
Generalized test utilities for long-tail performance in extreme multi-label classification

Erik Schultheis
Aalto University
Helsinki, Finland
erik.schultheis@aalto.fi

Marek Wydmuch
Poznan University of Technology
Poznan, Poland
mwydmuch@cs.put.poznan.pl

Wojciech Kotłowski
Poznan University of Technology
Poznan, Poland
wkotlowski@cs.put.poznan.pl

Rohit Babbar
University of Bath / Aalto University
Bath, UK / Helsinki, Finland
rb2608@bath.ac.uk

Krzysztof Dembczyński
Yahoo Research / Poznan University of Technology
New York, USA / Poznan, Poland
krzysztof.dembczynski@yahooinc.com

Abstract

Extreme multi-label classification (XMLC) is the task of selecting a small subset of relevant labels from a very large set of possible labels. As such, it is characterized by long-tail labels, i.e., most labels have very few positive instances. With standard performance measures such as precision@k, a classifier can ignore tail labels and still report good performance. However, it is often argued that correct predictions in the tail are more “interesting” or “rewarding,” but the community has not yet settled on a metric capturing this intuitive concept. The existing propensity-scored metrics fall short on this goal by confounding the problems of long-tail and missing labels. In this paper, we analyze generalized metrics budgeted “at k” as an alternative solution. To tackle the challenging problem of optimizing these metrics, we formulate it in the *expected test utility* (ETU) framework, which aims to optimize the expected performance on a fixed test set. We derive optimal prediction rules and construct computationally efficient approximations with provable regret guarantees and robustness against model misspecification. Our algorithm, based on block coordinate ascent, scales effortlessly to XMLC problems and obtains promising results in terms of long-tail performance.

1 Introduction

Extreme multi-label classification (XMLC) is a challenging task with a wide spectrum of real-life applications, such as tagging of text documents [10], content annotation for multimedia search [14], or different type of recommendation [5, 1, 33, 44, 53, 28, 7]. Because of the nature of its applications, the typical approach in XMLC is to predict exactly k labels (e.g., corresponding to k slots in the user interface) which optimize a standard performance metric such as precision or (normalized) discounted cumulative gain. Given the enormous number of labels in XMLC tasks, which can reach millions or more, it is not surprising that many of them are very sparse, and hence make the label distribution strongly long-tailed [2]. It has been noticed that algorithms can achieve high performance on the standard metrics, but never predict any tail labels [42]. Therefore, there is a need to develop

Table 1: Performance measures (%) on AmazonCat-13k of a classifier trained on the full set of labels and a classifier trained with only 1k head labels.

Metric	full labels			head labels		
	@1	@3	@5	@1 (diff.)	@3 (diff.)	@5 (diff.)
Precision	93.03	78.51	63.74	93.08 (+0.05%)	76.42 (-2.66%)	58.21 (-8.67%)
nDCG	93.03	87.25	85.35	93.08 (+0.05%)	85.75 (-1.71%)	80.91 (-5.19%)
PS-Precision	49.76	62.63	70.35	49.07 (-1.39%)	57.71 (-7.84%)	57.41 (-18.40%)
Macro-Precision	13.28	32.65	44.16	4.31 (-67.54%)	5.28 (-83.82%)	4.32 (-90.21%)
Macro-Recall	1.38	11.06	30.57	0.47 (-65.61%)	2.69 (-75.71%)	4.10 (-86.59%)
Macro-F1	2.26	14.67	32.84	0.74 (-67.37%)	3.10 (-78.88%)	3.77 (-88.51%)
Coverage	15.19	40.53	60.88	5.11 (-66.32%)	7.37 (-81.82%)	7.52 (-87.65%)

a metric that prefers “rewarding” [48], “diverse” [3], and “rare and informative” [34] labels over frequently-occurring head labels. Currently, the XMLC community attempts to capture this need using *propensity-scored* performance metrics [15]. These metrics give increased weight to tail labels, but have been derived from the perspective of missing labels, and as such they are not really solving the problem of tail labels [41].

In Table 1 we compare different metrics for budgeted at k predictions. We train a PLT model [17] on the full AMAZONCAT-13K dataset [27] and a reduced version with the 1000 most popular labels only. The test is performed for both models on the full set of labels. The standard metrics are only slightly perturbed by reducing the label space to the head labels. This holds even for propensity-scored precision, which decreases by just 1%-20% despite discarding over 90% of the label space. In contrast, macro measures and coverage decrease between 60% and 90% if tail labels are ignored. These results show that budgeted-at- k macro measures might be very attractive in the context of long tails. Macro-averaging treats all the labels equally important, preventing the labels with a small number of positive examples to be ignored. Furthermore, the budget of k labels “requires” the presence of long-tail labels in a compact set of predicted labels.

While we can easily use these measures to evaluate and compare different methods, we also would like to make predictions that directly optimize these metrics. The existing approaches to macro-averaged metrics consider the unconstrained case, in which label-wise optimization is possible [47, 11, 23, 16, 22, 26]. Each binary problem can be then solved under one of two frameworks for optimizing complex performance measures, namely *population utility* (PU) or *expected test utility* (ETU) [49, 12]. The former aims at optimizing the performance on the population level. The latter optimizes the performance directly on a given test set. Interestingly, in both frameworks the optimal solution is based on thresholding conditional label probabilities [12], but the resulting thresholds are different with the discrepancy diminishing with the size of the test set. The threshold tuning for PU is usually performed on a validation set [47, 26], while the exact optimization for ETU is performed on a test set. It requires cubic time in a general case and quadratic time in some special cases [49, 32]. Approximate solutions can be obtained in linear time [25, 12].

These approaches cannot be directly applied if prediction of exactly k labels for each instance is required. In such case, the optimization problems for different labels are tightly coupled through this constraint, making the final problem much more difficult. Despite the fact that optimization of complex performance metrics is a well-established problem, considered not only in binary and multi-label classification as discussed above, but also in multi-class classification [30, 31], the results presented in this paper go beyond the state-of-the-art as budgeted-at- k predictions have not yet been analyzed in this context. Let us underline that the requirement of k predictions is natural for recommendation systems, in which exactly k slots are available in the user interface to display recommendations. Even in situations where this does not apply, requiring the prediction to be “at k ” can be advantageous, as it prevents trivial solutions such as predicting nothing (for precision) or everything (for recall).

In this paper, we investigate optimal solutions for the class of utility functions that can be linearly decomposed over labels into binary utilities, which includes both instance-wise weighted measures and macro-averages. We solve the problem in the ETU framework which is well-suited, for example, to recommendation tasks in which recommendations for all users or items are rebuilt in regular intervals. In this case, we can first obtain probability estimates of individual labels for each instance in the test set, and then provide optimal predictions for a given metric based on these estimates. We

derive optimal prediction rules and construct computationally efficient approximations with provable guarantees, formally quantifying the influence of the estimation error of the label probabilities on the suboptimality of the resulting classifier. This result is expressed in the form of a regret bound [4, 30, 22, 13]. It turns out that for most metrics of interest, a small estimation error results in at most a small drop of the performance, which confirms our method is viable for applications. Our general algorithm, based on block coordinate ascent, scales effortlessly to XMLC problems and obtains promising empirical results.

2 Setup and notation

Let $\mathbf{x} \in \mathcal{X}$ denote an input instance, and $\mathbf{y} \in \{0, 1\}^m =: \mathcal{Y}$ the vector indicating the relevant labels, distributed according to $\mathbb{P}(\mathbf{y}|\mathbf{x})$. We consider the prediction problem in the *expected test utility* (ETU) framework, that is, we assume that we are given a known set of n instances $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathcal{X}^n$ with unknown labels, on which we have to make predictions.¹ Our goal is to assign each instance \mathbf{x}_i a set of exactly k (out of m) labels represented as a k -hot vector $\hat{\mathbf{y}}_i \in \mathcal{Y}_k := \{\mathbf{y} \in \mathcal{Y} : \|\mathbf{y}\|_1 = k\}$, and we let $\hat{\mathbf{Y}} = [\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n]^\top$ denote the entire $n \times m$ prediction matrix for a set of instances \mathbf{X} .

In the ETU framework, we treat \mathbf{X} as given and only make an assumption about the labeling process for the test sample: the labels $\mathbf{y}_i \in \mathcal{Y}$ corresponding to $\mathbf{x}_i \in \mathcal{X}$ do not depend on any other instances, that is $\mathbb{P}(\mathbf{Y}|\mathbf{X}) = \prod_{i=1}^n \mathbb{P}(\mathbf{y}_i|\mathbf{x}_i)$, where we use $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^\top \in \mathcal{Y}^n$ to denote the entire label matrix. We assume the quality of predictions $\hat{\mathbf{Y}}$ is jointly evaluated against the observed labels \mathbf{Y} by a *task utility* $\Psi(\mathbf{Y}, \hat{\mathbf{Y}})$, and define the *optimal (Bayes) prediction* $\hat{\mathbf{Y}}^*$ as the one maximizing the *expected task utility* Ψ_{ETU} :

$$\hat{\mathbf{Y}}^* = \underset{\hat{\mathbf{Y}} \in \mathcal{Y}_k^n}{\operatorname{argmax}} \mathbb{E}_{\mathbf{Y}|\mathbf{X}} [\Psi(\mathbf{Y}, \hat{\mathbf{Y}})] =: \underset{\hat{\mathbf{Y}} \in \mathcal{Y}_k^n}{\operatorname{argmax}} \Psi_{\text{ETU}}(\hat{\mathbf{Y}}). \quad (1)$$

We consider task utilities $\Psi(\mathbf{Y}, \hat{\mathbf{Y}})$ that linearly decompose over labels, i.e., there exists ψ^j such that

$$\Psi(\mathbf{Y}, \hat{\mathbf{Y}}) = \sum_{j=1}^m \psi^j(\mathbf{y}_{:,j}, \hat{\mathbf{y}}_{:,j}). \quad (2)$$

We allow the functions ψ^j to be non-linear themselves and different for each label j . This is a large class of functions, which encompasses weighted instance-wise and macro-averaged utilities, the two groups of functions which we thoroughly analyze in the next sections.

Let us next define the binary confusion matrix $\mathbf{C}(\mathbf{y}, \hat{\mathbf{y}})$ for a vector \mathbf{y} of n ground truth labels and a corresponding vector $\hat{\mathbf{y}}$ of binary predictions:

$$\mathbf{C}(\mathbf{y}, \hat{\mathbf{y}}) := \begin{pmatrix} \frac{1}{n} \sum_{i=1}^n (1 - y_i)(1 - \hat{y}_i) & \frac{1}{n} \sum_{i=1}^n (1 - y_i)\hat{y}_i \\ \frac{1}{n} \sum_{i=1}^n y_i(1 - \hat{y}_i) & \frac{1}{n} \sum_{i=1}^n y_i\hat{y}_i \end{pmatrix}. \quad (3)$$

By indexing from 0, the entry c_{00} corresponds to true negatives, c_{01} to false positives, c_{10} to false negatives, and c_{11} to true positives. We define the *multi-label confusion tensor*² $\mathbf{C}(\mathbf{Y}, \hat{\mathbf{Y}}) := [\mathbf{C}(\mathbf{y}_{:,1}, \hat{\mathbf{y}}_{:,1}), \dots, \mathbf{C}(\mathbf{y}_{:,m}, \hat{\mathbf{y}}_{:,m})]$ being the concatenation of binary confusion matrices of all m labels.

Assuming the utility function (2) to be invariant under instance reordering, i.e., its value does not change if rows of both matrices are re-ordered using the same permutation, we can define Ψ in terms of confusion matrices, instead of ground-truth labels and predictions (shown in [Appendix A.1](#)):

$$\Psi(\mathbf{Y}, \hat{\mathbf{Y}}) = \Psi(\mathbf{C}(\mathbf{Y}, \hat{\mathbf{Y}})) = \sum_{j=1}^m \psi^j(\mathbf{C}(\mathbf{y}_{:,j}, \hat{\mathbf{y}}_{:,j})). \quad (4)$$

Finally, we assume that we have access to a *label probability estimator* (LPE) $\hat{\boldsymbol{\eta}}(\mathbf{x})$ that estimates the marginal probability of each label given the instance, $\boldsymbol{\eta}(\mathbf{x}) = (\eta_1(\mathbf{x}), \dots, \eta_m(\mathbf{x})) := \mathbb{E}_{\mathbf{y}|\mathbf{x}}[\mathbf{y}]$. Such an LPE can be attained by fitting a model on an additional training set of n' examples $(\mathbf{x}_i, \mathbf{y}_i)_{i=1}^{n'}$ using a proper composite loss function [37], which is a common approach in XMLC, e.g., [18].

¹We use calligraphic letters for sets \mathcal{S} , bold font for vectors \mathbf{v} with entries v_i , bold capital letters for matrices \mathbf{Y} with entries y_{ij} , rows \mathbf{y}_i , and columns $\mathbf{y}_{:,j}$. $\mathbb{1}[S]$ denotes the indicator of event S , and $[s] := \{1, \dots, s\}$

²Notice that the confusion matrix can be computed either for multi-label predictions for a given instance \mathbf{x} or binary predictions for label j obtained on a set of instances \mathbf{X} . In the following, we focus on the latter.

3 Performance measures for tail labels

3.1 Instance-wise weighted utility functions

By assigning utility (or cost) to each correct/wrong prediction for each label, we can construct an *instance-wise weighted utility* $u_w: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ with labels $\mathbf{y} \in \mathcal{Y}$ and predictions $\hat{\mathbf{y}} \in \mathcal{Y}$ as

$$u_w(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{j=1}^m w_{00}^j (1 - y_j)(1 - \hat{y}_j) + w_{01}^j (1 - y_j)\hat{y}_j + w_{10}^j y_j(1 - \hat{y}_j) + w_{11}^j y_j \hat{y}_j. \quad (5)$$

We use w_{00}^j , w_{01}^j , w_{10}^j , and w_{11}^j to express the utility of true negatives, false positives, false negatives, and true positives, respectively. The corresponding task loss results from summing over all instances. By interchanging the order of summation, we can see that it is of form (4):

$$\Psi(\mathbf{Y}, \hat{\mathbf{Y}}) = \sum_{i=1}^n u_w(\mathbf{y}_i, \hat{\mathbf{y}}_i) = \sum_{j=1}^m w_{00}^j c_{00}^j + w_{01}^j c_{01}^j + w_{10}^j c_{10}^j + w_{11}^j c_{11}^j = \sum_{j=1}^m \psi^j(\mathbf{C}(\mathbf{y}_{:j}, \hat{\mathbf{y}}_{:j})). \quad (6)$$

In other words, the instance-wise weighted utilities can be seen as *linear* confusion-based metrics. By choosing $w_{00}^j = w_{11}^j = m^{-1}$ and $w_{01}^j = w_{10}^j = 0$, the expression (5) reduces to the @ k -variant of the *Hamming utility*. Similarly, $w_{11}^j = k^{-1}$ and $w_{00}^j = w_{01}^j = w_{10}^j = 0$ yield precision@ k .

Another example is the popular *propensity-scoring* approach [15], commonly used as a tail-performance metric in XMLC. Here, the weights are computed based on training data through

$$w_{11}^{j,\text{prop}} = k^{-1} (1 + (\log n' - 1)(b + 1)^a (n' \hat{\pi}_j + b)^{-a}), \quad (7)$$

where n' is the number of training instances, $\hat{\pi}_j$ is the empirical prior of label j , and parameters a and b are (potentially) dataset-dependent. This form of weighting has been derived from a missing-labels perspective, so its application to tail labels is not fully justified [41]. Also, it introduces two more hyperparameters, which makes the interpretation and comparison of its values rather difficult. It is not less heuristical than other approaches like power-law or logarithmic weighting, given by:

$$w_{11}^{j,\text{pl}} \propto \hat{\pi}_j^{-\beta}, \quad w_{11}^{j,\text{log}} \propto -\log(\hat{\pi}_j). \quad (8)$$

3.2 Macro-average of non-decomposable utilities

Macro-averaging usually concerns non-decomposable binary utilities such as the F-measure. In this case, we set up $\psi^j(\cdot, \cdot) = m^{-1} \psi(\cdot, \cdot)$ for all labels j in (2), yielding:

$$\Psi(\mathbf{Y}, \hat{\mathbf{Y}}) = m^{-1} \sum_{j=1}^m \psi(\mathbf{y}_{:j}, \hat{\mathbf{y}}_{:j}) = m^{-1} \sum_{j=1}^m \psi(\mathbf{C}(\mathbf{y}_{:j}, \hat{\mathbf{y}}_{:j})). \quad (9)$$

By using $\psi_{\text{pr}}(\mathbf{C}) := c_{11}/(c_{11} + c_{01})$, this becomes *macro-precision*, for $\psi_{\text{rec}}(\mathbf{C}) := c_{11}/(c_{11} + c_{10})$ we get *macro-recall*, and for $\psi_{F_\beta}(\mathbf{C}) := (\beta + 1)c_{11}/((1 + \beta)c_{11} + \beta^2 c_{10} + c_{01})$ the *macro-F-measure*.

Another measure that is promising for the evaluation of long-tailed performance is *coverage*. It is sometimes used as an auxiliary measure in XMLC [15, 3, 42, 41]. This metric detects for how many different labels the classifier is able to make *at least one* correct prediction. In our framework, this is achieved by using an indicator function on the true positives, $\psi_{\text{cov}}(\mathbf{C}) := \mathbb{1}[c_{11} > 0]$.

4 Optimal predictions

Let us start with the observation that the label probabilities $\boldsymbol{\eta}$ are sufficient to make optimal predictions. With the assumption that $\mathbb{P}(\mathbf{Y}|\mathbf{X}) = \prod_{i=1}^n \mathbb{P}(\mathbf{y}_i|\mathbf{x}_i)$, we obtain (cf. Appendix A.2):

$$\mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\Psi(\mathbf{Y}, \hat{\mathbf{Y}})] = \sum_{j=1}^m \sum_{\mathbf{y}' \in \{0,1\}^n} \left(\prod_{i=1}^n \eta_j(\mathbf{x}_i) y'_i + (1 - \eta_j(\mathbf{x}_i))(1 - y'_i) \right) \psi^j(\mathbf{y}', \hat{\mathbf{y}}_{:j}). \quad (10)$$

This equation lays out a daunting optimization task, as it requires summing over 2^n summands \mathbf{y}' . In case of binary classification, there exist methods to solve the problem exactly in $\mathcal{O}(n^3)$, or in $\mathcal{O}(n^2)$

in some special cases [32]. By using *semi-empirical* quantities (defined below), [12] provides an approximate algorithm that runs in $O(n)$. Following this approach, we construct a semi-empirical ETU approximation. If this approximation results in a *linear* function of the predictions, the problem decomposes over instances and can be solved easily. Otherwise, we use an algorithm that leads to locally optimal predictions. A minor modification of this algorithm can be used for coverage.

4.1 Semi-empirical ETU approximation

Since the entries of the confusion matrix are linearly dependent, it suffices to use three independent combinations. More precisely, we parameterize the confusion matrix by the true positives $t = c_{11}$, predicted positives $q = c_{11} + c_{01}$, and ground-truth positives $p = c_{11} + c_{10}$, and use $\mathbf{t} = (t_1, \dots, t_m)$, $\mathbf{q} = (q_1, \dots, q_m)$, and $\mathbf{p} = (p_1, \dots, p_m)$ to reformulate the ETU objective:

$$\Psi_{\text{ETU}}(\hat{\mathbf{Y}}) = \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\Psi(\mathbf{C}(\mathbf{Y}, \hat{\mathbf{Y}}))] = \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\Psi(\mathbf{t}, \mathbf{q}, \mathbf{p})] = \mathbb{E}_{\mathbf{Y}|\mathbf{X}}\left[\sum_{j=1}^m \psi^j(t_j, q_j, p_j)\right]. \quad (11)$$

In order to compute Ψ_{ETU} , one needs to take into account every possible combination of confusion-matrix values, and calculate the corresponding value of Ψ , which is then averaged according to the respective probabilities. A computationally easier approach is to take the expectation over the labels first, leading to *semi-empirical* quantities:

$$\tilde{\mathbf{t}} := \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\mathbf{t}], \quad \tilde{\mathbf{q}} := \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\mathbf{q}] = \mathbf{q}, \quad \tilde{\mathbf{p}} := \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\mathbf{p}], \quad (12)$$

where $\tilde{\mathbf{q}} = \mathbf{q}$ follows because the number of predicted positives depends only on the predictions $\hat{\mathbf{Y}}$. This allows us to define the semi-empirical ETU risk

$$\tilde{\Psi}_{\text{ETU}}(\hat{\mathbf{Y}}) := \Psi(\tilde{\mathbf{t}}, \mathbf{q}, \tilde{\mathbf{p}}) \approx \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\Psi(\mathbf{t}, \mathbf{q}, \mathbf{p})] = \Psi_{\text{ETU}}(\hat{\mathbf{Y}}). \quad (13)$$

In particular, the third argument to Ψ , $\tilde{\mathbf{p}}$, is a constant that does not depend on predictions.

Note that, if Ψ is *linear* in all arguments *depending on the random variable \mathbf{Y}* , then the approximation is exact, due to the linearity of expectations. Aside from instance-wise measures, which we showed to be linear above, the approximation is also exact for the more general class of functions of the form

$$\psi(t, q, p) = f_t(q) \cdot t + f_q(q) + f_p(q) \cdot p. \quad (14)$$

An important example is macro-precision, with $f_t(q) = q^{-1}$ and $f_q = f_p = 0$. In the general case, $\tilde{\Psi}_{\text{ETU}}$ as a surrogate for Ψ_{ETU} leads only to $\mathcal{O}(1/\sqrt{n})$ error as will be shown in [Theorem 5.2](#), while substantially simplifying the optimization process.

4.2 Linear confusion-matrix measures

We start the discussion on optimization of [\(13\)](#) with a special case in which $\tilde{\Psi}_{\text{ETU}}$ is *linear in the prediction-dependent arguments t, q* , that is, if

$$\psi(t, q, p) = f_t(p) \cdot t + f_q(p) \cdot q, \quad (15)$$

i.e., both $f_t(p)$ and $f_q(p)$ depend on p only, and $f_p(p) \cdot p$ can be dropped as it is a constant.

Aside from instance-wise weighted utilities (cf. [Appendix A.3](#)), which are linear in all arguments, this form also holds for weights dependent on the (empirical) label priors, e.g., power law weights of the form $f_t(p) = p^{-\alpha}$ and $f_q = 0$, which reduce to macro-recall for $\alpha = 1$. If one defines the weights with respect to externally determined label priors, i.e., approximations to $\mathbb{E}[y_j]$, which are fixed and thus independent of the test sample \mathbf{X} , then the power-law metrics turn into instance-wise weighted utilities.

From [\(6\)](#) we know that we can reformulate the optimization problem using an instance-wise weighted utility u_w with weights:

$$w_{11}^j = f_t(p_j) + f_q(p_j), \quad w_{01}^j = f_q(p_j), \quad w_{10}^j = 0, \quad w_{00}^j = 0. \quad (16)$$

Hence, the optimal predictions can be derived for each instance $\mathbf{x} \in \mathcal{X}$ separately, leading to

$$\hat{\mathbf{y}}^* = \operatorname{argmax}_{\hat{\mathbf{y}} \in \mathcal{Y}_k} \mathbb{E}_{\mathbf{y}|\mathbf{x}}[u_w(\mathbf{y}, \hat{\mathbf{y}})]. \quad (17)$$

Plugging in the definition of u_w from [\(5\)](#), and collecting terms, the expected loss is of the form

$$\mathbb{E}_{\mathbf{y}|\mathbf{x}}[u_w(\mathbf{y}, \hat{\mathbf{y}})] = \sum_{j=1}^m \mathbb{E}_{\mathbf{y}|\mathbf{x}}\left[w_{11}^j y_j \hat{y}_j + w_{01}^j (1 - y_j) \hat{y}_j\right] = \sum_{j=1}^m \hat{y}_j (\eta_j(\mathbf{x}) f_t(p_j) + f_q(p_j)), \quad (18)$$

where we call $g_j(\mathbf{x}) := \eta_j(\mathbf{x})f_t(p_j) + f_q(p_j)$ the *gain* of predicting label j for a given instance \mathbf{x} . The optimal prediction is to select the k labels with the largest values $g_j(\mathbf{x})$,

$$\hat{\mathbf{y}}^* = \operatorname{argmax}_{\hat{\mathbf{y}} \in \mathcal{Y}_k} \sum_{j=1}^m g_j(\mathbf{x}) \hat{y}_j = \text{select-top-}k(\mathbf{g}(\mathbf{x})). \quad (19)$$

Here, $\text{select-top-}k(\mathbf{g})$ denotes an operation that maps a vector of scores to a binary vector that contains 1s only at the positions of the k largest elements of \mathbf{g} .

4.3 General non-decomposable macro measures

Finally, we turn to the general problem for budgeted-at- k macro-measures based on non-linear Ψ . As this can be a very hard discrete optimization problem in general, we use an iterative approach based on *block-coordinate ascent* that constructs a sequence of predictions, $\hat{\mathbf{Y}}^0, \hat{\mathbf{Y}}^1, \dots$, with non-decreasing utility, so that we end up with a solution that is locally optimal.

Assume the predictions are fixed for all instances except \mathbf{x}_s , where they are given by \mathbf{z} . In that case, we can write the semi-empirical quantities from (12) as

$$\tilde{\mathbf{t}} = \frac{1}{n} \left(\eta_j(\mathbf{x}_s) z_j + \sum_{i \in [n] \setminus \{s\}} \eta_j(\mathbf{x}_i) \hat{y}_{ij} \right), \quad (20)$$

with analog expansions for \mathbf{q} and $\tilde{\mathbf{p}}$. Plugging into (13) leads to the following optimization:

$$\max_{\mathbf{z} \in \mathcal{Y}_k} \sum_{j=1}^m \psi^j \left(\frac{1}{n} \eta_j(\mathbf{x}_s) z_j + \frac{1}{n} \sum_{i \in [n] \setminus \{s\}} \eta_j(\mathbf{x}_i) \hat{y}_{ij}, \frac{1}{n} z_j + \frac{1}{n} \sum_{i \in [n] \setminus \{s\}} \hat{y}_{ij}, \frac{1}{n} \sum_{i=1}^n \eta_j(\mathbf{x}_i) \right). \quad (21)$$

As everything except $z_j \in \{0, 1\}$ is given, we can interpret ψ^j as a linear function of z_j , and define a gain vector with elements $g_j = \psi^j(1) - \psi^j(0)$. The optimal prediction \mathbf{z}^* is then given by $\mathbf{z}^* = \text{select-top-}k(\mathbf{g})$, in a similar form as in case of linear confusion-matrix measures. We get $\hat{\mathbf{Y}}^{t+1}$ by replacing the s^{th} row of $\hat{\mathbf{Y}}^t$ with \mathbf{z}^* , and know that $\tilde{\Psi}_{\text{ETU}}(\hat{\mathbf{Y}}^{t+1}) \geq \tilde{\Psi}_{\text{ETU}}(\hat{\mathbf{Y}}^t)$. Then we switch to the next instance $s \leftarrow s + 1$, and repeat this process until no more progress is made.

The algorithm starts by predicting k random labels for each instance. To speed up the computations we cache two quantities, \tilde{t}_j^t and q_j^t , for each label:

$$\begin{aligned} \tilde{t}_j^0 &:= \frac{1}{n} \sum_{i=1}^n \eta_j(\mathbf{x}_i) \hat{y}_{ij}^0, & \tilde{t}_j^{t+1} &:= \tilde{t}_j^t + \frac{1}{n} \eta_j(\mathbf{x}_s) (\hat{y}_{sj}^{t+1} - \hat{y}_{sj}^t), \\ q_j^0 &:= \frac{1}{n} \sum_{i=1}^n \hat{y}_{ij}^0, & q_j^{t+1} &:= q_j^t + \frac{1}{n} (\hat{y}_{sj}^{t+1} - \hat{y}_{sj}^t). \end{aligned} \quad (22)$$

With this, we can compute (21) in $\mathcal{O}(m)$ time using the following formulas:

$$\begin{aligned} \psi^j(1) &:= \psi^j \left(\tilde{t}_j^t + \frac{1}{n} \eta_j(\mathbf{x}_s) (1 - \hat{y}_{sj}^t), q_j^t + \frac{1}{n} (1 - \hat{y}_{sj}^t), \tilde{p}_j \right), \\ \psi^j(0) &:= \psi^j \left(\tilde{t}_j^t - \frac{1}{n} \eta_j(\mathbf{x}_s) \hat{y}_{sj}^t, q_j^t - \frac{1}{n} \hat{y}_{sj}^t, \tilde{p}_j \right). \end{aligned} \quad (23)$$

The block coordinate ascent (BCA) procedure is shown in more detail in Algorithm 1. Obviously, the actual implementation cannot use the unknown values $\boldsymbol{\eta}$, but instead has to rely on the LPE estimates $\hat{\boldsymbol{\eta}}(\mathbf{x}_i)$. The stopping criterion for the algorithm is whether, after going over a whole set of instances, the improvement in the objective value over the previous value is lower than ϵ , which ensures that the algorithm terminates for every bounded utility in $\mathcal{O}(1/\epsilon)$. In practice, even with small ϵ , the algorithm usually terminates after a few iterations. The time and space complexity of the single iteration are both $\mathcal{O}(nm)$. If ψ is a linear function, corresponding to an instance-wise weighted utility (6) such as macro-recall, the algorithm recovers the optimal solution in the first iteration, stopping after the second.

In general, Algorithm 1 requires multiple iterations. This can be computationally expensive and requires all of the data to be available at once. We thus propose a greedy algorithm that takes into account only statistics of *previously seen* instances while performing a *single pass* over the dataset, which allows for semi-online optimization. The greedy algorithm is outlined in Appendix B.4.

Algorithm 1 $\text{BCA}(\mathbf{X}, \hat{\boldsymbol{\eta}}, k, \epsilon)$

```

1:  $\hat{\mathbf{y}}_i \leftarrow \text{select-random-}k(m)$  for all  $i \in [n]$ 
2:  $\tilde{\mathbf{t}} \leftarrow \frac{1}{n} \sum_{i=1}^n \hat{\boldsymbol{\eta}}(\mathbf{x}_i) \odot \hat{\mathbf{y}}_i$ 
3:  $\mathbf{q} \leftarrow \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{y}}_i$ ,  $\tilde{\mathbf{p}} \leftarrow \frac{1}{n} \sum_{i=1}^n \hat{\boldsymbol{\eta}}(\mathbf{x}_i)$ 
4:  $u_{\text{old}} \leftarrow -\infty$ ,  $u_{\text{new}} \leftarrow \Psi(\tilde{\mathbf{t}}, \mathbf{q}, \tilde{\mathbf{p}})$ 
5: while  $u_{\text{new}} > u_{\text{old}} + \epsilon$  do
6:   for  $s \in \text{shuffle}([n])$  do
7:      $\tilde{\mathbf{t}} \leftarrow \tilde{\mathbf{t}} - \frac{1}{n} \hat{\boldsymbol{\eta}}(\mathbf{x}_s) \odot \hat{\mathbf{y}}_s$ ,  $\mathbf{q} \leftarrow \mathbf{q} - \frac{1}{n} \hat{\mathbf{y}}_s$ 
8:     for  $j \in [m]$  do
9:        $\psi^j(1) \leftarrow \psi(\tilde{t}_j + \frac{1}{n} \hat{\eta}_j(\mathbf{x}_s), q_j + \frac{1}{n}, \tilde{p}_j)$ 
10:       $\psi^j(0) \leftarrow \psi(\tilde{t}_j, q_j, \tilde{p}_j)$ 
11:       $g_j \leftarrow \psi^j(1) - \psi^j(0)$ 
12:       $\hat{\mathbf{y}}_s \leftarrow \text{select-top-}k(\mathbf{g})$ 
13:       $\tilde{\mathbf{t}} \leftarrow \tilde{\mathbf{t}} + \frac{1}{n} \hat{\boldsymbol{\eta}}(\mathbf{x}_s) \odot \hat{\mathbf{y}}_s$ ,  $\mathbf{q} \leftarrow \mathbf{q} + \frac{1}{n} \hat{\mathbf{y}}_s$ 
14:     $u_{\text{old}} \leftarrow u_{\text{new}}$ ,  $u_{\text{new}} \leftarrow \Psi(\tilde{\mathbf{t}}, \mathbf{q}, \tilde{\mathbf{p}})$ 
15: return  $\hat{\mathbf{Y}}$ 
```

4.4 Optimization of coverage

One measure for which the ETU approximation is not exact is coverage as $\psi_{\text{cov}}(t, q, p) := \mathbb{1}[t > 0]$ is *nonlinear*. In this case, we can do better than Algorithm 1 by reformulating Ψ_{ETU} for coverage as

$$\Psi_{\text{ETU}}(\hat{\mathbf{Y}}) = \mathbb{E}_{\mathbf{Y}|\mathbf{X}} \left[m^{-1} \sum_{j=1}^m \mathbb{1}[t_j > 0] \right] = 1 - m^{-1} \sum_{j=1}^m \prod_{i=1}^n (1 - \eta_j(\mathbf{x}_i) \hat{y}_{ij}) , \quad (24)$$

and performing block coordinate ascent directly on this expression, as detailed in Appendix B.3.

5 Regret bounds

Computing optimal predictions relies on an access to the conditional marginal probabilities $\boldsymbol{\eta}(\mathbf{x})$. In practice, however, $\boldsymbol{\eta}(\mathbf{x})$ are unknown, and are replaced by the LPE $\hat{\boldsymbol{\eta}}(\mathbf{x})$ to do inference (*plug-in* approach). Furthermore, we replaced the ETU objective Ψ_{ETU} with an approximation $\tilde{\Psi}_{\text{ETU}}$. As generally $\hat{\boldsymbol{\eta}}(\mathbf{x}) \neq \boldsymbol{\eta}(\mathbf{x})$ and $\Psi_{\text{ETU}} \neq \tilde{\Psi}_{\text{ETU}}$, this procedure may result in sub-optimal predictions, the errors of which we would like to control.

If we are able to control the change of the utility under small changes in its arguments, we can show that the suboptimality of the plug-in predictor relative to the optimal predictor (called Ψ -regret) decreases with increasing test size n and decreasing probability estimation error.

To this end we assume that each ψ^j is a p -Lipschitz function, which allows us to bound the approximation error.

Definition 5.1 (p -Lipschitz [13]). A binary classification metric $\psi(t, q, p)$ is said to be p -Lipschitz if

$$|\psi(t, q, p) - \psi(t', q', p')| \leq L_t(p)|t - t'| + L_q(p)|q - q'| + L_p(p)|p - p'|, \quad (25)$$

for any $q, q' \in [0, 1]$, $p, p' \in (0, 1)$, $0 \leq t \leq \min(p, q)$, and $0 \leq t' \leq \min(p', q')$. The constants $L_t(p)$, $L_q(p)$, $L_p(p)$ are allowed to depend on p , in contrast to the standard Lipschitz functions.

As shown in Appendix C.3, most of metrics of interest satisfy p -Lipschitz assumption, including the linear confusion-matrix measures (6) with fixed weights (e.g., Hamming utility, precision), macro-recall, macro-F-measure, etc., with macro-precision and coverage being notable exceptions.

Theorem 5.2. Let each ψ^j be p -Lipschitz with constants $L_t^j(p)$, $L_q^j(p)$, $L_p^j(p)$. For any $\hat{\mathbf{Y}}$ it holds:

$$|\Psi_{\text{ETU}}(\hat{\mathbf{Y}}; \mathbf{X}) - \tilde{\Psi}_{\text{ETU}}(\hat{\mathbf{Y}}; \mathbf{X})| \leq \frac{1}{2\sqrt{n}} \left(\sum_{j=1}^m (L_t^j(\tilde{p}_j) + L_p^j(\tilde{p}_j)) \right). \quad (26)$$

Thus, using $\tilde{\Psi}_{\text{ETU}}$ as a surrogate for Ψ_{ETU} leads only to $\mathcal{O}(1/\sqrt{n})$ error, diminishing with the test size, while substantially simplifying the optimization process.

Given a decomposable metric of the form (4), let $\tilde{\mathbf{Y}}^\dagger$ be the plug-in prediction matrix optimizing the semi-empirical ETU with plugged-in probability estimates, $\Psi(\mathbb{E}_{\mathbf{y} \sim \hat{\boldsymbol{\eta}}(\mathbf{x})}[t], \mathbf{q}, \mathbb{E}_{\mathbf{y} \sim \hat{\boldsymbol{\eta}}(\mathbf{x})}[\mathbf{p}])$.

Theorem 5.3. Let $\tilde{\mathbf{Y}}^\dagger$ be defined as above. Under the assumptions of Theorem 5.2:

$$\Psi_{\text{ETU}}(\hat{\mathbf{Y}}^*; \mathbf{X}) - \Psi_{\text{ETU}}(\tilde{\mathbf{Y}}^\dagger; \mathbf{X}) \leq \frac{m}{\sqrt{n}} B + 2 \frac{\sqrt{m}}{n} B \sum_{i=1}^n \|\boldsymbol{\eta}(\mathbf{x}_i) - \hat{\boldsymbol{\eta}}(\mathbf{x}_i)\|_2, \quad (27)$$

where $B := \sqrt{m^{-1} \sum_{j=1}^m (L_t^j(\tilde{p}_j) + L_p^j(\tilde{p}_j))^2}$ is the quadratic mean of the Lipschitz constants.

A similar statement, presented in Appendix C.4, can be made for the unapproximated ETU case.

Thus, the methods described in Section 4 can be used with probability estimates replacing the true marginals, and as long as the estimator is reliable, the resulting predictions will have small Ψ -regret. This also justifies the plug-in approach used in the experiments in Section 7.

6 Efficient inference

The optimization algorithms introduced so far need to obtain $\hat{\eta}_j(\mathbf{x}_i)$ for *all* labels and instances first. Then, the BCA inference is of $\mathcal{O}(nm)$ time and $\mathcal{O}(nm)$ space complexity for a single iteration. This

is problematic in the setting of XMLC, where many methods aim to predict probabilities only for top labels in time sublinear in m . Fortunately, this characteristic of XMLC algorithms can be combined with the introduced algorithms to efficiently obtain an approximate solution. Instead of predicting all η , we can predict probabilities only for top- k' labels with highest $\hat{\eta}_j$, where $k \ll k' \ll m$. For all other labels we then assume $\hat{\eta}_j = 0$. Under the natural assumption that ψ is *non-decreasing* in true positives and *non-increasing* in predicted positives, we can leverage the sparsity and consider labels with non-zero $\hat{\eta}_j$ only (using sparse vectors to represent $\hat{\eta}(x_i)$) to reduce the time and space complexity to $\mathcal{O}(n(k' + k \log k))$ and $\mathcal{O}(n(k' + k))$, respectively. As in real-world datasets the number of relevant labels $\|\mathbf{y}\|_1$ is much lower than m , and most $\eta_j(x_i)$ are close to 0, with reasonably selected k' , according to Theorem 5.3, we should only slightly increase the regret. A pseudocode for the sparse variant of Algorithm 1 can be found in Appendix D.

Alternatively, we can leverage probabilistic label trees (PLTs) [16], a popular approach in XMLC [35, 45, 19, 50, 7], to efficiently search for “interesting” labels. Originally, PLTs find top labels with the highest $\hat{\eta}_j$. In [46], an A^* -search algorithm has been introduced for finding k labels with highest value of $g_j \eta_j(x)$, where $g_j \in [0, \infty)$ is a gain assigned to label j . In Appendix F, we present a more general version of this procedure that can be used to efficiently obtain an exact solution of presented BCA or Greedy algorithms for some of the metrics considered in this work.

7 Experiments

To empirically test the introduced framework, we use popular benchmarks from the XMLC repository [6]. We train the LIGHTXML [18] model (with suggested default hyper-parameters) on provided training sets to obtain $\hat{\eta}$ for all test instances. We then plug these estimates into different inference strategies and report the results across the discussed measures. To run the optimization algorithm efficiently, we use $k' = 100$ or $k' = 1000$ to pre-select for each instance the top k' labels with the highest $\hat{\eta}_j$ as described in Section 6³.

We use the following inference strategies:

- TOP-K– the optimal strategy for precision@ k : selection of k labels with the highest η_j (default prediction strategy in many XMLC methods).
- PS-K– the optimal strategy for propensity-scored precision@ k : selection of k labels with the highest $w_{11}^{j, \text{prop}} \eta_j$, with $w_{11}^{j, \text{prop}}$ given by the empirical model of Jain et al. [15] (Equation 7) with values a and b recommended by the authors.
- POW-K, LOG-K– the optimal strategy for power-law and log weighted instance-wise utilities: selection of k labels with the highest $w_{11}^{j, \text{pl}} \eta_j$ or $w_{11}^{j, \text{log}} \eta_j$. For power-law, we use $\beta = 0.5$.
- MACRO-P_{BCA}, MACRO-R_{BCA}, MACRO-F1_{BCA}, COV_{BCA}– the block coordinate ascent (Algorithm 1) for optimizing macro-precision, -recall, -F1, and coverage,

We expect a strategy suited for a given metric to obtain the best results on this metric. Nevertheless, this might not always be the case, as in the derivation of our algorithms, we needed to apply different types of approximation to scale them to XMLC problems. We are mainly interested in the performance on the general non-decomposable macro measures since they seem to be well-tailored to long tails, and their optimization is the most challenging. The results are presented in Table 2. Notice that in almost all cases, the specialized inference strategies are indeed the best on the measure they aim to optimize and achieve substantial gains on corresponding metrics compared to the basic TOP-K inference. The other weighted strategies, PS-K, POW-K, LOG-K, usually provide a much smaller improvement over TOP-K and never beat strategies designed for specific macro measures. As the reported performance depends on three things: the inherent difficulty of the data, the success of the inference algorithm, and the quality of the provided




Figure 1: Results of an inference strategy with a mixed utility on AMAZONCAT-13K and $k = 3$. The green line shows the results for different interpolations between two measures.

³A code to reproduce all the experiments: <https://github.com/mwydmuch/xCOLUMNS>

Table 2: Results of different inference strategies on @k measures with $k \in \{3, 5, 10\}$. Notation: P—precision, R—recall, F1—F1-measure, Cov—Coverage. The green color indicates cells in which the strategy matches the metric. The best results are in **bold** and the second best are in *italic*.

Inference strategy	Instance @3		Macro @3				Instance @5		Macro @5				Instance @10		Macro @10			
	P	R	P	R	F1	Cov	P	R	P	R	F1	Cov	P	R	P	R	F1	Cov
EURLEX-4K, $k' = 100$																		
TOP-K	74.20	44.21	24.85	16.45	18.63	31.27	62.31	60.59	27.43	25.99	25.50	39.76	39.29	75.11	22.15	36.59	<i>26.01</i>	47.55
PS-K	69.01	41.06	30.12	23.72	25.08	40.29	60.77	59.17	29.10	29.89	28.17	44.03	38.98	74.54	21.36	38.27	25.78	49.47
POW-K $\beta=0.5$	65.35	38.95	30.49	25.26	25.93	42.11	58.53	57.10	28.88	<i>31.25</i>	28.53	45.48	38.48	73.66	20.77	38.85	25.43	50.13
LOG-K	<i>71.03</i>	<i>42.30</i>	29.04	22.10	23.76	37.99	<i>61.51</i>	<i>59.87</i>	28.48	28.47	27.26	42.32	<i>39.12</i>	<i>74.80</i>	21.60	37.49	25.83	48.61
MACRO-P _{BCA}	41.13	24.29	38.78	16.77	20.63	40.60	31.05	30.33	38.70	18.46	21.59	41.66	18.45	35.65	37.94	20.28	21.82	42.89
MACRO-R _{BCA}	41.47	24.61	30.36	<i>25.16</i>	24.37	44.26	38.32	37.54	28.97	31.27	27.05	<i>49.04</i>	29.80	57.58	22.99	38.14	25.99	52.15
MACRO-F1 _{BCA}	61.96	36.94	<i>34.42</i>	25.13	27.59	<i>44.52</i>	56.18	54.70	<i>33.82</i>	29.52	30.13	45.93	33.55	64.29	<i>32.78</i>	32.21	30.47	47.56
COV _{BCA}	24.00	14.03	31.46	21.30	19.05	47.19	16.40	15.95	29.86	24.53	18.81	50.09	9.31	18.05	26.99	28.28	17.41	52.84
AMAZONCAT-13K, $k' = 100$																		
TOP-K	83.35	63.01	25.71	11.37	13.83	33.05	68.01	79.01	33.92	31.34	29.28	52.60	41.72	89.41	24.63	51.76	28.84	68.61
PS-K	<i>79.09</i>	<i>60.15</i>	43.89	38.64	38.56	62.06	<i>66.63</i>	<i>77.89</i>	38.41	45.00	38.34	65.82	<i>41.50</i>	<i>89.11</i>	22.56	58.69	28.59	74.29
POW-K $\beta=0.5$	66.52	50.75	38.18	46.42	39.26	67.76	56.94	68.14	31.65	<i>54.06</i>	36.90	71.95	37.86	83.91	19.26	<i>64.05</i>	26.29	77.27
LOG-K	75.30	56.79	41.09	33.92	35.04	56.05	64.07	75.21	37.40	41.86	36.78	61.84	40.66	88.05	23.09	55.85	28.91	71.37
MACRO-P _{BCA}	54.97	41.61	64.27	29.22	35.76	76.18	41.53	49.66	63.81	30.43	35.99	76.75	25.32	57.33	61.84	33.06	<i>35.08</i>	78.17
MACRO-R _{BCA}	47.74	37.13	31.40	58.32	34.68	<i>80.71</i>	38.93	48.26	25.17	65.47	30.37	<i>82.91</i>	26.50	61.99	17.53	72.63	<i>23.22</i>	<i>84.84</i>
MACRO-F1 _{BCA}	61.70	53.86	<i>51.95</i>	<i>48.22</i>	47.93	<i>77.83</i>	60.70	71.93	<i>50.89</i>	52.32	49.48	79.63	37.80	82.24	<i>49.43</i>	52.17	49.21	81.39
COV _{BCA}	4.53	2.29	34.93	35.16	15.91	82.67	3.20	2.63	29.40	39.05	14.23	84.39	2.13	3.36	19.02	44.28	10.74	85.40
WIKI-31K, $k' = 100$																		
TOP-K	77.17	13.48	2.01	0.50	0.72	2.54	68.31	19.59	2.66	0.92	1.24	3.72	52.00	29.05	3.55	2.07	2.34	6.14
PS-K	<i>67.95</i>	<i>11.89</i>	3.89	1.47	1.93	5.47	<i>62.65</i>	<i>18.07</i>	4.21	2.14	2.56	6.54	<i>50.16</i>	28.20	4.64	3.62	3.61	9.01
POW-K $\beta=0.5$	55.06	9.59	4.49	2.14	2.59	6.98	50.12	14.40	4.83	3.13	3.37	8.53	40.15	22.57	5.07	5.10	4.47	11.37
LOG-K	65.20	11.34	3.40	1.24	1.66	4.84	57.74	16.49	3.74	1.86	2.25	5.99	43.82	24.37	4.24	3.46	3.37	8.61
MACRO-P _{BCA}	32.86	5.66	9.13	2.55	3.38	9.21	30.52	8.56	9.68	2.86	3.74	9.82	24.81	13.62	9.79	2.98	3.85	10.04
MACRO-R _{BCA}	13.78	2.36	4.77	3.24	3.10	7.26	14.01	3.96	5.79	4.71	4.05	10.72	13.88	7.79	5.72	7.21	<i>5.02</i>	<i>15.72</i>
MACRO-F1 _{BCA}	37.32	6.52	<i>7.81</i>	3.15	4.10	<i>10.42</i>	38.98	11.24	<i>8.30</i>	4.14	5.08	<i>11.88</i>	39.31	22.28	<i>8.18</i>	5.11	5.78	13.07
COV _{BCA}	27.00	4.63	6.24	<i>3.18</i>	<i>3.41</i>	11.04	20.80	5.88	6.42	<i>4.60</i>	<i>4.18</i>	13.23	13.95	7.78	6.23	6.88	4.81	16.39
WIKIPEDIA-LARGE-500K, $k' = 1000$																		
TOP-K	56.02	45.70	20.15	18.87	17.14	32.24	43.12	54.28	20.51	25.86	19.84	40.62	<i>27.01</i>	63.13	17.16	34.79	19.57	50.15
PS-K	<i>54.91</i>	<i>45.61</i>	23.33	22.68	20.41	37.85	<i>42.86</i>	54.51	21.84	29.15	21.77	45.03	27.05	63.43	16.70	37.58	19.87	53.67
POW-K $\beta=0.5$	51.81	43.94	23.73	23.67	21.13	39.36	40.55	52.72	21.87	30.03	22.18	46.21	26.09	62.26	16.48	38.17	19.85	54.38
LOG-K	54.82	45.21	21.43	20.18	18.38	34.28	42.66	54.09	20.97	26.81	20.49	41.86	26.97	<i>63.19</i>	17.08	35.51	19.69	51.04
MACRO-P _{BCA}	25.19	21.81	37.69	20.18	23.44	45.08	16.25	22.54	37.88	21.12	<i>24.10</i>	46.23	8.48	22.97	37.77	21.46	<i>24.19</i>	46.65
MACRO-R _{BCA}	43.40	39.60	25.35	27.55	23.72	46.30	33.58	47.34	22.08	33.56	23.63	<i>51.91</i>	21.81	56.27	15.92	40.97	20.07	<i>58.42</i>
MACRO-F1 _{BCA}	43.82	36.40	35.42	23.69	26.01	<i>46.36</i>	32.96	41.21	35.79	26.19	27.64	49.20	19.32	44.17	35.55	27.43	28.15	50.67
COV _{BCA}	27.31	24.55	25.92	26.76	21.60	50.16	19.46	28.06	23.14	<i>32.13</i>	21.08	55.40	11.59	32.09	18.45	38.36	18.08	61.15
AMAZON-670K, $k' = 1000$																		
TOP-K	41.71	24.08	10.77	9.68	9.51	14.35	37.71	35.23	14.26	15.16	13.76	20.41	25.35	<i>46.74</i>	14.83	21.56	16.20	26.85
PS-K	41.05	23.77	11.86	10.54	10.42	15.50	37.54	35.13	14.86	15.73	14.31	21.16	25.28	46.65	14.89	21.85	16.34	27.22
POW-K $\beta=0.5$	40.90	23.71	11.99	10.65	10.53	15.64	37.47	35.07	14.98	15.84	14.42	21.30	25.24	46.60	14.90	21.92	16.36	27.30
LOG-K	<i>41.54</i>	<i>24.00</i>	11.34	10.10	9.98	14.94	<i>37.69</i>	<i>35.22</i>	14.48	15.37	13.97	20.68	25.35	46.75	14.85	21.64	16.24	26.94
MACRO-P _{BCA}	33.79	19.75	17.27	10.53	<i>12.12</i>	17.75	27.52	26.09	21.09	14.38	<i>15.96</i>	21.96	15.02	28.23	23.21	16.36	<i>17.90</i>	24.13
MACRO-R _{BCA}	39.42	22.92	13.71	11.17	11.47	17.12	36.43	34.19	16.48	16.35	15.40	<i>22.61</i>	24.72	45.70	16.30	22.23	17.46	<i>28.08</i>
MACRO-F1 _{BCA}	37.34	21.70	<i>16.49</i>	10.75	12.17	17.65	31.97	30.04	20.39	15.15	16.37	22.25	18.48	34.28	22.79	17.95	18.89	25.12
COV _{BCA}	35.38	20.32	14.04	<i>10.85</i>	11.23	<i>17.70</i>	30.27	28.39	16.44	<i>15.94</i>	14.85	22.93	20.08	37.23	16.20	21.52	16.61	28.15

marginal probabilities, the results might diverge from expectations in some cases. In [Appendix E](#) we provide more details and conduct further experiments to investigate the impact of randomness, stopping conditions, quality of probability estimates, and shortlisting on the results. We also present similar experiments with probabilistic label trees used as an LPE.

Unfortunately, our results show that optimization of macro-measures comes with the cost of a significant drop in performance on instance-wise measures. Ideally, we would like to improve the performance on tail labels without sacrificing too much of general performance. To achieve such a trade-off, we can use straight-forward interpolation of instance-wise precision-at-k, as it is covered by our framework, and a selected macro-measure, since the considered class of utility functions is closed under linear combinations. Such an objective can be optimized by the proposed block-coordinate algorithm without any modification. As an example, we plot in [Table 7](#) the results of optimizing a linear combination of instance-wise precision and the macro-F1 measure: $\Psi(\mathbf{Y}, \hat{\mathbf{Y}}) = (1 - \alpha)\Psi_{\text{Instance-P}}(\mathbf{Y}, \hat{\mathbf{Y}}) + \alpha\Psi_{\text{Macro-F1}}(\mathbf{Y}, \hat{\mathbf{Y}})$, using different values of $\alpha \in [0, 1]$. In [Appendix E.5](#), we provide more plots for different utilities, datasets, and values of k . All the plots show that the instance-vs-macro curve has a nice concave shape that dominates simple baselines. In particular, the initial significant improvement on macro-measures comes with a minor drop in instance-measures, and only if one wants to optimize more strongly for macro-measures, the drop on instance-wise measures becomes more severe.

8 Discussion

The advantage of the ETU framework is that one can use multiple inference strategies without re-training a model. This is especially useful in cases where it is not a-priori clear which performance measure should be optimized, or predictions for different purposes are needed at the same time. On the other hand, as this framework optimizes the performance directly on a given test set, it is not designed to make predictions for single instances independently. The ETU framework has mainly been studied in the case of binary classification problems. We are not aware of any work that focuses on optimizing the complex performance metrics in the ETU framework for multi-label classification. One could try to generalize the results from binary classification, but the existing algorithms might not scale well to the extreme number of labels and they do not take the budget of k labels into account.

Our paper gives novel and non-trivial results regarding this challenging optimization problem. We have thoroughly analyzed the ETU framework for a wide class of performance metrics, derived optimal prediction rules and constructed their computationally efficient approximations with provable regret guarantees and being robust against model misspecification. Our algorithm, based on block coordinate descent, scales effortlessly to XMLC problems and obtains promising empirical results.

Overall, we identified four categories of utilities, that differ in the complexity of the optimization algorithm—whether to use instance-wise optimization as in Section 4.2, or the block coordinate-ascent (Algorithm 1)—and the guarantees for the result—whether semi-empirical quantities (Section 4.1) lead to an optimal solution, or a suboptimal with error bounded by Theorem 5.2. These are given as follows and summarized in Table 3:

Table 3: Four different classes of utilities

Linear in	Approx.	Algorithm
t, q, p	No	Instance-wise
t, q	Yes	Instance-wise
t, p	No	Block-Coordinate
—	Yes	Block-Coordinate

- **Fully linear:** Optimal predictions for metrics that are linear in all entries of the confusion matrix, as (6), can be solved *exactly* in an *instance-wise* manner. Examples are classical metrics such as instance-wise precision@ k , or propensity-scored precision@ k .
- **Linear in predictions:** *Approximately optimal* predictions for metrics that are linear in the predictions as given in (15) can be obtained using *instance-wise* optimization, by switching from Ψ_{ETU} to $\tilde{\Psi}_{\text{ETU}}$. An example is macro recall@ k .
- **Linear in labels:** If a metric is linear in the label variables as given in (14), then $\tilde{\Psi}_{\text{ETU}} \equiv \Psi_{\text{ETU}}$. However, the resulting combinatorial optimization problem for $\tilde{\Psi}_{\text{ETU}}$ is still complex enough, and we can solve it *only locally*. An example is macro precision@ k .
- **Nonlinear metrics:** If none of the above apply, we have $\tilde{\Psi}_{\text{ETU}} \neq \Psi_{\text{ETU}}$, and have to solve it *locally* using *block-coordinate ascent*. This is the case of macro F -measure@ k , or coverage@ k .

The macro-averaged metrics budgeted at k are attractive for measuring the long-tail performance. Macro-averaging treats all the labels as equally important, while the budget of k predictions requires the prediction algorithm to choose the labels “wisely.” We believe that this approach is substantially better than the one based on propensity-scored metrics. It is important to make the distinction between a metric used as a surrogate for training or inference, and its use as a performance metric in itself. While the former might be well justified for many metrics, the latter not necessarily as a metric might not have a clear interpretation of the calculated numbers.

Finally, the proposed framework is a plug-in approach that works on estimates of marginal probabilities and can be seamlessly applied to many existing state-of-the-art XMLC algorithms that are able to output such predictions, including XR-Transformer [51], CascadeXML [20], and algorithmic approaches for dealing with missing labels [40, 36]. It can also be combined with recently proposed methods that try to improve predictive performance on the tail labels by, e.g., leveraging labels features to estimate labels-correlations [29, 39, 9, 52] or to do data augmentation and generate new training points for tail labels [43, 21, 8].

Acknowledgments

A part of computational experiments for this paper had been performed in Poznan Supercomputing and Networking Center. We want to acknowledge the support of Academy of Finland via grants 347707 and 348215, and also thank Mohammadreza Qaraei for providing pre-generated LPEs, $\hat{\eta}$, obtained by LightXML using computational resources of CSC – IT Center for Science, Finland.

References

- [1] Rahul Agrawal, Archit Gupta, Yashoteja Prabhu, and Manik Varma. Multi-label learning with millions of labels: Recommending advertiser bid phrases for web pages. In *22nd International World Wide Web Conference, WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013*, pages 13–24. International World Wide Web Conferences Steering Committee / ACM, 2013.
- [2] Rohit Babbar and Bernhard Schölkopf. Dismec: Distributed sparse machines for extreme multi-label classification. In *Proceedings of the tenth ACM international conference on web search and data mining*, pages 721–729, 2017.
- [3] Rohit Babbar and Bernhard Schölkopf. Data scarcity, robustness and extreme multi-label classification. *Machine Learning*, 108(8):1329–1351, 2019.
- [4] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [5] Alina Beygelzimer, John Langford, Yury Lifshits, Gregory B. Sorkin, and Alexander L. Strehl. Conditional probability tree estimation analysis and algorithms. In *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21, 2009*, pages 51–58. AUAI Press, 2009.
- [6] Kush. Bhatia, Kunal. Dahiya, Himanshu Jain, Anshul Mittal, Yashoteja Prabhu, and Manik Varma. The extreme classification repository: Multi-label datasets and code, 2016. URL <http://manikvarma.org/downloads/XC/XMLRepository.html>.
- [7] Wei-Cheng Chang, Hsiang-Fu Yu, Kai Zhong, Yiming Yang, and Inderjit S. Dhillon. Taming pretrained transformers for extreme multi-label text classification. In Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash, editors, *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 3163–3171. ACM, 2020.
- [8] Eli Chien, Jiong Zhang, Cho-Jui Hsieh, Jyun-Yu Jiang, Wei-Cheng Chang, Olgica Milenkovic, and Hsiang-Fu Yu. Pina: Leveraging side information in extreme multi-label classification via predicted instance neighborhood aggregation. *arXiv preprint arXiv:2305.12349*, 2023.
- [9] Kunal Dahiya, Ananye Agarwal, Deepak Saini, K Gururaj, Jian Jiao, Amit Singh, Sumeet Agarwal, Purushottam Kar, and Manik Varma. Siamesexml: Siamese networks meet extreme classifiers with 100m labels. In *International Conference on Machine Learning*, pages 2330–2340. PMLR, 2021.
- [10] Ofer Dekel and Ohad Shamir. Multiclass-multilabel classification with more classes than examples. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pages 137–144. JMLR.org, 2010.
- [11] Krzysztof Dembczynski, Arkadiusz Jachnik, Wojciech Kotłowski, Willem Waegeman, and Eyke Hüllermeier. Optimizing the f-measure in multi-label classification: Plug-in rule approach versus structured loss minimization. In *International conference on machine learning*, pages 1130–1138. PMLR, 2013.
- [12] Krzysztof Dembczynski, Wojciech Kotłowski, Oluwasanmi Koyejo, and Nagarajan Natarajan. Consistency analysis for binary classification revisited. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 961–969. PMLR, 2017.

- [13] Krzysztof Dembczyński, Wojciech Kotłowski, Oluwasanmi Koyejo, and Nagarajan Natarajan. Consistency analysis for binary classification revisited. In *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, volume 70 of *Proceedings of Machine Learning Research*, pages 961–969. PMLR, 2017.
- [14] Jia Deng, Sanjeev Satheesh, Alexander C. Berg, and Fei-Fei Li. Fast and balanced: Efficient label tree learning for large scale object recognition. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain.*, pages 567–575, 2011.
- [15] Himanshu Jain, Yashoteja Prabhu, and Manik Varma. Extreme multi-label loss functions for recommendation, tagging, ranking and other missing label applications. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 935–944. Association for Computing Machinery, 2016. ISBN 9781450342322.
- [16] Kalina Jasinska, Krzysztof Dembczynski, Róbert Busa-Fekete, Karlson Pfannschmidt, Timo Klerx, and Eyke Hüllermeier. Extreme F-measure maximization using sparse probability estimates. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1435–1444. JMLR.org, 2016.
- [17] Kalina Jasinska-Kobus, Marek Wydmuch, Krzysztof Dembczyński, Mikhail Kuznetsov, and Róbert Busa-Fekete. Probabilistic label trees for extreme multi-label classification. *CoRR*, abs/2009.11218, 2020.
- [18] Ting Jiang, Deqing Wang, Leilei Sun, Huayi Yang, Zhengyang Zhao, and Fuzhen Zhuang. Lightxml: Transformer with dynamic negative sampling for high-performance extreme multi-label text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7987–7994, 2021.
- [19] Sujay Khandagale, Han Xiao, and Rohit Babbar. Bonsai: diverse and shallow trees for extreme multi-label classification. *Machine Learning*, 109(11):2099–2119, 2020.
- [20] Siddhant Kharbanda, Atmadeep Banerjee, Erik Schultheis, and Rohit Babbar. Cascadexml: Rethinking transformers for end-to-end multi-resolution training in extreme multi-label classification. *Advances in Neural Information Processing Systems*, 35:2074–2087, 2022.
- [21] Siddhant Kharbanda, Devaansh Gupta, Erik Schultheis, Atmadeep Banerjee, Vikas Verma, and Rohit Babbar. Gandalf: Data augmentation is all you need for extreme classification. 2022.
- [22] Wojciech Kotłowski and Krzysztof Dembczyński. Surrogate regret bounds for generalized classification performance metrics. *Machine Learning*, 10:549–572, 2017.
- [23] Oluwasanmi Koyejo, Nagarajan Natarajan, Pradeep Ravikumar, and Inderjit S. Dhillon. Consistent multilabel classification. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3321–3329, 2015.
- [24] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.
- [25] David D. Lewis. Evaluating and optimizing autonomous text classification systems. In *SIGIR'95, Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. Seattle, Washington, USA, July 9-13, 1995 (Special Issue of the SIGIR Forum)*, pages 246–254. ACM Press, 1995.
- [26] Yu-Jen Lin and Chih-Jen Lin. On the thresholding strategy for infrequent labels in multi-label classification. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, page 1441–1450, New York, NY, USA, 2023. Association for Computing Machinery.

- [27] Julian McAuley, Rahul Pandey, and Jure Leskovec. Inferring networks of substitutable and complementary products. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 785–794, 2015.
- [28] Tharun Kumar Reddy Medini, Qixuan Huang, Yiqiu Wang, Vijai Mohan, and Anshumali Shrivastava. Extreme classification in log memory using count-min sketch: A case study of amazon search with 50m products. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’ Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 13265–13275. Curran Associates, Inc., 2019.
- [29] Anshul Mittal, Naveen Sachdeva, Sheshansh Agrawal, Sumeet Agarwal, Purushottam Kar, and Manik Varma. Eclare: Extreme classification with label graph correlations. In *Proceedings of the Web Conference 2021*, pages 3721–3732, 2021.
- [30] Harikrishna Narasimhan, Harish Ramaswamy, Aadirupa Saha, and Shivani Agarwal. Consistent multiclass algorithms for complex performance measures. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2398–2407, Lille, France, 07 2015. PMLR.
- [31] Harikrishna Narasimhan, Harish G. Ramaswamy, Shiv Kumar Tavker, Drona Khurana, Praneeth Netrapalli, and Shivani Agarwal. Consistent multiclass algorithms for complex metrics and constraints, 2022. URL <https://arxiv.org/abs/2210.09695>.
- [32] Nagarajan Natarajan, Oluwasanmi Koyejo, Pradeep Ravikumar, and Inderjit Dhillon. Optimal classification with multivariate losses. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1530–1538, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [33] Yashoteja Prabhu and Manik Varma. Fastxml: a fast, accurate and stable tree-classifier for extreme multi-label learning. In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’14, New York, NY, USA - August 24 - 27, 2014*, pages 263–272. ACM, 2014.
- [34] Yashoteja Prabhu, Anil Kag, Shilpa Gopinath, Kunal Dahiya, Shrutendra Harsola, Rahul Agrawal, and Manik Varma. Extreme multi-label learning with label features for warm-start tagging, ranking & recommendation. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 441–449, 2018.
- [35] Yashoteja Prabhu, Anil Kag, Shrutendra Harsola, Rahul Agrawal, and Manik Varma. Parabel: Partitioned label trees for extreme classification with application to dynamic search advertising. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, pages 993–1002. ACM, 2018.
- [36] Mohammadreza Qaraei, Erik Schultheis, Priyanshu Gupta, and Rohit Babbar. Convex surrogates for unbiased loss functions in extreme classification with missing labels. In *Proceedings of The Web Conference 2021, WWW ’21, New York, NY, USA, 2021*. Association for Computing Machinery. doi:[10.1145/3442381.3450139](https://doi.org/10.1145/3442381.3450139). URL <https://doi.org/10.1145/3442381.3450139>.
- [37] Mark Reid and Robert Williamson. Convexity of proper composite binary losses. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 637–644, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <https://proceedings.mlr.press/v9/reid10a.html>.
- [38] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: a modern approach*. Pearson, 3 edition, 2009.
- [39] Deepak Saini, Arnav Kumar Jain, Kushal Dave, Jian Jiao, Amit Singh, Ruofei Zhang, and Manik Varma. Galaxc: Graph neural networks with labelwise attention for extreme classification. In *Proceedings of the Web Conference 2021*, pages 3733–3744, 2021.

- [40] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. Unbiased recommender learning from missing-not-at-random implicit feedback. In *Proceedings of the 13th International Conference on Web Search and Data Mining, WSDM '20*, page 501–509, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450368223. doi:[10.1145/3336191.3371783](https://doi.org/10.1145/3336191.3371783). URL <https://doi.org/10.1145/3336191.3371783>.
- [41] Erik Schultheis, Marek Wydmuch, Rohit Babbar, and Krzysztof Dembczynski. On missing labels, long-tails and propensities in extreme multi-label classification. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1547–1557, 2022.
- [42] Tong Wei and Yu-Feng Li. Does tail label help for large-scale multi-label learning? *IEEE Transactions on Neural Networks and Learning Systems*, 31(7):2315–2324, 2020.
- [43] Tong Wei, Wei-Wei Tu, Yu-Feng Li, and Guo-Ping Yang. Towards robust prediction on tail labels. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1812–1820, 2021.
- [44] Jason Weston, Ameet Makadia, and Hector Yee. Label partitioning for sublinear ranking. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, pages 181–189. JMLR.org, 2013.
- [45] Marek Wydmuch, Kalina Jasinska, Mikhail Kuznetsov, Róbert Busa-Fekete, and Krzysztof Dembczynski. A no-regret generalization of hierarchical softmax to extreme multi-label classification. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 6355–6366. Curran Associates, Inc., 2018.
- [46] Marek Wydmuch, Kalina Jasinska-Kobus, Rohit Babbar, and Krzysztof Dembczynski. Propensity-scored probabilistic label trees. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 2252–2256, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380379.
- [47] Yiming Yang. A study of thresholding strategies for text categorization. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '01*, page 137–145, New York, NY, USA, 2001. Association for Computing Machinery. ISBN 1581133316.
- [48] Hui Ye, Zhiyu Chen, Da-Han Wang, and Brian Davison. Pretrained generalized autoregressive model with adaptive probabilistic label clusters for extreme multi-label text classification. In *International Conference on Machine Learning*, pages 10809–10819. PMLR, 2020.
- [49] Nan Ye, Kian Ming Adam Chai, Wee Sun Lee, and Hai Leong Chieu. Optimizing F-measure: A tale of two approaches. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012.
- [50] Ronghui You, Zihan Zhang, Ziyi Wang, Suyang Dai, Hiroshi Mamitsuka, and Shanfeng Zhu. Attentionxml: Label tree-based attention-aware deep model for high-performance extreme multi-label text classification. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5812–5822. Curran Associates, Inc., 2019.
- [51] Jiong Zhang, Wei-Cheng Chang, Hsiang-Fu Yu, and Inderjit Dhillon. Fast multi-resolution transformer fine-tuning for extreme multi-label text classification. *Advances in Neural Information Processing Systems*, 34:7267–7280, 2021.
- [52] Ruohong Zhang, Yau-Shian Wang, Yiming Yang, Donghan Yu, Tom Vu, and Likun Lei. Long-tailed extreme multi-label text classification with generated pseudo label descriptions. *arXiv preprint arXiv:2204.00958*, 2022.
- [53] Jingwei Zhuo, Ziru Xu, Wei Dai, Han Zhu, Han Li, Jian Xu, and Kun Gai. Learning optimal tree models under beam search. In *Proceedings of the 37th International Conference on Machine Learning*, Vienna, Austria, 2020. PMLR.