

DEEP LEARNING FROM CROWDSOURCED LABELS: COUPLED CROSS-ENTROPY MINIMIZATION, IDENTIFIABILITY, AND REGULARIZATION

Shahana Ibrahim, Tri Nguyen, and Xiao Fu *

School of Electrical Engineering and Computer Science
Oregon State University
Corvallis, OR 97330, USA

ABSTRACT

Using noisy crowdsourced labels from multiple annotators, a deep learning-based end-to-end (E2E) system aims to learn the label correction mechanism and the neural classifier simultaneously. To this end, many E2E systems concatenate the neural classifier with multiple annotator-specific “label confusion” layers and co-train the two parts in a parameter-coupled manner. The formulated *coupled cross-entropy minimization* (CCEM)-type criteria are intuitive and work well in practice. Nonetheless, theoretical understanding of the CCEM criterion has been limited. The contribution of this work is twofold: First, performance guarantees of the CCEM criterion are presented. Our analysis reveals for the first time that the CCEM can indeed correctly identify the annotators’ confusion characteristics and the desired “ground-truth” neural classifier under realistic conditions, e.g., when only incomplete annotator labeling and finite samples are available. Second, based on the insights learned from our analysis, two regularized variants of the CCEM are proposed. The regularization terms provably enhance the identifiability of the target model parameters in various more challenging cases. A series of synthetic and real data experiments are presented to showcase the effectiveness of our approach.

1 INTRODUCTION

The success of deep learning has escalated the demand for labeled data to an unprecedented level. Some learning tasks can easily consume millions of labeled data (Najafabadi et al., 2015; Goodfellow et al., 2016). However, acquiring data labels is a nontrivial task—it often requires a pool of annotators with sufficient domain expertise to manually label the data items. For example, the popular Microsoft COCO dataset contains 2.5 million images and around 20,000 work hours aggregated from multiple annotators were used for its category labeling (Lin et al., 2014).

Crowdsourcing is considered an important working paradigm for data labeling. In crowdsourcing platforms, e.g., Amazon Mechanical Turk (Buhrmester et al., 2011), Crowdfunder (Wazny, 2017), and ClickWork (Vakharia & Lease, 2013), data items are dispatched and labeled by many annotators; the annotations are then integrated to produce reliable labels. A notable challenge is that annotator-output labels are sometimes considerably noisy. Training machine learning models using noisy labels could seriously degrade the system performance (Arpit et al., 2017; Zhang et al., 2016a). In addition, the labels provided by individual annotators are often largely incomplete, as a dataset is often divided and dispatched to different annotators.

Early crowdsourcing methods often treat annotation integration and downstream operations, e.g., classification, as separate tasks; see, (Dawid & Skene, 1979; Karger et al., 2011a; Whitehill et al., 2009; Snow et al., 2008; Welinder et al., 2010; Liu et al., 2012; Zhang et al., 2016b; Ibrahim et al., 2019; Ibrahim & Fu, 2021). This pipeline estimates the annotators’ confusion parameters (e.g., the confusion matrices under the Dawid & Skene (DS) model (Dawid & Skene, 1979)) in the first stage.

*Contact Information: Shahana Ibrahim, Tri Nguyen, and Xiao Fu: {ibrahish, nguyetr9, xiao.fu}@oregonstate.edu

Then, the corrected and integrated labels along with the data are used for training the downstream tasks’ classifiers. However, simultaneously learning the annotators’ confusions and the classifier in an end-to-end (E2E) manner has shown substantially improved performance in practice. (Raykar et al., 2010; Khetan et al., 2018; Tanno et al., 2019; Rodrigues & Pereira, 2018; Chu et al., 2021; Guan et al., 2018; Cao et al., 2019; Li et al., 2020; Wei et al., 2022; Chen et al., 2020).

To achieve the goal of E2E-based learning under crowdsourced labels, a class of methods concatenate a “confusion layer” for each annotator to the output of a neural classifier, and jointly learn these parts via a parameter-coupled manner. This gives rise to a *coupled cross entropy minimization* (CCEM) criterion (Rodrigues & Pereira, 2018; Tanno et al., 2019; Chen et al., 2020; Chu et al., 2021; Wei et al., 2022). In essence, the CCEM criterion models the observed labels as annotators’ confused outputs from the “*ground-truth*” predictor (GTP)—i.e., the classifier as if being trained using the noiseless labels and with perfect generalization—which is natural and intuitive. A notable advantage of the CCEM criteria is that they often lead to computationally convenient optimization problems, as the confusion characteristics are modeled as just additional structured layers added to the neural networks. Nonetheless, the seemingly simple CCEM criterion and its variants have shown promising performance. For example, the *crowdlayer* method is a typical CCEM approach, which has served as a widely used benchmark for E2E crowdsourcing since its proposal (Rodrigues & Pereira, 2018). In (Tanno et al., 2019), a similar CCEM-type criterion is employed, with a trace regularization added. More CCEM-like learning criteria are seen in (Chen et al., 2020; Chu et al., 2021; Wei et al., 2022), with some additional constraints and considerations.

Challenges. Apart from its empirical success, understanding of the CCEM criterion has been limited. Particularly, it is often unclear if CCEM can correctly identify the annotators’ confusion characteristics (that are often modeled as “confusion matrices”) and the GTP under reasonable settings (Rodrigues & Pereira, 2018; Chu et al., 2021; Chen et al., 2020; Wei et al., 2022), but identifiability stands as the key for ensured performance. The only existing identifiability result on CCEM was derived under restricted conditions, e.g., the availability of infinite samples and the assumption that there exists annotators with diagonally-dominant confusion matrices (Tanno et al., 2019). To our best knowledge, model identifiability of CCEM has not been established under realistic conditions, e.g., in the presence of incomplete annotator labeling and non-experts under finite samples.

We should note that a couple of non-CCEM approaches proposed identifiability-guaranteed solutions for E2E crowdsourcing. The work in (Khetan et al., 2018) showed identifiability under their *expectation maximization* (EM)-based learning approach, but the result is only applicable to binary classification. An information-theoretic loss-based learning approach proposed in (Cao et al., 2019) presents some identifiability results, but under the presence of a group of independent expert annotators and infinite samples. These conditions are hard to meet or verify. In addition, the computation of these methods are often more complex relative to the CCEM-based methods.

Contributions. Our contributions are as follows:

- **Identifiability Characterizations of CCEM-based E2E Crowdsourcing.** In this work, we show that the CCEM criterion can indeed provably identify the annotators’ confusion matrices and the GTP up to inconsequential ambiguities under mild conditions. Specifically, we show that, if the number of annotator-labeled items is sufficiently large and some other reasonable assumptions hold, the two parts can be recovered with bounded errors. Our result is the first finite-sample identifiability result for CCEM-based crowdsourcing. Moreover, our analysis reveals that the success of CCEM does *not* rely on conditional independence among the annotators. This is favorable (as annotator independence is a stringent requirement) and surprising, since conditional independence is often used to derive the CCEM criterion; see, e.g., (Tanno et al., 2019; Rodrigues & Pereira, 2018; Chu et al., 2021).

- **Regularization Design for CCEM With Provably Enhanced Identifiability.** Based on the key insights revealed in our analysis, we propose two types of regularizations that can provide enhanced identifiability guarantees under challenging scenarios. To be specific, the first regularization term ensures that the confusion matrices and the GTP can be identified without having any expert annotators if one has sufficiently large amount of data. The second regularization term ensures identifiability when class specialists are present among annotators. These identifiability-enhanced approaches demonstrate promising label integration performance in our experiments.

Notation. The notations are summarized in the supplementary material.

2 BACKGROUND

Crowdsourcing is a core working paradigm for labeling data, as individual annotators are often not reliable enough to produce quality labels. Consider a set of N items $\{\mathbf{x}_n\}_{n=1}^N$. Each item belongs to one of the K classes. Here, $\mathbf{x}_n \in \mathbb{R}^D$ denote the feature vector of n th data item. Let $\{y_n\}_{n=1}^N$ denote the set of *ground-truth* labels, where $y_n \in [K]$, $\forall n$. The ground-truth labels $\{y_n\}_{n=1}^N$ are *unknown*. We ask M annotators to provide their labels for their assigned items. Consequently, each item \mathbf{x}_n is labeled by a subset of annotators indexed by $\mathcal{S}_n \subseteq [M]$. Let $\{\hat{y}_n^{(m)}\}_{(m,n) \in \mathcal{S}}$ denote the set of labels provided by all annotators, where $\mathcal{S} = \{(m, n) \mid m \in \mathcal{S}_n, n \in [N]\}$ and $\hat{y}_n^{(m)} \in [K]$. Note that $|\mathcal{S}| \ll NM$ usually holds and that the annotators may often *incorrectly* label the items, leading to an incomplete and noisy $\{\hat{y}_n^{(m)}\}_{(m,n) \in \mathcal{S}}$. Here, “incomplete” means that \mathcal{S} does not cover all possible combinations of (m, n) —i.e., not all annotators label all items. The goal of an E2E crowdsourcing system is to train a reliable machine learning model (often a classifier) using such noisy and incomplete annotations. This is normally done via blindly estimating the confusion characteristics of the annotators from the noisy labels.

Classic Approaches. Besides the naive majority voting method, one of the most influential crowdsourcing model is by Dawid & Skene (Dawid & Skene, 1979). Dawid & Skene models the confusions of the annotators using the conditional probabilities $\Pr(\hat{y}_n^{(m)} | y_n)$. Consequently, annotator m ’s confusion is fully characterized by a matrix defined as follows: $\mathbf{A}_m(k, k') := \Pr(\hat{y}_n^{(m)} = k | y_n = k')$, $\forall k, k' \in [K]$. Dawid & Skene also proposed an EM algorithm that estimate \mathbf{A}_m ’s under a naive Bayes generative model. Many crowdsourcing methods follow Dawid & Skene’s modeling ideas (Ghosh et al., 2011; Dalvi et al., 2013; Karger et al., 2013; Zhang et al., 2016b; Traganitis et al., 2018; Ibrahim et al., 2019; Ibrahim & Fu, 2021). Like (Dawid & Skene, 1979), many of these methods do not exploit the data features \mathbf{x}_n while learning the confusion matrices of the annotators. They often treat confusion estimation and classifier training as two sequential tasks. Such two-stage approaches may suffer from error propagation.

E2E Crowdsourcing and The Identifiability Challenge. The E2E approaches in (Rodrigues & Pereira, 2018; Tanno et al., 2019; Khetan et al., 2018) model the probability of m th annotator’s response to the data item \mathbf{x}_n as follows:

$$\Pr(\hat{y}_n^{(m)} = k | \mathbf{x}_n) = \sum_{k'=1}^K \Pr(\hat{y}_n^{(m)} = k | y_n = k') \Pr(y_n = k' | \mathbf{x}_n), \quad k \in [K], \quad (1)$$

where the assumption that the annotator confusion is data-independent has been used to derive the right-hand side. In the above, the distribution $\Pr(\hat{y}_n^{(m)} | y_n)$ models the confusions in annotator responses given the true label y_n and $\Pr(y_n | \mathbf{x}_n)$ denotes the true label distribution of the data item \mathbf{x}_n . The distribution $\Pr(y_n | \mathbf{x}_n)$ can be represented using a mapping $\mathbf{f} : \mathbb{R}^D \rightarrow [0, 1]^K$, $[\mathbf{f}(\mathbf{x}_n)]_k = \Pr(y_n = k | \mathbf{x}_n)$. Define a probability vector $\mathbf{p}_n^{(m)} \in [0, 1]^K$ for every m and n , such that $[\mathbf{p}_n^{(m)}]_k = \Pr(\hat{y}_n^{(m)} = k | \mathbf{x}_n)$. With these notations, the following model holds:

$$\mathbf{p}_n^{(m)} = \mathbf{A}_m \mathbf{f}(\mathbf{x}_n), \quad \forall m, n. \quad (2)$$

In practice, one does not observe $\mathbf{p}_n^{(m)}$. Instead, the annotations $\hat{y}_n^{(m)}$ are categorical realizations of the distribution $\mathbf{p}_n^{(m)}$; i.e.,

$$\hat{y}_n^{(m)} \sim \text{Cart}(\mathbf{p}_n^{(m)}). \quad (3)$$

The goal of E2E crowdsourcing is to identify \mathbf{A}_m and \mathbf{f} in (2) from noisy labels $\{\hat{y}_n^{(m)}\}_{(m,n) \in \mathcal{S}}$ and data $\{\mathbf{x}_n\}_{n=1}^N$. Note that even if $\mathbf{p}_n^{(m)}$ in (2) is observed, \mathbf{A}_m and $\mathbf{f}(\cdot)$ are in general not *identifiable*. The reason is that one can easily find nonsingular matrices $\mathbf{Q} \in \mathbb{R}^{K \times K}$ such that $\mathbf{p}_n^{(m)} = \tilde{\mathbf{A}}_m \tilde{\mathbf{f}}(\mathbf{x}_n)$, where $\tilde{\mathbf{A}}_m = \mathbf{A}_m \mathbf{Q}$ and $\tilde{\mathbf{f}}(\mathbf{x}_n) = \mathbf{Q}^{-1} \mathbf{f}(\mathbf{x}_n)$. If an E2E approach does not ensure the identifiability of \mathbf{A}_m and \mathbf{f} , but outputs $\tilde{\mathbf{f}}(\mathbf{x}_n) = \mathbf{Q}^{-1} \mathbf{f}(\mathbf{x}_n)$, then the learned predictor is not likely to work reasonably. Identifying the ground-truth \mathbf{f} not only identifies the true labels for the training data (data for which the annotations have been collected), but is useful to predict the true labels for unseen test data items. In addition, correctly identifying \mathbf{A}_m ’s helps quantify how reliable each annotator is.

CCEM-based Approaches and Challenges. A number of existing E2E methods are inspired by the model in (2) to formulate CCEM-based learning losses.

- *Crowdlayer* (Rodrigues & Pereira, 2018): For each $(m, n) \in \mathcal{S}$, the generative model in (2) naturally leads to a cross-entropy based learning loss, i.e., $\text{CE}(\mathbf{A}_m \mathbf{f}(\mathbf{x}_n), \hat{\mathbf{y}}_n^{(m)})$.

As some annotators co-label items, the components $\mathbf{f}(\mathbf{x}_n)$ associated with the co-labeled items are shared by these annotators. With such latent component coupling and putting the cross-entropy losses for all $(m, n) \in \mathcal{S}$ together, the work in (Rodrigues & Pereira, 2018) formulate the following CCEM loss:

$$\underset{\mathbf{f} \in \mathcal{F}, \{\mathbf{A}_m \in \mathbb{R}^{K \times K}\}}{\text{minimize}} \frac{1}{|\mathcal{S}|} \sum_{(m,n) \in \mathcal{S}} \text{CE}(\text{softmax}(\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)), \hat{\mathbf{y}}_n^{(m)}), \quad (4)$$

where $\mathcal{F} \subseteq \{\mathbf{f}(\mathbf{x}) \in \mathbb{R}^K \mid \mathbf{f}(\mathbf{x}) \in \Delta_K, \forall \mathbf{x}\}$ is a function class and the softmax operator is applied as the output of $\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)$ is supposed to be a probability mass function (PMF). The method based on the formulation in (4) was called *crowdlayer* (Rodrigues & Pereira, 2018). The name comes from the fact that the confusion matrices of the “crowd” \mathbf{A}_m ’s can be understood as additional layers if \mathbf{f} is represented by a neural network. The loss function in (4) is relatively easy to optimize, as any off-the-shelf neural network training algorithms can be directly applied. This seemingly natural and simple CCEM approach demonstrated substantial performance improvement upon classic methods. However, there is no performance characterization for (4).

- *TraceReg* (Tanno et al., 2019): The work in (Tanno et al., 2019) proposed a similar loss function—with a regularization term in order to establish identifiability of the \mathbf{A}_m ’s and \mathbf{f} :

$$\underset{\mathbf{f} \in \mathcal{F}, \{\mathbf{A}_m \in \mathcal{A}\}}{\text{minimize}} \frac{1}{|\mathcal{S}|} \sum_{(m,n) \in \mathcal{S}} \text{CE}(\mathbf{A}_m \mathbf{f}(\mathbf{x}_n), \hat{\mathbf{y}}_n^{(m)}) + \lambda \sum_{m=1}^M \text{trace}(\mathbf{A}_m), \quad (5)$$

where \mathcal{A} is the constrained set $\{\mathbf{A} \in \mathbb{R}^{K \times K} \mid \mathbf{A} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{A} = \mathbf{1}^\top\}$. It was shown that (i) if $\mathbf{p}_n^{(m)}$ instead of $\hat{\mathbf{y}}_n^{(m)}$ is observed for every m and n and (ii) if the mean of the \mathbf{A}_m ’s are diagonally dominant, then solving (5) with $\text{CE}(\mathbf{A}_m \mathbf{f}(\mathbf{x}_n), \hat{\mathbf{y}}_n^{(m)})$ replaced by enforcing $\mathbf{p}_n^{(m)} = \mathbf{A}_m \mathbf{f}(\mathbf{x}_n)$ ensures identifying \mathbf{A}_m for all m , up to column permutations. Even though, the result provides some justification for using CCEM criterion under trace regularization, the result is clearly unsatisfactory, as it requires many stringent assumptions to establish identifiability.

- *SpeeLFC* (Chen et al., 2020), *Union-Net* (Wei et al., 2022), *CoNAL* (Chu et al., 2021): A number of other variants of the CCEM criterion are also proposed for E2E learning with crowdsourced labels. The work in (Chen et al., 2020) uses a similar criterion as in (4), but with the softmax operator applied on the columns of \mathbf{A}_m ’s, instead on the output of $\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)$. The work in (Wei et al., 2022) concatenates \mathbf{A}_m ’s vertically and applies the softmax operator on the columns of such concatenated matrix, while formulating the CCEM criterion. The method in (Chu et al., 2021) models a common confusion matrix in addition to the annotator-specific confusion matrices \mathbf{A}_m ’s and employs a CCEM-type criterion to learn both in a coupled fashion. Nonetheless, none of these approaches have successfully tackled the identifiability aspect. The work (Wei et al., 2022) claims theoretical support for their approach, but the proof is flawed. In fact, understanding to the success of CCEM-based methods under practical E2E learning settings is still unclear.

Other Related Works. A couple of non-CCEM approaches have also considered the identifiability aspects (Khetan et al., 2018; Cao et al., 2019). The work in (Khetan et al., 2018) presents an EM-inspired alternating optimization strategy. Such procedure involves training the neural classifier multiple times, making it a computationally intensive approach. In addition, its identifiability guarantees are under certain strict assumptions, e.g., binary classification, identical confusions and conditional independence for annotators—which are hard to satisfy in practice. The approach in (Cao et al., 2019) employs an information-theoretic loss function to jointly train a predictor network and annotation aggregation network. The identifiability guarantees are again under restrictive assumptions, e.g., infinite data, no missing annotations and existence of mutually independent expert annotators.

We should also note that there exist works considering data-dependent annotator confusions as well, e.g., (Zhang et al., 2020) (multiple annotator case) and (Cheng et al., 2021; Xia et al., 2020; Zhu et al., 2022) (single annotator case)—leveraging more complex models and learning criteria. In our work, we focus on the criterion designed using a model in (1) as it is shown to be effective in practice.

3 IDENTIFIABILITY OF CCEM-BASED E2E CROWDSOURCING

In this section, we offer identifiability analysis of the CCEM learning loss under realistic settings. We consider the following re-expressed objective function of CCEM:

$$\begin{aligned} \underset{\mathbf{f}, \{\mathbf{A}_m\}}{\text{minimize}} \quad & -\frac{1}{|\mathcal{S}|} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{1}[\hat{y}_n^{(m)} = k] \log[\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)]_k & (6a) \\ \text{subject to} \quad & \mathbf{f} \in \mathcal{F}, \mathbf{A}_m \in \mathcal{A}, \forall m, & (6b) \end{aligned}$$

where $\mathcal{F} \subseteq \{\mathbf{f}(\mathbf{x}) \in \mathbb{R}^K \mid \mathbf{f}(\mathbf{x}) \in \Delta_K, \forall \mathbf{x}\}$ is a function class and \mathcal{A} is the constrained set $\{\mathbf{A} \in \mathbb{R}^{K \times K} \mid \mathbf{A} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{A} = \mathbf{1}^\top\}$, since each column of \mathbf{A}_m 's are conditional probability distributions.

The rationale is that the confusion matrices \mathbf{A}_m 's act as correction terms for annotator's labeling noise to output a true label classifier \mathbf{f} . Intuitively, the objective in (6a) encourages the output estimates $\hat{\mathbf{f}}, \{\hat{\mathbf{A}}_m\}$ to satisfy the relation $\mathbf{p}_n^{(m)} = \hat{\mathbf{A}}_m \hat{\mathbf{f}}(\mathbf{x}_n)$, $\forall m, n$. One can easily note that the identifiability guarantees for \mathbf{A}_m and \mathbf{f} (i.e., when does the ground-truth \mathbf{A}_m and the ground-truth \mathbf{f} can be identified) are not so straightforward since there exists an infinite number of nonsingular matrices $\mathbf{Q} \in \mathbb{R}^{K \times K}$ such that $\mathbf{p}_n^{(m)} = (\hat{\mathbf{A}}_m \mathbf{Q})(\mathbf{Q}^{-1} \hat{\mathbf{f}}(\mathbf{x}_n))$.

To proceed, we make the following assumptions:

Assumption 1 Each data item \mathbf{x}_n , $n \in [N]$ is drawn from a distribution \mathcal{D} independently at random and has a bounded ℓ_2 norm. The observed index pairs (m, n) are included in \mathcal{S} uniformly at random.

Assumption 2 The neural network function class \mathcal{F} has a complexity measure denoted as $\mathcal{R}_{\mathcal{F}}$.

Here, we use the *sensitive complexity* parameter introduced in (Lin & Zhang, 2019) as the complexity measure $\mathcal{R}_{\mathcal{F}}$. In essence, $\mathcal{R}_{\mathcal{F}}$ gets larger as the neural network function class \mathcal{F} gets deeper and wider—also see supplementary material Sec. K for more discussion.

Assumption 3 There exists a GTP $\mathbf{f}^\natural : \mathbb{R}^D \rightarrow [0, 1]^K$, such that $[\mathbf{f}^\natural(\mathbf{x}_n)]_k = \Pr(y_n = k \mid \mathbf{x}_n)$, where y_n denotes the true label of the data item \mathbf{x}_n . In addition, there exists $\tilde{\mathbf{f}} \in \mathcal{F}$ such that $\|\tilde{\mathbf{f}}(\mathbf{x}_n) - \mathbf{f}^\natural(\mathbf{x}_n)\|_2 \leq \nu$ for all $\mathbf{x}_n \sim \mathcal{D}$, where $0 \leq \nu < \infty$.

Assumption 4 (Near-Class Specialist Assumption) Let \mathbf{A}_m^\natural denote the ground-truth confusion matrix for annotator m . For each class k , there exists a near-class specialist, indexed by m_k , such that $\|\mathbf{A}_{m_k}^\natural(k, \cdot) - \mathbf{e}_k^\top\|_2 \leq \xi_1$, where $0 \leq \xi_1 < \infty$.

Assumption 5 (Near-Anchor Point Assumption) For each class k , there exists a near-anchor point, \mathbf{x}_{n_k} , such that $\|\mathbf{f}^\natural(\mathbf{x}_{n_k}) - \mathbf{e}_k\|_2 \leq \xi_2$, where $0 \leq \xi_2 < \infty$.

Assumptions 1-3 are standard conditions for analyzing performance of neural models under finite sample settings. The Near-Class Specialist Assumption (NCSA) implies that there exist annotators for each class $k \in [K]$ who can correctly distinguish items belonging to class k from those belonging to other classes. It is important to note that NCSA assumption is much milder compared to the commonly used assumptions such as the existence of all-class expert annotators (Cao et al., 2019) or diagonally dominant \mathbf{A}_m^\natural 's (Tanno et al., 2019). Instead, the NCSA only requires that for each class, there exists one annotator who is specialized for it (i.e., class-specialist)—but the annotator needs not be an all-class expert; see Fig. 1 in the supplementary material for illustration. The Near-Anchor Point Assumption (NAPA) is the relaxed version of the exact anchor point assumption commonly used in provable label noise learning methods (Xia et al., 2019; Li et al., 2021). Under Assumptions 1-5, we have the following result:

Theorem 1 Assume that each $[\mathbf{A}_m^\natural \mathbf{f}^\natural(\mathbf{x}_n)]_k$ and $[\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)]_k$, $\forall \mathbf{A}_m \in \mathcal{A}, \forall \mathbf{f} \in \mathcal{F}$ are at least $(1/\beta)$. Also assume that $\sigma_{\max}(\mathbf{A}_m^\natural) \leq \sigma$, $\forall m$, for a certain $\sigma > 0$. Then, for any $\alpha > 0$, with

probability greater than $1 - K/N^\alpha$, any optimal solution $\widehat{\mathbf{A}}_m$ and $\widehat{\mathbf{f}}$ of the problem (6) satisfies the following relations:

$$\min_{\mathbf{\Pi}} \|\widehat{\mathbf{A}}_m - \mathbf{A}_m^{\natural} \mathbf{\Pi}\|_{\mathbb{F}}^2 = K\sigma^2(\eta + \xi_1 + \xi_2), \quad \forall m \in [M], \quad (7a)$$

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min_{\mathbf{\Pi}} \|\widehat{\mathbf{f}}(\mathbf{x}) - \mathbf{\Pi}^\top \mathbf{f}^{\natural}(\mathbf{x})\|_2^2 \right] = K(\eta + \xi_1 + \xi_2), \quad (7b)$$

where $\eta^2 = \mathcal{O}\left(\beta MN^\alpha / \sqrt{S} \left(\sqrt{M \log S} + (\|\mathbf{X}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{F}})^{\frac{1}{4}}\right) + \beta \sqrt{K} MN^\alpha \nu\right)$, $\mathbf{\Pi} \in \{0, 1\}^K$ is a permutation matrix, and $\mathbf{X} = [\mathbf{x}_{n_1}, \dots, \mathbf{x}_{n_S}]$, $(m_s, n_s) \in \mathcal{S}$, if conditions $\xi_1, \xi_2 \leq 1/K$, $\nu \leq 1/\beta K^2 M^2 N^\alpha$, and $S = |\mathcal{S}| = \Omega\left(\beta^2 M^2 N^{2\alpha} K^2 \max\left(M \log S, \sqrt{\|\mathbf{X}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{F}}}\right)\right)$ hold.

The proof of Theorem 1 is relegated to the supplementary material in Sec. B. The takeaways are as follows: First, the CCEM can provably identify the \mathbf{A}_m^{\natural} 's and the GTP \mathbf{f}^{\natural} under finite samples—and such finite-sample identifiability of CCEM was not established before. Second, some restrictive conditions (e.g., the existence of all-class experts (Tanno et al., 2019; Cao et al., 2019)) used in the literature for establishing CCEM's identifiability are actually not needed. Third, many works derived the CCEM criterion under the maximum likelihood principle via assuming that the annotators are conditionally independent; see (Rodrigues & Pereira, 2018; Chen et al., 2020; Tanno et al., 2019; Chu et al., 2021). Interestingly, as we revealed in our analysis, the identifiability of \mathbf{A}_m^{\natural} and \mathbf{f}^{\natural} under the CCEM criterion does *not* rely on the annotators' independence. These new findings support and explain the effectiveness of CCEM in a wide range of scenarios observed in the literature.

4 ENHANCING IDENTIFIABILITY VIA REGULARIZATION

Theorem 1 confirms that CCEM is a provably effective criterion for identifying \mathbf{A}_m^{\natural} 's and \mathbf{f}^{\natural} . The caveats lie in Assumptions 4 and 5. Essentially, these assumptions require that some rows of the collection $\{\mathbf{A}_m^{\natural}\}_{m=1}^M$ and some vectors of the collection $\{\mathbf{f}^{\natural}(\mathbf{x}_n)\}_{n=1}^N$ are close to the canonical vectors. These are reasonable assumptions, but satisfying them simultaneously may not always be possible. A natural question is if these assumptions can be relaxed, at least partially.

In this section, we propose variants of CCEM that admits enhanced identifiability. To this end, we use a condition that is often adopted in the nonnegative matrix factorization (NMF) literature:

Definition 1 (SSC)(Fu et al., 2019; Gillis, 2020) Consider the second-order cone $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^K \mid \sqrt{K-1} \|\mathbf{x}\|_2 \leq \mathbf{1}^\top \mathbf{x}\}$. A nonnegative matrix $\mathbf{Z} \in \mathbb{R}_+^{L \times K}$ satisfies the sufficiently scattered condition (SSC) if (i) $\mathcal{C} \subseteq \text{cone}(\mathbf{Z}^\top)$ and (ii) $\text{cone}\{\mathbf{Z}^\top\} \subseteq \text{cone}\{\mathbf{Q}\}$ does not hold for any orthogonal $\mathbf{Q} \in \mathbb{R}^{K \times K}$ except for the permutation matrices.

Geometrically, the matrix \mathbf{Z} satisfying the SSC means that its rows span a large “area” within the nonnegative orthant, so that the conic hull spanned by the rows contains the second order cone \mathcal{C} as a subset. Note that the SSC condition subsumes the the NCSA (Assumption (4) where $\xi_1 = 0$) and the NAPA (Assumption (5) where $\xi_1 = 0$) as special cases.

Using SSC, we first show that the CCEM criterion in (6) attains identifiability of \mathbf{A}_m^{\natural} 's and \mathbf{f}^{\natural} without using the NCSA or the NAPA, when the problem size grows sufficiently large.

Theorem 2 Suppose that the incomplete labeling paradigm in Assumption 1 holds with $S = \Omega(t)$ and $N = O(t^3)$, for a certain $t > 0$. Also, assume that there exists $\mathcal{Z} = \{n_1, \dots, n_T\}$ such that $\mathbf{F}_{\mathcal{Z}}^{\natural} = [\mathbf{f}^{\natural}(\mathbf{x}_{n_1}), \dots, \mathbf{f}^{\natural}(\mathbf{x}_{n_T})]$ satisfies the SSC and $\mathbf{W}^{\natural} = [\mathbf{A}_1^{\natural\top}, \dots, \mathbf{A}_M^{\natural\top}]$ satisfies the SSC as well. Then, at the limit of $t \rightarrow \infty$, if $\mathbf{f}^{\natural} \in \mathcal{F}$, the optimal solutions $(\{\widehat{\mathbf{A}}_m\}, \widehat{\mathbf{f}})$ of (6) satisfies $\widehat{\mathbf{A}}_m = \mathbf{A}_m^{\natural} \mathbf{\Pi}$ for all m and $\widehat{\mathbf{f}}(\mathbf{x}) = \mathbf{\Pi}^\top \mathbf{f}^{\natural}(\mathbf{x})$ for all $\mathbf{x} \sim \mathcal{D}$, where $\mathbf{\Pi}$ is a permutation matrix.

The proof is via connecting the CCEM problem to the classic nonnegative matrix factorization (NMF) problem at the limit of $t \rightarrow \infty$; see Sec. D. It is also worth to note that Theorem 2 does not require complete observation of the annotations to guarantee identifiability. Another remark is that \mathbf{W}^{\natural} and \mathbf{F}^{\natural} satisfying the SSC is more relaxed compared to that they satisfying the NCSA and NAPA,

respectively. For example, one may need class specialists with much higher expertise to satisfy NCSA than to satisfy SSC on \mathbf{W}^{\natural} —also see geometric illustration in Sec. L. Compared to the result in Theorem 1, Theorem 2 implies that when N is very large, the CCEM should work under less stringent conditions relative to the NCSA and NAPA. Nonetheless, the assumption that both the GTP and the annotators’ confusion matrices should meet certain requirements may not always hold. In the following, we further show that the SSC on either \mathbf{W}^{\natural} or \mathbf{F}^{\natural} can be relaxed:

Theorem 3 *Suppose that the incomplete labeling paradigm in Assumption 1 holds with $S = \Omega(t)$ and $N = O(t^3)$, for a certain $t > 0$. Then, at the limit of $t \rightarrow \infty$, if $\mathbf{f}^{\natural} \in \mathcal{F}$, we have $\mathbf{A}_m^* = \mathbf{A}_m^{\natural} \mathbf{\Pi}$ for all m and $\mathbf{f}^*(\mathbf{x}) = \mathbf{\Pi}^{\top} \mathbf{f}^{\natural}(\mathbf{x})$ for all $\mathbf{x} \sim \mathcal{D}$, when any of the following conditions hold:*

(a) *If there exists $\mathcal{Z} = \{n_1, \dots, n_T\}$ such that $(\mathbf{F}_{\mathcal{Z}}^{\natural})^{\top} = [\mathbf{f}^{\natural}(\mathbf{x}_{n_1}), \dots, \mathbf{f}^{\natural}(\mathbf{x}_{n_T})]^{\top}$ satisfies the SSC, $\text{rank}(\mathbf{W}^{\natural}) = K$, and $(\{\mathbf{A}_m^*\}, \mathbf{f}^*)$ is the optimal solution of (6) with the maximal $\log \det(\mathbf{F}^* \mathbf{F}^{*\top})$.*

(b) *If $\mathbf{W}^{\natural} = [\mathbf{A}_1^{\natural\top}, \dots, \mathbf{A}_M^{\natural\top}]^{\top}$ satisfies the SSC, $\text{rank}(\mathbf{F}^{\natural}) = K$, and $(\mathbf{A}_m^*, \mathbf{f}^*)$ is the optimal solution of (6) with the maximal $\log \det((\mathbf{W}^*)^{\top} \mathbf{W}^*)$.*

Theorem 3(a) shows that if the number of annotated data items is large and the GTP-outputs associated with \mathbf{x}_n are diverse enough to satisfy the SSC, then \mathbf{f}^{\natural} and \mathbf{A}_m^{\natural} are identifiable even if there are no class specialists available. Compared to Theorem 2, the identifiability is clearly enhanced, as the conditions on the annotators have been relaxed. Theorem 3(b) implies that if \mathbf{W}^{\natural} satisfies SSC, identifiability can be established irrespective of the geometry of \mathbf{F}^{\natural} . Intuitively, when there are more annotators available, \mathbf{W}^{\natural} is more likely to satisfy SSC (the relation between the size of \mathbf{W}^{\natural} and SSC was shown in (Ibrahim et al., 2019, Theorem 4) under reasonable conditions). Hence, in practice, Theorem 3(b) may offer ensured performance in cases where more annotators are available.

Similar to the proof in Theorem 2, the proof here connects the CCEM-based E2E crowdsourcing problem to the simplex-structured matrix factorization (SSMF) problem (see (Fu et al., 2015; 2018)) at the limit of $t \rightarrow \infty$. The detailed proof is provided in Sec. E. We should mention that various connections between matrix factorization models and data labeling problems were also observed in related prior works, e.g, the classic approaches based on Dawid & Skene model (Dawid & Skene, 1979; Zhang et al., 2016b; Ibrahim et al., 2019; Ibrahim & Fu, 2021) and E2E learning under noisy labels (Li et al., 2021). Nonetheless, the former does not consider cross-entropy losses or involve any feature extractor, whereas, the latter does not consider multiple annotators or incomplete labeling.

Implementation via Regularization Theorem 3 naturally leads to regularized versions of CCEM. Correspondingly, we have the following two cases:

Following Theorem 3(a), when the GTP is believed to be diverse over different \mathbf{x}_n ’s, we propose to employ the following regularized CCEM:

$$\underset{\mathbf{A}_m \in \mathcal{A}, \forall m, \mathbf{f} \in \mathcal{F}}{\text{minimize}} \quad -\frac{1}{|\mathcal{S}|} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\hat{y}_n^{(m)} = k] \log[\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)]_k - \lambda \log \det \mathbf{F} \mathbf{F}^{\top}, \quad (8)$$

where $\mathbf{F} = [\mathbf{f}(\mathbf{x}_1), \dots, \mathbf{f}(\mathbf{x}_N)]^{\top}$. In similar spirit, following Theorem 3(b), we propose the following regularized CCEM when M is large (so that \mathbf{W}^{\natural} likely satisfies SSC):

$$\underset{\mathbf{A}_m \in \mathcal{A}, \forall m, \mathbf{f} \in \mathcal{F}}{\text{minimize}} \quad -\frac{1}{|\mathcal{S}|} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\hat{y}_n^{(m)} = k] \log[\mathbf{A}_m \mathbf{f}(\mathbf{x}_n)]_k - \lambda \log \det \mathbf{W}^{\top} \mathbf{W}, \quad (9)$$

where $\mathbf{W} = [\mathbf{A}_1^{\top}, \dots, \mathbf{A}_M^{\top}]^{\top}$. The implementations for (8) and (9) are called as geometry-regularized crowdsourcing network, abbreviated as GeoCrowdNet (F) and GeoCrowdNet (W), respectively.

Remark 1 Based on our analyses, our suggested *rule of thumb* for choosing regularization is as follows: When M is large but N is relatively small, using GeoCrowdNet (W) is recommended, as \mathbf{W} is more likely to satisfy the SSC compared to \mathbf{F}^{\top} . However, when N is large, using GeoCrowdNet (F) is expected to have a better performance as \mathbf{F}^{\top} is likely to satisfy the SSC in this case. Under big data settings, N is often large and the GTP outputs of \mathbf{x}_n ’s are often reasonably diverse. Hence, (8) oftentimes works well, as will be seen in our experiments. Nonetheless, for cases where the geometry of \mathbf{F}^{\top} violates the SSC, the employment of more annotators can help via using (9)—which is also an intuitive advantage of crowdsourcing.

Table 1: Average test accuracy (\pm std) of the proposed methods and the baselines on MNIST & Fashion-MNIST dataset under various (N, M) 's; labels are produced by machine annotators; $p = 0.1$.

Methods	MNIST				Fashion-MNIST			
	Case 1		Case 2		Case 1		Case 2	
	(1000, 15)	(1000, 20)	(5000, 5)	(10000, 5)	(1000, 15)	(1000, 30)	(5000, 5)	(10000, 5)
GeoCrowdNet (F)	79.89 \pm 3.08	82.18 \pm 3.48	85.92 \pm 2.73	87.21 \pm 2.47	78.98 \pm 2.83	84.47 \pm 1.64	80.60 \pm 0.46	83.68 \pm 2.17
GeoCrowdNet (W)	80.97 \pm 1.31	83.69 \pm 2.37	77.79 \pm 8.97	82.37 \pm 9.18	79.80 \pm 4.23	85.56 \pm 1.91	72.36 \pm 3.84	74.03 \pm 7.41
GeoCrowdNet ($\lambda = 0$)	71.15 \pm 6.73	69.17 \pm 2.61	71.66 \pm 4.48	60.29 \pm 7.91	70.92 \pm 4.14	81.88 \pm 4.41	69.31 \pm 4.77	73.04 \pm 7.56
TraceReg	70.14 \pm 6.93	78.09 \pm 3.52	63.71 \pm 10.76	64.45 \pm 10.14	75.06 \pm 3.43	83.79 \pm 2.15	69.82 \pm 6.46	72.21 \pm 8.44
CrowdLayer	65.72 \pm 4.81	72.90 \pm 2.31	51.25 \pm 6.24	53.12 \pm 14.02	63.91 \pm 2.11	76.50 \pm 3.68	64.73 \pm 6.89	63.18 \pm 3.55
MBEM	37.24 \pm 8.54	33.39 \pm 5.72	28.14 \pm 6.24	26.90 \pm 2.01	23.14 \pm 2.36	25.43 \pm 5.68	37.34 \pm 3.73	38.62 \pm 5.45
CoNAL	51.33 \pm 4.27	55.59 \pm 5.32	44.79 \pm 7.01	39.41 \pm 9.21	59.76 \pm 3.34	71.44 \pm 6.22	52.15 \pm 6.08	54.61 \pm 7.54
Max-MIG	81.32 \pm 1.31	83.34 \pm 3.36	83.71 \pm 1.23	81.05 \pm 0.65	70.95 \pm 5.61	81.27 \pm 5.37	73.45 \pm 5.21	73.62 \pm 3.12
NN-MV	78.98 \pm 3.21	82.78 \pm 1.45	83.44 \pm 1.98	83.10 \pm 2.86	62.61 \pm 6.38	73.61 \pm 3.88	61.95 \pm 3.73	71.64 \pm 2.25
NN-DSEM	73.95 \pm 4.90	77.57 \pm 4.06	65.81 \pm 2.74	69.09 \pm 3.21	55.89 \pm 5.97	73.64 \pm 4.44	31.40 \pm 3.86	35.10 \pm 4.89

5 EXPERIMENTS

Baselines. The proposed methods are compared with a number of existing E2E crowdsourcing methods, namely, TraceReg (Tanno et al., 2019), MBEM (Khetan et al., 2018), CrowdLayer (Rodrigues & Pereira, 2018), CoNAL (Chu et al., 2021), and Max-MIG (Cao et al., 2019). In addition to these baselines, we also learn neural network classifiers using the labels aggregated from *majority voting* and the classic EM algorithm proposed by Dawid & Skene (Dawid & Skene, 1979), denoted as NN-MV and NN-DSEM, respectively.

Real-Data Experiments with Machine Annotations. The settings are as follows:

Dataset. We use the MNIST dataset (Deng, 2012) and Fashion-MNIST dataset (Xiao et al., 2017); see more details in Sec. N.

Noisy Label Generation. We train a number of machine classifiers to produce noisy labels. Five types of classifiers, including support vector machines (SVM), k -Nearest Neighbour (kNN), logistic regression, convolutional neural network (CNN), and fully connected neural network are considered to act as annotators. To create more annotators, we change some of their parameters (e.g., the number of nearest neighbour parameter k of kNN and the number of epochs for training the CNN). We consider two different cases: (i) *Case 1: NCSA (Assumption 4) holding*: Each annotator is chosen to be an all-class expert with probability 0.1. If an annotator is chosen to be an all-class expert, it is trained carefully so that its classification accuracy exceeds a certain threshold; see details in Sec. N. Note that the existence of an expert implies that the NCSA and the SSC are likely satisfied by \mathbf{W}^{\natural} , but the reverse is not necessarily true. We use this strategy to enforce NCSA only for validation purpose. It does not mean that the CCEM criterion needs the existence of experts to work. (ii) *Case 2: NCSA not holding*: Each machine annotator is trained by randomly choosing a small subset of the training data (with 100 to 500 samples) so that the accuracy of these machine annotators are low. This way, any all-class expert or class-specialist is unlikely to exist.

We use the two cases to validate our theorems. Under our analyses, GeoCrowdNet (W) is expected to work better under Case 1, as NCSA approximately implies SSC. In addition, GeoCrowdNet (F) does not rely on the geometry of \mathbf{W}^{\natural} , and thus can work under both cases, if N is reasonably large (which implies that $(\mathbf{F}^{\natural})^{\top}$ is more likely to satisfy the SSC). Once the machine annotators are trained, we let them label unseen data items of size N . To evaluate the methods under incomplete labeling, an annotator labels any data item with probability $p = 0.1$. Under this labeling strategy, every data item is only labeled by a small subset of the annotators.

Settings. The neural network classifier architecture for MNIST dataset is chosen to be Lenet-5 (Lecun et al., 1998), which consists of two sets of convolutional and max pooling layers, followed by a flattening convolutional layer, two fully-connected layers and finally a softmax layer. For Fashion-MNIST dataset, we use the ResNet-18 architecture (He et al., 2016). Adam (Kingma & Ba, 2015) is used as an optimizer with weight decay of 10^{-4} and batch size of 128. The regularization parameter λ and the initial learning rate of the Adam optimizer are chosen via grid search method using the validation set from $\{0.01, 0.001, 0.0001\}$ and $\{0.01, 0.001\}$, respectively. We choose the same neural network structures for all the baselines. The confusion matrices are initialized with identity matrices of size K for proposed methods and the baselines TraceReg and CrowdLayer.

Results. Table 1 presents the average label prediction accuracy on the testing data of the MNIST and the Fashion-MNIST over 5 random trials, for various cases. One can observe that GeoCrowdNet (F)

Table 2: Average test accuracy of the proposed methods and the baselines on LabelMe ($M = 59, K = 8$) and Music ($M = 44, K = 10$) datasets.

Methods	LabelMe	Music
GeoCrowdNet (F)	85.85 ± 0.33	67.13 ± 1.27
GeoCrowdNet (W)	81.59 ± 0.73	66.46 ± 1.77
GeoCrowdNet ($\lambda=0$)	80.35 ± 0.40	65.93 ± 1.81
TraceReg	80.89 ± 0.73	66.40 ± 0.90
Crowdlayer	81.41 ± 0.90	65.53 ± 1.74
MBEM	80.05 ± 3.31	69.00 ± 2.15
CoNAL	81.68 ± 2.72	60.93 ± 0.61
Max-MIG	85.03 ± 0.44	64.33 ± 2.31
NN-MV	79.57 ± 1.24	63.40 ± 2.68
NN-DSEM	78.26 ± 1.18	64.66 ± 2.34

performs the best when there are more number of annotated items, even when there are no class specialist annotators present (i.e., case 2). On the other hand, GeoCrowdNet (W) starts showing its advantage over GeoCrowdNet (F) when there are more number of annotators and class specialists among them. Both are consistent with our theorems. In addition, the baseline MaxMIG performs competitively when there are all-class experts—which is consistent with their identifiability analysis. However, when there are no such experts (case 2), its performance drops compared to the proposed methods, especially GeoCrowdNet (F) whose identifiability does not rely on the existence of all-class experts or class specialists. Overall, GeoCrowdNet (F) exhibits consistently good performance.

Real-Data Experiments with Human Annotators. The settings are as follows:

Datasets. We consider two different datasets, namely, LabelMe ($M = 59$) (Rodrigues et al., 2017; Russell et al., 2007) and Music ($M = 44$) (Rodrigues et al., 2014), both having noisy and incomplete labels provided by human workers from AMT.

Settings. For LabelMe, we employ the pretrained VGG-16 embeddings followed by a fully connected layer as used in (Rodrigues & Pereira, 2018). For Music, we choose a similar architecture, but with batch normalization layers. We use a batch size of 128 for LabelMe and a batch size of 100 for Music dataset. All other settings are the same as the machine annotator case. More details of the datasets and architecture can be seen in Sec. N.

Results. Table 2 shows the average label prediction accuracies on the test data for 5 random trials. One can see that GeoCrowdNet ($\lambda=0$) works reasonably well relative to the machine annotation case in the previous experiment. This may be because the number of annotators for both LabelMe and Music are reasonably large ($M = 59$ and $M = 44$, respectively), which makes \mathbf{W}^\dagger satisfying Assumption 4 become easier than before. This validates our claims in Theorem 1, but also shows the limitations of the plain-vanilla CCEM—that is, it may require a fairly large number of annotators to start showing promising results. Similar to the previous experiment, both GeoCrowdNet (F) and GeoCrowdNet (W) outperform the unregularized version GeoCrowdNet ($\lambda=0$), showing the advantages of the identifiability-enhancing regularization terms. More experiments and details are presented in Sec. N, due to page limitations.

6 CONCLUSION

In this work, we revisited the CCEM criterion—one of the most popular E2E learning criteria for crowdsourced label integration. We provided the first finite-sample identifiability characterization of the confusion matrices and the neural classifier which are learned using the CCEM criterion. Compared to many exiting identifiability results, our guarantees are under more relaxed and more realistic settings. In particular, a take-home point revealed in our analysis is that CCEM can provably identify the desired model parameters even if the annotators are dependent, which is a surprising but favorable result. We also proposed two regularized variants of the CCEM, based on the insights learned from our identifiability analysis for the plain-vanilla CCEM. The regularized CCEM criteria provably enhance the identifiability of the confusion matrices and the neural classifier under more challenging scenarios. We evaluated the proposed approaches on various synthetic and real datasets. The results corroborate our theoretical claims.

Acknowledgement. This work is supported in part by the National Science Foundation under Project NSF IIS-2007836.

REFERENCES

- Devansh Arpit, Stanisław Jastrzundefinedbski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *Proceedings of International Conference on Machine Learning*, pp. 233–242, 2017.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Michael Buhrmester, Tracy Kwang, and Samuel Gosling. Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6:3–5, 02 2011.
- Peng Cao, Yilun Xu, Yuqing Kong, and Yizhou Wang. Max-MIG: An information theoretic approach for joint learning from crowds. In *Proceedings of International Conference on Learning Representations*, 2019.
- Zhijun Chen, Huimin Wang, Hailong Sun, Pengpeng Chen, Tao Han, Xudong Liu, and Jie Yang. Structured probabilistic end-to-end learning from crowds. In *Proceedings of International Joint Conference on Artificial Intelligence*, pp. 1512–1518, 7 2020.
- Hao Cheng, Zhaowei Zhu, Xingyu Li, Yifei Gong, Xing Sun, and Yang Liu. Learning with instance-dependent label noise: A sample sieve approach. In *Proceedings of International Conference on Learning Representations*, 2021.
- Zhendong Chu, Jing Ma, and Hongning Wang. Learning from crowds by modeling common confusions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7):5832–5840, May 2021.
- Nilesh Dalvi, Anirban Dasgupta, Ravi Kumar, and Vibhor Rastogi. Aggregating crowdsourced binary ratings. In *Proceedings of International Conference on World Wide Web*, pp. 285–294, 2013.
- Alexander Philip Dawid and Allan M Skene. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied statistics*, pp. 20–28, 1979.
- Li Deng. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Process. Mag.*, 29(6):141–142, 2012.
- A.A. Fedotov, P. Harremoës, and F. Topsøe. Refinements of pinsker’s inequality. *IEEE Trans. Inf. Theory*, 49(6):1491–1498, 2003.
- Xiao Fu, Wing-Kin Ma, Kejun Huang, and Nicholas D. Sidiropoulos. Blind separation of quasi-stationary sources: Exploiting convex geometry in covariance domain. *IEEE Trans. Signal Process.*, 63(9):2306–2320, 2015.
- Xiao Fu, Kejun Huang, and Nicholas D Sidiropoulos. On identifiability of nonnegative matrix factorization. *IEEE Signal Process. Lett.*, 25(3):328–332, 2018.
- Xiao Fu, Kejun Huang, Nicholas D Sidiropoulos, and Wing-Kin Ma. Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications. *IEEE Signal Process. Mag.*, 36(2):59–80, 2019.
- Arpita Ghosh, Satyen Kale, and Preston McAfee. Who moderates the moderators?: Crowdsourcing abuse detection in user-generated content. In *Proceedings of the ACM conference on Electronic commerce*, pp. 167–176, 2011.
- Nicolas Gillis. *Nonnegative Matrix Factorization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2020.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

- Melody Guan, Varun Gulshan, Andrew Dai, and Geoffrey Hinton. Who said what: Modeling individual labelers improves classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), 2018.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- K. Huang, N. Sidiropoulos, and A. Swami. Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition. *IEEE Trans. Signal Process.*, 62(1):211–224, 2014.
- Kejun Huang, Nicholas D. Sidiropoulos, Evangelos E. Papalexakis, Christos Faloutsos, Partha Pratim Talukdar, and Tom M. Mitchell. Principled neuro-functional connectivity discovery. In *Proceedings of the SIAM International Conference on Data Mining*, 2015.
- Kejun Huang, Xiao Fu, and Nikolaos D Sidiropoulos. Anchor-free correlated topic modeling: Identifiability and algorithm. In *Advances in Neural Information Processing Systems*, pp. 1786–1794, 2016.
- Shahana Ibrahim and Xiao Fu. Crowdsourcing via annotator co-occurrence imputation and provable symmetric nonnegative matrix factorization. In *Proceedings of International Conference on Machine Learning*, volume 139, pp. 4544–4554, 2021.
- Shahana Ibrahim, Xiao Fu, Nikos Kargas, and Kejun Huang. Crowdsourcing via pairwise co-occurrences: Identifiability and algorithms. In *Advances in Neural Information Processing Systems*, volume 32, pp. 7847–7857, 2019.
- David Karger, Sewoong Oh, and Devavrat Shah. Iterative learning for reliable crowdsourcing systems. In *Advances in Neural Information Processing Systems*, volume 24, pp. 1953–1961, 2011a.
- David R Karger, Sewoong Oh, and Devavrat Shah. Efficient crowdsourcing for multi-class labeling. *ACM Sigmetrics Performance Evaluation Review*, 41(1):81–92, 2013.
- Ashish Khetan, Zachary Chase Lipton, and Anima Anandkumar. Learning from noisy singly-labeled data. In *Proceedings of International Conference on Learning Representations*, 2018.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representations*, 2015.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technical Report, University of Toronto*, 2009.
- Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Shikun Li, Shiming Ge, Yingying Hua, Chunhui Zhang, Hao Wen, Tengfei Liu, and Weiqiang Wang. Coupled-view deep classifier learning from multiple noisy annotators. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:4667–4674, 2020.
- Xuefeng Li, Tongliang Liu, Bo Han, Gang Niu, and Masashi Sugiyama. Provably end-to-end label-noise learning without anchor points. In *Proceedings of International Conference on Machine Learning*, pp. 6403–6413, 2021.
- Shan Lin and Jingwei Zhang. Generalization bounds for convolutional neural networks. *arXiv preprint*, arXiv:1910.01487, 2019.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, 2014.
- Qiang Liu, Jian Peng, and Alexander T Ihler. Variational inference for crowdsourcing. In *Advances in Neural Information Processing Systems*, volume 25, pp. 692–700, 2012.
- Shuying Liu and Weihong Deng. Very deep convolutional neural network based image classification using small training sample size. In *Asian Conference on Pattern Recognition (ACPR)*, pp. 730–734, 2015.

- Andreas Maurer. A vector-contraction inequality for rademacher complexities. In *International Conference on Algorithmic Learning Theory*, pp. 7–13, 2016.
- Maryam M Najafabadi, Flavio Villanustre, Taghi M Khoshgoftaar, Naeem Seliya, Randall Wald, and Edin Muharemagic. Deep learning applications and challenges in big data analytics. *Journal of Big Data*, 2(1):1–21, 2015.
- Vikas C. Raykar, Shipeng Yu, Linda H. Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. *Journal of Machine Learning Research*, 11: 1297–1322, 2010.
- Filipe Rodrigues and Francisco Pereira. Deep learning from crowds. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), 2018.
- Filipe Rodrigues, Francisco Pereira, and Bernardete Ribeiro. Gaussian process classification and active learning with multiple annotators. In *Proceedings of International Conference on Machine Learning*, volume 32, pp. 433–441, 2014.
- Filipe Rodrigues, Mariana Lourenco, Bernardete Ribeiro, and Francisco C. Pereira. Learning supervised topic models for classification and regression from crowds. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(12):2409–2422, 2017.
- Bryan C. Russell, Antonio Torralba, Kevin P. Murphy, and William T. Freeman. Labelme: A database and web-based tool for image annotation. *International Journal of Computer Vision*, 77:157–173, 2007.
- Pinsker M S. The information stability of gaussian random variables and processes (in russian). *Dokl. Akad. Nauk SSSR*, 133:28–30, 1960.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew Y Ng. 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, 2008.
- Ryutaro Tanno, Ardavan Saeedi, Swami Sankaranarayanan, Daniel C. Alexander, and Nathan Silberman. Learning from noisy labels by regularized estimation of annotator confusion. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11236–11245, 2019.
- Panagiotis A Traganitis, Alba Pages-Zamora, and Georgios B Giannakis. Blind multiclass ensemble classification. *IEEE Trans. Signal Process.*, 66(18):4737–4752, 2018.
- Donna Vakharia and Matthew Lease. Beyond AMT: An analysis of crowd work platforms. *Computing Research Repository*, 2013.
- Roman Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, pp. 210–268. Cambridge University Press, 2012.
- Kerri Wazny. “Crowdsourcing” ten years in: A review. *Journal of Global Health*, 7:020602, 2017.
- Hongxin Wei, Renchunzi Xie, Lei Feng, Bo Han, and Bo An. Deep learning from multiple noisy annotators as a union. *IEEE Trans. Neural Netw. Learn. Syst.*, pp. 1–11, 2022.
- Peter Welinder, Steve Branson, Pietro Perona, and Serge J Belongie. The multidimensional wisdom of crowds. In *Advances in Neural Information Processing Systems*, pp. 2424–2432, 2010.
- Jacob Whitehill, Ting fan Wu, Jacob Bergsma, Javier R. Movellan, and Paul L. Ruvolo. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *Advances in Neural Information Processing Systems*, volume 22, pp. 2035–2043. 2009.
- Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *Advances in Neural Information Processing Systems*, volume 32, 2019.

- Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu, Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent label noise. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7597–7610, 2020.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint*, arXiv:1708.07747, 2017.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *Proceedings of International Conference on Learning Representations*, 2016a.
- Le Zhang, Ryutaro Tanno, Mou-Cheng Xu, Chen Jin, Joseph Jacob, Olga Cicarrelli, Frederik Barkhof, and Daniel Alexander. Disentangling human error from ground truth in segmentation of medical images. In *Advances in Neural Information Processing Systems*, volume 33, pp. 15750–15762, 2020.
- Yuchen Zhang, Xi Chen, Dengyong Zhou, and Michael I. Jordan. Spectral methods meet EM: A provably optimal algorithm for crowdsourcing. *Journal of Machine Learning Research*, 17(102): 1–44, 2016b.
- Zhaowei Zhu, Jialu Wang, and Yang Liu. Beyond images: Label noise transition matrix estimation for tasks with lower-quality features. In *Proceedings of International Conference on Machine Learning*, volume 162, pp. 27633–27653, 2022.

Supplementary Material of “Deep Learning From Crowdsourced Labels: Coupled Cross-entropy Minimization, Identifiability, and Regularization”

A NOTATION

x , \mathbf{x} , and \mathbf{X} represent a scalar, a vector, and a matrix, respectively. $[\mathbf{x}]_i$ and $\mathbf{x}(i)$ both denote the i th entry of the vector \mathbf{x} . $\mathbf{X}(:, j)$ denotes the j th column vector of \mathbf{X} and $\mathbf{X}(i, :)$ denotes the i th row vector of \mathbf{X} . $[\mathbf{X}]_{i,j}$ and $\mathbf{X}(i, j)$ both mean the (i, j) th entry of \mathbf{X} . $\|\mathbf{x}\|_2$ and $\|\mathbf{X}\|_F$ mean the Euclidean (Frobenius) norm of the argument. $[I]$ means an integer set $\{1, 2, \dots, I\}$. \top denote transpose. $\Delta_K = \{\mathbf{x} \in \mathbb{R}^K : \sum_i \mathbf{x}(i) = 1, \mathbf{x} \geq \mathbf{0}\}$ denotes the probability simplex. $\mathbf{X} \geq \mathbf{0}$ implies that all the entries of the matrix \mathbf{X} are nonnegative. $\mathbb{1}[A]$ denotes an indicator function for the event A such that $\mathbb{1}[A] = 1$ if the event A happens, otherwise $\mathbb{1}[A] = 0$. $\text{CE}(\mathbf{x}, y) = -\sum_{k=1}^K \mathbb{1}[y = k] \log(\mathbf{x}(k))$ denotes the cross entropy function. \mathbf{I}_K denotes an identity matrix of size $K \times K$. $\sigma_{\max}(\mathbf{X})$ and $\sigma_{\min}(\mathbf{X})$ denote the largest and the smallest singular values of the matrix \mathbf{X} , respectively. $\text{cone}(\mathbf{X})$ denotes the conic hull formed by the columns of the matrix \mathbf{X} . $\det(\mathbf{X})$ represents the determinant of \mathbf{X} . $\text{vec}(\mathbf{X})$ denotes the vectorization operation that concatenates the columns of \mathbf{X} . $\text{trace}(\mathbf{X})$ outputs the trace (the sum of the diagonal elements) of \mathbf{X} .

B PROOF OF THEOREM 1

If the vectors $\mathbf{p}_n^{(m)}$ for all m, n are available, then one can represent the model in (2) as a low-rank *nonnegative matrix factorization* (NMF) model as follows:

$$\underbrace{\begin{bmatrix} \mathbf{p}_1^{(1)} & \dots & \mathbf{p}_N^{(1)} \\ \vdots & \ddots & \vdots \\ \mathbf{p}_1^{(M)} & \dots & \mathbf{p}_N^{(M)} \end{bmatrix}}_{\mathbf{P} \in \mathbb{R}^{MK \times N}} = \underbrace{\begin{bmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_M \end{bmatrix}}_{\mathbf{W} \in \mathbb{R}^{MK \times K}} \underbrace{\begin{bmatrix} \mathbf{f}(\mathbf{x}_1) & \dots & \mathbf{f}(\mathbf{x}_N) \end{bmatrix}}_{\mathbf{F} \in \mathbb{R}^{K \times N}}, \quad (10)$$

where the factors \mathbf{W} and \mathbf{F} are both nonnegative per their physical meaning.

Let \mathbf{P}^\natural denote the ground-truth and $\mathbf{p}_n^{\natural(m)} \in \mathbb{R}^K$ the (m, n) th block in \mathbf{P}^\natural , following the representation in (10). Note that we do not observe the entire \mathbf{P}^\natural but only the $\hat{\mathbf{y}}_n^{(m)}$ sampled from $\mathbf{p}_n^{\natural(m)}$ for $(m, n) \in \mathcal{S}$. We first show that the CCEM criterion in (6) *implicitly* estimates \mathbf{P}^\natural from incomplete labels indexed by \mathcal{S} . To this end, we consider the objective function in (6a):

$$D_{\mathcal{S}}(\mathbf{P}; \hat{\mathcal{Y}}) \triangleq -\frac{1}{S} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{1}[\hat{\mathbf{y}}_n^{(m)} = k] \log \mathbf{P}((m-1)K + k, n), \quad (11)$$

where $\hat{\mathcal{Y}}$ denotes the set of observed noisy labels, i.e., $\{\hat{\mathbf{y}}_n^{(m)}\}_{(m,n) \in \mathcal{S}}$. Let $\{\hat{\mathbf{A}}_m\}_{m=1}^M$ and $\hat{\mathbf{f}}$ denote the estimates given by the learning criterion in (6). Then, the criterion (6) helps define the following term:

$$\hat{\mathbf{P}} \triangleq \arg \min_{\mathbf{P} \in \mathcal{P}} D_{\mathcal{S}}(\mathbf{P}; \hat{\mathcal{Y}}), \quad (12)$$

where $\hat{\mathbf{P}} = \hat{\mathbf{W}} [\hat{\mathbf{f}}(\mathbf{x}_1) \dots \hat{\mathbf{f}}(\mathbf{x}_N)]$, $\hat{\mathbf{W}} = [\hat{\mathbf{A}}_1^\top, \dots, \hat{\mathbf{A}}_M^\top]^\top$,

$$\mathcal{P} \triangleq \{\mathbf{P} \in \mathbb{R}^{MK \times N} \mid \mathbf{P} = \mathbf{W} [\mathbf{f}(\mathbf{x}_1) \dots \mathbf{f}(\mathbf{x}_N)], \mathbf{W} \in \mathcal{W}, \mathbf{f} \in \mathcal{F}\},$$

$$\mathcal{W} \triangleq \{\mathbf{W} = [\mathbf{A}_1^\top, \dots, \mathbf{A}_M^\top]^\top \in \mathbb{R}^{MK \times K} \mid \mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top, \mathbf{A}_m \geq \mathbf{0}, \forall m\}, \text{ and}$$

$$\mathcal{F} \subset \{\mathbf{f}(\mathbf{x}) \in \mathbb{R}^K \mid \mathbf{f}(\mathbf{x}) \in \Delta_K, \forall \mathbf{x} \in \mathbb{R}^D\}$$
 is the neural network function class.

Our main goal is to characterize the estimation errors of $\{\hat{\mathbf{A}}_m\}_{m=1}^M$ and $\hat{\mathbf{f}}$. This can be done via first bounding the estimation errors $\|\mathbf{p}_n^{\natural(m)} - \hat{\mathbf{p}}_n^{(m)}\|_2^2$, where $\mathbf{p}_n^{\natural(m)}$ is given by:

$$\mathbf{p}_n^{\natural(m)} = \mathbf{A}_m^\natural \mathbf{f}^\natural(\mathbf{x}_n), \forall m, n,$$

in which \mathbf{A}_m^\natural and \mathbf{f}^\natural denote the m th ground-truth confusion matrix and the GTP, respectively, and $\hat{\mathbf{p}}_n^{(m)}$ denote the (m, n) th block in $\hat{\mathbf{P}}$.

B.1 ESTIMATING $\|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2$

We first show the following proposition:

Proposition 1 *Under the assumptions in Theorem 1, the following result hold with probability at least $1 - \delta$:*

$$\frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \leq 4\sqrt{2}K\beta\mathfrak{R}_S(\mathcal{P}) + 10 \log(\beta) \frac{\sqrt{2 \log(4/\delta)}}{\sqrt{S}} + 2\beta\sqrt{K}\nu, \quad (13)$$

where $\mathfrak{R}_S(\mathcal{P}) = \frac{16}{\sqrt{S}} \sqrt{MK \log(4S\sqrt{K})} + (2\|\mathbf{X}\|_{\mathbb{F}} \mathfrak{R}_{\mathcal{F}})^{\frac{1}{2}}$ and is related to the empirical Rademacher complexity of the set \mathcal{P} under the observed samples \mathcal{S} .

The proof is provided in Sec. C. Proposition 1 provides an upper-bound for $\|\mathbf{P}^{\natural} - \widehat{\mathbf{P}}\|_{\mathbb{F}}^2$. To achieve the main goal, i.e., characterizing the estimation errors of $\{\widehat{\mathbf{A}}_m\}_{m=1}^M$ and $\widehat{\mathbf{f}}$, we require an upper-bound for $\sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2$ as well. The following result gives a tighter upper bound for $\sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2$:

Proposition 2 *Under the assumptions in Theorem 1, for any $\alpha > 0$ and $\delta \in (0, 1)$, the following result holds with probability at least $1 - \frac{1}{N^\alpha}$:*

$$\begin{aligned} & \frac{1}{M} \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \\ & \leq N^\alpha \left(4\sqrt{2}K\beta\mathfrak{R}_S(\mathcal{P}) + 10 \log(\beta) \frac{\sqrt{2 \log(4/\delta)}}{\sqrt{S}} + 2\beta\sqrt{K}\nu + 4\delta \right), \quad \forall n, \end{aligned} \quad (14)$$

where $\mathfrak{R}_S(\mathcal{P}) = \frac{16}{\sqrt{S}} \sqrt{MK \log(4S\sqrt{K})} + (2\|\mathbf{X}\|_{\mathbb{F}} \mathfrak{R}_{\mathcal{F}})^{\frac{1}{2}}$.

The proof is provided in Sec. I. Note that the bound in (14) looks divergent at the first glance, as the second term in the R.H.S. could go to infinity. However, given a large enough S , one can choose appropriate α and δ such that the bound in (14) does not explode.

B.2 ESTIMATION ERRORS FOR $\widehat{\mathbf{A}}_m$ AND $\widehat{\mathbf{f}}$

In this section, we characterize the estimation error of the confusion matrices and the GTP. We start with using simplified notations to represent the results in Propositions 1-2, i.e.,

$$\|\mathbf{P}^{\natural} - \widehat{\mathbf{P}}\|_{\mathbb{F}}^2 \leq \zeta^2, \quad (15)$$

$$\sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \leq \varphi^2, \quad \forall n. \quad (16)$$

with probability greater than $1 - \delta$ and $1 - \frac{1}{N^\alpha}$, respectively.

Assumptions 4-5 imply that there exists index sets $\Lambda = \{\tilde{m}_1, \dots, \tilde{m}_K\}$ and $\Psi = \{\tilde{n}_1, \dots, \tilde{n}_K\}$ such that:

$$\begin{aligned} \mathbf{W}^{\natural}(\Lambda, \cdot) &= \mathbf{I}_K + \mathbf{N}_W \\ \mathbf{F}^{\natural}(\cdot, \Psi) &= \mathbf{I}_K + \mathbf{N}_F, \end{aligned}$$

where $\|\mathbf{N}_W\|_{\mathbb{F}} \leq \sqrt{K}\xi_1$, and $\|\mathbf{N}_F\|_{\mathbb{F}} \leq \sqrt{K}\xi_2$. Note that the conditions do not require any confusion matrices to be near identity; i.e., we do not need any annotators to be all-class specialists.

Nonetheless, for notation simplicity, we let $\Lambda = \{1, \dots, K\}$ and $\Psi = \{1, \dots, K\}$, which is without loss of generality for our proof in the sequel. Under this simplification, the following holds:

$$\mathbf{W}^{\natural} = \begin{bmatrix} \mathbf{A}_1^{\natural} \\ \mathbf{A}_2^{\natural} \\ \vdots \\ \mathbf{A}_M^{\natural} \end{bmatrix}, \quad \mathbf{F}^{\natural} = [\mathbf{f}^{\natural}(\mathbf{x}_1) \quad \dots \quad \mathbf{f}^{\natural}(\mathbf{x}_N)] = [\mathbf{F}_1^{\natural} \quad \mathbf{F}_2^{\natural}],$$

where $\mathbf{A}_1^{\natural} = \mathbf{I}_K + \mathbf{N}_W$, $\mathbf{F}_1^{\natural} = \mathbf{I}_K + \mathbf{N}_F$, and $\mathbf{F}_2^{\natural} \in \mathbb{R}^{K \times (N-K)}$, $\|\mathbf{N}_W\|_F \leq \sqrt{K}\xi_1$, and $\|\mathbf{N}_F\|_F \leq \sqrt{K}\xi_2$.

Let $\widehat{\mathbf{A}}_m$ and $\widehat{\mathbf{f}}$ denote the estimates of \mathbf{A}_m^{\natural} and \mathbf{f}^{\natural} , respectively, using the learning criterion (6) and construct

$$\widehat{\mathbf{W}} = \begin{bmatrix} \widehat{\mathbf{A}}_1 \\ \widehat{\mathbf{A}}_2 \\ \vdots \\ \widehat{\mathbf{A}}_M \end{bmatrix}, \quad \widehat{\mathbf{F}} = [\widehat{\mathbf{f}}(\mathbf{x}_1) \quad \dots \quad \widehat{\mathbf{f}}(\mathbf{x}_N)] = [\widehat{\mathbf{F}}_1 \quad \widehat{\mathbf{F}}_2],$$

where $\widehat{\mathbf{A}}_m \in \mathbb{R}^{K \times K}$, $\widehat{\mathbf{F}}_1 \in \mathbb{R}^{K \times K}$, and $\widehat{\mathbf{F}}_2 \in \mathbb{R}^{K \times (N-K)}$. Then, we have:

$$\begin{aligned} \|\mathbf{P}^{\natural} - \widehat{\mathbf{P}}\|_F^2 &= \|\mathbf{W}^{\natural} \mathbf{F}^{\natural} - \widehat{\mathbf{W}} \widehat{\mathbf{F}}\|_F^2 \\ &= \|\mathbf{I}_K + \mathbf{N}_W + \mathbf{N}_F + \mathbf{N}_W \mathbf{N}_F - \widehat{\mathbf{A}}_1 \widehat{\mathbf{F}}_1\|_F^2 + \|\mathbf{F}_2^{\natural} + \mathbf{N}_W \mathbf{F}_2^{\natural} - \widehat{\mathbf{A}}_1 \widehat{\mathbf{F}}_2\|_F^2 \\ &\quad + \sum_{m=2}^M \|\mathbf{A}_m^{\natural} + \mathbf{A}_m^{\natural} \mathbf{N}_F - \widehat{\mathbf{A}}_m \widehat{\mathbf{F}}_1\|_F^2 + \sum_{m=2}^M \|\mathbf{A}_m^{\natural} \mathbf{F}_2^{\natural} - \widehat{\mathbf{A}}_m \widehat{\mathbf{F}}_2\|_F^2. \end{aligned} \quad (17)$$

Let us define the error matrices as below:

$$\mathbf{E}_A^{(1)} \triangleq \widehat{\mathbf{A}}_1 - \underbrace{(\mathbf{I}_K + \mathbf{N}_W)\mathbf{\Pi}}_{\mathbf{A}_1^{\natural}\mathbf{\Pi}}, \quad \mathbf{E}_A^{(m)} \triangleq \widehat{\mathbf{A}}_m - \mathbf{A}_m^{\natural}\mathbf{\Pi}, \quad \forall m > 1, \quad (18)$$

$$\mathbf{E}_F^{(1)} \triangleq \widehat{\mathbf{F}}_1 - \underbrace{\mathbf{\Pi}^\top (\mathbf{I}_K + \mathbf{N}_F)}_{\mathbf{\Pi}^\top \mathbf{F}_1^{\natural}}, \quad \mathbf{E}_F^{(2)} \triangleq \widehat{\mathbf{F}}_2 - \mathbf{\Pi}^\top \mathbf{F}_2^{\natural}, \quad (19)$$

where $\mathbf{\Pi} \in [0, 1]^{K \times K}$ is a column permutation matrix and is the same across all the error blocks. We hope to characterize the norm of these error matrices.

Upper bound for $\|\mathbf{E}_A^{(1)}\|_F$ and $\|\mathbf{E}_F^{(1)}\|_F$. We start by considering the first term on the R.H.S. of (17), i.e., $\|\mathbf{I}_K + \mathbf{N}_W + \mathbf{N}_F + \mathbf{N}_W \mathbf{N}_F - \widehat{\mathbf{A}}_1 \widehat{\mathbf{F}}_1\|_F^2$. From (16), we have the following with probability greater than $1 - \frac{K}{N^\alpha}$:

$$\begin{aligned} \varphi^2 &\geq \|\mathbf{I}_K + \mathbf{N}_W + \mathbf{N}_F + \mathbf{N}_W \mathbf{N}_F - \widehat{\mathbf{A}}_1 \widehat{\mathbf{F}}_1\|_F^2 \\ &= \sum_{k=1}^K |1 + \widetilde{\mathbf{N}}(k, k) - \widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, k)|^2 + \sum_{j=1}^K \sum_{k \neq j} |\widetilde{\mathbf{N}}(j, k) - \widehat{\mathbf{A}}_1(j, :) \widehat{\mathbf{F}}_1(:, k)|^2, \end{aligned} \quad (20)$$

where $\widetilde{\mathbf{N}} = \mathbf{N}_W + \mathbf{N}_F + \mathbf{N}_W \mathbf{N}_F$. The absolute value of the largest entry in $\widetilde{\mathbf{N}}$ can be bounded by $\xi_1 + \xi_2 + \sqrt{K}\xi_1\xi_2$. Let us denote $\kappa = \varphi + \xi_1 + \xi_2 + \sqrt{K}\xi_1\xi_2$ for conciseness. Then, from (20), we get the following conditions:

$$1 - \kappa \leq |\widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, k)| \leq 1 + \kappa, \quad \forall k \in [K] \quad (21a)$$

$$|\widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, j)| \leq \kappa, \quad \forall k \neq j. \quad (21b)$$

From the conditions (21a) and (21b), we have the following result:

Lemma 1 Assume that the conditions in (21a) and (21b) hold and that $\kappa \leq \frac{1}{K+1}$. Then, we have the following relations satisfied:

$$\arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) \neq \arg \max_{\ell} \widehat{\mathbf{A}}_1(j, \ell), \forall k \neq j \quad (22a)$$

$$\arg \max_q \widehat{\mathbf{F}}_1(q, k) \neq \arg \max_q \widehat{\mathbf{F}}_1(q, j), \forall k \neq j \quad (22b)$$

$$\arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) = \arg \max_q \widehat{\mathbf{F}}_1(q, k), \forall k. \quad (22c)$$

The proof of the lemma is given in Sec. F.

From the lower bound condition in (21a), we have the following result:

$$\begin{aligned} \sum_{k=1}^K (1 - \kappa) &\leq \sum_{k=1}^K \widehat{\mathbf{A}}_1(k, 1) \widehat{\mathbf{F}}_1(1, k) + \cdots + \widehat{\mathbf{A}}_1(k, K) \widehat{\mathbf{F}}_1(K, k) \\ &\leq \sum_{k=1}^K \max_{\ell \in [K]} \widehat{\mathbf{A}}_1(k, \ell) (\widehat{\mathbf{F}}_1(1, k) + \cdots + \widehat{\mathbf{F}}_1(K, k)) \\ &= \sum_{k=1}^K \max_{\ell \in [K]} \widehat{\mathbf{A}}_1(k, \ell) \end{aligned} \quad (23)$$

where the last equality is obtained due to the probability simplex constraints on the columns of $\widehat{\mathbf{F}}_1$.

From the lower bound condition in (21a), we have the below result as well:

$$\begin{aligned} \sum_{k=1}^K (1 - \kappa) &\leq \sum_{k=1}^K \left(\widehat{\mathbf{A}}_1(k, 1) \widehat{\mathbf{F}}_1(1, k) + \cdots + \widehat{\mathbf{A}}_1(k, K) \widehat{\mathbf{F}}_1(K, k) \right) \\ &\leq \sum_{k=1}^K \max_{\ell} \widehat{\mathbf{F}}_1(k, \ell) (\widehat{\mathbf{A}}_1(1, k) + \cdots + \widehat{\mathbf{A}}_1(K, k)) \\ &= \sum_{k=1}^K \max_{\ell} \widehat{\mathbf{F}}_1(k, \ell). \end{aligned} \quad (24)$$

where the last equality is obtained due to the probability simplex constraints on all the columns of $\widehat{\mathbf{A}}_1$.

Next, we characterize the term $\|\mathbf{E}_A^{(1)}\|_{\mathbb{F}}^2$. To achieve this, by employing the result (22a) in Lemma 1, we fix the column permutation $\mathbf{\Pi}$ as follows:

$$\mathbf{\Pi}(j, k) = \begin{cases} 1, & j = \arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) \\ 0, & \text{otherwise.} \end{cases} \quad (25)$$

We get the following set of relations:

$$\begin{aligned}
\|\widehat{\mathbf{A}}_1 - \mathbf{I}_K \mathbf{\Pi}\|_{\text{F}}^2 &= \sum_{k=1}^K \|\widehat{\mathbf{A}}_1(k, :) - \mathbf{e}_k^\top \mathbf{\Pi}\|_2^2 \\
&= \sum_{k=1}^K \|\widehat{\mathbf{A}}_1(k, :)\|_2^2 + K - 2 \sum_{k=1}^K \widehat{\mathbf{A}}_1(k, :)\mathbf{\Pi}^\top \mathbf{e}_k \\
&= \|\widehat{\mathbf{A}}_1\|_{\text{F}}^2 + K - 2 \sum_{k=1}^K \widehat{\mathbf{A}}_1(k, :)\mathbf{\Pi}^\top \mathbf{e}_k \\
&\leq 2K - 2 \sum_{k=1}^K \widehat{\mathbf{A}}_1(k, :)\mathbf{\Pi}^\top \mathbf{e}_k \\
&\stackrel{(a)}{=} 2K - 2 \sum_{k=1}^K \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) \\
&\stackrel{(b)}{\leq} 2K - 2 \sum_{k=1}^K (1 - \kappa) = 2K\kappa,
\end{aligned}$$

where the last inequality (b) is by (23) and the relation (a) is obtained by choosing $\mathbf{\Pi}$ as defined in (25).

Hence, we have the following with probability greater than $1 - \frac{K}{N^\alpha}$:

$$\|\mathbf{E}_A^{(1)}\|_{\text{F}} \leq \|\widehat{\mathbf{A}}_1 - \mathbf{I}_K \mathbf{\Pi}\|_{\text{F}} + \|\mathbf{N}_W\|_{\text{F}} \leq \sqrt{2K\kappa} + \sqrt{K}\xi_1. \quad (26)$$

We proceed to consider the error term $\|\mathbf{E}_F^{(1)}\|_{\text{F}}^2$. To achieve this, consider the following:

$$\begin{aligned}
\|\widehat{\mathbf{F}}_1 - \mathbf{\Pi}^\top \mathbf{I}_K\|_{\text{F}}^2 &= \sum_{k=1}^K \|\widehat{\mathbf{F}}_1(k, :) - \mathbf{e}_k^\top \mathbf{\Pi}^\top\|_2^2 \\
&= \sum_{k=1}^K \|\widehat{\mathbf{F}}_1(k, :)\|_2^2 + K - 2 \sum_{k=1}^K \widehat{\mathbf{F}}_1(k, :)\mathbf{\Pi}^\top \mathbf{e}_k \\
&= \|\widehat{\mathbf{F}}_1\|_{\text{F}}^2 + K - 2 \sum_{k=1}^K \widehat{\mathbf{F}}_1(k, :)\mathbf{\Pi}^\top \mathbf{e}_k \\
&\leq 2K - 2 \sum_{k=1}^K \widehat{\mathbf{F}}_1(k, :)\mathbf{\Pi}^\top \mathbf{e}_k \\
&\stackrel{(a)}{=} 2K - 2 \sum_{k=1}^K \max_{\ell} \widehat{\mathbf{F}}_1(k, \ell) \\
&\stackrel{(b)}{\leq} 2K - 2 \sum_{k=1}^K (1 - \kappa) = 2K\kappa,
\end{aligned}$$

where the last inequality (b) is by (24) and the relation (a) is by combining the definition of $\mathbf{\Pi}$ in (25) and (22c) in Lemma 1. Hence, we get the following with probability greater than $1 - \frac{K}{N^\alpha}$:

$$\|\mathbf{E}_F^{(1)}\|_{\text{F}} \leq \|\widehat{\mathbf{F}}_1 - \mathbf{I}_K \mathbf{\Pi}\|_{\text{F}} + \|\mathbf{N}_F\|_{\text{F}} \leq \sqrt{2K\kappa} + \sqrt{K}\xi_2. \quad (27)$$

Upper bound for $\|\mathbf{E}_F^{(2)}\|_F$. Next, we consider the second term in (17), i.e., $\|\mathbf{F}_2^\natural + \mathbf{N}_W \mathbf{F}_2^\natural - \widehat{\mathbf{A}}_1 \widehat{\mathbf{F}}_2\|_F$. From (15), we have the following with probability greater than $1 - \delta$:

$$\begin{aligned}
\zeta &\geq \|\mathbf{F}_2^\natural + \mathbf{N}_W \mathbf{F}_2^\natural - \widehat{\mathbf{A}}_1 \widehat{\mathbf{F}}_2\|_F = \|\mathbf{F}_2^\natural + \mathbf{N}_W \mathbf{F}_2^\natural - (\mathbf{I}_K \mathbf{\Pi} + \mathbf{N}_W \mathbf{\Pi} + \mathbf{E}_A^{(1)})(\mathbf{E}_F^{(2)} + \mathbf{\Pi}^\top \mathbf{F}_2^\natural)\|_F \\
&= \|(\mathbf{I}_K \mathbf{\Pi} + \mathbf{N}_W \mathbf{\Pi} + \mathbf{E}_A^{(1)})\mathbf{E}_F^{(2)} + \mathbf{E}_A^{(1)} \mathbf{\Pi}^\top \mathbf{F}_2^\natural\|_F \\
&= \|\widehat{\mathbf{A}}_1 \mathbf{E}_F^{(2)} + \mathbf{E}_A^{(1)} \mathbf{\Pi}^\top \mathbf{F}_2^\natural\|_F \\
&\stackrel{(a)}{\geq} \left| \|\widehat{\mathbf{A}}_1 \mathbf{E}_F^{(2)}\|_F - \|\mathbf{E}_A^{(1)} \mathbf{\Pi}^\top \mathbf{F}_2^\natural\|_F \right| \\
&\stackrel{(b)}{\geq} \left| \sigma_{\min}(\widehat{\mathbf{A}}_1) \|\mathbf{E}_F^{(2)}\|_F - \|\mathbf{E}_A^{(1)}\|_F \|\mathbf{F}_2^\natural\|_F \right| \\
&\stackrel{(c)}{\geq} \left| \sigma_{\min}(\widehat{\mathbf{A}}_1) \|\mathbf{E}_F^{(2)}\|_F - (\sqrt{2K\kappa} + \sqrt{K}\xi_1) \|\mathbf{F}_2^\natural\|_F \right| \tag{28}
\end{aligned}$$

where the inequality (a) is by using the triangle inequality, (b) is obtained by applying the following two relations for any two matrices $\mathbf{A} \in \mathbb{R}^{K \times K}$, $\mathbf{B} \in \mathbb{R}^{K \times L}$, $L \geq K$:

$$\|\mathbf{A}\mathbf{B}\|_F \geq \sigma_{\min}(\mathbf{A})\|\mathbf{B}\|_F \tag{29}$$

$$\|\mathbf{A}\mathbf{B}\|_F \leq \|\mathbf{A}\|_F \|\mathbf{B}\|_F. \tag{30}$$

The inequality (c) is obtained by applying (26).

Hence, (28) combined with the fact that $\|\mathbf{F}_2\|_F \leq \|\mathbf{F}^\natural\|_F$ gives the following relation with probability greater than $1 - \delta - \frac{K}{N^\alpha}$:

$$\|\mathbf{E}_F^{(2)}\|_F \leq \frac{\zeta + (\sqrt{2K\kappa} + \sqrt{K}\xi_1) \|\mathbf{F}^\natural\|_F}{\sigma_{\min}(\widehat{\mathbf{A}}_1)}. \tag{31}$$

Next, we use the following lemma to characterize $\sigma_{\min}(\widehat{\mathbf{A}}_1)$:

Lemma 2 Suppose that the matrix $\mathbf{X} \in \mathbb{R}^{K \times K}$ takes the form $\mathbf{X} = \mathbf{I}_K + \mathbf{E}_1 + \mathbf{E}_2$. Assume that $\|\mathbf{E}_1\|_F \leq v_1$ and $\|\mathbf{E}_2\|_F \leq v_2$, for a certain $v_1, v_2 > 0$. Then, we have

$$\sigma_{\min}(\mathbf{X}) \geq |1 - v_1 - v_2|.$$

The proof is relegated to Sec. G. Applying Lemma 2, we get the following with probability greater than $1 - \delta - \frac{K}{N^\alpha}$:

$$\|\mathbf{E}_F^{(2)}\|_F \leq \frac{\zeta + (\sqrt{2K\kappa} + \sqrt{K}\xi_1) \|\mathbf{F}^\natural\|_F}{|1 - \sqrt{2K\kappa} - 2\sqrt{K}\xi_1|} \leq \frac{c_1(\zeta + \sqrt{K\kappa} \|\mathbf{F}^\natural\|_F)}{|1 - \sqrt{K\kappa}|} \tag{32}$$

for certain constant $c_1 > 0$. where we used the fact that $\xi_1 \leq \sqrt{\kappa}$ since $\kappa = \varphi + \xi_1 + \xi_2 + \sqrt{K}\xi_1\xi_2$ and $\kappa \leq 1$.

Upper bound for $\|\mathbf{E}_A^{(m)}\|_F$, $m > 1$. We consider the third term on the R.H.S. of (17). From (16), for each m , we have the following with probability greater than $1 - \frac{K}{N^\alpha}$:

$$\begin{aligned}
\sqrt{K}\varphi &\geq \|\mathbf{A}_m^\natural + \mathbf{A}_m^\natural \mathbf{N}_F - \widehat{\mathbf{A}}_m \widehat{\mathbf{F}}_1\|_F = \|\mathbf{A}_m^\natural + \mathbf{A}_m^\natural \mathbf{N}_F - (\mathbf{A}_m^\natural \mathbf{\Pi} + \mathbf{E}_A^{(m)})(\mathbf{\Pi}^\top \mathbf{I}_K + \mathbf{\Pi}^\top \mathbf{N}_F + \mathbf{E}_F^{(1)})\|_F \\
&= \|\mathbf{E}_A^{(m)}(\mathbf{\Pi}^\top \mathbf{I}_K + \mathbf{\Pi}^\top \mathbf{N}_F + \mathbf{E}_F^{(1)}) + \mathbf{A}_m^\natural \mathbf{\Pi} \mathbf{E}_F^{(1)}\|_F \\
&= \|\mathbf{E}_A^{(m)} \widehat{\mathbf{F}}_1 + \mathbf{A}_m^\natural \mathbf{\Pi} \mathbf{E}_F^{(1)}\|_F \\
&\stackrel{(a)}{\geq} \left| \|\mathbf{E}_A^{(m)} \widehat{\mathbf{F}}_1\|_F - \|\mathbf{A}_m^\natural \mathbf{\Pi} \mathbf{E}_F^{(1)}\|_F \right| \\
&\stackrel{(b)}{\geq} \left| \sigma_{\min}(\widehat{\mathbf{F}}_1) \|\mathbf{E}_A^{(m)}\|_F - \sigma_{\max}(\mathbf{A}_m^\natural) \|\mathbf{E}_F^{(1)}\|_F \right| \\
&\stackrel{(c)}{\geq} \left| \sigma_{\min}(\widehat{\mathbf{F}}_1) \|\mathbf{E}_A^{(m)}\|_F - (\sqrt{2K\kappa} + \sqrt{K}\xi_2) \sigma_{\max}(\mathbf{A}_m^\natural) \right|,
\end{aligned}$$

where the relation (a) by the triangle inequality, (b) is by applying (29) and (30), and (c) is via (27).

Hence, $\forall m > 1$, we get the following with probability greater than $1 - \frac{2K}{N^\alpha}$:

$$\|\mathbf{E}_A^{(m)}\|_F \leq \frac{\sqrt{K}\varphi + (\sqrt{2K\kappa} + \sqrt{K}\xi_2)\sigma_{\max}(\mathbf{A}_m^{\natural})}{|1 - \sqrt{2K\kappa} - 2\sqrt{K}\xi_2|} \leq \frac{c_2\sqrt{K\kappa}\sigma_{\max}(\mathbf{A}_m^{\natural})}{|1 - \sqrt{K\kappa}|} \quad (33)$$

for certain constant $c_2 > 0$, where we used the fact that $\varphi \leq \sqrt{\kappa}$ and $\xi_2 \leq \sqrt{\kappa}$ since $\kappa = \varphi + \xi_1 + \xi_2 + \sqrt{K}\xi_1\xi_2$ and $\kappa \leq 1$. In the above, we have also applied $\sigma_{\min}(\widehat{\mathbf{F}}_1) \geq |1 - \sqrt{2K\kappa} - 2\sqrt{K}\xi_2|$ following Lemma 2 and (27).

Putting Together. From (33), we have the following with probability greater than $1 - \frac{2K}{N^\alpha}$:

$$\|\mathbf{E}_A^{(m)}\|_F^2 = \|\widehat{\mathbf{A}}_m - \mathbf{A}_m^{\natural}\mathbf{\Pi}\|_F^2 \leq \frac{c_2^2 K^2 \kappa}{|1 - \sqrt{K\kappa}|^2}, \quad \forall m. \quad (34)$$

where we have used the fact that $\sigma_{\max}(\mathbf{A}_m^{\natural}) \leq \|\mathbf{A}_m^{\natural}\|_F \leq \sqrt{K}$.

Similarly, by combining (27) and (32), we have the following with probability greater than $1 - \delta - \frac{2K}{N^\alpha}$

$$\begin{aligned} \|\widehat{\mathbf{F}} - \mathbf{\Pi}^\top \mathbf{F}^{\natural}\|_F^2 &= \|\mathbf{E}_F^{(1)}\|_F^2 + \|\mathbf{E}_F^{(2)}\|_F^2 \\ &\leq (\sqrt{2K\kappa} + \sqrt{K}\xi_2)^2 + \frac{c_1^2(\zeta + \sqrt{K\kappa}\|\mathbf{F}^{\natural}\|_F)^2}{|1 - \sqrt{K\kappa}|^2} \\ &\leq \frac{c_3(\zeta + \sqrt{NK\kappa})^2}{|1 - \sqrt{K\kappa}|^2} \end{aligned} \quad (35)$$

for certain constant $c_3 > 0$, where we have used the fact that $\|\mathbf{F}^{\natural}\|_F \leq \sqrt{N}$.

We hope to characterize the generalization performance of the predicted function $\widehat{\mathbf{f}}$ from (35). Towards this, we have the following result:

Lemma 3 *Under Assumptions 1 and 2, the following holds with probability greater than $1 - \delta$:*

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\|\widehat{\mathbf{f}}(\mathbf{x}) - \mathbf{\Pi}^\top \mathbf{f}^{\natural}(\mathbf{x})\|_2^2 \right] &\leq \frac{1}{N} \|\widehat{\mathbf{F}} - \mathbf{\Pi}^\top \mathbf{F}^{\natural}\|_F^2 + 64N^{-5/8} (2\|\mathbf{X}\|_F \mathfrak{R}_{\mathcal{F}})^{\frac{1}{4}} \\ &\quad + 16\sqrt{\frac{2\log(4/\delta)}{N}}. \end{aligned} \quad (36)$$

The proof is given in Sec. H. Hence, combining Lemma 3 with (35), we have the following with probability greater than $1 - 2\delta - \frac{2K}{N^\alpha}$

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\|\widehat{\mathbf{f}}(\mathbf{x}) - \mathbf{\Pi}^\top \mathbf{f}^{\natural}(\mathbf{x})\|_2^2 \right] \leq \frac{c_3(\zeta + \sqrt{NK\kappa})^2}{N|1 - \sqrt{K\kappa}|^2} + 64N^{-5/8} (2\|\mathbf{X}\|_F \mathfrak{R}_{\mathcal{F}})^{\frac{1}{4}} + 16\sqrt{\frac{2\log(4/\delta)}{N}}, \quad (37)$$

where

$$\begin{aligned} \kappa &\leq \varphi + 2\sqrt{K}(\xi_1 + \xi_2), \\ \zeta^2 &\leq 4\sqrt{2}NMK\beta\mathfrak{R}_{\mathcal{S}}(\mathcal{P}) + 10NM \log(\beta) \frac{\sqrt{2\log(4/\delta)}}{\sqrt{S}} + 2NM\beta\sqrt{K}\nu \\ \varphi^2 &\leq 4\sqrt{2}N^\alpha MK\beta\mathfrak{R}_{\mathcal{S}}(\mathcal{P}) + 10N^\alpha M \log(\beta) \frac{\sqrt{2\log(4/\delta)}}{\sqrt{S}} + 2N^\alpha M\beta\sqrt{K}\nu + 4N^\alpha M\delta, \\ \mathfrak{R}_{\mathcal{S}}(\mathcal{P}) &= \frac{16}{\sqrt{S}} \sqrt{MK \log(4S\sqrt{K}) + (2\|\mathbf{X}\|_F \mathfrak{R}_{\mathcal{F}})^{\frac{1}{2}}}. \end{aligned}$$

Also note that the final bounds in (34) and (37) are satisfied only if $\kappa \leq \frac{1}{K+1}$ as given by Lemma 1. This gives the final conditions on ξ_1 , ξ_2 , ν , and S . By letting $\delta = \frac{1}{S}$, we obtain the final results in Theorem 1.

C PROOF OF PROPOSITION 1

Let us consider the following notation:

$$\mathbf{P}(\omega) = \mathbf{p}_n^{(m)}, \quad (38)$$

where $\omega = (m, n) \in [M] \times [N]$; i.e., $\mathbf{P}(\omega)$ “reads out” the (m, n) th block in (10). We also define $\hat{\mathcal{Y}}(\omega) = \hat{y}_n^{(m)}$. With this notation, we have

$$D_S(\mathbf{P}; \hat{\mathcal{Y}}) = \frac{1}{S} \sum_{s=1}^S \text{CE}(\mathbf{P}(\omega_s), \hat{\mathcal{Y}}(\omega_s)),$$

where $\text{CE}(\mathbf{x}, y) = -\sum_{k=1}^K \mathbb{1}[y = k] \log(\mathbf{x}(k))$, $S = |\mathcal{S}|$, and $\omega_s = (m_s, n_s) \in \mathcal{S}$. Under Assumption 1, we can define:

$$D_{\Pi}(\mathbf{P}, \hat{\mathcal{Y}}) \triangleq \mathbb{E}_{S \sim \Pi}[\text{CE}(\mathbf{P}(\omega), \hat{\mathcal{Y}}(\omega))] = \sum_{m=1}^M \sum_{n=1}^N \pi_n^{(m)} \text{CE}(\mathbf{p}_n^{(m)}, \hat{y}_n^{(m)}),$$

where Π denotes the uniform distribution and $\pi_n^{(m)} = \frac{1}{NM}$ denotes the probability of observing the annotation for the index pair (m, n) .

Eq. (12) implies

$$D_S(\hat{\mathbf{P}}; \hat{\mathcal{Y}}) \leq D_S(\tilde{\mathbf{P}}; \hat{\mathcal{Y}}), \quad (39)$$

where $\tilde{\mathbf{P}}$ is defined using the following construction:

$$\tilde{\mathbf{P}} = \mathbf{W}^{\natural} [\tilde{\mathbf{f}}(\mathbf{x}_1) \quad \dots \quad \tilde{\mathbf{f}}(\mathbf{x}_N)],$$

and $\tilde{\mathbf{f}}$ is a learning function constructed under Assumption 3. To be specific, $\tilde{\mathbf{f}}$ satisfies

$$\|\tilde{\mathbf{f}}(\mathbf{x}) - \mathbf{f}^{\natural}(\mathbf{x})\|_2 \leq \nu, \quad \forall \mathbf{x} \sim \mathcal{D}.$$

Hence, by taking expectation w.r.t. $\hat{\mathcal{Y}}$ we have

$$\begin{aligned} \mathbb{E}[D_{\Pi}(\hat{\mathbf{P}}; \hat{\mathcal{Y}}) - D_{\Pi}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}})] &= \mathbb{E}[D_{\Pi}(\hat{\mathbf{P}}; \hat{\mathcal{Y}})] - D_S(\hat{\mathbf{P}}; \hat{\mathcal{Y}}) + D_S(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}})] \\ &\quad + D_S(\hat{\mathbf{P}}; \hat{\mathcal{Y}}) - D_S(\tilde{\mathbf{P}}; \hat{\mathcal{Y}}) + D_S(\tilde{\mathbf{P}}; \hat{\mathcal{Y}}) - D_S(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) \\ &\leq \mathbb{E}[D_{\Pi}(\hat{\mathbf{P}}; \hat{\mathcal{Y}})] - D_S(\hat{\mathbf{P}}; \hat{\mathcal{Y}}) + D_S(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}})] \\ &\quad + D_S(\tilde{\mathbf{P}}; \hat{\mathcal{Y}}) - D_S(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) \\ &\leq \sup_{\mathbf{P} \in \mathcal{P}} \left| D_S(\mathbf{P}; \hat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathbf{P}; \hat{\mathcal{Y}})] \right| + \left| D_S(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}})] \right| \\ &\quad + \left| D_S(\tilde{\mathbf{P}}; \hat{\mathcal{Y}}) - D_S(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) \right|, \end{aligned} \quad (40)$$

where the first inequality is obtained from (39).

Let us consider the L.H.S. of (40):

$$\begin{aligned} &\mathbb{E} \left[D_{\Pi}(\hat{\mathbf{P}}; \hat{\mathcal{Y}}) - D_{\Pi}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) \right] \\ &= \sum_{n=1}^N \sum_{m=1}^M \left(\pi_n^{(m)} \sum_{k=1}^K -\mathbf{P}^{\natural}((m-1)K + k, n) \log \hat{\mathbf{P}}((m-1)K + k, n) \right. \\ &\quad \left. + \mathbf{P}^{\natural}((m-1)K + k, n) \log \mathbf{P}^{\natural}((m-1)K + k, n) \right) \\ &= \sum_{n=1}^N \sum_{m=1}^M \pi_n^{(m)} \sum_{k=1}^K \mathbf{P}^{\natural}((m-1)K + k, n) \log \frac{\mathbf{P}^{\natural}((m-1)K + k, n)}{\hat{\mathbf{P}}((m-1)K + k, n)} \\ &= D_{\text{KL}}(\mathbf{P}^{\natural}, \hat{\mathbf{P}}), \end{aligned} \quad (41)$$

where expectation is taken w.r.t. $\widehat{\mathcal{Y}}$ (while taking expectation, we used the uniform probability $\pi_n^{(m)} = \frac{1}{NM}, \forall, n, m$ for observing each annotation $\widehat{y}_n^{(m)}$ and used the ground-truth probability $\mathbf{P}^{\natural}((m-1)K+k, n)$ for each event $\mathbb{1}[\widehat{y}_n^{(m)} = k]$) and $D_{\text{KL}}(\mathbf{P}^{\natural}, \widehat{\mathbf{P}})$ is the average Kullback–Leibler (KL) divergence between the entries of the matrices \mathbf{P}^{\natural} and $\widehat{\mathbf{P}} \in \mathbb{R}^{MK \times N}$, which is given by

$$D_{\text{KL}}(\mathbf{P}^{\natural}, \widehat{\mathbf{P}}) = \frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M D_{\text{KL}}(\mathbf{p}_n^{(m)}, \widehat{\mathbf{p}}_n^{(m)}).$$

Upper-bounding the first term on the R.H.S of (40). Next, we characterize the first term on the R.H.S. of (40). To achieve this, we invoke the following theorem (Theorem 26.5 in (Shalev-Shwartz & Ben-David, 2014)) :

Theorem 4 (Shalev-Shwartz & Ben-David, 2014, Theorem 26.5) *Assume that for all y and for all x , we have $|\text{CE}(x; y)| \leq z_{\max}$. Then for any $\mathcal{P} \in \mathcal{P}$, the following holds with probability greater than $1 - \delta$:*

$$\left| D_{\mathcal{S}}(\mathcal{P}; \widehat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathcal{P}; \widehat{\mathcal{Y}})] \right| \leq 2\mathfrak{R}_{\mathcal{S}}(\ell \circ \mathcal{P} \circ \mathcal{S}) + 4z_{\max} \sqrt{\frac{2 \log(4/\delta)}{S}}, \quad (42)$$

where $\ell \circ \mathcal{P} \circ \mathcal{S}$ denotes the set

$$\ell \circ \mathcal{P} \circ \mathcal{S} \triangleq \left\{ \left(\text{CE}(\mathbf{P}(\omega_1); \widehat{\mathcal{Y}}(\omega_1)), \dots, \text{CE}(\mathbf{P}(\omega_S); \widehat{\mathcal{Y}}(\omega_S)) \right) \mid \mathbf{P} \in \mathcal{P} \right\}$$

and $\mathfrak{R}_{\mathcal{S}}(\mathcal{X})$ denotes the empirical Rademacher complexity of the set \mathcal{X} .

To apply Theorem 4, we will characterize $\mathfrak{R}_{\mathcal{S}}(\ell \circ \mathcal{P} \circ \mathcal{S})$ which is defined as follows (Shalev-Shwartz & Ben-David, 2014):

$$\mathfrak{R}_{\mathcal{S}}(\ell \circ \mathcal{P} \circ \mathcal{S}) \triangleq \frac{1}{S} \mathbb{E} \left[\sup_{\mathbf{P} \in \mathcal{P}} \sum_{s=1}^S \sigma_s \text{CE}(\mathbf{P}(\omega_s); \widehat{\mathcal{Y}}(\omega_s)) \right], \quad (43)$$

where expectation is w.r.t. the independent Rademacher random variables $\sigma_s \in \{-1, 1\}$. Note that $\mathbf{P}(\omega)$ is a vector-valued [see Eq. (38)]. Hence, we invoke the following contraction result to upper bound $\mathfrak{R}_{\mathcal{S}}(\ell \circ \mathcal{P} \circ \mathcal{S})$:

Lemma 4 (Maurer, 2016) *Let \mathcal{P} be a class of mappings $\{\mathbf{P} : \mathcal{X} \rightarrow \mathbb{R}^K\}$, where \mathcal{X} be any set and $(\omega_1, \dots, \omega_S) \in \mathcal{X}^S$. Also assume that $\ell : \mathbb{R}^K \rightarrow \mathbb{R}$ has the Lipschitz constant L . Then*

$$\mathbb{E} \left[\sup_{\mathbf{P} \in \mathcal{P}} \sum_{s=1}^S \sigma_s \ell(\mathbf{P}(\omega_s)) \right] \leq \sqrt{2} L \mathbb{E} \left[\sup_{\mathbf{P} \in \mathcal{P}} \sum_{s,k} \sigma_{s,k} \mathbf{P}_k(\omega_s) \right]$$

where $\mathbf{P}_k(w)$ denotes the k th component of $\mathbf{P}(w)$, $\sigma_{s,k}$ is an independent (doubly indexed) Rademacher random variable and the expectations are taken w.r.t. the Rademacher random variables.

Let us define a vector $\mathbf{z} \triangleq (\mathbf{P}_1(\omega_1), \mathbf{P}_2(\omega_1), \dots, \mathbf{P}_{K-1}(\omega_S), \mathbf{P}_K(\omega_S)) \in \mathbb{R}^{SK}$ and the set $\mathcal{Z} \triangleq \{\mathbf{z} = (\mathbf{P}_1(\omega_1), \dots, \mathbf{P}_K(\omega_S)) \mid \mathbf{P} \in \mathcal{P}\}$. With these definitions, we apply Lemma 4 in (43) and obtain

$$\begin{aligned} \mathfrak{R}_{\mathcal{S}}(\ell \circ \mathcal{P} \circ \mathcal{S}) &\leq \frac{\sqrt{2}\beta}{S} \mathbb{E} \left[\sup_{\mathbf{P} \in \mathcal{P}} \sum_{s,k} \sigma_{s,k} \mathbf{P}_k(\omega_s) \right] \\ &= \frac{\sqrt{2}\beta}{S} \mathbb{E} \left[\sup_{\mathbf{z} \in \mathcal{Z}} \sum_i \sigma_i \mathbf{z}(i) \right] \\ &= \sqrt{2}\beta K \frac{1}{SK} \mathbb{E} \left[\sup_{\mathbf{z} \in \mathcal{Z}} \sum_i \sigma_i \mathbf{z}(i) \right] = \sqrt{2}\beta K \mathfrak{R}_{\mathcal{S}}(\mathcal{Z}), \end{aligned} \quad (44)$$

where β is an upper bound of the Lipschitz constant of the cross entropy loss function $\text{CE}(x; y) = -\sum_{k=1}^K \mathbb{1}[y = k] \log x(k)$ when $x \in \Delta_K$ with $x(k) > (1/\beta), \forall k$.

Next, we will characterize $\mathfrak{R}_{\mathcal{S}}(\mathcal{Z})$ using the covering number of the set \mathcal{Z} .

Definition 2 (Vershynin, 2012) *The ϵ -net covering of the set \mathcal{Z} (denoted as $\bar{\mathcal{Z}}$) is a finite subset of \mathcal{Z} (i.e., $\bar{\mathcal{Z}} \subseteq \mathcal{Z}$) such that for any $\mathbf{z} \in \mathcal{Z}$, there exists an $\bar{\mathbf{z}} \in \bar{\mathcal{Z}}$ satisfying $\|\mathbf{z} - \bar{\mathbf{z}}\|_2 \leq \epsilon$. The smallest cardinality of the ϵ -nets of $\bar{\mathcal{Z}}$ is known as the covering number of \mathcal{Z} , which is denoted as $\bar{N}(\epsilon, \mathcal{Z})$.*

Let us consider a pair of vectors $\mathbf{z}, \bar{\mathbf{z}} \in \mathcal{Z}$ as below:

$$\begin{aligned}\mathbf{z} &= (\mathbf{P}_1(\omega_1), \dots, \mathbf{P}_K(\omega_s)) \in \mathbb{R}^{SK}, \quad \mathbf{P}(\omega_s) = \mathbf{A}_{m_s} \mathbf{F}(:, n_s), \\ \bar{\mathbf{z}} &= (\bar{\mathbf{P}}_1(\omega_1), \dots, \bar{\mathbf{P}}_K(\omega_s)) \in \mathbb{R}^{SK}, \quad \bar{\mathbf{P}}(\omega_s) = \bar{\mathbf{A}}_{m_s} \bar{\mathbf{F}}(:, n_s),\end{aligned}$$

where $\omega_s = (m_s, n_s)$ and all $\mathbf{A}_m, \bar{\mathbf{A}}_m, \mathbf{F}$, and $\bar{\mathbf{F}}$ satisfy the nonnegativity constraints and have unit ℓ_1 -norm on the columns. Then, we have

$$\begin{aligned}\|\mathbf{z} - \bar{\mathbf{z}}\|^2 &= \sum_{s=1}^S \|\mathbf{A}_{m_s} \mathbf{F}(:, n_s) - \bar{\mathbf{A}}_{m_s} \bar{\mathbf{F}}(:, n_s)\|_2^2 \\ &\leq \sum_{s=1}^S \|\mathbf{A}_{m_s} \mathbf{F}(:, n_s) - \bar{\mathbf{A}}_{m_s} \bar{\mathbf{F}}(:, n_s)\|_2 \\ &= \sum_{s=1}^S \|\mathbf{A}_{m_s} \mathbf{F}(:, n_s) - \bar{\mathbf{A}}_{m_s} \mathbf{F}(:, n_s) + \bar{\mathbf{A}}_{m_s} \mathbf{F}(:, n_s) - \bar{\mathbf{A}}_{m_s} \bar{\mathbf{F}}(:, n_s)\|_2 \\ &\leq \sum_{s=1}^S \|\mathbf{A}_{m_s} - \bar{\mathbf{A}}_{m_s}\|_{\text{F}} \|\mathbf{F}(:, n_s)\|_2 + \|\bar{\mathbf{A}}_{m_s}\|_{\text{F}} \|\mathbf{F}(:, n_s) - \bar{\mathbf{F}}(:, n_s)\|_2 \\ &\leq \sum_{s=1}^S \|\mathbf{A}_{m_s} - \bar{\mathbf{A}}_{m_s}\|_{\text{F}} + \sqrt{K} \sum_{s=1}^S \|\mathbf{F}(:, n_s) - \bar{\mathbf{F}}(:, n_s)\|_2\end{aligned}$$

where the first inequality is by $\|\mathbf{x}\|_2^2 \leq \|\mathbf{x}\|_2$ if the entries of \mathbf{x} are smaller than 1. The second inequality is by triangle inequality. The last inequality uses the fact the Frobenius norm of $\bar{\mathbf{A}}_m$'s are bounded by \sqrt{K} and the ℓ_2 norm of any column of \mathbf{F} is bounded by 1. Hence, to obtain an ϵ -net covering for the set \mathcal{Z} (i.e., $\|\mathbf{z} - \bar{\mathbf{z}}\| \leq \epsilon$), we only need to show that there exists a $\frac{\epsilon^2}{2\sqrt{K}}$ -net covering for $\mathcal{F} \circ \mathcal{S}$ and a $\frac{\epsilon^2}{2S}$ -net covering for each \mathbf{A}_m 's since

$$\sum_{s=1}^S \frac{\epsilon^2}{2S} + \sqrt{K} \frac{\epsilon^2}{2\sqrt{K}} = \epsilon^2.$$

Here $\mathcal{F} \circ \mathcal{S}$ denotes

$$\mathcal{F} \circ \mathcal{S} = \{[\mathbf{f}(\mathbf{x}_{n_1}), \dots, \mathbf{f}(\mathbf{x}_{n_S})] \in \mathbb{R}^{K \times S} \mid \mathbf{f} \in \mathcal{F}\},$$

where $(m_s, n_s) = \omega_s \in \mathcal{S}$. Note that the full rank matrix $\mathbf{A}_m \in \mathbb{R}^{K \times K}$ can be represented as a K^2 -dimensional vector whose Euclidean norm is bounded by \sqrt{K} . Hence, the cardinality of the $\frac{\epsilon^2}{2S}$ -net covering for $\mathbf{A}_m \in \mathbb{R}^{K \times K}$ is at most $\left(\frac{4SK\sqrt{K}}{\epsilon^2}\right)^{K^2}$ (Shalev-Shwartz & Ben-David, 2014). Next, we consider the covering number corresponding to the function class $\mathcal{F} \circ \mathcal{S}$. Using Lemma 14 of (Lin & Zhang, 2019), we get the cardinality of the $\frac{\epsilon^2}{2\sqrt{K}}$ -net covering for $\mathcal{F} \circ \mathcal{S}$ as below:

$$\bar{N}\left(\frac{\epsilon^2}{2\sqrt{K}}, \mathcal{F} \circ \mathcal{S}\right) \leq \exp\left(\sqrt{\frac{2\sqrt{K}\|\mathbf{X}\|_{\text{F}}\mathcal{R}_{\mathcal{F}}}{\epsilon^2}}\right),$$

where $\mathbf{X} = [\mathbf{x}_{n_1}, \dots, \mathbf{x}_{n_S}] \in \mathbb{R}^{d \times S}$ and the parameter $\mathcal{R}_{\mathcal{F}}$ is from Assumption 2.

Using the covering number results, the cardinality of the ϵ -net covering of set \mathcal{Z} is bounded by the following:

$$\bar{N}(\epsilon, \mathcal{Z}) \leq \left(\frac{4SK\sqrt{K}}{\epsilon^2}\right)^{MK^2} \times \exp\left(\sqrt{\frac{2\sqrt{K}\|\mathbf{X}\|_{\text{F}}\mathcal{R}_{\mathcal{F}}}{\epsilon^2}}\right). \quad (45)$$

Now that we have characterized $\bar{N}(\epsilon, \mathcal{Z})$, we invoke the below lemma to obtain the Rademacher complexity $\mathfrak{R}_{\mathcal{S}}(\mathcal{Z})$:

Lemma 5 (Bartlett et al., 2017, Lemma A.5) *The empirical Rademacher complexity of the set \mathcal{Z} with respect to the observed set \mathcal{S} having size SK is upper bounded as follows:*

$$\mathfrak{R}_{\mathcal{S}}(\mathcal{Z}) \leq \inf_{a>0} \left(\frac{4a}{\sqrt{SK}} + \frac{12}{SK} \int_a^{\sqrt{SK}} \sqrt{\log \bar{N}(\mu, \mathcal{Z})} d\mu \right). \quad (46)$$

We apply (45) in Lemma 5 and obtain

$$\begin{aligned} \mathfrak{R}_{\mathcal{S}}(\mathcal{Z}) &\stackrel{(a)}{\leq} \inf_{a>0} \left(\frac{4a}{\sqrt{SK}} + \frac{12}{SK} \sqrt{SK} \sqrt{\log \bar{N}(a, \mathcal{Z})} \right) \\ &\stackrel{(b)}{\leq} \inf_{a>0} \left(\frac{4a}{\sqrt{SK}} + \frac{12}{\sqrt{SK}} \sqrt{MK^2 \log \left(\frac{4SK\sqrt{K}}{a^2} \right) + \left(\frac{2\sqrt{K} \|\mathbf{X}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{F}}}{a^2} \right)^{\frac{1}{2}}} \right) \\ &\stackrel{(c)}{\leq} \frac{4}{\sqrt{S}} + \frac{12}{\sqrt{SK}} \sqrt{MK^2 \log(4S\sqrt{K}) + (2\|\mathbf{X}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{F}})^{\frac{1}{2}}} \\ &\leq \frac{16}{\sqrt{S}} \sqrt{MK \log(4S\sqrt{K}) + (2\|\mathbf{X}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{F}})^{\frac{1}{2}}}. \end{aligned} \quad (47)$$

In the above, the first inequality (a) is obtained by using the relation

$$\int_a^{\sqrt{SK}} \sqrt{\log \bar{N}(\mu, \mathcal{Z})} d\mu \leq \sqrt{SK} \sqrt{\log \bar{N}(a, \mathcal{Z})},$$

which holds because $\sqrt{\log \bar{N}(\mu, \mathcal{Z})}$ decreases monotonically as μ increases. The inequality (b) is obtained by applying (45), and (c) is obtained by fixing $a = \sqrt{K}$ which is smaller than \sqrt{SK} .

Combining the upperbound of $\mathfrak{R}_{\mathcal{S}}(\mathcal{Z})$ given by (47) with the upper bound of $\mathfrak{R}_{\mathcal{S}}(\ell \circ \mathcal{P} \circ \mathcal{S})$ as given by (44) and with the result in (42), we get that with probability greater than $1 - \delta$,

$$\left| D_{\mathcal{S}}(\mathbf{P}; \hat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathbf{P}; \hat{\mathcal{Y}})] \right| \leq 2\sqrt{2}\beta K \mathfrak{R}_{\mathcal{S}}(\mathcal{Z}) + 4z_{\max} \sqrt{\frac{2 \log(4/\delta)}{S}}, \quad (48)$$

where $\mathfrak{R}_{\mathcal{S}}(\mathcal{Z})$ is upper bounded by (47) and z_{\max} is the upperbound of the value of the function $\text{CE}(\mathbf{x}, y)$ which can be characterized as below:

$$z_{\max} = \max_{\substack{\mathbf{x}(k) > \frac{1}{\beta} \\ y \in [K]}} \text{CE}(\mathbf{x}, y) \leq \max_{\substack{\mathbf{x}(k) > \frac{1}{\beta} \\ y \in [K]}} - \sum_{k=1}^K \mathbb{I}[y = k] \log \mathbf{x}(k) \leq \max_{u > \frac{1}{\beta}} - \log u = \log(\beta). \quad (49)$$

Upper-bounding the second term on the R.H.S of (40). Next, we proceed to upper bound the second term on the R.H.S of (40). Let us consider the Hoeffding's inequality

Lemma 6 *Let Z_1, \dots, Z_S be independent bounded random variables with $Z_s \in [z_{\min}, z_{\max}]$ for all s where $-\infty < z_{\min} \leq z_{\max} < \infty$. Then for all $t \geq 0$,*

$$\Pr \left(\frac{1}{S} \sum_{s=1}^S (Z_s - \mathbb{E}[Z_s]) \geq t \right) \leq \exp \left(-\frac{2St^2}{(z_{\max} - z_{\min})^2} \right).$$

To use Lemma 6, let us define the random variable $Z_n^{(m)}$ as follows:

$$Z_n^{(m)} \triangleq \text{CE}(\mathbf{p}_n^{\natural(m)}, \hat{y}_n^{(m)}),$$

where $\mathbf{p}_n^{\natural(m)} = \mathbf{A}_m^{\natural} \mathbf{f}^{\natural}(\mathbf{x}_n)$. The maximum and minimum values of $Z_n^{(m)}$ are $z_{\max} = \log(\beta)$ (see (49)) and $z_{\min} = 0$, respectively. Then, invoking Lemma 6, one can obtain

$$\Pr \left(D_{\mathcal{S}}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}}) - \mathbb{E}[D_{\Pi}(\mathbf{P}^{\natural}; \hat{\mathcal{Y}})] \geq t \right) \leq \exp \left(-\frac{2St^2}{(\log(\beta))^2} \right). \quad (50)$$

Hence, by substituting $t = \log(\beta) \sqrt{\frac{\log(\frac{1}{\delta})}{2S}}$, where $\delta \in (0, 1)$ in (50), we get that with probability greater than $1 - \delta$

$$D_S(\mathbf{P}^\natural; \widehat{\mathcal{Y}}) - \mathbb{E}[D_\Pi(\mathbf{P}^\natural; \widehat{\mathcal{Y}})] \leq \log(\beta) \sqrt{\frac{\log(\frac{1}{\delta})}{2S}}. \quad (51)$$

Upper-bounding the third term on the R.H.S of (40).

$$\begin{aligned} \left| D_S(\widetilde{\mathbf{P}}; \widehat{\mathcal{Y}}) - D_S(\mathbf{P}^\natural; \widehat{\mathcal{Y}}) \right| &= \left| -\frac{1}{S} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\widehat{y}_n^{(m)} = k] \log \widetilde{\mathbf{P}}((m-1)K + k, n) \right. \\ &\quad \left. + \frac{1}{S} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\widehat{y}_n^{(m)} = k] \log \mathbf{P}^\natural((m-1)K + k, n) \right| \\ &\leq \frac{1}{S} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\widehat{y}_n^{(m)} = k] \left| \log[\mathbf{A}_m^\natural \mathbf{f}^\natural(\mathbf{x}_n)]_k - \log[\mathbf{A}_m^\natural \widetilde{\mathbf{f}}(\mathbf{x}_n)]_k \right| \\ &\leq \frac{1}{S} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\widehat{y}_n^{(m)} = k] \beta \left| [\mathbf{A}_m^\natural \mathbf{f}^\natural(\mathbf{x}_n)]_k - [\mathbf{A}_m^\natural \widetilde{\mathbf{f}}(\mathbf{x}_n)]_k \right| \\ &\leq \frac{1}{S} \sum_{(m,n) \in \mathcal{S}} \sum_{k=1}^K \mathbb{I}[\widehat{y}_n^{(m)} = k] \beta \|\mathbf{A}_m^\natural(k, \cdot)\|_2 \|\mathbf{f}^\natural(\mathbf{x}_n) - \widetilde{\mathbf{f}}(\mathbf{x}_n)\|_2 \\ &\leq \beta \sqrt{K} \nu, \end{aligned} \quad (52)$$

where the first inequality uses the triangle inequality, the second inequality uses the Lipschitz continuity of log function, the third inequality is via Cauchy Schwartz inequality, and the last inequality employs Assumption 3.

Putting Together. Hence, by combining the result in (48) with (41), (40), (51), and (52), we get that with probability greater than $1 - 2\delta$,

$$D_{\text{KL}}(\mathbf{P}^\natural, \widehat{\mathbf{P}}) \leq 2\sqrt{2}\beta K \mathfrak{R}_S(\mathcal{Z}) + 4 \log(\beta) \sqrt{\frac{2 \log(4/\delta)}{S}} + \log(\beta) \sqrt{\frac{\log(\frac{1}{\delta})}{2S}} + \beta \sqrt{K} \nu. \quad (53)$$

Using Pinsker's inequality (S, 1960; Fedotov et al., 2003), we get

$$\begin{aligned} D_{\text{KL}}(\mathbf{P}^\natural, \widehat{\mathbf{P}}) &= \frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M D_{\text{KL}}(\mathbf{p}_n^{\natural(m)}, \widehat{\mathbf{p}}_n^{(m)}) \geq \frac{1}{2NM} \sum_{n=1}^N \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_1^2 \\ &\geq \frac{1}{2NM} \sum_{n=1}^N \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \end{aligned}$$

where the last inequality uses the fact that $\|\mathbf{x}\|_1 \geq \|\mathbf{x}\|_2$. The above relation combined with (53), implies that with probability greater than $1 - 2\delta$:

$$\frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \leq 4\sqrt{2}K\beta \mathfrak{R}_S(\mathcal{Z}) + 10 \log(\beta) \frac{\sqrt{2 \log(4/\delta)}}{\sqrt{S}} + 2\beta \sqrt{K} \nu, \quad (54)$$

where $\mathfrak{R}_S(\mathcal{Z})$ is upper-bounded in (47)

D PROOF OF THEOREM 2

The proof of Theorem 2 utilizes some key results from the proof of Theorem 1. We start by employing the result from Proposition 1. Let us fix $\delta = \frac{1}{8}$ and $\alpha = \frac{1}{8}$ in the result of Proposition 1. Then, it

implies that when $\mathbf{f}^\natural \in \mathcal{F}$, i.e., $\nu = 0$, $S \geq C_1 t$, $N \leq C_2 t^3$, for certain constants $C_1, C_2 > 0$ and at the limit of $t \rightarrow \infty$, we get the following:

$$\|\mathbf{P}^\natural - \widehat{\mathbf{P}}\|_{\mathbb{F}}^2 = 0,$$

with probability 1, i.e.,

$$\widehat{\mathbf{P}} = \mathbf{W}^\natural \mathbf{F}^\natural. \quad (55)$$

On the other hand, the matrix $\widehat{\mathbf{P}}$ can be constructed using the estimates of the CCEM criterion (6) via

$$\widehat{\mathbf{P}} = \widehat{\mathbf{W}} \widehat{\mathbf{F}}. \quad (56)$$

Hence, in order to identify \mathbf{W}^\natural and \mathbf{F}^\natural from the NMF model in (55), we invoke the following result:

Lemma 7 (Huang et al., 2014) Consider the matrix factorization model $\mathbf{Z} = \mathbf{X}\mathbf{Y}$, where $\mathbf{X} \in \mathbb{R}^{I \times K}$, $\mathbf{Y} \in \mathbb{R}^{K \times J}$ and $\text{rank}(\mathbf{X}) = \text{rank}(\mathbf{Y}) = K$. If $\mathbf{X}, \mathbf{Y} \geq \mathbf{0}$ and both \mathbf{X} and \mathbf{Y} satisfy SSC, then any $\widehat{\mathbf{X}}$ and $\widehat{\mathbf{Y}}$ that satisfy $\mathbf{Z} = \widehat{\mathbf{X}}\widehat{\mathbf{Y}}$ must have the following form:

$$\widehat{\mathbf{X}} = \mathbf{X}\mathbf{\Pi}\mathbf{\Sigma}, \quad \widehat{\mathbf{Y}} = \mathbf{\Sigma}^{-1}\mathbf{\Pi}^\top\mathbf{Y},$$

where $\mathbf{\Pi}$ is a column permutation matrix and $\mathbf{\Sigma}$ is a diagonal nonnegative and scaling matrix.

Next, we show that Lemma 7 can be applied given the conditions in Theorem 2. The matrix $\mathbf{F}_Z^\natural = [\mathbf{f}^\natural(\mathbf{x}_{n_1}), \dots, \mathbf{f}^\natural(\mathbf{x}_{n_T})]$ includes a subset of columns of \mathbf{F}^\natural , i.e., $\text{cone}(\mathbf{F}_Z^\natural) \subseteq \text{cone}(\mathbf{F}^\natural)$. Hence, we have

$$\mathcal{C} \subseteq \text{cone}(\mathbf{F}_Z^\natural) \implies \mathcal{C} \subseteq \text{cone}(\mathbf{F}^\natural), \quad (57)$$

where $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^K \mid \sqrt{K-1}\|\mathbf{x}\|_2 \leq \mathbf{1}^\top \mathbf{x}\}$ is the second-order cone.

Also for any orthogonal matrix $\mathbf{Q} \in \mathbb{R}^{K \times K}$ except for the permutation matrices, the below holds:

$$\text{cone}(\mathbf{F}_Z^\natural) \not\subseteq \text{cone}\{\mathbf{Q}\} \implies \text{cone}(\mathbf{F}^\natural) \not\subseteq \text{cone}\{\mathbf{Q}\}. \quad (58)$$

Eqs (57) and (58) imply that \mathbf{F}^\natural satisfies SSC, given the assumption that \mathbf{F}_Z^\natural satisfies SSC. In addition, the column scaling does not affect the conic hull, i.e.,

$$\begin{aligned} \text{cone}(\mathbf{F}^\natural) &= \text{cone}(\mathbf{\Sigma}^{-1}\mathbf{F}^\natural), \\ \text{cone}(\mathbf{W}^{\natural\top}) &= \text{cone}(\mathbf{\Sigma}\mathbf{W}^{\natural\top}). \end{aligned}$$

Since both \mathbf{W}^\natural and \mathbf{F}^\natural satisfy SSC, $\text{rank}(\mathbf{F}^\natural) = \text{rank}(\mathbf{W}^\natural) = K$ hold (see (Huang et al., 2016)). Hence, the conditions in Lemma 7 holds for (55). Comparing (55) and (56) and invoking Lemma 7, we get

$$\widehat{\mathbf{W}} = \mathbf{W}^\natural \mathbf{\Pi}, \quad (59a)$$

$$\widehat{\mathbf{F}} = \mathbf{\Pi}^\top \mathbf{F}^\natural, \quad (59b)$$

where the scaling $\mathbf{\Sigma}$ is automatically removed due to the sum-to-one constraints on the columns of $\widehat{\mathbf{A}}_m$'s and $\widehat{\mathbf{F}}$.

The first result (59a) implies that

$$\widehat{\mathbf{A}}_m = \mathbf{A}_m^\natural \mathbf{\Pi}, \forall m.$$

Applying the second result (59b) in Lemma 3 along with the assumption that $N \rightarrow \infty$, we get

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\|\widehat{\mathbf{f}}(\mathbf{x}) - \mathbf{\Pi}^\top \mathbf{f}^\natural(\mathbf{x})\|_2^2 \right] = 0. \quad (60)$$

Fact 1 Let X be a nonnegative random variable with $\mathbb{E}[X] = 0$. Then, X is zero almost surely, i.e.,

$$\Pr(X = 0) = 1.$$

Employing Fact 1 in (60), we get that

$$\|\widehat{\mathbf{f}}(\mathbf{x}) - \mathbf{\Pi}^\top \mathbf{f}^\natural(\mathbf{x})\|_2^2 = 0, \forall \mathbf{x} \sim \mathcal{D}, \implies \widehat{\mathbf{f}}(\mathbf{x}) = \mathbf{\Pi}^\top \mathbf{f}^\natural(\mathbf{x}), \forall \mathbf{x} \sim \mathcal{D},$$

due to the nonnegativity of \mathbf{f}^\natural and $\widehat{\mathbf{f}}$.

E PROOF OF THEOREM 3

To prove Theorem 3, let us first consider the CCEM criterion in (6). Let $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{F}}$ be any optimal solution of (6) and

$$\widehat{\mathbf{P}} = \widehat{\mathbf{W}}\widehat{\mathbf{F}}. \quad (61)$$

From Proposition 1, by fixing $\delta = \frac{1}{S}$ and $\alpha = \frac{1}{8}$ and using the conditions in the statement of Theorem 3, i.e., $\mathbf{f}^\natural \in \mathcal{F}$ implying $\nu = 0$, $S \geq C_1 t$, $N \leq C_2 t^3$, for certain constants $C_1, C_2 > 0$, we get the following at the limit of $t \rightarrow \infty$:

$$\|\mathbf{P}^\natural - \widehat{\mathbf{P}}\|_{\text{F}}^2 = 0,$$

with probability 1, i.e.,

$$\widehat{\mathbf{P}} = \mathbf{W}^\natural \mathbf{F}^\natural. \quad (62)$$

Proof of Theorem 3(a): We consider the following result which is distilled and summarized from the proof of Theorem 1 in (Fu et al., 2015):

Lemma 8 *Suppose a matrix $\mathbf{Y} \in \mathbb{R}^{K \times J}$ satisfies $\mathbf{Y} \geq \mathbf{0}$, $\mathbf{1}^\top \mathbf{Y} = \mathbf{1}^\top$, $\text{rank}(\mathbf{Y}) = K$, and SSC. Then, for any $\widehat{\mathbf{Y}} = \mathbf{Q}\mathbf{Y}$ satisfying $\widehat{\mathbf{Y}} \geq \mathbf{0}$, $\mathbf{1}^\top \widehat{\mathbf{Y}} = \mathbf{1}^\top$, the following holds:*

$$|\det(\mathbf{Q})| \leq 1,$$

The equality holds only if \mathbf{Q} is a permutation matrix.

Let us start by considering the following criterion:

$$\underset{\mathbf{W}, \mathbf{F}}{\text{maximize}} \det(\mathbf{F}\mathbf{F}^\top) \quad (63a)$$

$$\text{s.t. } \widehat{\mathbf{P}} = \mathbf{W}\mathbf{F} \quad (63b)$$

$$\mathbf{F} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{F} = \mathbf{1}^\top \quad (63c)$$

where $\widehat{\mathbf{P}}$ is the optimal solution from the CCEM criterion (6) and hence, as shown before, $\widehat{\mathbf{P}}$ satisfies (62).

Let \mathbf{W}^* and \mathbf{F}^* be optimal solutions of (63), then the following holds:

$$\det(\mathbf{F}^* \mathbf{F}^{*\top}) \geq \det(\mathbf{F}^\natural \mathbf{F}^{\natural\top}). \quad (64)$$

From (62), we observe that \mathbf{W}^* and \mathbf{F}^* satisfies $\mathbf{W}^* = \mathbf{W}^\natural \mathbf{Q}^{-1}$, $\mathbf{F}^* = \mathbf{Q}\mathbf{F}^\natural$, for a certain invertible matrix \mathbf{Q} . Let us assume that \mathbf{Q} is not a permutation matrix, Then, we have

$$\begin{aligned} \det(\mathbf{F}^* \mathbf{F}^{*\top}) &= \det(\mathbf{Q}^\top \mathbf{F}^\natural \mathbf{F}^{\natural\top} \mathbf{Q}) \\ &= |\det(\mathbf{Q})|^2 \det(\mathbf{F}^\natural \mathbf{F}^{\natural\top}) \\ &< \det(\mathbf{F}^\natural \mathbf{F}^{\natural\top}) \end{aligned}$$

where the last inequality is from Lemma 8 using the SSC condition on \mathbf{F}^\natural . Note that the result is a contradiction from (64). Hence, \mathbf{Q} must be a permutation matrix. This implies that the optimal solution \mathbf{W}^* and \mathbf{F}^* of Problem 63 satisfies the following:

$$\mathbf{W}^* = \mathbf{W}^\natural \mathbf{\Pi}, \quad (65a)$$

$$\mathbf{F}^* = \mathbf{\Pi}^\top \mathbf{F}^\natural, \quad (65b)$$

Following the last part of Theorem 2, by using Lemma 3 and Fact 1, we have

$$\begin{aligned} \mathbf{A}_m^* &= \mathbf{A}_m^\natural \mathbf{\Pi}, \forall m. \\ \mathbf{f}^*(\mathbf{x}) &= \mathbf{\Pi}^\top \mathbf{f}^\natural(\mathbf{x}), \forall \mathbf{x} \sim \mathcal{D}. \end{aligned}$$

Note that since log is a monotonically increasing function, by using $\log \det(\mathbf{F}\mathbf{F}^\top)$ in (63) in place of $\det(\mathbf{F}\mathbf{F}^\top)$ does not change the optimal solution, yet it keeps the objective function in differentiable domain (this is because, $\det(\mathbf{X})$ becomes zero if \mathbf{X} is singular). Hence, we can conclude that the following criterion

$$\underset{\mathbf{W}, \mathbf{F}}{\text{maximize}} \log \det(\mathbf{F}\mathbf{F}^\top) \quad (66a)$$

$$\text{s.t. } \widehat{\mathbf{P}} = \mathbf{W}\mathbf{F} \quad (66b)$$

$$\mathbf{F} \geq \mathbf{0}, \quad (66c)$$

also results in the same optimal solutions \mathbf{W}^* and \mathbf{F}^* satisfying (65).

Proof of Theorem 3(b): We consider the following result which is a modified version of Lemma A.1 from (Huang et al., 2015):

Lemma 9 Suppose a matrix $\mathbf{X} \in \mathbb{R}^{I \times K}$ satisfies $\mathbf{X} \geq \mathbf{0}$, $\mathbf{1}^\top \mathbf{X} = \rho \mathbf{1}^\top$, $\rho > 0$, $\text{rank}(\mathbf{X}) = K$, and SSC. Then, for any $\widehat{\mathbf{X}} = \mathbf{X}\mathbf{Q}$ satisfying $\widehat{\mathbf{X}} \geq \mathbf{0}$, $\mathbf{1}^\top \widehat{\mathbf{X}} = \rho \mathbf{1}^\top$, the following holds:

$$|\det(\mathbf{Q})| \leq 1,$$

The equality holds only if \mathbf{Q} is a permutation matrix.

Note that, in Lemma 9, we are given that $\mathbf{1}^\top \mathbf{X} = \rho \mathbf{1}^\top$. Then, by multiplying \mathbf{Q} on both sides, we obtain

$$\mathbf{1}^\top \mathbf{X}\mathbf{Q} = \rho \mathbf{1}^\top \mathbf{Q}. \quad (67)$$

Since $\mathbf{1}^\top \mathbf{X}\mathbf{Q} = \rho \mathbf{1}^\top$ also holds by the assumption in Lemma 9, combining with (67), we get

$$\rho \mathbf{1}^\top \mathbf{Q} = \rho \mathbf{1}^\top \implies \mathbf{1}^\top \mathbf{Q} = \mathbf{1}^\top. \quad (68)$$

The result in (68) is used in order to conclude Lemma 9 from Lemma A.1 of (Huang et al., 2015).

We consider the following criterion:

$$\underset{\mathbf{W}, \mathbf{F}}{\text{maximize}} \det(\mathbf{W}^\top \mathbf{W}) \quad (69a)$$

$$\text{s.t. } \widehat{\mathbf{P}} = \mathbf{W}\mathbf{F} \quad (69b)$$

$$\mathbf{W} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{W} = \mathbf{M}\mathbf{1}^\top, \quad (69c)$$

where we used the constraint $\mathbf{1}^\top \mathbf{W} = \mathbf{M}\mathbf{1}^\top$ since $\mathbf{1}^\top \mathbf{A}_m = \mathbf{1}^\top$ for all m .

Let \mathbf{W}^* and \mathbf{F}^* be optimal solutions of (69), then we have:

$$\det(\mathbf{W}^{*\top} \mathbf{W}^*) \geq \det(\mathbf{W}^{\natural\top} \mathbf{W}^{\natural}). \quad (70)$$

Applying (62), \mathbf{W}^* and \mathbf{F}^* satisfies $\mathbf{W}^* = \mathbf{W}^{\natural}\mathbf{Q}$, $\mathbf{F}^* = \mathbf{Q}^{-1}\mathbf{F}^{\natural}$, for a certain invertible matrix \mathbf{Q} . Let us assume that \mathbf{Q} is not a permutation matrix, Then, we have

$$\begin{aligned} \det(\mathbf{W}^{*\top} \mathbf{W}^*) &= \det(\mathbf{Q}^\top \mathbf{W}^{\natural\top} \mathbf{W}^{\natural} \mathbf{Q}) \\ &= |\det(\mathbf{Q})|^2 \det(\mathbf{W}^{\natural\top} \mathbf{W}^{\natural}) \\ &< \det(\mathbf{W}^{\natural\top} \mathbf{W}^{\natural}) \end{aligned}$$

where the last inequality is from Lemma 9 using the SSC condition on \mathbf{W}^{\natural} which is a contradiction from (70). Hence, \mathbf{Q} must be a permutation matrix. This implies that the optimal solution \mathbf{W}^* and \mathbf{F}^* of Problem 69 satisfies the following:

$$\mathbf{W}^* = \mathbf{W}^{\natural}\mathbf{\Pi}, \quad (71a)$$

$$\mathbf{F}^* = \mathbf{\Pi}^\top \mathbf{F}^{\natural}, \quad (71b)$$

Similar to the last part of Theorem 2, by using Lemma 3 and Fact 1, we have

$$\mathbf{A}_m^* = \mathbf{A}_m^{\natural}\mathbf{\Pi}, \forall m.$$

$$\mathbf{f}^*(\mathbf{x}) = \mathbf{\Pi}^\top \mathbf{f}^{\natural}(\mathbf{x}), \forall \mathbf{x} \sim \mathcal{D}.$$

In this case as well, the optimal solutions do not change if we employ $\log \det(\mathbf{W}^\top \mathbf{W})$ in (69) in place of $\det(\mathbf{W}^\top \mathbf{W})$.

Remark 2 Geometrically, Theorem 3 seeks maximum volume solutions w.r.t. the conic hulls of \mathbf{F} and \mathbf{W}^\top , respectively. One can also note that in both cases of Theorem 3, the corresponding optimal solutions do not change if we minimize $\log \det(\mathbf{W}^\top \mathbf{W})$ in (63) and minimize $\log \det(\mathbf{F} \mathbf{F}^\top)$ in (69). However, relying on the SSC of \mathbf{F}^\natural (\mathbf{W}^\natural) and minimizing the volume of conic hull of \mathbf{W}^\top (\mathbf{F}) may be inefficient since $\mathbf{P}^\natural = \mathbf{W}^\natural \mathbf{F}^\natural$ does not hold exactly in practice.

F PROOF OF LEMMA 1

Let us define some notations:

$$\begin{aligned}\ell_k^* &= \arg \max_{\ell \in [K]} \widehat{\mathbf{A}}_1(k, \ell) \\ q_k^* &= \arg \max_{q \in [K]} \widehat{\mathbf{F}}_1(q, k).\end{aligned}$$

Consider the following conditions given in Lemma 1

$$1 - \kappa \leq |\widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, k)| \leq 1 + \kappa, \forall k \in [K] \quad (72a)$$

$$|\widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, j)| \leq \kappa, \forall k \neq j. \quad (72b)$$

F.1 PROVING $\ell_k^* \neq q_j^*, \forall k \neq j$

We begin by noticing the below relation from (72a):

$$\begin{aligned}(1 - \kappa) &\leq \widehat{\mathbf{A}}_1(k, 1) \widehat{\mathbf{F}}_1(1, k) + \cdots + \widehat{\mathbf{A}}_1(k, K) \widehat{\mathbf{F}}_1(K, k) \\ &\leq \max_{\ell \in [K]} \widehat{\mathbf{A}}_1(k, \ell) (\widehat{\mathbf{F}}_1(1, k) + \cdots + \widehat{\mathbf{F}}_1(K, k)) \\ &= \max_{\ell \in [K]} \widehat{\mathbf{A}}_1(k, \ell) \\ &= \widehat{\mathbf{A}}_1(k, \ell_k^*), \forall k,\end{aligned} \quad (73)$$

where the first equality employs the probability simplex constraints on the columns of $\widehat{\mathbf{F}}$.

We further proceed to prove by using contradiction. First, assume the below for certain $k \neq j$,

$$\ell_k^* = \arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) = \arg \max_q \widehat{\mathbf{F}}_1(q, j) = q_j^*. \quad (74)$$

Under the assumption (74), consider (72b). For (72b) to hold, we need to satisfy the below for a certain $k \neq j$:

$$\begin{aligned}\widehat{\mathbf{A}}_1(k, \ell_k^*) \widehat{\mathbf{F}}_1(q_j^*, j) &< \kappa \\ \implies \widehat{\mathbf{F}}_1(q_j^*, j) &< \frac{\kappa}{1 - \kappa}\end{aligned}$$

where the last relation is obtained from (73). This also implies that

$$\begin{aligned}1 = \sum_{q=1}^K \widehat{\mathbf{F}}_1(q, k) &< \frac{K\kappa}{1 - \kappa} \\ \implies \kappa &> \frac{1}{K + 1}.\end{aligned}$$

However, Lemma 1 assumes that $\kappa \leq \frac{1}{K+1}$. Hence, the assumption (74) does not hold and we get

$$\arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) \neq \arg \max_q \widehat{\mathbf{F}}_1(q, j), k \neq j. \quad (75)$$

F.2 PROVING $\ell_k^* = q_k^*, \forall k$

From the result in (75), we consider (72b) for a certain $k \neq j$.

$$\begin{aligned}\widehat{\mathbf{A}}_1(k, \ell_k^*) \widehat{\mathbf{F}}_1(q, j) &\leq \kappa \\ \implies \widehat{\mathbf{F}}_1(q, j) &\leq \frac{\kappa}{1 - \kappa}, q \neq q_j^*.\end{aligned} \quad (76)$$

where the last relation is obtained from (73).

Since $\sum_q \widehat{\mathbf{F}}_1(q, j) = 1$, we get

$$\widehat{\mathbf{F}}_1(q_j^*, j) \geq \frac{1 - K\kappa}{1 - \kappa}, \forall j. \quad (77)$$

From the result in (75) and the condition in (72b), we also have the following for a certain $k \neq j$

$$\begin{aligned} \widehat{\mathbf{A}}_1(k, \ell) \widehat{\mathbf{F}}_1(q_j^*, j) &\leq \kappa \\ \implies \widehat{\mathbf{A}}_1(k, \ell) &\leq \frac{\kappa(1 - \kappa)}{1 - K\kappa}, \ell \neq \ell_k^*, \end{aligned} \quad (78)$$

where the last relation by (77).

Next, we proceed to prove by contradiction. First, assume the below for certain k

$$\ell_k^* \neq q_k^*. \quad (79)$$

Hence, for each k , we have

$$\begin{aligned} |\widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, k)| &= \widehat{\mathbf{A}}_1(k, 1) \widehat{\mathbf{F}}_1(1, k) + \dots + \widehat{\mathbf{A}}_1(k, K) \widehat{\mathbf{F}}_1(K, k) \\ &\leq \sum_{q \neq q_k^*} \widehat{\mathbf{F}}_1(q, k) + \widehat{\mathbf{A}}_1(k, \ell), \ell \neq \ell_k^* \\ &\leq \frac{(K-1)\kappa}{1 - \kappa} + \frac{\kappa(1 - \kappa)}{1 - K\kappa} \end{aligned} \quad (80)$$

where we have used the fact that all the entries of $\widehat{\mathbf{A}}_1$ and $\widehat{\mathbf{F}}_1$ are smaller than 1, in the first inequality and have applied (76) and (78) in the last inequality.

The lower-bound condition in (72a) gives that for each k ,

$$|\widehat{\mathbf{A}}_1(k, :) \widehat{\mathbf{F}}_1(:, k)| \geq 1 - \kappa. \quad (81)$$

Then, comparing both (80) and (81), we hope to have

$$\begin{aligned} \frac{(K-1)\kappa}{1 - \kappa} &\geq 1 - \kappa \\ \implies \kappa &\geq \frac{K+1 - \sqrt{(K+1)^2 - 4}}{2} \end{aligned}$$

However, $\frac{K+1 - \sqrt{(K+1)^2 - 4}}{2} \geq \frac{1}{K+1}$, which leads to a contradiction to our assumption that $\kappa \leq \frac{1}{K+1}$. Hence, the assumption (79) does not hold and we get

$$\arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) = \arg \max_q \widehat{\mathbf{F}}_1(q, k), \forall k. \quad (82)$$

Combining both (82) and (82), we get

$$\begin{aligned} \arg \max_{\ell} \widehat{\mathbf{A}}_1(k, \ell) &\neq \arg \max_{\ell} \widehat{\mathbf{A}}_1(j, \ell), \forall k \neq j \\ \arg \max_q \widehat{\mathbf{F}}_1(q, k) &\neq \arg \max_q \widehat{\mathbf{F}}_1(q, j), \forall k \neq j. \end{aligned}$$

G PROOF OF LEMMA 2

We have

$$\begin{aligned} \sigma_{\min}(\mathbf{X}) &= \sigma_{\min}(\mathbf{I}_K + \mathbf{E}_1 + \mathbf{E}_2) = \min_{\|\mathbf{x}\|_2=1} \|(\mathbf{I}_K + \mathbf{E}_1 + \mathbf{E}_2)\mathbf{x}\|_2 \\ &= \min_{\|\mathbf{x}\|_2=1} \|\mathbf{I}_K \mathbf{x} + (\mathbf{E}_1 + \mathbf{E}_2)\mathbf{x}\|_2 \stackrel{(a)}{\geq} \min_{\|\mathbf{x}\|_2=1} \left| \|\mathbf{I}_K \mathbf{x}\|_2 - \|(\mathbf{E}_1 + \mathbf{E}_2)\mathbf{x}\|_2 \right| \\ &\geq \left| \min_{\|\mathbf{x}\|_2=1} \|\mathbf{I}_K \mathbf{x}\|_2 - \max_{\|\mathbf{x}\|_2=1} \|(\mathbf{E}_1 + \mathbf{E}_2)\mathbf{x}\|_2 \right| = \left| 1 - \max_{\|\mathbf{x}\|_2=1} \|(\mathbf{E}_1 + \mathbf{E}_2)\mathbf{x}\|_2 \right| \\ &= |1 - \|\mathbf{E}_1 + \mathbf{E}_2\|_2| \stackrel{(b)}{\geq} |1 - \|\mathbf{E}_1\|_F - \|\mathbf{E}_2\|_F| \\ &\stackrel{(c)}{\geq} |1 - v_1 - v_2| \end{aligned} \quad (83)$$

where the inequality (a) employs the triangle inequality, (b) is obtained via the matrix norm equivalence relation $\|\mathbf{E}_1 + \mathbf{E}_2\|_2 \leq \|\mathbf{E}_1 + \mathbf{E}_2\|_F \leq \|\mathbf{E}_1\|_F + \|\mathbf{E}_2\|_F$, (c) is obtained by the assumption that $\|\mathbf{E}_1\|_F \leq \nu_1$ and $\|\mathbf{E}_2\|_F \leq \nu_2$.

H PROOF OF LEMMA 3

First, let us define the following w.r.t the loss function $g(\mathbf{f}, \mathbf{y}) = \|\mathbf{f} - \mathbf{y}\|_2^2$, $\forall \mathbf{f} \in \mathcal{F}, \mathbf{y} \in \mathbb{R}^K$ and the input data $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ where each \mathbf{x}_n is sampled i.i.d. from the distribution \mathcal{D} under Assumption 1:

$$L_{\mathbf{X}}(\hat{\mathbf{f}}) \triangleq \frac{1}{N} \sum_{n=1}^N g(\hat{\mathbf{f}}(\mathbf{x}_n), \mathbf{f}^{\natural}(\mathbf{x}_n)) = \frac{1}{N} \sum_{n=1}^N \|\hat{\mathbf{f}}(\mathbf{x}_n) - \Pi^{\top} \mathbf{f}^{\natural}(\mathbf{x}_n)\|_2^2 = \frac{1}{N} \|\hat{\mathbf{F}} - \Pi^{\top} \mathbf{F}^{\natural}\|_F^2, \quad (84a)$$

$$L_{\mathcal{D}}(\hat{\mathbf{f}}) \triangleq \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [g(\hat{\mathbf{f}}(\mathbf{x}), \mathbf{f}^{\natural}(\mathbf{x}))] = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\|\hat{\mathbf{f}}(\mathbf{x}) - \Pi^{\top} \mathbf{f}^{\natural}(\mathbf{x})\|_2^2]. \quad (84b)$$

We invoke Theorem 26.5 in (Shalev-Shwartz & Ben-David, 2014) and get that with probability greater than $1 - \delta$:

$$\begin{aligned} L_{\mathcal{D}}(\hat{\mathbf{f}}) &\leq L_{\mathbf{X}}(\hat{\mathbf{f}}) + 2\mathfrak{R}_N(g \circ \mathcal{F}) + 4\bar{c} \sqrt{\frac{2 \log(4/\delta)}{N}} \\ \implies \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\|\hat{\mathbf{f}}(\mathbf{x}) - \Pi^{\top} \mathbf{f}^{\natural}(\mathbf{x})\|_2^2] &\leq \frac{1}{N} \|\hat{\mathbf{F}} - \Pi^{\top} \mathbf{F}^{\natural}\|_F^2 + 4\mathfrak{R}_N(\mathcal{F}) + 16\sqrt{\frac{2 \log(4/\delta)}{N}} \end{aligned} \quad (85)$$

where the last inequality utilizes the definitions in (84), the contraction lemma (Lemma 26.9 from (Shalev-Shwartz & Ben-David, 2014)) and also applied $\bar{c} = 4$ since $|g(\mathbf{f}, \mathbf{y})| \leq 4$ in our case. The term $\mathfrak{R}_N(\mathcal{F})$ denotes the empirical Rademacher complexity of the neural network function class \mathcal{F} which is upperbounded via the sensitive complexity parameter $\mathcal{R}_{\mathcal{F}}$ as follows (Lin & Zhang, 2019):

$$\mathfrak{R}_N(\mathcal{F}) \leq 16N^{-5/8} (2\|\mathbf{X}\|_F \mathcal{R}_{\mathcal{F}})^{\frac{1}{4}}.$$

I PROOF OF PROPOSITION 2

Our goal is to bound $\frac{1}{M} \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \hat{\mathbf{p}}_n^{(m)}\|_2^2$ for each n . To achieve this, let us define the random variable $U := \frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \hat{\mathbf{p}}_n^{(m)}\|_2^2$. From Proposition 1, the following result hold:

$$\Pr(U \leq \vartheta(\delta)) \geq 1 - \delta,$$

where randomness is due to all \mathbf{x}_n 's, $\hat{\mathbf{y}}_n^{(m)}$'s, and \mathcal{S} and $\vartheta(\delta)$ is given by the R.H.S. of (54).

Then we have

$$\begin{aligned} \mathbb{E}[U] &= \int_0^{u_{\max}} h(u) u du = \int_0^{\vartheta(\delta)} h(u) u du + \int_{\vartheta(\delta)}^{u_{\max}} h(u) u du \\ &\leq \vartheta(\delta) \int_0^{\vartheta(\delta)} h(u) du + u_{\max} \int_{\vartheta(\delta)}^{u_{\max}} h(u) du \\ &\leq \vartheta(\delta) + u_{\max} \delta \\ &\leq \vartheta(\delta) + 4\delta \end{aligned} \quad (86)$$

where u_{\max} denotes the maximum value of the random variable U and is given by $u_{\max} \leq 4$ and $h(u)$ denotes the probability density function of the random variable U .

Also, we have

$$\mathbb{E}[U] = \frac{1}{N} \sum_{n=1}^N \mathbb{E}[U_n] \quad (87)$$

where $U_n \triangleq \frac{1}{M} \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \hat{\mathbf{p}}_n^{(m)}\|_2^2$.

To proceed, we have the following lemma:

Lemma 10 Let $U_n \triangleq \frac{1}{M} \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2$. Then

$$\mathbb{E}[U_n] = \mathbb{E}[U_{n'}], \forall n, n',$$

where expectation is taken w.r.t. all \mathbf{x}_n 's, $\widehat{\mathbf{y}}_n^{(m)}$'s, and \mathcal{S} .

The proof is provided in Sec. J.

Combining Lemma 10 with (86) and (87), we get

$$\mathbb{E}[U_n] \leq \vartheta(\delta) + 4\delta.$$

Applying Markov inequality, we get the following for any $\tau > 0$

$$\begin{aligned} \Pr(U_n \leq \tau \mathbb{E}[U_n]) &\geq 1 - \frac{1}{\tau} \\ \implies \Pr(U_n \leq \tau \vartheta(\delta) + 4\tau\delta) &\geq 1 - \frac{1}{\tau} \end{aligned}$$

Letting $\tau = N^\alpha$, where $0 < \alpha < 1$, we get that with probability greater than $1 - \frac{1}{N^\alpha}$

$$\frac{1}{M} \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \leq N^\alpha \vartheta(\delta) + 4N^\alpha \delta.$$

J PROOF OF LEMMA 10

Let us define the following:

$$\begin{aligned} u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N}) &\triangleq \frac{1}{M} \sum_{m=1}^M \|\mathbf{p}_n^{\natural(m)} - \widehat{\mathbf{p}}_n^{(m)}\|_2^2 \\ &= \frac{1}{M} \sum_{m=1}^M \|\mathbf{A}_m^{\natural} \mathbf{f}^{\natural}(\mathbf{x}_n) - \widehat{\mathbf{A}}_m \widehat{\mathbf{f}}(\mathbf{x}_n)\|_2^2 \end{aligned}$$

where $\mathbf{v}_n = \{\mathbf{x}_n, \{\widehat{\mathbf{y}}_n^{(m)}\}_{m \in \mathcal{S}_n}, \mathcal{S}_n\}$ and $\widehat{\boldsymbol{\theta}}$ denotes estimates $\{\widehat{\mathbf{A}}_m\}$ and $\widehat{\mathbf{f}}$ which are obtained using the data $\{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N\} = \{\mathbf{x}_1, \dots, \mathbf{x}_N, \widehat{\mathcal{Y}}, \mathcal{S}\}$.

One can observe that for any $n' \neq n$,

$$\widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N} = \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N} \quad (88)$$

since the order of \mathbf{v}_n does not affect the value of the estimates. Then we have

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_n, \mathbf{x}_{n'}, \mathbf{x}_j, j \neq n, n'} [u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N})] &= \mathbb{E}_{\mathbf{x}_n, \mathbf{x}_{n'}, \mathbf{x}_j, j \neq n, n'} [u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N})] \\ &= \mathbb{E}_{\mathbf{x}_n, \mathbf{x}_{n'}, \mathbf{x}_j, j \neq n, n'} [u(\mathbf{x}_{n'}, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N})] \quad (89) \end{aligned}$$

where the first equality is by (88) and the last equality is obtained since \mathbf{x}_n and $\mathbf{x}_{n'}$ are identically distributed.

From (89), we can obtain that

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_N} [u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N})] &= \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_N} [u(\mathbf{x}_{n'}, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N})] \\ \implies \mathbb{E}_{\widehat{\mathcal{Y}}} \left[\mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_N} [u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N})] \right] &= \mathbb{E}_{\widehat{\mathcal{Y}}} \left[\mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_N} [u(\mathbf{x}_{n'}, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N})] \right] \\ \implies \mathbb{E}_{\mathcal{S}} \left[\mathbb{E}_{\widehat{\mathcal{Y}}} \left[\mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_N} [u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N})] \right] \right] &= \mathbb{E}_{\mathcal{S}} \left[\mathbb{E}_{\widehat{\mathcal{Y}}} \left[\mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_N} [u(\mathbf{x}_{n'}, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N})] \right] \right] \\ \implies \mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_N} [u(\mathbf{x}_n, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_n, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_N})] &= \mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_N} [u(\mathbf{x}_{n'}, \widehat{\boldsymbol{\theta}}_{\mathbf{v}_1, \dots, \mathbf{v}_{n'}, \dots, \mathbf{v}_n, \dots, \mathbf{v}_N})]. \end{aligned}$$

$$\mathbf{W}^{\natural} = \begin{bmatrix} \vdots \\ \mathbf{A}_{m_1}^{\natural}(1,:) \\ \vdots \\ \mathbf{A}_{m_2}^{\natural}(2,:) \\ \vdots \\ \mathbf{A}_{m_3}^{\natural}(3,:) \\ \vdots \end{bmatrix} \approx \begin{bmatrix} \mathbf{e}_1^{\top} \\ \mathbf{e}_2^{\top} \\ \mathbf{e}_3^{\top} \end{bmatrix} \quad \mathbf{W}^{\natural} = \begin{bmatrix} \vdots \\ \mathbf{A}_{m_*}^{\natural} \\ \vdots \\ \vdots \\ \vdots \end{bmatrix} \approx \mathbf{I}_3$$

Figure 1: (left) $\mathbf{W}^{\natural} = [\mathbf{A}_1^{\natural\top}, \dots, \mathbf{A}_M^{\natural\top}]^{\top}$ satisfying NCSA with $K = 3$, meaning that there exists specialists for each class $k \in [K]$ and (right) all-class expert annotator m_* among M annotators, meaning $\mathbf{A}_{m_*}^{\natural} \approx \mathbf{I}_K$.




Figure 2: The dots denote the rows of \mathbf{W}^{\natural} and the circle denote the second-order cone \mathcal{C} .

K MORE DETAILS ON THE COMPLEXITY MEASURE $\mathcal{R}_{\mathcal{F}}$

The work in (Lin & Zhang, 2019) introduces a complexity measure for measuring the expressiveness of different family of neural network classes. For example, consider a general convolution neural network (CNN) architecture with L layers such that

$$\mathbf{f}(\mathbf{x}) = \sigma_L(\gamma_L(\mathbf{H}_L)\sigma_{L-1}(\gamma_{L-1}(\mathbf{H}_{L-1})\dots\sigma_1(\gamma_1(\mathbf{H}_1)\mathbf{x}))), \quad \forall \mathbf{f} \in \mathcal{F},$$

where $\sigma_i(\cdot) = [\sigma_i(\cdot), \dots, \sigma_i(\cdot)] \in \mathbb{R}^{d_i}$ represents an element-wise activation function at i th layer satisfying $\sigma_i(0) = 0$ and σ_i being ρ -Lipschitz continuous for all i . Let $\mathcal{G}_{\mathcal{F}}$ denote the set of layer indices with fully connected layers and $\mathcal{G}_{\mathcal{C}}$ denote the layer indices with convolutional layers. If $i \in \mathcal{G}_{\mathcal{F}}$, we have $\gamma_i(\mathbf{H}_i) = \mathbf{H}_i \in \mathbb{R}^{d_i \times d_{i-1}}$. Suppose $i \in \mathcal{G}_{\mathcal{C}}$, then we have $\gamma_i(\mathbf{H}_i)\mathbf{z}_{i-1} = \mathbf{H}_i \otimes \mathbf{z}_{i-1}$, where the matrix $\mathbf{H}_i \in \mathbb{R}^{c_i \times r_i}$ contains c_i convolutional filters each of which having dimension r_i , \mathbf{z}_{i-1} denotes the output of the $(i-1)$ th layer, and \otimes denotes the convolution operation. Here $d_0 = d$ is the dimension of the of the input data items \mathbf{x}_n 's. Then, its sensitive complexity $\mathcal{R}_{\mathcal{F}}$ is defined as follows:

$$\mathcal{R}_{\mathcal{F}} = 2\rho^L L^2 \left(\sum_{i \in \mathcal{G}_{\mathcal{F}}} d_i^2 d_{i-1}^2 + \sum_{i \in \mathcal{G}_{\mathcal{C}}} c_i^2 r_i^2 \sqrt{d_i/c_i} \right).$$

Note that such as a general CNN architecture covers many popular neural networks, e.g., fully connected neural networks, CNNs such as Lenet-5 (Lecun et al., 1998), VGG-16 (Liu & Deng, 2015), and so on.

L NEAR-CLASS SPECIALIST ASSUMPTION AND SSC

NCSA is a relaxed condition compared to having all-class expert annotators or having diagonally dominant annotators—see Fig. 1. It can be understood that NCSA assumption does not require any single annotator to be specialists with respect to all classes. Nonetheless, having an all-class expert annotator m_* satisfies NCSA, but not vice versa.

Algorithm 1: GeoCrowdNet (F)

```

input :Data  $\{\mathbf{x}_n, \hat{y}_n^{(m)}, m \in \mathcal{S}_n\}_{n=1}^N$ 
1 Initialize  $\mathbf{A}_m$ 's to identity matrices  $\mathbf{I}_K$ ;
2 Initialize the parameters  $\theta$  of the neural network function class  $\mathcal{F}$ ;
3  $\psi = [\text{vec}(\mathbf{A}_1)^\top, \dots, \text{vec}(\mathbf{A}_M)^\top, \text{vec}(\theta)^\top]^\top$ ;
4 while stopping criterion is not reached do
5   while stopping criterion is not reached do
6     Draw a random batch  $\mathcal{B}$ ;
7     Compute  $\nabla_\psi \mathcal{L}_F(\psi, \mathcal{B})$ ;
8      $\psi \leftarrow \text{AdamOptimizer}(\psi, \nabla_\psi \mathcal{L}_F(\psi, \mathcal{B}))$ ;
9     for  $m = 1$  to  $M$  do
10       $\mathbf{A}_m \leftarrow \text{softmax}(\mathbf{A}_m)$  ▷ softmax operation on the columns of  $\mathbf{A}_m$ 
11    end
12  end
13 end
output :Estimates  $\hat{\mathbf{A}}_m, \forall m, \hat{\theta}$ 

```

The SSC (Definition 1) is a relaxed condition compared to the NCSA (Assumption 4). To illustrate this, we consider the geometry of the matrix \mathbf{W}^{\sharp} satisfying the NCSA and the SSC as shown in Fig. 2.

M ALGORITHM DESCRIPTION

In this section, we detail the implementation of the regularized criteria in (8) and (9). The neural network predictor function $\mathbf{f} \in \mathcal{F}$ is parameterized using θ and can be denoted as $\mathbf{f}_\theta : \mathbb{R}^d \rightarrow \Delta$. Let $\psi = [\text{vec}(\mathbf{A}_1)^\top, \dots, \text{vec}(\mathbf{A}_M)^\top, \text{vec}(\theta)^\top]^\top$.

Let us also define the following for certain batch $\mathcal{B} \subset [N]$:

$$\mathcal{L}_F(\psi, \mathcal{B}) = -\frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \sum_{m \in \mathcal{S}_n} \sum_{k=1}^K \mathbb{1}[\hat{y}_n^{(m)} = k] \log[\mathbf{A}_m \mathbf{f}_\theta(\mathbf{x}_n)]_k - \lambda \log \det \tilde{\mathbf{F}} \tilde{\mathbf{F}}^\top,$$

$$\mathcal{L}_W(\psi, \mathcal{B}) = -\frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \sum_{m \in \mathcal{S}_n} \sum_{k=1}^K \mathbb{1}[\hat{y}_n^{(m)} = k] \log[\mathbf{A}_m \mathbf{f}_\theta(\mathbf{x}_n)]_k - \lambda \log \det \mathbf{W}^\top \mathbf{W},$$

where $\tilde{\mathbf{F}} = [\mathbf{f}(\mathbf{x}_{n_1}), \dots, \mathbf{f}(\mathbf{x}_{n_{|\mathcal{B}|}})]$, $n_1, \dots, n_{|\mathcal{B}|} \in \mathcal{B}$ and $\mathbf{W} = [\mathbf{A}_1^\top, \dots, \mathbf{A}_M^\top]^\top$.

Let $\nabla_\psi \mathcal{L}_F(\psi, \mathcal{B})$ denote the stochastic gradient of $\mathcal{L}_F(\psi, \mathcal{B})$ w.r.t. ψ and $\nabla_\psi \mathcal{L}_W(\psi, \mathcal{B})$ denote the stochastic gradient of $\mathcal{L}_W(\psi, \mathcal{B})$ w.r.t. ψ . Under these notations, the proposed methods GeoCrowdNet (F) and GeoCrowdNet (W) are described in Algorithm ?? and ??, respectively. The Python implementation of the algorithms is available in GitHub¹.

N ADDITIONAL EXPERIMENTS AND IMPLEMENTATION DETAILS

N.1 SYNTHETIC ANNOTATORS

Dataset. We consider the CIFAR-10 (Krizhevsky, 2009) dataset for synthetic noisy label experiments. CIFAR-10 consists of 60,000 labeled color images of animals, vehicles and so on, each having a size of 32×32 and belonging $K = 10$ different classes. We use 45,000 images for training, 5,000 images for validation, and 10,000 images for testing.

Noisy Label Generation. In order to produce noisy annotations for the images of the training data, we simulate $M = 5$ synthetic annotators. We randomly choose an annotator m^* among M annotators and its confusion matrix is generated by $\mathbf{A}_{m^*}^\sharp = \mathbf{I}_K + \gamma \text{rand}(K, K)$, followed by normalization w.r.t. the ℓ_1 norm of the corresponding columns. Here, \mathbf{I}_K denotes the identity matrix of size K and

¹<https://github.com/shahanaibrahimosu/end-to-end-crowdsourcing>

Algorithm 2: GeoCrowdNet (\bar{W})

```

input :Data  $\{\mathbf{x}_n, \hat{y}_n^{(m)}, m \in \mathcal{S}_n\}_{n=1}^N$ 
1 Initialize  $\mathbf{A}_m$ 's to identity matrices  $\mathbf{I}_K$ ;
2 Initialize the parameters  $\theta$  of the neural network function class  $\mathcal{F}$ ;
3  $\psi = [\text{vec}(\mathbf{A}_1)^\top, \dots, \text{vec}(\mathbf{A}_M)^\top, \text{vec}(\theta)^\top]^\top$ ;
4 while stopping criterion is not reached do
5   while stopping criterion is not reached do
6     Draw a random batch  $\mathcal{B}$ ;
7     Compute  $\nabla_\psi \mathcal{L}_W(\psi, \mathcal{B})$ ;
8      $\psi \leftarrow \text{AdamOptimizer}(\psi, \nabla_\psi \mathcal{L}_W(\psi, \mathcal{B}))$ ;
9     for  $m = 1$  to  $M$  do
10       $\mathbf{A}_m \leftarrow \text{softmax}(\mathbf{A}_m)$  ▷ softmax operation on the columns of  $\mathbf{A}_m$ 
11    end
12  end
output :Estimates  $\hat{\mathbf{A}}_m, \forall m, \hat{\theta}$ 
13 end

```

$\text{rand}(K, K)$ denotes the $K \times K$ matrix with its entries randomly chosen from a uniform distribution between 0 and 1. The parameter γ controls how well the annotator m^* correctly identifies the ground-truth labels. The confusion matrices of the remaining $M - 1$ annotators are generated such that they provide labels uniformly at random—i.e., the $M - 1$ annotators are all unreliable. This type of modeling for the confusion matrices resembles the *hammer-spammer* model as employed in the works (Rodrigues & Pereira, 2018; Tanno et al., 2019). Using the ground-truth label y_n 's provided by the dataset and the generated confusion matrices, the noisy annotations $\hat{y}_n^{(m)}$'s are produced. Each annotation $\hat{y}_n^{(m)}$ is observed independently at random such that only 20% of the total annotations are available for training. Note that, the true labels are not accessible by any of the methods.

Neural Network Architecture and Settings. For the CIFAR-10 dataset, we choose the ResNet-9 architecture (He et al., 2016). Adam (Kingma & Ba, 2015) is used as the optimizer with weight decay of 10^{-4} and a batch size of 128. The regularization parameter λ and the initial learning rate of the Adam optimizer are chosen via grid search method over the validation set from $\{0.01, 0.001, 0.0001\}$ and $\{0.01, 0.001\}$, respectively. We choose the same neural network structures for all the baselines. The confusion matrices are initialized with identity matrices for the proposed method and the baselines TraceReg and CrowdLayer.

Results. Table 3 presents the results under various values of γ —when γ is smaller, the chance of annotator m^* correctly labeling the data items is better. One can see that the proposed approach, namely, GeoCrowdNet (F), outperforms the baselines in all the scenarios under test. When annotator m^* 's labeling accuracy drops drastically (i.e., when γ becomes larger), GeoCrowdNet (F) performs the best. This implies that even if there are no class specialists, GeoCrowdNet (F) works well under the large N cases.

Figs. 3, 4 and 5 show the ground-truth confusion matrices and the corresponding estimates when $\gamma = 0.01$ for the GeoCrowdNet methods. One can see all methods estimate the confusion matrices reasonably well, corroborating our identifiability claims.

N.2 REAL DATA EXPERIMENTS - MACHINE ANNOTATIONS

Here, we provide additional details of the real data experiments. For experiments with machine annotations, we consider the MNIST and the Fashion-MNIST datasets. Each image in these datasets is of size 28×28 in the grey-scale format. We consider 57,000 images as training data, 3000 images for validation, and 10,000 images for testing.

For machine annotation, we train a number of machine classifiers to act as annotators. In Table 1, under Case 2, we choose $M = 5$ annotators. Specifically, Linear SVM, logistic regression, k -NN with $k = 5$, CNN with two convolution layers followed by a max pooling layer, a fully connected layer and a softmax layer, and a fully connected neural network (FCNN) with 1 hidden layer and 128 hidden units are trained. The CNN is trained for 5 epochs and the FCNN is trained for 10 epochs. The classifiers are intentionally not well trained, so that they can mimic error-prone annotators.

Table 3: Average test accuracy of the proposed methods and the baselines on the CIFAR-10 dataset ($K = 10$), labeled by $M = 5$ synthetic annotators.

Methods	$\gamma = 0.01$	$\gamma = 0.2$	$\gamma = 0.3$
GeoCrowdNet (F)	84.82 ± 0.43	71.66 ± 1.08	67.97 ± 1.44
GeoCrowdNet (W)	82.76 ± 0.46	69.20 ± 0.59	63.86 ± 1.64
GeoCrowdNet ($\lambda=0$)	82.50 ± 0.38	69.14 ± 1.29	63.73 ± 2.64
TraceReg	82.54 ± 0.34	70.01 ± 1.16	65.01 ± 2.30
Crowdlayer	83.01 ± 0.23	69.91 ± 1.48	65.70 ± 2.32
MBEM	81.01 ± 0.32	69.45 ± 1.56	64.15 ± 2.21
CoNAL	81.41 ± 0.43	68.56 ± 1.78	63.21 ± 1.98
Max-MIG	82.01 ± 0.78	69.96 ± 1.23	62.01 ± 2.46
NN-MV	50.96 ± 0.83	37.90 ± 1.19	33.24 ± 1.17
NN-DSEM	52.02 ± 0.80	37.95 ± 1.41	33.64 ± 0.69

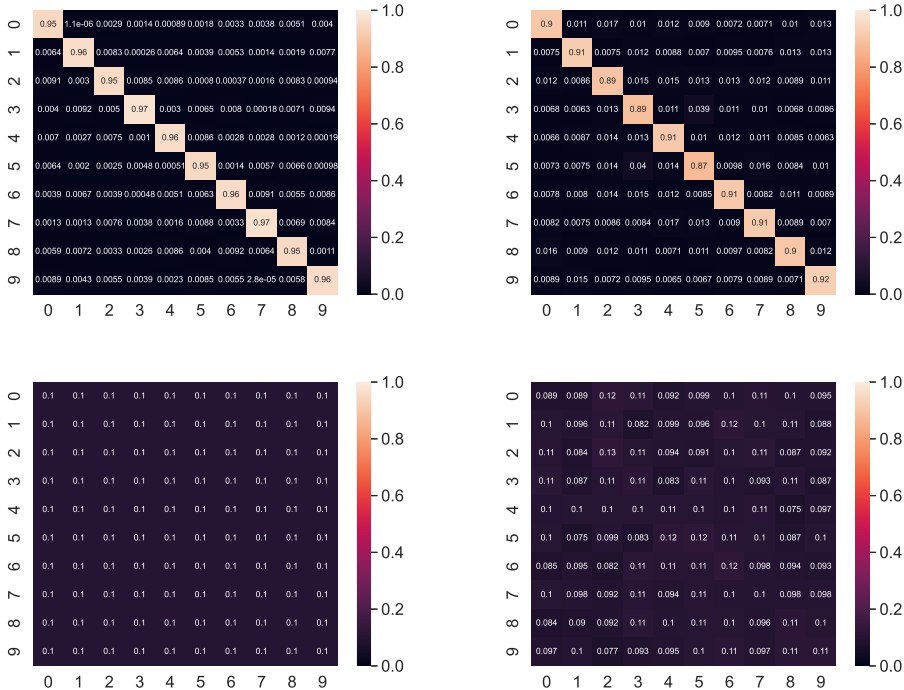


Figure 3: The illustration of the ground-truth confusion matrices and the estimated confusion matrices by the proposed approach GeoCrowdNet (F) for CIFAR-10 dataset, with $M = 5$ machine annotators and $\gamma = 0.01$. The ground-truth $A_{m^*}^h$ (top left) and the estimate \hat{A}_{m^*} (top right). The ground-truth A_m^h (bottom left) and the estimate \hat{A}_m (bottom right), $m \neq m^*$. The average mean squared error (MSE) of the confusion matrices is 0.088.

For example, for the MNIST dataset, the individual label prediction accuracy of these annotators on unseen data items of size N ranges from 15.88% to 82.23% (averaged over 5 random trials) and 2 annotators have less than 50% accuracy on average. Also note that our algorithm does not know which annotator is more accurate.

For Case 2, we train more annotators by changing some parameters of the above mentioned machine classifiers. Specifically, we use linear SVM, polynomial kernel-based SVM, and Gaussian SVM as different variants of SVM. For k -NN, we choose k from $\{3, 5, 7, 10\}$. FCNN, CNN, and logistic regression are trained with different number of epochs $\{10, 15, 20, 25\}$ and with random initializations for each case. Under case 1, if an annotator is chosen to be a specialist, it is trained with more number

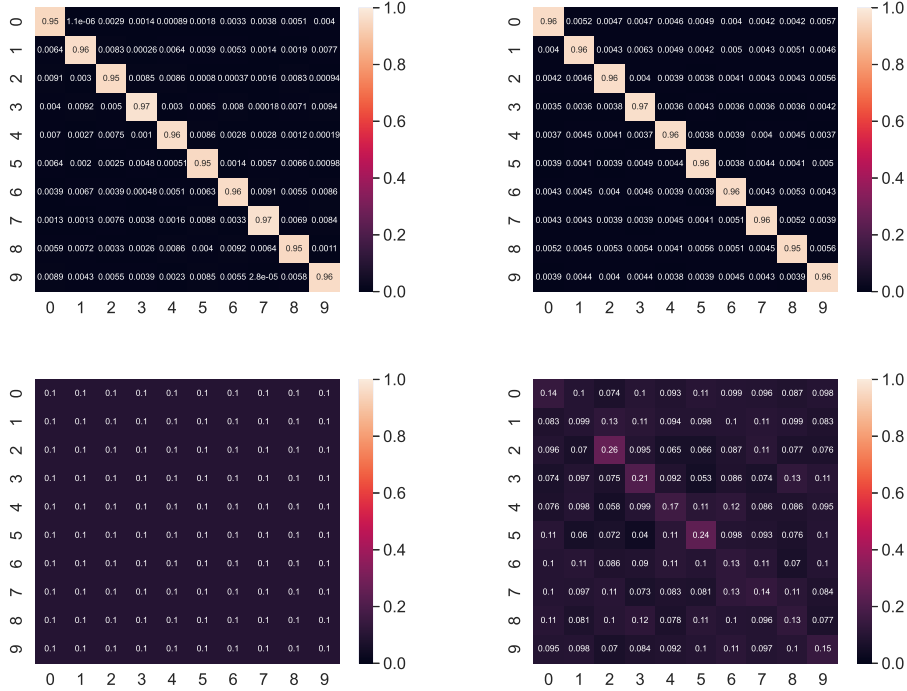


Figure 4: The illustration of the ground-truth confusion matrices and the estimated confusion matrices by the proposed approach GeoCrowdNet (W) for CIFAR-10 dataset, with $M = 5$ machine annotators and $\gamma = 0.01$. The ground-truth $A_{m^*}^h$ (top left) and the estimate $\hat{A}_{m^*}^h$ (top right). The ground-truth A_m^h (bottom left) and the estimate \hat{A}_m^h (bottom right), $m \neq m^*$. The average mean squared error (MSE) of the confusion matrices is 0.102.

Table 4: Average test accuracy (\pm std) of the proposed methods and the baselines on Fashion-MNIST under various (N, M) 's on Case 1; labels are produced by machine annotators; $p = 0.1$.

Methods	(1000, 25)	(750, 30)
GeoCrowdNet (F)	82.81 \pm 2.05	82.40 \pm 1.51
GeoCrowdNet (W)	83.33 \pm 3.39	83.65 \pm 4.11
GeoCrowdNet ($\lambda = 0$)	79.43 \pm 4.07	79.07 \pm 2.91
TraceReg	80.79 \pm 3.87	81.12 \pm 4.03
CrowdLayer	73.54 \pm 6.58	75.56 \pm 8.19
MBEM	26.19 \pm 3.45	26.67 \pm 3.45
CoNAL	63.13 \pm 8.53	66.02 \pm 5.65
Max-MIG	77.05 \pm 3.06	72.59 \pm 8.38
NN-MV	68.48 \pm 3.15	71.74 \pm 2.21
NN-DSEM	68.84 \pm 2.46	73.08 \pm 4.09

of samples (10,000 samples) with more number of epochs such that its label prediction accuracy is higher than 90%. Under this strategy, in Table 1 for the case $M = 15$, the individual label prediction accuracy of the annotators on unseen data items of size N ranges from 3.15% to 95.27% (2 annotators with more than 90% accuracy). For the case with $M = 20$, the individual label prediction accuracy of the annotators on unseen data items of size N ranges from 3.24% to 95.73% (3 annotators with more than 90% accuracy).

Additional experiment results using the MNIST and Fashion-MNIST datasets can be found in Tables 4-6. As expected, GeoCrowdNet (W) and GeoCrowdNet (F) offer the best performance in Case 1 and Case 2, respectively.

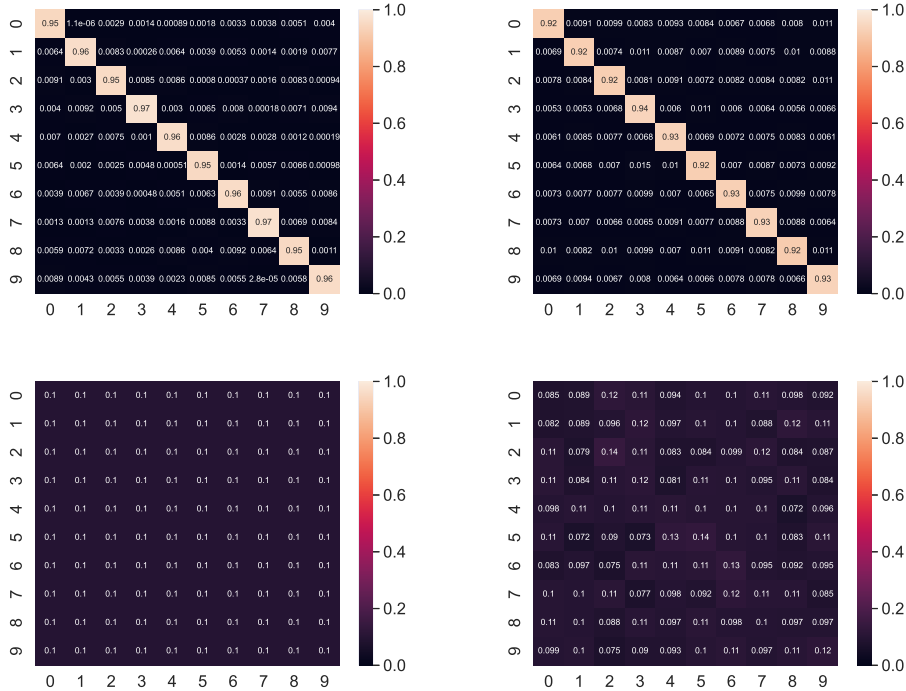


Figure 5: The illustration of the ground-truth confusion matrices and the estimated confusion matrices by GeoCrowdNet ($\lambda = 0$) for CIFAR-10 dataset, with $M = 5$ machine annotators and $\gamma = 0.01$. The ground-truth $A_{m^*}^h$ (top left) and the estimate \hat{A}_{m^*} (top right). The ground-truth A_m^h (bottom left) and the estimate \hat{A}_m (bottom right), $m \neq m^*$. The average mean squared error (MSE) of the confusion matrices is 0.097.

Table 5: Average test accuracy (\pm std) of the proposed methods and the baselines on MNIST under various (N, M) 's on Case 2; labels are produced by machine annotators; $p = 0.1$.

Methods	(5000, 25)	(10000, 25)
GeoCrowdNet (F)	87.26 \pm 0.70	87.66 \pm 1.23
GeoCrowdNet (W)	84.64 \pm 4.04	86.53 \pm 3.24
GeoCrowdNet ($\lambda = 0$)	80.08 \pm 3.14	85.51 \pm 2.35
TraceReg	78.06 \pm 6.72	86.85 \pm 5.45
CrowdLayer	57.39 \pm 6.93	56.66 \pm 6.34
MBEM	35.56 \pm 7.92	21.82 \pm 6.54
CoNAL	46.35 \pm 6.75	46.86 \pm 5.43
Max-MIG	83.34 \pm 1.54	86.84 \pm 2.13
NN-MV	83.46 \pm 0.67	83.32 \pm 1.34
NN-DSEM	84.37 \pm 0.70	83.78 \pm 1.11

N.3 REAL DATA EXPERIMENTS - HUMAN ANNOTATORS

For experiments with annotations collected from AMT, we use the LabelMe dataset (Rodrigues et al., 2017; Russell et al., 2007) and the Music dataset (Rodrigues et al., 2014).

The LabelMe dataset (Rodrigues et al., 2017; Russell et al., 2007) is an image classification dataset consisting of 2688 images from $K = 8$ different classes, namely, highway, inside city, tall building, street, forest, coast, mountain, and open country. From the available images, 1000 images are annotated by $M = 59$ AMT workers. In total, about 2547 image annotations are obtained by the AMT workers whose labeling accuracy ranges from 0% to 100% with a mean accuracy of 69.2%. In

Table 6: Average test accuracy (\pm std) of the proposed methods and the baselines on Fashion-MNIST under various (N, M) 's on Case 2; labels are produced by machine annotators; $p = 0.1$.

Methods	(5000, 25)	(10000, 25)
GeoCrowdNet (F)	87.47 \pm 0.86	88.39 \pm 1.92
GeoCrowdNet (W)	85.80 \pm 1.26	87.09 \pm 2.71
GeoCrowdNet ($\lambda = 0$)	85.58 \pm 2.23	85.37 \pm 2.59
TraceReg	85.20 \pm 2.36	84.02 \pm 2.74
CrowdLayer	70.07 \pm 7.19	74.92 \pm 6.19
MBEM	42.15 \pm 3.65	44.34 \pm 3.48
CoNAL	65.13 \pm 4.94	62.65 \pm 7.02
Max-MIG	85.33 \pm 2.09	85.94 \pm 1.93
NN-MV	79.85 \pm 2.88	83.53 \pm 1.53
NN-DSEM	81.34 \pm 1.65	84.67 \pm 1.15

order to enrich the training dataset, standard augmentation techniques such as rescaling, cropping, horizontal flips, etc., are employed and accordingly, the training dataset consists of 10,000 images annotated by 59 workers; see more details in (Rodrigues & Pereira, 2018; Chu et al., 2021). The validation set consists of 500 images and the remaining 1188 images are used for testing.

The Music dataset (Rodrigues et al., 2014) consists of audio samples of 10 different genres of music such as classical, country, disco, hiphop, jazz, rock, blues, reggae, pop, and metal. The dataset has about 1000 samples of songs each having a duration of 30 seconds. About 700 samples are annotated by 44 AMT workers (overall, about 2946 annotations are observed) and the remaining 300 are allocated for testing. Out of the annotated 700 samples, we consider 595 samples for training, and the remaining 105 samples are used for validation.

For the LabelMe dataset, we employ the settings used in (Rodrigues & Pereira, 2018). The pretrained VGG-16 embeddings for the images are given as inputs to a fully connected (FC) neural network with one hidden layer having 128 hidden units and ReLU activation functions. A dropout layer with rate 50% is also used, followed by a softmax layer. For the Music dataset, we choose the same FC layer and the softmax layer as employed in the LabelMe settings, but with batch normalization layers before each of these layers.