

NOVELTY DETECTION UNDER MULTI-INSTANCE MULTI-LABEL FRAMEWORK

Qi Lou, Raviv Raich, Forrest Briggs and Xiaoli Z. Fern

School of EECS, Oregon State University, Corvallis, OR 97331-5501, USA

{lou, raich, briggsf, xfern}@eecs.oregonstate.edu

ABSTRACT

Novelty detection plays an important role in machine learning and signal processing. This paper studies novelty detection in a new setting where the data object is represented as a bag of instances and associated with multiple class labels, referred to as multi-instance multi-label (MIML) learning. Contrary to the common assumption in MIML that each instance in a bag belongs to one of the known classes, in novelty detection, we focus on the scenario where bags may contain novel-class instances. The goal is to determine, for any given instance in a new bag, whether it belongs to a known class or a novel class. Detecting novelty in the MIML setting captures many real-world phenomena and has many potential applications. For example, in a collection of tagged images, the tag may only cover a subset of objects existing in the images. Discovering an object whose class has not been previously tagged can be useful for the purpose of soliciting a label for the new object class. To address this novel problem, we present a discriminative framework for detecting new class instances. Experiments demonstrate the effectiveness of our proposed method, and reveal that the presence of unlabeled novel instances in training bags is helpful to the detection of such instances in testing stage.

Index Terms— Novelty detection, multi-instance multi-label, kernel method

1. INTRODUCTION

Novelty detection is the identification of new or unknown data that is not labeled during training [1]. In the traditional setting, only training examples from a nominal distribution are provided and the goal is to determine for a new example whether it comes from the nominal distribution or not. Much work has been done in this field. Early work is generally divided into two categories [1, 2]. One category includes statistical approaches such as some density estimation methods. The other category consists of neural network based approaches, e.g., multi-layer perceptrons. Several new approaches have been introduced in recent years. In [3], geo-

metric entropy minimization is introduced for anomaly detection. An efficient anomaly detection method using bipartite k -NN graphs is presented in [4]. In [5], an anomaly detection algorithm is proposed based on score functions. Each point gets scores from its nearest neighbors. This algorithm can be directly applied to novelty detection. In [6], SVMs are applied to novelty detection to learn a function f that is positive on a subset S of the input space and negative outside S .

In this paper, we consider novelty detection in a new setting where the data follows a multi-instance multi-label (MIML) format. The MIML framework has been primarily studied for supervised learning [7] and widely used in applications where data is associated with multiple classes and can be naturally represented as bags of instances (i.e., collections of parts). For example, a document can be viewed as a bag of words and associated with multiple tags. Similarly, an image can be represented as a bag of pixels or patches, and associated with multiple classes corresponding to the objects that it contains. Formally speaking, the training data in MIML consists of a collection of labeled bags $\{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$, where $X_i \subset X$ is a set of instances and $Y_i \subset Y$ is a set of labels. In the traditional MIML applications the goal is to learn a bag-level classifier $f : 2^X \rightarrow 2^Y$ that can reliably predict the label set of a previously unseen bag.

It is commonly assumed in MIML that every instance we observe in the training set belongs to one of the known classes. However, in many applications, this assumption is violated. For example, in a collection of tagged images, the tag may only cover a subset of objects present in the images. The goal of novelty detection in the MIML setting is to determine whether a given instance comes from an unknown class given only a set of bags labeled with the known classes. This setup has several advantages compared to the more well-known setup in novelty detection: First, the labeled bags allow us to apply an approach that takes into account the presence of multiple known classes. Second, frequently the training set would contain some novel class instances. The presence of such instances, although never explicitly labeled as novel instances, can in a way serve as “implicit” negative examples for the known classes, which can be helpful for identifying novel instances in new bags.

This work was partially supported by the National Science Foundation grant CCF-1254218.

The work presented in this paper is inspired by a real world bioacoustics application. In this application, the annotation of individual bird vocalization is often a time consuming task. As an alternative, experts identify from a list of focal bird species the ones that they recognize in a given recording. Such labels are associated with the entire recording and not with a specific vocalization in the recording. Based on a collection of such labeled recordings, the goal is to annotate each vocalization in a new recording [8]. An implicit assumption here is that each vocalization in the recording must come from one of the focal species, which can be incomplete. Under this assumption, vocalizations of new species outside of the focal list will not be discovered. Instead, such vocalizations will be annotated with a label from the existing species list. The setup proposed in this paper allows for novel instances to be observed in the training data without being explicitly labeled, and hence should enable the annotation of vocalizations from novel species. In turn, such novel instances can be presented back to the experts for further inspection.

To the best of our knowledge, novelty detection in the MIML setting has not been investigated. Our main contributions are: (i) We propose a new problem – novelty detection in the MIML setting. (ii) We offer a framework based on score functions to solve the problem. (iii) We illustrate the efficacy of our method on a real-world MIML bioacoustics data.

2. PROPOSED METHODS

Suppose we are given a collection of labeled bags $\{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$, where the i th bag $X_i \subset X$ is a set of instances from the feature space $X \subset \mathbb{R}^d$, and Y_i is a subset of the know label set $Y = \bigcup_{i=1}^N Y_i$. For any label $y_{im} \in Y_i$, there is at least one instance $x_{im} \in X_i$ belonging to this class. We consider the scenario where an instance in X_i has no label in Y_i related to it, which extends the traditional MIML learning framework. Our goal is to determine for a given instance $x \in X$ whether it belongs to a known class in Y or not.

To illustrate the intuition behind our general strategy, consider the toy problem shown in Table 1. The known label set is $\{I, II\}$. We have four labeled bags available. According to the principle that one instance must belong to one class and one bag-level label must have at least one corresponding instance, we conclude that Δ is drawn from class I, \square belongs to class II, and \diamond doesn't come from the existing classes. ∇ cannot be fully determined based on current data.

To express this observation mathematically, we calculate the rate of co-occurrence of an instance and a label. For example, Δ appears with label I together in bag 1, 2, 4 and they are both missing in bag 3. So, the co-occurrence rate $p(\Delta, I) = 1$. All the other rates are listed in Table 2. If we detect an instance based on the maximal co-occurrence rate with respect to all classes and set a threshold to be $3/4$, we will reach a result that can generally reflect our previous observation.

Table 1. Toy problem with two known classes

#	Bags (X_i)	Labels (Y_i)
1	$\{\Delta \Delta \square \nabla\}$	$\{I, II\}$
2	$\{\Delta \Delta \diamond \nabla\}$	$\{I\}$
3	$\{\square \square \diamond \diamond\}$	$\{II\}$
4	$\{\Delta \Delta \square\}$	$\{I, II\}$

Table 2. Concurrence rates for the toy problem

	Δ	∇	\square	\diamond
I	1	3/4	1/2	1/4
II	1/2	1/4	1	1/4

This example inspires us to devise a general strategy for detection. We introduce a set of score functions, each of which corresponds to one class, i.e., for each label $c \in Y$, we assign a function f_c to class c . Generally, for an instance from a specific known class, the value of the score function corresponding to this class should be large. If all scores of an instance are below a prescribed threshold, it would not be considered to belong to any known class. The decision principle is: If $\max_{c \in \{1, \dots, |Y|\}} f_c(x) < \varepsilon$ then return 'unknown', otherwise return 'known'.

There are many possible choices for the set of score functions. Generally, the score functions are expected to enable us to achieve a high true positive rate with a given false positive (Type I error) rate, which can be measured by the area under the curve (AUC) of ROC.

2.1. Kernel Based Scoring Functions

We define the score function for class c as follows:

$$\begin{aligned} f_c(x) &= \sum_{x_l \in \bigcup_i X_i} \alpha_{cl} k(x, x_l) \\ &= \alpha_c^T k(x) \end{aligned} \quad (1)$$

where X_i 's are training bags, x_l 's are training instances from training bags, $k(\cdot, \cdot)$ is the kernel function such that $k(x) = (k(x, x_1), \dots, k(x, x_L))^T$, and α_{cl} 's are the components of the weight vector $\alpha_c = (\alpha_{c1}, \dots, \alpha_{cL})^T$.

We encourage f_c to take positive values on instances in class c and negative values on instances from other classes. Hence, we define the objective function \mathcal{OBF} as

$$\frac{\lambda}{2} \sum_{c=1}^{|Y|} \alpha_c^T \mathbf{K} \alpha_c + \frac{1}{N|Y|} \sum_{i=1}^N \sum_{c=1}^{|Y|} F_c(X_i) \quad (2)$$

where

$$F_c(X_i) = \max\{0, 1 - y_{ic} \max_{x_{ij} \in X_i} f_c(x_{ij})\}, y_{ic} \in \{-1, +1\}$$

λ is a regularization parameter, \mathbf{K} is the kernel matrix with (i, j) -th entry $k(x_i, x_j)$, $x_i, x_j \in \bigcup_k X_k$, and $y_{ic} = +1$ if and only if Y_i contains the label for class c .

In fact, we define an objective function for each class separately and sum over all these objective functions to construct \mathcal{OBJ} . The first term of \mathcal{OBJ} controls model complexity. $F_c(\cdot)$ in the second term of \mathcal{OBJ} can be viewed as a bag-level hinge loss for class c , which is a generalization of the single-instance case. If c is a bag-level label of bag X_i , we expect $\max_{x_{ij} \in X_i} f_c(x_{ij})$ to give a high score because there is at least one instance in X_i is from class c . Other loss functions such as rank loss [8] have already been introduced for MIML learning.

Our goal is to minimize the objective function which is unfortunately non-convex. However, if we fix the term $\max_{x_{ij} \in X_i} f_c(x_{ij})$, i.e., find the support instance x_{ic} such that $x_{ic} = \operatorname{argmax}_{x_{ij} \in X_i} \alpha_c^T k(x_{ij})$ and substitute back to the objective function, the resulted objective function \mathcal{OBJ}_* will be convex with respect to α_c 's. To solve this convex problem, we deploy the L-BFGS [9] algorithm. The subgradient along α_c used in L-BFGS is computed as follows:

$$\nabla_c = \lambda \mathbf{K} \alpha_c - \frac{1}{N|Y|} \sum_{i=1}^N y_{ic} k(x_{ic}) \mathbb{1}_{\{1 - y_{ic} f_c(x_{ic}) > 0\}} \quad (3)$$

Details can be found in Algorithm 1. This descent method can be applied to any choice of kernel function and according to our experience it works very well (usually converges within 30 steps). Note that many algorithms [8, 10] for MIML learning that attempt to learn an instance-level score functions including the proposed approach are based on a non-convex objective. Consequently, no global optimum is guaranteed. To reduce the effect induced by randomness, we usually rerun the algorithm multiple times with independent random initializations and adopt the result with the smallest value of the objective function.

Algorithm 1 Descent Method

Require: $\{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$, λ , T .
 Randomly initialize all α_c 's s.t. $\|\alpha_c^1\| = 1$
for $t = 1$ to T **do**
 Set $x_{ic}^t = \operatorname{argmax}_{x_{ij} \in X_i} (\alpha_c^t)^T k(x_{ij})$,
 $\mathbb{1}_{ic}^t = \mathbb{1}_{\{1 - y_{ic} f_c(x_{ic}^t) > 0\}}$,
 $\nabla_c^t = \lambda \mathbf{K} \alpha_c - \frac{1}{N|Y|} \sum_{i=1}^N y_{ic} k(x_{ic}^t) \mathbb{1}_{ic}^t$.
 Plug $\{x_{ic}^t\}$ into \mathcal{OBJ} to get a convex surrogate \mathcal{OBJ}_*^t .
 Run L-BFGS with inputs \mathcal{OBJ}_*^t , ∇_c^t to return $\{\alpha_c^{t+1}\}$
 and \mathcal{OBJ}^{t+1}
end for
return $\{\alpha_c^{T+1}\}$ and \mathcal{OBJ}^{T+1} .

2.2. Parameter Tuning

In our experiment, we use Gaussian kernel, i.e., $k(x_i, x_j) = e^{-\gamma \|x_i - x_j\|^2}$, where $\|\cdot\|$ is the Euclidean norm. The parameter γ controls the bandwidth of the kernel. Hence, there are a pair of parameters λ and γ in the objective function required to be determined.

While training, we search in a wide range of values for the parameter pair, and select the pair with corresponding α_c 's that minimizes

$$\sum_{i=1}^N \sum_{c=1}^{|Y|} g(y_{ic} \max_{x_{ij} \in X_i} f_c(x_{ij}))$$

where $g(x) = \mathbb{1}_{x < 0}$ is the zero-one loss function. Note that $\mathbb{1}_{x < 0}$ is a lower bound of the hinge loss $\max\{0, 1 - x\}$.

We vary the value of threshold to generate ROCs while testing. The values of threshold are derived from training examples.

3. EXPERIMENTAL RESULTS

In this section, we provide a number of experimental results based on both synthetic data and real-world data to show the effectiveness of our algorithm. Additionally, we present a comparison to one-class SVM, a notable anomaly detection algorithm.

3.1. MNIST Handwritten Digits Dataset

We generated the synthetic data based on the MNIST handwritten digits data set¹. Each image in the data set is a 28 by 28 bitmap, i.e., a vector of 784 dimensions. By using PCA, we reduced the dimension of instances to 20.

Table 3. Bag examples for the handwritten digits data. We take the first four digits ‘0’, ‘1’, ‘2’, ‘3’ as known classes, i.e., $Y = \{‘0’, ‘1’, ‘2’, ‘3’\}$. In each bag, some instances are without associated labels. For example, in bag 1 examples for ‘5’ and ‘9’ are considered from unknown classes.

number	bags	labels
1		‘1’
2		‘0’, ‘1’
3		‘2’, ‘3’
4		‘0’, ‘1’
5		‘0’, ‘2’

We created training and testing bags from the MNIST instances. Some examples for handwritten digits bags are

¹Available on-line <http://www.cs.nyu.edu/~roweis/data.html>

shown in Table 3. Two processes for generating bags are listed in Algorithm 2 and Algorithm 3. The only difference between these two procedures is that Algorithm 3 rules out the possibility of a label set for a bag being empty, i.e., a bag including purely novel examples. For Dirichlet process used in our simulation, we assigned relatively small concentration parameters $\beta = (\beta_1, \beta_2, \dots, \beta_{10})$ to the Dirichlet distribution in order to encourage a sparse label set for a bag, which is common in real-world scenarios. We set all $\beta_i = 0.1$ and the bag size $M = 20$. Typical examples of bags generated from Dirichlet distribution are shown in Table 4.

Table 4. Examples for numbers of each digit in 5 bags when each component of β is 0.1. The bag size is set to be 20.

'0'	'1'	'2'	'3'	'4'	'5'	'6'	'7'	'8'	'9'
0	1	0	0	5	5	2	3	4	0
6	0	4	8	0	0	1	0	0	1
0	17	0	2	0	0	1	0	0	0
0	0	1	0	2	16	1	0	0	0
0	0	15	0	0	0	0	5	0	0

Algorithm 2 Bag generation procedure for handwritten digits data

Require: N, M, Y, β .

for $i = 1$ to N **do**

Draw M instances $\{x_{ij}\}$ according to the proportion given by Dirichlet (β) distribution.

Extract labels from x_{ij} 's to form Y'_i and set $Y_i = Y \cap Y'_i$.

end for

Algorithm 3 Bag generation procedure with filtration for handwritten digits data.

Require: N, M, Y, β .

for $i = 1$ to N **do**

Set $Y_i = \emptyset$.

while $Y_i == \emptyset$ **do**

Draw M instances $\{x_{ij}\}$ according to the proportion given by Dirichlet (β) distribution.

Extract labels from x_{ij} 's to form Y'_i and set $Y_i = Y \cap Y'_i$.


end while

end for


We provided our method with bags generated in two different ways:

1. Generate both training and testing bags according to Algorithm 2.
2. Generate training bags according to Algorithm 3 while generate testing bags by applying Algorithm 2.

In our experiments, we consider various sizes of known label sets and different combinations of labels in these two settings. Two typical examples of ROCs from the two setting are shown in Figure 1.



(a)



(b)

Fig. 1. Typical examples of ROCs from the handwritten digit data. The subfigure (a) shows a ROC example from the first setting and the subfigure (b) gives an example from the second setting.

Table 5 shows the average AUCs of ROCs over multiple runs from the first setting. We observe that average AUCs are all above 0.85 for the known label sets of size 4. For the known label sets of size 8, the average AUCs are all larger than 0.8. The results are fairly stable with different combinations of labels. This demonstrates the effectiveness of our algorithm.

Table 6 shows the average AUCs of ROCs for the setting which does not contain bags with an empty label set. The label sets in these two tables are the same. The results in the two tables are comparable but those in Table 5 are always better. This demonstrates that it is beneficial to include bags with an empty label set. The reason could be that those bags contain purely novel examples and hence training on those

bags is very reliable.

Table 5. Average AUCs for handwritten digits data. Y is the known label set. Training bags and testing bags are both generated according to Algorithm 2, i.e., without bag filtration.

Y	AUC	Y	AUC
{'0','1','3','7'}	0.89	{'0','1','2','3','4','5','6','7'}	0.85
{'2','4','7','8'}	0.87	{'2','3','4','5','6','7','8','9'}	0.88
{'2','5','6','7'}	0.91	{'0','1','4','5','6','7','8','9'}	0.84
{'3','5','7','9'}	0.85	{'0','1','2','3','6','7','8','9'}	0.85
{'3','6','8','9'}	0.89	{'0','1','2','3','4','5','8','9'}	0.83

Table 6. Average AUCs for handwritten digits data. Y is the known label set. Training bags are generated according to Algorithm 3, i.e., with bag filtration, while testing bags are generated by Algorithm 2, i.e., without bag filtration.

Y	AUC	Y	AUC
{'0','1','3','7'}	0.86	{'0','1','2','3','4','5','6','7'}	0.85
{'2','4','7','8'}	0.86	{'2','3','4','5','6','7','8','9'}	0.84
{'2','5','6','7'}	0.88	{'0','1','4','5','6','7','8','9'}	0.82
{'3','5','7','9'}	0.83	{'0','1','2','3','6','7','8','9'}	0.84
{'3','6','8','9'}	0.86	{'0','1','2','3','4','5','8','9'}	0.80

3.2. HJA Birdsong Dataset

We tested our algorithm on the real-world dataset - HJA birdsong dataset², which has been used in [11, 12]. This dataset consists of 548 bags, each of which contains several 38-dimensional instances. The bag size, i.e., the number of instances in a bag, varies from 1 to 26, the average of which is approximately 9. The dataset includes 4998 instances from 13 species. Species names and the numbers of instances for those species are listed in Table 7. Each species corresponds to a class in the complete label set $\{1, 2, \dots, 13\}$. We took a subset of the complete label set as the known label set and conducted experiment with various choices of the known label set. Table 8 shows the average AUCs of different known label sets. Specifically, we intentionally made each species appear at least once in those known sets. From Table 8, we observe that most all of the values of AUCs are above 0.85 and some even reach 0.9. The results are quite stable with different label settings despite the imbalance in the instance population of the species. These results illustrate the potential of the approach as a utility for novel species discovery.

Table 7. Names of bird species and the number of total instances for each species. Each species corresponds to one class.

Class	Species	No. of Instances
1	Brown Creeper	602
2	Winter Wren	810
3	Pacific-slope Flycatcher	501
4	Red-breasted Nuthatch	494
5	Dark-eyed Junco	82
6	Olive-sided Flycatcher	277
7	Hermit Thrush	32
8	Chestnut-backed Chickadee	345
9	Varied Thrush	139
10	Hermit Warbler	120
11	Swainson's Thrush	190
12	Hammond's Flycatcher	1280
13	Western Tanager	126

Table 8. Average AUCs for birdsong data. Y is the known label set.

Y	AUC	Y	AUC
{1,2,4,8}	0.90	{1,2,3,4,5,6,7,8}	0.89
{3,5,7,9}	0.85	{3,4,5,6,7,8,9,10}	0.85
{4,6,8,10}	0.88	{5,6,7,8,9,10,11,12}	0.89
{5,7,9,11}	0.90	{1,7,8,9,10,11,12,13}	0.84
{6,10,12,13}	0.89	{1,2,3,9,10,11,12,13}	0.85

3.3. Comparison with One-Class SVM

Our algorithm deals with detection problem with MIML setting, which is different from the traditional setting for anomaly detection. We argue that traditional anomaly detection algorithms cannot be directly applied to our problem. To make comparison, we adopt one-class SVM [13–15], a well known algorithm for anomaly detection. To apply one-class SVM, we construct a normal class training data consisting of examples from the known label set. The parameter ν vary from 0 to 1 with step size 0.02 to generate ROCs. The Gaussian kernel is used for one-class SVM. We search the parameter γ for the kernel in a wide range and select the best one for on-class SMV post-hoc. We present this unfair advantage to one-class SVM for two reasons: (i) It is unclear how to optimize the parameter in the absence of novel instances. (ii) We would like to illustrate the point that even given such unfair advantage, one-class SVM cannot outperform our algorithm.

Table 9 and 10 show the average AUCs for handwritten digits data and birdsong data respectively. Compared to Tables 5 and 8, the proposed algorithm outperforms 1-class SVM

²Available on-line <http://web.engr.oregonstate.edu/~briggs/> and <http://www.birds.org>

in terms of AUC not only in absolute value but also in stability. This also demonstrates that training with unlabeled instances are beneficial to the detection.

Table 9. Average AUCs for the handwritten digits data by applying one-class SVM with Gaussian kernel. Y is the known label set.

Y	AUC	Y	AUC
{‘0’,‘1’,‘3’,‘7’}	0.66	{‘0’,‘1’,‘2’,‘3’,‘4’,‘5’,‘6’,‘7’}	0.59
{‘2’,‘4’,‘7’,‘8’}	0.66	{‘2’,‘3’,‘4’,‘5’,‘6’,‘7’,‘8’,‘9’}	0.63
{‘2’,‘5’,‘6’,‘7’}	0.57	{‘0’,‘1’,‘4’,‘5’,‘6’,‘7’,‘8’,‘9’}	0.68
{‘3’,‘5’,‘7’,‘9’}	0.65	{‘0’,‘1’,‘2’,‘3’,‘6’,‘7’,‘8’,‘9’}	0.62
{‘3’,‘6’,‘8’,‘9’}	0.63	{‘0’,‘1’,‘2’,‘3’,‘4’,‘5’,‘8’,‘9’}	0.65

Table 10. Average AUCs for the birdsong data by applying one-class SVM. Y is the known label set.

Y	AUC	Y	AUC
{1,2,4,8}	0.78	{1,2,3,4,5,6,7,8}	0.85
{3,5,7,9}	0.79	{3,4,5,6,7,8,9,10}	0.82
{4,6,8,10}	0.82	{5,6,7,8,9,10,11,12}	0.75
{5,7,9,11}	0.73	{1,7,8,9,10,11,12,13}	0.70
{6,10,12,13}	0.78	{1,2,3,9,10,11,12,13}	0.60

4. CONCLUSION

In this paper, we proposed a new problem – novelty detection in the MIML setting and offered a framework based on score functions to solve the problem. A large number of simulations show that our algorithm not only works well on synthetic data but also on real-world data. We also demonstrate that the presence of unlabeled examples in the training set is useful to detect new class examples while testing. We present the advantage in the MIML setting for novelty detection. Even though positive examples for the novelty that are not directly labeled, their presence provides a clear advantage over methods that rely on data that does not include novel class examples.

There are many relative problems call for investigation. One will be on how to use the information of bag-level labels in detection if bag-level labels are available, which will possibly improve the performance of our algorithm since we did not make use of such information in our experiment.

5. REFERENCES

- [1] Markos Markou and Sameer Singh, “Novelty detection: A review - part 1: Statistical approaches,” *Signal Processing*, vol. 83, pp. 2481–2497, 2003.
- [2] Markos Markou and Sameer Singh, “Novelty detection: A review - part 2: Neural network based approaches,” *Signal Processing*, vol. 83, pp. 2499–2521, 2003.
- [3] Alfred O. Hero, “Geometric entropy minimization (gem) for anomaly detection and localization,” in *NIPS*, 2006, pp. 585–592, MIT Press.
- [4] Kumar Sricharan and Alfred O. Hero, “Efficient anomaly detection using bipartite k-nn graphs,” in *NIPS*, 2011, pp. 478–486.
- [5] Manqi Zhao and Venkatesh Saligrama, “Anomaly detection with score functions based on nearest neighbor graphs,” in *NIPS*, 2009, pp. 2250–2258.
- [6] Bernhard Schölkopf, Robert C. Williamson, Alex J. Smola, John Shawe-Taylor, and John C. Platt, “Support vector method for novelty detection,” in *NIPS*, 1999, pp. 582–588.
- [7] Zhi-Hua Zhou, Min-Ling Zhang, Sheng-Jun Huang, and Yu-Feng Li, “Multi-instance multi-label learning,” *Artif. Intell.*, vol. 176, no. 1, pp. 2291–2320, 2012.
- [8] Forrest Briggs, Xiaoli Z. Fern, and Raviv Raich, “Rank-loss support instance machines for miml instance annotation,” in *KDD*, 2012.
- [9] Richard H. Byrd, Jorge Nocedal, and Robert B. Schnabel, “Representations of quasi-newton matrices and their use in limited memory methods,” 1994.
- [10] Oksana Yakhnenko and Vasant Honavar, “Multi-instance multi-label learning for image classification with large vocabularies,” in *Proceedings of the British Machine Vision Conference*. 2011, pp. 59.1–59.12, BMVA Press.
- [11] Forrest Briggs, Xiaoli Z. Fern, Raviv Raich, and Qi Lou, “Instance annotation for multi-instance multi-label learning,” *Transactions on Knowledge Discovery from Data (TKDD)*, 2012.
- [12] Li-Ping Liu and Thomas G. Dietterich, “A conditional multinomial mixture model for superset label learning,” in *NIPS*, 2012, pp. 557–565.
- [13] Bernhard Schölkopf, John C. Platt, John Shawe-taylor, Alex J. Smola, and Robert C. Williamson, “Estimating the support of a high-dimensional distribution,” 1999.
- [14] Chih-Chung Chang and Chih-Jen Lin, “LIBSVM: A library for support vector machines,” *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.
- [15] Larry M. Manevitz, Malik Yousef, Nello Cristianini, John Shawe-taylor, and Bob Williamson, “One-class svms for document classification,” *Journal of Machine Learning Research*, vol. 2, pp. 139–154, 2001.