



Published in final edited form as:

Proc AAAI Conf Artif Intell. 2019 ; 33: 4763–4771. doi:10.1609/aaai.v33i01.33014763.

Training Complex Models with Multi-Task Weak Supervision

Alexander Ratner[†], Braden Hancock[†], Jared Dunnmon[†], Frederic Sala[†], Shreyash Pandey[†], Christopher Ré[†]

[†]Department of Computer Science, Stanford University

Abstract

As machine learning models continue to increase in complexity, collecting large hand-labeled training sets has become one of the biggest roadblocks in practice. Instead, weaker forms of supervision that provide noisier but cheaper labels are often used. However, these weak supervision sources have diverse and unknown accuracies, may output correlated labels, and may label different tasks or apply at different levels of granularity. We propose a framework for integrating and modeling such weak supervision sources by viewing them as labeling different related sub-tasks of a problem, which we refer to as the *multi-task weak supervision* setting. We show that by solving a matrix completion-style problem, we can recover the accuracies of these *multi-task* sources given their dependency structure, but without any labeled data, leading to higher-quality supervision for training an end model. Theoretically, we show that the generalization error of models trained with this approach improves with the number of *unlabeled* data points, and characterize the scaling with respect to the task and dependency structures. On three fine-grained classification problems, we show that our approach leads to average gains of 20.2 points in accuracy over a traditional supervised approach, 6.8 points over a majority vote baseline, and 4.1 points over a previously proposed weak supervision method that models tasks separately.

1 Introduction

One of the greatest roadblocks to using modern machine learning models is collecting hand-labeled training data at the massive scale they require. In real-world settings where domain expertise is needed and modeling goals change frequently, hand-labeling training sets is prohibitively slow, expensive, and static. For these reasons, practitioners are increasingly turning to weak supervision techniques wherein noisier, often programmatically-generated labels are used instead. Common *weak supervision sources* include external knowledge bases [24; 37; 8; 31], heuristic patterns [14; 27], feature annotations [23; 36], and noisy crowd labels [17; 11]. The use of these sources has led to state-of-the-art results in a range of domains [37; 35]. A theme of weak supervision is that using the full diversity of available sources is critical to training high-quality models [27; 37].

The key technical difficulty of weak supervision is determining how to combine the labels of multiple sources that have different, unknown accuracies, may be correlated, and may label

at different levels of granularity. In our experience with users in academia and industry, the complexity of real world weak supervision sources makes this integration phase the key time sink and stumbling block. For example, if we are training a model to classify entities in text, we may have one available source of high-quality but coarse-grained labels (e.g. “Person” vs. “Organization”) and one source that provides lower-quality but finer-grained labels (e.g. “Doctor” vs. “Lawyer”); moreover, these sources might be correlated due to some shared component or data source [2; 33]. Handling such diversity requires addressing a core technical challenge: estimating the unknown accuracies of multi-granular and potentially correlated supervision sources without any labeled data.

To overcome this challenge, we propose MeTaL, a framework for modeling and integrating weak supervision sources with different unknown accuracies, correlations, and granularities. In MeTaL, we view each source as labeling one of several related sub-tasks of a problem—we refer to this as the *multi-task weak supervision* setting. We then show that given the dependency structure of the sources, we can use their observed agreement and disagreement rates to recover their unknown accuracies. Moreover, we exploit the relationship structure between tasks to observe additional cross-task agreements and disagreements, effectively providing extra signal from which to learn. In contrast to previous approaches based on sampling from the posterior of a graphical model directly [28; 2], we develop a simple and scalable matrix completion-style algorithm, which we are able to analyze by applying strong matrix concentration bounds [32]. We use this algorithm to learn and model the accuracies of diverse weak supervision sources, and then combine their labels to produce training data that can be used to supervise arbitrary models, including increasingly popular multi-task learning models [5; 29].

Compared to previous methods which only handled the single-task setting [28; 27], and generally considered conditionally-independent sources [1; 11], we demonstrate that our multi-task aware approach leads to average gains of 4.1 points in accuracy in our experiments, and has at least three additional benefits. First, many dependency structures between weak supervision sources may lead to non-identifiable models of their accuracies, where a unique solution cannot be recovered. We provide a compiler-like check to establish identifiability—i.e. the existence of a unique set of source accuracies—for arbitrary dependency structures, without resorting to the standard assumption of non-adversarial sources [11], alerting users to this potential stumbling block that we have observed in practice. Next, we provide sample complexity bounds that characterize the benefit of adding additional unlabeled data and the scaling with respect to the user-specified task and dependency structure. While previous approaches required thousands of sources to give non-vacuous bounds, we capture regimes with small numbers of sources, better reflecting the real-world uses of weak supervision we have observed. Finally, we are able to solve our proposed problem directly with SGD, leading to over 100× faster runtimes compared to prior Gibbs-sampling based approaches [28; 26], and enabling simple implementation using libraries like PyTorch.

We validate our framework on three fine-grained classification tasks in named entity recognition, relation extraction, and medical document classification, for which we have diverse weak supervision sources at multiple levels of granularity. We show that by

modeling them as labeling hierarchically-related sub-tasks and utilizing unlabeled data, we can get an average improvement of 20.2 points in accuracy over a traditional supervised approach, 6.8 points over a basic majority voting weak supervision baseline, and 4.1 points over data programming [28], an existing weak supervision approach in the literature that is not multi-task-aware. We also extend our framework to handle unipolar sources that only label one class, a critical aspect of weak supervision in practice that leads to an average 2.8 point contribution to our gains over majority vote. From a practical standpoint, we argue that our framework represents an efficient way for practitioners to supervise modern machine learning models, including new multi-task variants, for complex tasks by opportunistically using the diverse weak supervision sources available to them. To further validate this, we have released an open-source implementation of our framework.¹

2 Related Work

Our work builds on and extends various settings studied in machine learning.

Weak Supervision:

We draw motivation from recent work which models and integrates weak supervision using generative models [28; 27; 2] and other methods [13; 18]. These approaches, however, do not handle multi-granularity or multi-task weak supervision, require expensive sampling-based techniques that may lead to non-identifiable solutions, and leave room for sharper theoretical characterization of weak supervision scaling properties. More generally, our work is motivated by a wide range of specific weak supervision techniques, which include traditional distant supervision approaches [24; 8; 37; 15; 31], co-training methods [4], pattern-based supervision [14; 37], and feature-annotation techniques [23; 36; 21].

Crowdsourcing:

Our approach also has connections to the crowdsourcing literature [17; 11], and in particular to spectral and method of moments-based approaches [38; 9; 12; 1]. In contrast, the goal of our work is to support and explore settings not covered by crowdsourcing work, such as sources with correlated outputs, the proposed multi-task supervision setting, and regimes wherein a small number of labelers (weak supervision sources) each label a large number of items (data points). Moreover, we theoretically characterize the generalization performance of an end model trained with the weakly labeled data.

Multi-Task Learning:

Our proposed approach is motivated by recent progress on multi-task learning models [5; 29; 30], in particular their need for multiple large hand-labeled training datasets. We note that the focus of our paper is on generating supervision for these models, not on the particular multi-task learning model being trained, which we seek to control for by fixing a simple architecture in our experiments.

¹ github.com/HazyResearch/metal

Our work is also related to recent techniques for estimating classifier accuracies without labeled data in the presence of structural constraints [26]. We use matrix structure estimation [22] and concentration bounds [32] for our core results.

3 Programming Machine Learning with Weak Supervision

As modern machine learning models become both more complex and more performant on a range of tasks, developers increasingly interact with them by programmatically generating noisier or *weak* supervision. These approaches of effectively *programming* machine learning models have recently been formalized by the following pipeline [28; 27]: First, users provide one or more *weak supervision sources*, which are applied to unlabeled data to generate a set of noisy labels. These labels may overlap and conflict; we model and combine them via a *label model* in order to produce a final set of training labels. These labels are then used to train some discriminative model, which we refer to as the *end model*. This programmatic weak supervision approach can utilize sources ranging from heuristic rules to other models, and in this way can also be viewed as a pragmatic and flexible form of multi-source *transfer learning*.

In our experiences with users from science and industry, we have found it critical to utilize all available sources of weak supervision for complex modeling problems, including ones which label at multiple levels of *granularity*. However, this diverse, multi-granular weak supervision does not easily fit into existing paradigms. We propose a formulation where each weak supervision source labels some sub-task of a problem, which we refer to as the *multi-task weak supervision* setting. We consider an example:

Example 1

A developer wants to train a fine-grained Named Entity Recognition (NER) model to classify mentions of entities in the news (Figure 2). She has a multitude of available weak supervision sources which she believes have relevant signal for her problem—for example, pattern matchers, dictionaries, and pre-trained generic NER taggers. However, it is unclear how to properly use and combine them: some of them label phrases coarsely as PERSON versus ORGANIZATION, while others classify specific fine-grained types of people or organizations, with a range of unknown accuracies. In our framework, she can represent them as labeling tasks of different granularities—e.g. $Y_1 = \{Person, Org\}$, $Y_2 = \{Doctor, Lawyer, N/A\}$, $Y_3 = \{Hospital, Office, N/A\}$, where the label N/A applies, for example, when the type-of-person task is applied to an organization.

In our proposed multi-task supervision setting, the user specifies a set of structurally-related *tasks*, and then provides a set of *weak supervision sources* which are user-defined functions that either label each data point or abstain for each task, and may have some user-specified dependency structure. These sources can be arbitrary black-box functions, and can thus subsume a range of weak supervision approaches relevant to both text and other data modalities, including use of pattern-based heuristics, distant supervision [24], crowd labels, other weak or biased classifiers, declarative rules over unsupervised feature extractors [33], and more. Our goal is to estimate the unknown accuracies of these sources, combine their outputs, and use the resulting labels to train an end model.

4 Modeling Multi-Task Weak Supervision

The core technical challenge of the *multi-task weak supervision* setting is recovering the unknown *accuracies* of weak supervision sources given their dependency structure and a schema of the tasks they label, but without any ground-truth labeled data. We define a new algorithm for recovering the accuracies in this setting using a matrix completion-style optimization objective. We establish conditions under which the resulting estimator returns a unique solution. We then analyze the sample complexity of our estimator, characterizing its scaling with respect to the amount of *unlabeled data*, as well as the task schema and dependency structure, and show how the estimation error affects the generalization performance of the end model we aim to train. Finally, we highlight how our approach handles abstentions and *unipolar* sources, two critical scenarios in the weak supervision setting.

4.1 A Multi-Task Weak Supervision Estimator

Problem Setup—Let $X \in \mathcal{X}$ be a data point and $\mathbf{Y} = [Y_1, Y_2, \dots, Y_t]^T$ be a vector of categorical *task labels*, $Y_i \in \{1, \dots, k_i\}$, corresponding to t tasks, where (X, \mathbf{Y}) is drawn i.i.d. from a distribution \mathcal{D} (for a glossary of all variables used, see Appendix A.1).

The user provides a specification of how these tasks relate to each other; we denote this schema as the *task structure* G_{task} . The task structure expresses logical relationships between tasks, defining a *feasible set* of label vectors \mathcal{Y} , such that $\mathbf{Y} \in \mathcal{Y}$. For example, Figure 2 illustrates a hierarchical task structure over three tasks of different granularities pertaining to a fine-grained entity classification problem. Here, the tasks are related by logical subsumption relationships: for example, if $Y_2 = \text{DOCTOR}$, this implies that $Y_1 = \text{PERSON}$, and that $Y_3 = \text{N/A}$, since the task label Y_3 concerns types of organizations, which is inapplicable to persons. Thus, in this task structure, $\mathbf{Y} = [\text{PERSON}, \text{DOCTOR}, \text{N/A}]^T$ is in \mathcal{Y} while $\mathbf{Y} = [\text{PERSON}, \text{N/A}, \text{HOSPITAL}]^T$ is not. While task structures are often simple to define, as in the previous example, or are explicitly defined by existing resources—such as ontologies or graphs—we note that if no task structure is provided, our approach becomes equivalent to modeling the t tasks separately, a baseline we consider in the experiments.

In our setting, rather than observing the true label \mathbf{Y} , we have access to m *multi-task weak supervision* sources $s_j \in \mathcal{S}$ which emit label vectors $\boldsymbol{\lambda}_j$ that contain labels for some subset of the t tasks. Let 0 denote a null or abstaining label, and let the *coverage set* $\tau_j \subseteq \{1, \dots, t\}$ be the fixed set of tasks for which the j th source emits non-zero labels, such that $\lambda_i \in \mathcal{Y}_{\tau_j}$. For convenience, we let $\tau_0 = \{1, \dots, t\}$ so that $\mathcal{Y}_{\tau_0} = \mathcal{Y}$. For example, a source from our previous example might have a coverage set $\tau_j = \{1, 3\}$, emitting coarse-grained labels such as $\lambda_i = [\text{PERSON}, 0, \text{N/A}]^T$. Note that sources often label multiple tasks implicitly due to the constraints of the task structure; for example, a source that labels types of people (Y_2) also implicitly labels people vs. organizations ($Y_1 = \text{PERSON}$), and types of organizations (as $Y_3 = \text{N/A}$). Thus sources tailored to different tasks still have agreements and disagreements; we use this additional *cross-task* signal in our approach.

The user also provides the conditional dependency structure of the sources as a graph $G_{\text{source}} = (V, E)$, where $V = \{\mathbf{Y}, \lambda_1, \lambda_2, \dots, \lambda_m\}$ (Figure 3). Specifically, if (λ_i, λ_j) is not an edge in G_{source} , this means that λ_i is independent of λ_j conditioned on \mathbf{Y} and the other source labels. Note that if G_{source} is unknown, it can be estimated using statistical techniques such as [2]. Importantly, we do not know anything about the strengths of the correlations in G_{source} , or the sources' accuracies.

Our overall goal is to apply the set of weak supervision sources $S = \{s_1, \dots, s_m\}$ to an unlabeled dataset \mathcal{X}_U consisting of n data points, then use the resulting weakly-labeled training set to supervise an *end model* $f_w: \mathcal{X} \mapsto \mathcal{Y}$ (Figure 1). This weakly-labeled training set will contain overlapping and conflicting labels, from sources with unknown accuracies and correlations. To handle this, we will learn a *label model* $P_\mu(\mathbf{Y}|\boldsymbol{\lambda})$, parameterized by a vector of source correlations and accuracies μ , which for each data point X takes as input the noisy labels $\boldsymbol{\lambda} = \{\lambda_1, \dots, \lambda_m\}$ and outputs a single probabilistic label vector $\tilde{\mathbf{Y}}$. Succinctly, given a user-provided tuple $(\mathcal{X}_U, S, G_{\text{source}}, G_{\text{task}})$, our key technical challenge is recovering the parameters μ without access to ground truth labels \mathbf{Y} .

Modeling Multi-Task Sources—To learn a label model over multi-task sources, we introduce sufficient statistics over the random variables in G_{source} . Let \mathcal{C} be the set of cliques in G_{source} , and define an indicator random variable for the event of a clique $C \in \mathcal{C}$ taking on a set of values y_C :

$$\psi(C, y_C) = \mathbb{1}\{\cap_{i \in C} V_i = (y_C)_i\},$$

where $(y_C)_i \in \mathcal{Y}_{\tau_i}$. We define $\psi(C) \in \{0, 1\}^{\prod_{i \in C} (|\mathcal{Y}_{\tau_i}| - 1)}$ as the vector of indicator random variables for all combinations of all but one of the labels emitted by each variable in clique C —thereby defining a minimal set of statistics—and define $\psi(\mathbf{C})$ accordingly for any set of cliques $\mathbf{C} \subseteq \mathcal{C}$. Then $\mu = \mathbb{E}[\psi(\mathcal{C})]$ is the vector of sufficient statistics for the label model we want to learn.

We work with two simplifying conditions in this section. First, we consider the setting where G_{source} is *triangulated* and has a junction tree representation with singleton separator sets. If this is not the case, edges can always be added to G_{source} to make this setting hold; otherwise, we describe how our approach can directly handle non-singleton separator sets in Appendix A.3.3.

Second, we use a simplified *class-conditional* model of the noisy labeling process, where we learn one accuracy parameter for each label value λ_i that each source s_j emits. This is equivalent to assuming that a source may have a different accuracy on each different class, but that if it emits a certain label incorrectly, it does so uniformly over the different true labels \mathbf{Y} . This is a more expressive model than the commonly considered one, where each source is modeled by a single accuracy parameter, e.g. in [11; 28], and in particular allows us to capture the *unipolar* setting considered later on. For further details, see Appendix A.3.4.

Our Approach—The chief technical difficulty in our problem is that we do not observe \mathbf{Y} . We overcome this by analyzing the covariance matrix of an observable subset of the cliques in G_{source} , leading to a matrix completion-style approach for recovering μ . We leverage two pieces of information: (i) the observability of *part of* $\mathbf{Cov} [\psi(\mathcal{C})]$, and (ii) a result from [22] which states that the inverse covariance matrix $\mathbf{Cov} [\psi(\mathcal{C})]^{-1}$ is structured according to G_{source} , i.e., if there is no edge between λ_i and λ_j in G_{source} , then the corresponding entries are 0.

We start by considering two disjoint subsets of \mathcal{C} : the set of observable cliques, $O \subseteq \mathcal{C}$ —i.e., those cliques not containing \mathbf{Y} —and the separator set cliques of the junction tree, $\mathcal{S} \subseteq \mathcal{C}$. In the setting we consider in this section, $\mathcal{S} = \{\mathbf{Y}\}$ (see Figure 3). We can then write the covariance matrix of the indicator variables for $O \cup \mathcal{S}$, $\mathbf{Cov} [\psi(O \cup \mathcal{S})]$, in block form, similar to [6], as:

$$\mathbf{Cov} [\psi(O \cup \mathcal{S})] \equiv \Sigma = \begin{bmatrix} \Sigma_O & \Sigma_{O\mathcal{S}} \\ \Sigma_{O\mathcal{S}}^T & \Sigma_{\mathcal{S}} \end{bmatrix} \quad (1)$$

and similarly define its inverse:

$$\mathbf{K} = \Sigma^{-1} = \begin{bmatrix} K_O & K_{O\mathcal{S}} \\ K_{O\mathcal{S}}^T & K_{\mathcal{S}} \end{bmatrix} \quad (2)$$

Here, $\Sigma_O = \mathbf{Cov} [\psi(O)] \in \mathbb{R}^{d_O \times d_O}$ is the observable block of Σ , where $d_O = \sum_{c \in O} \prod_{i \in c} (|\mathcal{Y}_i| - 1)$. Next, $\Sigma_{O\mathcal{S}} = \mathbf{Cov} [\psi(O), \psi(\mathcal{S})]$ is the unobserved block which is a function of μ , the label model parameters that we wish to recover. Finally, $\Sigma_{\mathcal{S}} = \mathbf{Cov} [\psi(\mathcal{S})] = \mathbf{Cov} [\psi(\mathbf{Y})]$ is a function of the class balance $P(\mathbf{Y})$.

We make two observations about $\Sigma_{\mathcal{S}}$. First, while the full form of $\Sigma_{\mathcal{S}}$ is the covariance of the $|\mathcal{Y}| - 1$ indicator variables for each individual value of \mathbf{Y} but one, given our simplified class-conditional label model, we in fact only need a single indicator variable for \mathbf{Y} (see Appendix A.3.4); thus, $\Sigma_{\mathcal{S}}$ is a scalar. Second, $\Sigma_{\mathcal{S}}$ is a function of the class balance $P(\mathbf{Y})$, which we assume is either known, or has been estimated according to the unsupervised approach we detail in Appendix A.3.5. Thus, given Σ_O and $\Sigma_{\mathcal{S}}$, our goal is to recover the vector $\Sigma_{O\mathcal{S}}$ from which we can recover μ .

Algorithm 1

Source Accuracy Estimation for Multi-Task Weak Supervision

Input: Observed labeling rates
 $\widehat{\mathbb{E}}[\psi(O)]$
 and covariance
 $\widehat{\Sigma}_O$;
 class balance
 $\widehat{\mathbb{E}}[\psi(\mathbf{Y})]$
 and variance Σ

\mathcal{S}
 ; correlation sparsity structure Ω
 $\hat{z} \leftarrow \operatorname{argmin}_z \|\hat{\Sigma}_o^{-1} + zz^T\|_{\Omega}$
 $\hat{c} \leftarrow \Sigma_s^{-1} (1 + \hat{z}^T \hat{\Sigma}_o \hat{z}), \hat{\Sigma}_{os} \leftarrow \hat{\Sigma}_o \hat{z} / \sqrt{\hat{c}}$
 $\hat{\mu}' \leftarrow \hat{\Sigma}_{os} + \hat{\mathbb{E}}[\psi(\mathbf{Y})] \hat{\mathbb{E}}[\psi(O)]$
return ExpandTied
 $(\hat{\mu}')$

Applying the block matrix inversion lemma, we have:

$$K_o = \Sigma_o^{-1} + c \Sigma_o^{-1} \Sigma_{os} \Sigma_{os}^T \Sigma_o^{-1}, \quad (3)$$

where $c = (\Sigma_s - \Sigma_{os}^T \Sigma_o^{-1} \Sigma_{os})^{-1} \in \mathbb{R}^+$. Let $z = \sqrt{c} \Sigma_o^{-1} \Sigma_{os}$; we can then express (3) as:

$$K_o = \Sigma_o^{-1} + zz^T \quad (4)$$

The right hand side of (4) consists of an empirically observable term, Σ_o^{-1} , and a rank-one term, zz^T , which we can solve for to directly recover μ . For the left hand side, we apply an extension of Corollary 1 from [22] (see Appendix A.3.2) to conclude that K_o has graph-structured sparsity, i.e., it has zeros determined by the structure of dependencies between the sources in G_{source} . This suggests an algorithmic approach of estimating z as a matrix completion problem in order to recover an estimate of μ (Algorithm 1). In more detail: let Ω be the set of indices (i, j) where $(K_o)_{i,j} = 0$, determined by G_{source} , yielding a system of equations,

$$0 = (\Sigma_o^{-1})_{i,j} + (zz^T)_{i,j} \text{ for } (i, j) \in \Omega, \quad (5)$$

which is now a matrix completion problem. Define $\|A\|_{\Omega}$ as the Frobenius norm of A with entries not in Ω set to zero; then we can rewrite (5) as $\|\Sigma_o^{-1} + zz^T\|_{\Omega} = 0$. We solve this equation to estimate z , and thereby recover Σ_{os} , from which we can directly recover the label model parameters μ algebraically.

Checking for Identifiability—A first question is: which dependency structures G_{source} lead to unique solutions for μ ? This question presents a stumbling block for users, who might attempt to use non-identifiable sets of correlated weak supervision sources.

We provide a simple, testable condition for identifiability. Let G_{inv} be the inverse graph of G_{source} ; note that Ω is the edge set of G_{inv} expanded to include all indicator random variables $\psi(\mathcal{C})$. Then, let M_{Ω} be a matrix with dimensions $|\Omega| \times d_o$ such that each row in M_{Ω} corresponds to a pair $(i, j) \in \Omega$ with 1's in positions i and j and 0's elsewhere.

Taking the log of the squared entries of (5), we get a system of linear equations $M_{\Omega} I = q_{\Omega}$, where $I_i = \log(z_i^2)$ and $q_{(i,j)} = \log\left(\left((\Sigma_o^{-1})_{i,j}\right)^2\right)$. Assuming we can solve this system (which we can always ensure by adding sources; see Appendix), we can uniquely recover the z_i^2 , meaning our model is identifiable *up to sign*.

Given estimates of the z_i^2 , we can see from (5) that the sign of a single z_j determines the sign of all other z_j reachable from z_j in G_{inv} . Thus to ensure a unique solution, we only need to pick a sign for each connected component in G_{inv} . In the case where the sources are assumed to be independent, e.g., [10; 38; 11], it suffices to make the assumption that the sources are *on average* non-adversarial; i.e., select the sign of the z_j that leads to higher average accuracies of the sources. Even a single source that is conditionally independent from all the other sources will cause G_{inv} to be fully connected, meaning we can use this symmetry breaking assumption in the majority of cases even with correlated sources. Otherwise, a sufficient condition is the standard one of assuming non-adversarial sources, i.e. that all sources have greater than random accuracy. For further details, see Appendix B.1.

Source Accuracy Estimation Algorithm—Now that we know when a set of sources with correlation structure G_{source} is identifiable, yielding a unique z , we can estimate the accuracies μ using Algorithm 1. We also use the function `ExpandTied`, which is a simple algebraic expansion of tied parameters according to the simplified class-conditional model used in this section; see Appendix A.3.4 for details. In Figure 4, we plot the performance of our algorithm on synthetic data, showing its scaling with the number of unlabeled data points n , the density of pairwise dependencies in G_{source} , and the runtime performance as compared to a prior Gibbs sampling-based approach. Next, we theoretically analyze the scaling of the error $\|\hat{\mu} - \mu^*\|$.

4.2 Theoretical Analysis: Scaling with Diverse Multi-Task Supervision

Our ultimate goal is to train an *end model* using the source labels, denoised and combined by the label model $\hat{\mu}$ we have estimated. We connect the generalization error of this end model to the estimation error of Algorithm 1, ultimately showing that the generalization error scales as $n^{-\frac{1}{2}}$, where n is the number of unlabeled data points. This key result establishes the same asymptotic scaling as traditionally supervised learning methods, but with respect to *unlabeled* data points.

Let $P_{\hat{\mu}}(\tilde{\mathbf{Y}}|\lambda)$ be the probabilistic label (i.e. distribution) predicted by our label model, given the source labels λ as input, which we compute using the estimated $\hat{\mu}$. We then train an *end* multi-task discriminative model $f_w: \mathcal{X} \mapsto \mathcal{Y}$ parameterized by w , by minimizing the expected loss with respect to the label model over n unlabeled data points. Let $l(w, X, \mathbf{Y}) = \frac{1}{T} \sum_{s=1}^T l_i(w, X, \mathbf{Y}_s)$ be a bounded multi-task loss function such that without loss of generality $l(w, X, \mathbf{Y}) \in [0, 1]$; then we minimize the empirical *noise aware loss*:

$$\hat{w} = \operatorname{argmin}_w \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\tilde{\mathbf{Y}} \sim P_{\hat{\mu}}(\cdot|\lambda)} [l(w, X_i, \tilde{\mathbf{Y}})], \quad (6)$$

and let \tilde{w} be the w that minimizes the true noise-aware loss. This minimization can be performed by standard methods and is not the focus of our paper; let the solution \hat{w} satisfy $\mathbb{E}[\|\hat{w} - \tilde{w}\|^2] \leq \gamma$. We make several assumptions, following [28]: (1) that for some label model parameters μ^* , sampling $(\lambda, \mathbf{Y}) \sim P_{\mu^*}(\cdot)$ is the same as sampling from the true

distribution, $(\mathbf{A}, \mathbf{Y}) \sim \mathcal{D}$; and (2) that the task labels Y_s are independent of the features of the end model given \mathbf{A} sampled from $P_{\mu^*}(\cdot)$, that is, the output of the optimal label model provides sufficient information to discern the true label. Then we have the following result:

Theorem 1—Let \tilde{w} minimize the expected noise aware loss, using weak supervision source parameters $\hat{\mu}$ estimated with Algorithm 1. Let \hat{w} minimize the empirical noise aware loss with $\mathbb{E}[\|\hat{w} - \tilde{w}\|^2] \leq \gamma$, $w^* = \min_w l(w, X, Y)$, and let the assumptions above hold. Then the generalization error is bounded by:

$$\mathbb{E}[l(\hat{w}, X, Y) - l(w^*, X, Y)] \leq \gamma + 4|\mathcal{Y}|\|\hat{\mu} - \mu^*\|.$$

Thus, to control the generalization error, we must control $\|\hat{\mu} - \mu^*\|$, which we do in Theorem 2:

Theorem 2—Let $\hat{\mu}$ be an estimate of μ^* produced by Algorithm 1 run over n unlabeled

data points. Let $a := \left(\frac{d_o}{\Sigma_s} + \left(\frac{d_o}{\Sigma_s} \right)^2 \lambda_{\max}(K_o) \right)^{\frac{1}{2}}$ and $b := \frac{\|\Sigma_o^{-1}\|^2}{(\Sigma_o^{-1})_{\min}}$. Then, we have:

$$\mathbb{E}[\|\hat{\mu} - \mu^*\|] \leq 16(r-1)d_o^2 \sqrt{\frac{32\pi}{n}} ab \sigma_{\max}(M_{\Omega}^{\dagger}) (3\sqrt{d_o} a \lambda_{\min}^{-1}(\Sigma_o) + 1) (\kappa(\Sigma_o) + \lambda_{\min}^{-1}(\Sigma_o)).$$

Interpreting the Bound—We briefly explain the key terms controlling the bound in Theorem 2; more detail is found in Appendix B. Our primary result is that the estimation error scales as $n^{-\frac{1}{2}}$. Next, $\sigma_{\max}(M_{\Omega}^{\dagger})$, the largest singular value of the pseudoinverse M_{Ω}^{\dagger} , has a deep connection to the density of the graph G_{inv} . The smaller this quantity, the more information we have about G_{inv} , and the easier it is to estimate the accuracies. Next, $\lambda_{\min}(\Sigma_o)$, the smallest eigenvalue of the observed covariance matrix, reflects the conditioning of Σ_o ; better conditioning yields easier estimation, and is roughly determined by how far away from random guessing the worst weak supervision source is, as well as how conditionally independent the sources are. $\lambda_{\max}(K_o)$, the largest eigenvalue of the upper-left block of the inverse covariance matrix, similarly reflects the overall conditioning of Σ . Finally, $(\Sigma_o^{-1})_{\min}$, the smallest entry of the inverse observed matrix, reflects the smallest non-zero correlation between source accuracies; distinguishing between small correlations and independent sources requires more samples.

4.3 Extensions: Abstentions & Unipolar Sources

We briefly highlight two extensions handled by our approach which we have found empirically critical: handling *abstentions*, and modeling *unipolar* sources.

Handling Abstentions.—One fundamental aspect of the weak supervision setting is that sources may abstain from labeling a data point entirely—that is, they may have incomplete and differing coverage [27; 10]. We can easily deal with this case by extending the coverage

ranges \mathcal{Y}_{τ_i} of the sources to include the vector of all zeros, $\vec{0}$, and we do so in the experiments.

Handling Unipolar Sources.—Finally, we highlight the fact that our approach models *class conditional* source accuracies, in particular motivated by the case we have frequently observed in practice of *unipolar* weak supervision sources, i.e., sources that each only label a single class or abstain. In practice, we find that users most commonly use such unipolar sources; for example, a common template for a heuristic-based weak supervision source over text is one that looks for a specific pattern, and if the pattern is present emits a specific label, else abstains. As compared to prior approaches that did not model class-conditional accuracies, e.g. [28], we show in our experiments that we can use our class-conditional modeling approach to yield an improvement of 2.8 points in accuracy.

5 Experiments

We validate our approach on three fine-grained classification problems—entity classification, relation classification, and document classification—where weak supervision sources are available at both coarser and finer-grained levels (e.g. as in Figure 2). We evaluate the predictive accuracy of end models supervised with training data produced by several approaches, finding that our approach outperforms traditional hand-labeled supervision by 20.2 points, a baseline majority vote weak supervision approach by 6.8 points, and a prior weak supervision denoising approach [28] that is not multi-task-aware by 4.1 points.

Datasets

Each dataset consists of a large (3k-63k) amount of unlabeled training data and a small (200–350) amount of labeled data which we refer to as the *development set*, which we use for (a) a traditional supervision baseline, and (b) for hyperparameter tuning of the end model (see Appendix C). The average number of weak supervision sources per task was 13, with sources expressed as Python functions, averaging 4 lines of code and comprising a mix of pattern matching heuristics, external knowledge base or dictionary lookups, and pre-trained models. In all three cases, we choose the decomposition into sub-tasks so as to align with weak supervision sources that are either available or natural to express.

Named Entity Recognition (NER): We represent a fine-grained named entity recognition problem—tagging entity mentions in text documents—as a hierarchy of three sub-tasks over the OntoNotes dataset [34]: $Y_1 \in \{\text{Person, Organization}\}$, $Y_2 \in \{\text{Businessperson, Other Person, } N/A\}$, $Y_3 \in \{\text{Company, Other Org, } N/A\}$, where again we use *N/A* to represent “not applicable”.

Relation Extraction (RE): We represent a relation extraction problem—classifying entity-entity relation mentions in text documents—as a hierarchy of six sub-tasks which either concern labeling the subject, object, or subject-object pair of a possible or *candidate* relation in the TACRED dataset [39]. For example, we might label a relation as having a Person subject, Location object, and Place-of-Residence relation type.

Medical Document Classification (Doc): We represent a radiology report triaging (i.e. document classification) problem from the OpenI dataset [25] as a hierarchy of three sub-tasks: $Y_1 \in \{\text{Acute, Non-Acute}\}$, $Y_2 \in \{\text{Urgent, Emergent, N/A}\}$, $Y_3 \in \{\text{Normal, Non-Urgent, N/A}\}$.

End Model Protocol

Our goal was to test the performance of a basic multi-task end model using training labels produced by various different approaches. We use an architecture consisting of a shared bidirectional LSTM input layer with pre-trained embeddings, shared linear intermediate layers, and a separate final linear layer (“task head”) for each task. Hyperparameters were selected with an initial search for each application (see Appendix), then fixed.

Core Validation

We compare the accuracy of the end multi-task model trained with labels from our approach versus those from three baseline approaches (Table 1):

- *Traditional Supervision [Gold (Dev)]:* We train the end model using the small hand-labeled development set.
- *Hierarchical Majority Vote [MV]:* We use a hierarchical majority vote of the weak supervision source labels: i.e. for each data point, for each task we take the majority vote and proceed down the task tree accordingly. This procedure can be thought of as a hard decision tree, or a cascade of if-then statements as in a rule-based approach.
- *Data Programming [DP]:* We model each task separately using the data programming approach for denoising weak supervision [27].

In all settings, we used the same end model architecture as described above. Note that while we choose to model these problems as consisting of multiple sub-tasks, we evaluate with respect to the broad primary task of fine-grained classification (for subtask-specific scores, see Appendix). We observe in Table 1 that our approach of leveraging multi-granularity weak supervision leads to large gains—20.2 points over traditional supervision with the development set, 6.8 points over hierarchical majority vote, and 4.1 points over data programming.

Ablations

We examine individual factors:

Unipolar Correction: Modeling unipolar sources (Sec 4.3), which we find to be especially common when fine-grained tasks are involved, leads to an average gain of 2.8 points of accuracy in MeTaL performance.

Joint Task Modeling: Next, we use our algorithm to estimate the accuracies of sources for each task separately, to observe the empirical impact of modeling the multi-task setting jointly as proposed. We see average gains of 1.3 points in accuracy (see Appendix).

End Model Generalization: Though not possible in many settings, in our experiments we can directly apply the label model to make predictions. In Table 6, we show that the end model improves performance by an average 3.4 points in accuracy, validating that the models trained do indeed learn to generalize beyond the provided weak supervision. Moreover, the largest generalization gain of 7 points in accuracy came from the dataset with the most available unlabeled data ($n=63k$), demonstrating scaling consistent with the predictions of our theory (Fig. 5). This ability to leverage additional unlabeled data and more sophisticated end models are key advantages of the weak supervision approach in practice.

6 Conclusion

We presented MeTaL, a framework for training models with weak supervision from diverse, *multi-task* sources having different granularities, accuracies, and correlations. We tackle the core challenge of recovering the unknown source accuracies via a scalable matrix completion-style algorithm, introduce theoretical bounds characterizing the key scaling with respect to unlabeled data, and demonstrate empirical gains on real-world datasets. In future work, we hope to learn the task relationship structure and cover a broader range of settings where labeled training data is a bottleneck.

Acknowledgements

We gratefully acknowledge the support of DARPA under Nos. FA87501720095 (D3M) and FA86501827865 (SDH), NIH under No. N000141712266 (Mobilize), NSF under Nos. CCF1763315 (Beyond Sparsity), CCF1563078 (Volume to Velocity), and DGE-114747 (NSF GRF), ONR under No. N000141712266 (Unifying Weak Supervision), the Moore Foundation, NXP, Xilinx, LETI-CEA, Intel, Google, NEC, Toshiba, TSMC, ARM, Hitachi, BASF, Accenture, Ericsson, Qualcomm, Analog Devices, the Okawa Foundation, and American Family Insurance, the Stanford Interdisciplinary Graduate and Bio-X fellowships, the Intelligence Community Postdoctoral Fellowship, and members of the Stanford DAWN project: Intel, Microsoft, Teradata, Facebook, Google, Ant Financial, NEC, SAP, and VMWare. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views, policies, or endorsements, either expressed or implied, of DARPA, NIH, ONR, or the U.S. Government.

A: Problem Setup & Modeling Approach

In Section A, we review our problem setup and modeling approach in more detail, and for more general settings than in the body. In Section B, we provide an overview, additional interpretation, and the proofs of our main theoretical results. Finally, in Section C, we go over additional details of our experimental setup.

We begin in Section A.1 with a glossary of the symbols and notation used throughout this paper. Then, in Section A.2 we present the setup of our multi-task weak supervision problem, and in Section A.3 we present our approach for modeling multi-task weak supervision, and the matrix completion-style algorithm used to estimate the model parameters. Finally, in Section A.4, we present in more detail the subcase of hierarchical tasks considered in the main body of the paper.

A.1 Glossary of Symbols

Table 2:

Glossary of variables and symbols used in this paper.

Symbol	Used for
X	Data point, $X \in \mathcal{X}$
n	Number of data points
Y_s	Label for one of the t classification tasks, $Y_s \in \{1, \dots, k_s\}$
t	Number of tasks
\mathbf{Y}	Vector of task labels $\mathbf{Y} = [Y_1, Y_2, \dots, Y_t]^T$
r	Cardinality of the output space, $r = \mathcal{Y} $
G_{task}	Task structure graph
\mathcal{Y}	Output space of allowable task labels defined by G_{task} , $\mathbf{Y} \in \mathcal{Y}$
\mathcal{D}	Distribution from which we assume (X, \mathbf{Y}) data points are sampled i.i.d.
s_j	Weak supervision source, a function mapping X to a label vector
λ_i	Label vector $\lambda_i \in \mathcal{Y}$ output by the i th source for X
m	Number of sources
λ	$m \times t$ matrix of labels output by the m sources for X
\mathcal{Y}_0	Source output space, which is \mathcal{Y} augmented to include elements set to zero
τ_i	Coverage set of λ_i - the tasks s_j gives non-zero labels to; for convenience, $\tau_0 = \{1, \dots, t\}$
\mathcal{Y}_{τ_i}	The output space for λ_i given coverage set τ_i
$\mathcal{Y}_{\tau_i}^{\min}$	The output space \mathcal{Y}_{τ_i} with all but the first value, for defining a minimal set of statistics
G_{source}	Source dependency graph, $G_{\text{source}} = (V, E)$, $V = \{\mathbf{Y}, \lambda_1, \dots, \lambda_m\}$
\mathcal{C}	Cliqueset (maximal and non-maximal) of G_{source}
$\tilde{\mathcal{C}}, \mathcal{S}$	The maximal cliques (nodes) and separator sets of the junction tree over G_{source}
$\mathcal{I}(C, y_C)$	The indicator variable for the variables in clique $C \in \mathcal{C}$ taking on values y_C , $(y_C)_i \in \mathcal{Y}_{\tau_i}$
μ	The parameters of our label model we aim to estimate; $\mu = \mathbb{E}[\psi]$

Symbol	Used for
O	The set of observable cliques, i.e. those corresponding to cliques without \mathbf{Y}
Σ	Generalized covariance matrix of $O \subseteq \mathcal{S}$ $\Sigma \equiv \text{Cov} [\Psi(O \subseteq \mathcal{S})]$
K	The inverse generalized covariance matrix $K = \Sigma^{-1}$
$d_O, d_{\mathcal{S}}$	The dimensions of O and \mathcal{S} respectively
G_{aug}	The augmented source dependencies graph $G_{\text{aug}} = (\mathcal{Y}, E_{\text{aug}})$
Ω	The edge set of the inverse graph of G_{aug}
P	Diagonal matrix of class prior probabilities, $P(\mathbf{Y})$
$P_{\mu}(\mathbf{Y}, \boldsymbol{\lambda})$	The <i>label model</i> parameterized by μ
$\tilde{\mathbf{Y}}$	The probabilistic training label, i.e. $P_{\mu}(\mathbf{Y} \boldsymbol{\lambda})$
$f_w(X)$	The <i>end model</i> trained using $(X, \tilde{\mathbf{Y}})$

A.2 Problem Setup

Let $X \in \mathcal{X}$ be a data point and $\mathbf{Y} = [Y_1, Y_2, \dots, Y_t]^T$ be a vector of *task labels* corresponding to t tasks. We consider categorical task labels, $Y_i \in \{1, \dots, k_i\}$ for $i \in \{1, \dots, t\}$. We assume (X, \mathbf{Y}) pairs are sampled i.i.d. from distribution \mathcal{D} ; to keep the notation manageable, we do not place subscripts on the sample tuples.

Task Structure

The tasks are related by a *task graph* G_{task} . Here, we consider schemas expressing logical relationships between tasks, which thus define *feasible sets* of label vectors \mathcal{Y} , such that $\mathbf{Y} \in \mathcal{Y}$. We let $r = |\mathcal{Y}|$ be the number of feasible task vectors. In section A.4, we consider the particular subcase of a *hierarchical* task structure as used in the experiments section of the paper.

Multi-Task Sources

We now consider *multi-task* weak supervision sources $s_j \in \mathcal{S}$, which represent noisy and potentially incomplete sources of labels, which have unknown accuracies and correlations. Each source s_j outputs label vectors λ_i , which contain non-zero labels for *some* of the tasks, such that λ_i is in the feasible set \mathcal{Y} but potentially with some elements set to zero, denoting a null vote or abstention for that task. Let \mathcal{Y}_0 denote this extended set which includes certain task labels set to zero.

We also assume that each source has a fixed *task coverage set* τ_j , such that $(\lambda_i)_s = 0$ for $s \in \tau_j$ and $(\lambda_i)_s = 0$ for $s \notin \tau_j$; let $\mathcal{Y}_{\tau_j} \subseteq \mathcal{Y}_0$ be the range of λ_i given coverage set τ_j . For convenience, we let $\tau_0 = \{1, \dots, t\}$ so that $\mathcal{Y}_{\tau_0} = \mathcal{Y}$. The intuitive idea of the task coverage set is that some labelers may choose not to label certain tasks; Example 2 illustrates this notion.

Note that sources can also *abstain* for a data point, meaning they emit no label (which we denote with a symbol $\vec{0}$); we include this in \mathcal{Y}_{τ_j} . Thus we have $s_j: \mathcal{X} \mapsto \mathcal{Y}_{\tau_j}$ where, again, λ_j denotes the output of the function s_j .

Problem Statement

Our overall goal is to use the noisy or *weak, multi-task* supervision from the set of m sources, $\mathcal{S} = \{s_1, \dots, s_m\}$, applied to an unlabeled dataset \mathcal{X}_U consisting of n data points, to supervise an *end model* $f_w: \mathcal{X} \mapsto \mathcal{Y}$. Since the sources have unknown accuracies, and will generally output noisy and incomplete labels that will overlap and conflict, our intermediate goal is to learn a *label model* $P_\mu: \lambda \mapsto [0, 1]^{\mathcal{Y}}$ which takes as input the source labels and outputs a set of probabilistic label vectors, $\tilde{\mathbf{Y}}$, for each X , which can then be used to train the end model. Succinctly, given a user-provided tuple $(\mathcal{X}_U, \mathcal{S}, G_{\text{source}}, G_{\text{task}})$, our goal is to recover the parameters μ .

The key technical challenge in this approach then consists of learning the parameters of this label model—corresponding to the conditional accuracies of the sources (and, for technical reasons we shall shortly explain, cliques of correlated sources)—given that *we do not have access to the ground truth labels* \mathbf{Y} . We discuss our approach to overcoming this core technical challenge in the subsequent section.

A.3 Our Approach: Modeling Multi-Task Sources

Our goal is to estimate the parameters μ of a *label model* that produces probabilistic training labels given the observed source outputs, $\tilde{\mathbf{Y}} = P_\mu(\mathbf{Y}|\lambda)$, *without access to the ground truth labels* \mathbf{Y} . We do this in three steps:

1. We start by defining a graphical model over the weak supervision source outputs and the true (latent) variable \mathbf{Y} , $(\lambda_1, \dots, \lambda_m, \mathbf{Y})$, using the conditional independence structure G_{source} between the sources.
2. Next, we analyze the *generalized covariance matrix* Σ (following Loh & Wainwright [22]), which is defined over binary indicator variables for each value of each clique (or specific subsets of cliques) in G_{source} . We consider two specific subsets of the cliques in G_{source} , the observable cliques O and the separator sets \mathcal{S} , such that:

$$\Sigma = \begin{bmatrix} \Sigma_O & \Sigma_{O\mathcal{S}} \\ \Sigma_{O\mathcal{S}}^T & \Sigma_{\mathcal{S}} \end{bmatrix} \quad \Sigma^{-1} = K = \begin{bmatrix} K_O & K_{O\mathcal{S}} \\ K_{O\mathcal{S}}^T & K_{\mathcal{S}} \end{bmatrix},$$

where Σ_O is the block of Σ that we can observe, and $\Sigma_{O\mathcal{S}}$ is a function of μ , the parameters (corresponding to source and clique accuracies) we wish to recover. We then apply a result by Loh and Wainwright [22] to establish the sparsity pattern of $K = \Sigma^{-1}$. This allows us to apply the block-matrix inversion lemma to reformulate our problem as solving a matrix completion-style objective.

3. Finally, we describe how to recover the class balance $P(\mathbf{Y})$; with this and the estimate of μ , we then describe how to compute the probabilistic training labels $\tilde{Y} = P_\mu(\mathbf{Y}|\lambda)$.

We start by focusing on the setting where G_{source} has a junction tree with singleton separator sets; we note that a version of G_{source} where this holds can always be formed by adding edges to the graph. We then discuss how to handle graphs with non-singleton separator sets, and finally describe different settings where our problem reduces to rank-one matrix completion. In Section B, we introduce theoretical results for the resulting model and provide our model estimation strategy.



Figure 7:

A simple example of a weak supervision source dependency graph G_{source} (left) and its junction tree representation (right). Here \mathbf{Y} is as a vector-valued variable with a feasible set of values, $\mathbf{Y} \in |\mathcal{Y}|$, and the output of sources 1 and 2 are modeled as dependent conditioned on \mathbf{Y} . This results in a junction tree with singleton separator sets \mathbf{Y} . Here, the observable cliques are $O = \{\lambda_1, \lambda_2, \lambda_3, \lambda_4, \{\lambda_1, \lambda_2\}\} \subset \mathcal{C}$.

A.3.1 Defining a Multi-Task Source Model

We consider a model $G_{\text{source}} = (V, E)$, where $V = \{\mathbf{Y}, \lambda_1, \dots, \lambda_m\}$, and E consists of pairwise interactions (i.e. we consider an Ising model, or equivalently, a graph rather than a hypergraph of correlations). We assume that G_{source} is provided by the user. However, if G_{source} is unknown, there are various techniques for estimating it statistically [2] or even from static analysis if the sources are heuristic functions [33]. We provide an example G_{source} with singleton separator sets in Figure 7.

Augmented Sufficient Statistics—Finally, we extend the random variables in V by defining a matrix of indicator statistics over all cliques in G_{source} , in order to estimate all the parameters needed for our label model P_μ . We assume that the provided G_{source} is *chordal*, meaning it has no chordless cycles of length greater than three; if not, the graph can easily be *triangulated* to satisfy this property, in which case we work with this augmented version.

Let \mathcal{C} be the set of maximal and non-maximal cliques in the chordal graph G_{source} . We start by defining a binary indicator random variable for the event of a clique $C \in \mathcal{C}$ in the graph $G_{\text{source}} = (V, E)$ taking on a set of values y_C :

$$\psi(C, y_C) = \mathbb{1}\{\cap_{i \in C} V_i = (y_C)_i\},$$

where $(y_c)_i \in \mathcal{Y}_{\tau_i}^{\min}$ and $\mathcal{Y}_{\tau_i}^{\min}$ contains all but one values of \mathcal{Y}_{τ_i} thereby leading to a minimal set of statistics. Note that in our notation, $V_0 = \mathbf{Y}$, $\mathcal{Y}_{\tau_0} = \mathcal{Y}$, and $V_{i>0} = \lambda_i$. Accordingly, we define $\psi(C) \in \{0, 1\}^{\prod_{i \in C} (|\mathcal{Y}_{\tau_i}| - 1)}$ as the vector of indicator random variables for all combinations of all but one of the labels emitted by each variable in clique C , and define $\Psi(C)$ accordingly for any set of cliques $C \subseteq \mathcal{C}$. Then $\mu = \mathbb{E}[\psi(\mathcal{C})]$ is the vector of sufficient statistics for the label model we want to learn. Our model estimation goal is now stated simply: we wish to estimate μ , *without access to the ground truth labels \mathbf{Y}* .

A.3.2 Model Estimation without Ground Truth Using Inverse Covariance Structure

Our goal is to estimate $\mu = \mathbb{E}[\psi(\mathcal{C})]$; this, along with the class balance $P(\mathbf{Y})$ (which we assume we know, or else estimate using the approach in Section A.3.5), is sufficient information to compute $P_{\mu}(\mathbf{Y}|\boldsymbol{\lambda})$. If we had access to a large enough set of ground truth labels \mathbf{Y} , we could simply take the empirical expectation $\hat{\mathbb{E}}[\psi]$; however in our setting we cannot directly observe this. Instead, we proceed by analyzing a sub-block of the covariance matrix of $\Psi(\mathcal{C})$, which corresponds to the *generalized covariance matrix* of our graphical model as in [22], and leverage two key pieces of information:

- A sub-block of this generalized covariance matrix is observable, and
- By a simple extension of Corollary 1 in [22], we know the sparsity structure of the inverse generalized covariance matrix Σ^{-1} , i.e. we know that it will have elements equal to zero according to the structure of G_{source} .

Since G_{source} is triangulated, it admits a *junction tree* representation [19], which has maximal cliques (nodes) $\tilde{\mathcal{C}}$ and separator sets \mathcal{S} . Note that we follow the convention that \mathcal{S} includes the full powerset of separator set cliques, i.e. all subset cliques of separator set cliques are also included in \mathcal{S} . We proceed by considering two specific subsets of the cliques of our graphical model G_{source} : those that are observable (i.e. not containing \mathbf{Y}), $O = \{C | \mathbf{Y} \notin C, C \in \mathcal{C}\}$, and the set of separator set cliques (which will always contain \mathbf{Y} , and thus be unobservable).

For simplicity of exposition, we start by considering graphs G_{source} which have singleton separator sets; given our graph structure, this means that $\mathcal{S} = \{\{\mathbf{Y}\}\}$. Note that in general we will write single-element sets without braces when their type is obvious from context, so we have $\mathcal{S} = \{\mathbf{Y}\}$. Intuitively, this corresponds to models where weak supervision sources are correlated in fully-connected clusters, corresponding to real-world settings in which sources are correlated due to shared data sources, code, or heuristics. However, we can always either (i) add edges to G_{source} such that this is the case, or (ii) extend our approach to many settings where G_{source} does not have singleton separator sets (see Section A.3.3).

In this singleton separator set setting of $\mathcal{S} = \{\mathbf{Y}\}$, we now have:

$$O = \{C | \mathbf{Y} \notin C, C \in \mathcal{C}\} \quad \mathcal{S} = \{\mathbf{Y}\}.$$

where $\Psi(O)$ and $\Psi(\mathbf{Y})$ are the corresponding vectors of minimal indicator variables. We define corresponding dimensions d_O and $d_{\mathcal{S}}$:

$$d_o = \sum_{C \in \mathcal{O}} \prod_{i \in C} (|\mathcal{Y}_{\tau_i}| - 1) \quad d_s = r - 1.$$

We now decompose the generalized covariance matrix and its inverse as:

$$\mathbf{Cov}[\psi(\mathcal{O} \cup \mathcal{S})] \equiv \Sigma = \begin{bmatrix} \Sigma_o & \Sigma_{os} \\ \Sigma_{os}^T & \Sigma_s \end{bmatrix} \quad \Sigma^{-1} = K = \begin{bmatrix} K_o & K_{os} \\ K_{os}^T & K_s \end{bmatrix}, \quad (7)$$

This is similar to the form used in [6], but with several important differences: we consider discrete (rather than Gaussian) random variables and have additional knowledge of the graph structure. Here, Σ_o is the observable block of the generalized covariance matrix Σ , and Σ_{os} is the unobserved block which is a function of μ , the parameters (corresponding to source and source clique accuracies) we wish to recover. Note that with the singleton separator sets we are considering, Σ_s is a function of the class balance $P(\mathbf{Y})$, which we assume is either known, or has been estimated according to the unsupervised approach we detail in Section A.3.5. Therefore, we assume that Σ_s is also known. Concretely then, our goal is to recover Σ_{os} given Σ_o, Σ_s .

We start by applying the block matrix inversion lemma to get the equation:

$$K_o = \Sigma_o^{-1} + \Sigma_o^{-1} \Sigma_{os} (\Sigma_s - \Sigma_{os}^T \Sigma_o^{-1} \Sigma_{os})^{-1} \Sigma_{os}^T \Sigma_o^{-1}. \quad (8)$$

Next, let $JJ^T = (\Sigma_s - \Sigma_{os}^T \Sigma_o^{-1} \Sigma_{os})^{-1}$. We justify this decomposition by showing that this term is positive semidefinite. We start by applying the Woodbury matrix inversion lemma:

$$(\Sigma_s - \Sigma_{os}^T \Sigma_o^{-1} \Sigma_{os})^{-1} = \Sigma_s^{-1} + \Sigma_s^{-1} \Sigma_{os}^T (\Sigma_o + \Sigma_{os} \Sigma_s^{-1} \Sigma_{os}^T)^{-1} \Sigma_{os} \Sigma_s^{-1}. \quad (9)$$

Now, note that Σ_o and Σ_s are both covariance matrices themselves and are therefore PSD. Furthermore, from [22] we know that Σ^{-1} must exist, which implies that Σ_o and Σ_s are invertible (and thus in fact positive definite). Therefore we also have that $\Sigma_{os} \Sigma_s^{-1} \Sigma_{os}^T > 0 \Rightarrow (\Sigma_o + \Sigma_{os} \Sigma_s^{-1} \Sigma_{os}^T)^{-1} > 0$, and therefore (9) is positive definite, and can therefore always be expressed as JJ^T for some J . Therefore, we can write (8) as:

$$K_o = \Sigma_o^{-1} + \Sigma_o^{-1} \Sigma_{os} J J^T \Sigma_{os}^T \Sigma_o^{-1}.$$

Finally, define $Z = \Sigma_o^{-1} \Sigma_{os} J$; we then have:

$$K_o = \Sigma_o^{-1} + Z Z^T. \quad (10)$$

Note that $Z \in \mathbb{R}^{d_o \times d_H}$, where $d_H = r - 1$, and therefore ZZ^T is a rank- $(r - 1)$ matrix. Therefore, we now have a form (10) that appears close to being a matrix completion-style problem. We complete the connection by leveraging the known sparsity structure of K_O .

Define $G_{\text{aug}} = (\mathcal{Y}, E_{\text{aug}})$ to be the augmented version of our graph G_{source} . In other words, let $i = (C_1, y_{C_1})$ and $j = (C_2, y_{C_2})$ according to the indexing scheme of our augmented indicator variables; then, $(i, j) \in E_{\text{aug}}$ if C_1, C_2 are subsets of the same maximal clique in G_{source} . Then, let $G_{\text{inv-aug}} = (\mathcal{Y}, \Omega)$ be the inverse graph of G_{aug} , such that $(i, j) \in E_{\text{aug}} \iff (i, j) \notin \Omega$ and vice-versa.

We start with a result that extends Corollary 1 in Loh & Wainwright [22] to our specific setting where we consider a set of the variables that contains all observable cliques, O , and all separator sets \mathcal{S} (note that this result holds for all \mathcal{S} , not just $\mathcal{S} = \{\mathbf{Y}\}$):

Corollary 1—*Let $U = O \cup \mathcal{S}$. Let Σ_U be the generalized covariance matrix for U . Then $(\Sigma_U^{-1})_{i,j} = 0$ whenever i, j correspond to cliques C_1, C_2 respectively such that C_1, C_2 are not subsets of the same maximal clique.*

Proof: We partition the cliques \mathcal{C} into two sets, U and $W = \mathcal{C} \setminus U$. Let Σ be the full generalized covariance matrix (i.e. including all maximal and non-maximal cliques) and $\Gamma = \Sigma^{-1}$. Thus we have:

$$\Sigma = \begin{bmatrix} \Sigma_U & \Sigma_{UW} \\ \Sigma_{UW}^T & \Sigma_W \end{bmatrix} \quad \Sigma^{-1} = \Gamma = \begin{bmatrix} K_U & K_{UW} \\ K_{UW}^T & K_W \end{bmatrix}.$$

By the block matrix inversion lemma we have:

$$\Sigma_U^{-1} = K_U - K_{UW} K_W^{-1} K_{UW}^T.$$

We now follow the proof structure of Corollary 1 of [22]. We know K_U is graph structured by Theorem 1 of [22]. Next, using the same argument as in the proof of Corollary 1 of [22], we know that K_W , and therefore K_W^{-1} , is block-diagonal. Intuitively, because the set U contains all of the separator set cliques, and due to the running intersection property of a junction tree, each clique in W belongs to precisely one maximal clique-leading to block diagonal structure of K_W . We thus need only to show that the following quantity is zero for two cliques C_i, C_j that are not subsets of the same maximal clique, with corresponding indices i, j :

$$(K_{UW} K_W^{-1} K_{UW}^T)_{i,j} = \sum_B (K_{UW})_{i,B} (K_W^{-1})_{B,B} (K_{UW}^T)_{B,j}$$

where B are the indices corresponding to the blocks in K_W^{-1} , which correspond to maximal cliques. Our argument follows again as in Corollary 1 of [22]: since U contains the separator sets, if the two cliques C_1, C_2 are not subsets of the same maximal clique, then for each B , either $(K_{UW})_{i,B}$ or $(K_{UW}^T)_{B,j}$ must be zero, completing the proof.

Now, by Corollary 1, we know that $K_{i,j} = 0$ if $(i, j) \in \Omega$. Let A_Ω denote a matrix A with all entries $(i, j) \notin \Omega$ masked to zero. Then, we have:

$$(\Sigma_o^{-1})_\Omega + (ZZ^T)_\Omega = 0. \quad (11)$$

Thus, given the dependency graph G_{source} , we can solve for Z as a rank- $(r-1)$ matrix completion problem, with mask Ω . Defining the semi-norm $\|A\|_\Omega = \|A_\Omega\|_F$, we can solve:

$$\hat{Z} = \operatorname{argmin}_Z \|\Sigma_o^{-1} + ZZ^T\|_\Omega. \quad (12)$$

Now, we have an estimate of Z . Note that at this point, we can only recover Z up to orthogonal transformations. We proceed by considering a reduced rank-one model, detailed in Section A.3.4, and in Section B.1 establish concrete conditions under which this model is uniquely identifiable.

We denote this rank-one setting by switching to writing Z as $z \in \mathbb{R}^{d_o \times 1}$, in which case we now have:

$$\hat{z} = \operatorname{argmin}_z \|\Sigma_o^{-1} + zz^T\|_\Omega. \quad (13)$$

Once we have recovered z uniquely (see Section B.1), we next need to recover $\Sigma_{o,s} = c^{-1/2} \Sigma_{oz}$. We use the fact that $c = \Sigma_s^{-1} (1 + z^T \Sigma_{oz})$, which we can confirm explicitly below, starting from the definition of c :

$$\begin{aligned} c &= (\Sigma_s - \Sigma_{os}^T \Sigma_o^{-1} \Sigma_{os})^{-1} \\ &= \left(\Sigma_s - \left(c^{-1/2} \Sigma_{oz} \right)^T \Sigma_o^{-1} \left(c^{-1/2} \Sigma_{oz} \right) \right)^{-1} \\ &= (\Sigma_s - c^{-1} z^T \Sigma_{oz})^{-1} \end{aligned}$$

$$\Rightarrow c^{-1} = \Sigma_s - c^{-1} z^T \Sigma_{oz}$$

$$\Rightarrow c^{-1} (1 + z^T \Sigma_{oz}) = \Sigma_s$$

$$\Rightarrow c = \Sigma_s^{-1} (1 + z^T \Sigma_{oz})$$

Thus, we can directly recover an estimate of $\Sigma_{o,s}$ from the observed Σ_o , known Σ_s , and estimated z . Finally, we have:

$$\Sigma_{o_s} + \mathbb{E}[\psi(O)]\mathbb{E}[\psi(\mathcal{S})]^T = \mathbb{E}[\psi(O)\psi(\mathcal{S})^T]. \quad (14)$$

Here, we can clearly observe $\mathbb{E}[\psi(O)]$, and given that we know the class balance $P(\mathbf{Y})$, we also have $\mathbb{E}[\psi(\mathcal{S})]$ therefore we can compute $\mathbb{E}[\psi(O)\psi(\mathcal{S})^T]$. Our goal now is to recover the columns $\mathbb{E}[\psi(O)\psi(\mathbf{Y}_i)]$, which together make up μ ; we can do this based on the constraints of our rank-one model (Section A.3.4), thus recovering an estimate of μ , which given the uniqueness of \hat{z} (Section B.1) is also unique. The overall procedure is described in the main body, in Algorithm 1.

A.3.3 Handling Non-Singleton Separator Sets

Now, we consider the setting where G_{source} has arbitrary separator sets. Let $d_s = \sum_{\mathcal{S} \in \mathcal{S}} \prod_{i \in \mathcal{S}} (|\mathcal{Y}_i| - 1)$. We see that we could solve this using our standard approach—this time, involving a rank- d_s matrix completion problem—except for the fact that we do not know $\Sigma_{\mathcal{S}}$, as it now involves terms besides the class balance.

Note first of all that we can always add edges between sources to G_{source} such that it has singleton separator sets (intuitively, this consists of “completing the clusters”), and as long as our problem is still identifiable (see Section B.1), we can simply solve this instance as above.

Instead, we can also take a multi-step approach, wherein we first consider one or more subgraphs of G_{source} that contain only singleton separator sets, and contain the cliques in \mathcal{S} . We can then solve this problem as before, which then gives us the needed information to identify the elements of $\Sigma_{\mathcal{S}}$ in our full problem, which we can then solve. In particular, we see that this multi-step approach is possible whenever the graph G_{source} has at least three components that are disconnected except for through \mathbf{Y} .

A.3.4 Rank-One Settings

We now consider settings where we can estimate the parameters of our label model, μ , involving only a rank-one matrix completion problem.

First, in the simplest setting of a single-task problem with binary class variable, $\mathbf{Y} \in \{0,1\}$ and G_{source} with singleton separator sets, $d_H = r - 1 = 1$ and our problem is directly a rank-one instance.

Next, we consider the setting of general \mathbf{Y} , with $|\mathcal{Y}| = r$ and G_{source} with singleton separator sets. By default, our problem now involves a rank- $(r - 1)$ matrix completion problem. However, we can reduce this to involving only a rank-one matrix completion problem by adding one simplifying assumption to our model: namely, that sources emit different incorrect labels with uniform conditional probability. Concretely, we add the assumption that:

$$(\lambda_c)_i = \mathbf{Y} \Leftrightarrow (\lambda'_c)_i = \mathbf{Y} \quad \forall i \in C \Rightarrow P(\lambda_c | \mathbf{Y}) = P(\lambda'_c | \mathbf{Y}) \quad (15)$$

Note that this is the same assumption as in the main body, but expressed more explicitly with respect to a clique C . For example, under this assumption, $P(\lambda_i = y' | \mathbf{Y} = y)$ is the same for all y' such that $y' \neq y$. As another example, $P(\lambda_i = y, \lambda_j = y' | \mathbf{Y} = y)$ is the same for all y' such that $y' \neq y$. Intuitively, under this commonly-used model, we are not modeling the different class-wise errors a source makes, but rather just whether it is correct or not given the correctness of other sources it is correlated with. The idea then is that with assumption (15) even though $|H| = r - 1$ (and thus $\Sigma O\mathcal{S}$ has $r - 1$ columns), we only actually need to solve for a single parameter per element of O .

We can operationalize this by forming a new graph with a binarized version of \mathbf{Y} , $\mathbf{Y}_B \in \{0,1\}$, such that the r classes are mapped to either 0 or 1. We see that this new variable still results in the same structure of dependencies G_{source} , and still allows us to recover the parameters α_y (and thus μ). We now have:

$$\mathcal{S} = \{\mathbf{Y}_B\}$$

We now solve in the same rank-one way as in the binary \mathbf{Y} case. Now, for singleton cliques, $\{\lambda_i, \mathbf{Y}\}$, given that we know $P(\mathbf{Y})$, we can directly recover $P(\lambda_i = y' | \mathbf{Y} = y')$ for all y' , given our simplified model.

For non-singleton cliques $\{\lambda_C, \mathbf{Y}\}$, note that we can directly recover $P(\cap_{i \in C} \lambda_i = y | \mathbf{Y} = y')$ in the exact same way. From these, computed for all cliques, we can then recover any probability in our model. For example, for $y' \neq y$:

$$\begin{aligned} P(\lambda_i = y, \lambda_j = y' | \mathbf{Y} = y) &= P(\lambda_i = y | \mathbf{Y} = y) - \sum_{y'' \neq y'} P(\lambda_i = y, \lambda_j = y'' | \mathbf{Y} = y) \\ &= P(\lambda_i = y | \mathbf{Y} = y) - P(\lambda_i = y, \lambda_j = y | \mathbf{Y} = y) - \\ &\quad \times (r - 2)P(\lambda_i = y, \lambda_j = y' | \mathbf{Y} = y) \\ &= > P(\lambda_i = y, \lambda_j = y' | \mathbf{Y} = y) = \frac{1}{r - 1} (P(\lambda_i = y | \mathbf{Y} = y) - P(\lambda_i = y, \lambda_j = y | \mathbf{Y} = y)). \end{aligned}$$

In this way, we can recover all of the parameters μ while only involving a rank-one matrix completion problem. Note that this also suggests a way to solve for the more general model, i.e. without (15), using a hierarchical classification approach.

A.3.5 Recovering the Class Balance P & Computing $P(\mathbf{Y}|\lambda)$

We now turn to the task of recovering the class balance $P(\mathbf{Y})$, for $\mathbf{Y} \in \mathcal{Y}$. In many practical settings, $P(\mathbf{Y})$ can be estimated from a small labeled sample, or may be known in advance. However here, we consider using a subset of conditionally independent sources, s_1, \dots, s_k to estimate $P(\mathbf{Y})$. We note first of all that simply taking the majority vote of these sources is a biased estimator.

Instead, we consider a simplified version of the matrix completion-based approach taken so far. Here, we consider a subset of the sources s_1, \dots, s_k such that they are conditionally

independent given G_{Source} , i.e. $\lambda_i \perp \lambda_j | \mathbf{Y}$, and consider only the unary indicator statistics. Denote the vector of these unary indicator statistics over the conditionally independent subset of sources as ϕ , and let the observed overlaps matrix between sources i and j be $A_{i,j} = \mathbb{E}[\phi_i \phi_j^T]$. Note that due to the conditional independence of λ_i and λ_j , for any k, l we have:

$$\begin{aligned} (A_{i,j})_{k,l} &= \mathbb{E}[(\phi_i)_k (\phi_j)_l] \\ &= P(\lambda_i = y_k, \lambda_j = y_l) \\ &= \sum_{y \in \mathcal{Y}} P(\lambda_i = y_k, \lambda_j = y_l | \mathbf{Y} = y) P(\mathbf{Y} = y) \\ &= \sum_{y \in \mathcal{Y}} P(\lambda_i = y_k | \mathbf{Y} = y) P(\lambda_j = y_l | \mathbf{Y} = y) P(\mathbf{Y} = y). \end{aligned}$$

Letting B_j be the $|\mathcal{Y}_{\tau_j}| \times |\mathcal{Y}|$ matrix of conditional probabilities, $(B_j)_{j,k} = P(\lambda_i = y_j | \mathbf{Y} = y_k)$, and P be the diagonal matrix such that $P_{i,i} = P(\mathbf{Y} = y_i)$, we can re-express the above as:

$$A_{i,j} = B_i P B_j^T.$$

Since P is composed of strictly positive elements, and is diagonal (and thus PSD), we re-express this as:

$$A_{i,j} = \tilde{B}_i \tilde{B}_j^T, \quad (16)$$

where $\tilde{B}_i = B_i \sqrt{P}$. We could now try to recover P by decomposing the observed $A_{i,j}$ to recover the \tilde{B}_i , and from there recover P via the relation:

$$P = \text{diag}(\tilde{B}_i^{-T} \tilde{B}_i)^2, \quad (17)$$

since summing the column of \tilde{B}_i corresponding to label \mathbf{Y} is equal to $\sqrt{P(\mathbf{Y})} \sum_{y \in \mathcal{Y}_i} P(\lambda_i = y | \mathbf{Y}) = \sqrt{P(\mathbf{Y})}$ by the law of total probability. However, note that $\tilde{B}_i U$ for any orthogonal matrix U also satisfies (16), and could thus lead to a potentially infinite number of incorrect estimates of P .

Class Balance Identifiability with Three-Way View Constraint—A different approach involves considering the three-way overlaps observed as $A_{i,j,k}$. This is equivalent to performing a tensor decomposition. Note that above, the problem is that matrix decomposition is typically invariant to rotations and reflections; tensor decompositions have easier-to-meet uniqueness conditions (and are thus more rigid).

Specifically, we apply Kruskal's classical identifiability condition for unique 3-tensor decomposition. Consider some tensor

$$T = \sum_{r=1}^R X_r \otimes Y_r \otimes Z_r,$$

where X_r, Y_r, Z_r are column vectors that make up the matrices X, Y, Z . The Kruskal rank k_X of X is the largest k such that any k columns of X are linearly independent. Then, the decomposition above is unique if $k_X + k_Y + k_Z \geq 2R + 2$ [20; 3]. In our case, our triple views have $R = |Y|$, and we have

$$A_{i,j,k} = \tilde{B}_i \otimes \tilde{B}_j \otimes \tilde{B}_k. \quad (18)$$

Thus, if $k_{\tilde{B}_i} + k_{\tilde{B}_j} + k_{\tilde{B}_k} \geq 2|Y| + 2$, we have identifiability. Thus, it is sufficient to have the columns of each of the \tilde{B}_i 's be linearly independent. Note that each of the B_i 's have columns with the same sum, so these columns are only linearly dependent if they are equal, which would only be the case if the sources were random voters.

Thus, we can use (18) to recover the \tilde{B}_i in a stable fashion, and then use (17) to recover the $P(\mathbf{Y})$.




Figure 8:

Example task hierarchy G_{task} for a three-task classification problem. Task Y_1 classifies a data point X as a PERSON or BUILDING. If Y_1 classifies X as a PERSON, Y_2 is used to distinguish between DOCTOR and NON-DOCTOR. Similarly, if Y_1 classifies X as a BUILDING, Y_3 is used to distinguish between HOSPITAL and NON-HOSPITAL. Tasks Y_2, Y_3 are more specific, or *finer-grained* tasks, constrained by their parent task Y_1 .

A.3.6 Predicting Labels with the Label Model

Once we have an estimate of μ , we can make predictions with the label model—i.e. generate our *probabilistic* training labels $P_{\mu}(\mathbf{Y} | \lambda)$ —using the junction tree we have already defined

over G_{source} . Specifically, let $\tilde{\mathcal{C}}$ be the set of maximal cliques (nodes) in the junction tree, and let \mathcal{S} be the set of separator sets. Then we have:

$$P_{\mu}(\mathbf{Y}, \boldsymbol{\lambda}) = \frac{\prod_{C \in \tilde{\mathcal{C}}} P(V_C)}{\prod_{S \in \mathcal{S}} P(V_S)} = \frac{\prod_{C \in \tilde{\mathcal{C}}} \mu(C, (\mathbf{Y}, \boldsymbol{\lambda}_C))}{\prod_{S \in \mathcal{S}} \mu(S, (\mathbf{Y}, \boldsymbol{\lambda}_S))},$$

where again, $V_C = \{V_i\}_{i \in C}$, where $V_0 = \mathbf{Y}$ and $V_{i>0} = \boldsymbol{\lambda}_i$. Thus, we can directly compute the predicted labels $P_{\mu}(\mathbf{Y} | \boldsymbol{\lambda})$ based on the estimated parameters μ .

A.4 Example: Hierarchical Multi-Task Supervision

We now consider the specific case of *hierarchical* multi-task supervision, which can be thought of as consisting of coarser- and finer-grained labels, or alternatively higher- and lower-level labels, and provides a way to supervise e.g. fine-grained classification tasks at multiple levels of granularity. Specifically, consider a task label vector $\mathbf{Y} = [Y_1, \dots, Y_l]^T$ as before, this time with $Y_s \in \{N/A, 1, \dots, k_s\}$, where we will explain the meaning of the special value *N/A* shortly. We then assume that the tasks Y_s are related by a *task hierarchy* which is a hierarchy $G_{\text{task}} = (V, E)$ with vertex set $V = \{Y_1, Y_2, \dots, Y_l\}$ and directed edge set E . The task structure reflects constraints imposed by higher level (more general) tasks on lower level (more specific) tasks. The following example illustrates a simple tree task structure:

Example 2

Let Y_1 classify a data point X as either a *PERSON* ($Y_1 = 1$) or *BUILDING* ($Y_1 = 2$). If $Y_1 = 1$, indicating that X represents a *PERSON*, then Y_2 can further label X as a *DOCTOR* or *NON-DOCTOR*. Y_3 is used to distinguish between *HOSPITAL* and *NON-HOSPITAL* in the case that $Y_1 = 2$. The corresponding graph G_{task} is shown in Figure 8. If $Y_1 = 2$, then task Y_2 is not applicable, since Y_2 is only suitable for persons; in this case, Y_2 takes the value *N/A*. In this way the task hierarchy defines a feasible set of task vector values: $\mathbf{Y} = [1, 1, N/A]^T$, $[1, 2, N/A]^T$, $[2, N/A, 1]^T$, $[2, N/A, 2]^T$ are valid, while e.g. $\mathbf{Y} = [1, 1, 2]^T$ is not.

As in the example, for certain configurations of \mathbf{Y} 's, the parent tasks logically constrain the one or more of the children tasks to be irrelevant, or rather, to have inapplicable label values. In this case, the task takes on the value *N/A*. In Example 2, we have that if $Y_1 = 1$, representing a building, then Y_2 is inactive (since X corresponds to a building). We define the symbol *N/A* (for incompatible) for this scenario. More concretely, let $\mathcal{N}(Y_j) = \{Y_i : (Y_j, Y_i) \in E\}$ be the in-neighborhood of Y_j . Then, the values of the members of $\mathcal{N}(Y_j)$ determine whether $Y_j = N/A$, i.e., $\mathbb{1}\{Y_j = N/A\}$ is deterministic conditioned on $\mathcal{N}(Y_j)$.

Hierarchical Multi-Task Sources

Observe that in the mutually-exclusive task hierarchy just described, the value of a descendant task label Y_d determines the values of all other task labels in the hierarchy besides its descendants. For example, in Example 2, a label $Y_2 = 1 \Rightarrow (Y_1 = 1, Y_3 = N/A)$; in other words, knowing that X is a *DOCTOR* also implies that X is a *PERSON* and not a *BUILDING*.

For a source λ_i with coverage set t_i , the label it gives to the lowest task in the task hierarchy which is non-zero and non- N/A determines the entire label vector output by λ_i . E.g. if the lowest task that λ_i labels in the hierarchy is $Y_1 = 1$, then this implies that it outputs vector $[1, 0, N/A]^T$. Thus, in this sense, we can think of each sources λ_i as labeling one specific task in the hierarchy, and thus can talk about coarser- and finer-grained sources.

Reduced-Rank Form: Modeling Local Accuracies

In some cases, we can make slightly different modeling assumptions that reflect the nature of the task structure, and additionally can result in reduced-rank forms of our model. In particular, for the hierarchical setting introduced here, we can divide the statistics μ into *local* and *global* subsets, and for example focus on modeling only the *local* ones to once again reduce to rank-one form.

To motivate with our running example: a finer-grained source that labels DOCTOR versus NON-DOCTOR probably is not accurate on the building type subtask; we can model this source using one accuracy parameter for the former label set (the *local* accuracy) and a different (or no parameter) for the *global* accuracy on irrelevant tasks. More specifically, for cliques involving λ_i , we can model $P(\lambda_i, \mathbf{Y})$ for all \mathbf{Y} with only non- N/A values in the coverage set of λ_i using a single parameter, and call this the *local* accuracy; and we can either model μ for the other \mathbf{Y} using one or more other parameters, or simply set it to a fixed value and not model it, to reduce to rank one form, as we do in the experiments. In particular, this allows us to capture our observation in practice that if a developer is writing a source to distinguish between labels at one sub-tree, they are probably not designing or testing it to be accurate on any of the other subtrees.

B: Theoretical Results

In this section, we focus on theoretical results for the basic rank-one model considered in the main body of the paper. In Section B.1, we start by going through the conditions for identifiability in more detail for the rank-one case. In Section B.2, we provide additional interpretation for the expression of our primary theoretical result bounding the estimation error of the label model. In Section B.3, we then provide the proof of Theorem 1, connecting this estimation error to the generalization error of the end model; and in Section B.4, we provide the full proof of the main bound.

B.1 Conditions for Identifiability

We consider the rank-one setting as in the main body, where we have

$$-(\Sigma_o^{-1})_{\Omega} = (zz^T)_{\Omega}, \quad (19)$$

where Ω is the inverse augmented edge set, i.e. a pair of indices (i, j) , corresponding to elements of $\mathcal{Y}(\mathcal{C})$, and therefore to cliques $A, B \in \mathcal{C}$, is in Ω if A, B are not part of the same maximal clique in G_{source} (and therefore $(K_O)_{i,j} = 0$). This defines a set of $|\Omega|$ equations, which we can encode using a matrix M_{Ω} , where if (i, j) is the $(r-1)$ th entry in Ω , then

$$(M_{\Omega})_{r,s} = \begin{cases} 1 & s \in \{i, j\}, \\ 0 & \text{else.} \end{cases} \quad (20)$$

Let $l_i = \log(z_i^2)$ and $q_{(i,j)} = \log((\Sigma_{\sigma}^{-1})_{i,j})$; then by squaring and taking the log of both sides of 19, we get a system of linear equations:

$$M_{\Omega} l = q_{\Omega}. \quad (21)$$

Thus, we can identify z (and therefore μ) *up to sign* if the system of linear equations (21) has a solution.

Notes on Invertibility of M_{Ω}

Note that if the inverse augmented edge graph consists of a connected triangle (or any odd-numbered cycle), e.g. $\Omega = \{(i, j), (j, k), (i, k)\}$, then we can solve for the z_i up to sign, and therefore M_{Ω} must be invertible:

$$z_i^2 = \frac{(\Sigma_{\sigma}^{-1})_{i,j} (\Sigma_{\sigma}^{-1})_{i,k}}{(\Sigma_{\sigma}^{-1})_{j,k}},$$

and so on for z_j, z_k . Note additionally that if other z_i are connected to this triangle, then we can also solve for them up to sign as well. Therefore, if Ω contains at least one triangle (or odd-numbered cycle) per connected component, then M_{Ω} is invertible.

Also note that this is all in reference to the *inverse* source dependency graph, which will generally be dense (assuming the correlation structure between sources is generally sparse). For example, note that if we have one source λ_i that is conditionally independent of all the other sources, then Ω is fully connected, and therefore if there is a triangle in Ω , then M_{Ω} is invertible.

Identifying the Signs of the z_i

Finally, note that if we know the sign of one z_i , then this determines the signs of every other z_j in the same connected component. Therefore, for z to be uniquely identifiable, we need only know the *sign* of one of the z_i in each connected component. As noted already, if even one source λ_i is conditionally independent of all the other sources, then Ω is fully connected; in this case, we can simply assume that the average source is better than random, and therefore identify the signs of z without any additional information.

B.2 Interpreting the Main Bound

We re-state Theorem 2, which bounds the average error on the estimate of the label model parameters, providing more detail on and interpreting the terms of the bound.

Theorem 2

Let $\hat{\mu}$ be an estimate of μ^* produced by Algorithm 1 run over n unlabeled data points. Let

$$a := \left(\frac{d_o}{\Sigma_s} + \left(\frac{d_o}{\Sigma_s} \right)^2 \lambda_{\max}(K_o) \right)^{\frac{1}{2}} \text{ and } b := \frac{\|\Sigma_o^{-1}\|^2}{(\Sigma_o^{-1})_{\min}}. \text{ Then, we have:}$$

$$\mathbb{E}[\|\hat{\mu} - \mu^*\|] \leq 16(r-1)d_o^2 \sqrt{\frac{32\pi}{n}} ab \sigma_{\max}(M_\Omega^+) (3\sqrt{d_o} a \lambda_{\min}^{-1}(\Sigma_o) + 1) (\kappa(\Sigma_o) + \lambda_{\min}^{-1}(\Sigma_o)).$$

Influence of $\sigma_{\max}(M_\Omega^+)$ the largest singular value of the pseudoinverse M_Ω^+ . Note that $\|M_\Omega^+\|^2 = (\lambda_{\min}(M_\Omega^T M_\Omega))^{-1}$. As we shall see below, $\lambda_{\min}(M_\Omega^T M_\Omega)$ measures a quantity related to the structure of the graph G_{inv} . The smaller this quantity, the more information we have about G_{inv} , and the easier it is to estimate the accuracies. The smallest value of $\|M_\Omega^+\|^2$ (corresponding to the largest value of the eigenvalue) is $\sim \frac{1}{\sqrt{m}}$; the square of this quantity in the bound reduces the m^2 cost of estimating the covariance matrix to m .

It is not hard to see that

$$M_\Omega^T M_\Omega = \text{diag}(\text{deg}(G_{\text{inv}})) + \text{Adj}(G_{\text{inv}}).$$

Here, $\text{deg}(G_{\text{inv}})$ are the degrees of the nodes in G_{inv} and $\text{Adj}(G_{\text{inv}})$ is its adjacency matrix. This form closely resembles the graph Laplacian, which differs in the sign of the adjacency matrix term: $\mathcal{L}(G) = \text{diag}(\text{deg}(G)) - \text{Adj}(G)$. We bound

$$\sigma_{\max}(M_\Omega^+) \leq (d_{\min} + \lambda_{\min}(\text{Adj}(G_{\text{inv}})))^{-1},$$

where d_{\min} is the lowest-degree node in G_{inv} (that is, the source s with fewest appearances in Ω). In general, computing $\lambda_{\min}(\text{Adj}(G_{\text{inv}}))$ can be challenging. A closely related task can be done via *Cheeger inequalities*, which state that

$$2h_G \geq \lambda_{\min}(\mathcal{L}(G)) \geq \frac{1}{2}h_G^2,$$

where $\lambda_{\min}(\mathcal{L}(G))$ is the smallest non-zero eigenvalue of $\mathcal{L}(G)$ and

$$h_G = \min_X \frac{|E(X, \bar{X})|}{\min\{\sum_{x \in X} d_x, \sum_{y \in \bar{X}} d_y\}}$$

is the *Cheeger constant* of the graph [7]. The utility of the Cheeger constant is that it measures the presence of a bottleneck in the graph; the presence of such a bottleneck limits the graph density and is thus beneficial when estimating the structure in our case. Our Cheeger-constant like term $\sigma_{\max}(M_\Omega^+)$ acts the same way.

Now, in the easiest and most common case is that of conditionally independent sources [9; 38; 9; 17]., $\text{Adj}(G_{\text{inv}})$ has 1's everywhere but the diagonal, and we can compute explicitly that

$$\sigma_{\max}(M_{\Omega}^+) = \frac{1}{\sqrt{m-2}}.$$

In the general setting, we must compute the minimal eigenvalue of the adjacency matrix, which is tractable, for example, for tree structures.

Influence of $\lambda_{\min}(\Sigma_{\mathcal{O}})$ the smallest eigenvalue of the observed matrix. This quantity reflects the conditioning of the observed (correlation) matrix; the better conditioned the matrix, the easier it is to estimate $\Sigma_{\mathcal{O}}$.

Influence of $(\Sigma_{\mathcal{O}}^{-1})_{\min}$ the smallest entry of the inverse observed matrix. This quantity contributes to Σ^{-1} , the generalized precision matrix that we centrally use; it is a measure of the smallest non-zero correlation between source accuracies (that is, the smallest correlation between non-independent source accuracies). Note that the tail bound of Theorem 2 scales as $\exp(-(\Sigma_{\mathcal{O}}^{-1})_{\min})^2$. This is natural, as distinguishing between small correlations and independencies requires more samples.

B.3 Proof of Theorem 1

Let \mathcal{D} be the true data generating distribution, such that $(X, \mathbf{Y}) \sim \mathcal{D}$. Let $P_{\mu}(\mathbf{Y}|\boldsymbol{\lambda})$ be the label model parameterized by μ and conditioned on the observed source labels $\boldsymbol{\lambda}$. Furthermore, assume that:

1. For some optimal label model parameters μ^* , $P_{\mu^*}(\boldsymbol{\lambda}, \mathbf{Y}) = P(\boldsymbol{\lambda}, \mathbf{Y})$;
2. The label \mathbf{Y} is independent of the features of our end model given the source labels $\boldsymbol{\lambda}$.

That is, we assume that (i) the *optimal* label model, parameterized by μ^* , correctly matches the true distribution of source labels $\boldsymbol{\lambda}$ drawn from the true distribution, $(s(X), \mathbf{Y}) \sim \mathcal{D}$; and (ii) that these labels $\boldsymbol{\lambda}$ provide sufficient information to discern the label \mathbf{Y} . We note that these assumptions are the ones used in prior work [28], and are intended primarily to illustrate the connection between the estimation accuracy of $\hat{\mu}$, which we bound in Theorem 2, and the end model performance.

Now, suppose that we have an end model parameterized by w , and that to learn these parameters we minimize a normalized bounded loss function $\mathcal{L}(w, X, \mathbf{Y})$, such that without loss of generality, $\mathcal{L}(w, X, \mathbf{Y}) \in [0, 1]$. Normally our goal would be to find parameters that minimize the expected loss, which we denote w^* :

$$L(w) = \mathbb{E}_{(X, \mathbf{Y}) \sim \mathcal{D}}[\mathcal{L}(w, X, \mathbf{Y})] \quad (22)$$

However, since we do not have access to the true labels \mathbf{Y} , we instead minimize the expected noise-aware loss, producing an estimate \hat{w} :

$$L_\mu(w) = \mathbb{E}_{(X, \mathbf{Y}) \sim \mathcal{D}} [\mathbb{E}_{\tilde{\mathbf{Y}} \sim P_{\mu^*}(\cdot | \lambda(X))} [l(w, X, \tilde{\mathbf{Y}})]] . \quad (23)$$

In practice, we actually minimize the *empirical* version of the noise aware loss over an unlabeled dataset $U = \{X^{(1)}, \dots, X^{(n)}\}$, producing an estimate \hat{w} :

$$\hat{L}_\mu(w) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\tilde{\mathbf{Y}} \sim P_{\mu^*}(\cdot | \lambda(X^{(i)}))} [l(w, X^{(i)}, \tilde{\mathbf{Y}})] . \quad (24)$$

Let w^* be the expected loss L , let \tilde{w} be the minimizer of the noise-aware loss for estimated label model parameters μ , L_{μ^*} , and let \hat{w} be the minimize of the empirical noise aware loss \hat{L}_μ . Our goal is to bound the *generalization risk*- the difference between the expected loss of our empirically estimated parameters and of the optimal parameters,

$$L(\hat{w}) - L(w^*) . \quad (25)$$

Additionally, since analyzing the empirical risk minimization error is standard and not specific to our setting, we simply assume that error $|L_\mu(\tilde{w}) - L_\mu(\hat{w})| \leq \gamma(n)$, where $\gamma(n)$ is a decreasing function of the number of unlabeled data points n .

To start, using the law of total expectation first, followed by our assumption (2) about conditional independence, and finally using our assumption (1) about our optimal label model μ^* , we have that:

$$\begin{aligned} L(w) &= \mathbb{E}_{(X', \mathbf{Y}') \sim \mathcal{D}} [L(w)] \\ &= \mathbb{E}_{(X', \mathbf{Y}') \sim \mathcal{D}} [\mathbb{E}_{(X, \mathbf{Y}) \sim \mathcal{D}} [l(w, X', \mathbf{Y}) | X = X']] \\ &= \mathbb{E}_{(X', \mathbf{Y}') \sim \mathcal{D}} [\mathbb{E}_{(X, \mathbf{Y}) \sim \mathcal{D}} [l(w, X', \mathbf{Y}) | s(X) = s(X')]] \\ &= \mathbb{E}_{(X', \mathbf{Y}') \sim \mathcal{D}} [\mathbb{E}_{(\lambda, \tilde{\mathbf{Y}}) \sim \mu^*} [l(w, X', \tilde{\mathbf{Y}}) | \lambda = s(X')]] \\ &= L_{\mu^*}(w) . \end{aligned}$$

Now, we have:

$$\begin{aligned} L(\hat{w}) - L(w^*) &= L_{\mu^*}(\hat{w}) + L_\mu(\hat{w}) - L_\mu(\hat{w}) + L_\mu(\tilde{w}) - L_\mu(\tilde{w}) - L_{\mu^*}(w^*) \\ &\leq L_{\mu^*}(\hat{w}) + L_\mu(\tilde{w}) - L_\mu(\hat{w}) + L_\mu(w^*) - L_\mu(\tilde{w}) - L_{\mu^*}(w^*) \\ &\leq |L_\mu(\hat{w}) - L_\mu(\tilde{w})| + |L_{\mu^*}(\hat{w}) - L_\mu(\hat{w})| + |L_\mu(w^*) - L_{\mu^*}(w^*)| \\ &\leq \gamma(n) + 2 \max_{w'} |L_{\mu^*}(w') - L_\mu(w')| , \end{aligned}$$

where in the first step we use our result that $L = L_{\mu^*}$ as well as add and subtract terms; and in the second step we use the fact that $L_\mu(\tilde{w}) \leq L_\mu(w^*)$. We now have our generalization risk controlled primarily by $|L_{\mu^*}(w') - L_\mu(w')|$, which is the difference between the expected noise aware losses given the estimated label model parameters μ and the true label model parameters μ^* . Next, we see that, for any w' :

$$\begin{aligned}
|L_{\mu^*}(w') - L_{\mu}(w')| &= \left| \mathbb{E}_{(X, Y) \sim \mathcal{D}} [\mathbb{E}_{\tilde{Y} \sim P_{\mu^*}(\cdot | \lambda)} [l(w, X, \tilde{Y})] - \mathbb{E}_{\tilde{Y} \sim P_{\mu}(\cdot | \lambda)} [l(w, X, \tilde{Y})]] \right| \\
&= \left| \mathbb{E}_{(X, Y) \sim \mathcal{D}} \left[\sum_{\mathbf{Y}' \in \mathcal{Y}} l(w, X, \mathbf{Y}') (P_{\mu^*}(\mathbf{Y}' | \lambda) - P_{\mu}(\mathbf{Y}' | \lambda)) \right] \right| \\
&\leq \sum_{\mathbf{Y}' \in \mathcal{Y}} \mathbb{E}_{(X, Y) \sim \mathcal{D}} [|P_{\mu^*}(\mathbf{Y}' | \lambda) - P_{\mu}(\mathbf{Y}' | \lambda)|] \\
&\leq \left| \mathcal{Y} \right| \max_{\mathbf{Y}'} \mathbb{E}_{(X, Y) \sim \mathcal{D}} [|P_{\mu^*}(\mathbf{Y}' | \lambda) - P_{\mu}(\mathbf{Y}' | \lambda)|],
\end{aligned}$$

where we have now bounded $|L_{\mu^*}(w') - L_{\mu}(w')|$ by the size of the structured output space $|\mathcal{Y}|$, and a term having to do with the difference between the probability distributions of μ and μ^* .

Now, we use the result from [16] (Lemma 19) which establishes that the log probabilities of discrete factor graphs with indicator features (such as our model $P_{\mu}(\boldsymbol{\lambda}, \mathbf{Y})$) are $(L_{\infty}, 2)$ -Lipschitz with respect to their parameters, and the fact that for x, y s.t. $|x|, |y| \leq 1, |x - y| \leq |\log(x) - \log(y)|$, to get:

$$\begin{aligned}
|P_{\mu^*}(\mathbf{Y}' | \lambda) - P_{\mu}(\mathbf{Y}' | \lambda)| &\leq |\log P_{\mu^*}(\mathbf{Y}' | \lambda) - \log P_{\mu}(\mathbf{Y}' | \lambda)| \\
&\leq |\log P_{\mu^*}(\boldsymbol{\lambda}, \mathbf{Y}') - \log P_{\mu}(\boldsymbol{\lambda}, \mathbf{Y}')| + |\log P_{\mu^*}(\boldsymbol{\lambda}) - \log P_{\mu}(\boldsymbol{\lambda})| \\
&\leq 2\|\mu^* - \mu\|_{\infty} + 2\|\mu^* - \mu\|_{\infty} \\
&\leq 4\|\mu^* - \mu\|,
\end{aligned}$$

where we use the fact that the statement of Lemma 19 also holds for every marginal distribution as well. Therefore, we finally have:

$$L(\hat{w}) - L(w^*) \leq \gamma(n) + 4|\mathcal{Y}|\|\mu^* - \mu\|.$$

B.4 Proof of Theorem 2

Proof. First we briefly provide a roadmap of the proof of Theorem 2. We consider estimating $\tilde{\mu}$ with our procedure in the rank-one setting, and we seek a tail bound on $\|\tilde{\mu} - \mu\|$. The challenge here is that the observed matrix Σ_o we see is itself constructed from a series of observed i.i.d. samples $\psi(O)^{(1)}, \dots, \psi(O)^{(n)}$. We bound (through a matrix concentration inequality) the error $\Delta_o = \tilde{\Sigma}_o - \Sigma_o$, and view $\tilde{\Sigma}_o$ as a perturbation of Σ_o . Afterwards, we use a series of perturbation analyses to ultimately bound $\|\tilde{\Sigma}_{os} - \Sigma_{os}\|$, and then use this directly to bound $\|\tilde{\mu} - \mu\|$; each of the perturbation results is in terms of Δ_o .

We begin with some notation. We write the following perturbations (note that all the terms written with $\tilde{\cdot}$ are additive, while the S term is relative)

$$\begin{aligned}
\tilde{\Sigma}_{os} &= \Sigma_{os} + \Delta_{os}, \\
\tilde{\Sigma}_o &= \Sigma_o + \Delta_o, \\
\tilde{\ell} &= \ell + \Delta_{\ell}, \\
\tilde{z} &= (I + \text{diag}(\delta_z))z.
\end{aligned}$$

Now we start our perturbation analysis:

$$\begin{aligned}\tilde{\Sigma}_{os} &= \frac{1}{\sqrt{\tilde{c}}}\tilde{\Sigma}_o\tilde{z} = \frac{1}{\sqrt{\tilde{c}}}(\Sigma_o + \Delta_o)(I + \text{diag}(\delta_z))z \\ &= \frac{1}{\sqrt{\tilde{c}}}(\Sigma_o z + \Sigma_o \text{diag}(\delta_z)z + \Delta_o(I + \text{diag}(\delta_z))z).\end{aligned}$$

Subtracting $\Sigma_{os} = \frac{1}{\sqrt{c}}\Sigma_o z$, we get

$$\Delta_{os} = \left(\frac{1}{\sqrt{\tilde{c}}} - \frac{1}{\sqrt{c}}\right)\Sigma_o z + \frac{1}{\sqrt{\tilde{c}}}(\Sigma_o \text{diag}(\delta_z)z + \Delta_o(I + \text{diag}(\delta_z))z). \quad (26)$$

The rest of the analysis requires us to bound the norms for each of these terms.

Left-most term.

We have that

$$\left\|\left(\frac{1}{\sqrt{\tilde{c}}} - \frac{1}{\sqrt{c}}\right)\Sigma_o z\right\| = \left|\frac{\sqrt{c}}{\sqrt{\tilde{c}}} - 1\right| \left\|\frac{1}{\sqrt{c}}\Sigma_o z\right\| = \left|\frac{\sqrt{c}}{\sqrt{\tilde{c}}} - 1\right| \|\Sigma_{os}\| \leq \sqrt{d_o} \left|\frac{\sqrt{c}}{\sqrt{\tilde{c}}} - 1\right| \leq \sqrt{d_o}|\tilde{c} - c|.$$

Here, we bounded $\|\Sigma_{os}\|$ by $\sqrt{d_o}$, since $\Sigma_{os} \in [-1, 1]^{d_o}$. Then, note that

$c = \Sigma_s^{-1}(1 + z^T \Sigma_o z) \geq 0$, since $\Sigma_s < 1$ and $\Sigma_o \geq 0 \Rightarrow z^T \Sigma_o z \geq 0$, so therefore $c, \tilde{c} \geq 1$. In the last inequality, we use this to imply that $|\sqrt{c}/\sqrt{\tilde{c}} - 1| \leq |\sqrt{c} - \sqrt{\tilde{c}}| \leq |\tilde{c} - c|$. Next we work on bounding $|\tilde{c} - c|$. We have

$$\begin{aligned}|\tilde{c} - c| &= |\Sigma_s^{-1}| \left| z^T \tilde{\Sigma}_o \tilde{z} - z^T \Sigma_o z \right| \\ &= |\Sigma_s^{-1}| \left| z^T (I + \text{diag}(\delta_z))^T (\Sigma_o + \Delta_o)(I + \text{diag}(\delta_z))z - z^T \Sigma_o z \right| \\ &= |\Sigma_s^{-1}| \left| z^T \Sigma_o \text{diag}(\delta_z)z + z^T \Delta_o (I + \text{diag}(\delta_z))z + z^T \text{diag}(\delta_z)^T (\Sigma_o + \Delta_o)(I + \text{diag}(\delta_z))z \right| \\ &\leq |\Sigma_s^{-1}| \left\| z \right\|^2 \left(\|\Sigma_o\| (2\|\delta_z\| + \|\delta_z\|^2) + \|\Delta_o\| (2\|\delta_z\| + \|\delta_z\|^2 + 1) \right) \\ &\leq \|z\|^2 \left(\|\Sigma_o\| (2\|\delta_z\| + \|\delta_z\|^2) + \|\Delta_o\| (2\|\delta_z\| + \|\delta_z\|^2 + 1) \right).\end{aligned}$$

Thus,

$$\left\|\left(\frac{1}{\sqrt{\tilde{c}}} - \frac{1}{\sqrt{c}}\right)\Sigma_o z\right\| \leq \sqrt{d_o} \left\| z \right\|^2 \left(\|\Sigma_o\| (2\|\delta_z\| + \|\delta_z\|^2) + \|\Delta_o\| (2\|\delta_z\| + \|\delta_z\|^2 + 1) \right). \quad (27)$$

Bounding c .

We will need a bound on c to bound z . We have that

$$c = (\Sigma_s - \Sigma_{os}^T \Sigma_o^{-1} \Sigma_{os})^{-1}.$$

Applying the Woodbury matrix inversion lemma, we have:

$$c = \Sigma_s^{-1} + \Sigma_s^{-1} \Sigma_{O_s}^T (\Sigma_O - \Sigma_{O_s} \Sigma_s^{-1} \Sigma_{O_s}^T)^{-1} \Sigma_{O_s} \Sigma_s^{-1}$$

Now, by the blockwise inversion lemma, we know that

$$K_O = (\Sigma_O - \Sigma_{O_s} \Sigma_s^{-1} \Sigma_{O_s}^T)^{-1}$$

So we then have:

$$c = \Sigma_s^{-1} + \Sigma_s^{-1} \Sigma_{O_s}^T K_O \Sigma_{O_s} \Sigma_s^{-1} \leq \Sigma_s^{-1} + (\Sigma_s^{-1})^2 \|\Sigma_{O_s}\|^2 \|K_O\|$$

Bounding z .

We'll use our bound on c , since $z = \sqrt{c} \Sigma_O^{-1} \Sigma_{O_s}$.

$$\begin{aligned} \|z\| &= \|\sqrt{c} \Sigma_O^{-1} \Sigma_{O_s}\| \\ &\leq \left(\Sigma_s^{-1} + (\Sigma_s^{-1})^2 \|\Sigma_{O_s}\|^2 \|K_O\| \right)^{\frac{1}{2}} \|\Sigma_O^{-1}\| \|\Sigma_{O_s}\| \\ &\leq \left(\Sigma_s^{-1} + (\Sigma_s^{-1})^2 d_O \|K_O\| \right)^{\frac{1}{2}} \|\Sigma_O^{-1}\| \sqrt{d_O} \\ &= \frac{d_O}{\Sigma_s} \left(\frac{\Sigma_s}{d_O} + \lambda_{\max}(K_O) \right)^{\frac{1}{2}} \lambda_{\min}^{-1}(\Sigma_O) \end{aligned}$$

In the last inequality, we used the fact that $\|\Sigma_{O_s}\|^2 \leq d_O$. Now we want to control $\|\Delta_\epsilon\|$.

Perturbation bound.

We have the perturbation bound

$$\|\Delta_\epsilon\| \leq \|M_\Omega^\dagger\| \|\tilde{q}_s - q_s\|. \quad (28)$$

We need to work on the term $\|\tilde{q}_s - q_s\|$. To avoid overly heavy notation, we write

$P = \Sigma_O^{-1}$, $\tilde{P} = \tilde{\Sigma}_O^{-1}$, and $\Delta_P = P - \tilde{P}$. Then we have:

$$\begin{aligned}
\|\tilde{q}_s - q_s\|^2 &= \sum_{(i,j) \in \mathcal{S}} (\log(\tilde{P}_{i,j}^2) - \log(P_{i,j}^2))^2 \\
&= 4 \sum_{(i,j) \in \mathcal{S}} (\log(|\tilde{P}_{i,j}|) - \log(|P_{i,j}|))^2 \\
&= 4 \sum_{(i,j) \in \mathcal{S}} (\log(|P_{i,j} + (\Delta p)_{i,j}|) - \log(|P_{i,j}|))^2 \\
&\leq 4 \sum_{(i,j) \in \mathcal{S}} \left[\log\left(1 + \left|\frac{(\Delta p)_{i,j}}{P_{i,j}}\right|\right) \right]^2 \\
&\leq 8 \sum_{(i,j) \in \mathcal{S}} \left(\frac{|(\Delta p)_{i,j}|}{|P_{i,j}|} \right)^2 \\
&\leq \frac{8}{P_{\min}^2} \sum_{(i,j) \in \mathcal{S}} (\Delta p)_{i,j}^2 \\
&\leq \frac{8 \|\tilde{\Sigma}_o^{-1} - \Sigma_o^{-1}\|^2}{((\Sigma_o^{-1})_{\min})^2}.
\end{aligned}$$

Here, the second inequality uses $(\log(1+x))^2 \leq x^2$, and the fourth inequality sums over squared values. Next, we use the perturbation bound $\|\tilde{\Sigma}_o^{-1} - \Sigma_o^{-1}\| \leq \|\Sigma_o^{-1}\|^2 \|\Delta_o\|$, so that we have

$$\|\tilde{q}_s - q_s\| \leq \frac{2\sqrt{2}\|\Sigma_o^{-1}\|^2 \|\Delta_o\|}{(\Sigma_o^{-1})_{\min}}.$$

Then, plugging this into (28), we get that

$$\|\Delta_r\| \leq \sigma_{\max}(M_{\Omega}^+) \frac{2\sqrt{2}\|\Sigma_o^{-1}\|^2 \|\Delta_o\|}{(\Sigma_o^{-1})_{\min}}. \quad (29)$$

Bounding δ_z

Note also that $\|\Delta_r\|^2 = \sum_{i=1}^m (\log(\tilde{z}_i^2) - \log(z_i^2))$. We have that

$$\begin{aligned}
\|\Delta_r\|^2 &= \sum_{i=1}^m \log\left(\frac{\tilde{z}_i^2}{z_i^2}\right) \\
&= 2 \sum_{i=1}^m \log\left(\frac{|\tilde{z}_i|}{|z_i|}\right) \\
&= 2 \sum_{i=1}^m \log(1 + |(\delta_z)_i|), \\
&\geq 2 \sum_{i=1}^m (\delta_z)_i^2 \\
&= 2\|\delta_z\|^2,
\end{aligned}$$

where in the fourth step, we used the bound $\log(1+a) \geq a^2$ for small a . Then, we have

$$\|\delta_z\| \leq \frac{\sqrt{2}\|\Sigma_o^{-1}\|^2\|\Delta_o\|}{(\Sigma_o^{-1})_{\min}}\sigma_{\max}(M_{\hat{\Omega}}^+). \quad (30)$$

Putting it together.

Using (26), we have that

$$\begin{aligned} \|\Delta_{os}\| &= \left\| \left(\frac{1}{\sqrt{\tilde{c}}} - \frac{1}{\sqrt{c}} \right) \Sigma_o z + \frac{1}{\sqrt{\tilde{c}}} (\Sigma_o \text{diag}(\delta_z) z + \Delta_o (I + \text{diag}(\delta_z)) z) \right\| \\ &\leq \left\| \left(\frac{1}{\sqrt{\tilde{c}}} - \frac{1}{\sqrt{c}} \right) \Sigma_o z \right\| + (\|\Sigma_o \text{diag}(\delta_z)\| + \|\Delta_o (I + \text{diag}(\delta_z))\|) \|z\| \\ &\leq \sqrt{d_o} \|z\|^2 \left(\|\Sigma_o\| (2\|\delta_z\| + \|\delta_z\|^2) + \|\Delta_o\| (2\|\delta_z\| + \|\delta_z\|^2 + 1) \right) \\ &\quad + \|\Sigma_o\| \|\delta_z\| \|z\| + \|\Delta_o\| \|z\| (1 + \|\delta_z\|) \\ &\leq \sqrt{d_o} \|z\|^2 (3\|\Sigma_o\| \|\delta_z\| + 3\|\Delta_o\| \|\delta_z\| + \|\Delta_o\|) \\ &\quad + \|\Sigma_o\| \|\delta_z\| \|z\| + \|\Delta_o\| \|z\| (1 + \|\delta_z\|) \\ &\leq \|z\| \left(3\sqrt{d_o} \|z\| + 1 \right) (\|\Sigma_o\| + \|\Delta_o\|) \|\delta_z\| + \|\Delta_o\| \end{aligned}$$

Where in the first inequality, we use the triangle inequality and the fact that $\tilde{c} > 1$, and in the third inequality, we relied on the fact that we can control $\|\delta_z\|$ (through $\|\Delta_o\|$) so that we can make it small enough and thus take $\|\delta_z\|^2 \leq \|\delta_z\|$. Now we can plug in our bounds on $\|z\|$ and $\|\delta_z\|$ from before:

$$\begin{aligned} \|\Delta_{os}\| &\leq \left(\frac{d_o}{\Sigma_s} \left(\frac{\Sigma_s}{d_o} + \lambda_{\max}(K_o) \right)^{\frac{1}{2}} \lambda_{\min}^{-1}(\Sigma_o) \right) \left(3\sqrt{d_o} \left(\frac{d_o}{\Sigma_s} \left(\frac{\Sigma_s}{d_o} + \lambda_{\max}(K_o) \right)^{\frac{1}{2}} \lambda_{\min}^{-1}(\Sigma_o) \right) + 1 \right) \\ &\quad \times \left((\|\Sigma_o\| + \|\Delta_o\|) \left(\frac{\sqrt{2}\|\Sigma_o^{-1}\|^2\|\Delta_o\|}{(\Sigma_o^{-1})_{\min}} \sigma_{\max}(M_{\hat{\Omega}}^+) + \|\Delta_o\| \right) \right) \end{aligned}$$

For convenience, we set $\|\Delta_o\| = t$. Recall that

$$a = \left(\frac{d_o}{\Sigma_s} + \left(\frac{d_o}{\Sigma_s} \right)^2 \lambda_{\max}(K_o) \right)^{\frac{1}{2}}$$

and

$$b = \frac{\|\Sigma_o^{-1}\|^2}{(\Sigma_o^{-1})_{\min}}.$$

Then, we have

$$\|\Delta_{os}\| \leq \left(3\sqrt{d_o} a \lambda_{\min}^{-1}(\Sigma_o) + 1 \right) \left(\sqrt{2} a b \kappa(\Sigma_o) \sigma_{\max}(M_{\hat{\Omega}}^+) t + \sqrt{2} a b \frac{\sigma_{\max}(M_{\hat{\Omega}}^+)}{\lambda_{\min}(\Sigma_o)} t^2 + a \lambda_{\min}^{-1}(\Sigma_o) t \right).$$

Again we can take t small so that $t^2 \ll t$. Simplifying further, we have

$$\|\Delta_{O_S}\| \leq (3\sqrt{d_O}a\lambda_{\min}^{-1}(\Sigma_O) + 1)(\sqrt{2ab}\sigma_{\max}(M_{\Omega}^+)[\kappa(\Sigma_O) + \lambda_{\min}^{-1}(\Sigma_O)] + a\lambda_{\min}^{-1}(\Sigma_O))t.$$

Finally, since the $a\lambda_{\min}^{-1}(\Sigma_O)$ is smaller than the left-hand term inside the parentheses, we can write

$$\|\Delta_{O_S}\| \leq (3\sqrt{d_O}a\lambda_{\min}^{-1}(\Sigma_O) + 1)(2\sqrt{2ab}\sigma_{\max}(M_{\Omega}^+)[\kappa(\Sigma_O) + \lambda_{\min}^{-1}(\Sigma_O)])t. \quad (31)$$

Concentration bound.

We need to bound $t = \|\Sigma_O - \tilde{\Sigma}_O\|$, the error when estimating Σ_O from observations $\psi(O)^{(1)}, \dots, \psi(O)^{(n)}$ over n unlabeled data points.

To start, recall that O is the set of observable cliques, $\psi(O) \in \{0, 1\}^{d_O}$ is the corresponding vector of minimal statistics, and $\Sigma_O = \mathbf{Cov}[\psi(O)]$. For notational convenience, let $R = \mathbb{E}[\psi(O)\psi(O)^T]$, $r = \mathbb{E}[\psi(O)]$, and $r_k = \psi(O)^{(k)}$, and $\Delta_r = \frac{1}{n} \sum_{i=1}^n r_k - r$. Then we have:

$$\begin{aligned} \|\Delta_O\| &= \|\Sigma_O - \tilde{\Sigma}_O\| = \left\| (R - rr^T) - \left(\frac{1}{n} \sum_{i=1}^n r_i r_i^T - (r + \Delta_r)(r + \Delta_r)^T \right) \right\| \\ &\leq \underbrace{\left\| R - \frac{1}{n} \sum_{i=1}^n r_i r_i^T \right\|}_{\Delta_R} + \underbrace{\left\| rr^T - (r + \Delta_r)(r + \Delta_r)^T \right\|}_{\Delta_r}. \end{aligned}$$

We start by applying the matrix Hoeffding inequality [32] to bound the first term, Δ_R . Let $S_k = \frac{1}{n}(R - R_k)$, and thus clearly $\mathbb{E}[S_k] = 0$. We seek a sequence of symmetric matrices A_k s.t. $S_k^2 \preceq A_k^2$. First, note that, for some vectors x, v ,

$$x^T (\|v\|^2 I - vv^T) x = \|v\|^2 \|x\|^2 - \langle x, v \rangle^2 \geq 0$$

using Cauchy-Schwarz; therefore $\|v\|^2 I \succeq vv^T$, so that

$$d_O^2 I \succeq \|r_k\|^4 I \succeq \|r_k\|^2 r_k r_k^T = (r_k r_k^T)^2.$$

Next, note that $(r_k r_k^T + R)^2 \succeq 0$. Now, we use this to see that:

$$(nS_k)^2 = (r_k r_k^T - R)^2 \preceq (r_k r_k^T - R)^2 + (r_k r_k^T + R)^2 = 2((r_k r_k^T)^2 + R^2) \preceq 2(d_O^2 I + R^2).$$

Therefore, let $A_k^2 = \frac{2}{n^2}(d_O^2 I + R^2)$, and note that $\|R^2\| \leq \|R\|^2 \leq (d_O \|R\|_{\max})^2 = d_O^2$. We then have

$$\sigma^2 = \left\| \sum_{k=1}^n A_k^2 \right\| \leq \frac{2}{n} (d_o^2 + \|R^2\|) \leq \frac{4d_o^2}{n}.$$

And thus,

$$P(\|\Delta_R\| \geq \gamma) \leq 2d_o \exp\left(-\frac{n\gamma^2}{32d_o^2}\right). \quad (32)$$

Next, we bound $\|r\|$. We see that:

$$\begin{aligned} \|\Delta_r\| &= \|rr^T - (r + \Delta_r)(r + \Delta_r)^T\| \\ &= \|r\Delta_r^T + \Delta_r r^T + \Delta_r \Delta_r^T\| \\ &\leq \|r\Delta_r^T\| + \|\Delta_r r^T\| + \|\Delta_r \Delta_r^T\| \\ &\leq 2\|r\| \|\Delta_r\| + \|\Delta_r\|^2 \\ &\leq 3\|r\| \|\Delta_r\| \\ &\leq 3\|r\|_1 \|\Delta_r\|_1 \\ &\leq 3d_o^2 \|\Delta_r\|_1, \end{aligned}$$

where Δ_r is the perturbation for a single element of $\psi(O)$. We can then apply the standard Hoeffding's bound to get:

$$P(\|\Delta_r\| \geq \gamma) \leq 2\exp\left(-\frac{2n\gamma^2}{3d_o^2}\right).$$

Combining the bounds for $\|R\|$ and $\|r\|$, we get:

$$P(\|\Delta_o\| \geq \gamma) = P(t \geq \gamma) \leq 3d_o \exp\left(-\frac{n\gamma^2}{32d_o^2}\right). \quad (33)$$

Final steps

Now, we use the bound on t in (31) and the concentration bound above to write

$$\begin{aligned} P(\|\Delta_{o_s}\| \geq t') &\leq P(Vt \geq t') \\ &= P\left(t \geq \frac{t'}{V}\right) \\ &\leq 2d_o \exp\left(-\frac{nt'^2}{32V^2 d_o^2}\right), \end{aligned}$$

where $V = (3\sqrt{d_o} a \lambda_{\min}^{-1}(\Sigma_o) + 1) \left(2\sqrt{2} ab \sigma_{\max}(M_{\Omega}^+) \left[\kappa(\Sigma_o) + \frac{1}{\lambda_{\min}(\Sigma_o)}\right]\right)$.

Given $\tilde{\Sigma}_{o_s}$, we recover $\tilde{\mu}_1 = \tilde{\Sigma}_{o_s} + \mathbb{E}[\psi(H)] \widehat{\mathbb{E}}[\psi(O)]$. We assume $\mathbb{E}[\psi(H)]$ is known, and we can bound the error introduced by $\mathbb{E}[\psi(H)] \widehat{\mathbb{E}}[\psi(O)]$ as above, which we see can be folded into the looser bound for the error in $\tilde{\Sigma}_{o_s}$.

Finally, we expand the rank-one form $\tilde{\mu}_i$ into $\tilde{\mu}$ algebraically, according to our weight tying in the rank one model we use. Suppose in the rank one reduction (see Section A.3.4), we let $\mathbf{Y}_B = \mathbb{1}\{\mathbf{Y} = y_1\}$. Then each element of μ_1 that we track corresponds to either the probability of being correct, $\alpha_{c,y} = P(\cap_{i \in C} \{\lambda_i = y\}, \mathbf{Y} = y)$ or the probability of being incorrect, $\frac{1}{r-1}(1 - \alpha_{c,y})$, for each source clique C and label output combination y_C , and this value is simply copied $r-1$ times (for the other, weight-tied incorrect values), except for potentially one entry where it is multiplied by $(r-1)$ and then subtracted from 1 (to transform from incorrect to correct). Therefore, $\|\Delta_\mu\| = \|\mu - \tilde{\mu}\| \leq 2(r-1)\|\mu_1 - \tilde{\mu}_1\|$. Thus, we have:

$$P(\|\Delta_\mu\| \geq t') \leq 4(r-1)d_o \exp\left(-\frac{m'^2}{32V^2d_o^2}\right),$$

where V is defined as above. We only have one more step:

$$\begin{aligned} \mathbb{E}[\|\tilde{\mu} - \mu\|] &= \int_0^\infty P(\|\tilde{\mu} - \mu\| \geq \gamma) d\gamma \\ &\leq \int_0^\infty 4(r-1)d_o \exp\left(-\frac{n}{32V^2d_o^2}\gamma^2\right) d\gamma \\ &= \frac{4(r-1)d_o\sqrt{\pi}}{2\sqrt{\frac{n}{32V^2d_o^2}}} \\ &= 4(r-1)d_o\sqrt{\frac{32\pi}{n}}V. \end{aligned}$$

Here, we used the fact that $\int_0^\infty \exp(-a\gamma^2) d\gamma = \frac{\sqrt{\pi}}{2\sqrt{a}}$. \square

C: Experimental Details

C.1 Data Balancing and Label Model Training Procedure

For each application, rebalancing was applied via direct subsampling to the training set in the manner that was found to most improve development set micro-averaged accuracy. Specifically, we rebalance with respect to the median class for OpenI (i.e., removing examples from majority class such that none had more than the original median class), the minimum class for TACRED, and perform no rebalancing for OntoNotes. For generative model training, we use stochastic gradient descent with a step size, step number, and ℓ_2 penalty listed in Table 3 below. These parameters were found via 10-trial coarse random search, with all values determined via maximum micro-averaged accuracy evaluated on the development set.

C.2 End Model Training Procedure

Before training over multiple iterations to attain averaged results for reporting, a 10-trial random search over learning rate and ℓ_2 regularization with the Adam optimizer was

performed for each application based on micro-averaged development set accuracy. Learning rate was decayed by an order of magnitude if no increases in training loss improvement or development set accuracy were observed for 10 epochs, and the learning rate was frozen during the first 5 epochs. Models are reported using early stopping, wherein the best performing model on the development set is eventually used for evaluation on the held-out test set, and maximum epoch number is set for each application at a point beyond which minimal additional decrease in training loss was observed.

C.3 Dataset Statistics

We give additional detail in here (see Table 4) on the different datasets used for the experimental portion of this work. All data in the development and test sets is labeled with ground truth, while data in the training set is treated as unlabeled. Each dataset has a particular advantage in our study. The OntoNotes set, for instance, contains a particularly large number of relevant data points (over 63k), which enables us to investigate empirical performance scaling with the number of unlabeled data points. Further, the richness of the TACRED dataset allowed for the creation of an 8-class, 7-sub-task hierarchical classification problem, which demonstrates the utility of being able to supervise at each of the three levels of task granularity. Finally, the OpenI dataset represents a real-world, non-benchmark problem drawn from the domain of medical triage, and domain expert input was directly leveraged to create the relevant supervision sources. The fact that these domain expert weak supervision sources naturally occurred at multiple levels of granularity, and that they could be easily integrated to train an effective end model, demonstrates the utility of the MeTaL framework in practical settings.

Table 3:

Model architecture and training parameter details.

	OntoNotes	TACRED	OpenI
Label Model Training			
Step Size	5e-3	1e-2	5e-4
\mathcal{L}_2 Regularization	1e-4	4e-4	1e-3
Step Number	50	25	50
End Model Architecture			
Embedding Initialization	PubMed	FastText EN	Random
Embedding Size	100	300	200
LSTM Hidden Size	150	250	150
LSTM Layers	1	2	1
Intermediate Layer Dimensions	200, 50	200, 50, 25	200, 50
End Model Training			
Learning Rate	1e-2	1e-3	1e-3

	OntoNotes	TACRED	OpenI
ℓ_2 Regularization	1e-4	1e-4	1e-3
Epochs	20	30	50
Dropout	0.25	0.25	0.1

Table 4:

Dataset split sizes and sub-task structure for the three fine-grained classification tasks on which we evaluate MeTaL.

	# Train	#Dev	# Test	Tree Depth	# Tasks	# Sources/Task
OntoNotes (NER)	62,547	350	345	2	3	11
TACRED (RE)	9,090	350	2174	3	7	9
OpenI (Doc)	2,630	200	378	2	3	19

C.4 Task Accuracies

For clarity, we present in Table 5 the individual task accuracies of both the learned MeTaL model and MV for each experiment. These accuracies are computed from the output of evaluating each model on the test set with ties broken randomly.

C.5 Ablation Study: Unipolar Correction and Joint Modeling

We perform an additional ablation to demonstrate the relative gains of modeling unipolar supervision sources and jointly modeling accuracies across multiple tasks with respect to the data programming (DP) baseline [27]. Results of this investigation are presented in Table 6. We observe an average improvement of 2.8 points using the unipolar correction (DP-UI), and an additional 1.3 points from joint modeling within MeTaL, resulting in an aggregate gain of 4.1 accuracy points over the data programming baseline.

Table 5:

Label model task accuracies for each task for for both our approach and majority vote (MeTaL/MV)

	OntoNotes	TACRED	OpenI
<u>Task 1</u>			
MV	93.3	74.2	83.9
MeTaL	91.9	80.5	84.1
<u>Task 2</u>			
MV	73.3	46.2	77.8
MeTaL	75.6	65.9	83.7
<u>Task 3</u>			

	OntoNotes	TACRED	OpenI
MV	71.4	74.9	61.7
MeTaL	74.1	74.8	61.7
Task 4			
MV	-	34.4	-
MeTaL	-	60.2	-
Task 5			
MV	-	36.2	-
MeTaL	-	40.2	-
Task 6			
MV	-	56.3	-
MeTaL	-	49.9	-
Task 6			
MV	-	36.8	-
MeTaL	-	56.3	-

Table 6:
Effect of Unipolar Correction.

We compare the micro accuracy (avg. over 10 trials) with 95% confidence intervals of a model trained using data programming (DP), data program with a unipolar correction (DP-UI), and our approach (MeTaL).

	OntoNotes (NER)	TACRED (RE)	OpenI (Doc)	Average
DP [28]	78.4 ± 1.2	49.0 ± 2.7	75.8 ± 0.9	67.7
DP-UI	81.0 ± 1.2	54.2 ± 2.6	76.4 ± 0.5	70.5
MeTaL	82.2 ± 0.8	56.7 ± 2.1	76.6 ± 0.4	71.8

References

- [1]. Anandkumar A, Ge R, Hsu D, Kakade SM, and Telgarsky M Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- [2]. Bach SH, He B, Ratner AJ, and Ré C Learning the structure of generative models without labeled data, 2017.
- [3]. Bhaskara A, Charikar M, and Vijayaraghavan A Uniqueness of tensor decompositions with applications to polynomial identifiability, 2014.
- [4]. Blum A and Mitchell T Combining labeled and unlabeled data with co-training, 1998.
- [5]. Caruana R Multitask learning: A knowledge-based source of inductive bias, 1993.
- [6]. Chandrasekaran V, Parrilo PA, and Willsky AS Latent variable graphical model selection via convex optimization. In *Communication, Control, and Computing (Allerton)*, 2010 48th Annual Allerton Conference on, pages 1610–1613. IEEE, 2010.
- [7]. Chung FRK Laplacians of graphs and cheeger inequalities. 1996.
- [8]. Craven M and Kumlien J Constructing biological knowledge bases by extracting information from text sources, 1999.

- [9]. Dalvi N, Dasgupta A, Kumar R, and Rastogi V Aggregating crowdsourced binary ratings, 2013.
- [10]. Dalvi N, Dasgupta A, Kumar R, and Rastogi V Aggregating crowdsourced binary ratings, 2013.
- [11]. Dawid AP and Skene AM Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied statistics*, pages 20–28, 1979.
- [12]. Ghosh A, Kale S, and McAfee P Who moderates the moderators?: Crowdsourcing abuse detection in user-generated content, 2011.
- [13]. Guan MY, Gulshan V, Dai AM, and Hinton GE Who said what: Modeling individual labelers improves classification. *arXiv preprint arXiv:1703.08774*, 2017.
- [14]. Gupta S and Manning CD Improved pattern learning for bootstrapped entity extraction., 2014.
- [15]. Hoffmann R, Zhang C, Ling X, Zettlemoyer L, and Weld DS Knowledge-based weak supervision for information extraction of overlapping relations, 2011.
- [16]. Honorio J Lipschitz parametrization of probabilistic graphical models. *arXiv preprint arXiv:1202.3733*, 2012.
- [17]. Karger DR, Oh S, and Shah D Iterative learning for reliable crowdsourcing systems, 2011.
- [18]. Khetan A, Lipton ZC, and Anandkumar A Learning from noisy singly-labeled data. *arXiv preprint arXiv:1712.04577*, 2017.
- [19]. Koller D, Friedman N, and Bach F *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [20]. Kruskal JB Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear algebra and its applications*, 18(2):95–138, 1977.
- [21]. Liang P, Jordan MI, and Klein D Learning from measurements in exponential families, 2009.
- [22]. Loh P-L and Wainwright MJ Structure estimation for discrete graphical models: Generalized covariance matrices and their inverses, 2012.
- [23]. Mann GS and McCallum A Generalized expectation criteria for semi-supervised learning with weakly labeled data. *JMLR*, 11(Feb):955–984, 2010.
- [24]. Mintz M, Bills S, Snow R, and Jurafsky D Distant supervision for relation extraction without labeled data, 2009.
- [25]. National Institutes of Health. Open-i. 2017.
- [26]. Platanios E, Poon H, Mitchell TM, and Horvitz EJ Estimating accuracy from unlabeled data: A probabilistic logic approach, 2017.
- [27]. Ratner A, Bach S, Ehrenberg H, Fries J, Wu S, and Ré C Snorkel: Rapid training data creation with weak supervision, 2018.
- [28]. Ratner AJ, De Sa CM, Wu S, Selsam D, and Ré C Data programming: Creating large training sets, quickly, 2016.
- [29]. Ruder S An overview of multi-task learning in deep neural networks. *CoRR*, abs/1706.05098, 2017.
- [30]. Søgaard A and Goldberg Y Deep multi-task learning with low level tasks supervised at lower layers, 2016.
- [31]. Takamatsu S, Sato I, and Nakagawa H Reducing wrong labels in distant supervision for relation extraction, 2012.
- [32]. Tropp JA An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1–2):1–230, 2015.
- [33]. Varma P, He BD, Bajaj P, Khandwala N, Banerjee I, Rubin D, and Ré C Inferring generative model structure with static analysis, 2017.
- [34]. Weischedel R, Hovy E, Marcus M, Palmer M, Belvin R, Pradhan S, Ramshaw L, and Xue N Ontonotes: A large training corpus for enhanced processing. *Handbook of Natural Language Processing and Machine Translation*. Springer, 2011.
- [35]. Xiao T, Xia T, Yang Y, Huang C, and Wang X Learning from massive noisy labeled data for image classification, 2015.
- [36]. Zaidan OF and Eisner J Modeling annotators: A generative approach to learning from annotator rationales, 2008.

- [37]. Zhang C, Ré C, Cafarella M, De Sa C, Ratner A, Shin J, Wang F, and Wu S DeepDive: Declarative knowledgebase construction. *Commun. ACM*, 60(5):93–102, 2017.
- [38]. Zhang Y, Chen X, Zhou D, and Jordan MI Spectral methods meet em: A provably optimal algorithm for crowdsourcing, 2014.
- [39]. Zhang Y, Zhong V, Chen D, Angeli G, and Manning CD Position-aware attention and supervised data improve slot filling, 2017.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 1:

A schematic of the MeTaL pipeline. To generate training data for an *end model*, such as a multi-task model as in our experiments, the user inputs a *task graph* G_{task} defining the relationships between *task labels* Y_1, \dots, Y_i ; a set of *unlabeled data points* X ; a set of *multi-task weak supervision sources* s_i which each output a vector λ_i of task labels for X ; and the dependency structure between these sources, G_{source} . We train a *label model* to learn the accuracies of the sources, outputting a vector of probabilistic training labels \tilde{Y} for training the end model.

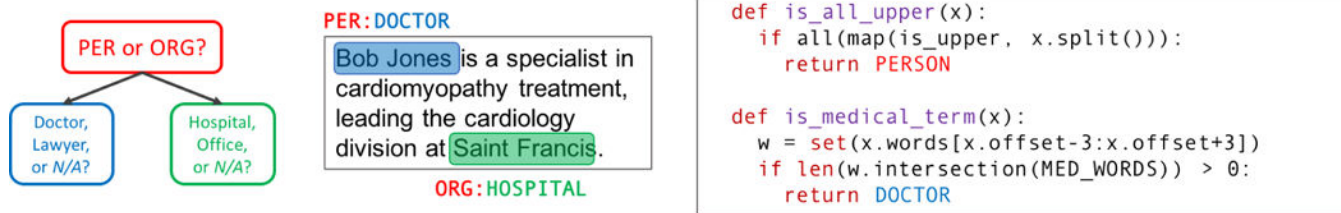


Figure 2:

An example fine-grained entity classification problem, where weak supervision sources label three sub-tasks of different granularities: (i) Person vs. Organization, (ii) Doctor vs. Lawyer (or *N/A*), (iii) Hospital vs. Office (or *N/A*). The example weak supervision sources use a pattern heuristic and dictionary lookup respectively.



Figure 3:

An example of a weak supervision source dependency graph G_{source} (left) and its junction tree representation (right), where \mathbf{Y} is a vector-valued random variable with a feasible set of values, $\mathbf{Y} \in \mathcal{Y}$. Here, the output of sources 1 and 2 are modeled as dependent conditioned on \mathbf{Y} . This results in a junction tree with singleton separator sets, \mathbf{Y} . Here, the observable cliques are $O = \{\lambda_1, \lambda_2, \lambda_3, \lambda_4, \{\lambda_1, \lambda_2\}\} \subset \mathcal{E}$.




Figure 4:

(Left) Estimation error $\|\hat{\mu} - \mu^*\|$ decreases with increasing n . (Middle) Given G_{source} , our model successfully recovers the source accuracies even with many pairwise dependencies among sources, where a naive conditionally-independent model fails. (Right) The runtime of MeTaL is independent of n after an initial matrix multiply, and can thus be multiple orders of magnitude faster than Gibbs sampling-based approaches [28].




Figure 5: In the OntoNotes dataset, end model accuracy scales with the amount of available *unlabeled* data.

	# Train	LM	EM	<i>Gain</i>
NER	62,547	75.2	82.2	7.0
RE	9,090	55.3	57.4	2.1
Doc	2,630	75.6	76.6	1.0

Figure 6:

Using the label model (LM) predictions directly versus using an end model trained on them (EM).

Table 1:
Performance Comparison of Different Supervision Approaches.

We compare the micro accuracy (avg. over 10 trials) with 95% confidence intervals of an end multi-task model trained using the training labels from the hand-labeled development set (Gold Dev), hierarchical majority vote (MV), data programming (DP), and our approach (MeTaL).

	NER	RE	Doc	Average
Gold (Dev)	63.7 ± 2.1	28.4 ± 2.3	62.7 ± 4.5	51.6
MV	76.9 ± 2.6	43.9 ± 2.6	74.2 ± 1.2	65.0
DP [28]	78.4 ± 1.2	49.0 ± 2.7	75.8 ± 0.9	67.7
MeTaL	82.2 ± 0.8	56.7 ± 2.1	76.6 ± 0.4	71.8

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript