

A Review on Multi-Label Learning Algorithms

Min-Ling Zhang and Zhi-Hua Zhou, *Fellow, IEEE*

Abstract

Multi-label learning studies the problem where each example is represented by a *single* instance while associated with a *set of* labels simultaneously. During the past decade, significant amount of progresses have been made towards this emerging machine learning paradigm. This paper aims to provide a timely review on this area, with emphasis on state-of-the-art multi-label learning algorithms. Firstly, fundamentals on multi-label learning including formal definition and evaluation metrics are given. Secondly and primarily, twelve representative multi-label learning algorithms are scrutinized under common notations, with corresponding analyses and discussions. Thirdly, several extended topics on multi-label learning are briefly summarized. As a conclusion, online resources and open research problems on multi-label learning are outlined for reference purposes.

Index Terms

Machine learning, multi-label learning, evaluation metrics, label correlations, problem transformation, algorithm adaptation.

I. INTRODUCTION

Traditional supervised learning is one of the mostly-studied machine learning paradigms, where each real-world object (example) is represented by a single instance (feature vector) and associated with a single label. Formally, let \mathcal{X} denote the instance space and \mathcal{Y} denote the label space, the task of traditional supervised learning is to learn a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ from the training set $\{(\mathbf{x}_i, y_i) \mid 1 \leq i \leq m\}$. Here, $\mathbf{x}_i \in \mathcal{X}$ is an instance characterizing

Min-Ling Zhang is with the School of Computer Science and Engineering, and the MOE Key Laboratory of Computer Network and Information Integration, Southeast University, Nanjing 210096, China. Email: zhangml@seu.edu.cn.

Zhi-Hua Zhou is with the National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China. Email: zhouzh@lamda.nju.edu.cn. (Corresponding author)

the properties (features) of an object and $y_i \in \mathcal{Y}$ is the corresponding label characterizing its semantics. Therefore, one fundamental assumption adopted by traditional supervised learning is that each example belongs to only one concept, i.e. having *unique* semantic meaning.

Although traditional supervised learning is prevailing and successful, there are many learning tasks where the above simplifying assumption does not fit well, as real-world objects might be complicated and have multiple semantic meanings simultaneously. To name a few, in text categorization, a news document could cover several topics such as *sports*, *London Olympics*, *ticket sales* and *torch relay*; In music information retrieval, a piece of symphony could convey various messages such as *piano*, *classical music*, *Mozart* and *Austria*; In automatic video annotation, one video clip could be related to some scenarios, such as *urban* and *building*, and so on.

To account for the multiple semantic meanings that one real-world object might have, one direct solution is to assign *a set of* proper labels to the object to explicitly express its semantics. Following the above consideration, the paradigm of multi-label learning naturally emerges [158]. In contrast to traditional supervised learning, in multi-label learning each object is also represented by a single instance while associated with a set of labels instead of a single label. The task is to learn a function which can predict the proper label sets for unseen instances.¹

Early researches on multi-label learning mainly focus on the problem of multi-label text categorization [51], [104], [129], [159]. During the past decade, multi-label learning has gradually attracted significant attentions from machine learning and related communities and has been widely applied to diverse problems from automatic annotation for multimedia contents including image [1], [7], [13], [93], [170], [171], [174], audio [96], [108], [127], [137], [138], [148], [172], and video [112], [133], [168], [169], to bioinformatics [25], [42], [190], web mining [84], [107], [144], rule mining [118], [146], [161], information retrieval [58], [182], [203], tag

¹In a broad sense, multi-label learning can be regarded as one possible instantiation of multi-target learning [158], where each object is associated with multiple target variables (multi-dimensional outputs) [5]. Different types of target variables would give rise to different instantiations of multi-target learning, such as multi-label learning (binary targets), multi-dimensional classification (categorical/multi-class targets), multi-output/multivariate regression (numerical targets), and even learning with combined types of target variables.

recommendation [83], [134], etc. Specifically, in recent six years (2007-2012), there are more than 60 papers with keyword *multi-label* (or *multilabel*) in the title appearing in major machine learning-related conferences (including ICML/ECML PKDD/IJCAI/AAAI/KDD/ICDM/NIPS).

This paper serves as a timely review on the emerging area of multi-label learning, where its state-of-the-art is presented in three parts.² In the first part (Section II), fundamentals on multi-label learning including formal definition (learning framework, key challenge, threshold calibration) and evaluation metrics (example-based, label-based, theoretical results) are given. In the second and primary part (Section III), technical details of up to twelve representative multi-label algorithms are scrutinized under common notations with necessary analyses and discussions. In the third part (Section IV), several extended topics on multi-label learning (noisy/weak labels, unlabeled data, dimensionality reduction, large-scale data, related learning settings) are briefly summarized. To conclude this review (Section V), online resources and possible lines of future researches on multi-label learning are discussed.

II. THE PARADIGM

A. Formal Definition

1) *Learning Framework*: Suppose $\mathcal{X} = \mathbb{R}^d$ (or \mathbb{Z}^d) denotes the d -dimensional instance space, and $\mathcal{Y} = \{y_1, y_2, \dots, y_q\}$ denotes the label space consisting of q possible class labels. The task of multi-label learning is to learn a function $h : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$ from the multi-label training set $\mathcal{D} = \{(\mathbf{x}_i, Y_i) \mid 1 \leq i \leq m\}$. For each multi-label example (\mathbf{x}_i, Y_i) , $\mathbf{x}_i \in \mathcal{X}$ is a d -dimensional feature vector $(x_{i1}, x_{i2}, \dots, x_{id})^\top$ and $Y_i \subseteq \mathcal{Y}$ is the set of labels associated with \mathbf{x}_i .³ For any

²Note that there have been some nice reviews on multi-label learning techniques [29], [152], [154]. Compared to earlier attempts in this regard, we strive to provide an enriched version with the following enhancements: a) In-depth descriptions on more algorithms; b) Comprehensive introductions on latest progresses; c) Succinct summarizations on extended topics.

³In this paper, the term “multi-label learning” is used in equivalent sense as “multi-label classification” since labels assigned to each instance are considered to be binary. Furthermore, there are alternative multi-label settings where other than a single instance each example is represented by a bag of instances [201], [202] or graphs [87], [89], or extra ontology knowledge might exist on the label space such as hierarchy structure [2], [4], [14], [17], [124], [125], [132], [162]. To keep the review comprehensive yet well-focused, examples are assumed to adopt *single-instance* representation and possess *flat* class labels.

TABLE I
SUMMARY OF MAJOR MATHEMATICAL NOTATIONS.

Notations	Mathematical Meanings
\mathcal{X}	d -dimensional instance space \mathbb{R}^d (or \mathbb{Z}^d)
\mathcal{Y}	label space with q possible class labels $\{y_1, y_2, \dots, y_q\}$
\mathbf{x}	d -dimensional feature vector $(x_1, x_2, \dots, x_d)^\top$ ($\mathbf{x} \in \mathcal{X}$)
Y	label set associated with \mathbf{x} ($Y \subseteq \mathcal{Y}$)
\bar{Y}	complementary set of Y in \mathcal{Y}
\mathcal{D}	multi-label training set $\{(\mathbf{x}_i, Y_i) \mid 1 \leq i \leq m\}$
\mathcal{S}	multi-label test set $\{(\mathbf{x}_i, Y_i) \mid 1 \leq i \leq p\}$
$h(\cdot)$	multi-label classifier $h : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$, where $h(\mathbf{x})$ returns the set of proper labels for \mathbf{x}
$f(\cdot, \cdot)$	real-valued function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, where $f(\mathbf{x}, y)$ returns the confidence of being proper label of \mathbf{x}
$rank_f(\cdot, \cdot)$	$rank_f(\mathbf{x}, y)$ returns the rank of y in \mathcal{Y} based on the descending order induced from $f(\mathbf{x}, \cdot)$
$t(\cdot)$	thresholding function $t : \mathcal{X} \rightarrow \mathbb{R}$, where $h(\mathbf{x}) = \{y \mid f(\mathbf{x}, y) > t(\mathbf{x}), y \in \mathcal{Y}\}$
$ \cdot $	$ \mathcal{A} $ returns the cardinality of set \mathcal{A}
$\llbracket \cdot \rrbracket$	$\llbracket \pi \rrbracket$ returns 1 if predicate π holds, and 0 otherwise
$\phi(\cdot, \cdot)$	$\phi(Y, y)$ returns +1 if $y \in Y$, and -1 otherwise
\mathcal{D}_j	binary training set $\{(\mathbf{x}_i, \phi(Y_i, y_j)) \mid 1 \leq i \leq m\}$ derived from \mathcal{D} for the j -th class label y_j
$\psi(\cdot, \cdot, \cdot)$	$\psi(Y, y_j, y_k)$ returns +1 if $y_j \in Y$ and $y_k \notin Y$, and -1 if $y_j \notin Y$ and $y_k \in Y$
\mathcal{D}_{jk}	binary training set $\{(\mathbf{x}_i, \psi(Y_i, y_j, y_k)) \mid \phi(Y_i, y_j) \neq \phi(Y_i, y_k), 1 \leq i \leq m\}$ derived from \mathcal{D} for the label pair (y_j, y_k)
$\sigma_{\mathcal{Y}}(\cdot)$	injective function $\sigma_{\mathcal{Y}} : 2^{\mathcal{Y}} \rightarrow \mathbb{N}$ mapping from the power set of \mathcal{Y} to natural numbers ($\sigma_{\mathcal{Y}}^{-1}$ being the corresponding inverse function)
$\mathcal{D}_{\mathcal{Y}}^\dagger$	multi-class (single-label) training set $\{(\mathbf{x}_i, \sigma_{\mathcal{Y}}(Y_i)) \mid 1 \leq i \leq m\}$ derived from \mathcal{D}
\mathcal{B}	binary learning algorithm [complexity: $\mathcal{F}_{\mathcal{B}}(m, d)$ for training; $\mathcal{F}'_{\mathcal{B}}(d)$ for (per-instance) testing]
\mathcal{M}	multi-class learning algorithm [complexity: $\mathcal{F}_{\mathcal{M}}(m, d, q)$ for training; $\mathcal{F}'_{\mathcal{M}}(d, q)$ for (per-instance) testing]

unseen instance $\mathbf{x} \in \mathcal{X}$, the multi-label classifier $h(\cdot)$ predicts $h(\mathbf{x}) \subseteq \mathcal{Y}$ as the set of proper labels for \mathbf{x} .

To characterize the properties of any multi-label data set, several useful multi-label indicators can be utilized [122], [158]. The most natural way to measure the degree of multi-labeledness is *label cardinality*: $LCard(\mathcal{D}) = \frac{1}{m} \sum_{i=1}^m |Y_i|$, i.e. the average number of labels per example; Accordingly, *label density* normalizes label cardinality by the number of possible labels in the label space: $LDen(\mathcal{D}) = \frac{1}{|\mathcal{Y}|} \cdot LCard(\mathcal{D})$. Another popular multi-labeledness measure is *label*

diversity: $LDiv(\mathcal{D}) = |\{Y \mid \exists \mathbf{x} : (\mathbf{x}, Y) \in \mathcal{D}\}|$, i.e. the number of distinct label sets appeared in the data set; Similarly, label diversity can be normalized by the number of examples to indicate the proportion of distinct label sets: $PLDiv(\mathcal{D}) = \frac{1}{|\mathcal{D}|} \cdot LDiv(\mathcal{D})$.

In most cases, the model returned by a multi-label learning system corresponds to a real-valued function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, where $f(\mathbf{x}, y)$ can be regarded as the *confidence* of $y \in \mathcal{Y}$ being the proper label of \mathbf{x} . Specifically, given a multi-label example (\mathbf{x}, Y) , $f(\cdot, \cdot)$ should yield larger output on the *relevant* label $y' \in Y$ and smaller output on the *irrelevant* label $y'' \notin Y$, i.e. $f(\mathbf{x}, y') > f(\mathbf{x}, y'')$. Note that the multi-label classifier $h(\cdot)$ can be derived from the real-valued function $f(\cdot, \cdot)$ via: $h(\mathbf{x}) = \{y \mid f(\mathbf{x}, y) > t(\mathbf{x}), y \in \mathcal{Y}\}$, where $t : \mathcal{X} \rightarrow \mathbb{R}$ acts as a *thresholding function* which dichotomizes the label space into relevant and irrelevant label sets.

For ease of reference, Table I lists major notations used throughout this review along with their mathematical meanings.

2) *Key Challenge*: It is evident that traditional supervised learning can be regarded as a degenerated version of multi-label learning if each example is confined to have only one single label. However, the generality of multi-label learning inevitably makes the corresponding learning task much more difficult to solve. Actually, the key challenge of learning from multi-label data lies in the overwhelming size of output space, i.e. the number of label sets grows *exponentially* as the number of class labels increases. For example, for a label space with 20 class labels ($q = 20$), the number of possible label sets would exceed one million (i.e. 2^{20}).

To cope with the challenge of exponential-sized output space, it is essential to facilitate the learning process by exploiting *correlations* (or dependency) among labels [158], [188]. For example, the probability of an image being labeled annotated with label *Brazil* would be high if we know it has labels *rainforest* and *soccer*; A document is unlikely to be labeled as *entertainment* if it is related to *politics*. Therefore, effective exploitation of the label correlations information is deemed to be crucial for the success of multi-label learning techniques. Existing strategies to label correlations exploitation could among others be roughly categorized into three families, based on the *order of correlations* that the learning techniques have considered [188]:

- *First-order strategy*: The task of multi-label learning is tackled in a *label-by-label* style and thus ignoring co-existence of the other labels, such as decomposing the multi-label learning problem into a number of independent binary classification problems (one per label) [7], [25], [27], [191]. The prominent merit of first-order strategy lies in its conceptual simplicity and high efficiency. On the other hand, the effectiveness of the resulting approaches might be suboptimal due to the ignorance of label correlations.
- *Second-order strategy*: The task of multi-label learning is tackled by considering *pairwise* relations between labels, such as the ranking between relevant label and irrelevant label [42], [49], [190], or interaction between any pair of labels [53], [112], [159], [203], etc. As label correlations are exploited to some extent by second-order strategy, the resulting approaches can achieve good generalization performance. However, there are certain real-world applications where label correlations go beyond the second-order assumption.
- *High-order strategy*: The task of multi-label learning is tackled by considering high-order relations among labels such as imposing all other labels' influences on each label [22], [55], [77], [176], or addressing connections among random subsets of labels [121], [122], [157], etc. Apparently high-order strategy has stronger correlation-modeling capabilities than first-order and second-order strategies, while on the other hand is computationally more demanding and less scalable.

In Section III, a number of multi-label learning algorithms adopting different strategies will be described in detail to better demonstrate the respective pros and cons of each strategy.

3) *Threshold Calibration*: As mentioned in Subsection II-A.1, a common practice in multi-label learning is to return some real-valued function $f(\cdot, \cdot)$ as the learned model [158]. In this case, in order to decide the proper label set for unseen instance \mathbf{x} (i.e. $h(\mathbf{x})$), the real-valued output $f(\mathbf{x}, y)$ on each label should be calibrated against the thresholding function output $t(\cdot)$.

Generally, threshold calibration can be accomplished with two strategies, i.e. setting $t(\cdot)$ as constant function or inducing $t(\cdot)$ from the training examples [45], [74]. For the first strategy, as $f(\mathbf{x}, y)$ takes value in \mathbb{R} , one straightforward choice is to use zero as the calibration constant

[7], [129]. Another popular choice for calibration constant is 0.5 when $f(\mathbf{x}, y)$ represents the posterior probability of y being a proper label of \mathbf{x} [25], [77], [191]. Furthermore, when all the unseen instances in the test set are available, the calibration constant can be set to minimize the difference on certain multi-label indicator between the training set and test set, notably the label cardinality [122], [123].

For the second strategy, a *stacking*-style procedure would be used to determine the thresholding function [42], [116], [190]. One popular choice is to assume a linear model for $t(\cdot)$, i.e. $t(\mathbf{x}) = \langle \mathbf{w}^*, \mathbf{f}^*(\mathbf{x}) \rangle + b^*$ where $\mathbf{f}^*(\mathbf{x}) = (f(\mathbf{x}, y_1), \dots, f(\mathbf{x}, y_q))^T \in \mathbb{R}^q$ is a q -dimensional stacking vector storing the learning system's real-valued outputs on each label. Specifically, to work out the q -dimensional weight vector \mathbf{w}^* and bias b^* , the following *linear least squares* problem is solved based on the training set \mathcal{D} :

$$\min_{\{\mathbf{w}^*, b^*\}} \sum_{i=1}^m (\langle \mathbf{w}^*, \mathbf{f}^*(\mathbf{x}_i) \rangle + b^* - s(\mathbf{x}_i))^2 \quad (1)$$

Here, $s(\mathbf{x}_i) = \arg \min_{a \in \mathbb{R}} (|\{y_j \mid y_j \in Y_i, f(\mathbf{x}_i, y_j) \leq a\}| + |\{y_k \mid y_k \in \bar{Y}_i, f(\mathbf{x}_i, y_k) \geq a\}|)$ represents the *target output* of the stacking model which bipartitions \mathcal{Y} into relevant and irrelevant labels for each training example with minimum misclassifications.

All the above threshold calibration strategies are *general-purpose* techniques which could be used as a post-processing step to any multi-label learning algorithm returning real-valued function $f(\cdot, \cdot)$. Accordingly, there also exist some *ad hoc* threshold calibration techniques which are specific to the learning algorithms [49], [157] and will be introduced as their inherent component in Section III. Instead of utilizing the thresholding function $t(\cdot)$, an equivalent mechanism to induce $h(\cdot)$ from $f(\cdot, \cdot)$ is to specify the number of relevant labels for each example with $t' : \mathcal{X} \rightarrow \{1, 2, \dots, q\}$ such that $h(\mathbf{x}) = \{y \mid \text{rank}_f(\mathbf{x}, y) \leq t'(\mathbf{x})\}$ [74], [144]. Here, $\text{rank}_f(\mathbf{x}, y)$ returns the rank of y when all class labels in \mathcal{Y} are sorted in descending order based on $f(\mathbf{x}, \cdot)$.

B. Evaluation Metrics

1) *Brief Taxonomy*: In traditional supervised learning, generalization performance of the learning system is evaluated with conventional metrics such as accuracy, F-measure, area under

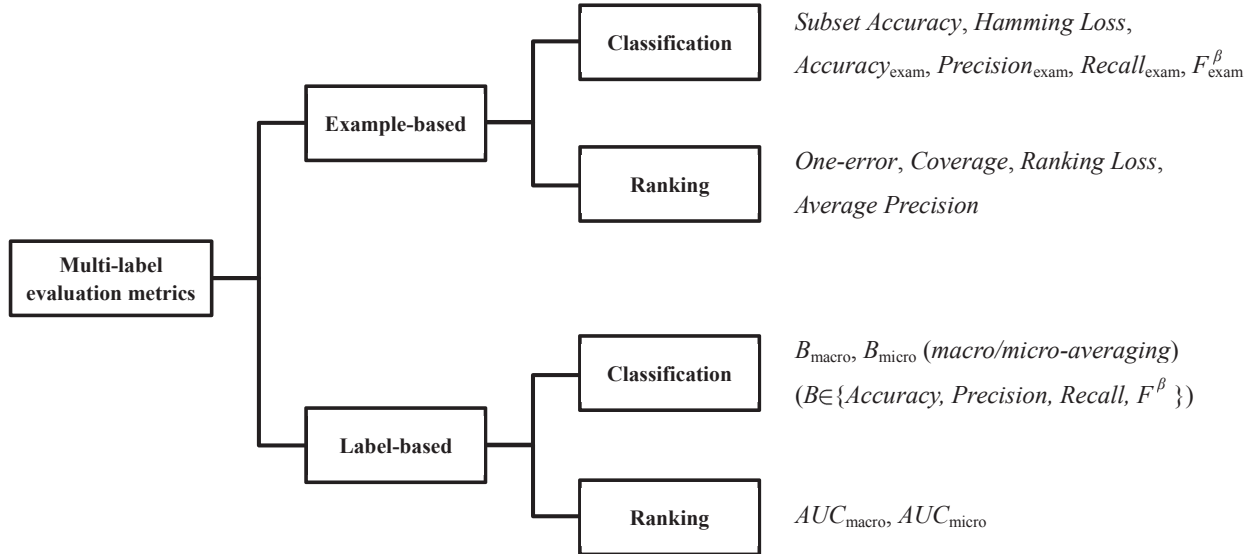


Fig. 1. Summary of major multi-label evaluation metrics.

the ROC curve (AUC), etc. However, performance evaluation in multi-label learning is much complicated than traditional single-label setting, as each example can be associated with multiple labels simultaneously. Therefore, a number of evaluation metrics specific to multi-label learning are proposed, which can be generally categorized into two groups, i.e. *example-based* metrics [53], [55], [129] and *label-based* metrics [157].

Following the notations in Table I, let $\mathcal{S} = \{(\mathbf{x}_i, Y_i) \mid 1 \leq i \leq p\}$ be the test set and $h(\cdot)$ be the learned multi-label classifier. Example-based metrics work by evaluating the learning system's performance on each test example separately, and then returning the *mean value* across the test set. Different to the above example-based metrics, label-based metrics work by evaluating the learning system's performance on each class label separately, and then returning the *macro/micro-averaged value* across all class labels.

Note that with respect to $h(\cdot)$, the learning system's generalization performance is measured from *classification* perspective. However, for either example-based or label-based metrics, with respect to the real-valued function $f(\cdot, \cdot)$ which is returned by most multi-label learning systems as a common practice, the generalization performance can also be measured from *ranking*

perspective. Fig. 1 summarizes the major multi-label evaluation metrics to be introduced in Subsections II-B.2 and II-B.3.

2) *Example-based Metrics*: Following the notations in Table I, six example-based classification metrics can be defined based on the multi-label classifier $h(\cdot)$ [53], [55], [129]:

- *Subset Accuracy*:

$$\text{subsetacc}(h) = \frac{1}{p} \sum_{i=1}^p \mathbb{I}[h(\mathbf{x}_i) = Y_i]$$

The subset accuracy evaluates the fraction of correctly classified examples, i.e. the predicted label set is identical to the ground-truth label set. Intuitively, subset accuracy can be regarded as a multi-label counterpart of the traditional accuracy metric, and tends to be overly strict especially when the size of label space (i.e. q) is large.

- *Hamming Loss*:

$$\text{hloss}(h) = \frac{1}{p} \sum_{i=1}^p |h(\mathbf{x}_i) \Delta Y_i|$$

Here, Δ stands for the symmetric difference between two sets. The hamming loss evaluates the fraction of misclassified instance-label pairs, i.e. a relevant label is missed or an irrelevant is predicted. Note that when each example in \mathcal{S} is associated with only one label, $\text{hloss}_{\mathcal{S}}(h)$ will be $2/q$ times of the traditional misclassification rate.

- *Accuracy_{exam}, Precision_{exam}, Recall_{exam}, F_{exam}^{β}* :

$$\text{Accuracy}_{\text{exam}}(h) = \frac{1}{p} \sum_{i=1}^p \frac{|Y_i \cap h(\mathbf{x}_i)|}{|Y_i \cup h(\mathbf{x}_i)|}; \quad \text{Precision}_{\text{exam}}(h) = \frac{1}{p} \sum_{i=1}^p \frac{|Y_i \cap h(\mathbf{x}_i)|}{|h(\mathbf{x}_i)|}$$

$$\text{Recall}_{\text{exam}}(h) = \frac{1}{p} \sum_{i=1}^p \frac{|Y_i \cap h(\mathbf{x}_i)|}{|Y_i|}; \quad F_{\text{exam}}^{\beta}(h) = \frac{(1 + \beta^2) \cdot \text{Precision}_{\text{exam}}(h) \cdot \text{Recall}_{\text{exam}}(h)}{\beta^2 \cdot \text{Precision}_{\text{exam}}(h) + \text{Recall}_{\text{exam}}(h)}$$

Here, Y_i and $h(\mathbf{x}_i)$ correspond to the *ground-truth* and *predicted* label set for \mathbf{x}_i respectively.

Furthermore, F_{exam}^{β} is an integrated version of $\text{Precision}_{\text{exam}}(h)$ and $\text{Recall}_{\text{exam}}(h)$ with balancing factor $\beta > 0$. The most common choice is $\beta = 1$ which leads to the *harmonic mean* of precision and recall.

When the intermediate real-valued function $f(\cdot, \cdot)$ is available, four example-based ranking metrics can be defined as well [129]:

- *One-error*:

$$\text{one-error}(f) = \frac{1}{p} \sum_{i=1}^p \mathbb{I}[\arg \max_{y \in \mathcal{Y}} f(\mathbf{x}_i, y) \notin Y_i]$$

The one-error evaluates the fraction of examples whose top-ranked label is not in the relevant label set.

- *Coverage*:

$$\text{coverage}(f) = \frac{1}{p} \sum_{i=1}^p \max_{y \in Y_i} \text{rank}_f(\mathbf{x}_i, y) - 1$$

The coverage evaluates how many steps are needed, on average, to move down the ranked label list so as to cover all the relevant labels of the example.

- *Ranking Loss*:

$$\text{rloss}(f) = \frac{1}{p} \sum_{i=1}^p \frac{1}{|Y_i| |\bar{Y}_i|} |\{(y', y'') \mid f(\mathbf{x}_i, y') \leq f(\mathbf{x}_i, y''), (y', y'') \in Y_i \times \bar{Y}_i\}|$$

The ranking loss evaluates the fraction of reversely ordered label pairs, i.e. an irrelevant label is ranked higher than a relevant label.

- *Average Precision*:

$$\text{avgprec}(f) = \frac{1}{p} \sum_{i=1}^p \frac{1}{|Y_i|} \sum_{y \in Y_i} \frac{|\{y' \mid \text{rank}_f(\mathbf{x}_i, y') \leq \text{rank}_f(\mathbf{x}_i, y), y' \in Y_i\}|}{\text{rank}_f(\mathbf{x}_i, y)}$$

The average precision evaluates the average fraction of relevant labels ranked higher than a particular label $y \in Y_i$.

For *one-error*, *coverage* and *ranking loss*, the smaller the metric value the better the system's performance, with optimal value of $\frac{1}{p} \sum_{i=1}^p |Y_i| - 1$ for *coverage* and 0 for *one-error* and *ranking loss*. For the other example-based multi-label metrics, the larger the metric value the better the system's performance, with optimal value of 1.

3) *Label-based Metrics*: For the j -th class label y_j , four basic quantities characterizing the binary classification performance on this label can be defined based on $h(\cdot)$:

$$TP_j = |\{\mathbf{x}_i \mid y_j \in Y_i \wedge y_j \in h(\mathbf{x}_i), 1 \leq i \leq p\}|; \quad FP_j = |\{\mathbf{x}_i \mid y_j \notin Y_i \wedge y_j \in h(\mathbf{x}_i), 1 \leq i \leq p\}|$$

$$TN_j = |\{\mathbf{x}_i \mid y_j \notin Y_i \wedge y_j \notin h(\mathbf{x}_i), 1 \leq i \leq p\}|; \quad FN_j = |\{\mathbf{x}_i \mid y_j \in Y_i \wedge y_j \notin h(\mathbf{x}_i), 1 \leq i \leq p\}|$$

In other words, TP_j , FP_j , TN_j and FN_j represent the number of *true positive*, *false positive*, *true negative*, and *false negative* test examples with respect to y_j . According to the above definitions, $TP_j + FP_j + TN_j + FN_j = p$ naturally holds.

Based on the above four quantities, most of the binary classification metrics can be derived accordingly. Let $B(TP_j, FP_j, TN_j, FN_j)$ represent some specific binary classification metric ($B \in \{Accuracy, Precision, Recall, F^\beta\}^4$), the label-based classification metrics can be obtained in either of the following modes [157]:

- *Macro-averaging*:

$$B_{\text{macro}}(h) = \frac{1}{q} \sum_{j=1}^q B(TP_j, FP_j, TN_j, FN_j)$$

- *Micro-averaging*:

$$B_{\text{micro}}(h) = B\left(\sum_{j=1}^q TP_j, \sum_{j=1}^q FP_j, \sum_{j=1}^q TN_j, \sum_{j=1}^q FN_j\right)$$

Conceptually speaking, macro-averaging and micro-averaging assume “equal weights” for labels and examples respectively. It is not difficult to show that both $Accuracy_{\text{macro}}(h) = Accuracy_{\text{micro}}(h)$ and $Accuracy_{\text{micro}}(h) + hloss(h) = 1$ hold. Furthermore, note that for any $B \in \{Accuracy, Precision, Recall, F^\beta\}$, its macro/micro-averaged version is different to the example-based version as shown in Subsection II-B.2.

When the intermediate real-valued function $f(\cdot, \cdot)$ is available, one label-based ranking metric, i.e. macro-averaged AUC, can be derived as:

$$AUC_{\text{macro}} = \frac{1}{q} \sum_{j=1}^q AUC_j = \frac{1}{q} \sum_{j=1}^q \frac{|\{(\mathbf{x}', \mathbf{x}'') \mid f(\mathbf{x}', y_j) \geq f(\mathbf{x}'', y_j), (\mathbf{x}', \mathbf{x}'') \in \mathcal{Z}_j \times \bar{\mathcal{Z}}_j\}|}{|\mathcal{Z}_j| |\bar{\mathcal{Z}}_j|} \quad (2)$$

⁴For example, $Accuracy(TP_j, FP_j, TN_j, FN_j) = \frac{TP_j + TN_j}{TP_j + FP_j + TN_j + FN_j}$, $Precision(TP_j, FP_j, TN_j, FN_j) = \frac{TP_j}{TP_j + FP_j}$, $Recall(TP_j, FP_j, TN_j, FN_j) = \frac{TP_j}{TP_j + FN_j}$, and $F^\beta(TP_j, FP_j, TN_j, FN_j) = \frac{(1+\beta^2) \cdot TP_j}{(1+\beta^2) \cdot TP_j + \beta^2 \cdot FN_j + FP_j}$.

Here, $\mathcal{Z}_j = \{\mathbf{x}_i \mid y_j \in Y_i, 1 \leq i \leq p\}$ ($\bar{\mathcal{Z}}_j = \{\mathbf{x}_i \mid y_j \notin Y_i, 1 \leq i \leq p\}$) corresponds to the set of test instances with (without) label y_j . The second line of Eq.(2) follows from the close relation between AUC and the Wilcoxon-Mann-Whitney statistic [64]. Correspondingly, the micro-averaged AUC can also be derived as:

$$AUC_{\text{micro}} = \frac{|\{(\mathbf{x}', \mathbf{x}'', y', y'') \mid f(\mathbf{x}', y') \geq f(\mathbf{x}'', y''), (\mathbf{x}', y') \in \mathcal{S}^+, (\mathbf{x}'', y'') \in \mathcal{S}^-\}|}{|\mathcal{S}^+| |\mathcal{S}^-|}$$

Here, $\mathcal{S}^+ = \{(\mathbf{x}_i, y) \mid y \in Y_i, 1 \leq i \leq p\}$ ($\mathcal{S}^- = \{(\mathbf{x}_i, y) \mid y \notin Y_i, 1 \leq i \leq p\}$) corresponds to the set of relevant (irrelevant) instance-label pairs.

For the above label-based multi-label metrics, the larger the metric value the better the system's performance, with optimal value of 1.

4) *Theoretical Results*: Based on the metric definitions, it is obvious that existing multi-label metrics consider the performance from diverse aspects and are thus of different natures. As shown in Section III, most multi-label learning algorithms actually learn from training examples by explicitly or implicitly optimizing one specific metric. In light of fair and honest evaluation, performance of the multi-label learning algorithm should therefore be tested on a broad range of metrics instead of only on the one being optimized. Specifically, recent theoretical studies show that classifiers aim at maximizing the *subset accuracy* would perform rather poor if evaluated in terms of *hamming loss*, and vice versa [36], [37].

As multi-label metrics are usually non-convex and discontinuous, in practice most learning algorithms resort to optimizing (convex) *surrogate* multi-label metrics [110], [111]. Recently, the *consistency* of multi-label learning [52] has been studied, i.e. whether the expected loss of a learned classifier converges to the Bayes loss as the training set size increases. Specifically, a necessary and sufficient condition for consistency of multi-label learning based on surrogate loss functions is given, which is intuitive and can be informally stated as that for a fixed distribution over $\mathcal{X} \times 2^{\mathcal{Y}}$, the set of classifiers yielding optimal surrogate loss must fall in the set of classifiers yielding optimal original multi-label loss.

In addition, by focusing on *ranking loss* and *hamming loss*, it is disclosed that none pairwise

convex surrogate loss defined on label pairs is consistent with the ranking loss and some recent multi-label approaches [65], [66] are inconsistent even for *deterministic* multi-label learning [52].⁵ Interestingly, in contrast to this negative result, a complementary positive result on consistent multi-label learning is reported for ranking loss minimization [35]. By using a reduction to the bipartite ranking problem [90], simple univariate convex surrogate loss (exponential or logistic) defined on single labels is shown to be consistent with the ranking loss with explicit regret bounds and convergence rates.

III. LEARNING ALGORITHMS

A. Simple Categorization

Algorithm development always stands as the core issue of machine learning researches, with multi-label learning being no exception. During the past decade, significant amount of algorithms have been proposed to learning from multi-label data. Considering that it is infeasible to go through all existing algorithms within limited space, in this review we opt for scrutinizing a total of twelve representative multi-label learning algorithms. Here, the *representativeness* of those selected algorithms are maintained with respect to the following criteria: a) Broad spectrum: each algorithm has unique characteristics covering a variety of algorithmic design strategies; b) Primitive impact: most algorithms lead to a number of follow-up or related methods along its line of research; and c) Favorable influence: each algorithm is among the highly-cited works in the multi-label learning field.⁶

As we try to keep the selection less biased with the above criteria, one should notice that the twelve algorithms to be detailed by no means exclude the importance of other methods. Furthermore, for the sake of notational consistency and mathematical rigor, we have chosen to rephrase

⁵Here, deterministic multi-label learning corresponds to the easier learning case where for any instance $\mathbf{x} \in \mathcal{X}$, there exists a label subset $Y \subseteq \mathcal{Y}$ such that the posteriori probability of observing Y given \mathbf{x} is greater than 0.5, i.e. $\mathbb{P}(Y | \mathbf{x}) > 0.5$.

⁶According to Google Scholar statistics (by November 2012), each paper for the twelve algorithms has been cited more than 40 times, with ten of them with citations more than 100 and four of them with citations more than 200.

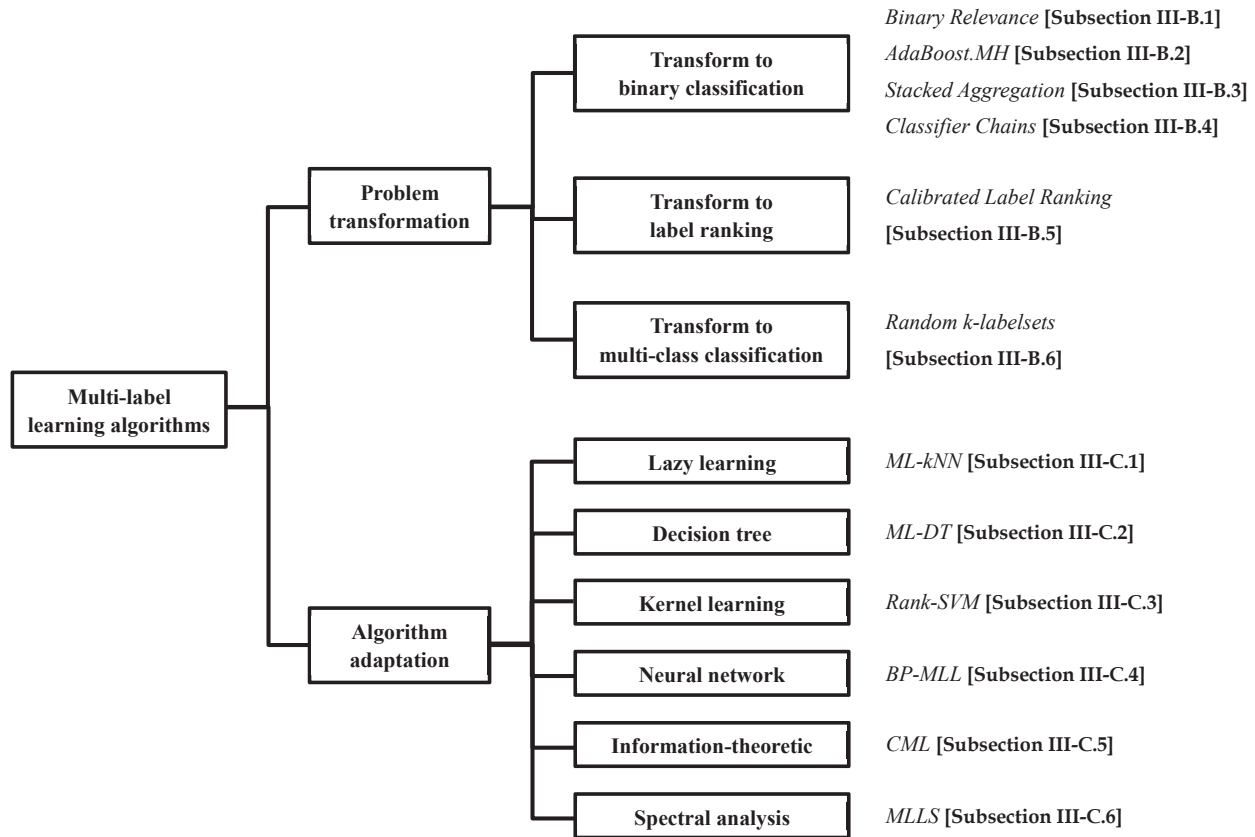


Fig. 2. Categorization of representative multi-label learning algorithms being reviewed.

and present each algorithm under common notations. In this paper, a simple categorization of multi-label learning algorithms is adopted [158]:

- Problem transformation methods*: This category of algorithms tackle multi-label learning problem by transforming it into other well-established learning scenarios. Representative algorithms include first-order approaches Binary Relevance [7] and high-order approach AdaBoost.MH [129], Stacked Aggregation [55], Classifier Chains [122] which transform the task of multi-label learning into the task of binary classification, second-order approach Calibrated Label Ranking [49] which transforms the task of multi-label learning into the task of label ranking, and high-order approach Random k -labelsets [157] which transforms the task of multi-label learning into the task of multi-class classification.
- Algorithm adaptation methods*: This category of algorithms tackle multi-label learning prob-

lem by adapting popular learning techniques to deal with multi-label data. Representative algorithms include first-order approach ML- k NN [201] adapting lazy learning techniques, first-order approach ML-DT [25] adapting decision tree techniques, second-order approach Rank-SVM [42] adapting kernel techniques, second-order approach BP-MLL [190] adapting neural network techniques, second-order approach CML [53] adapting information-theoretic techniques, and high-order approach MLLS [77] adapting spectral analysis techniques.

Briefly, the key philosophy of problem transformation methods is to *fit data to algorithm*, while the key philosophy of algorithm adaptation methods is to *fit algorithm to data*. Fig. 2 summarizes the above-mentioned algorithms to be detailed in the rest of this section.

B. Problem Transformation Methods

1) *Binary Relevance*: The basic idea of this algorithm is to decompose the multi-label learning problem into q *independent* binary classification problems, where each binary classification problem corresponds to a possible label in the label space [7].

Following the notations in Table I, for the j -th class label y_j , Binary Relevance firstly constructs a corresponding *binary* training set by considering the relevance of each training example to y_j :

$$\mathcal{D}_j = \{(\mathbf{x}_i, \phi(Y_i, y_j)) \mid 1 \leq i \leq m\} \quad (3)$$

$$\text{where } \phi(Y_i, y_j) = \begin{cases} +1, & \text{if } y_j \in Y_i \\ -1, & \text{otherwise} \end{cases}$$

After that, some binary learning algorithm \mathcal{B} is utilized to induce a binary classifier $g_j : \mathcal{X} \rightarrow \mathbb{R}$, i.e. $g_j \leftarrow \mathcal{B}(\mathcal{D}_j)$. Therefore, for any multi-label training example (\mathbf{x}_i, Y_i) , instance \mathbf{x}_i will be involved in the learning process of q binary classifiers. For relevant label $y_j \in Y_i$, \mathbf{x}_i is regarded as one *positive* instance in inducing $g_j(\cdot)$; On the other hand, for irrelevant label $y_k \in \bar{Y}_i$, \mathbf{x}_i is regarded as one *negative* instance. The above training strategy is termed as *cross-training* in [7].

For unseen instance \mathbf{x} , Binary Relevance predicts its associated label set Y by querying labeling relevance on each individual binary classifier and then combing relevant labels:

$$Y = \{y_j \mid g_j(\mathbf{x}) > 0, 1 \leq j \leq q\} \quad (4)$$

```

Y=BinaryRelevance( $\mathcal{D}$ ,  $\mathcal{B}$ ,  $\mathbf{x}$ )
1. for  $j = 1$  to  $q$  do
2.   Construct the binary training set  $\mathcal{D}_j$  according to Eq.(3);
3.    $g_j \leftarrow \mathcal{B}(\mathcal{D}_j)$ ;
4. endfor
5. Return  $Y$  according to Eq.(5);

```

Fig. 3. Pseudo-code of Binary Relevance.

Note that when all the binary classifiers yield negative outputs, the predicted label set Y would be empty. To avoid producing empty prediction, the following *T-Criterion* rule can be applied:

$$Y = \{y_j \mid g_j(\mathbf{x}) > 0, 1 \leq j \leq q\} \cup \{y_{j^*} \mid j^* = \arg \max_{1 \leq j \leq q} g_j(\mathbf{x})\} \quad (5)$$

Briefly, when none of the binary classifiers yield positive predictions, T-Criterion rule complements Eq.(4) by including the class label with *greatest* (least negative) output. In addition to T-Criterion, some other rules toward label set prediction based on the outputs of each binary classifier can be found in [7].

Remarks: The pseudo-code of Binary Relevance is summarized in Fig. 3. It is a *first-order* approach which builds classifiers for each label separately and offers the natural opportunity for parallel implementation. The most prominent advantage of Binary Relevance lies in its extremely straightforward way of handling multi-label data (Steps 1-4), which has been employed as the building block of many state-of-the-art multi-label learning techniques [34], [55], [122], [188]. On the other hand, Binary Relevance completely ignores potential correlations among labels, and the binary classifier for each label may suffer from the issue of class-imbalance when q is large and label density (i.e. $LDen(\mathcal{D})$) is low. As shown in Fig. 3, Binary Relevance has computational complexity of $\mathcal{O}(q \cdot \mathcal{F}_{\mathcal{B}}(m, d))$ for training and $\mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d))$ for testing.⁷

⁷In this paper, computational complexity is mainly examined with respect to three factors which are common for all learning algorithms, i.e.: m (number of training examples), d (dimensionality) and q (number of possible class labels). Furthermore, for binary (multi-class) learning algorithm $\mathcal{B}(\mathcal{M})$ embedded in problem transformation methods, we denote its training complexity as $\mathcal{F}_{\mathcal{B}}(m, d)$ ($\mathcal{F}_{\mathcal{M}}(m, d, q)$) and its (per-instance) testing complexity as $\mathcal{F}'_{\mathcal{B}}(d)$ ($\mathcal{F}'_{\mathcal{M}}(d, q)$). All computational complexity results reported in this paper are the worst-case bounds.

2) *AdaBoost.MH*: The basic idea of this algorithm is to transform the multi-label learning problem into *one* binary classification problem, where each multi-label example is mapped into q binary examples [129].

To instantiate the AdaBoost.MH algorithm, the original multi-label training set \mathcal{D} is converted into a binary training set by matching instances with class labels:

$$\mathcal{D}^* = \{([\mathbf{x}_i, y_j], \phi(Y_i, y_j)) \mid 1 \leq i \leq m, 1 \leq j \leq q\} \quad (6)$$

Here, $[\mathbf{x}_i, y_j]$ represents the $(d+1)$ -dimensional feature vector concatenating instance \mathbf{x}_i with label y_j . In other words, each multi-label training example (\mathbf{x}_i, Y_i) gives rise to q binary examples $([\mathbf{x}_i, y_j], \phi(Y_i, y_j))$ ($1 \leq j \leq q$). Based on the converted training set with a total of mq examples, some binary learning algorithm \mathcal{B} is utilized to induce a binary classifier $g^* : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, i.e. $g^* \leftarrow \mathcal{B}(\mathcal{D}^*)$. Obviously, $g^*(\cdot, \cdot)$ acts as the real-valued function $f(\cdot, \cdot)$ in Table I.

For unseen instance \mathbf{x} , AdaBoost.MH predicts its associated label set Y by referring to the confidence of being associated with each class label:

$$Y = \{y_j \mid g^*(\mathbf{x}, y_j) > 0, 1 \leq j \leq q\} \quad (7)$$

In [129], the popular AdaBoost algorithm [48] is chosen to act as the binary learning algorithm \mathcal{B} , where the minimization of AdaBoost on classification error naturally leads to the minimization of AdaBoost.MH on hamming loss.

Remarks: The pseudo-code of AdaBoost.MH is summarized in Fig. 4. Note that the main strategy of this algorithm lies in the transformation process specified in Step 1, and any binary learning algorithms besides AdaBoost could be applied in Step 2 as well [27]. Therefore, we categorize AdaBoost.MH as the *problem transformation* instead of algorithm adaptation method in this paper. It is a *high-order* approach where labels in \mathcal{Y} are treated as appending feature to \mathcal{X} and would be related to each other via the shared instance \mathbf{x} , as far as \mathcal{B} is capable of capturing dependencies among features. Compared to Binary Relevance [7] which works with q independent binary training set (in \mathcal{X}), AdaBoost.MH only deals with one unique transformed binary training set (in $\mathcal{X} \times \mathcal{Y}$). Furthermore, AdaBoost.MH has the advantage of exploiting

```

Y=AdaBoost.MH( $\mathcal{D}$ ,  $\mathcal{B}$ ,  $\mathbf{x}$ )
1. Construct the binary training set  $\mathcal{D}^*$  according to Eq.(6);
2.  $g^* \leftarrow \mathcal{B}(\mathcal{D}^*)$ ;
3. Return  $Y$  according to Eq.(7);

```

Fig. 4. Pseudo-code of AdaBoost.MH.

label correlations rather neatly (Step 1) while may also suffer from the issue of class-imbalance when label density is low. As shown in Fig. 4, AdaBoost.MH has computational complexity of $\mathcal{O}(\mathcal{F}_{\mathcal{B}}(mq, d + 1))$ for training and $\mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d + 1))$ for testing.

3) *Stacked Aggregation*: The basic idea of this algorithm is to transform the multi-label learning problem into *stacked* binary classification problems, where predictions of base binary classifiers are combined with instance space features to induce stacking binary classifiers [55].

Given the training set \mathcal{D} and binary learning algorithm \mathcal{B} , Stacked Aggregation firstly invokes the Binary Relevance training procedure to obtain a set of q base binary classifiers $g_j(\cdot)$ ($1 \leq j \leq q$). Based on the learned classifiers, any instance $\mathbf{z} \in \mathcal{X}$ can be augmented by a q -dimensional binary feature vector as $\nu(\mathbf{z}) = (\nu_1(\mathbf{z}), \nu_2(\mathbf{z}), \dots, \nu_q(\mathbf{z}))^\top$. Here, $\nu_j(\mathbf{z})$ takes the value of +1 if $g_j(\mathbf{z}) > 0$, and -1 otherwise.⁸

For the j -th class label y_j , Stacked Aggregation then constructs a corresponding stacking *binary* training set as follows:

$$\tilde{\mathcal{D}}_j = \{([\mathbf{x}_i, \nu(\mathbf{x}_i)], \phi(Y_i, y_j)) \mid 1 \leq i \leq m\} \quad (8)$$

After that, the binary learning algorithm \mathcal{B} used in base layer is again utilized to induce a stacking layer classifier $\tilde{g}_j : \mathcal{X} \times \{-1, +1\}^q \rightarrow \mathbb{R}$, i.e. $\tilde{g}_j \leftarrow \mathcal{B}(\tilde{\mathcal{D}}_j)$. Therefore, the stacking strategy amounts to augmenting each training example with q additional labeling features and then invoking the Binary Relevance training procedure once more.

For unseen instance \mathbf{x} , Stacked Aggregation predicts its associated label set Y by aggregating

⁸Similar to the T-Criterion, when all the binary classifiers yield negative outputs on \mathbf{x} , the binary bit in $\nu(\mathbf{x})$ corresponding to the least negative output is set to +1 [55].

```

Y=StackedAggregation( $\mathcal{D}$ ,  $\mathcal{B}$ ,  $\mathbf{x}$ )
1. for  $j = 1$  to  $q$  do
2.   Construct the base binary training set  $\mathcal{D}_j$  according to Eq.(3);
3.    $g_j \leftarrow \mathcal{B}(\mathcal{D}_j)$ ;
4. endfor
5. for  $j = 1$  to  $q$  do
6.   Construct the stacking binary training set  $\tilde{\mathcal{D}}_j$  according to Eq.(8);
7.    $\tilde{g}_j \leftarrow \mathcal{B}(\tilde{\mathcal{D}}_j)$ ;
8. endfor
9. Return  $Y$  according to Eq.(9);

```

Fig. 5. Pseudo-code of Stacked Aggregation.

the predictions of base and stacking classifiers:

$$Y = \{y_j \mid \tilde{g}_j([\mathbf{x}, \nu(\mathbf{x})]) > 0, 1 \leq j \leq q\} \quad (9)$$

Remarks: The pseudo-code of Stacked Aggregation is summarized in Fig. 5. It is a *high-order* approach where labeling predictions derived in base layer (Steps 1-4) are combined with instance space features to yield stacking layer predictions (Steps 5-8). Obviously, Stacked Aggregation shares similar advantages and disadvantages as Binary Relevance [7]. The major difference lies in that Stacked Aggregation does exploit correlations between labels while the risk of overfitting could be high due to the nature of stacked generalization [16], [147], [173]. Remedies to overfitting could be introducing label independent models [105] or pruning input features [151] in the stacking layer. As shown in Fig. 5, Stacked Aggregation has computational complexity of $\mathcal{O}(q \cdot \mathcal{F}_{\mathcal{B}}(m, d + q))$ for training and $\mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d + q))$ for testing.

4) *Classifier Chains*: The basic idea of this algorithm is to transform the multi-label learning problem into a *chain* of binary classification problems, where subsequent binary classifiers in the chain is built upon the predictions of preceding ones [122], [123].

For q possible class labels $\{y_1, y_2, \dots, y_q\}$, let $\tau : \{1, \dots, q\} \rightarrow \{1, \dots, q\}$ be a permutation function which is used to specify an *ordering* over them, i.e. $y_{\tau(1)} \succ y_{\tau(2)} \succ \dots \succ y_{\tau(q)}$. For the j -th label $y_{\tau(j)}$ ($1 \leq j \leq q$) in the ordered list, a corresponding binary training set is constructed by appending each instance with its relevance to those labels preceding $y_{\tau(j)}$:

$$\mathcal{D}_{\tau(j)} = \{([\mathbf{x}_i, \mathbf{pre}_{\tau(j)}^i], \phi(Y_i, y_{\tau(j)})) \mid 1 \leq i \leq m\} \quad (10)$$

where $\mathbf{pre}_{\tau(j)}^i = (\phi(Y_i, y_{\tau(1)}), \dots, \phi(Y_i, y_{\tau(j-1)}))^\top$

Here, $\mathbf{pre}_{\tau(j)}^i$ represents the binary assignment of those labels preceding $y_{\tau(j)}$ on \mathbf{x}_i (specifically, $\mathbf{pre}_{\tau(1)}^i = \emptyset$).⁹ After that, some binary learning algorithm \mathcal{B} is utilized to induce a binary classifier $g_{\tau(j)} : \mathcal{X} \times \{-1, +1\}^{j-1} \rightarrow \mathbb{R}$, i.e. $g_{\tau(j)} \leftarrow \mathcal{B}(\mathcal{D}_{\tau(j)})$. In other words, $g_{\tau(j)}(\cdot)$ determines whether $y_{\tau(j)}$ is a relevant label or not.

For unseen instance \mathbf{x} , its associated label set Y is predicted by traversing the classifier chain iteratively. Let $\lambda_{\tau(j)}^{\mathbf{x}} \in \{-1, +1\}$ represent the predicted binary assignment of $y_{\tau(j)}$ on \mathbf{x} , which are recursively derived as follows:

$$\begin{aligned} \lambda_{\tau(1)}^{\mathbf{x}} &= \text{sign} [g_{\tau(1)}(\mathbf{x})] \\ \lambda_{\tau(j)}^{\mathbf{x}} &= \text{sign} [g_{\tau(j)}([\mathbf{x}, \lambda_{\tau(1)}^{\mathbf{x}}, \dots, \lambda_{\tau(j-1)}^{\mathbf{x}}])] \quad (2 \leq j \leq q) \end{aligned} \quad (11)$$

Here, $\text{sign}[a]$ returns $+1$ if $a > 0$ and -1 otherwise. Accordingly, the predicted label set corresponds to:

$$Y = \{y_{\tau(j)} \mid \lambda_{\tau(j)}^{\mathbf{x}} = +1, 1 \leq j \leq q\} \quad (12)$$

It is obvious that for the classifier chain obtained as above, its effectiveness is largely affected by the ordering specified by τ . To account for the effect of ordering, an *Ensemble of Classifier Chains* can be built with n random permutations over the label space, i.e. $\tau^{(1)}, \tau^{(2)}, \dots, \tau^{(n)}$. For each permutation $\tau^{(r)}$ ($1 \leq r \leq n$), instead of inducing one classifier chain by applying $\tau^{(r)}$ directly on the original training set \mathcal{D} , a modified training set $\mathcal{D}^{(r)}$ is used by sampling \mathcal{D} *without* replacement ($|\mathcal{D}^{(r)}| = 0.67 \cdot |\mathcal{D}|$) [122] or *with* replacement ($|\mathcal{D}^{(r)}| = |\mathcal{D}|$) [123].

Remarks: The pseudo-code of Classifier Chains is summarized in Fig. 6. It is a *high-order* approach which considers correlations among labels in a random manner. Compared to Binary Relevance [7], Classifier Chains has the advantage of exploiting label correlations while

⁹In Classifier Chains [122], [123], binary assignment is represented by 0 and 1. Without loss of generality, binary assignment is represented by -1 and +1 in this paper for notational consistency.

```

Y=ClassifierChains( $\mathcal{D}$ ,  $\mathcal{B}$ ,  $\tau$ ,  $\mathbf{x}$ )
1. for  $j = 1$  to  $q$  do
2.   Construct the chaining binary training set  $\mathcal{D}_{\tau(j)}$  according to Eq.(10);
3.    $g_{\tau(j)} \leftarrow \mathcal{B}(\mathcal{D}_{\tau(j)})$ ;
4. endfor
5. Return  $Y$  according to Eq.(12) (in conjunction with Eq.(11));

```

Fig. 6. Pseudo-code of Classifier Chains.

loses the opportunity of parallel implementation due to its chaining property. Furthermore, both Classifier Chains and Stacked Aggregation [55] augment instance space with extra features during the training phase, where the former sets them according to ground-truth labeling (i.e. $\text{pre}_{\tau(j)}^i$ in Eq.(10)) while the latter sets them according to classifier predictions (i.e. $\nu(\mathbf{x}_i)$ in Eq.(8)). Instead of keeping extra features binary-valued, another possibility is to set them to the classifier's probabilistic outputs when the model returned by \mathcal{B} (e.g. Naive Bayes) is capable of yielding posteriori probability [34], [183]. As shown in Fig. 6, Classifier Chains has computational complexity of $\mathcal{O}(q \cdot \mathcal{F}_{\mathcal{B}}(m, d + q))$ for training and $\mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d + q))$ for testing.

5) *Calibrated Label Ranking*: The basic idea of this algorithm is to transform the multi-label learning problem into the *label ranking* problem, where ranking among labels is fulfilled by techniques of pairwise comparison [49].

For q possible class labels $\{y_1, y_2, \dots, y_q\}$, a total of $q(q - 1)/2$ binary classifiers can be generated by pairwise comparison, one for each label pair (y_j, y_k) ($1 \leq j < k \leq q$). Concretely, for each label pair (y_j, y_k) , pairwise comparison firstly constructs a corresponding binary training set by considering the *relative* relevance of each training example to y_j and y_k :

$$\mathcal{D}_{jk} = \{(\mathbf{x}_i, \psi(Y_i, y_j, y_k)) \mid \phi(Y_i, y_j) \neq \phi(Y_i, y_k), 1 \leq i \leq m\} \quad (13)$$

$$\text{where } \psi(Y_i, y_j, y_k) = \begin{cases} +1, & \text{if } \phi(Y_i, y_j) = +1 \text{ and } \phi(Y_i, y_k) = -1 \\ -1, & \text{if } \phi(Y_i, y_j) = -1 \text{ and } \phi(Y_i, y_k) = +1 \end{cases}$$

In other words, only instances with distinct relevance to y_j and y_k will be included in \mathcal{D}_{jk} . After that, some binary learning algorithm \mathcal{B} is utilized to induce a binary classifier $g_{jk} : \mathcal{X} \rightarrow \mathbb{R}$, i.e. $g_{jk} \leftarrow \mathcal{B}(\mathcal{D}_{jk})$. Therefore, for any multi-label training example (\mathbf{x}_i, Y_i) , instance \mathbf{x}_i will

be involved in the learning process of $|Y_i||\bar{Y}_i|$ binary classifiers. For any instance $\mathbf{x} \in \mathcal{X}$, the learning system votes for y_j if $g_{jk}(\mathbf{x}) > 0$ and y_k otherwise.

For unseen instance \mathbf{x} , Calibrated Label Ranking firstly feeds it to the $q(q-1)/2$ trained binary classifiers to obtain the overall votes on each possible class label:

$$\zeta(\mathbf{x}, y_j) = \sum_{k=1}^{j-1} \llbracket g_{kj}(\mathbf{x}) \leq 0 \rrbracket + \sum_{k=j+1}^q \llbracket g_{jk}(\mathbf{x}) > 0 \rrbracket \quad (1 \leq j \leq q) \quad (14)$$

Based on the above definition, it is not difficult to verify that $\sum_{j=1}^q \zeta(\mathbf{x}, y_j) = q(q-1)/2$. Here, labels in \mathcal{Y} can be ranked according to their respective votes (ties are broken arbitrarily).

Thereafter, some thresholding function should be further specified to bipartition the list of ranked labels into relevant and irrelevant label set. To achieve this within the pairwise comparison framework, Calibrated Label Ranking incorporates a *virtual label* y_V into each multi-label training example (\mathbf{x}_i, Y_i) . Conceptually speaking, the virtual label serves as an *artificial splitting point* between \mathbf{x}_i 's relevant and irrelevant labels [9]. In other words, y_V is considered to be ranked lower than $y_j \in Y_i$ while ranked higher than $y_k \in \bar{Y}_i$.

In addition to the original $q(q-1)/2$ binary classifiers, q auxiliary binary classifiers will be induced, one for each new label pair (y_j, y_V) ($1 \leq j \leq q$). Similar to Eq.(13), a binary training set corresponding to (y_j, y_V) can be constructed as follows:

$$\mathcal{D}_{jV} = \{(\mathbf{x}_i, \varphi(Y_i, y_j, y_V)) \mid 1 \leq i \leq m\} \quad (15)$$

$$\text{where } \varphi(Y_i, y_j, y_V) = \begin{cases} +1, & \text{if } y_j \in Y_i \\ -1, & \text{otherwise} \end{cases}$$

Based on this, the binary learning algorithm \mathcal{B} is utilized to induce a binary classifier corresponding to the virtual label $g_{jV} : \mathcal{X} \rightarrow \mathbb{R}$, i.e. $g_{jV} \leftarrow \mathcal{B}(\mathcal{D}_{jV})$. After that, the overall votes specified in Eq.(14) will be updated with the newly induced classifiers:

$$\zeta^*(\mathbf{x}, y_j) = \zeta(\mathbf{x}, y_j) + \llbracket g_{jV}(\mathbf{x}) > 0 \rrbracket \quad (1 \leq j \leq q) \quad (16)$$

Furthermore, the overall votes on virtual label can be computed as:

$$\zeta^*(\mathbf{x}, y_V) = \sum_{j=1}^q \llbracket g_{jV}(\mathbf{x}) \leq 0 \rrbracket \quad (17)$$

```

Y=CalibratedLabelRanking( $\mathcal{D}$ ,  $\mathcal{B}$ ,  $\mathbf{x}$ )
1. for  $j = 1$  to  $q - 1$  do
2.   for  $k = j + 1$  to  $q$  do
3.     Construct the binary training set  $\mathcal{D}_{jk}$  according to Eq.(13);
4.      $g_{jk} \leftarrow \mathcal{B}(\mathcal{D}_{jk})$ ;
5.   endfor
6. endfor
7. for  $j = 1$  to  $q$  do
8.   Construct the binary training set  $\mathcal{D}_{jV}$  according to Eq.(15);
9.    $g_{jV} \leftarrow \mathcal{B}(\mathcal{D}_{jV})$ ;
10. endfor
11. Return  $Y$  according to Eq.(18) (in conjunction with Eqs.(14)-(17));

```

Fig. 7. Pseudo-code of Calibrated Label Ranking.

Therefore, the predicted label set for unseen instance \mathbf{x} corresponds to:

$$Y = \{y_j \mid \zeta^*(\mathbf{x}, y_j) > \zeta^*(\mathbf{x}, y_V), 1 \leq j \leq q\} \quad (18)$$

By comparing Eq.(15) to Eq.(3), it is obvious that the training set \mathcal{D}_{jV} employed by Calibrated Label Ranking is identical the training set \mathcal{D}_j employed by Binary Relevance [7]. Therefore, Calibrated Label Ranking can be regarded as an enriched version of pairwise comparison, where the routine $q(q - 1)/2$ binary classifiers are enlarged with the q binary classifiers of Binary Relevance to facilitate learning [49].

Remarks: The pseudo-code of Calibrated Label Ranking is summarized in Fig. 7. It is a *second-order* approach which builds classifiers for any pair of class labels. Compared to previously introduced algorithms [7], [55], [122], [129] which construct binary classifiers in a *one-vs-rest* manner, Calibrated Label Ranking constructs binary classifiers (except those for virtual label) in a *one-vs-one* manner and thus has the advantage of mitigating the negative influence of the class-imbalance issue. On the other hand, the number of binary classifiers constructed by Calibrated Label Ranking grows from linear scale to *quadratic scale* in terms of the number class labels (i.e. q). Improvements on Calibrated Label Ranking mostly focus on reducing the quadratic number of classifiers to be queried in testing phase by exact pruning [100] or approximate pruning [54], [101], [102]. By exploiting idiosyncrasy of the underlying binary learning algorithm \mathcal{B} , such as dual representation for Perceptron [99], the quadratic number of

classifiers can be induced more efficiently in training phase [98]. As shown in Fig. 7, Calibrated Label Ranking has computational complexity of $\mathcal{O}(q^2 \cdot \mathcal{F}_{\mathcal{B}}(m, d))$ for training and $\mathcal{O}(q^2 \cdot \mathcal{F}'_{\mathcal{B}}(d))$ for testing.

6) *Random k-Labelsets*: The basic idea of this algorithm is to transform the multi-label learning problem into an ensemble of *multi-class classification* problems, where each component learner in the ensemble targets a random subset of \mathcal{Y} upon which a multi-class classifier is induced by the Label Powerset (LP) techniques [155], [157].

LP is a straightforward approach to transforming multi-label learning problem into multi-class (single-label) classification problem. Let $\sigma_{\mathcal{Y}} : 2^{\mathcal{Y}} \rightarrow \mathbb{N}$ be some *injective* function mapping from the power set of \mathcal{Y} to natural numbers, and $\sigma_{\mathcal{Y}}^{-1}$ be the corresponding *inverse* function. In the training phase, LP firstly converts the original multi-label training set \mathcal{D} into the following multi-class training set by treating every distinct label set appearing in \mathcal{D} as a *new class*:

$$\mathcal{D}_{\mathcal{Y}}^{\dagger} = \{(\mathbf{x}_i, \sigma_{\mathcal{Y}}(Y_i)) \mid 1 \leq i \leq m\} \quad (19)$$

where the set of new classes covered by $\mathcal{D}_{\mathcal{Y}}^{\dagger}$ corresponds to:

$$\Gamma(\mathcal{D}_{\mathcal{Y}}^{\dagger}) = \{\sigma_{\mathcal{Y}}(Y_i) \mid 1 \leq i \leq m\} \quad (20)$$

Obviously, $|\Gamma(\mathcal{D}_{\mathcal{Y}}^{\dagger})| \leq \min(m, 2^{|\mathcal{Y}|})$ holds here. After that, some multi-class learning algorithm \mathcal{M} is utilized to induce a multi-class classifier $g_{\mathcal{Y}}^{\dagger} : \mathcal{X} \rightarrow \Gamma(\mathcal{D}_{\mathcal{Y}}^{\dagger})$, i.e. $g_{\mathcal{Y}}^{\dagger} \leftarrow \mathcal{M}(\mathcal{D}_{\mathcal{Y}}^{\dagger})$. Therefore, for any multi-label training example (\mathbf{x}_i, Y_i) , instance \mathbf{x}_i will be re-assigned with the newly mapped single-label $\sigma_{\mathcal{Y}}(Y_i)$ and then participates in multi-class classifier induction.

For unseen instance \mathbf{x} , LP predicts its associated label set Y by firstly querying the prediction of multi-class classifier and then mapping it back to the power set of \mathcal{Y} :

$$Y = \sigma_{\mathcal{Y}}^{-1}(g_{\mathcal{Y}}^{\dagger}(\mathbf{x})) \quad (21)$$

Unfortunately, LP has two major limitations in terms of practical feasibility: a) Incompleteness: as shown in Eqs.(20) and (21), LP is confined to predict label sets appearing in the training set, i.e. unable to generalize to those outside $\{Y_i \mid 1 \leq i \leq m\}$; b) Inefficiency: when \mathcal{Y} is large,

there might be too many newly mapped classes in $\Gamma(\mathcal{D}_y^\dagger)$, leading to overly high complexity in training $g_y^\dagger(\cdot)$ and extremely few training examples for some newly mapped classes.

To keep LP's simplicity while overcoming its two major drawbacks, Random k -Labelsets chooses to combine ensemble learning [38], [200] with LP to learn from multi-label data. The key strategy is to invoke LP only on random k -labelsets (size- k subset in \mathcal{Y}) to guarantee computational efficiency, and then ensemble a number of LP classifiers to achieve predictive completeness.

Let \mathcal{Y}^k represent the collection of all possible k -labelsets in \mathcal{Y} , where the l -th k -labelset is denoted as $\mathcal{Y}^k(l)$, i.e. $\mathcal{Y}^k(l) \subseteq \mathcal{Y}$, $|\mathcal{Y}^k(l)| = k$, $1 \leq l \leq \binom{q}{k}$. Similar to Eq.(19), a multi-class training set can be constructed as well by shrinking the original label space \mathcal{Y} into $\mathcal{Y}^k(l)$:

$$\mathcal{D}_{\mathcal{Y}^k(l)}^\dagger = \{(\mathbf{x}_i, \sigma_{\mathcal{Y}^k(l)}(Y_i \cap \mathcal{Y}^k(l))) \mid 1 \leq i \leq m\} \quad (22)$$

where the set of new classes covered by $\mathcal{D}_{\mathcal{Y}^k(l)}^\dagger$ corresponds to:

$$\Gamma(\mathcal{D}_{\mathcal{Y}^k(l)}^\dagger) = \{\sigma_{\mathcal{Y}^k(l)}(Y_i \cap \mathcal{Y}^k(l)) \mid 1 \leq i \leq m\}$$

After that, the multi-class learning algorithm \mathcal{M} is utilized to induce a multi-class classifier $g_{\mathcal{Y}^k(l)}^\dagger : \mathcal{X} \rightarrow \Gamma(\mathcal{D}_{\mathcal{Y}^k(l)}^\dagger)$, i.e. $g_{\mathcal{Y}^k(l)}^\dagger \leftarrow \mathcal{M}(\mathcal{D}_{\mathcal{Y}^k(l)}^\dagger)$.

To create an ensemble with n component classifiers, Random k -Labelsets invokes LP on n random k -labelsets $\mathcal{Y}^k(l_r)$ ($1 \leq r \leq n$) each leading to a multi-class classifier $g_{\mathcal{Y}^k(l_r)}^\dagger(\cdot)$. For unseen instance \mathbf{x} , the following two quantities are calculated for each class label:

$$\begin{aligned} \tau(\mathbf{x}, y_j) &= \sum_{r=1}^n \mathbb{I}[y_j \in \mathcal{Y}^k(l_r)] \quad (1 \leq j \leq q) \\ \mu(\mathbf{x}, y_j) &= \sum_{r=1}^n \mathbb{I}\left[y_j \in \sigma_{\mathcal{Y}^k(l_r)}^{-1}\left(g_{\mathcal{Y}^k(l_r)}^\dagger(\mathbf{x})\right)\right] \quad (1 \leq j \leq q) \end{aligned} \quad (23)$$

Here, $\tau(\mathbf{x}, y_j)$ counts the *maximum* number of votes that y_j can be received from the ensemble, while $\mu(\mathbf{x}, y_j)$ counts the *actual* number of votes that y_j receives from the ensemble. Accordingly, the predicted label set corresponds to:

$$Y = \{y_j \mid \mu(\mathbf{x}, y_j) / \tau(\mathbf{x}, y_j) > 0.5, 1 \leq j \leq q\} \quad (24)$$

```

Y=Randomk-Labelsets( $\mathcal{D}$ ,  $\mathcal{M}$ ,  $k$ ,  $n$ ,  $\mathbf{x}$ )
1. for  $r = 1$  to  $n$  do
2.   Randomly choose a  $k$ -labelset  $\mathcal{Y}^k(l_r) \subseteq \mathcal{Y}$  with  $|\mathcal{Y}^k(l_r)| = k$ ;
3.   Construct the multi-class training set  $\mathcal{D}_{\mathcal{Y}^k(l_r)}^\dagger$  according to Eq.(22);
4.    $g_{\mathcal{Y}^k(l_r)}^\dagger \leftarrow \mathcal{M}(\mathcal{D}_{\mathcal{Y}^k(l_r)}^\dagger)$ ;
5. endfor
6. Return  $Y$  according to Eq.(24) (in conjunction with Eq.(23));

```

Fig. 8. Pseudo-code of Random k -Labelsets.

In other words, when the actual number of votes exceeds *half* of the maximum number of votes, y_j is regarded to be relevant. For an ensemble created by n k -labelsets, the maximum number of votes on each label is nk/q on average. A rule-of-thumb setting for Random k -Labelsets is $k = 3$ and $n = 2q$ [155], [157].

Remarks: The pseudo-code of Random k -Labelsets is summarized in Fig. 8. It is a *high-order* approach where the degree of label correlations is controlled by the size of k -labelsets. In addition to use k -labelset, another way to improve LP is to prune distinct label set in \mathcal{D} appearing less than a pre-specified counting threshold [119], [121]. Although Random k -Labelsets embeds ensemble learning as its *inherent* part to amend LP's major drawbacks, ensemble learning could be employed as a *meta-level* strategy to facilitate multi-label learning by encompassing homogeneous [122], [131] or heterogeneous [127], [142], [145] component multi-label learners. As shown in Fig. 8, Random k -Labelsets has computational complexity of $\mathcal{O}(n \cdot \mathcal{F}_{\mathcal{M}}(m, d, 2^k))$ for training and $\mathcal{O}(n \cdot \mathcal{F}'_{\mathcal{M}}(d, 2^k))$ for testing.

C. Algorithm Adaptation Methods

1) *Multi-Label k -Nearest Neighbor (ML- k NN)*: The basic idea of this algorithm is to adapt *k-nearest neighbor* techniques to deal with multi-label data, where maximum a posteriori (MAP) rule is utilized to make prediction by reasoning with the labeling information embodied in the neighbors [189], [191].

For unseen instance \mathbf{x} , let $\mathcal{N}(\mathbf{x})$ represent the set of its k nearest neighbors identified in \mathcal{D} . Generally, similarity between instances is measured with the Euclidean distance. For the j -th

class label, ML- k NN chooses to calculate the following statistics:

$$C_j = \sum_{(\mathbf{x}^*, Y^*) \in \mathcal{N}(\mathbf{x})} \llbracket y_j \in Y^* \rrbracket \quad (25)$$

Namely, C_j records the number of \mathbf{x} 's neighbors with label y_j .

Let H_j be the event that \mathbf{x} has label y_j , and $\mathbb{P}(H_j | C_j)$ represents the *posterior probability* that H_j holds under the condition that \mathbf{x} has exactly C_j neighbors with label y_j . Correspondingly, $\mathbb{P}(\neg H_j | C_j)$ represents the posterior probability that H_j doesn't hold under the same condition. According to the MAP rule, the predicted label set is determined by deciding whether $\mathbb{P}(H_j | C_j)$ is greater than $\mathbb{P}(\neg H_j | C_j)$ or not:

$$Y = \{y_j | \mathbb{P}(H_j | C_j) / \mathbb{P}(\neg H_j | C_j) > 1, 1 \leq j \leq q\} \quad (26)$$

Based on Bayes theorem, we have:

$$\frac{\mathbb{P}(H_j | C_j)}{\mathbb{P}(\neg H_j | C_j)} = \frac{\mathbb{P}(H_j) \cdot \mathbb{P}(C_j | H_j)}{\mathbb{P}(\neg H_j) \cdot \mathbb{P}(C_j | \neg H_j)} \quad (27)$$

Here, $\mathbb{P}(H_j)$ ($\mathbb{P}(\neg H_j)$) represents the *prior probability* that H_j holds (doesn't hold). Furthermore, $\mathbb{P}(C_j | H_j)$ ($\mathbb{P}(C_j | \neg H_j)$) represents the *likelihood* that \mathbf{x} has exactly C_j neighbors with label y_j when H_j holds (doesn't hold). As shown in Eqs.(26) and (27), it suffices to estimate the prior probabilities as well as likelihoods for making predictions.

ML- k NN fulfills the above task via the *frequency counting* strategy. Firstly, the prior probabilities are estimated by counting the number training examples associated with each label:

$$\mathbb{P}(H_j) = \frac{s + \sum_{i=1}^m \llbracket y_j \in Y_i \rrbracket}{s \times 2 + m}; \quad \mathbb{P}(\neg H_j) = 1 - \mathbb{P}(H_j) \quad (1 \leq j \leq q) \quad (28)$$

Here, s is a smoothing parameter controlling the effect of *uniform prior* on the estimation. Generally, s takes the value of 1 resulting in Laplace smoothing.

Secondly, the estimation process for likelihoods is somewhat involved. For the j -th class label y_j , ML- k NN maintains two *frequency arrays* κ_j and $\tilde{\kappa}_j$ each containing $k+1$ elements:

$$\begin{aligned} \kappa_j[r] &= \sum_{i=1}^m \llbracket y_j \in Y_i \rrbracket \cdot \llbracket \delta_j(\mathbf{x}_i) = r \rrbracket \quad (0 \leq r \leq k) \\ \tilde{\kappa}_j[r] &= \sum_{i=1}^m \llbracket y_j \notin Y_i \rrbracket \cdot \llbracket \delta_j(\mathbf{x}_i) = r \rrbracket \quad (0 \leq r \leq k) \end{aligned} \quad (29)$$

```

 $Y = \text{ML-}k\text{NN}(\mathcal{D}, k, \mathbf{x})$ 
1. for  $i = 1$  to  $m$  do
2.   Identify  $k$  nearest neighbors  $\mathcal{N}(\mathbf{x}_i)$  for  $\mathbf{x}_i$ ;
3. endfor
4. for  $j = 1$  to  $q$  do
5.   Estimate the prior probabilities  $\mathbb{P}(H_j)$  and  $\mathbb{P}(\neg H_j)$  according to Eq.(28);
6.   Maintain frequency arrays  $\kappa_j$  and  $\tilde{\kappa}_j$  according to Eq.(29);
7. endfor
8. Identify  $k$  nearest neighbors  $\mathcal{N}(\mathbf{x})$  for  $\mathbf{x}$ ;
9. for  $j = 1$  to  $q$  do
10.  Calculate statistic  $C_j$  according to Eq.(25);
11. endfor
12. Return  $Y$  according to Eq.(26) (in conjunction with Eqs.(27), (28) and (30));

```

Fig. 9. Pseudo-code of ML- k NN.

$$\text{where } \delta_j(\mathbf{x}_i) = \sum_{(\mathbf{x}^*, Y^*) \in \mathcal{N}(\mathbf{x}_i)} \llbracket y_j \in Y^* \rrbracket$$

Analogous to C_j in Eq.(25), $\delta_j(\mathbf{x}_i)$ records the number of \mathbf{x}_i 's neighbors with label y_j . Therefore, $\kappa_j[r]$ counts the number of training examples which have label y_j and have exactly r neighbors with label y_j , while $\tilde{\kappa}_j[r]$ counts the number of training examples which don't have label y_j and have exactly r neighbors with label y_j . Afterwards, the likelihoods can be estimated based on elements in κ_j and $\tilde{\kappa}_j$:

$$\begin{aligned} \mathbb{P}(C_j \mid H_j) &= \frac{s + \kappa_j[C_j]}{s \times (k + 1) + \sum_{r=0}^k \kappa_j[r]} & (1 \leq j \leq q, 0 \leq C_j \leq k) \\ \mathbb{P}(C_j \mid \neg H_j) &= \frac{s + \tilde{\kappa}_j[C_j]}{s \times (k + 1) + \sum_{r=0}^k \tilde{\kappa}_j[r]} & (1 \leq j \leq q, 0 \leq C_j \leq k) \end{aligned} \quad (30)$$

Thereafter, by substituting Eq.(28) (prior probabilities) and Eq.(30) (likelihoods) into Eq.(27), the predicted label set in Eq.(26) naturally follows.

Remarks : The pseudo-code of ML- k NN is summarized in Fig. 9. It is a *first-order* approach which reasons the relevance of each label separately. ML- k NN has the advantage of inheriting merits of both lazy learning and Bayesian reasoning: a) decision boundary can be adaptively adjusted due to the *varying neighbors* identified for each unseen instance; b) the class-imbalance issue can be largely mitigated due to the *prior probabilities* estimated for each class label. There are other ways to make use of lazy learning for handling multi-label data, such as combining k NN with ranking aggregation [10], [24], identifying k NN in a label-specific style [70], [163],

expanding k NN to cover the whole training set [23], [82], [175]. Considering that ML- k NN is ignorant of exploiting label correlations, several extensions have been proposed to provide patches to ML- k NN along this direction [22], [180], [181]. As shown in Fig. 9, ML- k NN has computational complexity of $\mathcal{O}(m^2d + qmk)$ for training and $\mathcal{O}(md + qk)$ for testing.

2) *Multi-Label Decision Tree (ML-DT)*: The basic idea of this algorithm is to adopt *decision tree* techniques to deal with multi-label data, where an information gain criterion based on multi-label entropy is utilized to build the decision tree recursively [25].

Information gain is one of the mostly-used splitting criteria for decision tree construction [117]. Given any multi-label data set $\mathcal{T} = \{(\mathbf{x}_i, Y_i) \mid 1 \leq i \leq n\}$ with n examples, the information gain achieved by dividing \mathcal{T} along the l -th feature at splitting value ϑ is:

$$\text{IG}(\mathcal{T}, l, \vartheta) = \text{MLEnt}(\mathcal{T}) - \sum_{\rho \in \{-, +\}} \frac{|\mathcal{T}^\rho|}{|\mathcal{T}|} \cdot \text{MLEnt}(\mathcal{T}^\rho) \quad (31)$$

where $\mathcal{T}^- = \{(\mathbf{x}_i, Y_i) \mid x_{il} \leq \vartheta, 1 \leq i \leq n\}$, $\mathcal{T}^+ = \{(\mathbf{x}_i, Y_i) \mid x_{il} > \vartheta, 1 \leq i \leq n\}$

Namely, \mathcal{T}^- (\mathcal{T}^+) consists of examples with values on the l -th feature less (greater) than ϑ .¹⁰

Starting from the root node (i.e. $\mathcal{T} = \mathcal{D}$), ML-DT identifies the feature and the corresponding splitting value which maximizes the information gain in Eq.(31), and then generates two child nodes with respect to \mathcal{T}^- and \mathcal{T}^+ . The above process is invoked recursively by treating either \mathcal{T}^- or \mathcal{T}^+ as the new root node, and terminates until some stopping criterion \mathcal{C} is met (e.g. size of the child node is less than the pre-specified threshold).

To instantiate ML-DT, the mechanism for computing multi-label entropy, i.e. $\text{MLEnt}(\cdot)$ in Eq.(31), needs to be specified. A straightforward solution is to treat each subset $Y \subseteq \mathcal{Y}$ as a *new class* and then resort to the conventional single-label entropy:

$$\widehat{\text{MLEnt}}(\mathcal{T}) = - \sum_{Y \subseteq \mathcal{Y}} \mathbb{P}(Y) \cdot \log_2(\mathbb{P}(Y)) \quad (32)$$

$$\text{where } \mathbb{P}(Y) = \frac{\sum_{i=1}^n \mathbb{I}[Y_i = Y]}{n}$$

¹⁰Without loss of generality, here we assume that features are real-valued and the data set is bi-partitioned by setting splitting point along each feature. Similar to Eq.(31), information gain with respect to discrete-valued features can be defined as well.

```

Y=ML-DT( $\mathcal{D}$ ,  $\mathcal{C}$ ,  $\mathbf{x}$ )
1. Create a decision tree with root node  $\mathcal{N}$  affiliated with the whole training set ( $\mathcal{T} = \mathcal{D}$ );
2. if stopping criterion  $\mathcal{C}$  is met then
3.   break and go to step 9;
4. else
5.   Identify the feature-value pair  $(l, \vartheta)$  which maximizes Eq.(31);
6.   Set  $\mathcal{T}^-$  and  $\mathcal{T}^+$  according to Eq.(31);
7.   Set  $\mathcal{N}.lsubtree$  and  $\mathcal{N}.rsubtree$  to the decision trees recursively constructed with  $\mathcal{T}^-$  and  $\mathcal{T}^+$  respectively;
8. endif
9. Traverse  $\mathbf{x}$  along the decision tree from the root node until a leaf node is reached;
10. Return  $Y$  according to Eq.(34);

```

Fig. 10. Pseudo-code of ML-DT.

However, as the number of new classes grows exponentially with respect to $|\mathcal{Y}|$, many of them might not even appear in \mathcal{T} and thus only have trivial estimated probability (i.e. $\mathbb{P}(Y) = 0$). To circumvent this issue, ML-DT assumes *independence* among labels and computes the multi-label entropy in a decomposable way:

$$\text{MLEnt}(\mathcal{T}) = \sum_{j=1}^q -p_j \log_2 p_j - (1 - p_j) \log_2 (1 - p_j) \quad (33)$$

$$\text{where } p_j = \frac{\sum_{i=1}^n \mathbb{1}[y_j \in Y_i]}{n}$$

Here, p_j represents the fraction of examples in \mathcal{T} with label y_j . Note that Eq.(33) can be regarded as a simplified version of Eq.(32) under the label independence assumption, and it holds that $\text{MLEnt}(\mathcal{T}) \geq \widehat{\text{MLEnt}}(\mathcal{T})$.

For unseen instance \mathbf{x} , it is fed to the learned decision tree by traversing along the paths until reaching a leaf node affiliated with a number of training examples $\mathcal{T} \subseteq \mathcal{D}$. Then, the predicted label set corresponds to:

$$Y = \{y_j \mid p_j > 0.5, 1 \leq j \leq q\} \quad (34)$$

In other words, if for one leaf node the majority of training examples falling into it have label y_j , any test instance allocated within the same leaf node will regard y_j as its relevant label.

Remarks : The pseudo-code of ML-DT is summarized in Fig. 10. It is a *first-order* approach which assumes label independence in calculating multi-label entropy. One prominent advantage

of ML-DT lies in its high efficiency in inducing the decision tree model from multi-label data. Possible improvements on multi-label decision trees include employing pruning strategy [25] or ensemble learning techniques [56], [85], [193]. As shown in Fig. 10, ML-DT has computational complexity of $\mathcal{O}(mdq)$ for training and $\mathcal{O}(mq)$ for testing.

3) *Ranking Support Vector Machine (Rank-SVM)*: The basic idea of this algorithm is to adapt *maximum margin* strategy to deal with multi-label data, where a set of linear classifiers are optimized to minimize the empirical ranking loss and enabled to handle nonlinear cases with kernel tricks [42].

Let the learning system be composed of q linear classifiers $\mathcal{W} = \{(\mathbf{w}_j, b_j) \mid 1 \leq j \leq q\}$, where $\mathbf{w}_j \in \mathbb{R}^d$ and $b_j \in \mathbb{R}$ are the *weight vector* and *bias* for the j -th class label y_j . Correspondingly, Rank-SVM defines the learning system's *margin* on (\mathbf{x}_i, Y_i) by considering its ranking ability on the example's relevant and irrelevant labels:

$$\min_{(y_j, y_k) \in Y_i \times \bar{Y}_i} \frac{\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k}{\|\mathbf{w}_j - \mathbf{w}_k\|} \quad (35)$$

Here, $\langle \mathbf{u}, \mathbf{v} \rangle$ returns the inner product $\mathbf{u}^\top \mathbf{v}$. Geometrically speaking, for each relevant-irrelevant label pair $(y_j, y_k) \in Y_i \times \bar{Y}_i$, their discrimination boundary corresponds to the hyperplane $\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x} \rangle + b_j - b_k = 0$. Therefore, Eq.(35) considers the *signed* L_2 -distance of \mathbf{x}_i to hyperplanes of every relevant-irrelevant label pair, and then returns the minimum as the margin on (\mathbf{x}_i, Y_i) . Therefore, the learning system's margin on the whole training set \mathcal{D} naturally follows:

$$\min_{(\mathbf{x}_i, Y_i) \in \mathcal{D}} \min_{(y_j, y_k) \in Y_i \times \bar{Y}_i} \frac{\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k}{\|\mathbf{w}_j - \mathbf{w}_k\|} \quad (36)$$

When the learning system is capable of properly ranking every relevant-irrelevant label pair for each training example, Eq.(36) will return *positive* margin. In this ideal case, we can rescale the linear classifiers to ensure: a) $\forall 1 \leq i \leq m$ and $(y_j, y_k) \in Y_i \times \bar{Y}_i$, $\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k > 1$; b) $\exists i^* \in \{1, \dots, m\}$ and $(y_{j^*}, y_{k^*}) \in Y_{i^*} \times \bar{Y}_{i^*}$, $\langle \mathbf{w}_{j^*} - \mathbf{w}_{k^*}, \mathbf{x}_{i^*} \rangle + b_{j^*} - b_{k^*} = 1$. Thereafter, the problem of maximizing the margin in Eq.(36) can be expressed as:

$$\max_{\mathcal{W}} \min_{(\mathbf{x}_i, Y_i) \in \mathcal{D}} \min_{(y_j, y_k) \in Y_i \times \bar{Y}_i} \frac{1}{\|\mathbf{w}_j - \mathbf{w}_k\|^2} \quad (37)$$

subject to : $\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k \geq 1 \quad (1 \leq i \leq m, (y_j, y_k) \in Y_i \times \bar{Y}_i)$

Suppose we have sufficient training examples such that for each label pair (y_j, y_k) ($j \neq k$), there exists $(\mathbf{x}, Y) \in \mathcal{D}$ satisfying $(y_j, y_k) \in Y \times \bar{Y}$. Thus, the objective in Eq.(37) becomes equivalent to $\max_{\mathcal{W}} \min_{1 \leq j < k \leq q} \frac{1}{\|\mathbf{w}_j - \mathbf{w}_k\|^2}$ and the optimization problem can be re-written as:

$$\min_{\mathcal{W}} \max_{1 \leq j < k \leq q} \|\mathbf{w}_j - \mathbf{w}_k\|^2 \quad (38)$$

subject to : $\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k \geq 1 \quad (1 \leq i \leq m, (y_j, y_k) \in Y_i \times \bar{Y}_i)$

To overcome the difficulty brought by the max operator, Rank-SVM chooses to simplify Eq.(38) by approximating the max operator with the sum operator:

$$\min_{\mathcal{W}} \sum_{j=1}^q \|\mathbf{w}_j\|^2 \quad (39)$$

subject to : $\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k \geq 1 \quad (1 \leq i \leq m, (y_j, y_k) \in Y_i \times \bar{Y}_i)$

To accommodate real-world scenarios where constraints in Eq.(39) can not be fully satisfied, *slack variables* can be incorporated into Eq.(39):

$$\min_{\{\mathcal{W}, \Xi\}} \sum_{j=1}^q \|\mathbf{w}_j\|^2 + C \sum_{i=1}^m \frac{1}{|Y_i| |\bar{Y}_i|} \sum_{(y_j, y_k) \in Y_i \times \bar{Y}_i} \xi_{ijk} \quad (40)$$

subject to : $\langle \mathbf{w}_j - \mathbf{w}_k, \mathbf{x}_i \rangle + b_j - b_k \geq 1 - \xi_{ijk}$

$$\xi_{ijk} \geq 0 \quad (1 \leq i \leq m, (y_j, y_k) \in Y_i \times \bar{Y}_i)$$

Here, $\Xi = \{\xi_{ijk} \mid 1 \leq i \leq m, (y_j, y_k) \in Y_i \times \bar{Y}_i\}$ is the set of slack variables. The objective in Eq.(40) consists of two parts balanced by the trade-off parameter C . Specifically, the first part corresponds to the *margin* of the learning system, while the second parts corresponds to the surrogate *ranking loss* of the learning system implemented in hinge form.

Note that Eq.(40) is a standard *quadratic programming* (QP) problem with convex objective and linear constraints, which can be tackled with any off-the-shelf QP solver. Furthermore, to

$$Y = \text{Rank-SVM}(\mathcal{D}, C, \mathbf{x})$$

1. Induce the classification system $\mathcal{W} = \{(\mathbf{w}_j, b_j) \mid 1 \leq j \leq q\}$ by solving the QP problem in Eq.(40);
 2. Induce (\mathbf{w}^*, b^*) for the thresholding function by solving the linear least square problem in Eq.(1);
 3. Return Y according to Eq.(41);
-
-

Fig. 11. Pseudo-code of Rank-SVM.

endow Rank-SVM with nonlinear classification ability, one popular way is to solve Eq.(40) in its dual form via *kernel trick*. More details on the dual formulation can be found in [41].

As discussed in Subsection II-A.3, Rank-SVM employs the stacking-style procedure to set the thresholding function $t(\cdot)$, i.e. $t(\mathbf{x}) = \langle \mathbf{w}^*, \mathbf{f}^*(\mathbf{x}) \rangle + b^*$ with $\mathbf{f}^*(\mathbf{x}) = (f(\mathbf{x}, y_1), \dots, f(\mathbf{x}, y_q))^T$ and $f(\mathbf{x}, y_j) = \langle \mathbf{w}_j, \mathbf{x} \rangle + b_j$. For unseen instance \mathbf{x} , the predicted label set corresponds to:

$$Y = \{y_j \mid \langle \mathbf{w}_j, \mathbf{x} \rangle + b_j > \langle \mathbf{w}^*, \mathbf{f}^*(\mathbf{x}) \rangle + b^*, 1 \leq j \leq q\} \quad (41)$$

Remarks: The pseudo-code of Rank-SVM is summarized in Fig. 11. It is a *second-order* approach which defines the margin over hyperplanes for relevant-irrelevant label pairs. Rank-SVM inherits the benefits of kernel learning to be amenable to nonlinear classification problems, while further improvements can be achieved along several directions. Firstly, as shown in [62], the empirical ranking loss considered in Eq.(40) can be replaced with other loss structures such as hamming loss, which can be cast as a general form of *structured output* classification [149], [150]. Secondly, the thresholding strategy can be accomplished with techniques other than stacking-style procedure [80]. Thirdly, to avoid the problem of kernel selection, multiple kernel learning techniques can be employed to learn from multi-label data [11], [76], [143]. As shown in Fig. 11, let $\mathcal{F}_{\text{QP}}(a, b)$ represent the time complexity for a QP solver to solve Eq.(40) with a variables and b constraints, Rank-SVM has computational complexity of $\mathcal{O}(\mathcal{F}_{\text{QP}}(dq + mq^2, mq^2) + q^2(q + m))$ for training and $\mathcal{O}(dq)$ for testing.

4) *Backpropagation for Multi-Label Learning (BP-MLL)*: The basic idea of this algorithm is to adapt *feed-forward neural networks* to deal with multi-label data, where the error backpropagation strategy is employed to minimize a global error function capturing label correlations [190].

Suppose the feed-forward neural network consists of d input neurons each corresponding to one feature, q output neurons each corresponding to one class label, and a hidden layer with s

hidden neurons. The input layer is fully connected to the hidden layer via a $d \times s$ weight matrix \mathbf{V} with elements v_{lu} ($1 \leq l \leq d, 1 \leq u \leq s$), and the hidden layer is also fully connected to the output layer via an $s \times q$ weight matrix \mathbf{W} with elements w_{uj} ($1 \leq u \leq s, 1 \leq j \leq q$). Furthermore, the bias γ_u (θ_j) of each hidden (output) neuron corresponds to a weight from an *extra* input (hidden) neuron with input value (activation value) fixed at 1.

To capture correlation among labels, BP-MLL defines the global error function which favors networks yielding larger outputs on relevant labels than irrelevant ones:

$$E = \sum_{i=1}^m \frac{1}{|Y_i||\bar{Y}_i|} \cdot \sum_{(y_j, y_k) \in Y_i \times \bar{Y}_i} \exp(-(c_j^i - c_k^i)) \quad (42)$$

Here, c_j^i represents the network's actual output on the i -th instance \mathbf{x}_i at the j -th class label y_j . According to Eq.(42), E can be regarded as the surrogate *ranking loss* of the learning system implemented in exponential form.

BP-MLL updates network weights and biases in incremental mode, i.e. training examples are fed to the network sequentially in one training epoch. The training procedure terminates when the maximum number of training epochs N is reached. With the ‘‘tanh’’ activation function $\eta(a) = (e^a - e^{-a})/(e^a + e^{-a})$, for each training example (\mathbf{x}_i, Y_i) , BP-MLL combines *gradient descent* with the classical *error backpropagation* strategy [67] to update weights and biases:

$$\begin{aligned} \Delta v_{lu} &= \alpha \cdot e_u \cdot a_l; & \Delta w_{uj} &= \alpha \cdot d_j \cdot b_u \\ \Delta \theta_j &= \alpha \cdot d_j; & \Delta \gamma_u &= \alpha \cdot e_u \quad (1 \leq l \leq d, 1 \leq u \leq s, 1 \leq j \leq q) \end{aligned} \quad (43)$$

Here, $\alpha \in [0, 1]$ is the learning rate for neural network training and the other quantities in Eq.(43) correspond to:

$$\begin{aligned} a_l &= x_{il} \\ b_u &= \eta \left(\sum_{l=1}^d a_l v_{lu} + \gamma_u \right) \\ d_j &= \begin{cases} \left(\frac{1}{|Y_i||\bar{Y}_i|} \sum_{y_k \in \bar{Y}_i} \exp(-(c_j^i - c_k^i)) \right) (1 + c_j^i)(1 - c_j^i) & \text{if } y_j \in Y_i \\ \left(-\frac{1}{|Y_i||\bar{Y}_i|} \sum_{y_k \in Y_i} \exp(-(c_k^i - c_j^i)) \right) (1 + c_j^i)(1 - c_j^i) & \text{if } y_j \in \bar{Y}_i \end{cases} \end{aligned} \quad (44)$$

$Y = \text{BP-MLL}(\mathcal{D}, N, s, \alpha, \mathbf{x})$

1. Initialize a feed-forward neural network with weights v_{lu} , w_{uj} , and biases γ_u , θ_j ($1 \leq l \leq d, 1 \leq u \leq s, 1 \leq j \leq q$);
2. **for** $n = 1$ to N **do**
3. **for** $i = 1$ to m **do**
4. Update network weights and biases according to Eqs.(43) and (44);
5. **endfor**
6. **endfor**
7. Induce (\mathbf{w}^*, b^*) for the thresholding function by solving the linear least square problem in Eq.(1);
8. Return Y according to Eq.(45);

Fig. 12. Pseudo-code of BP-MLL.

$$e_u = \left(\sum_{j=1}^q d_j w_{uj} \right) (1 + b_u)(1 - b_u)$$

For unseen instance \mathbf{x} , let c_j be the output of the trained network on \mathbf{x} at the j -th class label y_j , i.e. $c_j = \eta \left(\sum_{u=1}^s w_{uj} \cdot \eta \left(\sum_{l=1}^d v_{lu} x_l + \gamma_u \right) + \theta_j \right)$. Similar to Rank-SVM, let $t(\mathbf{x}) = \langle \mathbf{w}^*, \mathbf{c} \rangle + b^*$ be the stacking-style thresholding function with $\mathbf{c} = (c_1, c_2, \dots, c_q)^\top$. Therefore, the predicted label set corresponds to:

$$Y = \{y_j \mid c_j > \langle \mathbf{w}^*, \mathbf{c} \rangle + b^*, 1 \leq j \leq q\} \quad (45)$$

Remarks: The pseudo-code of BP-MLL is summarized in Fig. 12. It is a *second-order* approach which defines the error function with respect to relevant-irrelevant label pairs. Both BP-MLL and Rank-SVM [42] aim to minimize *surrogate* ranking loss, while the former one is implemented in exponential form and the latter one is implemented in hinge form. BP-MLL has strong ability in handling nonlinear classification cases while may get stuck in local minima due to the gradient descent-based minimization procedure. Direct improvement to BP-MLL can be achieved by modifying the global error function in Eq.(42) [59], while neural networks models other than BP can be applied to multi-label learning as well such as Random Access Memory (RAM) [30], Adaptive Resonance Theory (ART) [128] and Radial Basis Function (RBF) [186]. As shown in Fig. 12, BP-MLL has computational complexity of $\mathcal{O}(s(d+q)mN + q^2(q+m))$ for training and $\mathcal{O}(s(d+q))$ for testing.

5) *Collective Multi-Label Classifier (CML)*: The basic idea of this algorithm is to adapt *maximum entropy* principle to deal with multi-label data, where correlations among labels are encoded as constraints that the resulting distribution must satisfy [53].

For any multi-label example (\mathbf{x}, Y) , let (\mathbf{x}, \mathbf{y}) be the corresponding *random variables* representation using binary label vector $\mathbf{y} = (y_1, y_2, \dots, y_q)^\top \in \{-1, +1\}^q$, whose j -th component indicates whether Y contains the j -th label ($y_j = +1$) or not ($y_j = -1$). Statistically speaking, the task of multi-label learning is equivalent to learn a joint probability distribution $p(\mathbf{x}, \mathbf{y})$.

Let $\mathcal{H}_p(\mathbf{x}, \mathbf{y})$ represent the *information entropy* of (\mathbf{x}, \mathbf{y}) given their distribution $p(\cdot, \cdot)$. The *principle of maximum entropy* [75] assumes that the distribution best modeling the current state of knowledge is the one maximizing $\mathcal{H}_p(\mathbf{x}, \mathbf{y})$ subject to a collection \mathcal{K} of given facts:

$$\begin{aligned} \max_p \quad & \mathcal{H}_p(\mathbf{x}, \mathbf{y}) \\ \text{subject to:} \quad & \mathbb{E}_p[f_k(\mathbf{x}, \mathbf{y})] = F_k \quad (k \in \mathcal{K}) \end{aligned} \quad (46)$$

Generally, the fact is expressed as constraint on the *expectation* of some function over (\mathbf{x}, \mathbf{y}) , i.e. by imposing $\mathbb{E}_p[f_k(\mathbf{x}, \mathbf{y})] = F_k$. Here, $\mathbb{E}_p[\cdot]$ is the expectation operator with respect to $p(\cdot, \cdot)$, while F_k corresponds to the expected value estimated from training set, e.g. $\frac{1}{m} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} f_k(\mathbf{x}, \mathbf{y})$.

Together with the normalization constraint on $p(\cdot, \cdot)$ (i.e. $\mathbb{E}_p[1] = 1$), the constrained optimization problem of Eq.(46) can be carried out with standard *Lagrange Multiplier* techniques. Accordingly, the optimal solution is shown to fall within the *Gibbs* distribution family [3], [106]:

$$p(\mathbf{y} | \mathbf{x}) = \frac{1}{Z_\Lambda(\mathbf{x})} \exp \left(\sum_{k \in \mathcal{K}} \lambda_k \cdot f_k(\mathbf{x}, \mathbf{y}) \right) \quad (47)$$

Here, $\Lambda = \{\lambda_k | k \in \mathcal{K}\}$ is the set of parameters to be determined, and $Z_\Lambda(\mathbf{x})$ is the *partition function* serving as the normalization factor, i.e. $Z_\Lambda(\mathbf{x}) = \sum_{\mathbf{y}} \exp \left(\sum_{k \in \mathcal{K}} \lambda_k \cdot f_k(\mathbf{x}, \mathbf{y}) \right)$.

By assuming Gaussian prior (i.e. $\lambda_k \sim \mathcal{N}(0, \varepsilon^2)$), parameters in Λ can be found by maximizing the following *log-posterior* probability function:

$$\begin{aligned} l(\Lambda | \mathcal{D}) &= \log \left(\prod_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} p(\mathbf{y} | \mathbf{x}) \right) - \sum_{k \in \mathcal{K}} \frac{\lambda_k^2}{2\varepsilon^2} \\ &= \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left(\sum_{k \in \mathcal{K}} \lambda_k \cdot f_k(\mathbf{x}, \mathbf{y}) - \log Z_\Lambda(\mathbf{x}) \right) - \sum_{k \in \mathcal{K}} \frac{\lambda_k^2}{2\varepsilon^2} \end{aligned} \quad (48)$$

Note that Eq.(48) is a convex function over Λ , whose global maximum (though not in closed-form) can be found by any off-the-shelf unconstrained optimization method such as BFGS [12].

```

Y=CML( $\mathcal{D}$ ,  $\varepsilon^2$ ,  $\mathbf{x}$ )
1. for  $l = 1$  to  $d$  do
2.   for  $j = 1$  to  $q$  do
3.     Set constraint  $f_k(\mathbf{x}, \mathbf{y}) = x_l \cdot \llbracket y_j = 1 \rrbracket$  ( $k = (l, j) \in \mathcal{K}_1$ );
4.   endfor
5. endfor
6. for  $j_1 = 1$  to  $q - 1$  do
7.   for  $j_2 = j_1 + 1$  to  $q$  do
8.     Set constraint  $f_k(\mathbf{x}, \mathbf{y}) = \llbracket y_{j_1} = b_1 \rrbracket \cdot \llbracket y_{j_2} = b_2 \rrbracket$  ( $b_1, b_2 \in \{-1, +1\}$ ,  $k = (j_1, j_2, b_1, b_2) \in \mathcal{K}_2$ );
9.   endfor
10. endfor
11. Determine parameters  $\Lambda = \{\lambda_k \mid k \in \mathcal{K}_1 \cup \mathcal{K}_2\}$  by maximizing Eq.(48) (in conjunction with Eq.(49));
12. Return  $Y$  according to Eq.(50);

```

Fig. 13. Pseudo-code of CML.

Generally, gradients of $l(\Lambda \mid \mathcal{D})$ are needed by most numerical methods:

$$\frac{\partial l(\Lambda \mid \mathcal{D})}{\partial \lambda_k} = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left(f_k(\mathbf{x}, \mathbf{y}) - \sum_{\mathbf{y}} f_k(\mathbf{x}, \mathbf{y}) p(\mathbf{y} \mid \mathbf{x}) \right) - \frac{\lambda_k}{\varepsilon^2} \quad (k \in \mathcal{K}) \quad (49)$$

For CML, the set of constraints consists of two parts $\mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2$. Concretely, $\mathcal{K}_1 = \{(l, j) \mid 1 \leq l \leq d, 1 \leq j \leq q\}$ specifies a total of $d \cdot q$ constraints with $f_k(\mathbf{x}, \mathbf{y}) = x_l \cdot \llbracket y_j = 1 \rrbracket$ ($k = (l, j) \in \mathcal{K}_1$). In addition, $\mathcal{K}_2 = \{(j_1, j_2, b_1, b_2) \mid 1 \leq j_1 < j_2 \leq q, b_1, b_2 \in \{-1, +1\}\}$ specifies a total of $4 \cdot \binom{q}{2}$ constraints with $f_k(\mathbf{x}, \mathbf{y}) = \llbracket y_{j_1} = b_1 \rrbracket \cdot \llbracket y_{j_2} = b_2 \rrbracket$ ($k = (j_1, j_2, b_1, b_2) \in \mathcal{K}_2$). Actually, constraints in \mathcal{K} can be specified in other ways yielding variants of CML [53], [203].

For unseen instance \mathbf{x} , the predicted label set corresponds to:

$$Y = \arg \max_{\mathbf{y}} p(\mathbf{y} \mid \mathbf{x}) \quad (50)$$

Note that exact inference with $\arg \max$ is only tractable for small label space. Otherwise, pruning strategies need to be applied to significantly reduce the search space of $\arg \max$, e.g. only considering label sets appearing in the training set [53].

Remarks : The pseudo-code of CML is summarized in Fig. 13. It is a *second-order* approach where correlations between every label pair are considered via constraints in \mathcal{K}_2 . The second-order correlation studied by CML is more general than that of Rank-SVM [42] and BP-MLL [190] as the latter two approaches only consider relevant-irrelevant label pairs. As a conditional random field (CRF) model, CML is interested in using the *conditional* probability distribution

$p(\mathbf{y} | \mathbf{x})$ in Eq.(47) for classification. Interestingly, $p(\mathbf{y} | \mathbf{x})$ can be factored in various ways such as $p(\mathbf{y} | \mathbf{x}) = \prod_{j=1}^q p(y_j | \mathbf{x}, y_1, \dots, y_{j-1})$ where each term in the product can be modeled by one classifier in the classifier chain [34], [122], or $p(\mathbf{y} | \mathbf{x}) = \prod_{j=1}^q p(y_j | \mathbf{x}, \mathbf{pa}_j)$ where each term in the product can be modeled by node y_j and its parents \mathbf{pa}_j in a directed graph [61], [183], [188], and efficient algorithms exist when the directed graph corresponds to multi-dimensional Bayesian network with restricted topology [6], [31], [160]. Directed graphs are also found to be useful for modeling *multiple fault diagnosis* where y_j indicates the good/failing condition of one of the device's components [5], [33]. On the other hand, there have been some multi-label generative models which aim to model the *joint* probability distribution $p(\mathbf{x}, \mathbf{y})$ [104], [159], [167]. As shown in Fig. 13, let $\mathcal{F}_{\text{UNC}}(a, m)$ represent the time complexity for an unconstrained optimization method to solve Eq.(48) with a variables, CML has computational complexity of $\mathcal{O}(\mathcal{F}_{\text{UNC}}(dq + q^2, m))$ for training and $\mathcal{O}((dq + q^2) \cdot 2^q)$ for testing.

6) *Multi-Label Least Squares (MLLS)*: The basic idea of this algorithm is to adapt *spectral analysis* techniques to deal with multi-label data, where correlations among labels are modeled by a common subspace shared by q linear classifiers which are learned by solving a generalized eigenvalue problem [77], [78].

Following the notations in Table I, let the learning system be composed of q linear classifiers defined over an extended instance space with shared features:

$$f(\mathbf{x}, y_j) = (\mathbf{w}_j + \Theta^\top \mathbf{v}_j)^\top \mathbf{x} \quad (1 \leq j \leq q) \quad (51)$$

Here, Θ is an $r \times d$ linear transformation matrix parameterizing the r -dimensional *shared subspace* ($r \ll d$). Note that Θ is common for all labels and assumed to have orthonormal rows, i.e. $\Theta\Theta^\top = \mathbf{I}$ (\mathbf{I} : identity matrix). In addition, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_q] \in \mathbb{R}^{d \times q}$ and $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_q] \in \mathbb{R}^{r \times q}$ are the weight matrices for the *original* and *shared* subspace respectively.

To learn model parameters Θ , \mathbf{W} and \mathbf{V} , MLLS chooses to optimize the following regularized least squares problem:

$$\min_{\{\Theta, \mathbf{W}, \mathbf{V}\}} \sum_{j=1}^q \left(\frac{1}{n} \sum_{i=1}^n L((\mathbf{w}_j + \Theta^\top \mathbf{v}_j)^\top \mathbf{x}_i, \mathbf{Y}_{ij}) + \alpha \|\mathbf{w}_j\|^2 + \beta \|\mathbf{w}_j + \Theta^\top \mathbf{v}_j\|^2 \right) \quad (52)$$

$$\text{subject to: } \Theta\Theta^\top = \mathbf{I}$$

Here, \mathbf{Y} is an $n \times q$ labeling matrix with elements $\mathbf{Y}_{ij} = \phi(Y_i, y_j)$, and $L(\cdot, \cdot)$ is the squared loss function $L(a, b) = (a - b)^2$. The first term in the objective corresponds to *empirical risk* on training examples, while the second and third terms correspond to *regularizers* balanced by trade-off parameters α and β . Specifically, $\alpha\|\mathbf{w}_j\|^2$ controls the degree of information shared among labels while $\beta\|\mathbf{w}_j + \Theta^\top \mathbf{v}_j\|^2$ controls the model complexity.

By renaming $\mathbf{u}_j = \mathbf{w}_j + \Theta^\top \mathbf{v}_j$, Eq.(52) can be re-written in the following compact form:

$$\min_{\{\Theta, \mathbf{U}, \mathbf{V}\}} \frac{1}{n} \|\mathbf{X}\mathbf{U} - \mathbf{Y}\|_F^2 + \alpha \|\mathbf{U} - \Theta^\top \mathbf{V}\|_F^2 + \beta \|\mathbf{U}\|_F^2 \quad (53)$$

$$\text{subject to: } \Theta\Theta^\top = \mathbf{I}$$

Here, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$ is the instance matrix and $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_q] \in \mathbb{R}^{d \times q}$. Furthermore, $\|\cdot\|_F$ returns the matrix Frobenius norm [57].

Regarding the objective in Eq.(53), by fixing Θ and \mathbf{U} , the only term that involves \mathbf{V} is $\|\mathbf{U} - \Theta^\top \mathbf{V}\|_F^2$. Since $\|\mathbf{A}\|_F^2 = \text{trace}(\mathbf{A}^\top \mathbf{A})$ and $\Theta\Theta^\top = \mathbf{I}$, it follows that $\|\mathbf{U} - \Theta^\top \mathbf{V}\|_F^2 = \text{trace}(\mathbf{U}^\top \mathbf{U} - 2\mathbf{U}^\top \Theta^\top \mathbf{V} + \mathbf{V}^\top \mathbf{V})$. Furthermore, since $\frac{\partial \text{trace}(\mathbf{A}^\top \mathbf{A})}{\partial \mathbf{A}} = 2\mathbf{A}$ and $\frac{\partial \text{trace}(\mathbf{B}\mathbf{A})}{\partial \mathbf{A}} = \mathbf{B}^\top$, the derivative of $\|\mathbf{U} - \Theta^\top \mathbf{V}\|_F^2$ with respect to \mathbf{V} turns out to be $2\mathbf{V} - 2\Theta\mathbf{U}$. Setting the derivative to zero yields the optimal solution for \mathbf{V} as follows:

$$\mathbf{V} = \Theta\mathbf{U} \quad (54)$$

By substituting Eq.(54) into Eq.(53), the optimization problem reduces to:

$$\min_{\{\Theta, \mathbf{U}\}} \text{trace} \left(\mathbf{U}^\top \left(\frac{1}{n} \mathbf{X}^\top \mathbf{X} + (\alpha + \beta) \mathbf{I} - \alpha \Theta^\top \Theta \right) \mathbf{U} - \frac{2}{n} \mathbf{Y}^\top \mathbf{X} \mathbf{U} + \frac{1}{n} \mathbf{Y}^\top \mathbf{Y} \right) \quad (55)$$

$$\text{subject to: } \Theta\Theta^\top = \mathbf{I}$$

By fixing Θ , since $\frac{\partial \text{trace}(\mathbf{A}^\top \mathbf{B}\mathbf{A})}{\partial \mathbf{A}} = (\mathbf{B}^\top + \mathbf{B})\mathbf{A}$, the derivative of the objective in Eq.(55) with respect to \mathbf{U} turns out to be $2 \left(\mathbf{M} - \alpha \Theta^\top \Theta \right) \mathbf{U} - \frac{2}{n} \mathbf{X}^\top \mathbf{Y}$. Setting the derivative to zero yields the optimal solution for \mathbf{U} as follows:

$$\mathbf{U} = \frac{1}{n} \left(\mathbf{M} - \alpha \Theta^\top \Theta \right)^{-1} \mathbf{X}^\top \mathbf{Y} \quad \text{where } \mathbf{M} = \frac{1}{n} \mathbf{X}^\top \mathbf{X} + (\alpha + \beta) \mathbf{I} \quad (56)$$

 $Y = \text{MLLS}(\mathcal{D}, \alpha, \beta, \mathbf{x})$

1. Induce model parameters Θ by solving the optimization problem in Eq.(58);
 2. Induce model parameters U according to Eq.(56);
 3. Induce model parameters V according to Eq.(54);
 4. Set $W = U - \Theta^\top V$;
 5. Return Y according to Eq.(59);
-
-

Fig. 14. Pseudo-code of MLLS.

By substituting Eq.(56) into Eq.(55), the optimization problem further reduces to:

$$\max_{\Theta} \frac{1}{n^2} \text{trace} \left(\mathbf{Y}^\top \mathbf{X} (\mathbf{M} - \alpha \Theta^\top \Theta)^{-1} \mathbf{X}^\top \mathbf{Y} \right) \quad (57)$$

subject to : $\Theta \Theta^\top = \mathbf{I}$

The term $(\mathbf{M} - \alpha \Theta^\top \Theta)^{-1}$ in the objective is expanded via the Sherman-Woodbury-Morrison formula $(\mathbf{A} + \mathbf{B}\mathbf{C})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{B}(\mathbf{I} + \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{C}\mathbf{A}^{-1}$ [57]. Thereafter, since $\text{trace}(\mathbf{A}\mathbf{B}) = \text{trace}(\mathbf{B}\mathbf{A})$ and $\Theta \Theta^\top = \mathbf{I}$, Eq.(57) can be re-written in the following form:

$$\max_{\Theta} \text{trace} \left((\Theta \mathbf{S}_1 \Theta^\top)^{-1} \Theta \mathbf{S}_2 \Theta^\top \right) \quad (58)$$

subject to : $\Theta \Theta^\top = \mathbf{I}$

Here, $\mathbf{S}_1 = \mathbf{I} - \alpha \mathbf{M}^{-1}$ and $\mathbf{S}_2 = \mathbf{M}^{-1} \mathbf{X}^\top \mathbf{Y} \mathbf{Y}^\top \mathbf{X} \mathbf{M}^{-1}$. As shown in [179], the optimal value for Θ can be found by solving the generalized eigenvalue problem: $\mathbf{S}_1 \mathbf{z} = \lambda \mathbf{S}_2 \mathbf{z}$. Firstly, let $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_r] \in \mathbb{R}^{d \times r}$ be the matrix formed by the r eigenvectors corresponding to the largest r nonzero eigenvalues. Secondly, let $\mathbf{Z} = \mathbf{Z}_Q \mathbf{Z}_R$ be the reduced QR factorization of \mathbf{Z} , where $\mathbf{Z}_Q \in \mathbb{R}^{d \times r}$ has r orthonormal columns (i.e. $\mathbf{Z}_Q^\top \mathbf{Z}_Q = \mathbf{I}$) and $\mathbf{Z}_R \in \mathbb{R}^{r \times r}$ is an upper triangular matrix. Finally, the optimal solution for Θ corresponds to $\Theta = \mathbf{Z}_Q^\top$.

For unseen instance \mathbf{x} , the predicted label set corresponds to:

$$Y = \{y_j \mid (\mathbf{w}_j + \Theta^\top \mathbf{v}_j)^\top \mathbf{x} > 0, 1 \leq j \leq q\} \quad (59)$$

Remarks : The pseudo-code of MLLS is summarized in Fig. 14. It is a *high-order* approach which considers correlations among all labels via the shared feature subspace. One prominent advantage of MLLS lies in that the optimal solution to Eq.(52) can be identified in *closed-form*.

Exploiting label correlations by manipulating the input space can be realized in other ways such as Canonical Correlation Analysis (CCA) [139], [140], [194] or kernel feature mapping [166]. As shown in Fig. 14, let r be the rank of instance matrix \mathbf{X} , MLLS has computational complexity of $\mathcal{O}(dm^2 + rq^2 + dr^2)$ for training and $\mathcal{O}(dq)$ for testing.

D. Summary

Table II summarizes properties of the twelve multi-label learning algorithms investigated in Subsections III-B and III-C, including their basic idea, label correlations, computational complexity, tested domains, and optimized (surrogate) metric. As shown in Table II, (surrogate) *hamming loss* and *ranking loss* are among the most popular metrics to be optimized and theoretical analyses on them [35]–[37], [52] have been discussed in Subsection II-B.4. Furthermore, note that the *subset accuracy* optimized by Random k -Labelsets is only measured with respect to the k -labelset instead of the whole label space.

Domains reported in Table II correspond to data types on which the corresponding algorithm is shown to work well in the *original* literature. However, all those representative multi-label learning algorithms are *general-purpose* and can be applied to various data types. Nevertheless, the computational complexity of each learning algorithm does play a key factor on its suitability for different scales of data. Here, the data *scalability* can be studied in terms of three main aspects including the number of training examples (i.e. m), the dimensionality (i.e. d), and the number of possible class labels (i.e. q). Furthermore, algorithms which append class labels as extra features to instance space [55], [122], [129] might not benefit too much from this strategy when instance dimensionality is much larger than the number of class labels (i.e. $d \gg q$).

As arguably the mostly-studied supervised learning framework, a number of algorithms in Table II employ *binary classification* as the intermediate step to learn from multi-label data [7], [49], [55], [122], [129]. In addition to those transformation strategies being introduced, there are other ways to transform the task of multi-label learning into the task of binary classification such as Error-Correcting Output Codes (ECOC) [46], [195]. In addition, first-order algorithm

TABLE II
SUMMARY OF REPRESENTATIVE MULTI-LABEL LEARNING ALGORITHMS BEING REVIEWED.

Algorithm	Basic Idea	Order of Correlations	Complexity [Train/Test]	Tested Domains	Optimized Metric
Binary Relevance [7]	Fit multi-label data to q binary classifiers	first-order	$\mathcal{O}(q \cdot \mathcal{F}_{\mathcal{B}}(m, d)) / \mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d))$	image	classification (hamming loss)
AdaBoost.MH [129]	Fit multi-label data to one binary classifier	high-order	$\mathcal{O}(\mathcal{F}_{\mathcal{B}}(mq, d + 1)) / \mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d + 1))$	text	classification (hamming loss)
Stacked Aggregation [55]	Fit multi-label data to stacked binary classifiers	high-order	$\mathcal{O}(q \cdot \mathcal{F}_{\mathcal{B}}(m, d + q)) / \mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d + q))$	text	classification (hamming loss)
Classifier Chains [122]	Fit multi-label data to a chain of binary classifiers	high-order	$\mathcal{O}(q \cdot \mathcal{F}_{\mathcal{B}}(m, d + q)) / \mathcal{O}(q \cdot \mathcal{F}'_{\mathcal{B}}(d + q))$	image, video, text, biology	classification (hamming loss)
Calibrated Label Ranking [49]	Fit multi-label data to $\frac{q(q+1)}{2}$ binary classifiers	second-order	$\mathcal{O}(q^2 \cdot \mathcal{F}_{\mathcal{B}}(m, d)) / \mathcal{O}(q^2 \cdot \mathcal{F}'_{\mathcal{B}}(d))$	image, text, biology	Ranking (ranking loss)
Random k -Labelsets [157]	Fit multi-label data to n multi-class classifiers	high-order	$\mathcal{O}(n \cdot \mathcal{F}_{\mathcal{M}}(m, d, 2^k)) / \mathcal{O}(n \cdot \mathcal{F}'_{\mathcal{M}}(d, 2^k))$	image, text, biology	classification (subset accuracy)
ML- k NN [191]	Fit k -nearest neighbor to multi-label data	first-order	$\mathcal{O}(m^2 d + qmk) / \mathcal{O}(md + qk)$	image, text, biology	classification (hamming loss)
ML-DT [25]	Fit decision tree to multi-label data	first-order	$\mathcal{O}(mdq) / \mathcal{O}(mq)$	biology	classification (hamming loss)
Rank-SVM [42]	Fit kernel learning to multi-label data	second-order	$\mathcal{O}(\mathcal{F}_{\text{QP}}(dq + mq^2, mq^2) + q^2(q + m)) / \mathcal{O}(dq)$	biology	Ranking (ranking loss)
BP-MLL [190]	Fit neural network to multi-label data	second-order	$\mathcal{O}(s(d + q)mN + q^2(q + m)) / \mathcal{O}(s(d + q))$	text, biology	Ranking (ranking loss)
CML [53]	Fit conditional random field to multi-label data	second-order	$\mathcal{O}(\mathcal{F}_{\text{UNC}}(dq + q^2, m)) / \mathcal{O}((dq + q^2) \cdot 2^q)$	text	classification (subset accuracy)
MLLS [77]	Fit spectral analysis to multi-label data	high-order	$\mathcal{O}(dm^2 + rq^2 + dr^2) / \mathcal{O}(dq)$	text	classification (hamming loss)

adaptation methods can not be simply regarded as Binary Relevance [7] combined with specific binary learners. For example, ML- k NN [191] is more than Binary Relevance combined with k NN as Bayesian inference is employed to reason with neighboring information.

IV. EXTENDED TOPICS

A. Noisy/Weak Labels

In standard setting of multi-label learning, each training example (\mathbf{x}, Y) is assumed to have *valid* as well as *complete* label association. In other words, labels associated with \mathbf{x} , i.e. Y , are

its true relevant labels, while labels not associated with \mathbf{x} , i.e. \bar{Y} , are its true irrelevant labels. However, this assumption may not always hold in real-world applications.

On one case, labels assigned to an training example may be false positive [81], [107]. For example, in web applications an online object may be labeled by several labelers while it is quite possible that some of the labels are noisy due to the carelessness or inexpertise of the labelers. In the extreme case, only one label in the label set is valid and all the other labels in Y are false positive ones, and this setting has been recently termed as partial label learning [28].

On the other case, labels not assigned to an training example may be false negative [141]. For example, in medical applications the ground-truth disease for a patient may not be identified due to inaccuracy of medical examinations. When all the labels are independent, multi-label learning under this case can be decomposed into a series of PU-learning (Positive and Unlabeled data learning) problems [43], where instance without a specific label is considered to be unlabeled.

B. Unlabeled Data

In multi-label learning, labeling training examples is a labor-intensive task as each example should be assigned with *multiple* labels instead of a single one. In many real-world applications however, unlabeled data is usually readily available and easy to collect at lower cost. Therefore, it is natural to improve the effectiveness of multi-label learning system with unlabeled data.

Active learning [26] is one of the major frameworks to learn from labeled and unlabeled data, where the learning system can query the ground-truth labels of selected unlabeled data with an oracle. Earlier approaches on active multi-label learning assume independence among labels and design margin-based query criteria [8], [94]. Later works aim to further exploit label correlations with applications to image classification [113] and text categorization [44], [73], [177].

Semi-supervised learning [204] is another major framework to learn from labeled and unlabeled data, where the learning system automatically exploit unlabeled data in addition to labeled data without interaction with environment. Existing approaches on semi-supervised multi-label learning work under the transductive setting, and also take the common manifold assumption

that two examples with high similarity in the input space tend to have large overlap in the output space [20], [86], [95], [165], [168], [169], [184], [185].

C. Dimensionality Reduction

Multi-label data with high dimensionality widely exist in real-world applications, such as text categorization, gene functional analysis, multimedia content processing, etc. To mitigate the problem of *curse-of-dimensionality*, it is essential to design dimensionality reduction techniques for learning from high-dimensional multi-label data.

For filter-style dimensionality reduction techniques, the feature reduction process is independent of the learning algorithms to be used for classifier induction. Existing approaches within this category perform multi-label dimensionality reduction by extending popular feature extraction methods such as Latent Semantic Indexing (LSI) [182], Linear Discriminant Analysis (LDA) [92], [109], [164], Mutual Information [40], [91], and Hilbert-Schmidt Independence Criterion (HSIC) [196], [197].

For wrapper-style dimensionality reduction techniques, the feature reduction process is involved with the learning algorithms to be used for classifier induction. Existing approaches within this category perform multi-label dimensionality reduction by jointly learning a set of linear classifiers and the feature transformation matrix [60], [79]. There are some other wrapper-style approaches which conduct dimensionality reduction by combining with filter-style techniques [187] or unlabeled data [114].

D. Large-Scale Data

In multi-label learning, addressing the problem of large-scale data needs to be considered from two different facets. On one hand, large-scale data refers to the ordinary setting where the training set contains *large number of examples*. On the other hand, large-scale data refers to the specific setting where the label space contains *large number of labels*.

Most multi-label algorithms aim to exploit label correlations to improve the generalization ability of learning system. Nevertheless, when the training examples are abundant, there have

been some empirical results indicating that simple strategies without sophisticated label correlation modeling would suffice to achieve good performance, such as Binary Relevance with SVM [19] or with mild prior information [65], [66].

Another important issue for large-scale multi-label learning is to handle large output space. When the number of labels is huge, most existing multi-label learning algorithms would suffer from problems of computational inefficiency and class-imbalance [130], [153], especially for the second-order and high-order approaches. In this case, it is essential to make use of the sparsity nature of the training examples' labeling matrix to improve both efficiency and effectiveness [32], [69], [198], [199].

E. Related Learning Settings

There are several learning settings related to multi-label learning which are worth some discussion, such as multi-instance learning [39], ordinal classification [47], multi-task learning [15], and data streams classification [50].

Multi-instance learning [39] studies the problem where each example is described by a bag of instances while associated with a single (binary) label. A bag is regarded to be positive iff at least one of its constituent instances is positive. In contrast to multi-label learning which models the object's ambiguities (complicated semantics) in output (label) space, multi-instance learning can be viewed as modeling the object's ambiguities in input (instance) space [201], [202]. There are some initial attempt towards exploiting multi-instance representation for learning from multi-label data [192].

Ordinal classification [47] studies the problem where a natural ordering exists among all the class labels. In multi-label learning, we can accordingly assume an ordering of relevance on each class label to generalize the crisp membership ($y_j \in \{-1, +1\}$) into the graded membership ($y_j \in \{m_1, m_2, \dots, m_k\}$ where $m_1 < m_2 < \dots < m_k$). Therefore, graded multi-label learning accommodates the case where we can only provide vague (ordinal) instead of definite judgement on the label relevance. Existing work shows that graded multi-label learning can be solved by

transforming it into a set of ordinal classification problems (one for each class label), or a set of standard multi-label learning problems (one for each membership level) [21].

Multi-task learning [15] studies the problem where multiple tasks are trained in parallel such that training information of related tasks are used as an inductive bias to help improve the generalization performance of other tasks. Nonetheless, there are some essential differences between multi-task learning and multi-label learning to be noticed. Firstly, in multi-label learning all the examples share the same feature space, while in multi-task learning the tasks can be in the same feature space or different feature spaces. Secondly, in multi-label learning the goal is to predict the label subset associated with an object, while the purpose of multi-task learning is to have multiple tasks to be learned well simultaneously, and it does not concern on which task subset should be associated with an object (if we take a label as a task) since it generally assumes that every object is involved by all tasks. Thirdly, as discussed in Subsection IV-D, one major difficulty of multi-label learning occurs when the label space is large while in multi-task learning it is not reasonable to consider a large number of tasks. Nevertheless, techniques for multi-task learning might be used to benefit multi-label learning [97].

Data streams classification [50] studies the problem where real-world objects are generated online and processed in a real-time manner. Nowadays, streaming data with multi-label nature widely exist in real-world scenarios such as instant news, emails, microblogs, etc [120]. As a usual challenge for streaming data analysis, the key factor for effectively classifying multi-label data streams is how to deal with the concept drift problem. Existing works model concept drift by updating the classifiers significantly whenever a new batch of examples arrive [115], taking the fading assumption that the influence of past data gradually declines as time evolves [88], [135], or maintaining a change detector alerting whenever a concept drift is detected [120].

V. CONCLUSION

In this paper, the state-of-the-art of multi-label learning is reviewed in terms of paradigm formalization, learning algorithms and extended topics. In particular, instead of trying to go

TABLE III
ONLINE RESOURCES FOR MULTI-LABEL LEARNING.

Resource Type	Resource URL and Descriptions
Tutorial	http://www.ecmlpkdd2009.net/program/tutorials/learning-from-multi-label-data/ (In conjunction with ECML PKDD 2009)
	http://cig.fi.upm.es/index.php/presentations?download=4 (In conjunction with TAMIDA 2010)
Workshops	http://lpis.csd.auth.gr/workshops/ml09/ (MLD'09 : in conjunction with ECML PKDD 2009)
	http://cse.seu.edu.cn/conf/MLD10/ (MLD'10 : in conjunction with ICML/COLT 2010)
	http://cse.seu.edu.cn/conf/LAWS12/ (LAWS'12 : in conjunction with ACML 2012)
Special Issue	http://mlkd.csd.auth.gr/events/ml2010si.html (Machine Learning Journal Special Issue on Learning from Multi-Label Data [17], [37], [65], [68], [120], [126], [178], [199])
Software	http://mulan.sourceforge.net/index.html (The MULAN [156] open-source Java library)
	http://meqa.sourceforge.net/ (The MEKA project based on WEKA [63])
Data Sets	http://cse.seu.edu.cn/people/zhangml/Resources.htm#codes_mll (Matlab codes multi-label learning)
	http://mulan.sourceforge.net/datasets.html (Data sets from MULAN)
	http://meqa.sourceforge.net/#datasets (Data sets from MEKA)
	http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/multilabel.html (Data sets from LIBSVM [18])

through all the learning techniques within confined space, which would lead to only abridged introductions, we choose to elaborate the algorithmic details of twelve representative multi-label learning algorithms with references to other related works. Some online resources for multi-label learning are summarized in Table III, including academic activities (tutorial, workshops, special issue), publicly-available software and data sets.

As discussed in Section II-A.2, although the idea of exploiting label correlations have been employed by various multi-label learning techniques, there has not been any formal characterization on the underlying concept or any principled mechanism on the appropriate usage of label correlations. Recent researches indicate that correlations among labels might be *asymmetric*, i.e. the influence of one label to the other one is not necessarily be the same in the inverse direction [71], or *local*, i.e. different instances share different label correlations with few correlations being globally applicable [72]. Nevertheless, full understanding on label correlations, especially for scenarios with large output space, would remain as the holy grail for multi-label learning.

As reviewed in Section III, multi-label learning algorithms are introduced by focusing on their algorithmic properties. One natural complement to this review would be conducting thorough experimental studies to get insights on the pros and cons of different multi-label learning algorithms. A recent attempt towards extensive experimental comparison can be found in [103] where 12 multi-label learning algorithms are compared with respect to 16 evaluation metrics. Interestingly while not surprisingly, the best-performing algorithm for both classification and ranking metrics turns out to be the one based on *ensemble learning* techniques (i.e. random forest of predictive decision trees [85]). Nevertheless, empirical comparison across a broad range or within a focused type (e.g. [136]) are worthwhile topic to be further explored.

REFERENCES

- [1] B.-K. Bao, B. Ni, Y. Mu, and S. Yan, "Efficient region-aware large graph construction towards scalable multi-label propagation," *Pattern Recognition*, vol. 44, no. 3, pp. 598–606, 2011.
- [2] Z. Barutcuoglu, R. E. Schapire, and O. G. Troyanskaya, "Hierarchical multi-label prediction of gene function," *Bioinformatics*, vol. 22, no. 7, pp. 830–836, 2006.
- [3] A. L. Berger, V. J. Della Pietra, and S. A. Della Pietra, "A maximum entropy approach to natural language processing," *Computational Linguistics*, vol. 22, no. 1, pp. 39–71, 1996.
- [4] W. Bi and J. T. Kwok, "Multi-label classification on tree- and DAG-structured hierarchies," in *Proceedings of the 28th International Conference on Machine Learning*, Bellevue, WA, 2011, pp. 17–24.
- [5] C. Bielza, G. Li, and P. Larrañaga, "Multi-dimensional classification with Bayesian networks," *International Journal of Approximate Reasoning*, vol. 52, no. 6, pp. 705–727, 2011.
- [6] H. Borchani, C. Bielza, C. Toro, and P. Larrañaga, "Predicting human immunodeficiency virus type 1 inhibitors using multi-dimensional Bayesian network classifiers," *Artificial Intelligence in Medicine*, in press.
- [7] M. R. Boutell, J. Luo, X. Shen, and C. M. Brown, "Learning multi-label scene classification," *Pattern Recognition*, vol. 37, no. 9, pp. 1757–1771, 2004.
- [8] K. Brinker, "On active learning in multi-label classification," in *Proceedings of the 29th Annual Conference of the German Classification Society*, Magdeburg, Germany, 2005, pp. 206–213.
- [9] K. Brinker, J. Fürnkranz, and E. Hüllermeier, "A unified model for multilabel classification and ranking," in *Proceedings of the 17th European Conference on Artificial Intelligence*, Riva del Garda, Italy, 2006, pp. 489–493.
- [10] K. Brinker and E. Hüllermeier, "Case-based multilabel ranking," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, Hyderabad, India, 2007, pp. 702–707.
- [11] S. S. Bucak, R. Jin, and A. K. Jain, "Multi-label multiple kernel learning by stochastic approximation: Application to visual object recognition," in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. Cambridge, MA: MIT Press, 2010, pp. 325–333.
- [12] R. H. Byrd, J. Nocedal, and R. B. Schnabel, "Representations of quasi-newton matrices and their use in limited memory methods," *Mathematical Programming*, vol. 63, no. 1-3, pp. 129–156, 1994.

- [13] R. S. Cabral, F. De la Torre, J. P. Costeira, and A. Bernardino, "Matrix completion for multi-label image classification," in *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. Bartlett, F. C. N. Pereira, and K. Q. Weinberger, Eds. Cambridge, MA: MIT Press, 2011, pp. 190–198.
- [14] L. Cai and T. Hofmann, "Hierarchical document categorization with support vector machines," in *Proceedings of the 13th ACM International Conference on Information and Knowledge Management*, Washington, D.C., 2004, pp. 78–87.
- [15] R. Caruana, "Multitask learning," *Machine Learning*, vol. 28, no. 1, pp. 41–75, 1997.
- [16] R. Caruana, A. Niculescu-Mizil, G. Crew, and A. Ksikes, "Ensemble selection from libraries of models," in *Proceedings of the 21st International Conference on Machine Learning*, Banff, Canada, 2004.
- [17] N. Cesa-Bianchi, M. Re, and G. Valentini, "Synergy of multi-label hierarchical ensembles, data fusion, and cost-sensitive methods for gene functional inference," *Machine Learning*, vol. 88, no. 1-2, pp. 209–241, 2012.
- [18] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, no. 3, 2011, Article 27, Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [19] F. Chang and C.-C. Liu, "Solving large-scale multi-label svm problems with a tree decomposition approach," in *Poster Session of the 3rd Asian Conference on Machine Learning*, Taoyuan, Taiwan, 2011.
- [20] G. Chen, Y. Song, F. Wang, and C. Zhang, "Semi-supervised multi-label learning by solving a Sylvester equation," in *Proceedings of the 8th SIAM International Conference on Data Mining*, Atlanta, GA, 2008, pp. 410–419.
- [21] W. Cheng, K. Dembczyński, and E. Hüllermeier, "Graded multilabel classification: The ordinal case," in *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel, 2010, pp. 223–230.
- [22] W. Cheng and E. Hüllermeier, "Combining instance-based learning and logistic regression for multilabel classification," *Machine Learning*, vol. 76, no. 2-3, pp. 211–225, 2009.
- [23] —, "A simple instance-based approach to multilabel classification using the Mallows model," in *Working Notes of the First International Workshop on Learning from Multi-Label Data*, Bled, Slovenia, 2009, pp. 28–38.
- [24] T.-H. Chiang, H.-Y. Lo, and S.-D. Lin, "A ranking-based KNN approach for multi-label classification," in *Proceedings of the 4th Asian Conference on Machine Learning*, Singapore, 2012, pp. 81–96.
- [25] A. Clare and R. D. King, "Knowledge discovery in multi-label phenotype data," in *Lecture Notes in Computer Science 2168*, L. De Raedt and A. Siebes, Eds. Berlin: Springer, 2001, pp. 42–53.
- [26] D. Cohn, L. Atlas, and R. Ladner, "Improved generalization with active learning," *Machine Learning*, vol. 15, no. 2, pp. 201–221, 1994.
- [27] F. D. Comité, R. Gilleron, and M. Tommasi, "Learning multi-label alternating decision tree from texts and data," in *Lecture Notes in Computer Science 2734*, P. Perner and A. Rosenfeld, Eds. Berlin: Springer, 2003, pp. 35–49.
- [28] T. Cour, B. Sapp, and B. Taskar, "Learning from partial labels," *Journal of Machine Learning Research*, vol. 12, no. May, pp. 1501–1536, 2011.
- [29] A. de Carvalho and A. A. Freitas, "A tutorial on multi-label classification techniques," in *Studies in Computational Intelligence 205*, A. Abraham, A. E. Hassanien, and V. Snásel, Eds. Berlin: Springer, 2009, pp. 177–195.
- [30] A. F. De Souza, F. Pedroni, E. Oliveira, P. M. Ciarelli, W. F. Henrique, L. Veronese, and C. Badue, "Automated multi-label text categorization with VG-RAM weightless neural networks," *Neurocomputing*, vol. 72, no. 10-12, pp. 2209–2217, 2009.
- [31] P. R. de Waal and L. C. van der Gaag, "Inference and learning in multi-dimensional Bayesian network classifiers," in *Lecture Notes in Artificial Intelligence 4724*, K. Mellouli, Ed. Berlin: Springer, 2007, pp. 501–511.

- [32] O. Dekel and O. Shamir, “Multiclass-multilabel classification with more classes than examples,” in *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, Sardinia, Italy, 2010, pp. 137–144.
- [33] V. Delcroix, M.-A. Maalej, and S. Piechowiak, “Bayesian networks versus other probabilistic models for the multiple diagnosis of large devices,” *International Journal on Artificial Intelligence Tools*, vol. 16, no. 3, pp. 417–433, 2007.
- [34] K. Dembczyński, W. Cheng, and E. Hüllermeier, “Bayes optimal multilabel classification via probabilistic classifier chains,” in *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel, 2010, pp. 279–286.
- [35] K. Dembczyński, W. Kotłowski, and E. Hüllermeier, “Consistent multilabel ranking through univariate loss minimization,” in *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, UK, 2012, pp. 1319–1326.
- [36] K. Dembczyński, W. Waegeman, W. Cheng, and E. Hüllermeier, “Regret analysis for performance metrics in multi-label classification: The case of hamming loss and subset zero-one loss,” in *Lecture Notes in Artificial Intelligence 6321*, J. Balcázar, F. Bonchi, A. Gionis, and M. Sebag, Eds. Berlin: Springer, 2010, pp. 280–295.
- [37] —, “On label dependence and loss minimization in multi-label classification,” *Machine Learning*, vol. 88, no. 1-2, pp. 5–45, 2012.
- [38] T. G. Dietterich, “Ensemble methods in machine learning,” in *Proceedings of the 1st International Workshop on Multiple Classifier Systems*, Cagliari, Italy, 2000, pp. 1–15.
- [39] T. G. Dietterich, R. H. Lathrop, and T. Lozano-Pérez, “Solving the multiple-instance problem with axis-parallel rectangles,” *Artificial Intelligence*, vol. 89, no. 1-2, pp. 31–71, 1997.
- [40] G. Doquire and M. Verleysen, “Feature selection for multi-label classification problems,” in *Lecture Notes in Computer Science 6691*, J. Cabestany, I. Rojas, and G. Joya, Eds. Berlin: Springer, 2011, pp. 9–16.
- [41] A. Elisseeff and J. Weston, “Kernel methods for multi-labelled classification and categorical regression problems,” BIOwulf Technologies, Tech. Rep., 2001.
- [42] —, “A kernel method for multi-labelled classification,” in *Advances in Neural Information Processing Systems 14*, T. G. Dietterich, S. Becker, and Z. Ghahramani, Eds. Cambridge, MA: MIT Press, 2002, pp. 681–687.
- [43] C. Elkan and K. Noto, “Learning classifiers from only positive and unlabeled data,” in *Proceedings of the 14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Las Vegas, NV, 2008, pp. 213–220.
- [44] A. Esuli and F. Sebastiani, “Active learning strategies for multi-label text categorization,” in *Lecture Notes in Computer Science 5478*, M. Boughanem, C. Berrut, J. Mothe, and C. Soule-Dupuy, Eds. Berlin: Springer, 2009, pp. 102–113.
- [45] R.-E. Fan and C.-J. Lin, “A study on threshold selection for multi-label classification,” National Taiwan University, Tech. Rep., 2007.
- [46] C. S. Ferng and H. T. Lin, “Multi-label classification with error-correcting codes,” in *Proceedings of the 3rd Asian Conference on Machine Learning*, Taoyuan, Taiwan, 2011, pp. 281–295.
- [47] E. Frank and M. Hall, “A simple approach to ordinal classification,” in *Lecture Notes in Computer Science 2167*, L. De Raedt and P. Flach, Eds. Berlin: Springer, 2001, pp. 145–156.
- [48] Y. Freund and R. E. Schapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [49] J. Fürnkranz, E. Hüllermeier, E. Loza Mencía, and K. Brinker, “Multilabel classification via calibrated label ranking,” *Machine Learning*, vol. 73, no. 2, pp. 133–153, 2008.
- [50] M. M. Gaber, A. Zaslavsky, and S. Krishnaswamy, “A survey of classification methods in data streams,” in *Data Streams: Models and Algorithms*, C. C. Aggarwal, Ed. Berlin: Springer, 2007, pp. 39–59.

- [51] S. Gao, W. Wu, C.-H. Lee, and T.-S. Chua, "A MFoM learning approach to robust multiclass multi-label text categorization," in *Proceedings of the 21st International Conference on Machine Learning*, Banff, Canada, 2004.
- [52] W. Gao and Z.-H. Zhou, "On the consistency of multi-label learning," in *Proceedings of the 24th Annual Conference on Learning Theory*, Budapest, Hungary, 2011, pp. 341–358.
- [53] N. Ghamrawi and A. McCallum, "Collective multi-label classification," in *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*, Bremen, Germany, 2005, pp. 195–200.
- [54] D. Gjorgjevikj and G. Madjarov, "Two stage classifier chain architecture for efficient pair-wise multi-label learning," in *Proceedings of the 21st IEEE International Workshop on Machine Learning for Signal Processing*, Beijing, China, 2011, pp. 1–6.
- [55] S. Godbole and S. Sarawagi, "Discriminative methods for multi-labeled classification," in *Lecture Notes in Artificial Intelligence 3056*, H. Dai, R. Srikant, and C. Zhang, Eds. Berlin: Springer, 2004, pp. 22–30.
- [56] K. Gold and A. Petrosino, "Using information gain to build meaningful decision forests for multilabel classification," in *Proceedings of the 9th IEEE International Conference on Development and Learning*, Ann Arbor, MI, 2010, pp. 58–63.
- [57] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD: Johns Hopkins University Press, 1996.
- [58] S. Gopal and Y. Yang, "Multilabel classification with meta-level features," in *Proceedings of the 33rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Geneva, Switzerland, 2010, pp. 315–322.
- [59] R. Grodzicki, J. Mańdziuk, and L. Wang, "Improved multilabel classification with neural networks," in *Lecture Notes in Computer Science 5199*, G. Rudolph, T. Jansen, S. Lucas, C. Poloni, and N. Beume, Eds. Berlin: Springer, 2008, pp. 409–416.
- [60] Q. Gu, Z. Li, and J. Han, "Correlated multi-label feature selection," in *Proceedings of the 20th ACM Conference on Information and Knowledge Management*, Glasgow, UK, 2011, pp. 1087–1096.
- [61] Y. Guo and S. Gu, "Multi-label classification using conditional dependency networks," in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, Barcelona, Spain, 2011, pp. 1300–1305.
- [62] Y. Guo and D. Schuurmans, "Adaptive large margin training for multilabel classification," in *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, San Francisco, CA, 2011, pp. 374–379.
- [63] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The WEKA data mining software: An update," *SIGKDD Explorations*, vol. 11, no. 1, pp. 10–18, 2009.
- [64] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, vol. 143, no. 1, pp. 29–36, 1982.
- [65] B. Hariharan, S. V. N. Vishwanathan, and M. Varma, "Efficient max-margin multi-label classification with applications to zero-shot learning," *Machine Learning*, vol. 88, no. 1-2, pp. 127–155, 2012.
- [66] B. Hariharan, L. Zelnik-Manor, S. V. N. Vishwanathan, and M. Varma, "Large scale max-margin multi-label classification with priors," in *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel, 2010, pp. 423–430.
- [67] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1999.
- [68] J. He, H. Gu, and Z. Wang, "Bayesian multi-instance multi-label learning using Gaussian process prior," *Machine Learning*, vol. 88, no. 1-2, pp. 273–295, 2012.
- [69] D. Hsu, S. Kakade, J. Langford, and T. Zhang, "Multi-label prediction via compressed sensing," in *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, Eds. Cambridge, MA: MIT Press, 2009, pp. 772–780.

- [70] K.-W. Huang and Z. Li, “A multilabel text classification algorithm for labeling risk factors in SEC form 10-K,” *ACM Transactions on Management Information Systems*, vol. 2, no. 3, 2011, Article 18.
- [71] S.-J. Huang, Y. Yu, and Z.-H. Zhou, “Multi-label hypothesis reuse,” in *Proceedings of the 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Beijing, China, 2012, pp. 525–533.
- [72] S.-J. Huang and Z.-H. Zhou, “Multi-label learning by exploiting label correlations locally,” in *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, Toronto, Canada, 2012, pp. 949–955.
- [73] C.-W. Hung and H.-T. Lin, “Multi-label active learning with auxiliary learner,” in *Proceedings of the 3rd Asian Conference on Machine Learning*, Taoyuan, Taiwan, 2011, pp. 315–330.
- [74] M. Ioannou, G. Sakkas, G. Tsoumakas, and I. Vlahavas, “Obtaining bipartition from score vectors for multi-label classification,” in *Proceedings of the 22nd IEEE International Conference on Tools with Artificial Intelligence*, Arras, France, 2010, pp. 409–416.
- [75] E. T. Jaynes, “Information theory and statistical mechanics,” *Physical Review*, vol. 106, no. 4, pp. 620–630, 1957.
- [76] S. Ji, L. Sun, R. Jin, and J. Ye, “Multi-label multiple kernel learning,” in *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds. Cambridge, MA: MIT Press, 2009, pp. 777–784.
- [77] S. Ji, L. Tang, S. Yu, and J. Ye, “Extracting shared subspace for multi-label classification,” in *Proceedings of the 14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Las Vegas, NV, 2008, pp. 381–389.
- [78] —, “A shared-subspace learning framework for multi-label classification,” *ACM Transactions on Knowledge Discovery from Data*, vol. 4, no. 2, 2010, Article 8.
- [79] S. Ji and J. Ye, “Linear dimensionality reduction for multi-label classification,” in *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, Pasadena, TX, 2009, pp. 1077–1082.
- [80] A. Jiang, C. Wang, and Y. Zhu, “Calibrated rank-svm for multi-label image categorization,” in *Proceedings of the International Joint Conference on Neural Networks*, Hong Kong, 2008, pp. 1450–1455.
- [81] R. Jin and Z. Ghahramani, “Learning with multiple labels,” in *Advances in Neural Information Processing Systems 15*, S. Becker, S. Thrun, and K. Obermayer, Eds. Cambridge, MA: MIT Press, 2003, pp. 897–904.
- [82] F. Kang, R. Jin, and R. Sukthankar, “Correlated label propagation with application to multi-label learning,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, New York, NY, 2006, pp. 1719–1726.
- [83] I. Katakis, G. Tsoumakas, and I. Vlahavas, “Multilabel text classification for automated tag suggestion,” in *Proceedings of the ECML PKDD 2008 Discovery Challenge*, Antwerp, Belgium, 2008, pp. 75–83.
- [84] H. Kazawa, T. Izumitani, H. Taira, and E. Maeda, “Maximal margin labeling for multi-topic text categorization,” in *Advances in Neural Information Processing Systems 17*, L. K. Saul, Y. Weiss, and L. Bottou, Eds. Cambridge, MA: MIT Press, 2005, pp. 649–656.
- [85] D. Kocev, C. Vens, J. Struyf, and S. Džeroski, “Ensembles of multi-objective decision trees,” in *Proceedings of the 18th European Conference on Machine Learning*, Warsaw, Poland, 2007, pp. 624–631.
- [86] X. Kong, M. Ng, and Z.-H. Zhou, “Transductive multi-label learning via label set propagation,” *IEEE Transactions on Knowledge and Data Engineering*, 2012 (in press).
- [87] X. Kong and P. S. Yu, “Multi-label feature selection for graph classification,” in *Proceedings of the 10th IEEE International Conference on Data Mining*, Sydney, Australia, 2010, pp. 274–283.
- [88] —, “An ensemble-based approach to fast classification of multi-label data streams,” in *Proceedings of the 7th*

- International Conference on Collaborative Computing: Networking, Applications and Worksharing*, Orlando, FL, 2011, pp. 95–104.
- [89] —, “gMLC: A multi-label feature selection framework for graph classification,” *Knowledge and Information Systems*, vol. 31, no. 2, pp. 281–305, 2012.
- [90] W. Kotłowski, K. Dembczyński, and E. Hüllermeier, “Bipartite ranking through minimization of univariate loss,” in *Proceedings of the 28th International Conference on Machine Learning*, Bellevue, WA, 2011, pp. 1113–1120.
- [91] G. Lastra, O. Luaces, J. R. Quevedo, and A. Bahamonde, “Graphical feature selection for multilabel classification tasks,” in *Lecture Notes in Computer Science 7014*, J. Gama, E. Bradley, and J. Hollmén, Eds. Berlin: Springer, 2011, pp. 246–257.
- [92] M. Lee and C. H. Park, “On applying dimension reduction for multi-labeled problems,” in *Lecture Notes in Artificial Intelligence 4571*, P. Perner, Ed. Berlin: Springer, 2007, pp. 131–143.
- [93] T. Li, S. Yan, T. Mei, X.-S. Hua, and I.-S. Kweon, “Image decomposition with multilabel context: Algorithms and applications,” *IEEE Transactions on Image Processing*, vol. 8, no. 20, pp. 2301–2314, 2011.
- [94] X. Li, L. Wang, and E. Sung, “Multi-label svm active learning for image classification,” in *Proceedings of the 2004 International Conference on Image Processing*, Singapore, 2004, pp. 2207–2210.
- [95] Y. Liu, R. Jin, and L. Yang, “Semi-supervised multi-label learning by constrained non-negative matrix factorization,” in *Proceedings of the 21st AAAI Conference on Artificial Intelligence*, Boston, MA, 2006, pp. 421–426.
- [96] H.-Y. Lo, J.-C. Wang, H.-M. Wang, and S.-D. Lin, “Cost-sensitive multi-label learning for audio tag annotation and retrieval,” *IEEE Transactions on Multimedia*, vol. 13, no. 3, pp. 518–529, 2011.
- [97] E. Loza Mencía, “Multilabel classification in parallel tasks,” in *Working Notes of the Second International Workshop on Learning from Multi-Label Data*, Haifa, Israel, 2010, pp. 20–36.
- [98] E. Loza Mencía and J. Fürnkranz, “Efficient pairwise multilabel classification for large-scale problems in the legal domain,” in *Lecture Notes in Artificial Intelligence 5212*, W. Daelemans, B. Goethals, and K. Morik, Eds. Berlin: Springer, 2008, pp. 50–65.
- [99] —, “Pairwise learning of multilabel classifications with perceptrons,” in *Proceedings of the International Joint Conference on Neural Networks*, Hong Kong, 2008, pp. 2899–2906.
- [100] E. Loza Mencía, S.-H. Park, and J. Fürnkranz, “Efficient voting prediction for pairwise multilabel classification,” *Neurocomputing*, vol. 73, no. 7-9, pp. 1164–1176, 2010.
- [101] G. Madjarov, D. Gjorgjevikj, and T. Delev, “Efficient two stage voting architecture for pairwise multi-label classification,” in *Lecture Notes in Computer Science 6464*, J. Li, Ed. Berlin: Springer, 2011, pp. 164–173.
- [102] G. Madjarov, D. Gjorgjevikj, and S. Džeroski, “Two stage architecture for multi-label learning,” *Pattern Recognition*, vol. 45, no. 3, pp. 1019–1034, 2012.
- [103] G. Madjarov, D. Kocev, D. Gjorgjevikj, and S. Džeroski, “An extensive experimental comparison of methods for multi-label learning,” *Pattern Recognition*, vol. 45, no. 9, pp. 3084–3104, 2012.
- [104] A. McCallum, “Multi-label text classification with a mixture model trained by EM,” in *Working Notes of the AAAI’99 Workshop on Text Learning*, Orlando, FL, 1999.
- [105] E. Montañés, J. R. Quevedo, and J. J. del Coz, “Aggregating independent and dependent models to learn multi-label classifiers,” in *Lecture Notes in Artificial Intelligence 6912*, D. Gunopulos, T. Hofmann, D. Malerba, and M. Vazirgiannis, Eds. Berlin: Springer, 2011, pp. 484–500.

- [106] K. Nigam, J. Lafferty, and A. McCallum, "Using maximum entropy for text classification," in *Working Notes of the IJCAI'99 Workshop on Machine Learning for Information Filtering*, Stockholm, Sweden, 1999, pp. 61–67.
- [107] K. Ozonat and D. Young, "Towards a universal marketplace over the web: Statistical multi-label classification of service provider forms with simulated annealing," in *Proceedings of the 15th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Paris, France, 2009, pp. 1295–1303.
- [108] F. Pachet and P. Roy, "Improving multilabel analysis of music titles: A large-scale validation of the correction approach," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 2, pp. 335–343, 2009.
- [109] C. H. Park and M. Lee, "On applying linear discriminant analysis for multi-labeled problems," *Pattern Recognition Letters*, vol. 29, no. 7, pp. 878–887, 2008.
- [110] J. Petterson and T. Caetano, "Reverse multi-label learning," in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. Cambridge, MA: MIT Press, 2010, pp. 1912–1920.
- [111] —, "Submodular multi-label learning," in *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. Bartlett, F. C. N. Pereira, and K. Q. Weinberger, Eds. Cambridge, MA: MIT Press, 2011, pp. 1512–1520.
- [112] G.-J. Qi, X.-S. Hua, Y. Rui, J. Tang, T. Mei, and H.-J. Zhang, "Correlative multi-label video annotation," in *Proceedings of the 15th ACM International Conference on Multimedia*, Augsburg, Germany, 2007, pp. 17–26.
- [113] G.-J. Qi, X.-S. Hua, Y. Rui, J. Tang, and H.-J. Zhang, "Two-dimensional multilabel active learning with an efficient online adaptation model for image classification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 10, pp. 1880–1897, 2009.
- [114] B. Qian and I. Davidson, "Semi-supervised dimension reduction for multi-label classification," in *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, Atlanta, GA, 2010, pp. 569–574.
- [115] W. Qu, Y. Zhang, J. Zhu, and Q. Qiu, "Mining multi-label concept-drifting data streams using dynamic classifier ensemble," in *Lecture Notes in Artificial Intelligence 5828*, Z.-H. Zhou and T. Washio, Eds. Berlin: Springer, 2009, pp. 308–321.
- [116] J. R. Quevedo, O. Luaces, and A. Bahamonde, "Multilabel classifiers with a probabilistic thresholding strategy," *Pattern Recognition*, vol. 45, no. 2, pp. 876–883, 2012.
- [117] J. R. Quinlan, *C4.5: Programs for Machine Learning*. San Mateo, CA: Morgan Kaufmann, 1993.
- [118] R. Rak, L. Kurgan, and M. Reformat, "Multi-label associative classification of medical documents from medline," in *Proceedings of the 4th International Conference on Machine Learning and Applications*, Los Angeles, CA, 2005, pp. 177–186.
- [119] J. Read, "A pruned problem transformation method for multi-label classification," in *Proceedings of the 2008 New Zealand Computer Science Research Student Conference*, Christchurch, New Zealand, 2008, pp. 143–150.
- [120] J. Read, A. Bifet, G. Holmes, and B. Pfahringer, "Scalable and efficient multi-label classification for evolving data streams," *Machine Learning*, vol. 88, no. 1-2, pp. 243–272, 2012.
- [121] J. Read, B. Pfahringer, and G. Holmes, "Multi-label classification using ensembles of pruned sets," in *Proceeding of the 8th IEEE International Conference on Data Mining*, Pisa, Italy, 2008, pp. 995–1000.
- [122] J. Read, B. Pfahringer, G. Holmes, and E. Frank, "Classifier chains for multi-label classification," in *Lecture Notes in Artificial Intelligence 5782*, W. Buntine, M. Grobelnik, and J. Shawe-Taylor, Eds. Berlin: Springer, 2009, pp. 254–269.
- [123] —, "Classifier chains for multi-label classification," *Machine Learning*, vol. 85, no. 3, pp. 333–359, 2011.
- [124] J. Rousu, C. Saunders, S. Szedmak, and J. Shawe-Taylor, "Learning hierarchical multi-category text classification models," in *Proceedings of the 22nd International Conference on Machine Learning*, Bonn, Germany, 2005, pp. 774–751.

- [125] —, “Kernel-based learning of hierarchical multilabel classification models,” *Journal of Machine Learning Research*, vol. 7, no. Jul, pp. 1601–1626, 2006.
- [126] T. N. Rubin, A. Chambers, P. Smyth, and M. Steyvers, “Statistical topic models for multi-label document classification,” *Machine Learning*, vol. 88, no. 1-2, pp. 157–208, 2012.
- [127] C. Sanden and J. Z. Zhang, “Enhancing multi-label music genre classification through ensemble techniques,” in *Proceedings of the 34th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Beijing, China, 2011, pp. 705–714.
- [128] E. P. Sapozhnikova, “ART-based neural networks for multi-label classification,” in *Lecture Notes in Computer Science 5772*, N. Adams, C. Robardet, and A. Siebes, Eds. Berlin: Springer, 2009, pp. 167–177.
- [129] R. E. Schapire and Y. Singer, “Boostexter: a boosting-based system for text categorization,” *Machine Learning*, vol. 39, no. 2/3, pp. 135–168, 2000.
- [130] K. Sechidis, G. Tsoumakas, and I. Vlahavas, “On the stratification of multi-label data,” in *Lecture Notes in Artificial Intelligence 6913*, D. Gunopulos, T. Hofmann, D. Malerba, and M. Vazirgiannis, Eds. Berlin: Springer, 2011, pp. 145–158.
- [131] C. Shi, X. Kong, P. S. Yu, and B. Wang, “Multi-label ensemble learning,” in *Lecture Notes in Artificial Intelligence 6913*, D. Gunopulos, T. Hofmann, D. Malerba, and M. Vazirgiannis, Eds. Berlin: Springer, 2011, pp. 223–239.
- [132] C. N. Silla and A. A. Freitas, “A survey of hierarchical classification across different application domains,” *Data Mining and Knowledge Discovery*, vol. 22, no. 1-2, pp. 31–72, 2011.
- [133] C. G. M. Snoek, M. Worring, J. C. van Gemert, J.-M. Geusebroek, and A. W. M. Smeulders, “The challenge problem for automated detection of 101 semantic concepts in multimedia,” in *Proceedings of the 14th Annual ACM International Conference on Multimedia*, Santa Barbara, CA, USA, 2006, pp. 421–430.
- [134] Y. Song, L. Zhang, and L. C. Giles, “A sparse gaussian processes classification framework for fast tag suggestions,” in *Proceeding of the 17th ACM Conference on Information and Knowledge Management*, Napa Valley, CA, 2008, pp. 93–102.
- [135] E. Spyromitros-Xioufis, M. Spiliopoulou, G. Tsoumakas, and I. Vlahavas, “Dealing with concept drift and class imbalance in multi-label stream classification,” in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, Barcelona, Spain, 2011, pp. 1583–1588.
- [136] E. Spyromitros-Xioufis, G. Tsoumakas, and I. Vlahavas, “An empirical study of lazy multilabel classification algorithms,” in *Proceedings of the 5th Hellenic Conference on Artificial Intelligence*, Syros, Greece, 2008, pp. 401–406.
- [137] —, “Multi-label learning approaches for music instrument recognition,” in *Lecture Notes in Artificial Intelligence 6804*, M. Kryszkiewicz, H. Rybinski, A. Skowron, and Z. W. Raś, Eds. Berlin: Springer, 2011, pp. 734–743.
- [138] A. P. Streich and J. M. Buhmann, “Classification of multi-labeled data: A generative approach,” in *Lecture Notes in Artificial Intelligence 5212*, W. Daelemans, B. Goethals, and K. Morik, Eds. Berlin: Springer, 2008, pp. 390–405.
- [139] L. Sun, S. Ji, and J. Ye, “Hypergraph spectral learning for multi-label classification,” in *Proceedings of the 14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Las Vegas, NV, 2008, pp. 668–676.
- [140] —, “Canonical correlation analysis for multilabel classification: A least-squares formulation, extensions, and analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 1, pp. 194–200, 2011.
- [141] Y.-Y. Sun, Y. Zhang, and Z.-H. Zhou, “Multi-label learning with weak label,” in *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, Atlanta, GA, 2010, pp. 593–598.
- [142] M. A. Tahir, J. Kittler, K. Mikolajczyk, and F. Yan, “Improving multilabel classification performance by using ensemble

- of multi-label classifiers,” in *Lecture Notes in Computer Science 5997*, N. El Gayar, J. Kittler, and F. Roli, Eds. Berlin: Springer, 2010, pp. 11–21.
- [143] L. Tang, J. Chen, and J. Ye, “On multiple kernel learning with multiple labels,” in *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, Pasadena, TX, 2009, pp. 1255–1266.
- [144] L. Tang, S. Rajan, and V. K. Narayanan, “Large scale multi-label classification via metalabeler,” in *Proceedings of the 19th International Conference on World Wide Web*, Madrid, Spain, 2009, pp. 211–220.
- [145] L. Tenenboim-Chekina, L. Rokach, and B. Shapira, “Identification of label dependencies for multi-label classification,” in *Working Notes of the Second International Workshop on Learning from Multi-Label Data*, Haifa, Israel, 2010, pp. 53–60.
- [146] F. A. Thabtah, P. Cowling, and Y. Peng, “MMAC: A new multi-class, multi-label associative classification approach,” in *Proceedings of the 4th IEEE International Conference on Data Mining*, Brighton, UK, 2004, pp. 217–224.
- [147] K. M. Ting and I. H. Witten, “Issues in stacked generalization,” *Journal of Artificial Intelligence Research*, vol. 10, pp. 271–289, 1999.
- [148] K. Trohidis, G. Tsoumakas, G. Kalliris, and I. Vlahavas, “Multilabel classification of music into emotions,” in *Proceedings of the 9th International Conference on Music Information Retrieval*, Philadelphia, PA, 2008, pp. 325–330.
- [149] I. Tsochantaridis, T. Hofmann, T. Joachims, and Y. Altun, “Support vector machine learning for interdependent and structured output spaces,” in *Proceedings of the 21st International Conference on Machine Learning*, Banff, Canada, 2004.
- [150] —, “Large margin methods for structured and interdependent output variables,” *Journal of Machine Learning Research*, vol. 6, no. Sep, pp. 1453–1484, 2005.
- [151] G. Tsoumakas, A. Dimou, E. Spyromitros, V. Mezaris, I. Kompatsiaris, and I. Vlahavas, “Correlation-based pruning of stacked binary relevance models for multi-label learning,” in *Working Notes of the First International Workshop on Learning from Multi-Label Data*, Bled, Slovenia, 2009, pp. 101–116.
- [152] G. Tsoumakas and I. Katakis, “Multi-label classification: An overview,” *International Journal of Data Warehousing and Mining*, vol. 3, no. 3, pp. 1–13, 2007.
- [153] —, “Effective and efficient multilabel classification in domains with large number of labels,” in *Working Notes of the ECML PKDD’08 Workshop on Mining Multidimensional Data*, Antwerp, Belgium, 2008.
- [154] G. Tsoumakas, I. Katakis, and I. Vlahavas, “Mining multi-label data,” in *Data Mining and Knowledge Discovery Handbook*, O. Maimon and L. Rokach, Eds. Berlin: Springer, 2010, pp. 667–686.
- [155] —, “Random k-labelsets for multi-label classification,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 23, no. 7, pp. 1079–1089, 2011.
- [156] G. Tsoumakas, E. Spyromitros-Xioufis, J. Vilcek, and I. Vlahavas, “MULAN: A java library for multi-label learning,” *Journal of Machine Learning Research*, vol. 12, no. Jul, pp. 2411–2414, 2011.
- [157] G. Tsoumakas and I. Vlahavas, “Random k-labelsets: an ensemble method for multilabel classification,” in *Lecture Notes in Artificial Intelligence 4701*, J. N. Kok, J. Koronacki, R. L. de Mantaras, S. Matwin, D. Mladenič, and A. Skowron, Eds. Berlin: Springer, 2007, pp. 406–417.
- [158] G. Tsoumakas, M.-L. Zhang, and Z.-H. Zhou, “Tutorial on learning from multi-label data,” in *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, Bled, Slovenia, 2009 [<http://www.ecmlpkdd2009.net/wp-content/uploads/2009/08/learning-from-multi-label-data.pdf>].
- [159] N. Ueda and K. Saito, “Parametric mixture models for multi-label text,” in *Advances in Neural Information Processing Systems 15*, S. Becker, S. Thrun, and K. Obermayer, Eds. Cambridge, MA: MIT Press, 2003, pp. 721–728.

- [160] L. C. van der Gaag and P. R. de Waal, “Multi-dimensional Bayesian network classifiers,” in *Proceedings of the 3rd European Workshop in Probabilistic Graphical Models*, Prague, Czech Republic, 2006, pp. 107–114.
- [161] A. Veloso, W. Meira Jr., M. Gonçalves, and M. Zaki, “Multi-label lazy associative classification,” in *Lecture Notes in Artificial Intelligence 4702*, J. N. Kok, J. Koronacki, R. L. de Mantaras, S. Matwin, D. Mladenič, and A. Skowron, Eds. Berlin: Springer, 2007, pp. 605–612.
- [162] C. Vens, J. Struyf, L. Schietgat, S. Džeroski, and H. Blockeel, “Decision trees for hierarchical multi-label classification,” *Machine Learning*, vol. 73, no. 2, pp. 185–214, 2008.
- [163] H. Wang, C. Ding, and H. Huang, “Multi-label classification: Inconsistency and class balanced k -nearest neighbor,” in *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, Atlanta, GA, 2010, pp. 1264–1266.
- [164] —, “Multi-label linear discriminant analysis,” in *Lecture Notes in Computer Science 6316*, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Berlin: Springer, 2010, pp. 126–139.
- [165] H. Wang, H. Huang, and C. Ding, “Image annotation using multi-label correlated Green’s function,” in *Proceedings of the 12th International Conference on Computer Vision*, Kyoto, Japan, 2009, pp. 2029–2034.
- [166] —, “Multi-label feature transform for image classifications,” in *Lecture Notes in Computer Science 6314*, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Berlin: Springer, 2010, pp. 793–806.
- [167] H. Wang, M. Huang, and X. Zhu, “A generative probabilistic model for multi-label classification,” in *Proceedings of the 8th IEEE International Conference on Data Mining*, Pisa, Italy, 2008, pp. 628–637.
- [168] J. Wang, Y. Zhao, X. Wu, and X.-S. Hua, “Transductive multi-label learning for video concept detection,” in *Proceedings of the ACM International Conference on Multimedia Information Retrieval*, Vancouver, Canada, 2008, pp. 298–304.
- [169] —, “A transductive multi-label learning approach for video concept detection,” *Pattern Recognition*, vol. 44, no. 10-11, pp. 2274–2286, 2011.
- [170] M. Wang, X. Zhou, and T.-S. Chua, “Automatic image annotation via local multi-label classification,” in *Proceedings of the 7th ACM International Conference on Image and Video Retrieval*, Niagara Falls, Canada, 2008, pp. 17–26.
- [171] Z. Wang, Y. Hu, and L.-T. Chia, “Multi-label learning by image-to-class distance for scene classification and image annotation,” in *Proceedings of the ACM International Conference on Image and Video Retrieval*, Xi’an, China, 2010, pp. 105–112.
- [172] A. Wiczorkowska, P. Synak, and Z. W. Raś, “Multi-label classification of emotions in music,” in *Advances in Soft Computing 35*, M. A. Kłopotek, S. T. Wierzchon, and K. Trojanowski, Eds. Berlin: Springer, 2006, pp. 307–315.
- [173] D. H. Wolpert, “Stacked generalization,” *Neural Networks*, vol. 5, no. 2, pp. 241–259, 1992.
- [174] F. Wu, Y. Han, Q. Tian, and Y. Zhuang, “Multi-label boosting for image annotation by structural grouping sparsity,” in *Proceedings of the 18th ACM International Conference on Multimedia*, Firenze, Italy, 2010, pp. 15–24.
- [175] S. Xiang, S. Chen, and L. Qiao, “Sparse representation: Extract adaptive neighborhood for multilabel classification,” in *Lecture Notes in Artificial Intelligence 6230*, B.-T. Zhang and M.-A. Orgun, Eds. Berlin: Springer, 2010, pp. 304–314.
- [176] R. Yan, J. Tešić, and J. R. Smith, “Model-shared subspace boosting for multi-label classification,” in *Proceedings of the 13th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, San Jose, CA, 2007, pp. 834–843.
- [177] B. Yang, J.-T. Sun, T. Wang, and Z. Chen, “Effective multi-label active learning for text categorization,” in *Proceedings of the 15th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Paris, France, 2009, pp. 917–925.
- [178] Y. Yang and S. Gopal, “Multilabel classification with meta-level features in a learning-to-rank framework,” *Machine Learning*, vol. 88, no. 1-2, pp. 47–68, 2012.

- [179] J. Ye, “Characterization of a family of algorithms for generalized discriminant analysis on undersampled problems,” *Journal of Machine Learning Research*, vol. 6, no. Apr, pp. 483–502, 2005.
- [180] Z. Younes, F. Abdallah, and T. Denoeux, “An evidence-theoretic k -nearest neighbor rule for multi-label classification,” in *Lecture Notes in Artificial Intelligence 5785*, L. Godo and A. Pugliese, Eds. Berlin: Springer, 2009, pp. 297–308.
- [181] Z. Younes, F. Abdallah, T. Denoeux, and H. Snoussi, “A dependent multilabel classification method derived from the k -nearest neighbor rule,” *EURASIP Journal on Advances in Signal Processing*, 2011, Article 645964.
- [182] K. Yu, S. Yu, and V. Tresp, “Multi-label informed latent semantic indexing,” in *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Salvador, Brazil, 2005, pp. 258–265.
- [183] J. H. Zaragoza, L. E. Sucar, E. F. Morales, C. Bielza, and P. Larrañaga, “Bayesian chain classifiers for multidimensional classification,” in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, Barcelona, Spain, 2011, pp. 2192–2197.
- [184] Z.-J. Zha, T. Mei, J. Wang, Z. Wang, and X.-S. Hua, “Graph-based semi-supervised learning with multi-label,” in *Proceedings of the 2008 IEEE International Conference on Multimedia & Expo*, Hannover, Germany, 2008, pp. 1321–1324.
- [185] —, “Graph-based semi-supervised learning with multiple labels,” *Journal of Visual Communication and Image Representation*, vol. 20, no. 2, pp. 97–103, 2009.
- [186] M.-L. Zhang, “ML-RBF: RBF neural networks for multi-label learning,” *Neural Processing Letters*, vol. 29, no. 2, pp. 61–74, 2009.
- [187] M.-L. Zhang, J. M. Peña, and V. Robles, “Feature selection for multi-label naive bayes classification,” *Information Sciences*, vol. 179, no. 19, pp. 3218–3229, 2009.
- [188] M.-L. Zhang and K. Zhang, “Multi-label learning by exploiting label dependency,” in *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Washington D. C., 2010, pp. 999–1007.
- [189] M.-L. Zhang and Z.-H. Zhou, “A k -nearest neighbor based algorithm for multi-label classification,” in *Proceedings of the 1st IEEE International Conference on Granular Computing*, Beijing, China, 2005, pp. 718–721.
- [190] —, “Multilabel neural networks with applications to functional genomics and text categorization,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 10, pp. 1338–1351, 2006.
- [191] —, “ML-kNN: A lazy learning approach to multi-label learning,” *Pattern Recognition*, vol. 40, no. 7, pp. 2038–2048, 2007.
- [192] —, “Multi-label learning by instance differentiation,” in *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, Vancouver, Canada, 2007, pp. 669–674.
- [193] X. Zhang, Q. Yuan, S. Zhao, W. Fan, W. Zheng, and Z. Wang, “Multi-label classification without the multi-label cost,” in *Proceedings of the 10th SIAM International Conference on Data Mining*, Columbus, OH, 2010, pp. 778–789.
- [194] Y. Zhang and J. Schneider, “Multi-label output codes using canonical correlation analysis,” in *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, 2011, pp. 873–882.
- [195] —, “Maximum margin output coding,” in *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, UK, 2012, pp. 1575–1582.
- [196] Y. Zhang and Z.-H. Zhou, “Multi-label dimensionality reduction via dependency maximization,” in *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, Chicago, IL, 2008, pp. 1503–1505.

- [197] ———, “Multi-label dimensionality reduction via dependency maximization,” *ACM Transactions on Knowledge Discovery from Data*, vol. 4, no. 3, 2010, Article 14.
- [198] T. Zhou and D. Tao, “Multi-label subspace ensemble,” in *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, Canary Islands, Spain, 2012, pp. 1444–1452.
- [199] T. Zhou, D. Tao, and X. Wu, “Compressed labeling on distilled labelsets for multi-label learning,” *Machine Learning*, vol. 88, no. 1-2, pp. 69–126, 2012.
- [200] Z.-H. Zhou, *Ensemble Methods: Foundations and Algorithms*. Boca Raton, FL: Chapman & Hall/CRC, 2012.
- [201] Z.-H. Zhou and M.-L. Zhang, “Multi-instance multi-label learning with application to scene classification,” in *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. Platt, and T. Hoffman, Eds. Cambridge, MA: MIT Press, 2007, pp. 1609–1616.
- [202] Z.-H. Zhou, M.-L. Zhang, S.-J. Huang, and Y.-F. Li, “Multi-instance multi-label learning,” *Artificial Intelligence*, vol. 176, no. 1, pp. 2291–2320, 2012.
- [203] S. Zhu, X. Ji, W. Xu, and Y. Gong, “Multi-labelled classification using maximum entropy method,” in *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Salvador, Brazil, 2005, pp. 274–281.
- [204] X. Zhu and A. B. Goldberg, “Introduction to semi-supervised learning,” in *Synthesis Lectures to Artificial Intelligence and Machine Learning*, R. J. Brachman and T. G. Dietterich, Eds. San Francisco, CA: Morgan & Claypool Publishers, 2009, pp. 1–130.