
Regret Bounds for Non-decomposable Metrics with Missing Labels

Nagarajan Natarajan
Microsoft Research, INDIA
t-nanata@microsoft.com

Prateek Jain
Microsoft Research, INDIA
prajain@microsoft.com

Abstract

We consider the problem of recommending relevant labels (items) for a given data point (user). In particular, we are interested in the practically important setting where the evaluation is with respect to non-decomposable (over labels) performance metrics like the F_1 measure, *and* training data has missing labels. To this end, we propose a generic framework that given a performance metric Ψ , can devise a regularized objective function and a threshold such that all the values in the predicted score vector above and only above the threshold are selected to be positive. We show that the regret or generalization error in the given metric Ψ is bounded ultimately by estimation error of certain underlying parameters. In particular, we derive regret bounds under three popular settings: a) collaborative filtering, b) multilabel classification, and c) PU (positive-unlabeled) learning. For each of the above problems, we can obtain precise non-asymptotic regret bound which is small even when a large fraction of labels is missing. Our empirical results on synthetic and benchmark datasets demonstrate that by explicitly modeling for missing labels and optimizing the desired performance metric, our algorithm indeed achieves significantly better performance (like F_1 score) when compared to methods that do not model missing label information carefully.

1 Introduction

Predicting relevant labels/items for a given data point is by now a standard task with applications in several domains like recommendation systems [Koren et al., 2009], document tagging, image tagging [Prabhu and Varma, 2014], etc. In many problem settings, e.g. collaborative filtering, features for the data points might not be available and one needs to predict labels only on the basis of past labels (e.g., existing likes/dislikes for various labels/items). In the presence of features, the prediction problem is the standard multi-label classification problem.

Design and analysis of algorithms for such tasks should counter two fundamental challenges: a) in practical scenarios, desired performance metrics for predictions are typically complex *non-decomposable* functions such as F_1 score or precision@ k ; standard metrics like Hamming loss or RMSE over the labels may not be useful, and b) any realistic system in this domain should be able to handle missing labels. Furthermore, often the locations of missing labels may not be available; an important special case is the positive-unlabeled learning setting [Hsieh et al., 2015]. Dealing with missing labels may necessitate imposition of certain regularization on the parameters such as low-rank regularization so as to exploit the correlations between labels.

Most of the existing solutions address only one of the two aspects. For example, Koyejo et al. [2015] establish that for a large class of performance metrics, the optimal solution is to compute a score vector over all the labels and selecting all the labels with score greater than a constant. Their algorithm treats each label as independent to estimate class-conditional probability separately for each label. Clearly, such methods ignore available information about other labels, and hence cannot handle

missing information effectively. Also, such methods do not even apply to the collaborative filtering setting. On the other hand, most of the existing collaborative filtering/matrix completion methods focus only on decomposable losses like RMSE, sum of logistic loss [Lafond, 2015, Yu et al., 2014], which are not effective in real-world systems with large number of labels [Prabhu and Varma, 2014].

In this work, we devise a simple and generic framework that addresses both the aforementioned issues; the framework leads to simple and efficient algorithms in several different settings and for a wide variety of performance metrics used in practice including the multi-label F -measure. Our framework is motivated by a simple observation that has been used in other contexts as well [Kotłowski and Dembczyński, 2015, Koyejo et al., 2015]: for a large class of metrics Ψ , simply thresholding the class probability vector leads to bayes-optimal estimators. Hence, the goal would be to estimate per-label class probabilities accurately. To this end, we show that using a strongly proper loss along with appropriate thresholding leads to bounded regret wrt. Ψ (Theorem 1), and that the threshold can be learned using cross-validation over a small fraction of the training data.

In more detail, strong convexity of the loss function ensures that by minimizing a nuclear-norm regularized ERM (with risk measured by the selected loss function) wrt. a parameter matrix $W \in \mathbb{R}^{d \times L}$, we can bound the regret in Ψ by regret in estimation of the optimal W (Theorem 1); here, d is the dimensionality of the data and is equal to number of users in case of recommender system. Hence, this result allows us to focus on estimation of W^* in various settings such as: a) one-bit matrix completion (Theorem 2), popularly used in recommender systems with only like/dislike information, b) one-bit matrix completion with PU learning (Theorem 4) applicable to recommender systems where only “likes” or positive feedback is observed, and c) general multi-label learning with missing labels (Theorem 3).

For one-bit matrix completion (and the related PU setting), we obtain our final regret bound by adapting existing results from Lafond [2015] and Hsieh et al. [2015], respectively. For general multilabel setting, a direct application of existing results, such as [Lafond, 2015] leads to *weak* bounds. A main technical contribution of our work is to analyze the parameter estimation problem in this setting and provide tight regret bounds. In fact, our result strictly generalizes the result by Lafond [2015], which is for matrix completion with exponential family noise, to the general *inductive matrix completion* setting [Jain and Dhillon, 2013] with exponential family noise. Hence, it should have applications beyond our framework as well. Finally, we illustrate our framework and algorithms on synthetic as well as real-world datasets. Our method exhibits significant improvement over a natural extension of the method by Koyejo et al. [2015] that optimizes Ψ directly but ignores label correlations, hence does not handle missing labels in a principled manner. For example, our method achieves 12% higher F_1 -measure on a benchmark dataset than that by Koyejo et al. [2015].

Related Work. We now highlight some related theoretical work in recommender systems and multi-label learning. Gao and Zhou [2013] study consistency and surrogate losses for two specific losses namely Hamming and expected (partial) ranking losses, and leave the other losses to future work. Dembczynski et al. [2012] consider expected pairwise ranking loss in multilabel learning, show that the problem decomposes into independent binary problems, and provide regret bound for the same. Yun et al. [2014] consider the learning-to-rank problem, where the goal is to rank the relevant labels for a given instance. They show that popular ranking losses like NDCG can be written as a generalization of certain robust binary loss functions, although they do not provide any explicit regret bounds. Existing theoretical guarantees for 1-bit matrix completion methods used in recommender systems focus solely on RMSE or 0-1 loss [Lafond, 2015, Hsieh et al., 2015].

2 Problem Setup and Background

Let $\mathbf{x}_i \in \mathcal{X} \subseteq \mathbb{R}^d$ denote instances and $\mathbf{y}_i \in \{0, 1\}^L$ denote label vectors. Let $Y \in \{0, 1\}^{n \times L}$ denote the label matrix, with \mathbf{y}_i 's as rows. In typical multi-label learning and recommender system settings a) the labeling process has some inherent uncertainty, which is usually captured by assuming a conditional distribution $\mathbb{P}(\mathbf{y}_i | \mathbf{x}_i)$, b) furthermore, we do not get to observe all the entries of \mathbf{y}_i , but only a small subset, say Ω_i . Formally, let $\Omega \subset [n] \times [L]$ denote a subset of indices sampled i.i.d. from a fixed distribution π over $[n] \times [L]$. We consider the following sampling model for observing labels:

$$Y_{ij} = \begin{cases} 1 & \text{with probability } g_j(\mathbf{x}_i; W^*) \\ 0 & \text{with probability } 1 - g_j(\mathbf{x}_i; W^*) \end{cases} \text{ for } (i, j) \in \Omega. \quad (1)$$

where \mathbf{W}^* parameterizes the underlying conditional distribution $\mathbb{P}(\mathbf{y}_i|\mathbf{x}_i)$. Following the low-rank *inductive* matrix completion model [Yu et al., 2014, Zhong et al., 2015], we let $\mathbf{W}^* \in \mathbb{R}^{d \times L}$ be the parameter matrix and $g_j(\mathbf{x}_i; \mathbf{W}^*) = g(\langle \mathbf{x}_i, \mathbf{w}_j^* \rangle)$ where \mathbf{w}_j^* is the j th column of \mathbf{W}^* corresponding to the j th label, for some differentiable function $g: \mathbb{R} \rightarrow [0, 1]$. A popular choice of g is given by $g(\langle \mathbf{x}_i, \mathbf{w}_j \rangle) = \frac{\exp(\langle \mathbf{x}_i, \mathbf{w}_j \rangle)}{1 + \exp(\langle \mathbf{x}_i, \mathbf{w}_j \rangle)}$, which corresponds to the logistic regression model. When we do not observe feature vectors \mathbf{x} , as in the classical recommender system or matrix completion setting, the above model (1) reduces to the widely studied 1-bit matrix completion model [Cai and Zhou, 2013, Davenport et al., 2014]:

$$\mathbf{Y}_{ij} = \begin{cases} 1 & \text{with probability } g(\mathbf{W}_{ij}^*) \\ 0 & \text{with probability } 1 - g(\mathbf{W}_{ij}^*) \end{cases} \text{ for } (i, j) \in \Omega, \quad (2)$$

where $\mathbf{W}^* \in \mathbb{R}^{n \times L}$ is the parameter matrix that captures user-item preferences.

The goal is to learn a multi-label classifier $f: \mathcal{X}^n \rightarrow \{0, 1\}^{n \times L}$ jointly over n instances, such that a performance metric of interest Ψ is maximized. The training data consists of input features $\mathbf{X} \in \mathbb{R}^{n \times d}$ where each row corresponds to an instance, drawn iid from some distribution $\mathbb{P}_{\mathcal{X}}$ over \mathcal{X} , and *partially observed* label matrix \mathbf{Y} using the sampling model (1) or (2). In this work, we consider a large family of non-decomposable metrics [Koyejo et al., 2015] that constitutes linear-fractional functions of (multi-label analogues of) true positives, false positives, false negatives and true negatives defined below. Let $\hat{\mathbf{Y}} \in \{0, 1\}^{n \times L}$ denote the predicted labels, i.e. $\hat{\mathbf{Y}} = f(\mathbf{X})$ for some f . Define the primitives:

$$\begin{aligned} \widehat{\text{TP}}_{ij}(\hat{\mathbf{Y}}, \mathbf{Y}) &= [[\hat{\mathbf{Y}}_{ij} = 1, \mathbf{Y}_{ij} = 1]], & \widehat{\text{FP}}_{ij}(\hat{\mathbf{Y}}, \mathbf{Y}) &= [[\hat{\mathbf{Y}}_{ij} = 1, \mathbf{Y}_{ij} = 0]], \\ \widehat{\text{TN}}_{ij}(\hat{\mathbf{Y}}, \mathbf{Y}) &= [[\hat{\mathbf{Y}}_{ij} = 0, \mathbf{Y}_{ij} = 0]], & \widehat{\text{FN}}_{ij}(\hat{\mathbf{Y}}, \mathbf{Y}) &= [[\hat{\mathbf{Y}}_{ij} = 0, \mathbf{Y}_{ij} = 1]]. \end{aligned}$$

For convenience, we drop the arguments and just write $\widehat{\text{TP}}_{ij}$ to denote $\widehat{\text{TP}}_{ij}(\hat{\mathbf{Y}}, \mathbf{Y})$ and so on.

1. **Micro-averaged metrics.** Define: $\widehat{\text{TP}}(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{|\Omega|} \sum_{(i,j) \in \Omega} \widehat{\text{TP}}_{ij}$ and $\widehat{\text{FP}}(\hat{\mathbf{Y}}, \mathbf{Y}), \widehat{\text{TN}}(\hat{\mathbf{Y}}, \mathbf{Y}), \widehat{\text{FN}}(\hat{\mathbf{Y}}, \mathbf{Y})$ similarly. Let $\text{TP} = \mathbb{E}[\widehat{\text{TP}}], \text{FP} = \mathbb{E}[\widehat{\text{FP}}]$ (and so on), where the expectation is defined wrt to the sampling distribution π over indices $[n] \times [L]$ as well as the joint distribution \mathbb{P} . Micro-averaged performance metric $\Psi: \{0, 1\}^{n \times L} \times \{0, 1\}^{n \times L} \rightarrow \mathbb{R}_+$ is given by:

$$\Psi(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{a_0 + a_{11}\text{TP} + a_{01}\text{FP} + a_{10}\text{FN} + a_{00}\text{TN}}{b_0 + b_{11}\text{TP} + b_{01}\text{FP} + b_{10}\text{FN} + b_{00}\text{TN}}. \quad (3)$$

for bounded constants a 's and b 's. Assume that Ψ is bounded, i.e. $\exists \gamma' > 0$ such that $b_0 + b_{11}\text{TP} + b_{01}\text{FP} + b_{10}\text{FN} + b_{00}\text{TN} > \gamma'$ for all $\hat{\mathbf{Y}}, \mathbf{Y}$.

2. **Instance-averaged metrics.** Define $\widehat{\text{TP}}_i(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{|\Omega_i|} \sum_{j \in \Omega_i} \widehat{\text{TP}}_{ij}$. Let $\text{TP}_i = \mathbb{E}[\widehat{\text{TP}}_i]$. Instance-averaged performance metric Ψ is given by:

$$\Psi(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n \frac{a_0 + a_{11}\text{TP}_i + a_{01}\text{FP}_i + a_{10}\text{FN}_i + a_{00}\text{TN}_i}{b_0 + b_{11}\text{TP}_i + b_{01}\text{FP}_i + b_{10}\text{FN}_i + b_{00}\text{TN}_i}. \quad (4)$$

for bounded constants a 's and b 's. Assume that Ψ is bounded, i.e. $\exists \gamma' > 0$ such that $b_0 + b_{11}\text{TP}_i + b_{01}\text{FP}_i + b_{10}\text{FN}_i + b_{00}\text{TN}_i > \gamma'$ for all $\hat{\mathbf{Y}}, \mathbf{Y}, i$.

3. **Macro-averaged metrics.** Let $\Omega^{(j)} = \{i : (i, j) \in \Omega\}$. Define: $\widehat{\text{TP}}_j(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{|\Omega^{(j)}|} \sum_{i \in \Omega^{(j)}} \widehat{\text{TP}}_{ij}$. Let $\text{TP}_j = \mathbb{E}[\widehat{\text{TP}}_j]$. Macro-averaged performance metric Ψ is given by:

$$\Psi(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{L} \sum_{j=1}^L \frac{a_0 + a_{11}\text{TP}_j + a_{01}\text{FP}_j + a_{10}\text{FN}_j + a_{00}\text{TN}_j}{b_0 + b_{11}\text{TP}_j + b_{01}\text{FP}_j + b_{10}\text{FN}_j + b_{00}\text{TN}_j}. \quad (5)$$

for bounded constants a 's and b 's. Assume that Ψ is bounded, i.e. $\exists \gamma' > 0$ such that $b_0 + b_{11}\text{TP}_j + b_{01}\text{FP}_j + b_{10}\text{FN}_j + b_{00}\text{TN}_j > \gamma'$ for all $\hat{\mathbf{Y}}, \mathbf{Y}, j$.

Example metrics:

$$1. \text{ Instance-averaged } F_1 \text{ metric defined as: } \Psi_{F_1}(\hat{\mathbf{Y}}, \mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n \frac{2\text{TP}_i}{2\text{TP}_i + \text{FP}_i + \text{FN}_i}.$$

2. Accuracy (equivalent to the Hamming loss): $\Psi_{\text{Ham}}(\hat{\mathbf{Y}}, \mathbf{Y}) = 1 - \frac{1}{n} \sum_{i=1}^n \text{FP}_i + \text{FN}_i$.

Remark 1. *The aforementioned definitions of performance metrics naturally apply to the recommender system setting, where data is observed via the 1-bit matrix completion sampling model (2). Note that in this case, the expectations are defined wrt the sampling distribution π and the inherent noise in 1-bit sampling $\mathbb{P}(Y_{ij}|W_{ij}^*)$.*

Let Ψ^* denote the Bayes optimal performance, i.e. $\Psi^* = \max_f \Psi(f(\mathbf{X}), \mathbf{Y})$. Our objective can be now stated learning \hat{f} such that the Ψ -regret, i.e. $\Psi^* - \Psi(\hat{f}(\mathbf{X}), \mathbf{Y})$, is provably bounded. Koyejo et al. [2015] showed that the Bayes optimal Ψ^* thresholds the conditional probability of each label j , i.e. $\mathbb{P}(y_j|\mathbf{x})$ at a certain value $\delta^* \in (0, 1)$, and that the value δ^* is shared across all the labels.¹

3 Algorithm

Our approach is based on estimating real-valued predictions and then thresholding the predictions optimally in order to maximize a given metric Ψ . Koyejo et al. [2015] proposed a simple consistent plug-in estimator algorithm, which first computes conditional marginals $\mathbb{P}(y_j|\mathbf{x})$ independently for each label j , and then estimates a threshold jointly to optimize Ψ . While the approach is provably consistent asymptotically, it is not clear if it admits a useful regret bound; in particular, we would like to characterize the behavior in the finite samples regime. In case of the sampling model (1), the approach translates to learning columns of the parameter matrix \mathbf{W}^* independently. In many cases, \mathbf{W} exhibits some structure, such as low-rankness, reflecting correlation between labels [Yu et al., 2014, Zhong et al., 2015, Davenport et al., 2014]. Statistically, capturing correlations via a low-rank structure could help improve the sample complexity for recovery, and computationally, it would help reduce space and time complexity of the learning procedure.

Our proposed algorithm is presented in Algorithm 1. In Step 1, we solve a trace-regularized minimization problem to estimate the parameter matrix \mathbf{W}^* , where the function ℓ can be any bounded loss such as the squared, the logistic or the squared Hinge loss. In particular, using the logistic loss corresponds to the maximum likelihood estimation of the sampling model (1). Yu et al. [2014] also solve essentially the same objective as (6), except for the additional bound constraint on entries of $\mathbf{X}\mathbf{W}$. The optimization problem (6) can be solved using a proximal gradient descent algorithm, with a fast proximal operator computation by storing the current solution in a low-rank form. We could also use fast non-convex procedure, by writing $\mathbf{W} = \mathbf{W}_1\mathbf{W}_2^T$, where \mathbf{W}_1 and \mathbf{W}_2 are low-rank matrices with $k \ll \min(d, L)$ columns each, and applying alternating minimization.

The real-valued estimator is given by $\mathbf{Z} = \mathbf{X}\hat{\mathbf{W}}$ in Step 2. To obtain binary-valued predictions, we solve a 1-dimensional optimization problem to compute the optimal threshold, on the training data. Note that this step can be done in $|\Omega|$ time.

Remark 2. *In the 1-bit matrix completion setting, we obtain a thresholded max-likelihood estimator of $\mathbf{W}^* \in \mathbb{R}^{n \times L}$ using identical procedure; where we interpret \mathbf{X} in Algorithm 1 as the identity matrix of size n .*

4 Analysis: Regret Bounds

Here, we first show that Ψ -regret can be bounded with the regret of a certain loss ℓ . Then, under various sampling models pertaining to different settings, namely, 1-bit matrix completion, multi-label learning, and PU (positive-unlabeled) learning, we show that the ℓ -regret can be bounded via recovering the underlying parameter matrix \mathbf{W}^* governing $\mathbb{P}(\mathbf{y}_i|\mathbf{x}_i)$.

4.1 Low ℓ -regret implies low Ψ -regret

Our first main result connects Ψ -regret to regret with respect to a strongly proper loss function ℓ [Agarwal, 2014]. Canonical examples of strongly proper losses include the logistic loss $\ell(t, y) = \log(1 + \exp(-(2y - 1)t))$, the exponential loss $\ell(t, y) = \exp(-(2y - 1)t)$ and the squared loss

¹The definitions in [Koyejo et al., 2015] do not include general sampling distribution π , but the results can be generalized in a straight-forward manner.

Algorithm 1 Thresholded Max-Likelihood Estimator

Input: Training data $\mathbf{X} \in \mathbb{R}^{n \times d}$, labels $\mathbf{Y}_\Omega \in \{0, 1\}^{n \times L}$ observed on indices Ω , metric Ψ , parameters γ, λ .

1. Obtain $\hat{\mathbf{W}}$ by solving the trace-constrained matrix completion:

$$\hat{\mathbf{W}} = \arg \min_{\mathbf{W}: \|\mathbf{X}\mathbf{W}\|_\infty \leq \gamma} \frac{1}{|\Omega|} \sum_{(i,j) \in \Omega} \ell(\langle \mathbf{x}_i, \mathbf{w}_j \rangle, Y_{ij}) + \lambda \|\mathbf{W}\|_*, \quad (6)$$

2. Let $\mathbf{Z} = \mathbf{X}\hat{\mathbf{W}}$. Define the thresholding operator $\hat{\mathbf{Y}} = \text{Thr}_\theta(\mathbf{Z})$, such that $\hat{Y}_{ij} = \llbracket \mathbf{Z}_{ij} \geq \theta \rrbracket$.

3. Let $\hat{\Psi}$ denote the empirical metric computed on a finite sample; i.e., (3), (4), (5) defined wrt. $\widehat{\text{TP}}_{ij}, \widehat{\text{FP}}_{ij}, \widehat{\text{TN}}_{ij}$ and $\widehat{\text{FN}}_{ij}$.

4. Return $\hat{\mathbf{Y}} = \text{Thr}_{\hat{\theta}}(\mathbf{Z})$, where

$$\hat{\theta} = \arg \max_{\theta} \hat{\Psi}(\text{Thr}_\theta(\mathbf{Z}_\Omega), \mathbf{Y}_\Omega),$$

$\ell(t, y) = (1 - (2y - 1)t)^2$, for $y \in \{0, 1\}$ and $t \in \mathbb{R}$. Define the ℓ -regret of $\mathbf{Z} \in \mathbb{R}^{n \times L}$ as:

$$\text{Reg}_\ell(\mathbf{Z}) = \mathbb{E}[\ell(\mathbf{Z}_{ij}, Y_{ij})] - \min_{\mathbf{Z}' \in \mathbb{R}^{n \times L}} \mathbb{E}[\ell(\mathbf{Z}'_{ij}, Y_{ij})],$$

where the expectation is wrt. draws from π and the joint distribution \mathbb{P} over instances and labels.

Theorem 1 (Main Result 1). *Let Ψ be a performance metric as defined in (3), (4) or (5). Let ℓ be a λ' -strongly proper loss function. Assume the input $\mathbf{X} \in \mathbb{R}^{n \times L}$ consists of iid instances sampled from marginal $\mathbb{P}_\mathcal{X}$, label matrix $\mathbf{Y}_\Omega \in \{0, 1\}^{n \times L}$, where Y_{ij} is sampled iid from $\mathbb{P}(Y_{ij} | \mathbf{x}_i)$, observed on a subset of indices Ω sampled iid from a fixed distribution π . Then, the output $\hat{\mathbf{Y}}$ obtained by thresholding the estimate \mathbf{Z} in Step 4 of Algorithm 1 satisfies the regret bound:*

$$\Psi^* - \Psi(\hat{\mathbf{Y}}, \mathbf{Y}) \leq C \sqrt{\frac{2}{\lambda'}} \sqrt{\text{Reg}_\ell(\mathbf{Z})} + O\left(\frac{1}{\sqrt{|\Omega|}}\right), \quad (7)$$

for some positive constant C .

We emphasize that the above result holds for arbitrary metric Ψ from the family (3), (4) or (5). Consider the RHS of (7): $1/\sqrt{|\Omega|}$ is the lower-order term, and independent of dimensionality; the first term makes the framework fairly powerful, as it can use any strongly proper loss. In the next subsection, we will provide precise instantiations of this term under various learning settings.

Proof Outline for Theorem 1. Proof technique is based on [Kotłowski and Dembczyński, 2015], where they derive similar bound in the binary classification setting. We first relate the Ψ -regret to weighted 0-1 loss regret (Lemma 2). Then, we show there exists a threshold θ^* such that the weighted loss regret of $\text{Thr}_{\theta^*}(\mathbf{Z}) \in \{0, 1\}^{n \times L}$ is bounded by its ℓ -regret corresponding to any strongly proper loss ℓ (Lemma 3). Finally, we argue that it suffices to estimate $\hat{\theta}$ from the training data (Lemma 4). Detailed proof and associated Lemmas are available in Appendix A.1. \square

4.2 Bounding ℓ -regret

Below, we provide the desired ℓ -regret bound under three different settings.

4.2.1 Collaborative Filtering

Consider the 1-bit matrix completion sampling model in (2). Then (6) reduces to the optimization problem considered by Lafond [2015]. We have the following regret bound for the estimator $\mathbf{Z} = \hat{\mathbf{W}}$ obtained in Step 2 of Algorithm 1 (Note that \mathbf{X} is just treated as identity in this setting).

Theorem 2. *Assume π is uniform, and consider the 1-bit matrix completion sampling model (2). Let ℓ denote a 1-Lipschitz, strongly proper loss (appearing in (7)), and \mathbf{Z} denote the output of Step 2 of*

Algorithm 1 with the choice of $\lambda = 2c_\gamma \sqrt{\frac{2 \log(n+L)}{\min(n,L)|\Omega|}}$. With probability at least $1 - \delta$, the following holds:

$$\text{Reg}_\ell(\mathbf{Z}) \leq \sqrt{\tilde{C} \max\left(\frac{\max(n, L) \text{rank}(W^*) \log(3/\delta)}{|\Omega|} (\sigma_\gamma^2 + 1), \gamma^2 \sqrt{\frac{\log(3/\delta)}{|\Omega|}}\right)},$$

where $\tilde{C}, c_\gamma, \sigma_\gamma$ are numerical constants, and $\gamma = \max_{ij} |W_{ij}^*|$.

Note that when $|\Omega| > \max(n, L)$, the RHS of the above bound converges; in particular, the second term within max is the lower-order term: $\gamma \approx O(\sqrt{1/nL})$. Theorem 2 can be extended to general distributions π beyond uniform, satisfying mild assumptions. See Appendix A.2.

4.2.2 Multi-label Learning

Consider the sampling model (1) with features. We have the following regret bound for the estimator $\mathbf{Z} = \mathbf{X}\hat{\mathbf{W}}$ obtained in Step 2 of Algorithm 1, under the following assumptions.

Assumption 1. *The marginal distribution over the features $\mathbb{P}_{\mathcal{X}}$ is sub-Gaussian with sub-Gaussian norm K and covariance $\Sigma \in \mathbb{R}^{d \times d}$.*

Assumption 2. *Let $\pi_{k,l}$ denote the probability of sampling the entry $(k, l) \in [n] \times [L]$; $\exists \mu \geq 1$ s.t. for all $n, L, \min_{k \in [n], l \in [L]} \pi_{k,l} \geq \frac{1}{\mu n L}$, and $\nu \geq 1$ s.t. $\max_{i', j'} (\sum_j \pi_{i'j}, \sum_i \pi_{ij'}) \leq \frac{\nu}{\min(n, L)}$.*

Theorem 3 (Main Result 2). *Assume 1, 2 and consider the sampling model (1). Also assume $L \geq d$. Let $\hat{\mathbf{W}}$ be the solution to the trace-norm regularized optimization problem (6) using logistic loss for ℓ , number of training data points $n \geq C'd$, number of observations $|\Omega| \geq L + d$, and setting the regularization parameter $\lambda = \frac{2c}{\sqrt{|\Omega|}}$. Then, with probability at least $1 - 3(n+L)^{-1} - 2(d+L)^{-1}$, the following holds:*

$$\frac{\|\hat{\mathbf{W}} - W^*\|_F^2}{dL} \leq \frac{C_2 \mu^2}{d} \max\left(\frac{L \text{rank}(W^*) \log(n+L)}{|\Omega|} (\sigma_\gamma^2 + 1), \frac{\gamma^2}{\mu} \sqrt{\frac{\log(n+L)}{|\Omega|}}\right),$$

where c, C', C_2 are numerical constants, $\sigma_\gamma \leq (1 + e^\gamma)^2 e^\gamma$, $\gamma = \max_{ij} |(XW^*)_{ij}|$ and μ is defined as in Assumption 2.

A few remarks of our result in the multi-label setting are in order:

Remark 3 (Generalization). *The result in Theorem 3, and Theorem 8 in Appendix B for general exponential distributions, is a key technical contribution of this work. In particular, our analysis applies to Y arising from general exponential distributions, including Gaussian when Y is real-valued and Poisson when Y models counts. See Appendix B for more details.*

Remark 4 (Comparing [Lafond, 2015]). *If we directly apply the method and the analysis of [Lafond, 2015], the resulting bounds are very weak; in fact, when $n \geq L$ and $|\Omega| = O(n)$, which is quite common in the multi-label scenario, the ensuing bound suggests that the estimator is not even consistent, even when π is uniform. See Appendix A.3 for details.*

Remark 5 (Comparing [Koyejo et al., 2015]). *The plugin-in estimator algorithm of [Koyejo et al., 2015] estimates \mathbf{w}_j^* for each label j independently, and learns a common threshold as in Step 4 of Algorithm 1. Let $\hat{\mathbf{w}}_j$ denote the estimator for label j . Then, using standard analysis we have, $\|\hat{\mathbf{w}}_j - \mathbf{w}_j^*\|_2 \leq \sigma \sqrt{\frac{d}{|\Omega^j|}}$, where $|\Omega^j|$ is the number of observations per label which is $O(\frac{|\Omega|}{L})$. Thus we have the bound: $\frac{\|W^* - \hat{\mathbf{W}}\|_F^2}{L} \leq \sigma O(\frac{Ld}{|\Omega|})$. This is how our bounds behave, when W^* is indeed full rank, up to constants. When $\text{rank}(W^*) \ll \min(d, L)$, we achieve much faster convergence.*

We now give the desired ℓ -regret bound as a corollary.

Corollary 1. *Assume the conditions and the choices in Theorem 3 hold. Let ℓ denote a 1-Lipschitz, strongly proper loss (appearing in (7)), and $\mathbf{Z} = \mathbf{X}\hat{\mathbf{W}}$ denote the output of Step 2 of Algorithm 1. With probability at least $1 - \delta - (d+L)^{-1}$, the following holds:*

$$\text{Reg}_\ell(\mathbf{Z}) \leq \sqrt{C_2 \mu^2 \max\left(\frac{L \text{rank}(W^*) \log(3/\delta)}{|\Omega|} (\sigma_\gamma^2 + 1), \frac{\gamma^2}{\mu} \sqrt{\frac{\log(3/\delta)}{|\Omega|}}\right)},$$

where $c, C', C_2, \sigma_\gamma, \gamma, \mu$ are defined as in Theorem 3.

Proof Outline for Theorem 3. We analyze the following general exponential noise model for Y :

$$Y_{ij} | \mathbf{x}_i, \mathbf{w}_j \sim \exp_{h,G}(\mathbf{x}_i, \mathbf{w}_j) := h(Y_{ij}) \exp(\langle \mathbf{x}_i, \mathbf{w}_j \rangle Y_{ij} - G(\langle \mathbf{x}_i, \mathbf{w}_j \rangle)), \quad (8)$$

where h and G are the base measure and log-partition functions associated with this canonical representation. Our proof sketch is based on Lafond [2015], but requires bounding certain quantities carefully. In particular, we prove a tight bound for $\|\mathbf{X}^T \nabla \Phi_Y(\mathbf{X}, \mathbf{W}^*)\|_2$ in terms of the regularization parameter λ , where $\Phi_Y(\mathbf{X}, \mathbf{W}^*)$ is the MLE wrt. general exponential distribution (reduces to (6), without regularization, when Y_{ij} 's are from (1)), as stated below.

Lemma 1. *Consider the sampling model (8). Assume (i) $d \leq L$, (ii) $|\Omega| \geq (L + d)$, (iii) Y_{ij} 's are sampled independently given \mathbf{x}_i , and (iv) $|Y_{ij} - G'(\langle \mathbf{x}_i, \mathbf{w}_j^* \rangle)| \leq \alpha$, for all $i, j \in [n] \times [L]$, for any n, L . Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ whose rows (\mathbf{x}_i 's) are iid samples from $\mathbb{P}_{\mathcal{X}}$ satisfying Assumption 1. Then, there exists numerical constant c such that, with probability at least $1 - (d + L)^{-1}$:*

$$\|\mathbf{X}^T \nabla \Phi_Y(\mathbf{X}, \mathbf{W}^*)\|_2 \leq c \cdot \frac{\alpha}{\sqrt{|\Omega|}}.$$

4.2.3 PU Learning

In many collaborative filtering and multi-label learning tasks, only the positive entries ($y_{ij} = 1$) are observed. In this setting, we can use the approach of [Hsieh et al., 2015], where they consider a two-stage sampling model: sample y_{ij} using (2) for all $i, j \in [n] \times [L]$ (or using (1) when features are available), and then flip a fraction ρ of the sampled 1's to 0's, resulting in \tilde{Y} . We would then use the unbiased estimator $\tilde{\ell}$ of loss ℓ in (6); $\tilde{\ell}$ satisfies $\mathbb{E}[\tilde{\ell}(\mathbf{Z}_{ij}, \tilde{Y}_{ij})] = \ell(\mathbf{Z}_{ij}, Y_{ij})$, where the expectation is wrt the flipping process, parameterized by ρ . For the estimator $\mathbf{Z} = \hat{\mathbf{W}}$ obtained thus, we have the following regret bound.

Theorem 4. *Let ℓ denote a 1-Lipschitz, strongly proper loss (appearing in (7)). Assume $\|\mathbf{W}^*\|_* \leq t$. Let $\mathbf{Z} = \mathbf{X}\tilde{\mathbf{W}}$, where $\tilde{\mathbf{W}}$ is obtained by solving the unbiased estimator objective of Hsieh et al. [2015]. There exists a numerical constant C_3 such that, with probability at least $1 - \delta$:*

$$\text{Reg}_\ell(\mathbf{Z}) \leq \sqrt{6 \frac{\sqrt{\log(2/\delta)}}{\sqrt{nL}(1-\rho)} + 2C_3 \cdot t \frac{\sqrt{n} + \sqrt{L}}{(1-\rho)nL}}.$$

The RHS of the bound above, when $n = L$, is of $O(\sqrt{\frac{1}{n(1-\rho)}})$, where $(1 - \rho)$ is the fraction of observed 1's in \tilde{Y} . Naturally, as ρ is large, we need more samples to achieve similar rates as in the other settings.

Remark 6. *This PU learning result is particularly very useful in extreme classification setting [Bhatia et al., 2015a, Prabhu and Varma, 2014]; where there are too many labels and is unrealistic to get feedback on every label, but possible to obtain a small subset of relevant labels for instances. Furthermore, the above result serves to attest to the utility of our framework.*

5 Experiments

We focus on multi-label datasets for experimental study. The goal is to show that the convergence happens as suggested by the theory, and that the proposed algorithm performs well on real-world datasets. To solve (6), we use an alternating minimization procedure by forming $\mathbf{W} = \mathbf{W}_1 \mathbf{W}_2^T$, such that $\mathbf{W}_1 \in \mathbb{R}^{d \times k}$ and $\mathbf{W}_2 \in \mathbb{R}^{L \times k}$, where k , the rank of \mathbf{W} , is an input parameter.

5.1 Synthetic data

We generate multi-label data as follows. We fix $n = 1000, L = 100$ and $d = 10$. First, we form $\mathbf{X} \in \mathbb{R}^{n \times d}$ using iid samples from multi-variate Gaussian $\mathcal{N}(0, I)$. Then, we generate \mathbf{W}^* of rank 5. The label matrix \mathbf{Y} is obtained by thresholding $\mathbf{X}\mathbf{W}^*$ at $\theta^* = 0$, i.e. $Y_{ij} = \text{sign}(\langle \mathbf{x}_i, \mathbf{w}_j^* \rangle)$. In this noise-free setting, we expect that our algorithm would recover both \mathbf{W}^* and θ^* accurately as it sees more and more observations. The results for maximizing micro F_1 and accuracy metrics are presented

in Figure 1. As the sampling ratio $\frac{|\Omega|}{nL}$ increases, we observe that the proposed estimator achieves optimal performance in both the cases. Furthermore, even when only 10% of the observations is revealed, we observe that the proposed method achieves very high F_1 as well as accuracy values, compared to learning the columns of W^* independently via the plugin estimator method proposed by [Koyejo et al., 2015] (followed by learning a threshold).

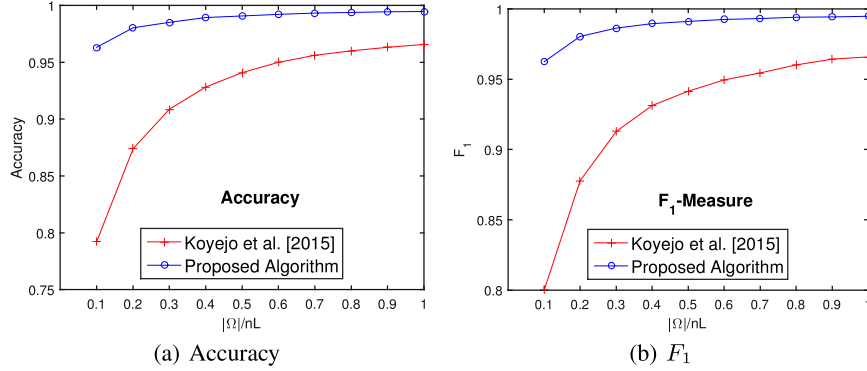


Figure 1: Convergence of the methods for Accuracy and micro- F_1 metrics on synthetic data.

DATASET	Koyejo et al. [2015] micro F_1	Algorithm 1 micro F_1	Koyejo et al. [2015] Accuracy	Algorithm 1 Accuracy
CAL500	0.4267 ± 0.0016	0.3855 ± 0.0005	0.8541 ± 0.0034	0.8493 ± 0.0002
AUTOFOOD	0.4897 ± 0.0103	0.5597 ± 0.0047	0.9307 ± 0.0064	0.9345 ± 0.0043
BIBTEX	0.2641 ± 0.0251	0.2398 ± 0.0133	0.9849 ± 0.0016	0.9856 ± 0.0003
COMPPHYS	0.2463 ± 0.0315	0.3510 ± 0.0293	0.9448 ± 0.0011	0.9466 ± 0.0012
COREL5K	0.1552 ± 0.0116	0.1642 ± 0.0001	0.9906 ± 0.0000	0.9906 ± 0.0000

Table 1: Comparison of proposed algorithm and plugin-estimator method of [Koyejo et al., 2015] on multi-label micro F_1 and Hamming (i.e. Accuracy) metrics. Reported values correspond to *micro-averaged* metric computed on test data. In all the cases, $\frac{|\Omega|}{nL}$ was fixed to 20% for training. The rank of W was set to $0.4L$ for Algorithm 1. We observe that the proposed algorithm which captures label correlations is competitive across datasets.

5.2 Real-world data

We consider five real-world multi-label datasets widely used as benchmarks [Bhatia et al., 2015a, Yu et al., 2014]. (i) CAL500: a music dataset with 400 training and 100 test instances, $L = 174$, $d = 68$, (ii) COREL5K: an image dataset with 4500 training and 500 test instances, $L = 374$, $d = 499$, (iii) BIBTEX: a text dataset with 4,880 training and 2,515 test instances, $L = 159$, $d = 1,836$, (iv) COMPPHYS dataset with 161 training and 40 test instances, $L = 208$, $d = 33,284$, and (v) AUTOFOOD dataset with 4,880 training and 2,515 test instances, $L = 162$, $d = 1,836$.

We set the rank k of W to $0.4L$ for all the datasets in our method, and set $\frac{|\Omega|}{nL} = 20\%$ to train the models in each method. The results are presented in Table 1. We observe that the proposed method is competitive in all the datasets, and achieves better micro- F_1 and accuracy values, with a small value of rank $0.4L$. We note that the label matrices of most of the datasets are very sparse (for instance, less than 8.5% of the test data are positive labels in AUTOFOOD), which explains high accuracy and low F_1 values. The learned model is much more compact than that of [Koyejo et al., 2015] ($k(d+L)$ vs dL parameters). While in theory our bounds hold for the case $L \geq d$ (Theorem 3), some of the datasets considered here have $d \gg L$ and yet the performance is competitive.

6 Conclusions

We presented a framework for optimizing general performance metrics applicable to multi-label as well as collaborative filtering settings. Our work complements recent results in this direction: on the theoretical front, we derive strong regret bounds for practically used metrics like F -measure, and on the algorithmic front, we provide simple and efficient procedure that works well in practice.

References

- Shivani Agarwal. Surrogate regret bounds for bipartite ranking via strongly proper losses. *The Journal of Machine Learning Research*, 15(1):1653–1674, 2014.
- Kush Bhatia, Himanshu Jain, Purushottam Kar, Manik Varma, and Prateek Jain. Sparse local embeddings for extreme multi-label classification. In *Advances in Neural Information Processing Systems*, pages 730–738, 2015a.
- Kush Bhatia, Prateek Jain, and Purushottam Kar. Robust regression via hard thresholding. In *Advances in Neural Information Processing Systems*, pages 721–729, 2015b.
- Tony Cai and Wen-Xin Zhou. A max-norm constrained minimization approach to 1-bit matrix completion. *The Journal of Machine Learning Research*, 14(1):3619–3647, 2013.
- Mark A Davenport, Yaniv Plan, Ewout van den Berg, and Mary Wootters. 1-bit matrix completion. *Information and Inference*, 3(3):189–223, 2014.
- Krzysztof Dembczynski, Wojciech Kotłowski, and Eyke Hüllermeier. Consistent multilabel ranking through univariate losses. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*, 2012.
- Wei Gao and Zhi-Hua Zhou. On the consistency of multi-label learning. *Artificial Intelligence*, 199: 22–44, 2013.
- Cho-jui Hsieh, Nagarajan Natarajan, and Inderjit Dhillon. PU learning for matrix completion. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2445–2453, 2015.
- Prateek Jain and Inderjit S Dhillon. Provable inductive matrix completion. *arXiv preprint arXiv:1306.0626*, 2013.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- Wojciech Kotłowski and Krzysztof Dembczyński. Surrogate regret bounds for generalized classification performance metrics. *arXiv preprint arXiv:1504.07272*, 2015.
- Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. Consistent binary classification with generalized performance metrics. In *Advances in Neural Information Processing Systems 27*, pages 2744–2752. 2014.
- Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. Consistent multilabel classification. In *Advances in Neural Information Processing Systems*, pages 3303–3311, 2015.
- Jean Lafond. Low rank matrix completion with exponential family noise. *arXiv preprint arXiv:1502.06919*, 2015.
- Yashoteja Prabhu and Manik Varma. Fastxml: a fast, accurate and stable tree-classifier for extreme multi-label learning. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 263–272. ACM, 2014.
- Mark D Reid and Robert C Williamson. Composite binary losses. *The Journal of Machine Learning Research*, 11:2387–2422, 2010.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit Dhillon. Large-scale multi-label learning with missing labels. In *Proceedings of The 31st International Conference on Machine Learning*, pages 593–601, 2014.
- Hyokun Yun, Parameswaran Raman, and S Vishwanathan. Ranking via robust binary classification. In *Advances in Neural Information Processing Systems*, pages 2582–2590, 2014.
- Kai Zhong, Prateek Jain, and Inderjit S. Dhillon. Efficient matrix sensing using rank-1 gaussian measurements. In *International Conference on Algorithmic Learning Theory (ALT)*, 2015.