
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Schultheis, Erik; Mohammadnia Qaraei, Mohammadreza; Gupta, Priyanshu; Babbar, Rohit
Unbiased Loss Functions for Extreme Classification With Missing Labels

Published: 01/07/2020

Please cite the original version:

Schultheis, E., Mohammadnia Qaraei, M., Gupta, P., & Babbar, R. (2020). *Unbiased Loss Functions for Extreme Classification With Missing Labels*. arXiv.org.

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Unbiased Loss Functions for Extreme Classification With Missing Labels

Erik Schultheis¹ Mohammadreza Qaraei¹ Priyanshu Gupta² Rohit Babbar¹

Abstract

The goal in extreme multi-label classification (XMC) is to tag an instance with a small subset of relevant labels from an extremely large set of possible labels. In addition to the computational burden arising from large number of training instances, features and labels, problems in XMC are faced with two statistical challenges, (i) large number of ‘tail-labels’ – those which occur very infrequently, and (ii) missing labels as it is virtually impossible to manually assign *every* relevant label to an instance. In this work, we derive an unbiased estimator for general formulation of loss functions which decompose over labels, and then infer the forms for commonly used loss functions such as hinge- and squared-hinge-loss and binary cross-entropy loss. We show that the derived unbiased estimators, in the form of appropriate weighting factors, can be easily incorporated in state-of-the-art algorithms for extreme classification, thereby scaling to datasets with hundreds of thousand labels. However, empirically, we find a slightly altered version that gives more relative weight to tail labels to perform even better. We suspect is due to the label imbalance in the dataset, which is not explicitly addressed by our theoretically derived estimator. Minimizing the proposed loss functions leads to significant improvement over existing methods (up to 20% in some cases) on benchmark datasets in XMC.

1. Introduction

Extreme Multi-label Classification (XMC) refers to supervised learning where each training/test instance is labeled with small subset of relevant labels that are chosen from a large set of possible target labels. Problems consisting of extremely large number of labels are common in various domains such as annotating large encyclopedia (Dekel & Shamir, 2010; Partalas et al., 2015), image-classification

(Deng et al., 2010) and next word prediction (Mikolov et al., 2013). It has also been noticed that the framework of XMC can be effectively leveraged to address learning problems arising in recommendation systems and web-advertising (Agrawal et al., 2013; Jain et al., 2019; Prabhu & Varma, 2014). For the case of recommendation systems, by learning from other similar users’ buying patterns, this framework can be used to recommend a small subset of relevant items from a large collection of all possible items.

With applications in a diverse range, designing effective machine learning algorithms to solve XMC has become a key research challenge. From the computational aspect of the learning problem, building effective extreme classifiers is faced with a scaling challenge arising due to large number of (i) output labels, (ii) input training instances, and (iii) input features. It is not unlikely to have the dimensionality of the above dataset statistics to be $\mathcal{O}(10^6)$. Two properties of datasets in XMC which pose further statistical challenges, (i) Long-tail distribution of instances among labels, and (ii) Missing labels, are discussed next.

1.1. Tail Labels

An important statistical feature of the datasets in XMC is that a large fraction of labels are tail labels, i.e., those which have very few training instances (also referred to as a fat-tailed distribution and Zipf’s law). Typically the label frequency distribution follows a power law, an example of which is shown in Figure 1 for a publicly available benchmark WikiLSHTC-325K dataset, consisting of approximately 325,000.

Concretely, let $n_{(r)}$ denote the number of occurrences of the r -th ranked label, when ranked in decreasing order of number of training instances that belong to that label, then

$$n_{(r)} = n_{(1)}r^{-\beta},$$

where $\beta > 0$ denotes the exponent of the power law.

Tail labels exhibit diversity of the label space, and contain informative content not captured by the head or torso labels. Indeed, by predicting well the head labels, yet omitting most of the tail labels, an algorithm can achieve high accuracy (Wei & Li, 2019). Such behavior is not typically desirable in real world applications (Babbar & Schölkopf, 2019). In movie recommendation systems, for instance,

¹Aalto University, Helsinki, Finland ²IIT, Kanpur, India. Correspondence to: Rohit Babbar <firstname.lastname@aalto.fi>.




Figure 1. Label frequency in WikiLSHTC-325K. X-axis shows the label IDs sorted by the frequency of positive instances and Y-axis gives the actual frequency (on log-scale). More than half of the labels have fewer than 10 training instances.

the head labels correspond to popular blockbusters—most likely, the user has already watched these. However, the tail of the distribution corresponds to less popular yet equally favored films, like independent movies. These are the movies that the recommendation engine should ideally focus on. A similar argument applies to search engine development (Radlinski et al., 2009) and hash-tag recommendation in social media (Denton et al., 2015). However, effectively predicting tail-labels can be an enormous challenge due to the extreme data imbalance problem, where a given tail label appears in only a couple of (positive) instances and does not appear in millions of other (negative) training instances.

1.2. Missing Labels

In addition to having unfavourable statistics, when learning to classify tail labels it has been shown that one also needs to account for missing labels in the training data. In a dataset where the labels for each example are chosen from a label space with thousands of elements, it is impossible to explicitly check for the presence or absence of each label, some examples will have missing labels. Even worse, the chance for a label to be missing is higher for tail labels than for head labels. In the movie example, this means that there are more people who would have liked an independent movie, but did not because never seeing it, than there are people who would have liked a blockbuster but never saw it. However, we can typically assume that most people who claim to like a movie actually do so, i.e. that we do not have significant amounts of spurious positive labels in the training set. This leads to the propensity model introduced in (Jain et al., 2016), formally presented in section 2.

In (Jain et al., 2016), it was shown that it is possible to calculate an unbiased estimate of certain loss functions and evaluation metrics, even if the available data has missing labels. They further showed that using such an estimate for training results in better performance on tail labels. However, the analysis left out some important loss functions, such as hinge- and squared-hinge-loss as well as binary cross-entropy. In this paper we derive a way of calculating unbiased estimates for all loss functions that separate over binary labels. However, it turns out that such estimates may lose important properties of the original loss function, e.g. an unbiased estimate of the (convex) hinge loss need no longer be convex. Therefore, we provide an alternative derivation, in which we consider the hinge loss as a piecewise-linear convex upper-bound on the 0-1 loss, and recommend using a piecewise-linear convex upper-bound on the unbiased estimate of the 0-1 loss.

We show that the derived unbiased estimators, in the form of appropriate weighting factors for loss functions, can be readily incorporated in state-of-the algorithms for extreme classification, and hence easily scale to datasets with hundreds of thousand labels. Empirically, the efficacy of the proposed loss functions is demonstrated by exhibiting superior performance to existing methods, with relative improvements of as much as 20% compared to some of the recently proposed tree-based methods for extreme classification (Prabhu et al., 2018).

2. Theory

In the extreme classification setting, it is not possible for a human annotator to consider every possible label when deciding which labels to assign to a given data point. Instead, they will look at an example and assign a set of fitting labels that comes to mind. It is reasonable to assume that any label assigned in such fashion will be correct, i.e. if the annotator where asked whether the label belonged to the example, they would confirm this. The converse is not necessarily true: If one were to ask the annotator for each label that was not chosen whether it was relevant for the example, it is likely that some would be considered relevant.

To capture this effect, the notion of propensity is introduced. The propensity of a label (for an example) is defined as the probability of the label being present, given that when explicitly asked the ground-truth annotator would confirm it to be present. An empirical model for estimating propensities from label frequencies is presented in (Jain et al., 2016).

For any given label i , we denote with $Y_i \in \{0, 1\}$ whether the label is present in the annotated dataset, and with $Y_i^* \in \{0, 1\}$ whether it should be present in the ground-truth. In the work of (Jain et al., 2016), the authors proposed to take into account the missing labels by replacing stochastic

estimates of the form $f^*(Y)$ by unbiased estimates g s.t. $\mathbb{E}[g(Y)] = \mathbb{E}[f^*(Y^*)]$. They derived expressions for cases in which $y = 0 \implies f^*(y) = 0$ (e.g. P@k), as well as for the hamming loss. In the following, we extend these results to all loss functions that decompose over labels.

2.1. Generic Results

We first derive an unbiased estimator for an arbitrary loss in the case of single-label binary classification [Theorem 1](#), and then extend this to all multi-label loss functions that separate over labels in [Corollary 1](#).

Theorem 1. *Let (Ω, \mathcal{F}, P) be a probability space with RVs $Y : \Omega \rightarrow \{0, 1\}$ and $Y^* : \Omega \rightarrow \{0, 1\}$. We assume a propensity model as above, i.e.*

$$\mathcal{P}\{Y = 1, Y^* = 0\} = 0. \quad (1)$$

$$\mathcal{P}\{Y = 1 | Y^* = 1\} =: p. \quad (2)$$

Let $l^* : \{0, 1\} \times \mathbb{R} \rightarrow \mathbb{R}$ be a function which can be decomposed into

$$l^*(y, \hat{y}) =: \begin{cases} l_+^*(\hat{y}) & y = 1 \\ l_-^*(\hat{y}) & y = 0 \end{cases}. \quad (3)$$

Then the function $l : \{0, 1\} \times \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$l_+(\hat{y}) := p^{-1} (l_+^*(\hat{y}) + (p-1)l_-^*(\hat{y})) \quad (4)$$

$$l(y, \hat{y}) = \begin{cases} l_+(\hat{y}) & y = 1 \\ l_-^*(\hat{y}) & y = 0 \end{cases} \quad (5)$$

allows to calculate an unbiased estimate of l^* :

$$\mathbb{E}[l^*(Y^*, \hat{y})] = \mathbb{E}[l(Y, \hat{y})]. \quad (6)$$

Proof. Combining the law of total probability with (1) and (2) gives

$$\begin{aligned} \mathcal{P}\{Y = 1\} &= \mathcal{P}\{Y = 1 | Y^* = 0\} \mathcal{P}\{Y^* = 0\} \\ &\quad + \mathcal{P}\{Y = 1 | Y^* = 1\} \mathcal{P}\{Y^* = 1\} \end{aligned} \quad (7)$$

$$= 0 \cdot \mathcal{P}\{Y^* = 0\} + p \mathcal{P}\{Y^* = 1\} \quad (8)$$

$$= p \cdot \mathcal{P}\{Y^* = 1\}. \quad (9)$$

The decomposition (3) of l^* can be rewritten as

$$l^*(y, \hat{y}) = \mathbb{1}\{y = 1\} (l_+^*(\hat{y}) - l_-^*(\hat{y})) + l_-^*(\hat{y}) \quad (10)$$

Let P and P^* denote the probability mass functions of Y and Y^* . Importance sampling ([Press et al., 2007](#)) lets us change the random variable by introducing the correspond-

ing importance factors

$$\mathbb{E}[f(Y^*)] = \sum_y [P^*(y) f(y)] \quad (11)$$

$$= \sum_y \left[P(y) \frac{P^*(y)}{P(y)} f(y) \right] \quad (12)$$

$$= \mathbb{E} \left[f(Y) \frac{P^*(Y)}{P(Y)} \right]. \quad (13)$$

Applied to l^* this results in

$$\begin{aligned} \mathbb{E}[l^*(Y^*, \hat{y})] &= \mathbb{E} [\mathbb{1}\{Y^* = 1\} (l_+^*(\hat{y}) - l_-^*(\hat{y})) + l_-^*(\hat{y})] \\ &= \mathbb{E} [\mathbb{1}\{Y^* = 1\} (l_+^*(\hat{y}) - l_-^*(\hat{y})) + l_-^*(\hat{y})] \\ &= \mathbb{E} \left[\mathbb{1}\{Y = 1\} \frac{P^*(Y)}{P(Y)} (l_+^*(\hat{y}) - l_-^*(\hat{y})) + l_-^*(\hat{y}) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{Y = 1\} \frac{P^*(1)}{P(1)} (l_+^*(\hat{y}) - l_-^*(\hat{y})) + l_-^*(\hat{y}) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{Y = 1\} \frac{l_+^*(\hat{y}) - l_-^*(\hat{y})}{p} + l_-^*(\hat{y}) \right], \end{aligned} \quad (14)$$

where we have used the indicator function to replace $P^*(Y)/P(Y) = P^*(1)/P(1) = p^{-1}$. Splitting

$$l_-^*(\hat{y}) = (\mathbb{1}\{Y = 1\} + \mathbb{1}\{Y = 0\}) l_-^*(\hat{y}), \quad (15)$$

we get the desired statement

$$\begin{aligned} \mathbb{E}[l^*(Y^*, \hat{y})] &= \mathbb{E} \left[\mathbb{1}\{Y = 1\} \frac{l_+^*(\hat{y}) - l_-^*(\hat{y})}{p} \right. \\ &\quad \left. + (\mathbb{1}\{Y = 1\} + \mathbb{1}\{Y = 0\}) l_-^*(\hat{y}) \right] \end{aligned} \quad (16)$$

$$\begin{aligned} &= \mathbb{E} \left[\mathbb{1}\{Y = 1\} \frac{l_+^*(\hat{y}) - l_-^*(\hat{y}) + p l_-^*(\hat{y})}{p} \right. \\ &\quad \left. + \mathbb{1}\{Y = 0\} l_-^*(\hat{y}) \right] \end{aligned} \quad (17)$$

$$= \mathbb{E} [\mathbb{1}\{Y = 1\} l_+(\hat{y}) + \mathbb{1}\{Y = 0\} l_-^*(\hat{y})]. \quad (18)$$

□

Corollary 1. *In an extreme classification setting with N labels, where for each label the conditions of [Theorem 1](#) are fulfilled, any loss function that decomposes over labels $\mathbf{Y} = \{Y_1, \dots, Y_N\}$ can be estimated without bias. Let $L : \{0, 1\}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$ such that*

$$L^*(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^N l_i^*(y_i, \hat{y}_i), \quad (19)$$

and let l_i be related to l_i^* as prescribed by [Theorem 1](#), equation (4). Defining

$$L(\mathbf{y}, \hat{\mathbf{y}}) := \sum_{i=1}^N l_i(y_i, \hat{y}_i), \quad (20)$$

then $L(\mathbf{Y}, \hat{\mathbf{y}})$ is an unbiased estimate of $L^*(\mathbf{Y}^*, \hat{\mathbf{y}})$.

Proof. The proof follows by linearity of expectation.

$$\mathbb{E}[L^*(\mathbf{Y}^*, \hat{\mathbf{y}})] = \mathbb{E}\left[\sum_{i=1}^N l_i^*(Y_i^*, \hat{y}_i)\right] \quad (21)$$

$$= \sum_{i=1}^N \mathbb{E}[l_i^*(Y_i^*, \hat{y}_i)] = \sum_{i=1}^N \mathbb{E}[l_i(Y_i, \hat{y}_i)] \quad (22)$$

$$= \mathbb{E}\left[\sum_{i=1}^N l_i(Y_i, \hat{y}_i)\right] = \mathbb{E}[L(\mathbf{Y}, \hat{\mathbf{y}})] \quad (23)$$

□

2.2. Examples of Loss Functions

We can now apply [Theorem 1](#) to the case of different loss functions. Due to the [Corollary](#) above, it suffices to show the following results for the case of a single binary label. Note that for binary labels with predictions in $\hat{y} \in \{0, 1\}$, squared error, 0-1-loss and Hamming loss are identical.

Squared Error The decomposition of

$$L^*(y, \hat{y}) = (y - \hat{y})^2 \quad (24)$$

is

$$l_+^*(\hat{y}) = 1 - 2\hat{y} + \hat{y}^2 \quad (25)$$

$$l_-^*(\hat{y}) = \hat{y}^2, \quad (26)$$

resulting in

$$l_+(\hat{y}) = \frac{1 - 2\hat{y} + \hat{y}^2 + (p-1)\hat{y}^2}{p} = \frac{1 - 2\hat{y}}{p} + \hat{y}^2 \quad (27)$$

Combining with l_-^* recovers the result of [\(Jain et al., 2016\)](#)

$$L(y, \hat{y}) = y \frac{1 - 2\hat{y}}{p} + y\hat{y}^2 + (1 - y)\hat{y}^2 = y \frac{1 - 2\hat{y}}{p} + \hat{y}^2. \quad (28)$$

Hinge Loss This loss is a convex upper-bound on the 0-1 loss, typically defined over $\{-1, +1\}$. Therefore, we introduce the notation $z = 2y - 1$, $\hat{z} = 2\hat{y} - 1$, such that the hinge loss is

$$L^*(z, \hat{z}) = \max(1 - z\hat{z}, 0). \quad (29)$$

This corresponds to

$$l_+^*(\hat{z}) = \max(1 - \hat{z}, 0), \quad (30)$$

$$l_-^*(\hat{z}) = \max(1 + \hat{z}, 0), \quad (31)$$

and thus

$$l_+(\hat{z}) = p^{-1} (\max(1 - \hat{z}, 0) + (p-1) \max(1 + \hat{z}, 0)). \quad (32)$$

Therefore, using

$$\mathbb{1}\{y = 1\} = (z + 1)/2, \quad (33)$$

$$\mathbb{1}\{y = 0\} = (1 - z)/2, \quad (34)$$

the re-weighted loss becomes

$$L(z, \hat{z}) = \frac{z + 1}{2} \frac{\max(1 - \hat{z}, 0) + (p-1) \max(1 + \hat{z}, 0)}{p} + \frac{1 - z}{2} \max(1 + \hat{z}, 0). \quad (35)$$

This is not a convex function (see [Figure 2](#)), and as such might be undesirable from an optimization perspective.

Thus we provide a different view, which instead of calculating an unbiased estimate of the convex upper-bound, calculates a convex upper-bound of an unbiased estimate of the 0-1 loss. For the 0-1 loss, we have

$$l_+^*(\hat{z}) = \begin{cases} 1 & \hat{z} < 0 \\ 0 & \hat{z} \geq 0 \end{cases} \quad l_-^*(\hat{z}) = \begin{cases} 0 & \hat{z} < 0 \\ 1 & \hat{z} \geq 0 \end{cases} \quad (36)$$

This results in the correction

$$l_+(\hat{z}) = \begin{cases} 1/p & \hat{z} < 0 \\ 1 - 1/p & \hat{z} \geq 0 \end{cases} \quad (37)$$

This will give an unbiased estimate of the 0-1-loss in the case of missing labels, but for interpretability it has the disadvantage that it can be negative. Therefore, we use the following equivalent (for optimization) loss function

$$\tilde{l}_+(\hat{z}) = \begin{cases} 2/p - 1 & \hat{z} < 0 \\ 0 & \hat{z} \geq 0 \end{cases}, \quad (38)$$

which is just shifted by a constant.

A convex upper-bound of this is given by

$$\tilde{l}_+^{\text{cv}}(\hat{z}) = \max((2/p - 1)(1 - \hat{z}), 0) \quad (39)$$

$$= (2/p - 1) \max(1 - \hat{z}, 0) \quad (40)$$

Combining with the unchanged negative part we get

$$L_{\text{cv}}(y, \hat{y}) = \left[\frac{z + 1}{2} \frac{2 - p}{p} + \frac{1 - z}{2} \right] \max(1 - z\hat{z}, 0) = \frac{z(1 - p) + 1}{p} \max(1 - z\hat{z}, 0) \quad (41)$$

Squared Hinge Loss To get the squared hinge loss, instead of bounding [\(38\)](#) piecewise-linearly, we derive a quadratic upper bound. The cusp of the parabola should be at $\hat{z} = 1$, thus we parameterize as $l_+(\hat{z}) =$




Figure 2. Visualization of 0-1-loss (gray), hinge loss (blue) and squared hinge loss (green) for $y = 1$. The dashed lines show the original loss functions without propensity re-weighting. The dotted gray line indicates the re-weighted 0-1-loss without the constant shift (37), the dotted blue line the unbiased, non-convex re-weighted hinge loss (32) (the red circle indicates the kink where the non-convexity appears) the dotted green line the square of the convex, re-weighted hinge loss. The solid lines indicate (38), (41) and (42) respectively.

$\alpha \max(\hat{z} - 1, 0)^2$, which leads to $\alpha = 2/p - 1$. Thus the loss function becomes

$$\begin{aligned} L_{cv}(y, \hat{y}) &= \left[\frac{z+1}{2} \frac{2-p}{p} + \frac{1-z}{2} \right] \max(1 - z\hat{z}, 0)^2 \\ &= \frac{z(1-p) + 1}{p} \max(1 - z\hat{z}, 0)^2 \end{aligned} \quad (42)$$

Note that this is a different result from squaring the reweighed hinge loss (41), which results in a quadratic dependency on the propensity.

Binary Cross-Entropy The BCE loss is given by

$$l_+^*(\hat{y}) = -\log \hat{y} \quad \text{and} \quad l_-^*(\hat{y}) = -\log(1 - \hat{y}). \quad (43)$$

Plugging into (4) gives

$$l_+(\hat{y}) = p^{-1} (-\log \hat{y} + (1-p) \log(1 - \hat{y})) \quad (44)$$

Combining gives

$$\begin{aligned} L(y, \hat{y}) &= -\frac{y}{p} \log \hat{y} + \frac{y(1-p) - p + py}{p} \log(1 - \hat{y}) \\ &= -\frac{y}{p} \log \hat{y} - \left(1 - \frac{y}{p}\right) \log(1 - \hat{y}) \end{aligned} \quad (45)$$

This result also follows directly from the fact that the BCE loss is linear in y .

$$\begin{aligned} \mathbb{E}[L^*(Y^*, \hat{y})] &= \mathbb{E}[-Y^* \log \hat{y} - (1 - Y^*) \log(1 - \hat{y})] \\ &= -\mathbb{E}[Y^*] \log \hat{y} + (1 - \mathbb{E}[Y^*]) \log(1 - \hat{y}) \end{aligned} \quad (46)$$

$$= -\mathbb{E}[Y/p] \log \hat{y} + (1 - \mathbb{E}[Y/p]) \log(1 - \hat{y}) \quad (47)$$

$$= -\mathbb{E}[Y/p] \log \hat{y} + (\mathbb{E}[Y/p] - 1) \log(1 - \hat{y}) \quad (48)$$

$$= \mathbb{E}[-Y/p \log \hat{y} - (1 - Y/p) \log(1 - \hat{y})] \quad (49)$$

To show that the re-weighted loss functions provide training benefits, we tested them in a controlled environment. We took the MNIST dataset, interpreted as a multi-label classification problem, and removed a percentage of the labels. Compared to the XMC settings presented in section 3, this has the advantage that one does not have to use an empirical model to estimate the propensities. Furthermore, data imbalance can be another motivator for re-weighting the loss function, and is a prominent feature of long-tailed extreme classification data sets. The results, in terms of loss function and classification accuracy, are presented in Figure 3.

3. Experimental Analysis

We evaluate the propensity weighting developed in the previous section on publicly available datasets from the extreme classification repository (Bhatia et al., 2016). We use the empirical model of (Jain et al., 2016, Section 5) for evaluating the label propensities

$$p_\ell = (1 + C \exp(-A \log(N_\ell + B)))^{-1}, \quad (50)$$

which is standard in the community. Here A , B and $C = (\log N - 1)(B + 1)^A$ are dataset dependent parameters, and N_ℓ denotes the number of positives for label ℓ . The data is represented as (sparse) bag-of-words features which is suitable for scaling with large-scale linear classification algorithms. Also, these datasets exhibit a long-tailed label distribution (Figure 1 for WikiLSHTC-325K dataset). The statistics of the datasets as well as A and B for each dataset are presented in Table 1

We apply the weighting scheme as part of a one-vs-rest implementation of linear SVM in multi-label setting, which is also referred to binary relevance in the literature (Luaces et al., 2012). The weight vector w_ℓ for each label ℓ is learnt by minimizing a combination of squared hinge loss and l_2 -regularization. Separating the squared hinge loss into the contributions from false negatives and false positives for




Figure 3. Comparison of loss functions for training MNIST with missing labels. In this case the MNIST data is interpreted as a multi-label learning problem, and true labels are dropped with a pre-selected propensity. A single-hidden-layer relu network is trained using different loss functions for 10 epochs. The resulting classifier is evaluated on MNIST test-data where all labels are present, based on test loss and binary classification accuracy. The plots to the left show different variations of binary cross-entropy with sigmoid activation, to the right the hinge loss.

label ℓ , this is given by the following optimization problem:

$$\begin{aligned} \min_{w_\ell} & \|w_\ell\|_2^2 + C_{p_\ell} \underbrace{\sum_{i \in \mathcal{L}_\ell^+} \max(0, 1 - (w_\ell^T x_i + b_\ell))^2}_{\text{FN}} \\ & + C_{n_\ell} \underbrace{\sum_{i \in \mathcal{L}_\ell^-} \max(0, 1 + (w_\ell^T x_i + b_\ell))^2}_{\text{FP}} \quad (51) \end{aligned}$$

where \mathcal{L}_ℓ^+ (\mathcal{L}_ℓ^-) denotes the set of positive (negative) training samples corresponding to the label ℓ .

The objective function for each label ℓ consists of two parts, FN (false negatives) and FP (false positives). The FN part penalizes samples that have not been assigned to a relevant label, whereas FP penalizes those that have been mistakenly assigned to an irrelevant label.

According to (42), we get a convex upper-bound on the 0-1 loss when $C_{p_\ell} = \left(\frac{2}{p} - 1\right)$ and $C_{n_\ell} = 1$. However, minimizing the 0-1 loss may not be the best goal in the XMC setting, due to the presence of tail labels. For those, most samples are negative, and thus most contribution to the loss will be from negatives, and recall might be bad for tail labels. To counteract this, we want to weight FN more strongly for tail labels, compared to head labels. Empirically,

we found a weighting of $C_{p_\ell} = \left(\frac{1}{p} - 1\right)$ to work well.

With the above weighting factors, the optimization in equation (3) generalizes the objective function optimized by one-vs-rest based extreme classification algorithm DiSMEC (Babbar & Schölkopf, 2017), where both the weights are set to 1. We thus call our approach Propensity Weighted DiSMEC (PW-DiSMEC).

3.1. Evaluation metrics

For evaluation of the methods, we use precision at k ($P@k$) and normalized discounted cumulative gain at k ($nDCG@k$), which are standard and commonly used metric for evaluating models in XMC. For each test sample with observed ground truth label vector $\mathbf{y} \in \{0, 1\}^N$ and predicted vector $\hat{\mathbf{y}} \in \mathbb{R}^N$, $P@k$ and $nDCG@k$ are computed as

$$P@k(\mathbf{y}, \hat{\mathbf{y}}) := \frac{1}{k} \sum_{\ell \in \text{rank}_k(\hat{\mathbf{y}})} \mathbf{y}_\ell, \quad (52)$$

$$nDCG@k(\mathbf{y}, \hat{\mathbf{y}}) := \frac{DCG@k}{\sum_{l=1}^{\min(k, \|\mathbf{y}\|_0)} \frac{1}{\log(l+1)}}, \quad (53)$$

$$DCG@k(\mathbf{y}, \hat{\mathbf{y}}) := \sum_{\ell \in \text{rank}_k(\hat{\mathbf{y}})} (\log(l+1))^{-1} \mathbf{y}_\ell, \quad (54)$$

| Dataset | # Training | # Test | # Labels | # Features | APpL | ALpP | A | B |
|-----------------------|------------|---------|----------------|------------|-------|------|------|-----|
| EURLex-4K | 15,539 | 3,809 | 3993 | 5,000 | 25.7 | 5.3 | 0.55 | 1.5 |
| AmazonCat-13K | 1,186,239 | 306,782 | 13,330 | 203,882 | 448.5 | 5.04 | 0.55 | 1.5 |
| Wikipedia-31K | 14,146 | 6,616 | 30,938 | 101,938 | 8.5 | 18.6 | 0.55 | 1.5 |
| WikiLSHTC-325K | 1,778,351 | 587,084 | 325,056 | 1,617,899 | 17.4 | 3.2 | 0.5 | 0.4 |
| Wikipedia-500K | 1,813,391 | 783,743 | 501,070 | 2,381,304 | 24.7 | 4.7 | 0.5 | 0.4 |
| Amazon-670K | 490,499 | 153,025 | 670,091 | 135,909 | 3.9 | 5.4 | 0.6 | 2.6 |

Table 1. The statistics of the multi-label datasets used in our experiments. APpL denotes the average points per label and ALpP is the average labels per point respectively. A and B refer to the parameters of the propensity model.

where $\text{rank}_k(\hat{\mathbf{y}})$ returns the k largest indices of $\hat{\mathbf{y}}$.

With applications of XMC arising in recommendation systems and web-advertising, the objective of an algorithm in this domain is to correctly recommend/advertise among the top- k slots. Hence evaluation metrics such as $\text{precision}@k$ and its rank-sensitive version $\text{nDCG}@k$ are commonly used to benchmark extreme classification algorithms. A collection of results from recent papers on datasets in Table 1 for algorithms developed over the last few years is given on the extreme classification repository (Bhatia et al., 2016).

For the case of missing labels, and imbalanced classification arising from the presence of tail-labels, propensity scored variants of $\text{P}@k$ and $\text{nDCG}@k$ are of main interest to us in this work and are given by :

$$\text{PSP}@k(\mathbf{y}, \hat{\mathbf{y}}) := \frac{1}{k} \sum_{\ell \in \text{rank}_k(\hat{\mathbf{y}})} \frac{\mathbf{y}_\ell}{p_\ell} \quad (55)$$

$$\text{PSnDCG}@k(\mathbf{y}, \hat{\mathbf{y}}) := \frac{\text{PSDCG}@k}{\sum_{\ell=1}^{\min(k, \|\mathbf{y}\|_0)} \frac{1}{\log(\ell+1)}} \quad (56)$$

$$\text{PSDCG}@k(\mathbf{y}, \hat{\mathbf{y}}) := \sum_{\ell \in \text{rank}_k(\hat{\mathbf{y}})} \frac{\mathbf{y}_\ell}{p_\ell \log(\ell+1)}. \quad (57)$$

To match against the best possible performance attainable by any system, as suggested in (Jain et al., 2016), we define, for M test samples, $\mathbb{G}(\{\hat{\mathbf{y}}\}) = \frac{-1}{M} \sum_{i=1}^M \mathbb{L}(\hat{\mathbf{y}}_i, \mathbf{y}_i)$, where $\mathbb{L}(\cdot, \cdot)$ and $\mathbb{G}(\cdot)$ signify loss and gain respectively. We use $100 * \mathbb{G}(\{\hat{\mathbf{y}}\}) / \mathbb{G}(\{\mathbf{y}\})$ as the performance metric. The loss $\mathbb{L}(\cdot, \cdot)$ can take two forms, (i) $\mathbb{L}(\hat{\mathbf{y}}_i, \mathbf{y}_i) = -\text{PSnDCG}@k$, and (ii) $\mathbb{L}(\hat{\mathbf{y}}, \mathbf{y}) = -\text{PSP}@k$. This leads to the two metrics which are finally used in our comparison in Table 2 (denoted $\text{PSP}@k$ and $\text{PSnDCG}@k$), and evaluated for $k = 1, 3, 5$.

3.2. Baseline methods

We compare the proposed method against seven state-of-the-art algorithms from three classes of XMC methods:

1. Label-embedding methods:

- **LEML** (Yu et al., 2014) - Uses a regularized least squares objective for learning to embed the label vectors to a lower-dimensional linear subspace,

map the feature vectors to the lower dimensional label space, and decompress the lower dimensional labels by a linear transformation.

- **SLEEC** (Bhatia et al., 2015) - Embeds the label vectors to a lower-dimensional space in which nearest neighbors are preserved, learns a regressor to map the feature vectors to the embedded space, and performs kNN over the regressor for decompression.
- #### 2. Tree-based methods:
- **PFastXML** (Jain et al., 2016) - The instances are recursively partitioned into two nodes, with the objective of maximizing a propensity scored metric in each node.
 - **Parabel** (Prabhu et al., 2018) - Partitions the labels into two balanced groups using 2-means and learns a one-vs-rest classifier at each node.
- #### 3. One-vs-all methods:
- **(P)PD-Sparse** (Yen et al., 2016; 2017) - These class of extreme classification algorithms exploit the sparsity in the primal and dual problem combined with elastic net regularization. PD-Sparse uses multi-class hinge loss while PPD-Sparse uses hinge loss for binary classification. Being amenable to distributed training, PPD-Sparse can scale to bigger datasets.
 - **DiSMEC** (Babbar & Schölkopf, 2017) - Optimizes Hamming loss with ℓ_2 regularization in a distributed environment which achieves state-of-the-art on vanilla $\text{P}@k$. DiSMEC prunes weights for model size reduction.
 - **ProXML** (Babbar & Schölkopf, 2019) - Improves tail-label detection by posing the learning problem as an instance of robustness optimization. It proposes to guard against small perturbations in the feature composition of the instances of the same class, leading to ℓ_1 regularization. Its key difference compared to our method is that it does not weigh the hinge loss but treats the regularization part.

Unbiased Loss Functions for Extreme Classification With Missing Labels

| Dataset | Ours | Embedding based | | Tree based | | Linear one-vs-rest | | |
|-----------------------|-------------|-----------------|------|-------------|------|--------------------|--------|-------------|
| | PW-DiSMEC | SLEEC | LEML | PfXML | Prbl | (P)PD | DiSMEC | ProXML |
| EURLex-4K | | | | | | | | |
| <i>PSP@1</i> | 44.9 | 35.4 | 24.1 | 39.9 | 37.7 | 38.2 | 41.2 | 45.2 |
| <i>PSP@3</i> | 49.2 | 39.8 | 27.2 | 43.0 | 44.7 | 42.7 | 45.4 | 48.2 |
| <i>PSP@5</i> | 52.2 | 42.7 | 29.1 | 44.5 | 48.8 | 44.8 | 49.3 | 51.0 |
| <i>PSnDCG@3</i> | 48.0 | 38.8 | 26.4 | 42.2 | 43.4 | 40.9 | 44.3 | 47.5 |
| <i>PSnDCG@5</i> | 50.1 | 40.3 | 27.7 | 43.2 | 46.1 | 42.8 | 46.9 | 49.1 |
| AmazonCat-13K | | | | | | | | |
| <i>PSP@1</i> | 67.1 | 46.8 | - | 68.0 | 50.9 | 49.6 | 51.4 | 52.3 |
| <i>PSP@3</i> | 71.3 | 58.5 | - | 72.3 | 64.0 | 61.6 | 61.0 | 62.0 |
| <i>PSP@5</i> | 73.2 | 65.0 | - | 74.6 | 72.1 | 68.2 | 65.9 | 66.2 |
| <i>PSnDCG@3</i> | 70.4 | 55.1 | - | 71.1 | 35.2 | 58.3 | 65.2 | 66.1 |
| <i>PSnDCG@5</i> | 72.4 | 60.1 | - | 72.6 | 38.1 | 62.7 | 68.8 | 69.5 |
| Wikipedia-31K | | | | | | | | |
| <i>P@1</i> | 15.6 | 11.1 | 9.4 | 9.4 | 11.7 | - | 13.4 | 13.6 |
| <i>P@3</i> | 16.2 | 11.9 | 10.1 | 9.9 | 12.7 | - | 12.9 | 13.0 |
| <i>P@5</i> | 17.3 | 12.4 | 10.6 | 10.4 | 13.7 | - | 13.6 | 13.2 |
| <i>PSnDCG@3</i> | 16.0 | 11.7 | 9.9 | 9.7 | 12.5 | - | 13.2 | 13.4 |
| <i>PSnDCG@5</i> | 16.8 | 12.1 | 10.2 | 10.1 | 13.1 | - | 13.6 | 13.7 |
| WikiLSHTC-325K | | | | | | | | |
| <i>PSP@1</i> | 35.3 | 20.5 | 3.4 | 25.4 | 28.7 | 28.3 | 29.1 | 34.8 |
| <i>PSP@3</i> | 38.4 | 23.3 | 3.7 | 26.8 | 35.0 | 33.5 | 35.6 | 37.7 |
| <i>PSP@5</i> | 41.7 | 25.2 | 4.2 | 28.3 | 38.6 | 36.6 | 39.4 | 41.0 |
| <i>PSnDCG@3</i> | 39.8 | 22.4 | 3.6 | 26.4 | 35.2 | 31.9 | 35.9 | 38.7 |
| <i>PSnDCG@5</i> | 42.6 | 23.5 | 3.9 | 27.2 | 38.1 | 33.6 | 39.4 | 41.5 |
| Wikipedia-500K | | | | | | | | |
| <i>PSP@1</i> | 33.9 | 21.1 | 3.2 | 22.2 | 28.8 | - | 31.2 | 33.1 |
| <i>PSP@3</i> | 36.7 | 21.0 | 3.4 | 21.3 | 31.9 | - | 33.4 | 35.0 |
| <i>PSP@5</i> | 39.9 | 20.8 | 3.5 | 21.6 | 34.6 | - | 37.0 | 39.4 |
| <i>PSnDCG@3</i> | 39.7 | 20.9 | 3.1 | 21.6 | 31.2 | - | 33.7 | 35.2 |
| <i>PSnDCG@5</i> | 42.7 | 23.1 | 3.3 | 21.8 | 35.5 | - | 37.1 | 39.0 |
| Amazon-670K | | | | | | | | |
| <i>PSP@1</i> | 30.9 | 20.6 | 2.0 | 27.1 | 27.6 | 26.6 | 27.8 | 30.8 |
| <i>PSP@3</i> | 33.1 | 23.3 | 2.2 | 28.2 | 31.0 | 30.7 | 30.6 | 32.8 |
| <i>PSP@5</i> | 35.2 | 26.0 | 2.4 | 29.3 | 34.1 | 34.7 | 34.2 | 35.1 |
| <i>PSnDCG@3</i> | 31.9 | 22.6 | 2.2 | 27.9 | 28.4 | - | 28.8 | 31.7 |
| <i>PSnDCG@5</i> | 32.9 | 24.4 | 2.3 | 28.6 | 29.9 | - | 30.7 | 32.6 |

Table 2. Comparison of PSP@k and PSnDCG@k on benchmark datasets for $k = 1, 3$ and 5. The columns PfXML, Prbl, (P)PD show the results for PfastXML, Parabel, and (P)PD-Sparse respectively. For each row, score of the best performing algorithm is highlighted in bold. Entries marked “-” imply the corresponding method could not scale to the particular dataset (LEML for AmazonCat-13K and PD-Sparse for Wikipedia-500K) or the scores are unavailable (PDSparse for Wikipedia-31K). The comparison of the models in terms of vanilla precision and nDCG can be found in Figure 4.

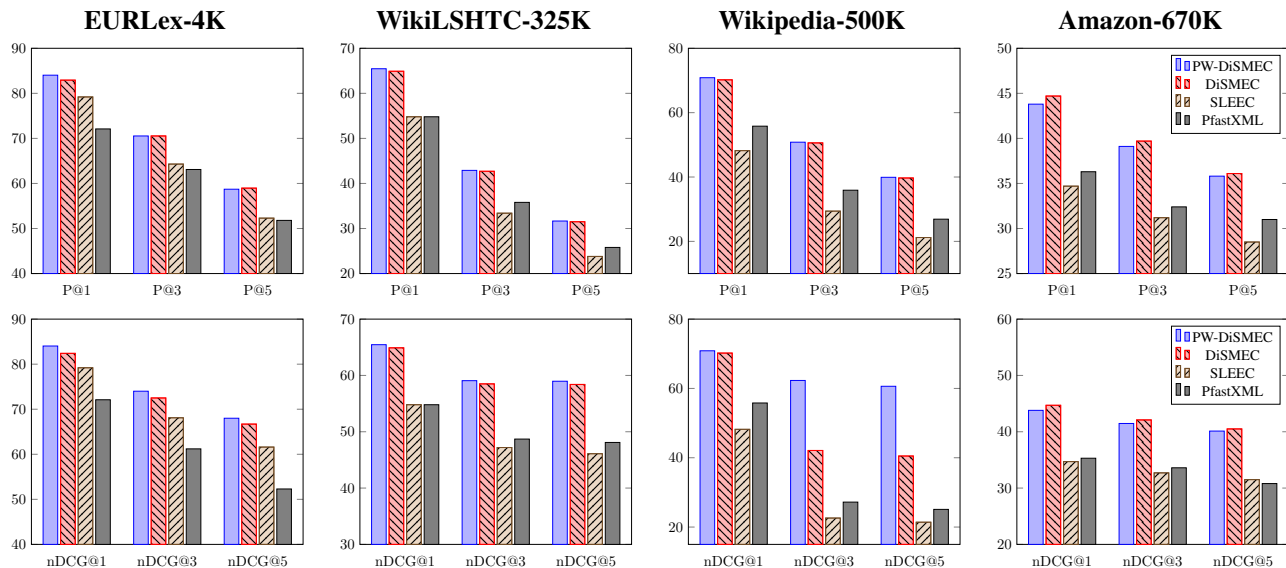


Figure 4. Comparison of $P@k$ (top row) and $nDCG@k$ (bottom row) on benchmark datasets. The comparison of the models in terms of the propensity scored metrics can be found in Table 2.

3.3. Experimental Results

Performance on propensity-scored metrics A comparison of the proposed method with the baselines in terms of the propensity scored precision is presented in Table 2. The proposed method outperforms the embedding-based methods LEML and SLEEC on all benchmarks.

In comparison to tree-based methods, our method achieves better results on five out of six datasets. For WikiLSHTC-325K and Wikipedia-500K, the improvements in $PSP@k$ over PfastXML, in which the objective is to maximize a propensity scored metric, are about 10% in absolute terms. Over Parabel, the relative improvement is almost 20-30% on most datasets.

It is clear that the proposed method outperforms PD-Sparse and vanilla DISMEC on all the datasets. The proposed method which is based on weighting the squared hinge loss function and keeping the regularization as l_2 also outperforms PROXML, which is specifically designed for better tail-label detection. Furthermore, being amenable to second order optimization scheme, it is almost two orders of magnitude faster to train than PROXML.

Performance on vanilla metrics While the proposed method aims at mitigating the impact of missing label and data-imbalance on the performance on tail labels, it is desirable that this should not come at an expense of accuracy in vanilla metrics, in equations (52) & (53). Figure 4 illustrates a comparison of the proposed method and the baselines in terms of precision and nDCG. In the interest of space, the results are compared to three other baseline methods. As

can be seen, the proposed method works at par or even in better in most cases on these metrics as well.

4. Other Related work

Various works in XMC can be broadly divided into five main categories as follows :

1. *Tree-based* : Tree-based methods implement a divide-and-conquer paradigm and scale to large label sets in XMC by partitioning the labels space. As a result, these scheme of methods have the computational advantage of enabling faster training and prediction. Apart from those discussed in previous sections, some of the other methods include recent works (Jasinska et al., 2016; Majzoubi & Choromanska, 2019; Wydmuch et al., 2018). However, tree-based methods suffer from error propagation in the tree cascade, and perform particularly worse on metrics which are sensitive for tail-labels. An approach based on shallower trees which mitigates the effect of error cascading has been recently demonstrated in (Khandagale et al., 2019). Approaches based on decision trees have also been proposed for multi-label classification and those tailored to XMC regime (Si et al., 2017).
2. *Label embedding* : Label-embedding approaches assume that, despite large number of labels, the label matrix is effectively low rank and therefore project it to a low-dimensional sub-space. These approaches have been at the fore-front in multi-label classification for small scale problems with few tens or hundred labels (Hsu et al., 2009; Lin et al., 2014; Weston et al., 2011).

In some of the works, it was argued that the low rank embedding may be insufficient for capturing the label diversity in XMC settings ((Bhatia et al., 2015; Tagami, 2017; Xu et al., 2016)), which has been questioned in the recent work (Guo et al., 2019).

3. *One-vs-rest* : As the name suggests, these methods learn a classifier per label which distinguishes it from rest of the labels. In terms of prediction accuracy, these methods have been shown to be among the best performing ones for XMC (Babbar & Schölkopf, 2017; Babbar & Schölkopf, 2019; Yen et al., 2017). However, due to their reliance on a distributed training framework, it remains challenging to employ them in resource constrained environments.
4. *Deep learning* : Deeper architectures on top of word-embeddings have also been explored in recent works. A convolutional network based approach, XML-CNN, for deep extreme multi-label classification was proposed in (Liu et al., 2017). Recently, successful application of attention mechanism, originally motivated from a machine translation perspective, has been demonstrated in the recent work AttentionXML (You et al., 2019). For dense feature vectors, AttentionXML performs the best among deep learning approaches significantly improving over XML-CNN. X-Bert, an approach based on pre-trained Bert language model ((Devlin et al., 2018)) has been presented in the work (Chang et al., 2019).
5. *Negative Sampling based methods* : Recently, there has also been a surge of methods in designing efficient negative sampling based methods for extreme classification (Bamler & Mandt, 2020; Jain et al., 2019; Reddi et al., 2019). The primary goal of these algorithms is to avoid computing the loss over all the samples which do not belong to a given label, and hence speed up training without any significant loss in prediction accuracy.

5. Conclusion

Towards modeling the missing labels and extreme data imbalance when learning with number of labels, we presented an derived and analyzed unbiased loss functions which decompose over the individual labels, which includes the commonly used variants such as hinge- and squared-hinge-loss and binary cross-entropy loss. For the setting of XMC, the suitability of the estimator for squared-hinge-loss function was demonstrated by incorporating in the one-vs-rest implementation of large-scale linear SVM. Empirically, it was observed that the resulting method outperforms all existing baselines on benchmark datasets.

References

- Agrawal, R., Gupta, A., Prabhu, Y., and Varma, M. Multi-label learning with millions of labels: Recommending advertiser bid phrases for web pages. In *WWW*, May 2013.
- Babbar, R. and Schölkopf, B. Dismec: Distributed sparse machines for extreme multi-label classification. In *WSDM*, pp. 721–729, 2017.
- Babbar, R. and Schölkopf, B. Data scarcity, robustness and extreme multi-label classification. *Machine Learning*, 2019.
- Bamler, R. and Mandt, S. Extreme classification via adversarial softmax approximation. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rJxe3xSYDS>.
- Bhatia, K., Jain, H., Kar, P., Varma, M., and Jain, P. Sparse local embeddings for extreme multi-label classification. In *NIPS*, 2015.
- Bhatia, K., Dahiya, K., Jain, H., Mittal, A., Prabhu, Y., and Varma, M. The extreme classification repository: Multi-label datasets and code, 2016. URL <http://manikvarma.org/downloads/XC/XMLRepository.html>.
- Chang, W.-C., Yu, H.-F., Zhong, K., Yang, Y., and Dhillon, I. A modular deep learning approach for extreme multi-label text classification. *arXiv preprint arXiv:1905.02331*, 2019.
- Dekel, O. and Shamir, O. Multiclass-multilabel classification with more classes than examples. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 137–144, 2010.
- Deng, J., Berg, A. C., Li, K., and Fei-Fei, L. What does classifying more than 10,000 image categories tell us? In *ECCV*, 2010.
- Denton, E., Weston, J., Paluri, M., Bourdev, L., and Fergus, R. User conditional hashtag prediction for images. In *KDD*, 2015.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Guo, C., Mousavi, A., Wu, X., Holtmann-Rice, D. N., Kale, S., Reddi, S., and Kumar, S. Breaking the glass ceiling for embedding-based classifiers for large output spaces. In *Advances in Neural Information Processing Systems*, pp. 4944–4954, 2019.

- Hsu, D., Kakade, S., Langford, J., and Zhang, T. Multi-label prediction via compressed sensing. In *Advances in neural information processing systems*, 2009.
- Jain, H., Prabhu, Y., and Varma, M. Extreme multi-label loss functions for recommendation, tagging, ranking and other missing label applications. In *KDD*, August 2016.
- Jain, H., Balasubramanian, V., Chunduri, B., and Varma, M. Slice: Scalable linear extreme classifiers trained on 100 million labels for related searches. In *WSDM*, pp. 528–536, 2019.
- Jasinska, K., Dembczynski, K., Busa-Fekete, R., Pfanschmidt, K., Klerx, T., and Hüllermeier, E. Extreme f-measure maximization using sparse probability estimates. In *ICML*, 2016.
- Khandagale, S., Xiao, H., and Babbar, R. Bonsai-diverse and shallow trees for extreme multi-label classification. *arXiv preprint arXiv:1904.08249*, 2019.
- Lin, Z., Ding, G., Hu, M., and Wang, J. Multi-label classification via feature-aware implicit label space encoding. In *International Conference on Machine Learning*, pp. 325–333, 2014.
- Liu, J., Chang, W.-C., Wu, Y., and Yang, Y. Deep learning for extreme multi-label text classification. In *SIGIR*, pp. 115–124. ACM, 2017.
- Luaces, O., Díez, J., Barranquero, J., del Coz, J. J., and Bahamonde, A. Binary relevance efficacy for multilabel classification. *Progress in Artificial Intelligence*, 1(4): 303–313, 2012.
- Majzoubi, M. and Choromanska, A. Ldsm: Logarithm-depth streaming multi-label decision trees. *arXiv preprint arXiv:1905.10428*, 2019.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- Partalas, I., Kosmopoulos, A., Baskiotis, N., Artieres, T., Paliouras, G., Gaussier, E., Androutsopoulos, I., Amini, M.-R., and Galinari, P. Lshtc: A benchmark for large-scale text classification. *arXiv preprint arXiv:1503.08581*, 2015.
- Prabhu, Y. and Varma, M. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *KDD*, pp. 263–272. ACM, 2014.
- Prabhu, Y., Kag, A., Harsola, S., Agrawal, R., and Varma, M. Parabel: Partitioned label trees for extreme classification with application to dynamic search advertising. In *WWW*, 2018.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- Radlinski, F., Bennett, P. N., Carterette, B., and Joachims, T. Redundancy, diversity and interdependent document relevance. In *ACM SIGIR Forum*, 2009.
- Reddi, S. J., Kale, S., Yu, F., Holtmann-Rice, D., Chen, J., and Kumar, S. Stochastic negative mining for learning with large output spaces. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1940–1949, 2019.
- Si, S., Zhang, H., Keerthi, S. S., Mahajan, D., Dhillon, I. S., and Hsieh, C.-J. Gradient boosted decision trees for high dimensional sparse output. In *ICML*, 2017.
- Tagami, Y. Annexml: Approximate nearest neighbor search for extreme multi-label classification. In *KDD*. ACM, 2017.
- Wei, T. and Li, Y.-F. Does tail label help for large-scale multi-label learning? *IEEE Transactions on Neural Networks and Learning Systems*, 2019.
- Weston, J., Bengio, S., and Usunier, N. Wsabie: Scaling up to large vocabulary image annotation. In *IJCAI*, 2011.
- Wydmuch, M., Jasinska, K., Kuznetsov, M., Busa-Fekete, R., and Dembczynski, K. A no-regret generalization of hierarchical softmax to extreme multi-label classification. In *Advances in Neural Information Processing Systems*, pp. 6355–6366, 2018.
- Xu, C., Tao, D., and Xu, C. Robust extreme multi-label learning. In *KDD*, 2016.
- Yen, I. E., Huang, X., Ravikumar, P., Zhong, K., and Dhillon, I. S. Pd-sparse: A primal and dual sparse approach to extreme multiclass and multilabel classification. In *ICML*, 2016.
- Yen, I. E., Huang, X., Dai, W., Ravikumar, P., Dhillon, I., and Xing, E. Ppdspare: A parallel primal-dual sparse method for extreme classification. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 545–553. ACM, 2017.
- You, R., Zhang, Z., Wang, Z., Dai, S., Mamitsuka, H., and Zhu, S. Attentionxml: Label tree-based attention-aware deep model for high-performance extreme multi-label text classification. In *Advances in Neural Information Processing Systems*, pp. 5812–5822, 2019.
- Yu, H.-F., Jain, P., Kar, P., and Dhillon, I. Large-scale multi-label learning with missing labels. In *ICML*, pp. 593–601, 2014.