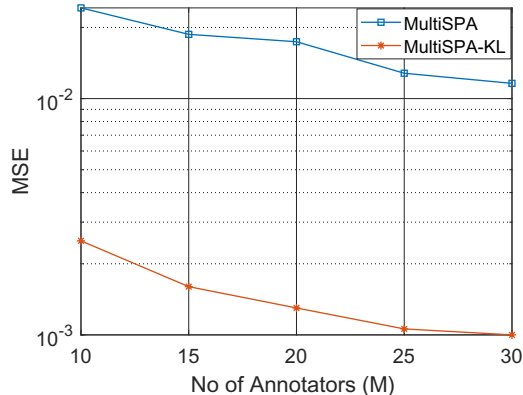


Figure 1: MSE of the confusion matrices for various values of M

Table 5: Classification Error(%) & Average run-time when $\mathbf{d} = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]^\top$

Algorithms	$p = 0.2$	$p = 0.3$	$p = 0.5$	Run-time(sec)
MultiSPA	37.24	26.39	19.21	0.049
MultiSPA-KL	31.71	21.10	12.79	18.07
MultiSPA-D&S	31.95	21.11	12.80	0.069
Spectral-D&S	46.37	23.92	12.89	27.17
TensorADMM	32.16	21.34	12.91	56.09
MV-D&S	66.91	57.92	13.09	0.096
Minmax-entropy	62.83	65.50	67.31	200.91
KOS	71.47	61.05	13.12	5.653
Majority Voting	67.57	68.37	71.39	-

Table 6: Classification Error(%) & Average run-time when $\mathbf{d} = [\frac{1}{6}, \frac{2}{3}, \frac{1}{6}]^\top$

Algorithms	$p = 0.2$	$p = 0.3$	$p = 0.5$	Run-time(sec)
MultiSPA	30.75	21.29	13.67	0.105
MultiSPA-KL	23.19	16.62	10.13	18.93
MultiSPA-D&S	40.12	32.1	21.46	0.122
Spectral-D&S	56.17	49.41	39.17	28.01
TensorADMM	34.17	25.53	11.97	152.76
MV-D&S	83.14	83.15	32.98	0.090
Minmax-entropy	83.04	63.08	74.29	232.82
KOS	70.79	67.55	78.00	6.19
Majority Voting	65.37	65.57	66.06	-

The proposed methods and the baselines are compared in terms of runtime for various datasets and the results are reported in Table 8. All the results are averaged from 10 different trials.

AMT data. The Amazon Mechanical Turk (AMT) datasets used in our crowdsourcing data experiments is given in Table 9. Specifically, the tasks involving the Bird dataset [39], the RTE dataset [35], and the TREC dataset [27], are binary classification tasks. The tasks associated with the Dog dataset [8] and the web dataset [41] are multi-class tasks (i.e., 4 and 5 classes, respectively).

We would like to add one remark regarding the two-stage approaches that involving an initial stage and a refinement stage (e.g., Spectral-D&S, MV-D&S, and MultiSPA-KL). Due to very high sparsity of the annotator responses in most of the AMT data, the estimated confusion matrices from the first stage may contain many zero entries, which may sometimes lead to numerical issues in the second stage, as observed in [40]. In our experiments, we follow an empirical thresholding strategy proposed in [40]. Specifically, the confusion matrix entries that are smaller than a threshold Δ are reset to Δ and the columns are normalized before initialization. In our experiments, we use $\Delta = 10^{-6}$ for most of the cases except the extremely large dataset TREC, which enjoys better performance of all methods using $\Delta = 10^{-5}$.

C Algorithm for Criterion (13)

In this section, the MultiSPA-KL algorithm is discussed in detail. To implement the identification criterion in (13), we lift the constraint (13b) and employ the following coupled matrix factorization

Table 7: Details of UCI Datasets.

UCI dataset name	# classes	# items	# annotators
Adult	2	7017	10
Mushroom	2	6358	10
Nursery	4	3575	10

Table 8: Average runtime (sec) for UCI dataset experiments.

Algorithms	Nursery	Mushroom	Adult
MultiSPA	0.021	0.012	0.018
MultiSPA-KL	1.112	0.663	0.948
MultiSPA-D&S	0.035	0.027	0.027
Spectral-D&S	10.09	0.496	0.512
TensorADMM	5.811	0.743	4.234
MV-D&S	0.009	0.007	0.008
Minmax-entropy	19.94	2.304	6.959
EigenRatio	–	0.005	0.007
KOS	0.768	0.085	0.118
Ghosh-SVD	–	0.081	0.115

Table 9: AMT Dataset description.

Dataset	# classes	# items	# annotators	# annotator labels
Bird	2	108	30	3240
RTE	2	800	164	8,000
TREC	2	19,033	762	88,385
Dog	4	807	52	7,354
Web	5	2,665	177	15,567

criterion:

$$\underset{\{\mathbf{A}_m\}_{m=1}^M, \mathbf{D}}{\text{minimize}} \sum_{m,\ell} \text{KL} \left(\widehat{\mathbf{R}}_{m,\ell} \parallel \mathbf{A}_m \mathbf{D} \mathbf{A}_\ell^\top \right), \quad (15a)$$

$$\text{subject to : } \mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top, \mathbf{A}_m \geq \mathbf{0}, \mathbf{1}^\top \mathbf{d} = 1, \mathbf{d} \geq \mathbf{0}, \quad (15b)$$

where $\mathbf{D} = \text{Diag}(\mathbf{d})$ and the Kullback-Leibler (KL) divergence is employed as the distance measure. The reason is that $\mathbf{R}_{m,\ell}$ is a joint PMF of two random variables, and the KL-divergence is the most natural distance measure under such circumstances. Problem (15) is a nonconvex optimization problem, but can be handled by a simple alternating optimization procedure.

Specifically, we propose to solve the following subproblems cyclically:

$$\mathbf{A}_m \leftarrow \arg \min_{\mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top, \mathbf{A}_m \geq \mathbf{0}} \sum_{\ell \in \mathcal{S}_m} \text{KL} \left(\widehat{\mathbf{R}}_{m,\ell} \parallel \mathbf{A}_m \mathbf{D} \mathbf{A}_\ell^\top \right) \quad (16a)$$

$$\mathbf{d} \leftarrow \arg \min_{\mathbf{1}^\top \mathbf{d} = 1, \mathbf{d} \geq \mathbf{0}} \sum_{\ell \in \mathcal{S}_m} \text{KL} \left(\widehat{\mathbf{R}}_{m,\ell} \parallel \mathbf{A}_m \mathbf{D} \mathbf{A}_\ell^\top \right) \quad (16b)$$

where \mathcal{S}_m denotes the index set of ℓ 's such that $\mathbf{R}_{m,\ell}$ is available. Both of the above problems are convex optimization problems, and thus can be effectively solved via a number of off-the-shelf optimization algorithms, e.g., ADMM [19] and mirror descent [2]. The detailed summarized

algorithm is in Algorithm 2. The alternating optimization algorithm is also guaranteed to converge to a stationary point under mild conditions [4, 32].

Algorithm 2 MultiSPA-KL

Input: Annotator Responses $\{X_m(\mathbf{f}_n)\}$.
Output: $\hat{\mathbf{A}}_m$ for $m = 1, \dots, M$, $\hat{\mathbf{d}}$.
 Estimate second order statistics $\hat{\mathbf{R}}_{m,\ell}$;
 get initial estimates of $\{\hat{\mathbf{A}}_m\}$ using MultiSPA
for $t = 1$ **to** MaxIter **do**
 for $m = 1$ **to** M **do**
 update $\mathbf{A}_m \leftarrow (16a)$;
 end for
 update $\mathbf{d} \leftarrow (16b)$;
end for

Note that this coupled factorization formulation bears some resemblance to the coupled tensor factorization formulation in [38]. However, the two are very different in essence. The formulation in [38] relies on the third-order statistics to establish identifiability, while the formulation in (15) establishes identifiability using nonnegativity of the confusion matrices and the prior. The KL-divergence based fitting criterion also fits the statistical learning problem better than the least squares based criterion in [38].

D Estimation of Prior Probability Vector

In this section, we discuss different methods to estimate the prior probability vector \mathbf{d} once the confusion matrices are estimated via MultiSPA algorithm.

It is to be noted that the SPA-estimated $\hat{\mathbf{A}}_m$ is up to column permutation, even if there is no noise, i.e., $\hat{\mathbf{A}}_m = \mathbf{A}_m \mathbf{\Pi}_m$ in the best case. Since our algorithm runs SPA separately for different $\bar{\mathbf{Z}}_m$'s, the permutation matrices resulted by each run of SPA need not to be identical; i.e., it is highly likely that $\mathbf{\Pi}_\ell \neq \mathbf{\Pi}_m$ for $m \neq \ell$. To estimate the prior PMF \mathbf{d} , one will need to use estimators such as

$$\hat{\mathbf{D}} = \hat{\mathbf{A}}_m^{-1} \mathbf{R}_{m,\ell} (\hat{\mathbf{A}}_\ell^\top)^{-1},$$

which cannot be applied before the permutation mismatch is fixed. In practice, the mismatch can be removed by a number of simple methods. For example, if annotator ℓ has co-labeled data with annotator m , then $\tilde{\mathbf{A}}_\ell = \mathbf{A}_\ell \mathbf{\Pi}_m$ can be estimated from $\mathbf{R}_{m,\ell}$ via $\tilde{\mathbf{A}}_\ell = \hat{\mathbf{A}}_m^{-1} \mathbf{R}_{m,\ell}$. We also have $\hat{\mathbf{A}}_\ell = \mathbf{A}_\ell \mathbf{\Pi}_\ell$ estimated from $\bar{\mathbf{Z}}_\ell$. Using a permutation matching algorithm, e.g., the Hungarian algorithm [21], one can easily remove the permutation mismatch between $\hat{\mathbf{A}}_\ell$ and $\tilde{\mathbf{A}}_\ell$. Another more heuristic yet more efficient way is to rearrange the columns of $\hat{\mathbf{A}}_m$ so that it is diagonally dominant—this makes a lot of sense if one believes that all the annotators are reasonably trained.

E Geometry of The Sufficiently Scattered Condition

In this section, we present more discussion on the sufficiently scattered condition that is used in Theorem 3. To simplify the notation, we omit the superscript of $\mathbf{H}^{(i)}$ for $i = 1, 2$ and use \mathbf{H} to denote these two matrices. The sufficiently scattered condition is geometrically intuitive. The key to understand this condition is the second-order cone \mathcal{C} , which is shown in Fig. 2. This cone is very special since it is tangent to all the facets of the nonnegative orthant.

A case where $\mathbf{H} \in \mathbb{R}^{N \times K}$ satisfies the sufficiently scattered condition is plotted in Fig. 3. One can see that the sufficiently scattered condition is much more relaxed compared to the condition that enables SPA (cf. Fig. 4). In order to apply SPA to $\bar{\mathbf{Z}}_m$, one needs that there are rows in $\bar{\mathbf{H}}_m$ that attain the extreme rays of the nonnegative orthant.




Figure 2: Illustration of the cone \mathcal{C} in an 3-dimensional space.




Figure 3: A case where \mathbf{H} satisfies the sufficiently scattered condition. The inner circle corresponds to \mathcal{C} , the dots correspond to $\mathbf{H}(q, :)$'s, and the triangle corresponds to the nonnegative orthant. The shaded region is cone $\{\mathbf{H}^\top\}$.

F Proof of Theorem 1

F.1 Identification Theory of SPA

To understand Theorem 1, let us start with a noisy nonnegative matrix factorization (NMF) model

$$\mathbf{X} = \mathbf{W}\mathbf{H}^\top + \mathbf{N}, \quad (17)$$

where $\mathbf{W} \in \mathbb{R}^{M \times K}$, $\mathbf{H} \in \mathbb{R}^{N \times K}$, $\mathbf{W} \geq \mathbf{0}$ and $\mathbf{H} \geq \mathbf{0}$, and \mathbf{N} represents the noise. Also assume that $\text{rank}(\mathbf{W}) = K$ and $\mathbf{H} = \mathbf{\Pi} \begin{bmatrix} \mathbf{I}_K \\ \mathbf{H}^* \end{bmatrix}$; i.e., there exists $\Lambda = \{q_1, \dots, q_K\}$ such that $\mathbf{H}(\Lambda, :) = \mathbf{I}$. Also assume that $\mathbf{H}\mathbf{1} = \mathbf{1}$.

The SPA algorithm under this model is as follows [17, 11, 2, 5]:

$$\begin{aligned} \hat{q}_1 &= \arg \max_{q \in \{1, \dots, N\}} \|\mathbf{X}(:, q)\|_2^2 \\ \hat{q}_k &= \arg \max_{q \in \{1, \dots, N\}} \left\| \mathbf{P}_{\widehat{\mathbf{W}}(:, 1:k-1)}^\perp \mathbf{X}(:, q) \right\|_2^2, \quad k > 1. \end{aligned}$$

where

$$\widehat{\mathbf{W}}(:, 1:k-1) = [\mathbf{X}(:, \hat{q}_1), \dots, \mathbf{X}(:, \hat{q}_{k-1})]$$

collects all the previously estimated columns of \mathbf{W} and $\mathbf{P}_{\widehat{\mathbf{W}}(:, 1:k-1)}^\perp$ is a projector onto the orthogonal complement of $\text{range}(\widehat{\mathbf{W}}(:, 1:k-1))$.




Figure 4: A case where $\overline{\mathbf{H}}_m$ satisfies the condition for applying SPA (aka. the separability condition). The inner circle corresponds to \mathcal{C} , the dots correspond to $\overline{\mathbf{H}}_m(q, \cdot)$'s, and the triangle corresponds to the nonnegative orthant. The shaded region is cone $\{\overline{\mathbf{H}}_m^\top\}$.

When there is no noise, it was shown in the literature that SPA readily identifies Λ [17, 2]. To see this, consider

$$\begin{aligned} \|\mathbf{X}(\cdot, q)\|_2 &= \left\| \sum_{k=1}^K \mathbf{W}(:, k) \mathbf{H}(q, k) \right\|_2 \leq \sum_{k=1}^K \|\mathbf{W}(:, k) \mathbf{H}(q, k)\|_2 \\ &= \sum_{k=1}^K \mathbf{H}(q, k) \|\mathbf{W}(:, k)\|_2 \leq \max_{k=1, \dots, K} \|\mathbf{W}(:, k)\|_2, \end{aligned}$$

where the two equalities hold simultaneously if and only if $\mathbf{H}(q, \cdot) = e_k^\top$ for a certain k , i.e., $q \in \Lambda$. After identifying the first index \hat{q}_1 in Λ , then by projecting all the data column onto the orthogonal complement of $\mathbf{X}(\cdot, \hat{q}_1)$, the same $\mathbf{W}(:, k)$ will not come up again. Hence, K steps of SPA identifies the whole \mathbf{W} .

A salient feature of SPA is that it is provably robust to noise. To be specific, Gillis and Vavasis have shown that:

Lemma 1. [17] *Under the described NMF model, assume that $\|\mathbf{N}(:, l)\|_2 \leq \delta$ for all l . If the below holds:*

$$\delta \leq \sigma_{\min}(\mathbf{W}) \min\left(\frac{1}{2\sqrt{K-1}}, \frac{1}{4}\right) (1 + 80\kappa^2(\mathbf{W}))^{-1},$$

then, SPA identifies an index set $\hat{\Lambda} = \{\hat{q}_1, \dots, \hat{q}_K\}$ such that

$$\max_{1 \leq j \leq K} \min_{\hat{q}_k \in \hat{\Lambda}} \|\mathbf{W}(:, j) - \mathbf{X}(\cdot, \hat{q}_k)\|_2 \leq \delta(1 + 80\kappa^2(\mathbf{W}))$$

where $\kappa(\mathbf{W}) = \frac{\sigma_{\max}(\mathbf{W})}{\sigma_{\min}(\mathbf{W})}$ is the condition number of \mathbf{W} .

F.2 Proof of The Theorem

Since $\mathbf{R}_{m,\ell}$ is obtained by sample averaging of a finite number of pairwise co-occurrences, the estimated $\hat{\mathbf{R}}_{m,\ell}$ is always noisy; i.e., we have

$$\hat{\mathbf{R}}_{m,\ell} = \mathbf{R}_{m,\ell} + \mathbf{N}_{m,\ell}, \quad (19)$$

where the noise matrix $\mathbf{N}_{m,\ell}$ has same dimension as $\hat{\mathbf{R}}_{m,\ell}$ or $\mathbf{R}_{m,\ell}$ and its norm can be bounded by Lemma 2.

Lemma 2. [1] *Let $\delta \in (0, 1)$ and let $\hat{\mathbf{R}}_{m,\ell}$ be the empirical average of S independent co-occurrences of random variables X_m and X_ℓ where $X_m, X_\ell \in \{1, \dots, K\}$, then the following holds*

$$\Pr \left[\|\hat{\mathbf{R}}_{m,\ell} - \mathbf{R}_{m,\ell}\|_F = \|\mathbf{N}_{m,\ell}\|_F \leq \frac{1 + \sqrt{\ln(1/\delta)}}{\sqrt{S}} \right] \geq 1 - \delta$$

Using the estimates $\hat{\mathbf{R}}_{m,\ell}$, $\hat{\mathbf{Z}}_m$ is constructed according to (5) (with $\mathbf{R}_{m,\ell}$'s replaced by $\hat{\mathbf{R}}_{m,\ell}$'s). The columns of $\hat{\mathbf{Z}}_m$ are normalized before performing MultiSPA, essentially normalizing the columns of $\hat{\mathbf{R}}_{m,\ell}$. Normalization complicates the analysis since the demonstrators used in this step are also noisy. We derive Lemma 3 to characterize the noise bound after column normalization.

Lemma 3. Assume that there exists at least S joint responses from each of the annotator pairs m, ℓ . Let $\eta \in (0, 1)$. If $\|\widehat{\mathbf{R}}_{m,\ell}(:, k)\|_1 \geq \eta$ and $\|\mathbf{N}_{m,\ell}(:, k)\|_1 < \|\mathbf{R}_{m,\ell}(:, k)\|_1, \forall k \in \{1, \dots, K\}, \forall m \neq \ell$, then with probability greater than $1 - \delta$, the below holds $\forall k \in \{1, \dots, K\}, \forall m \neq \ell$,

$$\frac{\widehat{\mathbf{R}}_{m,\ell}(:, k)}{\|\widehat{\mathbf{R}}_{m,\ell}(:, k)\|_1} = \frac{\mathbf{R}_{m,\ell}(:, k)}{\|\mathbf{R}_{m,\ell}(:, k)\|_1} + \overline{\mathbf{N}}_{m,\ell}(:, k)$$

where $\|\overline{\mathbf{N}}_{m,\ell}(:, k)\|_2 \leq \frac{2\sqrt{K}(1+\sqrt{\ln(1/\delta)})}{\sqrt{S}\eta}$.

Proof. For simpler representation, let us assign $\mathbf{x} := \mathbf{R}_{m,\ell}(:, k)$, $\widehat{\mathbf{x}} := \widehat{\mathbf{R}}_{m,\ell}(:, k)$ and $\mathbf{n} := \widehat{\mathbf{R}}_{m,\ell}(:, k) - \mathbf{R}_{m,\ell}(:, k) = \mathbf{N}_{m,\ell}(:, k)$.

Let $\mathbf{x} = [x_1, \dots, x_K]^\top$ and $\mathbf{n} = [n_1, \dots, n_K]^\top$. Note that $\mathbf{x} \geq 0$ and $\mathbf{x} + \mathbf{n} \geq 0$ —since \mathbf{x} is a legitimate PMF and $\mathbf{x} + \mathbf{n}$ is averaged from co-occurrence counts. Then, we have

$$\begin{aligned} \frac{\widehat{\mathbf{x}}}{\|\widehat{\mathbf{x}}\|_1} &= \frac{\mathbf{x} + \mathbf{n}}{\|\mathbf{x} + \mathbf{n}\|_1} = \frac{\mathbf{x} + \mathbf{n}}{\sum_i x_i + n_i} = \frac{\mathbf{x} + \mathbf{n}}{\sum_i x_i + \sum_i n_i} \\ &= \frac{\mathbf{x} + \mathbf{n}}{\sum_i x_i \left(1 + \frac{\sum_i n_i}{\sum_i x_i}\right)}. \end{aligned}$$

Let $\mu = \frac{\sum_i n_i}{\sum_i x_i}$. Using the assumption $\|\mathbf{N}_{m,\ell}(:, k)\|_1 < \|\mathbf{R}_{m,\ell}(:, k)\|_1$, then $|\mu| < 1$. From this,

$$\begin{aligned} \frac{\mathbf{x} + \mathbf{n}}{\|\mathbf{x} + \mathbf{n}\|_1} &= \frac{(\mathbf{x} + \mathbf{n})(1 + \mu)^{-1}}{\sum_i x_i} = \frac{(\mathbf{x} + \mathbf{n})}{\sum_i x_i} (1 - \mu + \mu^2 - \mu^3 + \dots) \\ &= \frac{\mathbf{x}}{\sum_i x_i} - \mu \frac{\mathbf{x}}{\sum_i x_i} (1 - \mu + \mu^2 - \mu^3 + \dots) + \frac{\mathbf{n}}{\sum_i x_i} (1 - \mu + \mu^2 - \mu^3 + \dots) \\ &= \frac{\mathbf{x}}{\sum_i x_i} + \frac{\mathbf{n}}{(1 + \mu) \sum_i x_i} - \frac{\mu \mathbf{x}}{(1 + \mu) \sum_i x_i} \\ &= \frac{\mathbf{x}}{\|\mathbf{x}\|_1} + \underbrace{\frac{\mathbf{n} - \mu \mathbf{x}}{(1 + \mu) \sum_i x_i}}_{\Gamma} \end{aligned} \tag{20}$$

Now let us bound the term Γ :

$$\begin{aligned} \|\Gamma\|_1 &:= \left\| \frac{\mathbf{n} - \mu \mathbf{x}}{(1 + \mu) \sum_i x_i} \right\|_1 \leq \frac{\|\mathbf{n}\|_1}{\|\mathbf{x} + \mathbf{n}\|_1} + \frac{\|\sum_i n_i\|}{\|\mathbf{x} + \mathbf{n}\|_1} \\ &\leq 2 \frac{\|\mathbf{n}\|_1}{\|\mathbf{x} + \mathbf{n}\|_1} \\ &\leq 2 \frac{\|\mathbf{n}\|_1}{\eta}. \end{aligned} \tag{21}$$

The first and second inequalities are due to Cauchy-Schwartz inequality and the last inequality is by $\|\widehat{\mathbf{R}}_{m,\ell}(:, k)\|_1 \geq \eta$.

From Lemma 2, with probability greater than $1 - \delta$, the below holds,

$$\sum_{k=1}^K \|\mathbf{N}_{m,\ell}(:, k)\|_2^2 = \|\mathbf{N}_{m,\ell}\|_F^2 \leq \frac{(1 + \sqrt{\ln(1/\delta)})^2}{S}.$$

By norm equivalence, $\frac{\|\mathbf{N}_{m,\ell}(:, k)\|_1}{\sqrt{K}} \leq \|\mathbf{N}_{m,\ell}(:, k)\|_2$. Therefore,

$$\begin{aligned} \sum_{k=1}^K \|\mathbf{N}_{m,\ell}(:, k)\|_1^2 &\leq \frac{K(1 + \sqrt{\ln(1/\delta)})^2}{S} \\ \implies \|\mathbf{N}_{m,\ell}(:, k)\|_1^2 &\leq \frac{K(1 + \sqrt{\ln(1/\delta)})^2}{S}, \quad \forall k \\ \implies \|\mathbf{N}_{m,\ell}(:, k)\|_1 &\leq \frac{\sqrt{K}(1 + \sqrt{\ln(1/\delta)})}{\sqrt{S}}, \quad \forall k \end{aligned} \tag{22}$$

From (21) and (22),

$$\|\Gamma\|_1 = \|\overline{\mathbf{N}}_{m,\ell}(:,k)\|_1 \leq \frac{2\sqrt{K}(1 + \sqrt{\ln(1/\delta)})}{\sqrt{S}\eta}$$

By norm equivalence, $\|\overline{\mathbf{N}}_{m,\ell}(:,k)\|_2 \leq \|\overline{\mathbf{N}}_{m,\ell}(:,k)\|_1$, then

$$\|\overline{\mathbf{N}}_{m,\ell}(:,k)\|_2 \leq \frac{2\sqrt{K}(1 + \sqrt{\ln(1/\delta)})}{\sqrt{S}\eta}$$

This completes the proof of Lemma 3. \square

With the above lemmas, we are ready to characterize the accuracy of applying SPA to identify the Dawid-Skene model given the assumptions in Eq. (11).

Eq. (11) indicates that there exists a set of indices $\Lambda_q = \{q_1, \dots, q_K\}$ such that

$$\overline{\mathbf{H}}_m(\Lambda_q, :) = \mathbf{I}_K + \mathbf{E} \quad (23)$$

where \mathbf{I}_K is the identity matrix of size K and \mathbf{E} is the error matrix with $\max_j |\mathbf{E}(l, j)| = \|\mathbf{E}(l, :)\|_\infty \leq \epsilon$. By norm equivalence, we have $\|\mathbf{E}(l, :)\|_2 \leq \sqrt{K} \|\mathbf{E}(l, :)\|_\infty \leq \sqrt{K}\epsilon$.

Without loss of generality, let us assume $\Lambda_q = \{1, \dots, K\}$ and

$$\overline{\mathbf{H}}_m = \begin{bmatrix} \mathbf{I}_K + \mathbf{E} \\ \mathbf{H}_m^* \end{bmatrix}.$$

Now we have,

$$\overline{\mathbf{Z}}_m = \mathbf{A}_m \overline{\mathbf{H}}_m^\top + \overline{\mathbf{N}} \quad (24)$$

$$= \mathbf{A}_m [\mathbf{I}_K + \mathbf{E}^\top \quad (\mathbf{H}_m^*)^\top] + \overline{\mathbf{N}} \quad (25)$$

$$= \mathbf{A}_m [\mathbf{I}_K \quad (\mathbf{H}_m^*)^\top] + [\mathbf{A}_m \mathbf{E}^\top \quad \mathbf{0}] + \overline{\mathbf{N}} \quad (26)$$

where $\overline{\mathbf{N}} = [\overline{\mathbf{N}}_{m,m_1}, \dots, \overline{\mathbf{N}}_{m,m_{T(m)}}]$ and the zero matrix $\mathbf{0}$ has the same size as that of \mathbf{H}_m^* .

This model is similar to the noisy NMF model, i.e., $\mathbf{X} = \mathbf{W}\mathbf{H}^\top + \mathbf{N}$, where the noise matrix $\mathbf{N} = [\mathbf{A}_m \mathbf{E}^\top \quad \mathbf{0}_m^*] + \overline{\mathbf{N}}$. To be specific, we have

$$\mathbf{N}(:,l) = \mathbf{A}_m \mathbf{E}(l, :)\top + \overline{\mathbf{N}}(:,l) \quad (27)$$

Therefore, one can see that

$$\|\mathbf{N}(:,l)\|_2 = \|\mathbf{A}_m \mathbf{E}(l, :)\top + \overline{\mathbf{N}}(:,l)\|_2 \quad (28)$$

$$\leq \|\mathbf{A}_m\|_2 \|\mathbf{E}(l, :)\|_2 + \|\overline{\mathbf{N}}(:,l)\|_2 \quad (29)$$

$$\leq \sigma_{\max}(\mathbf{A}_m) \sqrt{K}\epsilon + \frac{2\sqrt{K}(1 + \sqrt{\ln(1/\delta)})}{\sqrt{S}\eta} \quad (30)$$

where (29) is by the Cauchy-Schwartz inequality, (30) is by Lemma 3 and $\sigma_{\max}(\mathbf{A}_m)$ is the largest singular value of matrix \mathbf{A}_m .

Hence, we effectively have the same model as in (17). Applying Lemma 1, we see that if

$$\begin{aligned} \epsilon &\leq \frac{1}{\sqrt{K}\kappa(\mathbf{A}_m)} \min\left(\frac{1}{2\sqrt{K}-1}, \frac{1}{4}\right) (1 + 80\kappa^2(\mathbf{A}_m))^{-1} - \frac{2(1 + \sqrt{\ln(1/\delta)})}{\sigma_{\max}(\mathbf{A}_m)\sqrt{S}\eta} \\ &= \mathcal{O}\left(\max\left(K^{-1}\kappa^{-3}(\mathbf{A}_m), \sqrt{\ln(1/\delta)}(\sigma_{\max}(\mathbf{A}_m)\sqrt{S}\eta)^{-1}\right)\right), \end{aligned}$$

then, with probability at least $1 - \delta$, the SPA algorithm identifies the matrix \mathbf{A}_m with an error bound given by

$$\begin{aligned}
& \min_{\Pi} \|\widehat{\mathbf{A}}_m \Pi - \mathbf{A}_m\|_{2,\infty} \\
&= \max_{1 \leq j \leq K} \min_{\widehat{q}_k \in \widehat{\mathcal{A}}_q} \|\mathbf{A}_m(:, j) - \overline{\mathbf{Z}}_m(:, \widehat{q}_k)\|_2 \\
&\leq \left(\sigma_{\max}(\mathbf{A}_m) \sqrt{K} \epsilon + \frac{2\sqrt{K}(1 + \sqrt{\ln(1/\delta)})}{\sqrt{S\eta}} \right) (1 + 80\kappa^2(\mathbf{A}_m)) \\
&= \mathcal{O}\left(\sqrt{K}\kappa^2(\mathbf{A}_m) \max\left(\sigma_{\max}(\mathbf{A}_m)\epsilon, \sqrt{\ln(1/\delta)}(\sqrt{S\eta})^{-1} \right) \right).
\end{aligned}$$

This completes the proof.

G Proof of Theorem 2

Assuming that the rows of $\overline{\mathbf{H}}_m$ are generated from the probability simplex uniformly at random, we now analyze under what conditions vectors close to all K vertices of the probability simplex appear in the rows of $\overline{\mathbf{H}}_m$.

Let us denote the probability simplex as $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^K | \mathbf{x}^\top \mathbf{1} = 1, \mathbf{x} \geq 0\}$.

Let us consider an ϵ -neighbourhood of the k -th vertex \mathbf{e}_k denoted as $\mathcal{Q}_k(\epsilon)$ such that

$$\mathcal{Q}_k(\epsilon) := \{\mathbf{q} \in \mathcal{X} | \|\mathbf{q} - \mathbf{e}_k\|_2 \leq \epsilon\} \quad (31)$$

Geometrically, the continuous set $\mathcal{Q}_k(\epsilon)$ can be considered as the intersection of the probability simplex \mathcal{X} and the euclidean ball of radius ϵ centered at \mathbf{e}_k , i.e., $\mathcal{Q}_k(\epsilon) = \mathcal{X} \cap \mathcal{B}(\mathbf{e}_k, \epsilon)$ (see Fig. 5).




Figure 5: The big triangle represents the probability simplex \mathcal{X} when $K = 3$, the dotted circles denotes the euclidean balls $\mathcal{B}(\mathbf{e}_k, \epsilon)$, the shaded region denotes $\mathcal{X} \cap \mathcal{B}(\mathbf{e}_k, \epsilon)$. The small triangles near the vertices has the same volume as the simplex having edge lengths ϵ denoted as \mathcal{X}_ϵ

Suppose we are uniformly sampling a set \mathcal{P} of size s from the probability simplex \mathcal{X} such that $\mathcal{P} := \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_s\}$

Let us define an event J_i such that for every $i \in \{1, \dots, s\}$,

$$J_i = \begin{cases} 1, & \text{if } \mathbf{p}_i \in \mathcal{Q}_k(\epsilon) \\ 0, & \text{otherwise} \end{cases} \quad (32)$$

Consider the probability such that event J_i happens,

$$\Pr(J_i = 1) = \frac{\text{vol}(\mathcal{Q}_k(\epsilon))}{\text{vol}(\mathcal{X})} \quad (33)$$

$$= \frac{\text{vol}(\mathcal{X} \cap \mathcal{B}(e_k, \epsilon))}{\text{vol}(\mathcal{X})} \quad (34)$$

$$\geq \frac{\text{vol}(\mathcal{X}_\epsilon)}{\text{vol}(\mathcal{X})} \quad (35)$$

$$\geq \left(\frac{\epsilon}{\sqrt{2}}\right)^{K-1} \quad (36)$$

where \mathcal{X}_ϵ denotes the $(K-1)$ -dimensional simplex which intersects the co-ordinate axes at $\frac{\epsilon}{\sqrt{2}}e_k$, for every $k \in \{1, \dots, K\}$ and its volume is given by $\frac{(\epsilon/\sqrt{2})^{K-1}}{(K-1)!}$ [36]. (Note that the probability simplex \mathcal{X} intersects the co-ordinate axes at e_k for every k). The inequality in Eq.(35) uses the geometric property that the volume of $\mathcal{X} \cap \mathcal{B}(e_k, \epsilon)$ is greater than the volume of \mathcal{X}_ϵ (see Fig. 5).

Let us define the random variable $U = \sum_{i=1}^s J_i$. Then,

$$\mathbb{E}[U] = \mathbb{E}\left[\sum_{i=1}^s J_i\right] = \sum_{i=1}^s \mathbb{E}[J_i] \quad (37)$$

$$= \sum_{i=1}^s \Pr(J_i = 1) = s\Pr(J_i = 1) \quad (38)$$

$$\geq s \left(\frac{\epsilon}{\sqrt{2}}\right)^{K-1} \quad (39)$$

Now, if there exists at least one sample from set \mathcal{P} which is in the ϵ -neighbourhood of k -th vertex, ie the event J_i happens at least once, then $U = \sum_{i=1}^s J_i \geq 1$. We are interested in finding the below probability,

$$\Pr(U \geq 1) = 1 - \Pr(U < 1) \quad (40)$$

$$= 1 - \Pr(U \leq 0) \quad (41)$$

$$= 1 - \Pr(U = 0) \quad (42)$$

So, our goal boils down to finding $\Pr(U \leq 0)$ and we will achieve this using Chernoff-Hoeffding bound.

Lemma 4. [37] *Let J_1, \dots, J_s be independent bounded random variables such that J_i falls in the interval $[a_i, b_i]$ with probability one and let $U = \sum_{i=1}^s J_i$. Then for any $t > 0$,*

$$\Pr(U - \mathbb{E}[U] \leq -t) \leq e^{-2t^2 / \sum_{i=1}^s (b_i - a_i)^2} \quad (43)$$

It follows that

$$\Pr(U \leq \mathbb{E}[U] - t) \leq e^{-2t^2 / \sum_{i=1}^s (b_i - a_i)^2} \quad (44)$$

By assigning $\mathbb{E}[U] - t = 0$, we get $t = \mathbb{E}[U]$. Also, notice that in our case $b_i = 1$, $a_i = 0$, then

$$\Pr(U \leq 0) \leq e^{-\frac{2\mathbb{E}[U]^2}{s}} \quad (45)$$

$$\leq e^{-\frac{s\epsilon^2(K-1)}{2^{K-2}}} \quad (46)$$

Eq. (46) is obtained by using the inequality (39) and implies that, the probability such that the uniform sample \mathcal{P} does not contain any points from k -th vertex is less than $e^{-\frac{s\epsilon^2(K-1)}{2^{K-2}}}$.

Now we have to find the corresponding probability that considers all the K vertices.

For this, let us define events E_k as follows,

$$E_k = \{\text{There exists no point } \mathbf{p} \text{ in the uniform sample set } \mathcal{P} \text{ such that } \mathbf{p} \in \mathcal{Q}_k(\epsilon)\} \quad (47)$$

From Eq. (46), it is clear that $\Pr(E_k) \leq e^{-\frac{s\epsilon^{2(K-1)}}{2^{K-2}}}$. Since the points are uniformly sampled from the probability simplex \mathcal{X} , this bound is applicable for all $k \in \{1, \dots, K\}$.

Now let us define the event E as below

$$E = \{\text{there exists at least one point in the uniform sample set } \mathcal{P} \text{ such that } \mathbf{p} \in \mathcal{Q}_k(\epsilon) \text{ for each } k\}$$

We can observe that $E = \bigcap_{k=1}^K \overline{E_k}$ where $\overline{E_k}$ is the complement of the event E_k .

Therefore,

$$\begin{aligned} \Pr(E) &= \Pr\left(\bigcap_{k=1}^K \overline{E_k}\right) = \Pr(\overline{\bigcup_{k=1}^K E_k}) \\ &= 1 - \Pr\left(\bigcup_{k=1}^K E_k\right) \\ &\geq 1 - \sum_k \Pr(E_k) \\ &\geq 1 - Ke^{-\frac{s\epsilon^{2(K-1)}}{2^{K-2}}} \end{aligned} \quad (48)$$

Eq. (48) implies that with probability greater than or equal to $1 - Ke^{-\frac{s\epsilon^{2(K-1)}}{2^{K-2}}}$, the points from the ϵ -neighbourhood of all the vertices are contained by set \mathcal{P} .

If s represents the number of rows in $\overline{\mathbf{H}}_m$, then for $s \geq \frac{2^{K-2}}{\epsilon^{2(K-1)}} \log\left(\frac{K}{\rho}\right)$, with probability at least $1 - \rho$, a uniform sample from the probability simplex \mathcal{X} will contain ϵ -near-vertex points of all the K vertices. Note that $s = (M - 1)K$ where M is the number of annotators. This provides a bound on the number of annotators needed. Specifically, if there exists at least $1 + \frac{2^{K-2}}{K\epsilon^{2(K-1)}} \log\left(\frac{K}{\rho}\right)$ annotators, then we have the conclusion of Theorem 2

H Proof of Theorem 3

To show this theorem, we will use the following Lemma:

Lemma 5. [18] Consider a matrix factorization model $\mathbf{R} = \mathbf{P}_1 \mathbf{P}_2^\top$, where $\mathbf{R} \in \mathbb{R}^{M \times N}$, $\mathbf{P}_1 \in \mathbb{R}^{M \times K}$, $\mathbf{P}_2 \in \mathbb{R}^{N \times K}$, and $\text{rank}(\mathbf{P}_1) = \text{rank}(\mathbf{P}_2) = K$. If $\mathbf{P}_i \geq \mathbf{0}$ for $i = 1, 2$ and both \mathbf{P}_1 and \mathbf{P}_2 are sufficiently scattered, we have any $\hat{\mathbf{P}}_1 \geq \mathbf{0}$ and $\hat{\mathbf{P}}_2 \geq \mathbf{0}$ that satisfy $\mathbf{R} = \hat{\mathbf{P}}_1 \hat{\mathbf{P}}_2^\top$ must have the following form

$$\hat{\mathbf{P}}_1 = \mathbf{P}_1 \mathbf{\Pi} \mathbf{\Sigma}, \quad \hat{\mathbf{P}}_2 = \mathbf{P}_2 \mathbf{\Pi} \mathbf{\Sigma}^{-1}, \quad (49)$$

where $\mathbf{\Pi}$ is a permutation matrix and $\mathbf{\Sigma}^{-1}$ is a diagonal nonnegative singular matrix.

Lemma 5 addresses the identifiability of a conventional nonnegative matrix factorization (NMF) model. Simply speaking, if both latent factors of $\mathbf{R} = \mathbf{P}_1 \mathbf{P}_2^\top$ are sufficiently scattered, then the NMF of \mathbf{R} is unique up to column permutation and scaling of the latent factors.

Now we start proving Theorem 3. Let us consider the following matrix \mathbf{R} :

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{m_1, \ell_1} & \mathbf{R}_{m_1, \ell_2} & \cdots & \mathbf{R}_{m_1, \ell_{|\mathcal{P}_2|}} \\ \vdots & \vdots & \cdots & \vdots \\ \mathbf{R}_{m_{|\mathcal{P}_1|}, \ell_1} & \mathbf{R}_{m_{|\mathcal{P}_1|}, \ell_2} & \cdots & \mathbf{R}_{m_{|\mathcal{P}_1|}, \ell_{|\mathcal{P}_2|}} \end{bmatrix}$$

It is readily seen that

$$\begin{aligned} \mathbf{R} &= \begin{bmatrix} \mathbf{A}_{m_1} \\ \vdots \\ \mathbf{A}_{m_{|\mathcal{P}_1|}} \end{bmatrix} \mathbf{D}[\mathbf{A}_{\ell_1}^\top, \dots, \mathbf{A}_{\ell_{|\mathcal{P}_2|}}^\top] \\ &= \mathbf{H}^{(1)} \mathbf{D}(\mathbf{H}^{(2)})^\top. \end{aligned}$$

It suffices to show that $\mathbf{R} = \mathbf{H}^{(1)}\mathbf{D}(\mathbf{H}^{(2)})^\top$ is unique up to column permutations of $\mathbf{H}^{(i)}$ for $i = 1, 2$. The reason is that if such uniqueness holds, then \mathbf{A}_m for $m \notin \mathcal{P}_1 \cup \mathcal{P}_2$ can be identified with via solving

$$\mathbf{R}_{m,r} = \mathbf{A}_m \mathbf{D} \mathbf{A}_r^\top, \quad r \in \mathcal{P}_1 \cup \mathcal{P}_2,$$

up to the same column permutation.

Note that since \mathbf{D} and \mathbf{A}_r have been identified from \mathbf{R} , solving \mathbf{A}_m amounts to solving a system of linear equations, which has a unique solution under $\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{A}_m) = K$. Under the assumption that all $m \notin \mathcal{P}_1 \cup \mathcal{P}_2$ are connected to a certain $r \in \mathcal{P}_1 \cup \mathcal{P}_2$ via $\mathbf{R}_{m,r}$, all the \mathbf{A}_m 's for $m \notin \mathcal{P}_1 \cup \mathcal{P}_2$ can be identified up to the same column permutation.

To show the identifiability of $\mathbf{R} = \mathbf{H}^{(1)}\mathbf{D}(\mathbf{H}^{(2)})^\top$, consider re-writing the right hand side as

$$\begin{aligned} \mathbf{R} &= \underbrace{\mathbf{H}^{(1)}\mathbf{D}^{1/2}}_{\mathbf{P}_1} \underbrace{\mathbf{D}^{1/2}(\mathbf{H}^{(2)})^\top}_{\mathbf{P}_2^\top} \\ &= \mathbf{P}_1 \mathbf{P}_2^\top. \end{aligned}$$

It suffices to show that \mathbf{P}_1 and \mathbf{P}_2 are unique up to column permutation and scaling, since the constraints on $\mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top (\Rightarrow \mathbf{1}^\top \mathbf{H}^{(i)} = |\mathcal{P}_i|)$ removes the scaling ambiguity.

Assume that there is an alternative solution $\mathbf{R} = \hat{\mathbf{P}}_1 \hat{\mathbf{P}}_2^\top$. By the fact $\text{rank}(\mathbf{A}_m) = K$, we have

$$\text{rank}(\mathbf{H}^{(i)}) = K \Rightarrow \text{rank}(\mathbf{P}_i) = K.$$

Note that $\mathbf{H}^{(1)}$ and $\mathbf{H}^{(2)}$ are both sufficiently scattered. This directly implies that both \mathbf{P}_1 and \mathbf{P}_2 are sufficiently scattered, since column scaling of $\mathbf{H}^{(1)}$ and $\mathbf{H}^{(2)}$ does not affect the cone of their respective transposes, i.e.,

$$\text{cone} \left\{ (\mathbf{H}^{(i)})^\top \right\} = \text{cone} \left\{ (\mathbf{H}^{(i)} \boldsymbol{\Sigma})^\top \right\}$$

for any nonnegative, nonsingular and diagonal $\boldsymbol{\Sigma}$.

Then, by Lemma 5, it must hold that

$$\hat{\mathbf{P}}_1 = \mathbf{H}^{(1)} \mathbf{D}^{1/2} \boldsymbol{\Pi} \boldsymbol{\Sigma}, \quad \hat{\mathbf{P}}_2 = \mathbf{H}^{(2)} \mathbf{D}^{1/2} \boldsymbol{\Pi} \boldsymbol{\Sigma}^{-1},$$

for a certain $\boldsymbol{\Pi}$ and $\boldsymbol{\Sigma}$. Nevertheless, $\boldsymbol{\Sigma}$ is automatically removed by the constraints $\mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top$.

I Proof of Theorem 4

Let $\overline{\mathbf{H}}$ represent the row normalized matrix of \mathbf{H} . Geometrically, $\overline{\mathbf{H}}$ can be viewed as the projection of the rows of \mathbf{H} onto the $(K - 1)$ -probability simplex. (cf. Fig. 2).

The work [28] provides a characterization to the spread of the rows of $\overline{\mathbf{H}}$ in the probability simplex using a measure called as uniform pixel purity level (named in the context of hyperspectral imaging). The purity level is denoted by γ and defined as follows:

$$\gamma = \sup \{ r \leq 1 \mid \mathcal{R}(r) \subseteq \text{conv} \{ \overline{\mathbf{H}}^\top \} \} \quad (50)$$

where

$$\mathcal{R}(r) = \{ \mathbf{x} \in \mathbb{R}^K \mid \|\mathbf{x}\|_2 \leq r \} \cap \text{conv} \{ \mathbf{e}_1, \dots, \mathbf{e}_K \} \quad (51)$$

Geometrically, $\mathcal{R}(\gamma)$ is the ‘largest’ $\mathcal{R}(r)$ that can be inscribed inside $\text{conv} \{ \overline{\mathbf{H}}^\top \}$. There is an interesting link between γ and the sufficiently scattered condition shown in [12]:

Lemma 6. [12] Assume $K \geq 3$ holds. If $\gamma \geq \frac{1}{\sqrt{K-1}}$, then $\overline{\mathbf{H}}$ is sufficiently scattered.

In general cases, it is hard to check if $\gamma \geq \frac{1}{\sqrt{K-1}}$ holds [18]. However, in [28], a sufficient condition for $\gamma \geq \frac{1}{\sqrt{K-1}}$ is derived:

Lemma 7. [28] Suppose the following assumption holds true: for every $k, k' \in \{1, \dots, K\}, k \neq k'$, there exist a row index $q_{kk'}$ in $\overline{\mathbf{H}}$ such that

$$\overline{\mathbf{H}}(q_{kk'}, :) = \alpha_{kk'} \mathbf{e}_k^\top + (1 - \alpha_{kk'}) \mathbf{e}_{k'}^\top \quad (52)$$

where $\frac{1}{2} < \alpha_{kk'} < 1$ for $K \geq 4$, $\frac{2}{3} < \alpha_{kk'} < 1$ for $K = 3$ and $\alpha_{kk'} = 1$, for $K = 2$. Then $\gamma \geq \frac{1}{\sqrt{K-1}}$ and by Lemma 6, $\overline{\mathbf{H}}$ is sufficiently scattered.




Figure 6: (ε -sufficiently scattered) The big triangle represents the probability simplex \mathcal{X} when $K = 3$, the shaded region denotes the region which is ε near the edges and in the inset, the shaded region depicts the lower bound for the volume (area) of this region at each vertex.

Let us define $\alpha = \min_{k, k' \in \{1, \dots, K\}, k \neq k'} \alpha_{kk'}$. Lemma 7 states that for every edges of the probability simplex, if there exists at least one row in $\overline{\mathbf{H}}$ which belongs to certain range in the edge which is of length α (α -edge) (see Fig. 6), then the matrix $\overline{\mathbf{H}}$ is sufficiently scattered.

Let us denote the probability simplex as $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^K | \mathbf{x}^\top \mathbf{1} = 1, \mathbf{x} \geq 0\}$.

For each vertex, there exists $K - 1$ edges associated to it. Let us denote an ε -neighbourhood of α -edge connecting the vertices k and k' as $\tilde{\mathcal{Q}}_{k, k'}(\varepsilon, \alpha)$. By the conditions in Lemma 7 and the definition of ε -sufficiently scattered (cf. Def. 2), it can be seen that, for every edges connecting k and k' , if there exists at least one row in $\overline{\mathbf{H}}$ belonging to $\tilde{\mathcal{Q}}_{k, k'}(\varepsilon, \alpha)$, then $\overline{\mathbf{H}}$ is ε -sufficiently scattered.

For each vertex k , the union of $\tilde{\mathcal{Q}}_{k, k'}(\varepsilon, \alpha), k' = \{1, \dots, K - 1\}$ forms a continuous neighbourhood around the vertex k denoted as $\tilde{\mathcal{Q}}_k(\varepsilon, \alpha)$, i.e.,

$$\tilde{\mathcal{Q}}_k(\varepsilon, \alpha) = \bigcup_{k'=1}^{K-1} \tilde{\mathcal{Q}}_{k, k'}(\varepsilon, \alpha) \quad (53)$$

Geometrically, the volume of the continuous set $\tilde{\mathcal{Q}}_k(\varepsilon, \alpha)$ can be lower bounded as below (see Fig 6)

$$\text{vol}(\tilde{\mathcal{Q}}_k(\varepsilon, \alpha)) \geq \text{vol}(\mathcal{X}_\alpha) - \text{vol}(\mathcal{X}_{\alpha-2\varepsilon}) \quad (54)$$

where $\mathcal{X}_{\alpha'}$ is $(K - 1)$ -dimensional simplex which intersects the co-ordinate axes at $\frac{\alpha'}{\sqrt{2}} \mathbf{e}_k$ for every $k = \{1, \dots, K\}$ and thus has the edge lengths α' . The volume of $\mathcal{X}_{\alpha'}$ is given by $\frac{(\alpha'/\sqrt{2})^{K-1}}{(K-1)!}$ [36]

Eq. (54) can then be written as

$$\text{vol}(\tilde{\mathcal{Q}}_k(\varepsilon, \alpha)) \geq \frac{\alpha^{K-1}}{\sqrt{2}^{K-1}(K-1)!} - \frac{(\alpha - 2\varepsilon)^{K-1}}{\sqrt{2}^{K-1}(K-1)!} \quad (55)$$

$$= \frac{\alpha^{K-1}}{\sqrt{2}^{K-1}(K-1)!} \left(1 - \left(1 - \frac{2\varepsilon}{\alpha} \right)^{K-1} \right) \quad (56)$$

$$= \frac{\alpha^{K-1}}{\sqrt{2}^{K-1}(K-1)!} \left(\frac{2\varepsilon}{\alpha} \left(1 + \left(1 - \frac{2\varepsilon}{\alpha} \right) + \dots + \left(1 - \frac{2\varepsilon}{\alpha} \right)^{K-2} \right) \right) \quad (57)$$

$$= \frac{\alpha^{K-1}}{\sqrt{2}^{K-1}(K-1)!} \left(\frac{2\varepsilon}{\alpha} + \frac{2\varepsilon}{\alpha} \left(1 - \frac{2\varepsilon}{\alpha} \right) + \dots + \frac{2\varepsilon}{\alpha} \left(1 - \frac{2\varepsilon}{\alpha} \right)^{K-2} \right) \quad (58)$$

$$\geq \frac{\alpha^{K-1}}{\sqrt{2}^{K-1}(K-1)!} \frac{2\varepsilon}{\alpha} \quad (59)$$

$$= \frac{\alpha^{K-2}\varepsilon}{\sqrt{2}^{K-3}(K-1)!} \quad (60)$$

Eq. (57) uses the geometric series sum formula $1 - a^n = (1 - a)(1 + a + \dots + a^{n-1})$ and the assumption that $\alpha > 2\varepsilon$.

From Eq. (53) and (60), the volume of the set $\tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha)$ can be lower bounded as

$$\text{vol} \left(\bigcup_{k'=1}^{K-1} \tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha) \right) = \text{vol} \left(\tilde{\mathcal{Q}}_k(\varepsilon, \alpha) \right) \quad (61)$$

$$\implies (K-1)\text{vol} \left(\tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha) \right) \geq \text{vol} \left(\tilde{\mathcal{Q}}_k(\varepsilon, \alpha) \right) \quad (62)$$

$$\geq \frac{\alpha^{K-2}\varepsilon}{\sqrt{2}^{K-3}(K-1)!} \quad (63)$$

$$\implies \text{vol} \left(\tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha) \right) \geq \frac{\alpha^{K-2}\varepsilon}{\sqrt{2}^{K-3}(K-1)(K-1)!} \quad (64)$$

Suppose we are uniformly sampling a set \mathcal{P} of size s from the probability simplex \mathcal{X} such that $\mathcal{P} := \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_s\}$

Let us define an event J_i such that for every $i \in \{1, \dots, s\}$,

$$J_i = \begin{cases} 1, & \text{if } \mathbf{p}_i \in \tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha) \\ 0, & \text{otherwise} \end{cases} \quad (65)$$

Consider the probability such that event J_i happens,

$$\Pr(J_i = 1) = \frac{\text{vol}(\tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha))}{\text{vol}(\mathcal{X})} \quad (66)$$

$$\geq (K-1)! \frac{\alpha^{K-2}\varepsilon}{\sqrt{2}^{K-3}(K-1)(K-1)!} \quad (67)$$

$$= \frac{\alpha^{K-2}\varepsilon}{(K-1)\sqrt{2}^{K-3}} \quad (68)$$

Eq. (68) uses the fact that the volume of the $(K-1)$ -dimensional simplex \mathcal{X} is given by $\frac{1}{(K-1)!}$ [36].

Now, let us define the random variable $U = \sum_{i=1}^s J_i$. Then,

$$\mathbb{E}[U] = \mathbb{E}\left[\sum_{i=1}^s J_i\right] = \sum_{i=1}^s \mathbb{E}[J_i] \quad (69)$$

$$= \sum_{i=1}^s \Pr(J_i = 1) = s\Pr(J_i = 1) \quad (70)$$

$$\geq s \frac{\alpha^{K-2}\varepsilon}{(K-1)\sqrt{2}^{K-3}} \quad (71)$$

Now, if there exists at least one sample from set \mathcal{P} which is in the ε -neighbourhood of α -edge, i.e., the event J_i happens at least once, then $U = \sum_{i=1}^s J_i \geq 1$. Also,

$$\Pr(U \geq 1) = 1 - \Pr(U < 1) \quad (72)$$

$$= 1 - \Pr(U \leq 0) \quad (73)$$

$$= 1 - \Pr(U = 0) \quad (74)$$

$$(75)$$

So, our goal boils down to finding $\Pr(U \leq 0)$ and we will achieve this using Lemma 4.

From Lemma 4, it follows that

$$\Pr(U \leq \mathbb{E}[U] - t) \leq e^{-2t^2 / \sum_{i=1}^s (b_i - a_i)^2} \quad (76)$$

By assigning $\mathbb{E}[U] - t = 0$, we get $t = \mathbb{E}[U]$. Also, notice that in our case $b_i = 1$, $a_i = 0$, then

$$\Pr(U \leq 0) \leq e^{-\frac{2\mathbb{E}[U]^2}{s}} \quad (77)$$

$$\leq e^{-\frac{s\alpha^{2(K-2)}\varepsilon^2}{2^{K-4}(K-1)^2}} \quad (78)$$

Eq. (78) is obtained by using the inequality $\mathbb{E}[U] \geq s \frac{\alpha^{K-2}\varepsilon}{(K-1)\sqrt{2}^{K-3}}$ as in Eq. (71) and implies that, the probability such that the uniform sample \mathcal{P} does not contain any points from $\tilde{\mathcal{Q}}_{k,k'}(\varepsilon, \alpha)$ is less than $e^{-\frac{s\alpha^{2(K-2)}\varepsilon^2}{2^{K-4}(K-1)^2}}$.

Now we have to find the corresponding probability that considers all the $(K-1)$ edges for each vertex k .

For this, let us define events $\tilde{E}_{kk'}$ as follows,

$$\tilde{E}_{kk'} = \{\text{There exists no point } \mathbf{p} \text{ in the uniform sample set } \mathcal{P} \text{ such that } \mathbf{p} \in \mathcal{Q}_{k,k'}(\varepsilon, \alpha)\} \quad (79)$$

From Eq. (78), it is clear that $\Pr(\tilde{E}_{kk'}) \leq e^{-\frac{s\alpha^{2(K-2)}\varepsilon^2}{2^{K-4}(K-1)^2}}$. Since the points are uniformly sampled from the probability simplex \mathcal{X} , this bound is applicable for all $k, k' \in \{1, \dots, K\}$, $k \neq k'$.

Now let us define the event \tilde{E} as below

$$\tilde{E} = \{\text{there exists at least one point in the set } \mathcal{P} \text{ such that } \mathbf{p} \in \mathcal{Q}_{k,k'}(\varepsilon, \alpha) \text{ for all } k, k', k \neq k'\}$$

We can observe that $\tilde{E} = \bigcap_{\substack{k,k' \\ k \neq k'}} \overline{\tilde{E}_{kk'}}$ where $\overline{\tilde{E}_{kk'}}$ is the complement of the event $\tilde{E}_{kk'}$.

Therefore,

$$\begin{aligned}
\Pr(E) &= \Pr\left(\bigcap_{\substack{k,k' \\ k \neq k'}} \widetilde{E}_{kk'}\right) = \Pr\left(\overline{\bigcup_{\substack{k,k' \\ k \neq k'}} \widetilde{E}_{kk'}}\right) \\
&= 1 - \Pr\left(\bigcup_{\substack{k,k' \\ k \neq k'}} \widetilde{E}_{kk'}\right) \\
&\geq 1 - \sum_{k,k'} \Pr(\widetilde{E}_{kk'}) \\
&\geq 1 - K(K-1)e^{-\frac{s\alpha^2(K-2)\epsilon^2}{2^{K-4}(K-1)^2}} \tag{80}
\end{aligned}$$

Eq. (80) implies that with probability greater than or equal to $1 - K(K-1)e^{-\frac{s\alpha^2(K-2)\epsilon^2}{2^{K-4}(K-1)^2}}$, the points from the ϵ -neighbourhood of all the α -edges are contained by set \mathcal{P} . This essentially means that the rows of $\overline{\mathbf{H}}$ satisfy the assumption (52) with ϵ accuracy and thus the ϵ -sufficiently scattered condition is achieved.

From Lemma 7, we get the lower bounds for α under various values of K :

$$\alpha > \alpha_{min}, \quad \alpha_{min} = \begin{cases} 1, & \text{for } K = 2, \\ \frac{2}{3}, & \text{for } K = 3, \\ \frac{1}{2}, & \text{for } K > 3. \end{cases} \tag{81}$$

Therefore, Eq. (80) can be again bounded as,

$$\Pr(E) \geq 1 - K(K-1)e^{-\frac{s\alpha_{min}^2(K-2)\epsilon^2}{2^{K-4}(K-1)^2}} \tag{82}$$

If s represents the number of rows in $\overline{\mathbf{H}}$, then for $s \geq \frac{2^{K-4}(K-1)^2}{\alpha_{min}^2 \epsilon^2} \log\left(\frac{K(K-1)}{\rho}\right)$, with probability at least $1 - \rho$, \mathbf{H} is ϵ -sufficiently scattered. Note that $s = (M-1)K$ where M is the number of annotators. This provides a bound on the number of annotators needed.

Consequently, if there exists at least $1 + \frac{2^{K-4}(K-1)^2}{K\alpha_{min}^2 \epsilon^2} \log\left(\frac{K(K-1)}{\rho}\right)$ annotators, then we have the conclusion of Theorem 4.

J Sample complexity for second order and third order statistics

In this section, we compare the sample complexity needed to estimate the second order statistics of the annotator responses from m and ℓ denoted as $\mathbf{R}_{m,\ell}$ and the third order statistics of the annotator responses from m , n and ℓ denoted as $\mathbf{R}_{m,n,\ell}$ given a dataset of N samples to jointly label as one of the K classes.

In crowdsourcing, not all samples are labeled by an annotator. To be specific, an annotator m labels each sample with probability $p_m \in (0, 1]$ and in most of the practical cases, $p_m \ll 1$. For simpler analysis, let us take $p_m = p$, for all annotators. Then, this results to have an average of $\lceil Np^2 \rceil$ joint responses from annotators m and ℓ and $\lceil Np^3 \rceil$ joint responses from annotators m , n and ℓ . With this and using the matrix and tensor concentration results from [40], the estimation error for $\mathbf{R}_{m,\ell}$ and $\mathbf{R}_{m,n,\ell}$ can be re-stated as, with probability at least $1 - \delta$,

$$\|\mathbf{R}_{m,\ell} - \widehat{\mathbf{R}}_{m,\ell}\|_F \leq \frac{1 + \sqrt{\log(1/\delta)}}{p\sqrt{N}} \tag{83}$$

$$\|\mathbf{R}_{m,n,\ell} - \widehat{\mathbf{R}}_{m,n,\ell}\|_F \leq \frac{1 + \sqrt{\log(K/\delta)}}{p^{\frac{3}{2}}\sqrt{N/K}} \tag{84}$$

It is clear from Eq. (83) and (84) that in order to achieve the same accuracy, third order statistics need much higher number of samples compared to second order statistics when p is smaller and K is larger.