

Harnessing the Power of Beta Scoring in Deep Active Learning for Multi-Label Text Classification

Wei Tan¹, Ngoc Dang Nguyen¹, Lan Du^{1*}, Wray Buntine²

¹Faculty of IT, Monash University

²VinUniversity

wei.tan2@monash.edu, dan.nguyen2@monash.edu, lan.du@monash.edu, wray.buntine@vinuni.edu.vn

Abstract

Within the scope of natural language processing, the domain of multi-label text classification is uniquely challenging due to its expansive and uneven label distribution. The complexity deepens due to the demand for an extensive set of annotated data for training an advanced deep learning model, especially in specialized fields where the labeling task can be labor-intensive and often requires domain-specific knowledge. Addressing these challenges, our study introduces a novel deep active learning strategy, capitalizing on the Beta family of proper scoring rules within the Expected Loss Reduction framework. It computes the expected increase in scores using the Beta Scoring Rules, which are then transformed into sample vector representations. These vector representations guide the diverse selection of informative samples, directly linking this process to the model’s expected proper score. Comprehensive evaluations across both synthetic and real datasets reveal our method’s capability to often outperform established acquisition techniques in multi-label text classification, presenting encouraging outcomes across various architectural and dataset scenarios.

Introduction

Text classification is an essential task in natural language processing (NLP) with various applications, including sentiment analysis and topic classification (Allaouzi and Ben Ahmed 2019; Kobmann, Wilhelm, and Fink 2021). While deep learning has significantly improved multi-class text classification, yet its sub-task, multi-label text classification (MLTC), presents unique challenges. In MLTC, samples can be assigned multiple labels, navigating a vast, sparse label space with unresolved complexities (Liu et al. 2017; Wei and Li 2020). For every multi-label instance, the determination of appropriate labels adds to the intricacies and costs of annotation, making it even more resource-intensive than single-label learning. The increased demand for data becomes especially evident in specialized sectors, like the medical field, where domain-specific expertise and privacy considerations are crucial (Bustos et al. 2020). Recently, active learning (AL) has been explored as a promising solution to MLTC’s data scarcity issue, focusing on the

annotation of the most informative instances. However, a comprehensive AL framework for MLTC, which robustly tackles imbalanced label distributions, has yet to be fully realized (Cherman et al. 2019; Wu et al. 2020). In spite of these challenges, MLTC remains a critical task in NLP with diverse practical applications.

AL can alleviate the complexities MLTC, primarily due to its capacity to identify high-information samples, which subsequently optimizes human annotation and enhances model performance (Reyes, Morell, and Ventura 2018; Wu et al. 2020). Binary Relevance (BR) plays a key role in MLTC for its ability to segment multi-label tasks into individual binary classifications, facilitating the inclusion of advanced classifiers like deep neural networks and support vector machines (SVMs) (Zhang et al. 2018). Inspired by BR’s framework, earlier work in multi-label active learning (MLAL) adopted standard SVMs to establish binary classifiers tailored to individual labels (Yang et al. 2009; Li and Guo 2013; Cherman et al. 2019). Contrary to the notion that BR overlooks label correlation, when paired with deep neural networks, label correlations are implicitly handled, addressing a critical aspect of multi-label classification (Su et al. 2021). However, this focus on label correlation should not overshadow another crucial aspect of MLTC: data imbalance. The skewed distribution of labels in many real-world datasets affects instance selection strategies, potentially biasing models towards frequent labels (Wu, Lyu, and Ghanem 2016). Therefore, recent studies, emphasize the need to explore the inherent properties of MLTC, especially label imbalance, as a pathway to enhance generalization performance, making it imperative to address this issue precisely (Zhang et al. 2018).

In this paper, we propose a hybrid approach, named **Beta Scoring Rules for Deep Active Learning (BESRA¹)** that incorporates the Beta family of proper scoring rules within the Expected Loss Reduction (ELR) framework, applied to MLTC. Unlike expected loss used in other MLAL methods, BESRA pinpoints informative samples involving both uncertainty and diversity, drawing on expected score changes within a predictive model governed by the Beta scoring function. Through extensive experimentation on diverse ar-

*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Our implementation of BESRA can be downloaded from <https://github.com/davidtw999/BESRA>.

chitectures (e.g., TextCNN, TextRNN, BERT) and datasets (e.g., Eurler, RCV1, Bibtext, Delicious, Yahoo Health and TMC2007), we reveal that BESRA, even in the presence of varying label imbalances, consistently delivers robust performance, outperforming other AL methods. Our BESRA is adept handling of label imbalances sets a promising direction for future MLAL research and demonstrating robust performance in diverse data scenarios.

Related Works

AL is an iterative method that aims to reduce the cost and effort required for training high-quality prediction models by selecting the most informative instances from a pool of unlabeled data and querying an oracle for their labels (Settles 2009; Tan, Du, and Buntine 2021). From the literature, AL studies typically focus on multi-class classification (MCC) problems with two common strategies, including uncertainty-based and diversity-based sampling (Ren et al. 2021). Uncertainty-based methods choose the instance that is most uncertain to label for the current trained classification model and quantify uncertainty using entropy or disagreement from ensembles (Perona, Holub, and Burl 2008). Diversity-based methods acquire informative and diverse samples using sample representations generated by pre-trained models (Ash et al. 2020). However, relying solely on either uncertainty or diversity may not be adequate to select informative samples for annotation, as demonstrated by previous studies (Ren et al. 2021; Tan, Du, and Buntine 2023). To overcome this challenge, hybrid methods balance both uncertainty and diversity to acquire the most informative samples. These methods often generate sample representations from the loss that affects the model’s performance and includes predictive uncertainty (Ash et al. 2020; Tan, Du, and Buntine 2023). Hybrid methods have shown promise in enhancing AL for MCC by combining the benefits of uncertainty and diversity.

AL for multi-label classification (MLC) is an area of ongoing research, and current methods often involve transforming the multi-label problem into one or more single-label problems, allowing for the application of traditional single-label learning algorithms (Cherman et al. 2019; Reyes, Morell, and Ventura 2018). In MLC, Binary Relevance is a widely used method that decomposes the problem into a set of binary classification problems, enabling instance selection decisions through the independent exploitation of binary classifiers (Shi and Yu 2019). For example, Yang et al. (2009) presents a strategy called maximum loss reduction with maximal confidence (MMC), which uses multi-class logistic regression to predict the number of labels for an unlabeled instance. Subsequently, it computes the expected loss associated with the most confident result of label prediction from support vector machine classifiers across all labels. To improve the MMC framework, Hung and Lin (2011) proposed a soft Hamming loss reduction (SHLR) by extending MMC for MLAL by introducing a major and an auxiliary learner. The two learners are designed to evolve together during AL, and queries are made based on their disagreement. The framework provides flexibility in the choice of learners and query criteria and also includes MMC, Bin-

Min, and the random query algorithm as special cases (Chen, Sun, and Zhao 2018; Reyes, Morell, and Ventura 2018). Recently, ADAPTIVE (Li and Guo 2013) and CVIRS (Reyes, Morell, and Ventura 2018) consider label dependencies in their uncertainty sampling. The CSRPE (Yang et al. 2018) method introduces cost-sensitive coding to manage varying error costs. Additionally, the GP-B2M (Shi, Yu, and Yu 2021) method deals with label sparsity and seeks to reveal label relationships using a Bayesian Bernoulli mixture of label clusters. However, all these studies do not explicitly tackle the issues of label imbalance, which are frequently observed in MLTC, as indicated in Table 2. Our AL method offers a fresh approach to addressing the imbalance challenges in MLTC, displaying consistent performance across both synthetic and real datasets.

Beta Distribution-based Scoring Rules

ELR (Roy and McCallum 2001) quantifies the generalisation error that is expected to reduce when a new sample is added to the labeled pool L . Leveraging the insights from previous research, it uses an unlabeled pool U to represent potential test data, thereby quantifying the expected generalization error (Settles 2009). Mean Objective Cost of Uncertainty (MOCU) (Zhao et al. 2021) generalizes this approach and adds theoretical rigour. The Bayesian Estimate of Mean Proper Scores (BEMPS) improves these formulations by replacing error with proper scoring rules (Gneiting and Raftery 2007), leading to an AL strategy proven to converge (Tan, Du, and Buntine 2023). Scoring rules measure the quality of predictive distributions by rewarding calibrated predictive distributions. Different scoring rules, including Brier score, Logarithmic score, and Beta family (Gneiting and Raftery 2007), focus the quality measurement on different aspects of the model prediction. Unlike the Brier and Logarithmic scores, which are tailored for MCC, the Beta distribution inherently models variables limited between 0 and 1, making it more apt for BR in multi-label problems.

The optimization goal of our multi-label active learner is to label data that leads to the largest change in expected scores measured by the Beta family on the prediction of all binary classifiers. Let $P_T(x)$ be the (unknown) input distribution and $P_T(y|x)$ be the (unknown) conditional distribution on label sets $y \in \mathcal{Y}$, wherein \mathcal{Y} is the set of all label sets. This distribution $P_T(y|x)$ is used by the so-called Bayes classifier, an ideal because it is unknown. We denote the predicted probability of $y|x$ given the current running model θ as $f_L(x, \theta) = P(y | \theta, x)$, and the predicted probability of $y|x$ under a Bayes optimal classifier as $f_L(x) = P(y | L, x) = \int_{\theta} P(y | \theta, x) P(\theta|L) d\theta$.

A “quality” score for the active learner is formulated as the expected difference Q_L between the score of the true probability estimates and a Bayes classifier after training on L (Tan, Du, and Buntine 2023), following for instance Hung and Lin (2011, Sect. 4.1). However, the true conditional distribution $P_T(y | x)$ is unknown, so we make do with the current Bayesian estimate $f_L(x)$ and average the score dif-

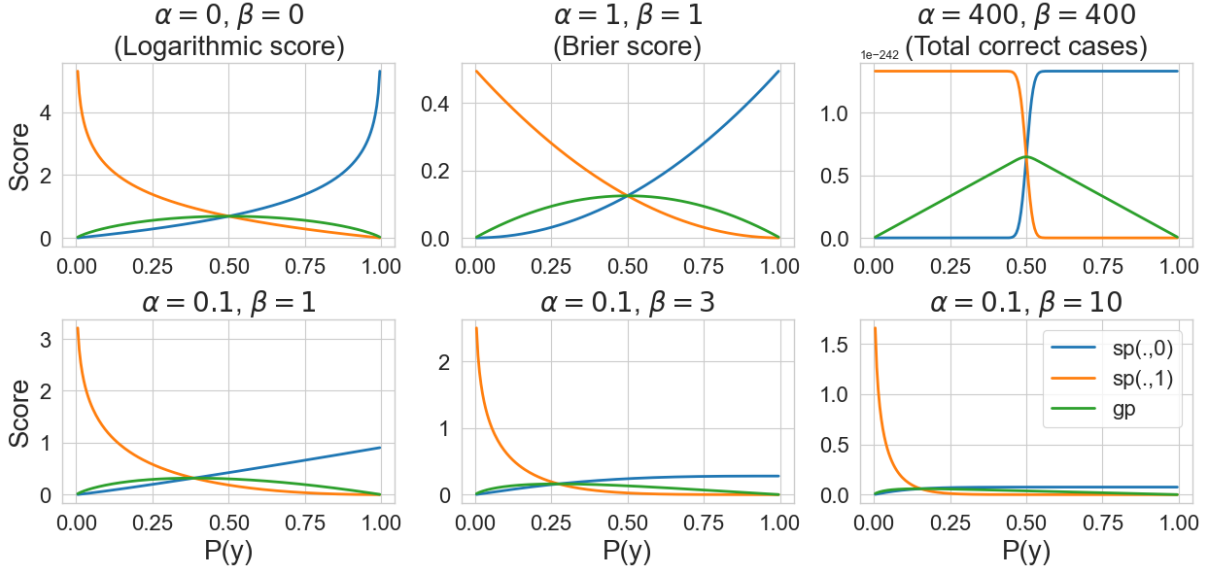


Figure 1: The graph depicts the Expected Score gp (green) and Scoring Functions from Eq (5) for the Beta family in blue ($sp(\cdot, 0)$ when $y = 0$ and orange ($sp(\cdot, 1)$ when $y = 1$). It covers six scenarios: three specific to Brier Score, Logarithmic Score, and total error approximations, and three emphasizing asymmetry with varied Beta values.

ferences across different models θ :

$$Q_L = \mathbb{E}_{P_T(x)} \mathbb{E}_{P(\theta|L)} \left[\mathbb{E}_{P(y|\theta, x)} [S(f_L(x, \theta), y) - S(f_L(x), y)] \right], \quad (1)$$

$$\begin{aligned} \Delta Q(x|L) &= Q_L - \mathbb{E}_{P(y|L, x)} [Q_{L+\{x, y\}}] \\ &= \mathbb{E}_{P(y|L, x)} \left[\mathbb{E}_{P(x')P(y'|L, (x, y), x')} [S(P(\cdot | L, (x, y), x'), y') - S(P(\cdot | L, x'), y')] \right] \end{aligned} \quad (2)$$

Here S is a scoring function which scores a predictive probability distribution given the true label y . The active learner should evaluate Eq (1)- with L given by $L + \{x, y\}$ for each possible unlabeled data $x \in U$ to find the optimal query x^* . Since labels y of any new sample of data x are unknown, the $Q_{L+\{x, y\}}$ should be averaged over labels again using the Bayesian estimate. The formula to maximise for unlabeled sample x then becomes the difference, $\Delta Q(x|L)$, as in Eq (2), which can be further simplified (Tan, Du, and Buntine 2023, Equation 6).

In the multi-label setting where there are K classes, and each class can be made a binary classification problem with $y_k \in \{0, 1\}$, the scoring function using BR is $S_{BR}(f_L(x), y) = \sum_{k=1}^K S^k(f_L(x), y_k)$ where $S^k(\cdot, \cdot)$ is the score for each binary class (Dawid and Musio 2014). Because of linearity, this works neatly with Eq (2).

For implementation, we follow the algorithm of BEMPS in (Tan, Du, and Buntine 2023) using our S_{BR} score. For this, $\mathbb{E}_{P_T(x)}$ is approximated as the average over a smaller estimation pool (sampled from U). Moreover, $\mathbb{E}_{P(\theta|L)}$ is approximated with an ensemble $\frac{1}{|\Theta|} \sum_{\theta \in \Theta}$ as is $\mathbb{E}_{P(y|L, x)}$ as $\frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{P(y|\theta, x)}$. The Bayesian update required for $P(y' | L + \{x, y\}, x')$ is done by reweighting the ensemble:

ble:

$$P(y' | L, (x, y), x') = \sum_{\theta \in \Theta^E} P(y' | \theta, x') P(\theta | L, (x, y)) \quad (3)$$

$$P(\theta | L, (x, y)) \approx \frac{P(\theta | L) P(y|\theta, x)}{\sum_{\theta \in \Theta^E} P(\theta | L) P(y|\theta, x)} \quad (4)$$

Now the sum of strictly proper scores leads to a strictly proper score (Dawid and Musio 2014). So if $S^k(\cdot, \cdot)$ is a strictly proper score, then it follows that our score $S_{BR}(\cdot, \cdot)$ will be strictly proper, and therefore all the theoretical benefits ensue for our MLAL framework, such as provable convergence. We propose using the Beta family introduced by (Buja, Stuetzle, and Shen 2005), a two-parameter family of proper scoring rules $\alpha, \beta > -1$ particularly useful for imbalanced classes or unequal costs.

$$S_{\alpha, \beta}^k(p, y_k) = \begin{cases} - \int_0^p c^\alpha (1-c)^{\beta-1} dc \\ = - \frac{\Gamma(\alpha+1)\Gamma(\beta)}{\Gamma(\alpha+\beta+1)} I_p(\alpha+1, \beta) & , y_k = 0 \\ - \int_p^1 c^{\alpha-1} (1-c)^\beta dc \\ = - \frac{\Gamma(\alpha)\Gamma(\beta+1)}{\Gamma(\alpha+\beta+1)} I_{1-p}(\beta+1, \alpha) & , y_k = 1 \end{cases} \quad (5)$$

Here, $I_x(a, b)$ denotes the Incomplete Beta Function for $a, b > 0$, and the probability vector $P(y) = (p, 1-p)$ given as p . If one of $a, b \leq 0$, an alternative evaluation is required.

The Beta family of scoring rules generalise the Logarithmic score and the Brier score used in BEMPS as follows (Merkle and Steyvers 2013): When $\alpha = \beta = 0$, Eq (5) reproduces the Logarithmic score and when $\alpha = \beta = 1$, it yields the Brier score. From this perspective, BESRA can be regarded as the generalization of BEMPS. When $\alpha = \beta \rightarrow \infty$, the scoring function becomes a step function, which approximates misclassification score. However, unlike the Brier and Logarithmic scores, which equally penalize both the low probability region (i.e., False Negatives (FN)) and the high probability region (i.e., False Positives (FP)),

Algorithm 1: Beta Scoring Rules For Deep Active Learning

Require: initial unlabeled pool U , initial labeled pool L , model ensemble $\Theta^E = \{\theta_1, \dots, \theta_E\}$ built from L by retraining, pre-computed values for $P(y_k | \theta, x)$ and $P(y_k | L, x)$ for all relevant y and x , acquire batch size B , acquire batch set \mathcal{A} , estimation pool X , number of acquisition iterations N , number of classes K .

```

1: Initialize:  $i = 0, L_0 \leftarrow L, U_0 \leftarrow U$ 
2: while  $i < N$  do
3:    $vec_{x,x'} = 0$ 
4:   for  $x \in U_i, x' \in X$  do
5:     for  $k \in K, y_k \in \{0, 1\}$  do
6:       Compute  $P(\theta|L, (x, y_k))$  with Eq (4) for each  $y_k$ 
7:       Compute  $p(y'_k|L + \{x, y'_k\}, x')$  with Eq (3) for each  $y'_k$ 
8:     end for
9:      $vec_{x,x'} += \Delta Q(x|L, x')$ , where  $\Delta Q(x|L, x')$  is computed by Eq (2) with the beta scoring function defined in Eq (5).
10:   end for
11:    $\mathcal{A} = \emptyset$ 
12:    $centroids = k$ -Means centers ( $vec_{x \in U_i}, B$ )
13:   for  $c \in centroids$  do
14:      $\mathcal{A} \cup = \{ \operatorname{argmin}_{x \in U_i} \|c - vec_x\| \}$ 
15:   end for
16:    $L_{i+1} \leftarrow L_i \cup \mathcal{A}$ 
17:    $U_{i+1} \leftarrow U_i \setminus \mathcal{A}$ 
18: end while

```

(FP)) (as seen in Figure 1), the Beta family allows for a more tailored approach. This adaptability is especially beneficial in scenarios with specific demands, such as imbalanced multi-label classifications (see Table 2). Adjusting the α and β values, the Beta Scoring Rules can be adjusted to differentially penalize false positives and negatives. For example, by fixing $\alpha = 0.1$ and tweaking β , we can incline the penalties more towards FN predictions. Conversely, selecting the right β and α is critical and varies based on the problem, as highlighted in (Merkle and Steyvers 2013; Gneiting and Raftery 2007). Our research explores how varying these parameter values for BESRA impacts MLAL’s performance in the Beta Parameters section below.

After incorporating the obtained samples into the labeled set, we proceed to retrain the multi-label classifier (e.g., TextCNN, TextRNN and BERT) using randomly initialized parameters. This retraining occurs after each iteration of acquisition and is carried out using the binary cross-entropy loss (BCE), a commonly employed loss function in multi-label learning (Rezaei-Dastjerdehei, Mijani, and Fatemizadeh 2020). Specifically, given a specific data point x , we compute the probability of x being affiliated with class k , $\forall k \in \{1, 2, \dots, K\}$ as $p(y_k = 1|x) = \frac{1}{1 + \exp(f_k(x))}$, where f_k is the Binary Relevance associated with class k . BESRA’s Algorithm 1 uses Bayes theorem to estimate $P(\theta|L, (x, y_k))$ from $P(\theta|L)$, leveraging a predefined estimation pool X , a randomly selected subset from unlabeled pool, for estimating expected values $\mathbb{E}_{P(x')}[\cdot]$ in Eq (2). Each query is now represented as a vector of expected scored changes, i.e., $vec_{x,x'}$. These vectors are then organized us-

Dataset	#Document Train/Test	#Vocabulary /#Label	MeanIR Train/Test
RCV1-10	1,200/600	23,759/10	10/50
RCV1-50	1,200/600	23,117/10	50/50
RCV1-200	1,200/600	21,074/10	200/50
RCV1-400	1,200/600	34,464/10	400/50

Table 1: Five datasets (including the test set) are created using the MeanIR values that represent the imbalance level of the multi-label datasets. Higher values suggest increased sparsity and imbalance.

Dataset	#Document Train/Test	#Vocabulary /#Label	MeanIR Train/Test
Bibtex	5,916/1,479	1,836/159	12/15
TMC2007	22,876/5,720	497/25	22/23
Delicious	12,872/3,219	501/982	72/73
RCV1	20,833/7,965	119,475/102	285/152
Eurlex	15,470/3,868	5,000/201	426/164
Yahoo	7,364/1,841	30,587/32	656/162

Table 2: Six benchmark datasets with their corresponding imbalance level MeanIR statistics.

ing k-Means clustering, and from each cluster, only the data points closest to the centroid are selected for precision and representation.

Experiments

We evaluated BESRA’s impact on synthetic and real-world datasets, particularly studying the influence of Beta Scoring Rules on different imbalance levels in synthetic data. Further, we benchmarked BESRA against state-of-the-art (SOTA) AL methods, affirming BESRA’s efficacy in MLTC.

Datasets

Synthetic Datasets. We created five synthetic datasets, each consisting of ten labels and associated with a unique mean imbalance ratio (MeanIR), a measure of the average level of imbalance in the multi-label dataset (Charte et al. 2015), as outlined in Table 1. Four of these datasets, aligned with MeanIR values of (10, 50, 200, 400), served training purposes. Additionally, a consistent test set with a MeanIR of 50 was established for standardized evaluations (Wang et al. 2023). To make the generation process for the synthetic dataset controllable, we reduced the label space to ten (i.e., $K = 10$) by selecting the large dataset RCV1, and sampled instances based on the MeanIR metric provided by Charte et al. (2015). In contrast, label cardinality measures the average labels per instance and label density evaluates the instance-label proportion. However, neither offers a holistic perspective on the overall label imbalance. As our main objective is to develop a general AL framework for MLTC that can handle the label imbalance issue, These synthetic datasets are crucial for evaluating the performance of AL methods with the Beta Proper Scores under varying




Figure 2: The average Micro F1-score of AL models with acquisition size 100 on BERT, which were run with 5 different random seeds on various synthetic datasets.

levels of label imbalance. The experiments used a greedy search to identify the optimal parameters for the Beta scoring function ($\alpha = 0.1, \beta = 3$) compared to the Brier score ($\alpha = 1, \beta = 1$) and the Logarithmic score ($\alpha = 0, \beta = 0$). By systematically manipulating the Alpha and Beta values, the study evaluates BESRA’s robustness and quantifies its influence across different magnitudes of dataset imbalance.

Real Datasets. We have selected six diverse benchmark multi-label text datasets (MLTDs): Eulex (Loza Mencía and Fürnkranz 2010); RCV1 (Lewis et al. 2004), comprising Reuters newswire stories; Bibtex (Katakis, Tsoumakas, and Vlahavas 2008), which contains bibliographic records in BibTeX format; Delicious (Tsoumakas, Katakis, and Vlahavas 2008), capturing user bookmarking data; Yahoo (Ueda and Saito 2002), detailing health-related user discussions; and TMC2007 (Srivastava and Zane-Ulman 2005), a collection of texts from varied domains. These datasets’ label spaces vary significantly, ranging from 25 to 982 labels. We used the MeanIR statistics to measure the imbalance level in each. For example, Bibtex possesses the lowest MeanIR of the training set, meaning it’s less imbalanced than the others. In contrast, Yahoo has the highest MeanIR, indicating it’s the most imbalanced dataset among those we considered. To ensure data integrity, we have kept both train and test sets in their original sizes. Table 2 offers a summarized overview of these datasets’ key characteristics for in-depth analysis and research.

Model Architectures

In our experiments, we use TextCNN (Kim 2014), TextRNN (Cai et al. 2018), and BERT (Devlin et al. 2019) as the backbone classifier. After each AL iteration, we fine-tune these models using random re-initialization (Frankle and Carbin 2019) for enhanced efficacy—a method found to be more effective than incremental fine-tuning with new samples (Gal, Islam, and Ghahramani 2017). The fine-tuning process, executed on RTX3090 GPUs, sets a maximum sequence length of 256, runs for up to 80 epochs, and uses the AdamW optimizer with a $1e-5$ learning rate (Nguyen et al. 2022, 2023a,b). Within the AL setup, the initial training and validation sets are composed of 100 and 1000 samples from the

training set. To ensure robustness, each AL method iterates five times on each dataset with the random seeds, sampling 100 instances in each acquisition. By incorporating a deep ensembles method (Lakshminarayanan, Pritzel, and Blundell 2017), we compute predictive distributions $P(y|L, x)$ using five ensemble models. Performance was assessed using metrics like micro-F1, macro-F1, precision, recall, precision@5, and recall@5, though the primary paper emphasizes the micro F1-score, and further metrics’ results can be found in the Appendix.

Active Learning Baselines

We compared BESRA with various AL methods, including a random baseline. For each method, we adapted its strategies to the MLTDs and using a consistent backbone classifier for all. Among the methods examined, the **MMC** method by Yang et al. (2009) selects instances by maximizing loss reduction; **ADAPTIVE** by Li and Guo (2013) chooses instances by balancing prediction uncertainty with label cardinality inconsistency; **AUDI** emphasizes label ranking and adjusted cardinality inconsistency, targeting uncertainty and diversity (Huang and Zhou 2013); **CVIRS** (Reyes, Morell, and Ventura 2018) leverages label vector inconsistency and score rankings to decide on instance labeling; **SHLR** by Hung and Lin (2011) employs an auxiliary learner to determine instance choices based on disagreement levels; **CSRPE**, as introduced by Yang et al. (2018), extends one-versus-one coding to a cost-sensitive approach, using code-bits to inform instance weights and improve performance; finally, **GPB2M** (Shi, Yu, and Yu 2021) combines a Gaussian Process with a Bayesian Bernoulli mixture to capture label correlations for instance acquisition.

Results

Synthetic Datasets. Figure 2 illustrates the micro F1-scores for three methods. Both the Brier ($\alpha = 1, \beta = 1$) and Logarithmic score ($\alpha = 0, \beta = 0$) equally penalize FNs and FPs. In contrast, the Beta score method with parameters ($\alpha = 0.1, \beta = 3$) emphasizes a more substantial penalty for FN predictions, leading to an enhanced AL performance. Average results from five AL experiments, with a 95% bootstrapped error band, underscore the Beta score’s

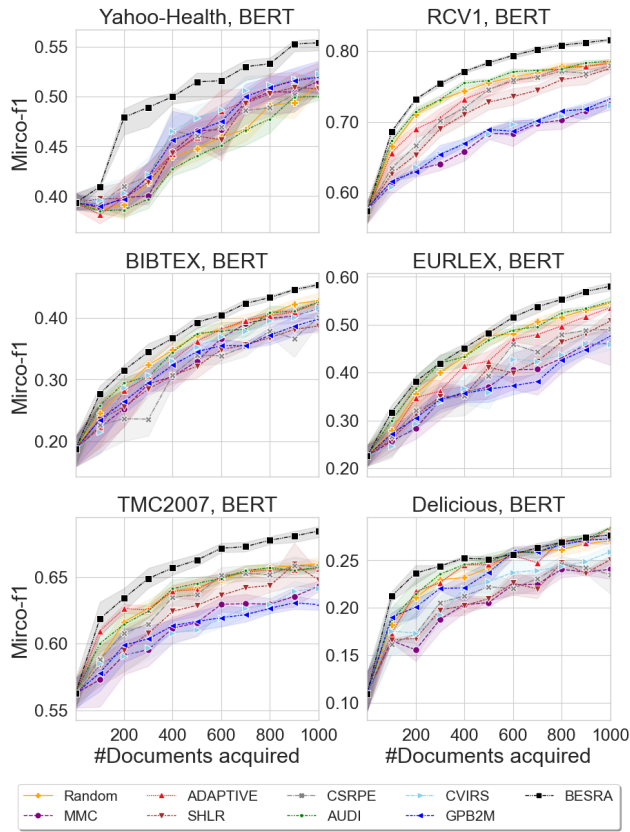


Figure 3: The average Micro F1-score of AL models with acquisition size 100 on BERT, which were run with 5 different random seeds on various datasets.

superior performance, especially on high imbalance datasets like RCV1-200 and RCV1-400. Although Brier shows a slight edge over the Logarithmic score in the RCV1-50 scenario, their differences diminish in other scenarios. Specifically, by optimizing for FN predictions, the Beta Proper Scores ($\alpha = 0.1$ and $\beta = 3$) have showcased their efficacy in handling varied imbalance distributions, especially as the imbalance ratio increases. However, these parameter values can be adjusted for specific applications based on desired scoring criteria, as discussed by (Buja, Stuetzle, and Shen 2005; Merkle and Steyvers 2013). Further insights on this topic, including an ablation study and detailed discussion, are available in subsequent sections, with a more exhaustive analysis presented in the Appendix.

Real dataset. Figure 3 shows a performance comparison between BESRA and the baseline methods across six MLTDs, as listed in Table 2. We use the Beta score with $\alpha = 0.1$ and $\beta = 3$ due to its superior performance compared to the Brier score and Logarithmic score. Remarkably, BESRA consistently outperforms other AL methods in all acquisition scenarios, regardless of the text domains. Among the evaluated methods, ADAPTIVE has the highest overall performance based on the micro F1-score. It consistently demonstrated strong performance across all datasets, showing its effectiveness in MLTC AL. Besides from ADAP-

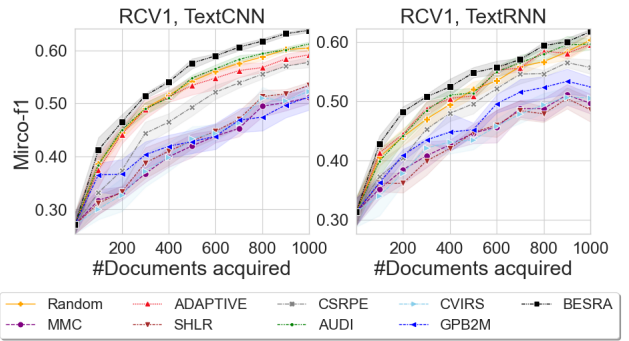


Figure 4: The average Micro F1-score of AL models with acquisition size 100 on TextCNN and TextRNN, which were run with 5 different random seeds on RCV1.

TIVE, AUDI also provided competitive performance on all datasets except for Yahoo (health), suggesting its potential in specific domains. Conversely, CVIRS and GPB2M exhibited inconsistent performance across the datasets, struggling particularly in EURLEX, RCV1, and TMC2007. CSRPE method yielded inadequate results in the Bibtex and Delicious datasets, indicating a lack of effectiveness in those specific contexts. MMC emerged as the least effective method across all datasets except for Yahoo (health). In addition, MMC and SHRL showed similar performance, as both methods focused on measuring the expected loss of the model. It is worth highlighting that SHRL showcased better results than MMC specifically in the RCV1 and TMC2007 datasets, this could be due to SHRL being the improvement of MMC by introducing the soft Hamming loss (Hung and Lin 2011) as an alternative to the mean of the expected loss used in MMC. GPB2M, the most recent work in ML AL, demonstrated superior performance in the Yahoo (health) and Delicious datasets, highlighting its efficacy in those specific contexts. However, the robustness of this AL strategy raises a significant concern, as its performance exhibits variability not only across diverse datasets but also among various model architectures. For a more comprehensive analysis, including additional metrics such as macro-F1, precision and recall, please refer to the Appendix.

Ablation Study

To assess the effectiveness and generalizability of BESRA, we conducted a comprehensive evaluation of three well-established neural network architectures commonly used in MLTC tasks: TextCNN, TextRNN, and BERT. These architectures have played a significant role in multilabel learning. Additionally, we investigated the impact of different Alpha and Beta values on the model’s performance. Our study aimed to answer two key research questions: (i) Does the performance advantage offered by BESRA generalise across diverse models and architectures? (ii) How do the Alpha and Beta values influence the scoring mechanism and subsequently affect the model’s performance across real-world datasets with the varied imbalance level? By addressing these questions, we aimed to gain insights of BESRA and its potential application in various MLTC tasks.

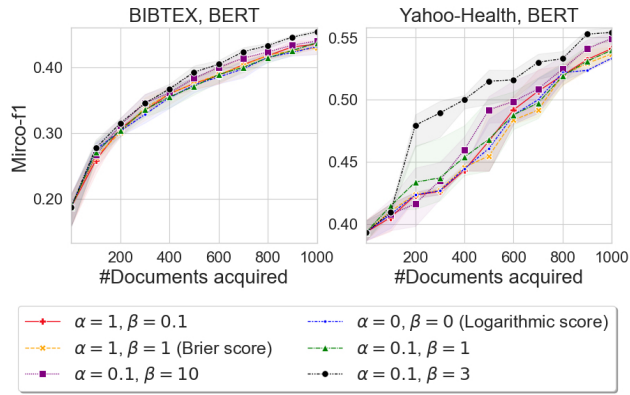


Figure 5: The average Micro F1-score of AL models with acquisition size 100 on BERT, which were run with 5 different random seeds on for Bibtex and Yahoo (health).

Model Generalizability

The evaluation of BESRA is being conducted in two other models to test its generalizability. As depicted in Figure 4, BESRA consistently outperforms other baselines across TextCNN and TextRNN tested on the RCV1 dataset. Additionally, BESRA demonstrates exceptional performance compared to these models on five additional datasets included in the Appendix. These results validate that BESRA demonstrates effectiveness across a diverse range of pre-trained language models, highlighting its applicability irrespective of the specific model architecture applied.




Figure 6: Left: Learning curves of ALs with batch size 50 on RCV1. Right: Learning curves for BESRA with batch sizes $B \in \{50, 100\}$ on RCV1. All results were run with 5 different random seeds

Batch Size

Figure 6 indicate that BESRA consistently outperforms other ALs when using a batch size of 50. Additionally, during the early stages of acquisition, BESRA performs more effectively with smaller batch sizes (50) than with larger batch sizes (100), aligning with the results reported in (Tan, Du, and Buntine 2023; Kirsch, van Amersfoort, and Gal 2019). We hypothesise that at the early acquisition stages, due to the lack of knowledge, the model trained with a limited number of samples has high uncertainty and lacks calibration, thus acquiring a large batch of samples can result in

noise. Additional experiments conducted on other MLTDs can be found in the Appendix.

Beta Parameters

In this subsection, we delve into the implications of varying Beta scores on the performance of AL, with a particular emphasis on how these scores influence the penalization behaviour of the active learner. We evaluated a range of α and β values, associated with distinct Beta scores as shown in Eq (5). Our evaluations spanned three representative datasets, ranging from relatively balanced (i.e., BIBTEX) to highly imbalanced (i.e., Yahoo). To gain a comprehensive understanding of how different α and β values impact the outcomes, we considered several scoring methods including the Brier score and Logarithmic score (which provide equal penalization), alongside four distinct scenarios of Beta score. These scenarios include (1) mild penalization on False Positives (FPs) where $\alpha = 1$ and $\beta = 0.1$, (2) light penalization on False Negatives (FNs) with $\alpha = 0.1$ and $\beta = 1$, (3) moderate penalization on FNs defined by $\alpha = 0.1$ and $\beta = 3$, and finally (4) stringent penalization on FNs denoted by $\alpha = 0.1$ and $\beta = 10$.

Insights from Figure 5 reveal that the Brier and Logarithmic scores, as well as the scenario with $\alpha = 1$ and $\beta = 0.1$, consistently underperform across various imbalance levels within MLTDs. Such results are consistent with our prior expectations, considering the inherent label imbalance challenges. A significant reason for this underperformance is the equal penalization rendered to both FN and FP outcomes by the Brier and Logarithmic scores. Additionally, the specific scenario of $\alpha = 1$ and $\beta = 0.1$ tends to disproportionately penalize FPs, thereby dampening performance. When evaluating the effect of α and β values focused on penalizing FNs, we note that while light penalization settings (i.e., $\alpha = 0.1$, $\beta = 1$) have a negligible impact on enhancing the active learner’s effectiveness, the more stringent configuration of $\alpha = 0.1$ and $\beta = 10$ offers notable improvements, especially in highly imbalanced datasets. However it does not necessarily culminate in the optimal active learner. Instead, a moderate penalization strategy with $\alpha = 0.1$ and $\beta = 3$ consistently stands out as the most effective across MLTDs.

Conclusion

We have introduced BESRA, a novel acquisition strategy for MLAL. This generalizes the recently published BEMPS using the Beta family of proper scoring rules, which allow customizable asymmetric scoring rules that effectively address the challenges such as imbalanced data associated with multi-label learning. Moreover, by our methodical construction, the use of BESRA provably converges to optimal solutions. Through empirical studies conducted on synthetic and real-world datasets, we have demonstrated the effectiveness of BESRA in acquiring highly informative samples for multi-label active learning, consistently surpassing seven existing acquisition strategies. This finding highlights the crucial role of Beta Scoring Rules and their great potential for AL with tailored acquisition strategies. Future research can further explore combinations of Alpha and Beta values for specific datasets, addressing a current limitation of BESRA.

References

- Allaouzi, I.; and Ben Ahmed, M. 2019. A Novel Approach for Multi-Label Chest X-Ray Classification of Common Thorax Diseases. *IEEE Access*, 7: 64279–64288.
- Ash, J. T.; Zhang, C.; Krishnamurthy, A.; Langford, J.; and Agarwal, A. 2020. Deep Batch Active Learning by Diverse, Uncertain Gradient Lower Bounds. In *Proc. 8th Int. Conf. Learn. Representations*.
- Buja, A.; Stuetzle, W.; and Shen, Y. 2005. Loss Functions for Binary Class Probability Estimation and Classification: Structure and Applications. Technical report, The Wharton School, University of Pennsylvania.
- Bustos, A.; Pertusa, A.; Salinas, J.-M.; and de la Iglesia-Vayá, M. 2020. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *Medical Image Analysis*, 66: 101797.
- Cai, J.; Li, J.; Li, W.; and Wang, J. 2018. Deep Learning Model Used in Text Classification. In *2018 15th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*, 123–126.
- Charte, F.; Rivera, A. J.; del Jesus, M. J.; and Herrera, F. 2015. Addressing imbalance in multilabel classification: Measures and random resampling algorithms. *Neurocomputing*, 163: 3–16. Recent Advancements in Hybrid Artificial Intelligence Systems and its Application to Real-World Problems Progress in Intelligent Systems Mining Humanistic Data.
- Chen, J.; Sun, S.; and Zhao, J. 2018. Multi-label Active Learning with Conditional Bernoulli Mixtures. In Geng, X.; and Kang, B.-H., eds., *PRICAI 2018: Trends in Artificial Intelligence*, 954–967. Cham: Springer International Publishing. ISBN 978-3-319-97304-3.
- Cherman, E. A.; Papanikolaou, Y.; Tsoumakas, G.; and Monard, M. C. 2019. Multi-label active learning: key issues and a novel query strategy. *Evolving Systems*, 10(1): 63–78.
- Dawid, A. P.; and Musio, M. 2014. Theory and applications of proper scoring rules. *Metron*, 72(2): 169–183.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Frankle, J.; and Carbin, M. 2019. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In *Proc. 7th Int. Conf. Learn. Representations*.
- Gal, Y.; Islam, R.; and Ghahramani, Z. 2017. Deep Bayesian Active Learning with Image Data. In *Proc. 34th Int. Conf. Mach. Learn.*, volume 70, 1183–1192.
- Gneiting, T.; and Raftery, A. E. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statistical Assoc.*, 102(477): 359–378.
- Huang, S.; and Zhou, Z. 2013. Active Query Driven by Uncertainty and Diversity for Incremental Multi-label Learning. In *2013 IEEE International Conference on Data Mining (ICDM)*, 1079–1084. Los Alamitos, CA, USA: IEEE Computer Society.
- Hung, C.-W.; and Lin, H.-T. 2011. Multi-label Active Learning with Auxiliary Learner. In Hsu, C.-N.; and Lee, W. S., eds., *Proceedings of the Asian Conference on Machine Learning Research*, 315–332. South Garden Hotels and Resorts, Taoyuan, Taiwan: PMLR.
- Katakis, I. M.; Tsoumakas, G.; and Vlahavas, I. P. 2008. Multilabel Text Classification for Automated Tag Suggestion. In *Proceedings of the ECML/PKDD 2008 Discovery Challenge*.
- Kim, Y. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1746–1751. Doha, Qatar: Association for Computational Linguistics.
- Kirsch, A.; van Amersfoort, J.; and Gal, Y. 2019. Batch-BALD: Efficient and Diverse Batch Acquisition for Deep Bayesian Active Learning. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Koßmann, D.; Wilhelm, T.; and Fink, G. A. 2021. Towards Tackling Multi-Label Imbalances in Remote Sensing Imagery. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 5782–5789.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 6405–6416. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.
- Lewis, D. D.; Yang, Y.; Rose, T. G.; and Li, F. 2004. RCV1: A New Benchmark Collection for Text Categorization Research. *JMLR*, 5: 361–397.
- Li, X.; and Guo, Y. 2013. Active Learning with Multi-Label SVM Classification. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, 1479–1485. AAAI Press. ISBN 9781577356332.
- Liu, J.; Chang, W.-C.; Wu, Y.; and Yang, Y. 2017. Deep Learning for Extreme Multi-Label Text Classification. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17*, 115–124. New York, NY, USA: Association for Computing Machinery. ISBN 9781450350228.
- Loza Mencía, E.; and Fürnkranz, J. 2010. Efficient Multilabel Classification Algorithms for Large-Scale Problems in the Legal Domain. In Francesconi, E.; Montemagni, S.; Peters, W.; and Tiscornia, D., eds., *Semantic Processing of Legal Texts: Where the Language of Law Meets the Law of Language*, 192–215. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-12837-0.

- Merkle, E. C.; and Steyvers, M. 2013. Choosing a strictly proper scoring rule. *Decision Analysis*, 10(4): 292–304.
- Nguyen, N. D.; Du, L.; Buntine, W.; Chen, C.; and Beare, R. 2022. Hardness-guided domain adaptation to recognise biomedical named entities under low-resource scenarios. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 4063–4071. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Nguyen, N. D.; Tan, W.; Du, L.; Buntine, W.; Beare, R.; and Chen, C. 2023a. AUC Maximization for Low-Resource Named Entity Recognition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11): 13389–13399.
- Nguyen, N. D.; Tan, W.; Du, L.; Buntine, W.; Beare, R.; and Chen, C. 2023b. Low-Resource Named Entity Recognition: Can One-vs-All AUC Maximization Help? arXiv:2311.04918.
- Perona, P.; Holub, A.; and Burl, M. C. 2008. Entropy-based active learning for object recognition. In *2008 IEEE Computer Society Conf. Computer Vision Pattern Recognition Workshops (CVPR Workshops)*, 1–8. IEEE Computer Society.
- Ren, P.; Xiao, Y.; Chang, X.; Huang, P.-Y.; Li, Z.; Gupta, B. B.; Chen, X.; and Wang, X. 2021. A Survey of Deep Active Learning. *ACM Comput. Surv.*, 54(9).
- Reyes, O.; Morell, C.; and Ventura, S. 2018. Effective active learning strategy for multi-label learning. *Neurocomputing*, 273: 494–508.
- Rezaei-Dastjerdehei, M. R.; Mijani, A.; and Fatemizadeh, E. 2020. Addressing Imbalance in Multi-Label Classification Using Weighted Cross Entropy Loss Function. In *2020 27th National and 5th International Iranian Conference on Biomedical Engineering (ICBME)*, 333–338.
- Roy, N.; and McCallum, A. 2001. Toward Optimal Active Learning through Sampling Estimation of Error Reduction. In *Proc. 18th Int. Conf. Mach. Learn.*, 441–448. ISBN 1558607781.
- Settles, B. 2009. Active Learning Literature Survey. Computer Sciences Technical Report 1648, Univ. of Wisconsin–Madison.
- Shi, W.; Yu, D.; and Yu, Q. 2021. A Gaussian Process-Bayesian Bernoulli Mixture Model for Multi-Label Active Learning. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 27542–27554. Curran Associates, Inc.
- Shi, W.; and Yu, Q. 2019. Fast Direct Search in an Optimally Compressed Continuous Target Space for Efficient Multi-Label Active Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 5769–5778. PMLR.
- Srivastava, A.; and Zane-Ulman, B. 2005. Discovering recurring anomalies in text reports regarding complex space systems. In *2005 IEEE Aerospace Conference*, 3853–3862.
- Su, Z.; Chow, J. K.; Tan, P. S.; Wu, J.; Ho, Y. K.; and Wang, Y.-H. 2021. Deep convolutional neural network-based pixel-wise landslide inventory mapping. *Landslides*, 18(4): 1421–1443.
- Tan, W.; Du, L.; and Buntine, W. 2021. Diversity Enhanced Active Learning with Strictly Proper Scoring Rules. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 10906–10918. Curran Associates, Inc.
- Tan, W.; Du, L.; and Buntine, W. 2023. Bayesian Estimate of Mean Proper Scores for Diversity-Enhanced Active Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–16.
- Tsoumakas, G.; Katakis, I. M.; and Vlahavas, I. P. 2008. Effective and Efficient Multilabel Classification in Domains with Large Number of Labels. In *Proceedings of the ECM-L/PKDD 2008 Discovery Challenge*.
- Ueda, N.; and Saito, K. 2002. Parametric Mixture Models for Multi-Labeled Text. In *NIPS’02*, 737–744. Cambridge, MA, USA: MIT Press.
- Wang, R.; Jia, X.; Wang, Q.; Wu, Y.; and Meng, D. 2023. Imbalanced Semi-supervised Learning with Bias Adaptive Classifier. arXiv:2207.13856.
- Wei, T.; and Li, Y.-F. 2020. Does Tail Label Help for Large-Scale Multi-Label Learning? *IEEE Transactions on Neural Networks and Learning Systems*, 31(7): 2315–2324.
- Wu, B.; Lyu, S.; and Ghanem, B. 2016. Constrained Submodular Minimization for Missing Labels and Class Imbalance in Multi-label Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1).
- Wu, J.; Sheng, V. S.; Zhang, J.; Li, H.; Dadakova, T.; Swisher, C. L.; Cui, Z.; and Zhao, P. 2020. Multi-Label Active Learning Algorithms for Image Classification: Overview and Future Promise. *ACM Comput. Surv.*, 53(2).
- Yang, B.; Sun, J.-T.; Wang, T.; and Chen, Z. 2009. Effective Multi-Label Active Learning for Text Classification. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’09, 917–926. New York, NY, USA: Association for Computing Machinery. ISBN 9781605584959.
- Yang, Y.-Y.; Huang, K.-H.; Chang, C.-W.; and Lin, H.-T. 2018. Cost-Sensitive Reference Pair Encoding for Multi-Label Learning. In Phung, D.; Tseng, V. S.; Webb, G. I.; Ho, B.; Ganji, M.; and Rashidi, L., eds., *Advances in Knowledge Discovery and Data Mining*, 143–155. Cham: Springer International Publishing.
- Zhang, M.-L.; Li, Y.-K.; Liu, X.-Y.; and Geng, X. 2018. Binary relevance for multi-label learning: an overview. *Frontiers of Computer Science*, 12(2): 191–202.
- Zhao, G.; Dougherty, E.; Yoon, B.-J.; Alexander, F.; and Qian, X. 2021. Uncertainty-aware Active Learning for Optimal Bayesian Classifier. In *Proc. 9th Int. Conf. Learn. Representations*.