A Novel Online Stacked Ensemble for Multi-Label Stream Classification

Alican Büyükçakır
Bilkent Information Retrieval Group
Computer Engineering Department
Bilkent University
alicanbuyukcakir@bilkent.edu.tr

Hamed Bonab
College of Information and Computer
Sciences
University of Massachusetts Amherst
bonab@cs.umass.edu

Fazli Can
Bilkent Information Retrieval Group
Computer Engineering Department
Bilkent University
canf@cs.bilkent.edu.tr

ABSTRACT

As data streams become more prevalent, the necessity for online algorithms that mine this transient and dynamic data becomes clearer. Multi-label data stream classification is a supervised learning problem where each instance in the data stream is classified into one or more pre-defined sets of labels. Many methods have been proposed to tackle this problem, including but not limited to ensemble-based methods. Some of these ensemble-based methods are specifically designed to work with certain multi-label base classifiers; some others employ online bagging schemes to build their ensembles. In this study, we introduce a novel online and dynamically-weighted stacked ensemble for multi-label classification, called GOOWE-ML, that utilizes spatial modeling to assign optimal weights to its component classifiers. Our model can be used with any existing incremental multi-label classification algorithm as its base classifier. We conduct experiments with 4 GOOWE-ML-based multi-label ensembles and 7 baseline models on 7 real-world datasets from diverse areas of interest. Our experiments show that GOOWE-ML ensembles yield consistently better results in terms of predictive performance in almost all of the datasets, with respect to the other prominent ensemble models.

CCS CONCEPTS

• **Information systems** → **Data stream mining**; • **Computing methodologies** → **Ensemble methods**; *Online learning settings*;

KEYWORDS

Multi-label; data stream; supervised learning; classification; ensemble learning; bagging; online learning

ACM Reference Format:

Alican Büyükçakır, Hamed Bonab, and Fazli Can. 2018. A Novel Online Stacked Ensemble for Multi-Label Stream Classification. In *The 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*, October 22–26, 2018, Torino, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3269206.3271774>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '18, October 22–26, 2018, Torino, Italy

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6014-2/18/10...\$15.00

<https://doi.org/10.1145/3269206.3271774>

1 INTRODUCTION

The traditional supervised learning task is single-label, i.e. a data instance is classified into one label λ among a disjoint set of labels $\mathcal{L}$. However, this may not be the case for some real-world data. For instance, *The Big Lebowski* can simultaneously be classified as a *crime*, *comedy*, and *cult* movie. In such settings where an instance can be classified into a subset of labels, $L^* \subseteq \mathcal{L}$, the learning paradigm is called Multi-label Learning (MLL).

As of 2017, it is estimated that around 4.9 billion connected devices are generating data and this number is expected to rise to 25 billion by 2020 [28]. With such a rate of increase in the number of data in the form of streams, it becomes more and more important to extract meaningful information from seemingly chaotic data. Some of these data streams are multi-label, which led to MLL algorithms that can cope with streaming settings (i.e. with time and memory constraints, as well as changes in the distribution of the data over time) being developed.

MLL algorithms have drawn considerable attention over the last decades by accomplishing strong results in diverse areas including bioinformatics [1, 9], text classification [17] and image scene classification [7]. A considerable number of MLL algorithms resort to ensemble methods to increase their predictive performances [21, 23, 24, 33]. However, these methods have usually employed online bagging schemes for ensemble construction (where some of them utilized change detection mechanisms [2] as an upgrade to these bagging-based ensembles). To the best of our knowledge, there are very few stacked ensembles for multi-label stream classification, and most of the stacked ensembles in the literature are designed for and can only work with specific types of MLL algorithms. In this paper, we propose a novel stacked ensemble that is agnostic of the type of multi-label classifier that is used within the ensemble.

The main contributions of this paper are as follows: We (1) introduce a batch-incremental, online stacked ensemble for multi-label stream classification, GOOWE-ML, that can work with any incremental multi-label classifier as its component classifiers; (2) construct an $|\mathcal{L}|$ dimensional space to represent the relevance scores of classifiers of the ensemble, and utilize this construction to assign optimum weights to the model's component classifiers; (3) conduct experiments on 7 real-world datasets to compare GOOWE-ML with 7 state-of-the-art ensemble methods; (4) apply statistical tests to show that our model outperforms the state-of-the-art multi-label ensemble models. Additionally, we discuss how and why some multi-label classifiers yield poor Hamming Scores while performing considerably well on the rest of the performance metrics (e.g. accuracy, F1 Score). All in all, we argue that GOOWE-ML is well-suited

Table 1: Symbols and Notation for Multi-Label Stream Classification

Symbol	Meaning
M	Number of attributes in a data instance
L	Number of labels in the labelset of a data instance
N	Number of instances in the data stream
$\mathcal{X}$	Input attribute space. $\mathcal{X} = \mathbb{R}^M$
x	A data instance. $x = \langle x_1, x_2, \dots, x_i, \dots, x_M \rangle \in \mathcal{X}$
$\mathcal{L}$	Set of all possible labels. $\mathcal{L} = \{\lambda_1, \lambda_2, \dots, \lambda_L\}$
y	Label relevance vector. $y = \langle y_1, y_2, \dots, y_j, \dots, y_L \rangle = \{0, 1\}^L$
$\hat{y}$	Predicted relevance vector. $\hat{y} = h(x) = \langle \hat{y}_1, \hat{y}_2, \dots, \hat{y}_j, \dots, \hat{y}_L \rangle = [0, 1]^L$
$d_t = (x^t, y^t)$	The data point that arrives at time t
$\mathcal{D}$	Possibly infinite data stream. $\mathcal{D} = d_0, d_1, \dots, d_t, \dots, d_N$

for the multi-label stream classification task, and it is a valuable addition to the present day models.

The rest of the paper is organized as follows: Section 2 gives preliminaries on multi-label stream classification. Section 3 introduces the most widely used multi-label algorithms and ensemble techniques in the literature. In Section 4, our ensemble, GOOWE-ML, is described with the theory behind it. After mentioning the experimental setup, evaluation metrics and datasets in Section 5, the results are presented and discussed in Section 6. Lastly, the paper is concluded with insights and possible future work.

2 PRELIMINARIES

MLL is considered to be a hard task by nature, as the output space increases exponentially with the number of labels, since there are 2^L possible outcomes of classification for the labelset of size L . The high dimensionality of the label space causes increased computational cost, execution time and memory consumption. *Multi-label stream classification* (MLSC [33]) is the version of this task that takes place on data streams.

Data stream $\mathcal{D}$ is the set of data that has a temporal dimension and is possibly infinite. $\mathcal{D} = d_0, d_1, \dots, d_t, \dots, d_N$ where d_t is the data point in time t and d_N is the lastly seen data point in the data stream. The knowledge of lastly seen data point d_N is not known a priori, and it is only there to indicate the end of the processed data instances for evaluation purposes. Each data point d_t is of form $d_t = (x, y)$ where x is a data instance and y is its labelset (label relevance vector). The data instance x is a vector represented as $x = \langle x_1, x_2, \dots, x_i, \dots, x_M \rangle$, and each $x_i \in \mathcal{X}$. The labelset y is a vector represented as $y = \langle y_1, y_2, \dots, y_j, \dots, y_L \rangle$, and each $y_j \in \{0, 1\}$. Here, $y_j = 1$ means the j th label is relevant, and 0 otherwise. A prediction (hypothesis) of a multi-label classifier is $\hat{y} = h(x)$ that is of form $\hat{y} = \langle \hat{y}_1, \hat{y}_2, \dots, \hat{y}_j, \dots, \hat{y}_L \rangle$ and $\hat{y} \in [0, 1]^L$ meaning that the prediction vector consists of relevance probabilities (relevance scores) for each label. For the final decision of classification and evaluation, the prediction vector is sent to *de-fuzzification*, typically done by thresholding the relevance scores [26].

3 RELATED WORK

Comprehensive reviews on multi-label learning can be found at [13, 26, 37], on ensemble learning for data streams at [14, 15] and ensemble of multi-label classifiers at [18]. In this paper, we discuss the state-of-the-art multi-label methods, and focus on how these methods are used in ensemble learners for data streams.

3.1 Multi-label Methods

As a widely accepted taxonomy in the field of MLL, there are two general methods [37] of tackling a multi-label classification problem:

3.1.1 Problem Transformation. In *Problem Transformation*, the multi-label problem is transformed into more well-understood and simpler problems.

The most widely used Problem Transformation method is the Binary Relevance (BR) [31] where the multi-label problem is transformed into $|\mathcal{L}|$ distinct binary classification problems. After the transformation is applied to the dataset, any off-the-shelf binary classification algorithm can be utilized to get individual outputs corresponding to each binary problem. It scales linearly with respect to the number of labels, which makes it an efficient choice for practical purposes. However, it has been discussed [24, 37] that BR inherently fails to capture label-wise interdependencies.

To capture dependencies among labels and overcome this weakness of BR, some other BR-based methods are developed; most notably Classifier Chains (CC) [24] where BR classifiers are randomly permuted and linked in a chain-like manner in which each BR classifier yields its output to its connected neighbor classifier as an attribute. It is claimed that this helps the classifiers to capture the label dependencies, as each classifier in the chain learns not only the data itself, but also the label associations of every previous classifier in the chain.

Another common method of Problem Transformation is Label Powerset (LP) [31] method where each possible subset of labels is treated as a single label to translate the initial problem into single-label classification task with a bigger set of labels (hence, having multi-class problem of size $2^{|\mathcal{L}|}$). Pruned Sets (PS) [23] is an LP-based technique where the instances with infrequent label sets are pruned from the dataset. This allows only the instances with the most important subsets of labels to be considered for classification. Afterwards, the pruned instances are recycled back into an auxiliary dataset for another phase of classification; but for every subset of their relevant labels, instead of their initial relevant labels.

3.1.2 Algorithm Adaptation. In *Algorithm Adaptation*, existing algorithms are modified to be compatible with the multi-label setting.

In ML-KNN [36], the k-Nearest Neighbor algorithm is modified by counting the number of relevant labels for each neighboring instance to acquire posterior relevance probabilities for labels.

In ML-DT [9], the split criterion of C4.5 decision trees is modified by introducing the concept of multi-label entropy. In streaming

Figure 1: A stacked multi-label ensemble for stream classification. For each data instance, associated labels are shown with the geometric shapes ($\square$, $\circ$, and so on). A shape is colored if that label is relevant. Component classifiers (C_1, C_2, C_3, C_4) generate their own predictions, and these predictions are combined by the combiner algorithm of the ensemble.

environments, however, Hoeffding Trees are the common choice for decision trees. Hoeffding Trees [11] are incremental decision trees that have a theoretical guarantee that their output will become asymptotically identical to that of a regular decision tree as more and more data instances arrive. Modifying the split criterion of Hoeffding Trees for multi-label entropy, Multi-label Hoeffding Trees [22] are developed. More recently, a novel decision tree based method, iSOUP-Trees (incremental Structured Output Prediction Tree) [19] are proposed where adaptive perceptrons are placed in the leaves of incremental trees and the perceptrons' weights are used in producing a prediction that is a linear combination of the input's attributes.

In a nutshell, "in Problem Transformation, data is modified to make it suitable for algorithms; whereas in Algorithm Adaptation, algorithms are modified to make them suitable for data" [37].

3.2 Ensembles in MLL and MLSC

One of the most commonly used ensemble methods is Bagging where each classifier in an ensemble is trained with a bootstrap sample (a data sample that has the same size with the dataset, but each data point is randomly drawn with replacement). This assumes that the whole dataset is available, which is not the case for data stream environments. However, observing that the probability of having K many of a certain data point in a bootstrap sample is approximately $\text{Poisson}(1)$ for big datasets, each incoming data instance in a data stream can be weighted proportional to $\text{Poisson}(1)$ distribution to mimic bootstrapping in an online setting [20]. This is called Online Bagging, or OzaBagging, and it has been widely used in MLSC. In fact, the phrase 'Ensemble of' in the field usually means that it is the OzaBagged version of the base classifier that is mentioned. EBR [24], ECC [24], EPS [23] and EBRT [19] (Ensembles of BR, CC, PS and iSOUP Regression Trees respectively) are examples of this convention.

Additionally, it is common for the ensembles that use OzaBagging to also use a concept change detection mechanism called ADWIN (Adaptive Windowing) [2]. ADWIN keeps a variable-length window of the most recent items in the data stream to detect diversions from the average of the values in the window. Therefore, whenever ADWIN detects a change, the worst classifier in the OzaBag is reset. This is called ADWIN Bagging [22].

To the best of our knowledge, stacked ensemble models in the field of MLSC are very rare. A general scheme for a stacked ensemble for MLSC is given in Figure 1. Predictions of the component classifiers of an ensemble should be combined by a function (a meta-classifier) which will generate the final prediction of the ensemble. One can use either raw confidence scores of the labels for each instance, or predictions for each label and their counts (majority voting scheme) as the contributions of each component. How to optimally combine the contributions of each classifier is still a question in MLSC. Stacked ensembles that are proposed in the field are as follows:

SWMEC [33] is a weighted ensemble that is designed for ML-KNN as its base classifier. Its weight adjustment scheme utilizes distances in ML-KNN to obtain a confidence coefficient. IBR (Improved BR) [21] employs a feature extension mechanism in which the outputs of a BR classifier is firstly weighted by the accuracy of that classifier, and then added as a new feature to the data instance. New BR classifiers are trained from the data with extended feature spaces. The functionalities of these models involve algorithm-specific properties and therefore cannot be extended to any other base classifier. Such models are constrained by the success of their base classifiers. In [34], the authors followed an unorthodox approach and created a *label-based* ensemble instead of a chunk-based one, which tackled the class imbalance problem that exists in the multi-label datasets as well as concept drifts. Recently, in ML-AMRules [27], multi-label classification task is interpreted as a rule learning task and the rule learners are combined in an ensemble that uses online bagging (called ML-Random Rules).

All in all, ensemble models in MLSC are not explored thoroughly. Base multi-label classifiers are either combined with Online Bagging or ADWIN Bagging, or with stacked combination schemes that depends on the type of the base classifier. There is a lack of online ensembles in the field that can work with any type of multi-label base classifier which also involve a smart combination scheme. GOOWE-ML addresses this inadequacy.

4 GOOWE-ML

We propose GOOWE-ML (Geometrically Optimum Online Weighted Ensemble for Multi-Label Classification): a batch-incremental (chunk-based) and dynamically-weighted online ensemble that can be used with any incremental multi-label learner that yields confidence outputs for predicting relevant labels for an incoming data instance.

Table 2: Additional Symbols and Notation for GOOWE-ML

Symbol	Meaning
K	Number of component classifiers in the ensemble, i.e. ensemble size
n	Number of data points in the instance window I
h	Maximum capacity of a data chunk DC
C_k	k th component classifier in the ensemble. $1 \leq k \leq K$
ξ	Ensemble of classifiers. $\xi = \{C_1, C_2, \dots, C_k, \dots, C_K\}$
w	Weight vector for the ensemble ξ . $w = \langle W_1, W_2, \dots, W_k, \dots, W_K \rangle$
s_k^i	Relevance scores for the k th classifier for i th instance in the ensemble. $s_k^i = \langle S_{k1}^i, S_{k2}^i, \dots, S_{kj}^i, \dots, S_{kL}^i \rangle$
S	Relevance scores matrix. Each relevance score s_{kj} is an element in this matrix. $S \in \mathbb{R}^{K \times L}$
I	Instance window of size n , having latest n data instances. $I = d_1, d_2, \dots, d_n$
DC	Data chunk that consists of the latest h data points. $DC = d_1, d_2, \dots, d_h$

Let the multi-label classifiers in the ensemble be $\{C_1, C_2, \dots, C_K\}$. For each incoming data instance, each classifier C_k generates relevance score vector s_k , which consists of the relevance scores of each label for that instance, i.e. $s_k = \langle S_{k1}, S_{k2}, \dots, S_{kL} \rangle$. The relevance score vectors of classifiers for each instance is stored in the rows of matrix S , which will be used to populate elements of the matrix A and the vector d (see Eqn.4 and 5, and Alg.2:7-8).

4.1 Ensemble Maintenance

Let ξ denote the ensemble that is initially empty. A new classifier is trained at each incoming data chunk, as well as the existing ones (if any). The ensemble grows as the new classifiers from incoming data chunks are introduced, until the maximum ensemble size is reached (i.e. ensemble is full). Then, the newly trained classifier replaces one of the old classifiers in the ensemble. This replacement is often times done by removing the temporally oldest component or the most poorly performed component with respect to some metric [16]. In GOOWE-ML, this replacement is done by re-weighting the component classifiers and removing the component with the lowest weight (Alg 1:7-12). Analogous model management systems are employed in both ensembles for single-label classification such as Accuracy Weighted Ensemble (AWE) [32] and Accuracy Updated Ensemble (AUE2) [8]; and for multi-label classification such as SWMEC [33].

Having a fixed number of base classifiers in the ensemble prevents the model to swell in terms of memory usage. Also, training new classifiers from each data chunk allows the ensemble to notice new trends in the distribution of the data, and thus, be more robust against concept drifts.

In addition to fixed-sized data chunks, GOOWE-ML also uses a sliding window for stream evaluation purposes, which consists of the most recently seen n instances. Size of the instance window can be smaller than the size of each data chunk, i.e. $n \leq h$, so that higher resolution can be obtained for the prequential evaluation. Prequential evaluation is discussed in more detail in Experimental Setup section.

4.2 Weight Assignment and Update

In our geometric framework, we represent the relevance scores s_k of each component classifier in our ensemble as vectors in an L -dimensional space. Previously, Tai & Lin [29] used a similar approach, which they called Principal Label-Space Transformation,

Figure 2: Transformation into label space in GOOWE-ML. Relevance scores of the components (red): $S_1 = \langle 0.65, 0.35 \rangle$ and $S_2 = \langle 0.82, 0.18 \rangle$. The optimal vector $\bar{y}$ (blue): $y = \langle 1, 1 \rangle$, generated from the ground truth. Weighted prediction of the ensemble: $S\bar{w}$ (green). The distance between $\bar{y}$ and $S\bar{w}$ is minimized.

to interpret the existing multi-label algorithms in a geometrical setting and reduce the high dimensionality of the multi-label data. Bonab & Can [5] adapted an analogous setting to investigate the optimal ensemble size for single-label data stream classification. In GOOWE-ML, this spatial modeling is used to assign optimal weights for component classifiers in the ensemble.

Geometrically, an intuitive explanation of our spatial modeling and weighting scheme is shown in Figure 2 for the 2-dimensional case, i.e. when $L = 2$. After representing relevance scores in the label space, GOOWE-ML minimizes the Euclidean distance between the combined relevance vector, $\hat{y}$, and the ideal vector that represents the ground truth, y , in the label space. Analogously, Wu & Crestani [35] utilized this approach in the field of Data Fusion to optimally combine the query results, and Bonab & Can [6] in the field of single-label data stream classification, both with successful results. This is equivalent to the following *linear least squares problem*.

$$\min_{\bar{w}} \|\bar{y} - S\bar{w}\|_2^2 \quad (1)$$

Here, S is the relevance scores matrix, w is the weight vector which is to be determined, and y is the vector representing the

ground truth for a given data point. In other words, our objective function to be minimized is the following:

$$f(W_1, W_2, \dots, W_K) = \sum_{i=1}^n \sum_{j=1}^L \left(\sum_{k=1}^K (W_k S_{kj}^i - y_j^i) \right)^2 \quad (2)$$

Taking a partial derivative of W and setting the gradient to zero, i.e. $\nabla f = 0$, we get:

$$\sum_{k=1}^K W_k \left(\sum_{i=1}^n \sum_{j=1}^L S_{qj}^i S_{kj}^i \right) = \sum_{i=1}^n \sum_{j=1}^L y_j^i S_{qj}^i \quad (3)$$

Equation 3 is of the form $Aw = d$ where A is a square matrix of size $K \times K$ with elements:

$$a_{qk}^i = \sum_{i=1}^n \sum_{j=1}^L S_{qj}^i S_{kj}^i \quad (1 \leq q, k \leq K) \quad (4)$$

and d is the remainder vector of size K with elements:

$$d_q^i = \sum_{i=1}^n \sum_{j=1}^L y_j^i S_{qj}^i \quad (1 \leq q \leq K) \quad (5)$$

Therefore, solving the equation $Aw = d$, for w , gives us the optimally adjusted weight vector. The weight vector w is updated at the end of each data chunk, where the components in the ensemble are trained from the instances in the data chunk, as well. Also, notice that this update operation resembles the Batch Gradient Descent in a way that w is updated at the end of each batch, having trained from the instances in the batch. However, unlike Batch Gradient Descent, this weight update scheme does not take steps towards better weights iteratively, but rather finds the optimal weights directly after solving the linear system $Aw = d$. As a consequence, the updated weights do not depend on the previous values in w , they only depend on the performance of the components on the latest chunk. This allows the ensemble to capture sudden changes in the distribution of the data.

4.3 Multi-label Prediction

The ensemble's prediction for the i th example, $\hat{y}^i$, is the weighted sum of its components' relevance scores, s_k .

$$\hat{y}_j^i(\xi) = \sum_{k=1}^K w_k S_{kj}^i \quad (1 \leq j \leq L) \quad (6)$$

Here, each relevance score, S_{kj}^i , is normalized beforehand into the range of $[0, 1]$ by the following normalization:

$$S_{kj}^i \leftarrow \frac{S_{kj}^i}{\sum_{j=1}^L S_{kj}^i} \quad (1 \leq j \leq L) \quad (7)$$

After normalization, the relevance scores sum up to 1. The final prediction of the classifier is obtained by thresholding the relevance scores by $(1/L)$, which is the expected prior relevance probability of a label of a data instance.

$$\hat{y}_j^i \leftarrow \begin{cases} 1, & \text{if } \hat{y}_j^i > \frac{1}{L} \\ 0, & \text{otherwise} \end{cases} \quad 1 \leq j \leq L \quad (8)$$

These three operations (Weighted Voting (Eqn.6), Normalization (Eqn.7) and Thresholding (Eqn.8)) are done consecutively and can be considered as one atomic operation in the algorithm, shown as *predict()* in the pseudocode (see Alg.1:5).

Algorithm 1 GOOWE-ML: Geometrically Optimum Online Weighted Ensemble for Multi-Label Classification

Require: $\mathcal{D}$: data stream, DC : latest data chunk, K : maximum number of classifiers in the ensemble, C : a multi-label classifier in the ensemble,

Ensure: ξ : ensemble of weighted classifiers, $\hat{y}$: multi-label prediction of the ensemble as combined score vector.

```

1:  $\xi \leftarrow \emptyset$ ;
2:  $A, d \leftarrow \text{null}, \text{null}$ 
3: while  $\mathcal{D}$  has more instances do
4:    $d_i \leftarrow$  current data instance
5:    $\hat{y} \leftarrow \text{predict}(d_i, \xi)$  {Eqn.6, 7 and 8}
6:   if  $DC$  is full then
7:      $C_{in} \leftarrow$  new component classifier built on  $DC$ ;
8:     if  $\xi$  has  $K$  classifiers then
9:        $A', d' \leftarrow \text{TrainOptimumWeights}(DC, \xi, \text{null}, \text{null})$ 
10:       $w \leftarrow \text{solve}(A'w = d')$ ;
11:       $C_{out} \leftarrow$  classifier  $C_k$  with minimum  $w_k$ 
12:       $\xi \leftarrow \xi - C_{out}$ 
13:    end if
14:     $\xi \leftarrow \xi \cup C_{in}$ 
15:    Train all classifiers  $C \in \xi - C_{in}$  with  $DC$ 
16:  end if
17: end while
```

Algorithm 2 GOOWE-ML: Train Optimum Weights

Require: DC : one or more data instances, ξ : Ensemble of Classifiers, A : square matrix, d : remainder vector

Ensure: The matrix A and the vector d , ready for optimum weight assignment

```

1: if  $A$  is null or  $d$  is null then
2:   Initialize square matrix  $A$  of size  $K \times K$ 
3:   Initialize remainder vector  $d$  of size  $K$ 
4: end if
5: for all instances  $x^t \in DC$  do
6:    $y^t \leftarrow$  true relevance vector of  $x^t$  {To be used in Eqn.5}
7:    $A \leftarrow A + A^t$ ; {Eqn.4}
8:    $d \leftarrow d + d^t$ ; {Eqn.5}
9: end for
```

4.4 Complexity Analysis

Let the prediction of a component classifier in the ensemble take $O(c)$ time. Also, notice that the ensemble size is of order $O(K)$, since ensemble is not fully formed only for the first K chunks and the size is always K afterwards.

For a data chunk, each component classifier predicts each data instance, which takes $O(hKc)$ time, as the size of each data chunk in the stream are the same and h . At the same time, the square matrix

A and the remainder vector d are filled using each pair of relevance scores of the components for each label and each instance, which takes $O(hK^2L)$ time. Then, the linear system $Aw = d$ is solved, where A is of size K . Solving this linear system with no complex optimization methods take at most $O(K^3)$ time [4] (where there are more complex but asymptotically better methods). This loop continues for N/h many chunks. Thus, the whole process has the complexity of:

$$O\left(\frac{N}{h}(hKc + hK^2L) + K^3\right) = O\left(N\left(Kc + K^2L + \frac{K^3}{h}\right)\right) \quad (9)$$

Here, the term with (Kc) , (K^2L) and (K^3/h) represents the time complexity of prediction, training and optimal weight assignment respectively.

c is generally small, since most of the models use Hoeffding Trees and its derivatives as their base classifiers. Therefore, the term with (K^2L) dominates the sum there in the Eqn. 9. When the terms (K^2L) and (K^3/h) are compared, it can be also noticed that the former always dominates the latter: h 's magnitude is of hundreds or thousands, whereas K 's magnitude is generally of tens. As a result, L (a whole number) will be always higher than (K/h) (a fraction that is < 1). As a consequence, the algorithm has overall $O(NK^2L)$ complexity.

5 EXPERIMENTAL DESIGN

5.1 Datasets

To understand how densely multi-labeled a dataset is, *Label Cardinality* and *Label Density* are used. Label Cardinality is the average number of relevant labels of the instances in $\mathcal{D}$; Label Density is the Label Cardinality per number of labels [30], indicating the percentage of labels that are relevant on average.

$$LC(\mathcal{D}) = \frac{1}{N} \sum_{i=1}^N |y^i| \quad LD(\mathcal{D}) = \frac{LC(\mathcal{D})}{L} = \frac{1}{LN} \sum_{i=1}^N |y^i|$$

Our experiments are conducted on 7 datasets¹ that are from diverse application domains (genes, newspapers, aviation safety reports and so on), given in Table 3. These datasets are extensively used in the literature [19, 22, 27].

Table 3: Table of Multi-Label Datasets

Source $\mathcal{D}$	Domain	N	M	L	LC($\mathcal{D}$)	LD($\mathcal{D}$)
20NG ^b	Text	19,300	1,006	20	1.020	0.051
Yeast ⁿ	Biology	2,417	103	14	4.237	0.303
Ohsumed ^b	Text	13,529	1,002	23	1.660	0.072
Slashdot ^b	Text	3,782	1,079	22	1.180	0.053
Reuters ⁿ	Text	6,000	500	101	2.880	0.028
IMDB ^b	Text	120,919	1,001	28	2.000	0.071
TMC2007 ^b	Text	28,596	500	22	2.160	0.098

The superscripts after the name of the dataset indicates that the features in that dataset is binary (^b) or numeric (ⁿ).

¹Datasets are downloaded from MEKA's webpage. Available at: <https://sourceforge.net/projects/meka/files/Datasets/>.

5.2 Evaluating Multi-label Learners

Multi-label evaluation metrics that are widely used throughout the studies in the field are divided into two groups [37]: (1) *Instance-Based Metrics*, (2) *Label-Based Metrics*. These two metrics indicate how well the algorithms perform. In addition to these, efficiency of the performing algorithms can be measured, which indicates how much resources they consume. Hence, (3) *Efficiency Metrics* is added to the evaluation. In the Tables 4 and 5, $\uparrow$ ($\downarrow$) next to the metric indicates that the corresponding metric's score is to be maximized (minimized).

5.2.1 Instance-Based Metrics. Instance-based metrics are evaluated for every instance and averaged over the whole dataset. Exact Match, Hamming Score, and Instance-Based {Accuracy, Precision, Recall, F1-Score} [37] are used in this study.

5.2.2 Label-Based Metrics. Label-based metrics are evaluated for every label and averaged over examples within each individual label. Macro and Micro-Averaged Precision, Recall and F1 Score [13] are used in this study.

5.2.3 Efficiency Metrics. Finally, to measure the efficiency of the algorithms, the execution time and memory consumption of each algorithm are monitored.

5.3 Experimental Setup

Experiments are implemented in MOA [3], utilizing multi-label methods in MEKA [25]. The evaluation of each algorithm is *prequential* [12]. An incoming data instance is first tested by classifiers (see Alg.1:5); evaluation measures corresponding the prediction are recorded, and then, that data instance is used to train classifiers, as well as the updated weighting scheme(see Alg.1:??). This is also called *Interleaved-Test-Then-Train (ITTT)* approach and is widely common in algorithms in streaming settings.

If an ensemble is batch-incremental, then the ensemble is trained at the end of each batch (i.e. whenever a data chunk is filled). The evaluation of ensembles are started after the first learner in the ensemble is formed. We used fixed number of 10 classifiers as the ensemble size, mimicking the previously conducted experiments to enable comparison [19, 22]. For incremental evaluation of the classifiers, we used window-based evaluation with the window size {100, 250, 500, 1000} according to the size of datasets².

All experiments are conducted on a machine with an Intel Xeon E3-1200 v3 @ 3.40GHz processor and 16GB DDR3 RAM.

We experimented with 4 GOOWE-ML models (referred with their abbreviations from now onward):

- **GOBR**: the components use BR Transformation.
- **GOCC**: the components use CC Transformation.
- **GOPS**: the components use PS Transformation.
- **GORT**: the components use ISOUP Regression Trees.

We have 7 baseline models. Four of the baselines use fixed-sized windows with no concept drift detecting mechanism: **EBR** [24], **ECC** [24], **EPS** [23], **EBRT** [19], whereas 3 of them use ADWIN as their concept drift detector: **EaBR** [22], **EaCC** [22], and **EaPS**

²The source code is available at <https://github.com/abuyukcakilr/gooweml>. The results can be reproduced using the aforementioned datasets. The program outputs the predictive performance measures, time and memory consumption, as well as incremental evaluations for each window.

[22]). In all models, BR and CC transformations use a Hoeffding Tree classifier whereas the PS transformation uses a Naive Bayes classifier.

5.4 Evaluation of Statistical Significance

We evaluated the aforementioned algorithms using multi-label example-based, and label-based evaluation metrics, as well as efficiency metrics. To check the statistical significance among the algorithms, we used *Friedman test with Nemenyi post-hoc analysis* [10]. We applied the Friedman test with $\alpha = 0.05$ where the null hypothesis is that all of the measurements come from the same distribution. If the null hypothesis is failed, then Nemenyi post-hoc analysis is applied to see which algorithms that performed statistically significantly better than which others.

The result of Friedman-Nemenyi Test can be seen in the *Critical Distance Diagrams* where each algorithm is sorted according to their average ranks for a given metric on a number line, and the algorithms that are within the critical distance of each other (that are not statistically significantly better than each other) are linked with a line. These diagrams compactly show Nemenyi Significance. Better models have lower average rank, and therefore on the right side of a Critical Distance Diagram. The Critical Distance for Nemenyi Significance is calculated as follows [10]:

$$CD = q_{\alpha, m} \sqrt{\frac{m(m+1)}{6|\mathcal{D}|}} \quad (10)$$

where m is the number of models that are being compared, and $|\mathcal{D}|$ number of datasets that are experimented on. Plugging in $m = 11$, $q_{\alpha=0.05, m=11} = 3.219$ (from Critical Values Table for Two-Tailed Nemenyi Test ³) and $|\mathcal{D}| = 7$, we get $CD = 5.707$ as our Critical Distance.

6 RESULTS AND DISCUSSION

6.1 Predictive Performance

Example-Based F1 Score, Micro-Averaged F1 Score, Hamming Score and Example-Based Accuracy are given in Table 4 for each model on each dataset. Winner models of each dataset for the metrics are shown in bold in the table. Precision and Recall scores are omitted, since we report the F1 Scores in Table 4, which is calculated as the harmonic mean of the two. Exact Match scores are omitted, since it is a very strict metric and the scores tend to be near zero for each algorithm especially when $|\mathcal{L}|$ is large.

Before starting to analyze models individually, let us look at the big picture: It is apparent that the predictive performance of a streaming multi-label model highly depends on the dataset. Looking at the results, no single model is clearly better than the rest, regardless of the dataset that it has run on. For instance, PS transformation-based ensembles (GOPS, EPS and EaPS) did relatively better in the Slashdot, Reuters and IMDB datasets; whereas the ensembles with BR and CC transformations were clearly superior in the Yeast, Ohsumed and TMC2007 datasets.

It can be observed in the Table 4 and Figure 3, GOOWE-ML-based classifiers performed better than Online Bagging and ADWIN Bagging models consistently over all datasets. Especially **GOCC**

and **GOPS** placed 1st and 2nd respectively, in every performance metric, except Hamming Score. More detailed discussion on the Hamming Score and its relation to the Precision and Recall scores of the models are provided below in a separate section.

Read et al. [22] previously claimed that instance-incremental methods are better than batch-incremental methods in the MLSC task. However, our experimental evidence shows that our batch-incremental ensemble performs better than the state-of-the-art instance-incremental models in almost every performance metric (again, except Hamming Score).

6.2 Efficiency

Results for the Execution Time and Memory Consumption of the models, and the corresponding Critical Distance Diagrams are given in Table 5 and Figure 4, respectively. It is clear that time and memory efficiency of an MLSC ensemble is highly correlated with the problem transformation method that its component classifiers use. Ensembles that used PS Transformation (GOPS, EPS, EaPS) are ranked consistently higher in terms of both time and memory efficiency. Indeed, as it can be seen in Figure 4, EPS and GOPS are among the top 3 for both of the metrics.

Models with iSOUP Tree are among the fastest, but their memory consumption is significantly high compared to the PS-based ensembles. Considering GORT and EBRT's relatively underwhelming predictive performance (see Figure 4), PS-based ensembles should be preferable over ensembles of iSOUP Regression Trees.

BR and CC Transformation-based models performed similarly within datasets across ensembling techniques. Their execution time and memory consumption is nearly identical with a few exceptions (where ADWIN Bagging models had significantly lower memory consumption due to resetting component classifiers many times). Having similar resource consumptions, GOCC can be preferred due to its greater predictive performance.

6.3 On Hamming Scores in Datasets with Large Labelsets

Consider the prediction and the ground truth vector of a given data instance. Let TP , FP , FN and TN denote the number of true positives, false positives, false negatives and true negatives. For instance, FP is the number of labels that are predicted as relevant but are not. Then, Hamming Score for that instance can be calculated as $\frac{TP+TN}{TP+FP+FN+TN}$.

For a multi-label dataset with a fairly large labelset and low label density, TN in the numerator and the denominator will dominate this score and Hamming Score will yield mis-interpretable results. Take IMDB dataset ($|\mathcal{L}| = 28$, $LD(\mathcal{D}) = 0.071$) for example: **GOPS** was the clear winner in terms of $F1_{ex}$, Acc_{ex} and $F1_{micro}$, yielding 20% better scores than its closest competitor (which was **GOBR**). Despite performing this well, GOPS had a considerably low Hamming Score (**0.836**) with respect to the Online Bagging-based models (all of them around **0.928**). Here, one can argue that perhaps the Hamming Score is the true indicator of success in MLL, and therefore Online Bagging-based models performed better. However, this hypothesis cannot be correct, since even a dummy classifier that predicts every single label as irrelevant (0) yields Hamming Score of **0.929**! Additionally, the reason why GOOWE-ML-based

³Available at: http://www.cin.ufpe.br/~fatc/AM/Nemenyi_critval.pdf

Table 4: Experimental Results: Example and Label-based Metrics

	20NG	Yeast	Ohsumed	Slashdot	Reuters	IMDB	TMC7		20NG	Yeast	Ohsumed	Slashdot	Reuters	IMDB	TMC7		
(a) Example-Based F1 Score ($F1_{ex}$) ↑								Avg. Rank	(b) Micro-Averaged F1 Score ($F1_{micro}$) ↑								Avg. Rank
GOBR	0.364	0.650	0.307	0.189	0.076	0.283	0.623	4.00	GOBR	0.237	0.638	0.291	0.187	0.076	0.276	0.584	4.86
GOCC	0.442	0.652	0.352	0.028	0.145	0.221	0.668	<u>2.57</u>	GOCC	0.516	0.640	0.410	0.050	0.196	0.228	0.634	<u>2.71</u>
GOPS	0.224	0.644	0.331	0.405	0.252	0.333	0.485	3.00	GOPS	0.206	0.629	0.298	0.315	0.210	0.314	0.447	3.43
GORT	0.196	0.607	0.297	0.189	0.078	0.283	0.452	5.71	GORT	0.153	0.598	0.270	0.187	0.077	0.277	0.439	6.57
EBR	0.365	0.638	0.23	0.023	0.106	0.075	0.654	4.71	EBR	0.499	0.631	0.294	0.041	0.141	0.099	0.638	4.29
ECC	0.349	0.632	0.217	0.020	0.098	0.016	0.643	6.43	ECC	0.486	0.625	0.280	0.037	0.134	0.025	0.631	6.14
EPS	0.096	0.584	0.213	0.269	0.148	0.133	0.330	6.71	EPS	0.115	0.584	0.216	0.286	0.162	0.138	0.342	7.00
EBRT	0.100	0.509	0.056	0.001	0.000	0.001	0.008	10.57	EBRT	0.174	0.519	0.076	0.001	0.000	0.001	0.008	10.58
EaBR	0.341	0.638	0.202	0.018	0.059	0.031	0.661	6.57	EaBR	0.477	0.632	0.266	0.033	0.081	0.041	0.640	5.71
EaCC	0.156	0.633	0.005	0.020	0.004	0.001	0.646	8.14	EaCC	0.262	0.627	0.007	0.037	0.007	0.002	0.632	7.71
EaPS	0.109	0.578	0.200	0.258	0.183	0.104	0.384	6.85	EaPS	0.180	0.580	0.205	0.278	0.200	0.118	0.378	6.71
(c) Hamming Score ↑								Avg. Rank	(d) Example-Based Accuracy (Acc_{ex}) ↑								Avg. Rank
GOBR	0.749	0.769	0.738	0.625	0.707	0.727	0.886	9.86	GOBR	0.239	0.508	0.184	0.106	0.040	0.164	0.457	4.57
GOCC	0.952	0.771	0.932	0.946	0.984	0.887	0.916	5.57	GOCC	0.391	0.509	0.277	0.025	0.120	0.138	0.515	<u>3.00</u>
GOPS	0.769	0.754	0.830	0.872	0.956	0.836	0.854	9.29	GOPS	0.137	0.504	0.211	0.299	0.160	0.204	0.327	3.29
GORT	0.624	0.716	0.730	0.644	0.720	0.732	0.815	10.57	GORT	0.115	0.454	0.178	0.107	0.040	0.164	0.298	6.71
EBR	0.961	0.786	0.936	0.946	0.986	0.925	0.934	2.14	EBR	0.352	0.502	0.191	0.020	0.098	0.055	0.520	4.29
ECC	0.961	0.786	0.936	0.947	0.986	0.928	0.934	<u>1.57</u>	ECC	0.337	0.493	0.180	0.018	0.093	0.012	0.511	6.14
EPS	0.924	0.764	0.918	0.937	0.985	0.919	0.911	7.29	EPS	0.094	0.460	0.180	0.260	0.143	0.105	0.246	6.29
EBRT	0.952	0.773	0.930	0.946	0.986	0.929	0.902	4.00	EBRT	0.100	0.372	0.049	0.001	0.000	0.001	0.007	10.57
EaBR	0.961	0.786	0.935	0.946	0.986	0.928	0.935	2.00	EaBR	0.330	0.502	0.169	0.016	0.056	0.024	0.529	6.14
EaCC	0.955	0.787	0.928	0.947	0.986	0.929	0.934	2.29	EaCC	0.152	0.495	0.004	0.018	0.004	0.001	0.516	7.71
EaPS	0.950	0.767	0.918	0.937	0.985	0.924	0.913	6.71	EaPS	0.108	0.455	0.170	0.250	0.179	0.083	0.290	6.43

Figure 3: Critical Distance Diagrams for the Predictive Performance Metrics (given in Table 4).

models have smaller Hamming Scores is that they have high FP (hence low TN) values in the contingency table. In other words, GOOWE-ML-based models are *Low Precision - High Recall* models. They eagerly predict labels as relevant. On the other hand, Online Bagging-based models are *High Precision - Low Recall* models. They predict few labels as relevant at each data instance. As a consequence, they are more confident about their predictions, but they

miss many relevant labels due to being more conservative. This dichotomy is shown for 3 datasets with low label densities in Table 6, where the higher value among Precision and Recall is shown in bold for each model and dataset.

Observing these, we claim that *in datasets with large labelset and low label density, High Recall models may have considerably low Hamming Scores due to the nature of the metric. Hence, Hamming*

Table 5: Experimental Results: Efficiency Metrics

	20NG	Yeast	Ohsumed	Slashdot	Reuters	IMDB	TMC7		20NG	Yeast	Ohsumed	Slashdot	Reuters	IMDB	TMC7		
(a) Execution Time (seconds) ↓								Avg. Rank	(b) Memory Consumption (MB) ↓								Avg. Rank
GOBR	2,631	28	2,310	537	2,366	31,769	1,942	8.86	GOBR	1,685.82	18.40	1,364.32	381.98	1,029.23	4,384.09	780.58	7.57
GOCC	2,591	33	2,314	544	2,555	34,348	1,990	9.71	GOCC	1,429.26	24.85	1,229.32	351.74	1,261.34	6,284.62	748.33	7.43
GOPS	670	8	522	129	115	5,098	181	3.14	GOPS	76.30	2.03	43.55	29.38	15.57	75.68	41.88	3.00
GORT	390	47	435	68	412	3,719	333	4.14	GORT	431.76	198.81	656.16	77.92	660.38	542.24	227.70	5.71
EBR	2,246	25	1,934	488	1,917	(*) 101,243	1,769	6.71	EBR	2,152.97	17.56	1,775.72	545.72	1,425.88	(*) 22,119.42	1,289.87	9.00
ECC	2,270	29	1,958	495	2,057	(*) 48,325	1,789	7.71	ECC	2,171.09	28.99	1,792.66	549.33	1,539.00	(*) 22,380.89	1,305.07	10.57
EPS	383	5	299	99	46	2,168	109	1.43	EPS	8.40	0.97	8.53	10.38	3.67	7.55	6.34	1.29
EBRT	338	63	404	64	389	3,919	264	3.43	EBRT	521.96	274.61	809.76	76.70	943.59	1,826.87	234.06	6.57
EaBR	2,376	35	1,997	488	1,968	20,675	2,220	7.86	EaBR	1,997.59	17.57	1,678.32	399.99	1,330.31	3,522.14	93.60	7.29
EaCC	2,041	40	1,622	503	2,062	17,148	2,292	7.71	EaCC	373.03	26.39	295.09	549.35	652.92	661.76	135.09	5.86
EaPS	2,393	24	1,862	363	402	14,361	574	5.29	EaPS	6.25	1.50	15.80	12.15	8.05	13.53	2.56	1.71

Note. Measurements that are marked with (*) are conducted on a machine with 256 GB RAM. Consistency among the rankings is preserved, as the obtained results do not change the rankings of the efficiency metrics for the IMDB dataset.

Figure 4: Critical Distance Diagrams for the Efficiency Metrics (given in Table 5).

Score may not be the true indicator of the predictive performance while evaluating multi-label models.

Table 6: Micro Precision (Prec) vs Recall (Rec), and Their Effect on Hamming Score (HS)

	20NG			Ohsumed			Reuters		
	Prec	Rec	HS	Prec	Rec	HS	Prec	Rec	HS
GOBR	0.140	0.757	0.749	0.181	0.743	0.738	0.040	0.848	0.707
GOPS	0.125	0.580	0.769	0.212	0.500	0.830	0.140	0.418	0.956
EBR	0.753	0.373	0.961	0.713	0.185	0.936	0.510	0.082	0.986
EPS	0.142	0.096	<u>0.924</u>	0.348	0.157	<u>0.918</u>	0.361	0.105	<u>0.985</u>

Two GOOWE-ML models and two Online Bagging Models with different problem transformation types are picked. Higher Precision and lower Recall resulted in better Hamming Scores consistently.

This hypothesis helps explaining why GOBR and GOCC performed poorly in terms of Hamming Scores in the datasets with relatively lower label densities (such as Slashdot, Reuters and IMDB datasets) whereas both of them were clear winners in the overall predictive performance.

6.4 Window-Based Evaluation

Figure 5 presents window-based evaluation for two datasets and three models, where the sliding window size is equal to the chunk size, i.e. $n = h$. To this end, we evaluate each window's performance using $F1_{ex}$ measurement. For each group of models, the best performing strategy is chosen for the given dataset—e.g. in Reuters

dataset, GOPS, EPS and EaPS performed the best among GOOWE-ML, Online Bagging and ADWIN Bagging models, respectively.

It can be seen that GOOWE-ML-based models do not predict in the first evaluation window, since no training has been done while waiting the first chunk to be filled. In both of the datasets, we observe the optimal weight assignment strategy in effect: after the first few chunks, GOOWE-ML-based model's predictions continually yield better performance than its competitors.

7 CONCLUSION

We present an online batch-incremental multi-label stacked ensemble, GOOWE-ML, that constructs a spatial model for the relevance scores of its classifiers, and uses this model to assign optimal weights to its component classifiers. Our experiments show that GOOWE-ML models outperform the most prominent Online Bagging and ADWIN Bagging models. Two of the GOOWE-ML-based ensembles especially stand out: **GOCC** is the clear winner in terms of overall predictive performance, ranking first in Acc_{ex} , $F1_{ex}$ and $F1_{micro}$ scores. **GOPS**, on the other hand, is the best compromise between predictive performance and resource consumption among all models, yielding strong performance with very conservative time and memory requirements. In addition, we argue that Hamming Score can be deceptively low for models with low Precision and high Recall and support this claim by experimental evidence.

In the future, we plan to investigate optimal ensemble size for MLSC in relation to the dimensionality of the feature set, number of labels, and label cardinality and density. We also plan to study

Figure 5: Window-Based Evaluation of Models: Example-Based F1 Score for Reuters and 20NG datasets.

the performance of GOOWE-ML for concept-evolving multi-label data streams, in which the labelset can be updated with new labels.

REFERENCES

- [1] Zafer Barutcuoglu, Robert E Schapire, and Olga G Troyanskaya. 2006. Hierarchical Multi-Label Prediction of Gene Function. *Bioinformatics* 22, 7 (2006), 830–836.
- [2] Albert Bifet and Ricard Gavalda. 2007. Learning from Time-changing Data with Adaptive Windowing. In *Proceedings of the 2007 SIAM International Conference on Data Mining*. SIAM, 443–448.
- [3] Albert Bifet, Geoff Holmes, Bernhard Pfahringer, Jesse Read, Philipp Kranen, Hardy Kremer, Timm Jansen, and Thomas Seidl. 2011. MOA: A Real-Time Analytics Open Source Framework. In *ECML PKDD*. Springer, 617–620.
- [4] A Bojańczyk. 1984. Complexity of Solving Linear Systems in Different Models of Computation. *SIAM J. Numer. Anal.* 21, 3 (1984), 591–603.
- [5] Hamed R Bonab and Fazli Can. 2016. A Theoretical Framework on the Ideal Number of Classifiers for Online Ensembles in Data Streams. In *Proceedings of the 25th ACM CIKM International Conference on Information and Knowledge Management*. ACM, 2053–2056.
- [6] Hamed R Bonab and Fazli Can. 2018. GOOWE: Geometrically Optimum and Online-Weighted Ensemble classifier for evolving data streams. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12, 2 (2018), 25.
- [7] Matthew R Boutell, Jiebo Luo, Xipeng Shen, and Christopher M Brown. 2004. Learning Multi-Label Scene Classification. *Pattern Recognition* 37, 9 (2004), 1757–1771.
- [8] Dariusz Brzezinski and Jerzy Stefanowski. 2014. Reacting to Different Types of Concept Drift: The Accuracy Updated Ensemble Algorithm. *IEEE Transactions on Neural Networks and Learning Systems* 25, 1 (2014), 81–94.
- [9] Amanda Clare and Ross D King. 2001. Knowledge Discovery in Multi-Label Phenotype Data. In *European Conference on Principles of Data Mining and Knowledge Discovery*. Springer, 42–53.
- [10] Janez Demšar. 2006. Statistical Comparisons of Classifiers over Multiple Data Sets. *Journal of Machine Learning Research* 7, Jan (2006), 1–30.
- [11] Pedro Domingos and Geoff Hulten. 2000. Mining High-speed Data Streams. In *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 71–80.
- [12] João Gama, Raquel Sebastião, and Pedro Pereira Rodrigues. 2009. Issues in Evaluation of Stream Learning Algorithms. In *the 15th ACM SIGKDD*. ACM, 329–338.
- [13] Eva Gibaja and Sebastián Ventura. 2014. Multi-Label Learning: A Review of the State of the Art and Ongoing Research. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 4, 6 (2014), 411–444.
- [14] Heitor Murilo Gomes, Jean Paul Barddal, Fabrício Enembreck, and Albert Bifet. 2017. A Survey on Ensemble Learning for Data Stream Classification. *ACM Comput. Surv.* 50, 2, Article 23 (March 2017), 36 pages.
- [15] Bartosz Krawczyk, Leandro L Minku, Joao Gama, Jerzy Stefanowski, and Michał Woźniak. 2017. Ensemble Learning for Data Stream Analysis: A Survey. *Information Fusion* 37 (2017), 132–156.
- [16] Ludmila I Kuncheva. 2004. Classifier Ensembles for Changing Environments. In *International Workshop on Multiple Classifier Systems*. Springer, 1–15.
- [17] David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. 2004. Rcv1: A New Benchmark Collection for Text Categorization Research. *Journal of Machine Learning Research* 5, Apr (2004), 361–397.
- [18] Jose M Moyano, Eva L Gibaja, Krzysztof J Cios, and Sebastián Ventura. 2018. Review of Ensembles of Multi-Label Classifiers: Models, Experimental Study and Prospects. *Information Fusion* 44 (2018), 33–45.
- [19] Aljaž Osojnik, Panče Panov, and Sašo Džeroski. 2017. Multi-label Classification via Multi-Target Regression on Data Streams. *Machine Learning* 106, 6 (2017), 745–770.
- [20] Nikunj C Oza. 2005. Online Bagging and Boosting. In *2005 IEEE International Conference on Systems, Man and Cybernetics*, Vol. 3. IEEE, 2340–2345.
- [21] Wei Qu, Yang Zhang, Junping Zhu, and Qiang Qiu. 2009. Mining Multi-Label Concept-Drifting Data Streams Using Dynamic Classifier Ensemble. In *Asian Conference on Machine Learning*. Springer, 308–321.
- [22] Jesse Read, Albert Bifet, Geoff Holmes, and Bernhard Pfahringer. 2012. Scalable and Efficient Multi-Label Classification for Evolving Data Streams. *Machine Learning* 88, 1-2 (2012), 243–272.
- [23] Jesse Read, Bernhard Pfahringer, and Geoff Holmes. 2008. Multi-label Classification Using Ensembles of Pruned Sets. In *The Eighth IEEE ICDM*. IEEE, 995–1000.
- [24] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. 2009. Classifier Chains for Multi-Label Classification. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 254–269.
- [25] Jesse Read, Peter Reutemann, Bernhard Pfahringer, and Geoff Holmes. 2016. MEKA: A Multi-label/Multi-target Extension to WEKA. *Journal of Machine Learning Research* 17, 21 (2016), 1–5.
- [26] Mohammad S Sorower. 2010. *A Literature Survey on Algorithms for Multi-Label Learning*. Technical Report. Oregon State University, Corvallis.
- [27] Ricardo Sousa and João Gama. 2018. Multi-label Classification from High-Speed Data Streams with Adaptive Model Rules and Random Rules. *Progress in Artificial Intelligence* (2018), 1–11.
- [28] CACM Staff. 2017. Big Data. *Commun. ACM* 60, 6 (May 2017), 24–25.
- [29] Farbound Tai and Hsuan-Tien Lin. 2012. Multilabel Classification with Principal Label Space Transformation. *Neural Computation* 24, 9 (2012), 2508–2542.
- [30] Grigoris Tsoumakas and Ioannis Katakis. 2006. Multi-Label Classification: An Overview. *International Journal of Data Warehousing and Mining* 3, 3 (2006).
- [31] Grigoris Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. 2011. Random k-Labelsets for Multilabel Classification. *IEEE Transactions on Knowledge and Data Engineering* 23, 7 (2011), 1079–1089.
- [32] Haixun Wang, Wei Fan, Philip S Yu, and Jiawei Han. 2003. Mining Concept-Drifting Data Streams using Ensemble Classifiers. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 226–235.
- [33] Lulu Wang, Hong Shen, and Hui Tian. 2017. Weighted Ensemble Classification of Multi-label Data Streams. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 551–562.
- [34] Peng Wang, Peng Zhang, and Li Guo. 2012. Mining Multi-Label Data Streams using Ensemble-based Active Learning. In *Proceedings of the 2012 SIAM international Conference on Data Mining*. SIAM, 1131–1140.
- [35] Shengli Wu and Fabio Crestani. 2015. A Geometric Framework for Data Fusion in Information Retrieval. *Information Systems* 50, Supplement C (2015), 20–35.
- [36] Min-Ling Zhang and Zhi-Hua Zhou. 2007. ML-KNN: A Lazy Learning Approach to Multi-Label Learning. *Pattern Recognition* 40, 7 (2007), 2038–2048.
- [37] Min-Ling Zhang and Zhi-Hua Zhou. 2014. A Review on Multi-Label Learning Algorithms. *IEEE Transactions on Knowledge and Data Engineering* 26, 8 (2014), 1819–1837.