A Robust Experimental Evaluation of Automated Multi-Label Classification Methods

Supplementary Material: Proposed *Search Spaces*

Alex G. C. de Sá

Computer Science Department
Universidade Federal de Minas Gerais
Belo Horizonte, MG, Brazil
alexgcsa@dcc.ufmg.br

Gisele L. Pappa

Computer Science Department
Universidade Federal de Minas Gerais
Belo Horizonte, MG, Brazil
glpappa@dcc.ufmg.br

Cristiano G. Pimenta

Computer Science Department
Universidade Federal de Minas Gerais
Belo Horizonte, MG, Brazil
cgpimenta@dcc.ufmg.br

Alex A. Freitas

School of Computing
University of Kent
Canterbury, Kent, United Kingdom
a.a.freitas@kent.ac.uk

ABSTRACT

This supplementary material aims to describe the proposed multi-label classification (MLC) *search spaces* based on the MEKA and WEKA softwares. First, we overview 26 MLC algorithms and meta-algorithms in MEKA, presenting their main characteristics, such as hyper-parameters, dependencies and constraints. Second, we review 28 single-label classification (SLC) algorithms, preprocessing algorithms and meta-algorithms in the WEKA software. These SLC algorithms were also studied because they are part of the proposed MLC *search spaces*. Fundamentally, this occurs due to the problem transformation nature of several MLC algorithms used in this work. These algorithms transform an MLC problem into one or several SLC problems in the first place and solve them with SLC model(s) in a next step. Therefore, understanding their main characteristics is crucial to this work. Finally, we present a formal description of the *search spaces* by proposing a context-free grammar that encompasses the 54 learning algorithms. This grammar basically comprehends the possible combinations, the constraints and dependencies among the learning algorithms.

1 STUDYING THE SEARCH SPACE OF MULTI-LABEL CLASSIFICATION ALGORITHMS

We perform a study of 26 multi-label and meta multi-label classification algorithms from the MEKA software [54], which are described in the following two sections.

Table 1 lists these 26 MLC algorithms and present their association to the different versions of the designed *search spaces*, showing their names, names' acronyms and their respective types. Table 1 also indicates if the MLC algorithm in the row belongs or not ('Y' for yes and 'N' for no) to the respective *search space* in the column (i.e., if the MLC *search space* comprehends that algorithm), and how many hyper-parameters (#HP) it encompasses (when it is used).

It is also important to mention that the algorithms in (this version of) MEKA can define a **threshold** to perform the classification using the model's confidence outputs (typically, class probabilities). For the general multi-label context, it is in general better to optimize the threshold than simply using an arbitrary threshold of 0.5 [18, 53].

This parameter (*pred_tshd*) [*-threshold*] could take the following values:

- Proportional cut method by instance (PCut1) [47]: It takes into account the label cardinality of the dataset, which is simply the average number of labels associated with each instance of this dataset. Thus, PCut1 automatically calibrates the prediction confidence threshold, by minimizing the difference between the label cardinality of the training set and the label cardinality obtained with a given set of predicted labels – where the latter set is determined by the threshold value. This does not require access to the true predictions in the test set.
- Proportional cut method by label (PCutL): It is used to calibrate the prediction confidence threshold the same way as PCut1, but for each label individually.
- The threshold could also take a unique real value between zero (0.0) and one (1.0) for all instances being classified. Formally, the threshold can also be defined by the following interval: $\{threshold \in \mathbb{R} \mid 0.0 < threshold < 1.0\}$.

Default value: PCut1.

2 SEARCH SPACE – TRADITIONAL MLC ALGORITHMS

This section covers the main traditional MLC algorithms in the MEKA software [54].

2.1 Binary Relevance

The standard binary relevance (BR) algorithm [61]. It creates a binary classification problem for each label and learns a model for each label individually. **Parameters:**

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.

There are no dependencies/constraints in BR.

2.2 The 'quick' version of Binary Relevance

The quick version of BR (BRq) [53] is a version of BR which is able to downsample the number of training instances across the binary

Table 1: Overview of the employed multi-label classification (MLC) algorithms from the MEKA software*.

id	Algorithm Name	Acronym	Type	Small		Medium		Large	
				Used?	#HP	Used?	#HP	Used?	#HP
1	Back Propagation Neural Network	ML-BPNN	AA	Y	4	Y	4	Y	4
2	Binary Relevance	BR	PT	Y	0	Y	0	Y	0
3	Classifier Chain	CC	PT	Y	0	Y	0	Y	0
4	Label Powerset	LP	PT	Y	0	Y	0	Y	0
5	Random K-Label Pruned Sets	RAkEL	PT	Y	4	Y	4	Y	4
6	Bayesian Classifier Chains	BCC	PT	N	-	Y	1	Y	1
7	Binary Relevance – Quick Version	BRq	PT	N	-	Y	1	Y	1
8	Classifier Chain – Quick Version	CCq	PT	N	-	Y	1	Y	1
9	Four-Class Pairwise Classification	FW	PT	N	-	Y	0	Y	0
10	Monte-Carlo Classifier Chains	MCC	PT	N	-	Y	3	Y	3
11	Probabilistic Classifier Chains	PCC	PT	N	-	Y	0	Y	0
12	Pruned Sets	PS	PT	N	-	Y	2	Y	2
13	Pruned Sets with Threshold	PSt	PT	N	-	Y	2	Y	2
14	Random K-Label Disjoint Pruned Sets	RAkELd	PT	N	-	Y	3	Y	3
15	Ranking and Threshold	RT	PT	N	-	Y	0	Y	0
16	Classifier Trellis	CT	PT	N	-	N	-	Y	6
17	Conditional Dependency Networks	CDN	PT	N	-	N	-	Y	2
18	Conditional Dependency Trellis	CDT	PT	N	-	N	-	Y	5
19	Population of MCC	PMCC	PT	N	-	N	-	Y	6
20	Bagging of Multi-Label Classifiers	BaggingML	Meta-MLC	N	-	N	-	Y	1
21	Bagging of Multi-Label Classifiers (Duplicate)	BaggingMLDup	Meta-MLC	N	-	N	-	Y	2
22	Classification Maximization	CM	Meta-MLC	N	-	N	-	Y	1
23	Ensemble of Multi-Label classifiers	EnsembleML	Meta-MLC	N	-	N	-	Y	2
24	Expectation Maximization	EM	Meta-MLC	N	-	N	-	Y	1
25	Random Subspace Multi-Label	RSML	Meta-MLC	N	-	N	-	Y	3
26	Subset Mapper	SM	Meta-MLC	N	-	N	-	Y	0

*All algorithm names are clickable links to their respective part in this supplementary material.

models. It is intended for use in an ensemble (but it works in a standalone fashion as well).

Parameters:

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.
- Down-sample ratio (*dsr*)[-P]: It is a ratio used to reduce the number of instances across the binary models. Low values mean more removals and high values mean less removals, as BRq uses the following formula $(1-dsr)*number\_of\_instances$ to calculate the number of instances to remove. This parameter is constrained by the interval:
 $\{dsr \in \mathbb{R} \mid 0.2 \leq dsr \leq 0.8\}$.

There is no explanation about this parameter in the original paper and in other papers in the multi-label classification literature. The justification – about the used interval – is that we would like to have (at least) 20% of the instances from the original data to construct the model (otherwise, the algorithm may not have sufficient instances to build the model). Additionally, we would like to have (at most) 80% of the instances from the original data in order to learn the classifier (otherwise, the algorithm would be very similar to BR).

Default value: 0.75.

There are no dependencies/constraints between the parameters of BRq.

2.3 Classifier Chain

The classifier chain (CC) algorithm [53] is also similar to BR, but the label outputs predicted by a classifier become new inputs for the next classifiers in the chain. It uses a single random order of labels in the chain.

Parameters:

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.

There are no dependencies/constraints in CC.

2.4 The ‘quick’ version of Classifier Chains

The quick version of CC (CCq) [53] is a version of CC which is able to down-sample the number of training instances across the binary models. It is also intended for use in an ensemble (but it works in a standalone fashion as well). **Parameters:**

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.
- Down-sample ratio (*dsr*)[-P]: It is a ratio used to reduce the number of instances across the binary models. Low values mean more removals and high values mean less removals, as CCq uses the following formula $(1-dsr)*number\_of\_instances$

to calculate the number of instances to remove. This parameter is constrained by the interval:

$$\{d_{sr} \in \mathbb{R} \mid 0.2 \leq d_{sr} \leq 0.8\}.$$

There is no explanation about this parameter in the original paper and in other papers in the multi-label classification literature. The justification – about the used interval – is that we would like to have (at least) 20% of the instances from the original data to construct the model (otherwise, the algorithm may not have sufficient instances to build the model). Additionally, we would like to have (at most) 80% of the instances from the original data in order to learn the classifier (otherwise, the algorithm would be very similar to CC).

Default value: 0.75.

There are no dependencies/constraints between the parameters of CCq.

2.5 Bayesian Classifier Chain

The Bayesian classifier chain (BCC) algorithm [65] creates a maximum spanning tree based on marginal dependencies, define a Bayesian network from it, and then employ a classifier chain (CC) using the order of the labels found in the Bayesian network model. The original paper used Naïve Bayes as a base classifier, but other types of classifiers can be used. **Parameters:**

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.
- Dependency type (*dp*)[-X]: The way to measure and find the dependencies. It may take ten categorical values: 1. C (co-occurrence counts); 2. I (mutual information); 3. Ib (mutual information using binary approximation); 4. IbF (Mutual information using fast binary approximation); 5. H (Conditional information); 6. HbF (Conditional information using fast binary approximation); 7. X (Chi-squared); 8. F (Frequencies); 9. No label dependence; 10. L (The “LEAD” method for finding conditional dependence).

Default value: IbF.

There are no dependencies/constraints between the parameters of BCC.

2.6 (Bayes Optimal) Probabilistic Classifier Chain

The probabilistic classifier chain (PCC) algorithm [15] acts exactly like CC at training time, but explores all possible paths as inference at test time (hence, “Bayes optimal”). **Parameters:**

- Base classifier (*bc*)[-W]: It can be any classifier from WEKA.

Dependencies/Constraints:

- (1) PCC has poor scalability, i.e., it is very slow when the number of labels is greater than a certain threshold. In the PCC’s original paper [15], it is said that this threshold should be 15 labels. In the future, we might consider it to scale the proposed solution to evolve multi-label learning algorithm. For instance, we must impose a constraint in the grammar that specifies the use of PCC only if the number of labels is less than 15. We can also specify a time budget for all MLC algorithms, depending on the size of the dataset. This will consequently limit the effectiveness of the PCC algorithm in more complex types of data.

2.7 Monte-Carlo Classifier Chains

The algorithms based on Monte-Carlo classifier chains, MCC and M2CC [48, 49], apply classifier chains with Monte Carlo optimization, using a maximum number of inference and chain-order trials. MCC has a tractable label prediction scheme only at the test time (MCC), whereas M2CC performs an additional search for the optimal chain sequence at the training time. **Parameters:**

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.
- Inference Iterations (*ii*)[-Iy]: The number of iterations to search the output space at test time. This parameter is bounded by the values in the interval: $\{ii \in \mathbb{Z} \mid 1 < ii \leq 100\}$.

Default value: 10.

- Chain Iterations (*chi*)[-Is]: The number of iterations to search the chain space at training time. This parameter is bounded by the values in the interval: $\{chi \in \mathbb{Z} \mid 1 < chi \leq 1500\}$. It can also take the value zero and the MCC algorithm is used instead of M2CC. This will happen with 50% of probability, i.e., MCC and M2CC have the same chances of being selected.

Default value: 0.

- Payoff function (*pof*)[-P]: It sets the payoff function to evaluate the chains when performing the search. It can take 23 values: 1. Accuracy; 2. Jaccard index; 3. Hamming score; 4. Exact match; 5. Jaccard distance; 6. Hamming loss; 7. Zero One loss; 8. Harmonic score; 9. One error; 10. Rank loss; 11. Average precision; 12. Log Loss limited by the number of labels; 13. Log loss limited by the number of instances; 14. Micro Precision; 15. Micro Recall; 16. Macro Precision; 17. Macro Recall; 18. F_1 micro averaged; 19. F_1 macro averaged by example; 20. F_1 macro averaged by label; 21. AUPRC macro averaged; 22. AUROC macro averaged; 23. Levenshtein distance.

Default value: Exact match.

There are no dependencies/constraints between the parameters in MCC and M2CC. Additionally, we studied the range of the parameters in the works of Read *et al.* [48, 49]. However, the authors did not employ a proper parameter tuning at the single-label level, neither at the multi-label level. In the work of [49], a search is performed to find the proper number of chain iterations in accordance to the payoff function. We are using part of this study to define the range of the parameters.

2.8 Population of Monte-Carlo Classifier Chains

The population of Monte-Carlo classifier chains (PMCC) [48, 49] is an algorithm that has similar properties when compared to MCC and M2CC. However, it is considered an extension of both algorithms. The difference is that PMCC creates a population of M chains at training time (from I_s candidate chains, using Monte Carlo sampling), and uses all of them at test time. This is not a typical majority-vote ensemble algorithm. The simulated annealing search [36] can also be applied to the chain structures (produced by MCC or M2CC) in order to find the best one.

Parameters:

- Base classifier (*bc*)[-W] : It can be any classifier from WEKA.
- Inference Iterations (*ii*)[-Iy]: The number of iterations to search the output space at test time. This parameter is bounded by

the values in the interval: $\{ii \in \mathbb{Z} \mid 1 < ii \leq 100\}$.

Default value: 10.

- Chain Iterations (*chi*)[-Is]: The number of iterations to search the chain space at training time. This parameter is bounded by the values in the interval: $\{chi \in \mathbb{Z} \mid 50 < chi \leq 1500\}$.

Default value: 50.

- Beta (β)[-B]: It sets the factor with which the temperature (and thus the acceptance probability of steps in the wrong direction in the search space) is decreased in each iteration of the simulated annealing search. This parameter is bounded by the interval: $\{\beta \in \mathbb{Z} \mid 0.01 \leq \beta \leq 0.99\}$.

Default value: 0.03.

- Temperature switch (*ts*)[-O]: It sets the use of simulated annealing search and, when it is activated, it cools the chain down over time (from the beginning of the chain). It may take the values zero (0) or one (1). The value zero (0) means that no temperature is used, i.e., the parameter β is ignored internally by PMCC. If using $ts = 1$, this sets the use of the β constant.

Default value: 0.

- Population size (*ps*)[-M]: It sets the population size. It should be always smaller than the total number of chains evaluated (Is). This parameter takes one of the values defined by the following interval:

$\{ps \in \mathbb{Z} \mid 1 \leq ps \leq 50\}$.

Default value: 10.

- Payoff function (*pof*)[-P]: It sets the payoff function to evaluate the chains when performing the search. It can take 23 values: 1. Accuracy; 2. Jaccard index; 3. Hamming score; 4. Exact match; 5. Jaccard distance; 6. Hamming loss; 7. Zero One loss; 8. Harmonic score; 9. One error; 10. Rank loss; 11. Average precisio; 12. Log Loss limited by the number of labels; 13. Log loss limited by the number of instances; 14. Micro Precision; 15. Micro Recall; 16. Macro Precision; 17. Macro Recall; 18. F_1 micro averaged; 19. F_1 macro averaged by example; 20. F_1 macro averaged by label; 21. AUPRC macro averaged; 22. AUROC macro averaged; 23. Levenshtein distance.

Default value: Exact match.

Dependencies/Constraints:

- (1) The parameter “population size” must be smaller than the parameter “chain iterations”.

Again, we studied the range of the parameters in the works of [49]. However, the authors did not employ a proper parameter tuning at the single-label level, neither at the multi-label level. During the work of [49], a search is performed to find the proper number of chain iterations in accordance to the payoff function. We are using part of this study to define the range of the parameters. Nevertheless, parameters β , temperature and population size are not properly studied for the multi-label scenario.

2.9 Classifier Trellis

The classifier trellis (CT) algorithm [50] builds classifier chains in a trellis structure (rather than a cascaded chain). It is possible to set the width and type/connectivity of the trellis, and optionally to change the payoff function which guides the placement of nodes (labels) within the trellis. **Parameters:**

- Base classifier (*bc*)[-W]: It can be any classifier from WEKA.
- Width (*w*)[-H]: it determines the width of the trellis (0 for chain, i.e., $w = L$; -1 for a square trellis, i.e., $w = \sqrt{L}$, always using the default floor function to convert it to an integer value). Thus, the trellis structure will always have w rows and L nodes, in total, connected using directed edges.

Default value: -1.

- Dependency type (*dp*)[-X]: The way to measure and find the label dependencies. It may take nine categorical values: 1. C (co-occurrence counts); 2. I (mutual information); 3. Ib (mutual information using binary approximation); 4. IbF (Mutual information using fast binary approximation); 5. H (Conditional information); 6. HbF (Conditional information using fast binary approximation); 7. X (Chi-squared); 8. F (Frequencies); 9. No label dependence.

Default value: IbF.

- Inference Iterations (*ii*)[-Iy]: The number of iterations to search the output space at test time. This parameter is bounded by the values in the interval: $\{ii \in \mathbb{Z} \mid 1 \leq ii \leq 100\}$.

Default value: 10.

- Chain Iterations (*chi*)[-Is]: The number of iterations to search the chain space at train time. This parameter is bounded by the values in the interval: $\{chi \in \mathbb{Z} \mid 1 < chi \leq 1500\}$.

Default value: 0.

- Density (*d*)[-L]: It determines the neighborhood density (the number of neighbors for each node in the trellis). The default value for the density parameter is one (1), and zero (0) indicates a BR classifier. Thus, this parameter is not allowed to take the value zero, being restricted by the interval: $\{d \in \mathbb{Z} \mid 1 \leq d \leq \sqrt{L} + 1\}$, where L is the total number of labels.

Default value: 1.

- Payoff function (*pof*)[-P]: It sets the payoff function to evaluate the chains when performing the search. It can take 23 values: 1. Accuracy; 2. Jaccard index; 3. Hamming score; 4. Exact match; 5. Jaccard distance; 6. Hamming loss; 7. Zero One loss; 8. Harmonic score; 9. One error; 10. Rank loss; 11. Average precisio; 12. Log Loss limited by the number of labels; 13. Log loss limited by the number of instances; 14. Micro Precision; 15. Micro Recall; 16. Macro Precision; 17. Macro Recall; 18. F_1 micro averaged; 19. F_1 macro averaged by example; 20. F_1 macro averaged by label; 21. AUPRC macro averaged; 22. AUROC macro averaged; 23. Levenshtein distance.

Default value: Exact match.

Dependencies/Constraints:

- (1) If the width $w = L$ ($w = 0$), the density $d = 1$. Otherwise, if $w = \sqrt{L}$ ($w = -1$), the density d should be $\sqrt{L} + 1$, at most, i.e., $d \leq \sqrt{L} + 1$.

2.10 Conditional Dependency Networks

The conditional dependency networks (CDN) algorithm [26] builds a fully connected undirected network, where each node (label) is connected to each other node (label). Each node is a binary classifier that predicts $p(y_j|x, y_1, \dots, y_{j-1}, \dots, y_L)$. Then, inference is done using the Gibbs Sampling method over I iterations. Additionally, the

final I_c iterations are used to collect the marginal probabilities, which become the prediction ($y[]$).

Parameters:

- Base classifier (bc)[-W] : It can be any classifier from WEKA.
- Iterations (i)[-I] : The total number of iterations to perform in CDT. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 100 < i \leq 1000\}$.
Default value: 1000.
- Collection iterations (ci)[-Ic] : The number of collection iterations used to compute the output class probabilities in the Gibbs Sampling method. The parameter ci is restricted by the interval: $\{ci \in \mathbb{Z} \mid 1 \leq ci \leq 100\}$.
Default value: 100.

Dependencies/Constraints:

- (1) The collections will happen just after $(i - ci)$ iterations. So, i should be substantially greater than ci in order to make the algorithm work properly.

2.11 Conditional Dependency Trellis

The conditional dependency trellis (CDT) algorithm [26, 50] is similar to the CDN approach. However, it constructs a trellis structure (like CT) instead of a fully connected network. **Parameters:**

- Base classifier (bc)[-W] : It can be any classifier from WEKA.
- Width (w)[-H] : it determines the width of the trellis (0 for chain, i.e., $w = L$; -1 for a square trellis, i.e., $w = \sqrt{L}$, always using the default floor function to convert it to an integer value). Thus, the trellis structure will always have w rows and L nodes, in total, connected using directed edges.
Default value: -1.
- Dependency type (dp)[-X] : The way to measure and find the label dependencies. It may take nine categorical values: 1. C (co-occurrence counts); 2. I (mutual information); 3. Ib (mutual information using binary approximation); 4. Ibf (Mutual information using fast binary approximation); 5. H (Conditional information); 6. Hbf (Conditional information using fast binary approximation); 7. X (Chi-squared); 8. F (Frequencies); 9. None (Using empty).
Default value: None.
- Density (d)[-L] : It determines the neighborhood density (the number of neighbors for each node in the trellis). The default value for the density parameter is one (1), and zero (0) indicates a BR classifier. Thus, this parameter is not allowed to take the value zero, being restricted by the interval: $\{d \in \mathbb{Z} \mid 1 \leq d \leq \sqrt{L} + 1\}$, where L is the total number of labels.
Default value: 1.
- Iterations (i)[-I] : The total number of iterations to perform in CDT. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 100 < i \leq 1000\}$.
Default value: 1000.
- Collection iterations (ci)[-Ic] : The number of collection iterations used to compute the output class probabilities in the Gibbs Sampling method. The parameter ci is restricted by the interval: $\{ci \in \mathbb{Z} \mid 1 \leq ci \leq 100\}$.
Default value: 100.

Dependencies/Constraints:

- (1) If the width $w = L$ ($w = 0$), the density $d = 1$. Otherwise, if $w = \sqrt{L}$ ($w = -1$), the density d should be $\sqrt{L} + 1$, at most, i.e., $d \leq \sqrt{L} + 1$.
- (2) The collections will happen just after $(i - ci)$ iterations. So, i should be substantially greater than ci in order to make the algorithm work properly.

2.12 Four-Class Pairwise Classification

The Four-class PairWise Classification (FW) algorithm [54] trains a multi-class base classifier for each pair of labels. Thus, the number of classifiers is $\frac{L*(L-1)}{2}$ in total (where L is the number of labels), each one with four possible class values (00,01,10,11) representing the possible combinations of relevant (1)/irrelevant (0) values for each label in the label pair. It uses a voting and a threshold scheme at testing time where, e.g., 01 from pair jk gives one vote to label k and any label with a number of votes above the threshold is considered relevant. It uses the same threshold specified in the Section 1 to define the relevance of a label. **Parameters:**

- Base classifier (bc)[-W] : It can be any classifier from WEKA.

There are no dependencies/constraints in FW.

2.13 Ranking and Threshold

The ranking and threshold (RT) algorithm [47] duplicates each multi-labeled example, and assigns one of the labels (only) to each copy. After that, it trains a regular multi-class base classifier. At test time, a threshold separates relevant from irrelevant labels using the posterior probability for each class value (i.e., label). It uses the same threshold specified in the Section 1 to define the relevance of a label. **Parameters:**

- Base classifier (bc)[-W] : It can be any classifier from WEKA.

There are no dependencies/constraints in RT.

2.14 Label Combination

The label combination (LC) algorithm [61], also known as label powerset (LP), treats each label combination as a single class in a multi-class learning scheme. The set of possible values of each class is the powerset of the set of labels.

- Base classifier (bc)[-W] : It can be any classifier from WEKA.

There are no dependencies/constraints in LC.

2.15 Pruned Sets

The pruned sets (PS) algorithm [46, 47] was created to use the power of LC's labelset-based paradigm, without the disadvantages of such algorithm. In order to do this, this algorithm has two important steps: a pruning step and a label-set subsampling step. The pruning step removes infrequently occurring label sets from the training data. This removes unnecessary complexity from the LC-transformed data by reducing the number of labelsets. Nevertheless, PS does not simply discard the pruned examples. Instead of doing that, PS subsamples the labelsets of these examples for label subsets which occur more frequently in the training data. It then attaches these label sets to the example, creating new examples and reintroducing them into the training. It subsamples these labelsets pv times to

produce pv new examples, where pv is the pruning value (defined in the followed items).

After these steps, it trains a standard LC classifier. The idea of the algorithm is to reduce the number of unique class values that would otherwise need to be learned by LC. PS achieves its best performance when used in an Ensemble (e.g., EnsembleML).

Parameters:

- Base classifier (bc)[-W] : It can be any classifier from WEKA.
- Pruning value (pv)[-P]: It defines an infrequent labelset as one which occurs less than p times in the data. $p = 0$ would mean that LC classifier is learned. Thus, this parameter is bounded by the following interval: $\{pv \in \mathbb{Z} \mid 1 \leq pv \leq 5\}$.

Default value: 0.

- Subsampling value (sv)[-N]: The label set of each pruned example (in accordance to the examples pruned by the use of the previous parameter, i.e., the pruning value) becomes a candidate for label-set subsampling. The PS algorithm subsamples the label sets of pruned examples to create examples which do meet the pruning criterion. So, the subsample value defines the (maximum) number of frequent labelsets to subsample from the infrequent labelsets. This parameter is bounded by the following interval: $\{sv \in \mathbb{Z} \mid 0 \leq sv \leq 5\}$.

Default value: 0.

There are no dependencies/constraints between the parameters of PS. Additionally, there is a proper study in the work of [47] about the range of these two parameters.

2.16 Pruned Sets with Threshold

The pruned sets algorithm with a threshold (PSt) [46, 47, 52], which is a modification of PS that can form new label sets at classification (i.e., test) time by using a threshold function. Given the posterior of the label classes (combinations) and the number of labels, it returns the distribution across labels. Using the threshold (defined in the Section 1) could make the algorithm to predict labelsets not seen in the training set, differently from PS. **Parameters:**

- Base classifier (bc)[-W] : It can be any classifier from WEKA.
- Pruning value (pv)[-P]: It defines an infrequent labelset as one which occurs less than p times in the data. $p = 0$ would mean that LC classifier is learned. Thus, this parameter is bounded by the following interval: $\{pv \in \mathbb{Z} \mid 1 \leq pv \leq 5\}$.

Default value: 0.

- Subsampling value (sv)[-N]: The label set of each pruned example (in accordance to the examples pruned by the use of the previous parameter, i.e., the pruning value) becomes a candidate for label-set subsampling. The PSt algorithm subsamples the label sets of pruned examples to create examples which do meet the pruning criterion. So, the subsample value defines the (maximum) number of frequent labelsets to subsample from the infrequent labelsets. This parameter is bounded by the following interval: $\{sv \in \mathbb{Z} \mid 0 \leq sv \leq 5\}$.

Default value: 0.

There are no dependencies/constraints between the parameters of PSt. Additionally, there is a proper study in the work of [47] about the range of these two parameters (pv and sv). The parameters are the same of PS. The main thing that is changed in PSt when compared to PS occurs at the test time.

2.17 Random k-Label Pruned Sets

The RANdom k-labEL Pruned Sets (RAkEL) algorithm [47, 62] randomly draws M subsets of labels, each with k labels, from the set of labels, and trains PS upon each one. Finally, it combines label votes from the PS classifiers to get a label-vector prediction. **Parameters:**

- Base classifier (bc)[-W] : It can be any classifier from WEKA.
- Pruning value (pv)[-P]: It prunes an infrequent labelset when it occurs less than pv times in the data. $pv = 0$ means that LC classifier is learned. Thus, this value is not allowed for RAkEL, which makes this parameter being bounded by the following interval:

$$\{pv \in \mathbb{Z} \mid 1 \leq pv \leq 5\}.$$

Default value: 0.

- Subsampling value (sv)[-N]: The label set of each pruned example (in accordance to the examples pruned by the use of the previous parameter, i.e., the pruning value) becomes a candidate for label-set subsampling. This version of RAkEL in MEKA subsamples the label sets of pruned examples to create examples which do meet the pruning criterion. So, the subsample value defines the (maximum) number of frequent labelsets to subsample from the infrequent labelsets. This parameter is bounded by the following interval: $\{sv \in \mathbb{Z} \mid 0 \leq sv \leq 5\}$.

Default value: 0.

- Number of labels for each subset (les)[-k]: It defines the number of labels in each label subset. This parameter should be bounded by the interval [42]: $\{les \in \mathbb{Z} \mid 1 \leq les \leq \frac{L}{2}\}$, where L is the number of labels.

Default value: 3.

- Number of subsets to run in an ensemble (sre)[-M]: This parameter controls the number of models to build in a ensemble and take values in accordance to the following interval [42]: $\{sre \in \mathbb{Z} \mid 2 \leq sre \leq \min(2 \cdot L, 100)\}$, where L is the number of labels.

Default value: 10.

There are no dependencies/constraints between the parameters of RAkEL. Additionally, we followed the work of [47] about the range of the subsampling and pruning values. The other two parameters (number of labels in each subset and number of models to build in a ensemble) were defined in accordance to the work of [42].

2.18 Random k-Label Disjoint Pruned Sets

The RANdom k-labEL Disjoint Pruned Sets (RAkELd) algorithm [47, 62] takes a random partition of labels, but unlike RAkEL the labelsets are disjoint/non-overlapping subsets. **Parameters:**

- Base classifier (bc)[-W] : It can be any classifier from WEKA.
- Pruning value (pv)[-P]: It prunes an infrequent labelset when it occurs less than p times in the data. $pv = 0$ means that LC classifier is learned. Thus, this value is not allowed for RAkEL, which makes this parameter being bounded by the following interval:

$$\{pv \in \mathbb{Z} \mid 1 \leq pv \leq 5\}.$$

Default value: 0.

- Subsampling value (sv)[-N]: The label set of each pruned example (in accordance to the examples pruned by the use

of the previous parameter, i.e., the pruning value) becomes a candidate for label-set subsampling. The version of RAKED in MEKA subsamples the label sets of pruned examples to create examples which do meet the pruning criterion. So, the subsample value defines the (maximum) number of frequent labelsets to subsample from the infrequent labelsets. This parameter is bounded by the following interval: $\{sv \in \mathbb{Z} \mid 0 \leq sv \leq 5\}$.

Default value: 0.

- **Number of subsets to run in an ensemble (*sre*)[-M]):** This parameter controls the number of models to build in an ensemble and take values in accordance to the following interval [42]: $\{sre \in \mathbb{Z} \mid 2 \leq sre \leq \min(2 \cdot L, 100)\}$, where *L* is the number of labels.

Default value: 10.

There are no dependencies/constraints between the parameters of RAKELd. Additionally, we followed the work of [47] again to set the range of the subsampling and pruning values. The other two parameters (number of labels in each subset and number of models to build in an ensemble) were defined in accordance to the work of [42].

2.19 Multi-Label Back Propagation Neural Network

The multi-label back propagation neural network (ML-BPNN) algorithm [51, 66] is a standard Back-Propagation Neural Network [55] with multiple outputs that correspond to multiple labels. That is, each node in the output layer corresponds to a different class label.

Parameters:

- **Number of epochs (*ne*)[-E]:** It is the number of iterations to train the neural network. It is restricted by the interval: $\{ne \in \mathbb{Z} \mid 10 \leq ne \leq 1000\}$.

Default value: 100.

- **Number of hidden units (*nhu*)[-H]:** It defines the number of hidden units in the neural network. It is import to mention that the version of ML-BPNN in MEKA is limited to one hidden layer with *nhu* hidden units. This parameter takes values in proportion to the number of attributes (received as input). Thus, the number of hidden units of the network can vary from 20% to 100% of the number of attributes: $\{nhu \in \mathbb{Z} \mid 0.2 \cdot \text{number\_of\_attributes} \leq nhu \leq \text{number\_of\_attributes}\}$. The proportion will always be rounded to the nearest integer.

Default value: 10.

- **Learning rate (*lr*)[-r]:** The amount by which the weights are updated during training. It is restricted by the interval: $\{lr \in \mathbb{R} \mid 0.001 \leq lr \leq 0.1\}$.

Default value: 0.1.

- **Momentum (*m*)[-m]:** It is applied to the weights during updating. It is restricted by the interval: $\{m \in \mathbb{R} \mid 0.1 \leq m \leq 0.8\}$.

Default value: 0.1.

There are no dependencies/constraints between the parameters of ML-BPNN. Additionally, the range of values for the parameters number of epochs, momentum and learning rate were set following

the work of [51]. The only parameter which was defined based on a different work [66] was the number of hidden units, *nhu*.

3 SEARCH SPACE – MLC META-ALGORITHMS

In this section, we describe the *search space* of multi-label meta-algorithms in MEKA. It is important to say that some of the multi-label classifiers (presented in the last section) do not perform very well when used as the multi-label base classifier in a meta classifier. This is due to the poor scalability of such combination (meta multi-label and base multi-label). Examples of algorithms that would not scale up well are: MCC, PCC, PMCC, CDN and CDT (these last two algorithms involve Gibbs sampling, which may be too expensive in an ensemble), RAKEL and RAKELd (these two algorithms are ensembles by themselves, and using an ensemble as base classifier would lead to a very slow ensemble of ensembles). This must be considered in the grammar or directly in the execution of the algorithm (i.e, setting a time budget for such algorithms when they are used at the multi-label base level).

3.1 Subset Mapper

The subset mapper (SM) algorithm [57] maps the output of a multi-label classifier to a known label combination using the Hamming distance, i.e., it checks what label combination (label subsets) from the training set has the closest distance to the predicted label combination on the test instance using probability distribution of the label subset for this instance. In order to do that, SM transforms the probability distribution array of the label subset in a binary array. For each label subset in the training set (also represented by a binary array), it calculates the Hamming distance to the binary probability distribution array, outputting the closest label subset to the predicted distribution array. SM will map this label subset to this particular test instance. **Parameters:**

- **Multi-label classifier(*mlc*)[-W]:** The multi-label algorithm that creates a model at the multi-label classification level.

Dependencies/Constraints:

- (1) The multi-label classification algorithm can be any one described in the Section 2.

3.2 Bagging of Multi-Label Classifiers

The bagging of multi-label classifiers (BaggingML) [47] is an algorithm that combines several multi-label classifiers using Bootstrap AGGREGatING (Bagging) [9]. It randomly sets weights higher than zero to certain instances, on only those instances are chosen for the bag. The parameter “bag percent size” is then not used as the number of instances in the bag is just based on the weight values. Thus, the members of the ensemble could have 100% of the instances if all of them have a weight assigned. **Parameters:**

- **Multi-label classifier(*mlc*)[-W]:** The multi-label algorithm that creates a model at the base multi-label classification level.
- **Number of iterations (*i*)[-I]:** The number of iterations to perform, i.e., the number of members in the ensemble. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 10 \leq i \leq 50\}$. The range of the number of iterations for this ensemble algorithm was defined by Read’s thesis [47].

Default value: 10.

Dependencies/Constraints:

- (1) The multi-label classification algorithm can be any one described in the Section 2, except for BCC, which is not suitable for this meta-learner.

3.3 Bagging of Multi-Label Classifiers with Duplicates

BaggingML with duplicates (BaggingMLDup) [47] is an algorithm that also combines several multi-label classifiers using Bootstrap AGGREGatING. However, it uses the parameter “bag size percent” to define a specific number of instances for each member (classifier) of the ensemble. After that, it randomly samples instances, being able to sample the same instance (*duplicates*) for the bag. This algorithm does not use any weight to select the instances for the members of the ensemble. **Parameters:**

- Multi-label classifier (*mle*)[-W]: The multi-label algorithm that creates a model at the multi-label classification level.
- Bag size percent (*bsp*)[-P]: The size of the bag in percentage of the training set size (number of training instances) and it is defined by the interval: $\{bsp \in \mathbb{Z} \mid 10 \leq bsp \leq 100\}$.
Default value: 67.
- Number of iterations (*i*)[-I]: The number of iterations to perform, i.e., the number of members in the ensemble. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 10 \leq i \leq 50\}$.
Default value: 10.

Dependencies/Constraints:

- (1) The multi-label classification algorithm can be any one described in the Section 2, except for BCC, which is not suitable for this meta-learner.

The range of the number of iterations for this ensemble algorithm was defined by Read’s thesis [47]. The parameter “bag size percent” is defined in the MEKA documentation.

3.4 Ensemble of Multi-Label Classifiers

The ensemble of multi-label classifiers (EnsembleML) [47] is an algorithm that combines several multi-label classifiers in a simple-subset ensemble. This algorithm is very similar to BaggingMLDup. The only difference is that BaggingMLDup allows sampling with replacement for each model, whereas EnsembleML uses sampling without replacement.

Parameters:

- Multi-label classifier (*mle*)[-W]: The multi-label algorithm that creates a model at the base multi-label classification level.
- Bag size percent (*bsp*)[-P]: The size of the bag in percentage of the training size (number of training instances) and it is defined by the interval: $\{bsp \in \mathbb{Z} \mid 52 \leq bsp \leq 72\}$.
Default value: 67.
- Number of iterations (*i*)[-I]: The number of iterations to perform, i.e., the number of members in the ensemble. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 10 \leq i \leq 50\}$.
Default value: 10.

Dependencies/Constraints:

- (1) The multi-label classification algorithm can be any one described in the Section 2, except for BCC, which is not suitable for this meta-learner.

The range of the number of iterations for this ensemble algorithm was defined by Read’s thesis [47]. Additionally, in his thesis, the author mentioned that they found that values around 62% are the best ones for the parameter “bag size percent” in an ensemble without replacement, which is the case. Thus, we are trying to set the range for this parameter introducing lower and upper bounds close to this value (10% smaller and 10% greater).

3.5 Random Subspace Multi-Label

The random subspace multi-label (RSML) algorithm [10] combines several multi-label classifiers in an ensemble where the attribute space and the instance space used for building each model are random subsets from the original space. In other words, RSML subsamples the attribute space and instance space randomly for each ensemble member. Basically, it is a generalized version of Random Forests. Additionally, it is computationally cheaper than EnsembleML for the same number of models in the ensemble and the same value of bag size percent.

Parameters:

- Multi-label classifier (*mle*)[-W]: The multi-label algorithm that creates a model at the base multi-label classification level.
- Bag size percent (*bsp*)[-P]: The size of the bag in percentage of the training set size (number of training instances), and it is defined by the interval: $\{bsp \in \mathbb{Z} \mid 10 \leq bsp \leq 100\}$.
Default value: 67.
- Number of iterations (*i*)[-I]: The number of iterations to perform, i.e., the number of members in the ensemble. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 10 \leq i \leq 50\}$.
Default value: 10.
- Attribute percent (*ap*)[-A]: The size of the attribute space, as a percentage of total attribute space size (number of attributes). This parameter is bounded by the following interval: $\{ap \in \mathbb{Z} \mid 10 \leq ap \leq 100\}$.
Default value: 50.

Dependencies/Constraints:

- (1) The multi-label classification algorithm can be any one described in the Section 2.

The range of the number of iterations for this ensemble algorithm was defined based on Read’s thesis [47]. The range of values for this parameter also considers scalability issues as we need to run a multi-label algorithm many times in an ensemble. The parameter “bag size percent” is defined in the MEKA documentation. The attribute percentage was set in accordance to the single-label version for the same algorithm. This was done because there is not any work that studies this algorithm for multi-label classification.

3.6 Expectation Maximization

In the expectation maximization (EM) algorithm [16], a specified multi-label classifier is built on the training data. This model is then used to classify the training data. The confidence with which instances are classified is used to reweight them. This data is then used to retrain the classifier. This cycle continues (‘EM’-style) for I

iterations. The final model is used to classify the test data. Because of the weighting, it is advised to use a classifier which gives good confidence (probabilistic) outputs. **Parameters:**

- Multi-label classifier (*mle*)[-W]: The multi-label algorithm that creates a model at the multi-label classification level.
- Number of iterations (*i*)[-I]: The number of iterations to perform. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 10 \leq i \leq 50\}$.

Default value: 10.

Dependencies/Constraints:

- (1) The classifier at the base multi-label classification level should be capable to produce probabilistic predictions. However, in our preliminary tests, most multi-label classification algorithms described in Section 2 were suitable for this meta-learner, except for PMCC. Thus, we will use all the suitable algorithms in accordance to these experiments at the base multi-label classification level.

The range of the number of iterations for this ensemble algorithm was defined based on the Read's thesis [47]. The range of values for this parameter also considers scalability issues as we need to run a multi-label algorithm many times in an ensemble. This was done because there is not an appropriate work that studies this algorithm for multi-label classification.

3.7 Classification Maximization

The classification maximization (CM) algorithm [16, 54] trains a classifier with labeled and unlabeled data (semi-supervised) learning using the Classification Expectation algorithm, which is a hard version of EM algorithm, as it does not update the instance weights using (a product factor of) the probability distribution produced by the classifier. Instead, it sets to zero (0.0) or one (1.0) the weight of any instance in the dataset. Unlike EM, it can use any classifier, not necessarily one that gives good probabilistic outputs. **Parameters:**

- Multi-label classifier (*mle*)[-W]: The multi-label algorithm that creates a model at the base multi-label classification level.
- Number of iterations (*i*)[-I]: The number of iterations to perform. This parameter is restricted by the interval: $\{i \in \mathbb{Z} \mid 10 \leq i \leq 50\}$.

Default value: 10.

Dependencies/Constraints:

- (1) The multi-label classification algorithm can be any one described in the Section 2, except for PMCC, which is not suitable for this meta-learner.

The range of the number of iterations for this ensemble algorithm was defined based on Read's thesis [47]. The range of values for this parameter also considers scalability issues as we need to run a multi-label algorithm many times in an ensemble. This was done because there is not any work that studies this algorithm for multi-label classification.

4 STUDYING THE SEARCH SPACE OF SINGLE-LABEL CLASSIFICATION ALGORITHMS

In this section, we study 28 traditional (single label) classification algorithms, preprocessing algorithms and meta-algorithms from the WEKA software [27]. This is done in order to understand the whole *search space* of multi-label algorithms. Most hyper-parameters in this section were set in accordance to the *search space* definition from Auto-WEKA [38, 39, 60]. The algorithms and their respective (hyper-)parameters were defined after studying the code, logs and configuration files of Auto-WEKA, which is considered a stable and robust approach for automatically selecting and configuring machine learning algorithms.

Table 2 shows the 28 SLC algorithms used, i.e., the possible algorithms at the MLC base level and their hyper-parameters. Similar to Table 1, Table 2 defines the SLC algorithms in terms of their names, names' acronyms and types (i.e., trees, rules, lazy, functions, Bayes, preprocessing and, meta-SLC). As before, we show if the SLC algorithm in the row is used or not ('Y' for yes and 'N' for no) by the respective *search space* in the column and how many hyper-parameters (#HP) it has (when it is used).

5 SEARCH SPACE – TRADITIONAL SLC ALGORITHMS

This section covers the main traditional SLC algorithms in the WEKA software [27]. As already mentioned and explained, this is done as most MLC algorithms, which are described in the last two sections of this supplementary material, are from the problem transformation type, i.e., they need to employ a single-label classifier at their base level to perform the multi-label classification.

5.1 C4.5

The algorithm for generating a C4.5 decision tree [45]. This algorithm can decide whether it will use the default C4.5's error-based pruning method [8, 56, 64] or not. If the algorithm decides to use pruning, the C4.5's pruning method is applied to the tree, and an estimation of the error rate of every subtree is done. After that, the pruning method will replace the subtree with a leaf node if the estimated error of the leaf is lower than a threshold [56]. **Parameters:**

- Confidence factor (*cf*)[-C]: It is used for C4.5's error-based pruning method (smaller values incur more pruning) and is defined by the interval: $\{cf \in \mathbb{R} \mid 0.0 \leq cf \leq 1.0\}$.
Default value: 0.25.
- Minimum number of objects (*mno*)[-M]: The minimum number of instances per leaf. It can take values in the interval: $\{mno \in \mathbb{Z} \mid 1 \leq mno \leq 64\}$.
Default value: 2.
- Collapse tree (*ct*)[-O]: It is used to decide if internal nodes will be collapsed to avoid overfitting. This parameter is used with C4.5's error-based pruning method to enhance the final decision tree. It collapses a subtree to a node only if training error of the subtree does not increase when compared to the entire tree. It is applied to every subtree in the tree, where subtrees are collapsed (pruned) if pruning does not increase its classification error. For example, if there is a subtree with

Table 2: Overview of the employed single-label classification (SLC) algorithms from the WEKA software.*

id	Algorithm Name	Acronym	Category	Small		Medium		Large	
				Used?	#HP	Used?	#HP	Used?	#HP
1	JRip	JRip	Rules	Y	4	Y	4	Y	4
2	K-Nearest Neighbors	KNN	Lazy	Y	3	Y	3	Y	3
3	Logistic Regression	LR	Functions	Y	1	Y	1	Y	1
4	Naïve Bayes	NB	Bayes	Y	2	Y	2	Y	2
5	Random Forest	RF	Trees	Y	3	Y	3	Y	3
6	Bayesian Network Classifier	BNC	Bayes	N	-	Y	1	Y	1
7	C4.5	C4.5	Trees	N	-	Y	8	Y	8
8	Decision Table	DT	Rules	N	-	Y	4	Y	4
9	K Star	K*	Lazy	N	-	Y	3	Y	3
10	Logistic Model Trees	LMT	Trees	N	-	Y	7	Y	7
11	Multi-Layer Perceptron	MLP	Functions	N	-	Y	6	Y	6
12	PART	PART	Rules	N	-	Y	4	Y	4
13	REPTree	REPTree	Trees	N	-	Y	3	Y	3
14	Stochastic Gradient Descent	SGD	Functions	N	-	Y	5	Y	5
15	Sequential Minimal Optimization	SMO	Functions	N	-	Y	6	Y	6
16	Decision Stump	DS	Trees	N	-	N	-	Y	0
17	Naïve Bayes Multinomial	NBM	Bayes	N	-	N	-	Y	0
18	One Rule	OneR	Rules	N	-	N	-	Y	1
19	Random Tree	RT	Trees	N	-	N	-	Y	4
20	Simple Logistic	SL	Functions	N	-	N	-	Y	3
21	Voted Perceptron	VP	Functions	N	-	N	-	Y	3
22	Zero Rules	ZeroR	Rules	N	-	N	-	Y	0
23	Attribute Selection Classifier	ASC	Preprocessing	N	-	N	-	Y	1
24	Ada boost M1	AdaM1	Meta-SLC	N	-	N	-	Y	3
25	Bagging of Single-Label Classifiers	Bagging	Meta-SLC	N	-	N	-	Y	3
26	Locally Weighted Learning	LWL	Meta-SLC	N	-	N	-	Y	2
27	Random Committee	RC	Meta-SLC	N	-	N	-	Y	1
28	Random Subspace	RSS	Meta-SLC	N	-	N	-	Y	3

*All algorithm names are clickable links to their description in this supplementary material.

two leaf nodes having the same classification on the training data, this subtree will be replaced by a single leaf. It can take Boolean values (true or false).

Default value: true.

- Unpruned (*u*)[-U]: It decides whether pruning is performed or not. It can take Boolean values (true or false).

Default value: false.

- Binary splits (*bs*)[-B]: It decides whether C4.5 will use binary splits on nominal attributes when building the trees. It can take Boolean values (true or false).

Default value: false.

- Use MDL correction (*umc*)[-J]: It decides whether the MDL correction is used when finding splits on numeric attributes. It can take Boolean values (true or false).

Default value: true.

- Use Laplace (*ul*)[-A]: It decides if the counts of instances at leaves are smoothed based on the Laplace correction. It can take Boolean values (true or false).

Default value: false.

- Subtree raising (*sr*)[-S]: It is used for C4.5's error-based pruning and decides whether the algorithm will consider the subtree raising operation when pruning. It can take Boolean values (true or false).

Default value: true.

Dependencies/Constraints:

- (1) If the parameter unpruned is set to "true", the parameters "confidence factor", "collapse tree" and "subset raising" are not used (omitted).

5.2 Logistic Model Trees

The algorithm for building logistic model trees (LMT) [40, 59], which are classification trees with logistic regression functions at the leaves. This is done by using the LogitBoost algorithm. In this case, boosting is used (aiming) to build very effective decision trees. The idea of LMT is to use LogitBoost to induce trees with linear-logistic regression models at the leaves. LogitBoost performs additive logistic regression. Thus, at each iteration of the boosting algorithm, it creates a simple regression model by going through all

the attributes, finding the simple regression function with the smallest error, and adding it into the additive model [64]. The algorithm can deal with binary and multi-class target variables, numeric and nominal attributes and missing values. **Parameters:**

- Minimum number of objects (*mno*)[-M]: The minimum number of instances per leaf. It can take values in the interval: $\{mno \in \mathbb{Z} \mid 1 \leq mno \leq 64\}$.
Default value: 15.
- Convert Nominal (*cn*)[-B]: It decides if the algorithm will convert all nominal attributes to binary ones before building the tree. This means that all splits in the final tree will be binary. It can take Boolean values (true or false).
Default value: false.
- Split on residuals (*sor*)[-R]: It decides whether the algorithm will set the splitting criterion based on the residuals of LogitBoost. There are two possible splitting criteria for LMT: the default is to use the C4.5 splitting criterion that uses information gain on the class variable. The other splitting criterion tries to improve the purity in the residuals produced when fitting the logistic regression functions. It can take Boolean values (true or false).
Default value: false.
- Fast Regression (*fr*)[-C]: It decides whether the algorithm will use a heuristic that avoids the use of cross-validation to optimize the number of LogitBoost iterations at every node. In the case of using this heuristic, LMT will fit the logistic regression functions at a leaf node using the LogitBoost algorithm, applying a 5-fold cross-validation procedure to determine how many iterations to run just once. Then, it employs the same number of iterations throughout the tree, instead of cross-validating at every node. This heuristic reduces the running time considerably, with little effect on accuracy [64]. It can take Boolean values (true or false).
Default value: true.
- Error on probabilities (*eop*)[-P]: It decides if the algorithm will minimize the error on classification probabilities instead of the misclassification error when cross-validating the number of LogitBoost iterations. When this parameter is set to 'true', the number of LogitBoost iterations that minimizes the error on classification probabilities instead of the misclassification error is chosen. It can take Boolean values (true or false).
Default value: false.
- Weight trim beta (*wtb*)[-W]: It sets the beta value used for weight trimming in LogitBoost. Only instances carrying $(1 - \text{beta})\%$ of the weight from the previous iteration are used in the next iteration. The value zero (0) means no weight trimming, which is the default value. The values are restricted to the interval: $\{wtb \in \mathbb{R} \mid 0.0 \leq wtb \leq 1.0\}$.
Default value: 0.0.
- Use AIC (*uaic*)[-A]: It decides if the algorithm will use the AIC (Akaike's Information Criterion) measure to determine when to stop LogitBoost's iterative process. More precisely, if *uaic* takes the value 'true', the best number of iterations will be defined by an information criterion measure (currently, AIC). If false, the stopping criterion will be determined by the best number of iterations in a 5-fold cross-validation

procedure. It can take Boolean values (true or false).

Default value: false.

There are no dependencies/constraints between the parameters of LMT.

5.3 Decision Stump

The algorithm for building and applying a decision stump (DS) model [64], which is considered a weak learner. Because of that, it is usually used in conjunction with a boosting algorithm.

The DS's classification is based on the entropy measure and a missing value is treated as a separate value. The DS algorithm constructs a simple decision tree that has only one level, i.e., a decision tree that has only one internal (root) node, that is directly linked to the leaves. It also creates an extra branch for missing values.

In the case of nominal attributes at the root node, there are two possibilities. The first possibility is to build a stump which contains a leaf for each possible feature value. The second possibility is to consider a stump with two leaves, one of them is mapped to some category, and the another to all other categories. The DS from WEKA employs the latter approach. This algorithm has no explicit parameters.

5.4 Random Forest

The Random Forest (RF) algorithm for constructing a forest of random trees [10]. **Parameters:**

- Number of trees (*nt*)[-I]: The number of trees to be generated by the algorithm. It is an integer value bounded by the interval: $\{nt \in \mathbb{Z} \mid 2 \leq nt \leq 256\}$.
Default value: 100.
- Number of features (*nf*)[-K]: It sets the number of randomly sampled attributes used as candidate attributes at each tree node. It is an integer value bounded by the interval: $\{nf \in \mathbb{Z} \mid 2 \leq nf \leq 32\}$. However, it may also take the value zero (0), which means *nf* will be just used as a flag to indicate that the real value produced by the equation $\log_2(\text{number\_of\_attributes} + 1)$ rounded to the nearest integer is automatically used for this parameter.
Default value: 0.
- Maximum depth (*md*)[-depth]: The maximum depth of the tree. It is bounded by the interval: $\{md \in \mathbb{Z} \mid 2 \leq md \leq 20\}$. However, it may also take the value zero (0) as a flag and, in this case, the depth of the tree can be unlimited.
Default value: 0.

There are no dependencies/constraints between the parameters of RF.

5.5 Random Tree

The algorithm for constructing a tree that considers K randomly sampled attributes as candidate attributes at each node, i.e., a random tree (RT) [64]. It is important to mention that this version of RT performs no pruning. **Parameters:**

- Minimum weight (*mw*)[-M]: The minimum total weight of the instances in a leaf. It is restricted by the interval: $\{mw \in \mathbb{Z} \mid 1 \leq mw \leq 64\}$.
Default value: 1.

- Number of features (nf)[-K]: It sets the number of randomly sampled attributes used as candidate attributes at each tree node. It is an integer value bounded by the interval: $\{nf \in \mathbb{Z} \mid 2 \leq nf \leq 32\}$. However, it may also take the value zero (0), which means nf will be just used as a flag to indicate that the real value produced by the equation $\log_2(\text{number\_of\_attributes} + 1)$ rounded to the nearest integer is automatically used for this parameter.
Default value: 0.

- Maximum depth (md)[-depth]: The maximum depth of the tree. It is bounded by the interval: $\{md \in \mathbb{Z} \mid 2 \leq md \leq 20\}$. However, it may also take the value zero (0) as a flag and, in this case, the depth of the tree can be unlimited.
Default value: 0.

- Number of folds for back-fitting and for growing the tree ($nfbgt$)[-N]: It determines the amount of data used for back-fitting and for growing the tree. One fold is used for back-fitting, i.e., for making a preliminary estimation of class probabilities based on a hold-out set. The others ($nf - 1$) folds are used for growing the tree. It is bounded by the interval: $\{nfbgt \in \mathbb{Z} \mid 2 \leq nfbgt \leq 5\}$. It can also use the value zero (0), which means no back-fitting will be performed in this case. It can not take the value one (1) because we would have zero folds for growing the tree. In the case of taking the value one, the algorithm returns an error and does not run. It is important to mention that Auto-WEKA allows this error, ignoring RT algorithm with this configuration (when it occurs), and continuing the search from this point.
Default value: 0.

There is no constraints/dependencies between the parameters of RT.

5.6 REPTree

The algorithm for the fast decision tree learner, which is well-known as REPTree [64]. It builds a decision tree using information gain and prunes it using reduced-error pruning (with back-fitting). It only sorts values for numeric attributes once, at the start of the algorithm. Missing values are dealt with by splitting the corresponding instances into pieces (i.e., as in C4.5).

Parameters:

- Minimum weight (mw)[-M]: The minimum total weight of the instances in a leaf. It is restricted by the interval: $\{mw \in \mathbb{Z} \mid 1 \leq mw \leq 64\}$.
Default value: 2.
- Maximum depth (md)[-L]: The maximum tree depth. It can take integer values considering the interval: $\{md \in \mathbb{Z} \mid 2 \leq md \leq 20\}$. However, it may also take the value -1 as a flag and, in this case, the depth of the tree the depth will not be restricted.
Default value: -1.
- Use pruning (up)[-P]: It decides whether REPTree will use reduced-error pruning or not. In the case of using this pruning method, a simple hold-out set ($\frac{1}{3}$ of the training data) is used to estimate the error of a node, instead of using cross-validation. It can take Boolean values (true or false).
Default value: false.

There are no dependencies/constraints between the parameters of REPTree.

5.7 Decision Table

The algorithm for building and using a simple decision table (DT) classifier [37]. **Parameters:**

- Evaluation Measure (em)[-E]: The measure used to evaluate the performance of attribute combinations used in the decision table. It can take one of the four categorical values: 1. accuracy (acc); 2. root mean squared error (rmse) of the the class probabilities; 3. mean absolute error (mae) of the class probabilities; 4. area under the ROC curve (auc). The two measures *rmse* and *mae* are adapted to be used in the classification context.
Default value: acc.
- Use IBk ($uibk$)[-I]: It sets whether a k-nearest neighbor ($k=1$) classifier should be used instead of the majority class in order to classify non-matching instances. It can take Boolean values (true or false).
Default value: false.
- Search method (sm)[-S]: It sets the search algorithm which will be used to find good attribute combinations for the decision table. It can take the values Greedy Stepwise or Best First.
Default value: Best First.
- Cross-Validation (crv)[-X]: It sets the number of folds for the internal cross validation procedure to evaluate the attribute sets. It may take the values one (1), two (2), three (3) or four (4). If the value one (1) is set for this hyper-parameter, a leave one out procedure is applied.
Default value: 1.

There are no dependencies/constraints between the parameters of DT.

5.8 JRip

The algorithm that implements a propositional rule learner algorithm, namely Repeated Incremental Pruning to Produce Error Reduction (RIPPER) [13]. **Parameters:**

- Minimum total weight (mtw)[-N]: This parameter determines the minimum total weight of the instances in a rule. It can take values considering the interval: $\{mtw \in \mathbb{R} \mid 1.0 \leq mtw \leq 5.0\}$.
Default value: 2.0.
- Check error rate (cer)[-E]: It decides whether JRip will consider the “error rate greater or equal than 0.5” as a stopping criterion. It can take Boolean values (true or false).
Default value: true.
- Use pruning (up)[-P]: It decides whether JRip will use reduced error pruning or not. In the case of using this pruning method, a 3-fold cross-validation procedure is applied to prune the rules. Otherwise, no pruning method is used. It can take Boolean values (true or false).
Default value: false.
- Optimizations (o)[-O]: The number of optimization runs. It can take integer values considering the interval: $\{o \in \mathbb{Z} \mid 1 \leq$

$o \leq 5$).

Default value: 2.

There are no dependencies/constraints between the parameters of JRip.

5.9 One Rule

The algorithm for building and using a one rule (OneR) classifier [32]. In other words, it uses the minimum-error attribute for prediction, discretizing numeric attributes. **Parameters:**

- Minimum bucket size (*mbz*)[-B]: It is used for discretizing numeric attributes. It is limited by the interval: $\{mbz \in \mathbb{Z} \mid 1 \leq mbz \leq 32\}$.

Default value: 6.

OneR has only one parameter and, consequently, there is no dependencies/constraints for it.

5.10 PART

The algorithm for generating a PART decision list [22]. PART uses the separate-and-conquer paradigm: It builds a partial C4.5 decision tree in each iteration and makes the “best” leaf into a rule.

Parameters:

- Minimum number of objects (*mno*)[-M]: The minimum number of instances per leaf. It can take values in the interval: $\{mno \in \mathbb{Z} \mid 1 \leq mno \leq 64\}$.

Default value: 2.

- Binary splits (*bs*)[-B]: It decides whether C4.5 will use binary splits on nominal attributes when building the trees. It can take Boolean values (true or false).

Default value: false.

- Reduced-error pruning (*rep*)[-R]: It is used to decide whether reduced-error pruning is used instead of C4.5’s default pruning (error-based pruning). If C4.5’s error-based pruning is chosen, a (default) confidence factor of 0.25 is used to prune the tree. If not (i.e., the reduced-error pruning is chosen), the algorithm will consider each node for pruning and the removal of a subtree at a node is done if the resulting tree performs no worse than the original one on the validation set. The size of the validation set is determined by the next parameter (*nr*). It can take Boolean values (true or false).

Default value: true.

- Number of folds (*nr*)[-N]: It determines the amount of data used for reduced-error pruning. One fold is used for pruning and the rest for growing the tree. It can take the values two (2), three (3), four (4) or five (5).

Default value: not used.

Dependencies/Constraints:

- (1) If the reduced-error pruning method is not set to “true”, the parameter “number of folds” is not used.

5.11 Zero Rule

The algorithm for building and using a zero rule (ZeroR) classifier [64]. The ZeroR classifier simply predicts the majority category (class), ignoring the predictor attributes. This algorithm has no explicit parameters.

5.12 K-Nearest Neighbors

The algorithm for k-nearest neighbors (KNN) classifier [3]. KNN can select an appropriate value of K based on internal leave-one-out evaluation and can also compute distances based on instance weighting. **Parameters:**

- Number of neighbors (*k*)[-K]: The number of neighbors to use. The value of *k* is bounded by the interval: $\{k \in \mathbb{Z} \mid 1 \leq k \leq 64\}$.

Default value: 1.

- Leave-one-out (*loo*)[-X]: It decides whether leave-one-out evaluation on the training data will be used or not to select the best *k* value between 1 and the value specified as the KNN parameter. If set as false, the selected *k* value is used. It can take only Boolean values (true or false).

Default value: false.

- Distance weighting (*dw*): It sets the used distance weighting method. It may take the unique following values:
  - -I: Weight neighbors by the inverse of their distance.
  - -F: Weight neighbors by one minus their distance.
  - None: No distance weighting method is applied.

Default value: None.

There are no dependencies/constraints between the parameters of KNN.

5.13 K*

The algorithm K* is an instance-based classification algorithm [12]. Thus, in order to classify a test instance, K* considers the class of those training instances similar to it, as determined by some similarity function. It differs from other instance-based learners by using an entropy-based distance function. **Parameters:**

- Global blending (*gb*)[-B]: The parameter is a percentage for global blending. This parameter controls the “sphere of influence” by specifying how many of the neighbors of the instance *i* should be considered important (although there is no hard cut off at the edge of the sphere – it is more related to a gradual decreasing of importance). The values are restricted to the interval $\{gb \in \mathbb{Z} \mid 1 \leq gb \leq 100\}$. Thus, selecting zero (0) for this parameter gives a nearest neighbor algorithm (this is why Auto-WEKA does not allow to choose it), and choosing 100 gives equally weighted instances. Intermediate values are interpolated linearly.

Default value: 20.

- Entropic auto-blending (*eab*)[-E]: It decides whether entropy-based blending will be used or not. It can take Boolean values (true or false).

Default value: false.

- Missing Mode (*mm*)[-M]: It determines how missing attribute values are treated. It can take one of the four categorical values: 1. average column entropy curves (a); 2. ignore the instances with missing values (d); 3. treat missing values as maximally different (m); 4. normalize over the attributes (n).

Default value: a.

There are no dependencies/constraints between the parameters of K*.

5.14 Voted Perceptron

The voted perceptron (VP) algorithm created by [24]. It globally replaces all missing values by their default values. More precisely, VP replaces all missing values for nominal and numeric attributes by the modes and the means from the training data, respectively. Additionally, it transforms nominal attributes into binary ones.

Parameters:

- Number of iterations (*i*)[-I]: The number of iterations to be performed by VP. This parameter varies in accordance to the interval:
 $\{i \in \mathbb{Z} \mid 1 \leq i \leq 10\}$.
Default value: 1.
- Max K(*mk*)[-M]: The maximum number of alterations to the perceptron, i.e., the maximum number of perceptrons used in the iterative process. It can take values of the interval:
 $\{mk \in \mathbb{Z} \mid 5,000 \leq mk \leq 50,000\}$
Default value: 1,000.
- Exponent (*e*)[-E]: The exponent for the polynomial kernel. It can take values of the interval: $\{e \in \mathbb{R} \mid 0.2 \leq e \leq 5.0\}$
Default value: 1.0.

There are no dependencies/constraints between the parameters of VP.

5.15 Multi-Layer Perceptron

Multi-Layer Perceptron (MLP) uses the traditional back-propagation algorithm [55] to create a neural model to classify the instances. MLP creates just one hidden layer (for now) and all its nodes use sigmoid activation functions (except for when the class is numeric in which case the the output nodes become unthresholded linear units). **Parameters:**

- Learning rate (*lr*)[-L]: The amount by which the weights are updated during training. It is restricted by the interval:
 $\{lr \in \mathbb{R} \mid 0.1 \leq lr \leq 1.0\}$.
Default value: 0.3.
- Momentum (*m*)[-M]: It is applied to the weights during updating. It is restricted by the interval: $\{m \in \mathbb{R} \mid 0.0 \leq m \leq 1.0\}$.
Default value: 0.2.
- Number of hidden nodes (*nhn*)[-H]: It defines the number of hidden nodes in the hidden layer of the neural network. This parameter may take four predefined nominal values (a, i, o and t), which represent the following integer values:
  - $a = \frac{(\text{number\_of\_attributes} + \text{number\_of\_classes})}{2}$, always using the default floor function to convert it to an integer value.
  - i = number of attributes.
  - o = number of classes.
  - $t = (\text{number\_of\_attributes} + \text{number\_of\_classes})$.**Default value:** a.
- Nominal to binary filter (*n2b*)[-B]: It decides whether the algorithm will transform nominal attributes to binary ones or not. This could help improve performance if there are nominal attributes in the data. It can take Boolean values (true or false).
Default value: true.

- Reset (*r*)[-R]: It decides whether the algorithm will use the reset approach. In this case, the algorithm will allow the network to reset with a lower learning rate. If the network diverges from the answer, this will automatically reset the network with a lower learning rate and begin training again. It can take Boolean values (true or false).
Default value: true.

- Decay (*d*)[-D]: It decides whether the algorithm will cause the learning rate to decrease. This will divide the starting value of the learning rate by the sequential number of the current epoch in order to determine what the current learning rate should be. This may help to stop the network from diverging from the target output, as well as improving general performance. It can take Boolean values (true or false).
Default value: false.

There are no dependencies/constraints between the parameters of MLP.

5.16 Stochastic Gradient Descent

The algorithm that implements the stochastic gradient descent (SGD) approach [64] for learning various linear models (binary class SVM, binary class logistic regression, squared loss, Huber loss and epsilon-insensitive loss linear regression). **Parameters:**

- Loss function (*lf*)[-F]: It sets the loss function to be minimized. It can take the following integer values associated to three approaches:
  - (0): hinge loss (SVM).
  - (1): log loss (logistic regression).
  - (2): squared loss (regression).**Default value:** 0.
- Learning rate (*lr*)[-L]: The learning rate. If normalization is turned off, then the default learning rate will need to be reduced. It is restricted by the interval: $\{lr \in \mathbb{R} \mid 0.00001 \leq lr \leq 1.0\}$.
Default value: 0.01.
- Ridge (*r*)[-R]: It sets the Ridge value in the log-likelihood. This parameter can take any value of the given set:
 $\{r \in \mathbb{R} \mid 10^{-12} \leq r \leq 10.0\}$
Default value: 0.0001
- Do not normalize (*nn*)[-N]: It decides whether normalization will be turned off or not. It can take Boolean values (true or false).
Default value: false.
- Do not replace missing values (*nrmv*)[-M]: It decides whether global replacement of missing values will be turned off or not. In the case of being turned off, the missing values will be ignored. Otherwise, SGD will replace all missing values for nominal and numeric attributes by the modes and the means from the training data, respectively. It can take Boolean values (true or false).
Default value: false.

There are no dependencies/constraints between the parameters of SGD.

5.17 Sequential Minimal Optimization

This algorithm implements John Platt's sequential minimal optimization (SMO) algorithm for training a support vector classifier (SVC) [28, 35, 44]. It globally replaces all missing values by their default values. More precisely, SMO (like VP) replaces all missing values for nominal and numeric attributes by the modes and the means from the training data, respectively. Additionally, it transforms nominal attributes into binary ones. **Parameters:**

- **Cost (c)[-C]:** It defines the complexity parameter, which is the penalty parameter of the error term and is defined by the interval: $\{c \in \mathbb{R} \mid 0.5 \leq c \leq 1.5\}$. This is a parameter that controls the trade-off between training error and model complexity. It is important to mention that a low value of c will increase the number of training errors, whereas a high value of c will lead to a behavior similar to that of a hard-margin SVM [33].

Default value: 1.0.

- **Filter type (ft)[-N]:** It determines how/if the data will be transformed. It may take the values zero (0, i.e., normalize the training data – it sets all the numeric attributes in the given dataset into the interval $[0,1]$), one (1, i.e., standardize the training data – it standardizes all numeric attributes in the given dataset to have zero mean and unit variance) or two (2, i.e., no normalization/standardization is applied to the data).

Default value: 0.

- **Build Calibration Models (bcm)[-M]:** It decides whether the model will fit calibration models to SVM's outputs (for proper probability estimates). It can take Boolean values (true or false).

Default value: false.

- **Kernel(k)[-K]:** The kernel to use. It can take one of the following possible kernels (and associated constrained parameters):

- **PolyKernel:** The standard polynomial kernel. **Parameters:**

- (1) **Exponent (exp)[-E]:** It determines the exponent value and is defined by the interval: $\{exp \in \mathbb{R} \mid 0.2 \leq exp \leq 5.0\}$.

- (2) **Use Lower-Order (ulo)[-L]:** It decides whether the algorithm will use lower-order terms or not. It can take Boolean values (true or false).

- **NormalizedPolyKernel:** The normalized polynomial kernel. **Parameters:**

- (1) **Exponent (exp)[-E]:** It determines the exponent value and is defined by the interval: $\{exp \in \mathbb{R} \mid 0.2 \leq exp \leq 5.0\}$.

- (2) **Use Lower-Order (ulo)[-L]:** It decides whether the algorithm will use lower-order terms or not. It can take Boolean values (true or false).

- **Puk:** The Pearson VII function-based universal kernel [63]. **Parameters:**

- (1) **Omega (om)[-O]:** The omega value. It is defined by the interval: $\{om \in \mathbb{R} \mid 0.1 \leq om \leq 1.0\}$.

- (2) **Sigma (sig)[-S]:** The sigma value. It is defined by the interval: $\{sig \in \mathbb{R} \mid 0.1 \leq sig \leq 10.0\}$.

- **RBF:** The RBF kernel. **Parameters:**

- (1) **Gamma (g)[-G]:** The gamma value. It is defined by the interval: $\{g \in \mathbb{R} \mid 0.0001 \leq g \leq 1.0\}$.

Default value: PolyKernel with 'Exponent' equals to 1.0 and 'Use Lower-Order' equals to true.

The constraints/dependencies for SMO are only in the selection of the kernel and its respective parameters.

5.18 Logistic Regression

The algorithm for building and using a multinomial logistic regression (LogR) model with a ridge estimator [11]. **Parameters:**

- **Ridge (r)[-R]:** It sets the Ridge value in the log-likelihood. This parameter can take any value of the given set:

$\{r \in \mathbb{R} \mid 10^{-12} \leq r \leq 10.0\}$

Default value: 0.00000001

LogR has one parameter and, consequently, there is no dependencies/constraints for it.

5.19 Simple Logistic

The algorithm for constructing simple logistic regression (SL) models [40, 59]. LogitBoost with simple regression functions as base learners is used for fitting the logistic models. **Parameters:**

- **Weight trim beta (wtb)[-W]:** It sets the beta value used for weight trimming in LogitBoost. Only instances carrying $(1 - \text{beta})\%$ of the weight from the previous iteration are used in the next iteration. The value zero (0) means no weight trimming, which is the default value. The values are restricted to the interval: $\{wtb \in \mathbb{R} \mid 0.0 \leq wtb \leq 1.0\}$. It also can be omitted and take the default value of zero.

Default value: 0.0.

- **Use Cross-Validation (ucv)[-S]:** It decides if SL will try to find the best number of LogitBoost iterations using an internal 5-fold cross-validation procedure or simply using the number of iterations that minimizes error on the training set. Thus, if not set to 'true', the number of LogitBoost iterations which is used is the one that minimizes the error on the training set (misclassification error). It can take Boolean values (true or false).

Default value: true.

- **Use AIC (uaic)[-A]:** It decides if the method will use the AIC (Akaike's Information Criterion) measure to determine when to stop the LogitBoost iterative process. More precisely, if *uaic* takes the value 'true', the best number of iterations will be defined by an information criterion measure (currently, AIC). If false, the stopping criterion will be determined by the best number of iterations in an internal 5-fold cross-validation procedure or simply in accordance to the error on the training set, as explained in the previous item. It can take Boolean values (true or false).

Default value: false.

There are no dependencies/constraints between the parameters of SL.

5.20 Naïve Bayes

The Naïve Bayes (NB) classifier using estimator classes [34]. This algorithm builds a fixed structure (model) given the attributes of the dataset. **Parameters:**

- Use kernel estimator (*uke*)[-K]: It decides whether NB will use a kernel estimator for numeric attributes rather than a (single) Gaussian distribution. In the case of using the kernel estimator, NB will apply one Gaussian kernel per observed data value (for more details, see Flexible Naïve Bayes' section in [34]). It can take Boolean values (true or false). It is important to mention that a discrete estimator is automatically used for nominal attributes, which is a simple discrete probability estimator based on nominal values' counts. This also means the Laplace correction is applied in order to perform the estimation.

Default value: false.

- Use supervised distribution (*usd*)[-D]: It decides whether NB will use supervised discretization to convert numeric attributes to nominal ones. Discretization is performed by the algorithm proposed in [19]. This method uses a criterion based on the minimum description length (MDL) principle to define the number of intervals produced over the continuous space [17]. It can take Boolean values (true or false).

Default value: false.

Dependencies/Constraints:

- (1) If the parameter "use kernel estimator" is activated, the parameter "use supervised distribution" must not be activated; and vice-versa. This constraint must be enforced in the grammar.

5.21 Bayesian Network Classifier

This algorithm is used to learn a Bayesian Network Classifier (BNC) [7] based on various search algorithms and a local Bayesian scoring metric [14, 29]. **Parameters:**

- Search Method (*sm*)[-Q]: For BNC algorithms, the optimization occurs just on the method used for searching network structures. Thus, the search method can be one of the following: 1. Tree Augmented Naïve Bayes (TAN) [25]; 2. K2 [14]; 3. Hill Climbing (HC) [6, 30]; 4. Look Ahead in Good Directions Hill Climbing (LAGDHC) [1]; 5. Simulated Annealing (SA) [6]; 6. Tabu Search (TS) [6]. All the search methods use the parameter "maximum number of parents" set to two (including the class node), except for TAN and SA which do not use the "maximum number of parents" as a parameter. In addition, the methods use the (default) Bayesian scoring metric to search for appropriate Bayesian networks to data.
- Default value:** there is not default value for this hyperparameter because all these search methods are important algorithms in the literature. Therefore, we use all search methods in our space.

BNC has one parameter and, consequently, there is no dependencies/constraints for it.

5.22 Naïve Bayes Multinomial

The algorithm for building and using Naïve Bayes Multinomial (NBM) [2, 20, 41, 43, 64]. This algorithm was particularly designed for text classification and, for this reason, it changes how the traditional Naïve Bayes calculates the probabilities. This is done to take into account the number of times a word appears in the document.

This algorithm has no explicit parameters and, consequently, there is no dependencies/constraints for it.

6 SEARCH SPACE – META CLASSIFICATION ALGORITHMS FROM WEKA

In this section, the *search space* of 5 traditional (single label) meta classification algorithms from WEKA [27] is studied. This is also done in order to extend and improve the *search space* of multi-label methods. All parameters in this section were also set in accordance to the *search space* definition from Auto-WEKA [38, 39, 60]. The methods and their respective (hyper-)parameters were defined after studying the code, logs and configuration files of Auto-WEKA, which is considered a stable and robust approach for automatically selecting and configuring machine learning algorithms.

6.1 Locally Weighted Learning

The locally weighted learning (LWL) method [5, 21]. It uses an instance-based algorithm to assign instance weights which are then used by a specified Weighted Instances Handler. In other words, LWL assigns weights using an instance-based method and, after this step, another classification algorithm is used to build a classifier from the weighted instances. For example, it can do the classification by using a naïve Bayes classifier or a decision stump (default) from these weighted instances. **Parameters:**

- Classifier (*c*)[-W]: The classifier to be used. It can be one classification algorithm from Section 5, except for the algorithms LMT, OneR, K*, SGD and VP, as these the classifiers produced by these algorithms do not handle weighted instances.
  - Number of neighbors (*k*)[-K]: It sets how many neighbors are used to determine the width of the weighting function. It may take the following values: $\{-1, 10, 30, 60, 90, 120\}$. A negative value means that all neighbors will be considered.
- Default value:** -1.
- Weighting kernel (*wk*)[-U]: It determines the weighting function and may take the five following integer values:
    - (0) Linear.
    - (1) Epnechnikov.
    - (2) Tricube.
    - (3) Inverse.
    - (4) Gaussian.

It can be omitted with 50% of probability, and then LWL will use the default value zero for this parameter, i.e., the linear function. The five (5) not omitted values jointly take the other 50% of probability, which represents at the end that the value 0 has 60% of probability to be chosen.

Default value: 0.

There are no dependencies/constraints between the parameters of LWL.

6.2 Random Subspace

This method constructs an ensemble classifier that consists of multiple models systematically constructed by randomly selecting subsets of components of the feature vector, i.e., the classification models are constructed according to random subspaces (RSS) [31].

More precisely, for each classifier, a certain percentage of the number of attributes is randomly sampled and then used to build the classifier. **Parameters:**

- Classifier (c)[-W]: The classifier to be used. It can be any classification algorithm from Section 5.
- Subspace size (sss)[-P]: It defines the size of each sub-space as a percentage of the number of attributes. It could take values in the range: $\{sss \in \mathbb{R} \mid 0.1 \leq sss \leq 1.0\}$.
Default value: 0.5.

- Number of iterations (ni)[-I]: It defines the number of iterations to be performed, i.e., the number of classifiers in the ensemble. It may take values in the range: $\{ni \in \mathbb{Z} \mid 2 \leq ni \leq 64\}$.
Default value: 10.

There are no dependencies/constraints between the parameters of RSS.

6.3 Bagging of Single-Label Classifiers

The method for bagging a classifier in order to reduce variance [9].

Parameters:

- Classifier (c)[-W]: The classifier to be used for each member of the ensemble. It can be any classification algorithm from Section 5.
- Bag size percent (bsp)[-P]: It defines the size of each bag, as a percentage of the training set size. It may take values in the range: $\{bsp \in \mathbb{Z} \mid 10 \leq bsp \leq 100\}$. It makes sampling with replacement. Thus, even if the bag size percent is 100%, it will sample different sets with the same size of the training set.

Default value: 100.

- Number of iterations (ni)[-I]: It defines the number of iterations to be performed, i.e., the number of classifiers in the ensemble. It may take values in the range: $\{ni \in \mathbb{Z} \mid 2 \leq ni \leq 128\}$.
Default value: 10.

- Calculate out-of-bag (coob)[-O]: It decides whether the out-of-bag error is calculated. It can take Boolean values (true or false).
Default value: false.

Dependencies/Constraints:

- (1) If the parameter “calculate out-of-bag” is activated (set to true), the parameter “bag size percent” must be equal to 100. This is a constraint of WEKA and only an internal modification in the WEKA code could suppress it. This can happen with 50% of probability. I.e., in half of the cases, the parameter “bag size percent” is set to 100. In the other part of the cases, “bag size percent” may take values between 10 and 100, because the parameter “calculate out-of-bag” is not activated (set to false).

6.4 Random Committee

The method for building an ensemble of randomizable base classifiers from WEKA [64], creating a random committee (RC) of classifiers. For this reason, the only classifiers (at the base level) that can be used in for this meta-algorithm are Random Forest (RF), Random Tree (RT), REP Tree (REPTree), Stochastic Gradient

Descent (SGD) and Multilayer Perceptron (MLP). The creation of a randomizable classifier is done by using an input (pseudo-random) seed. It is important to mention that the classifiers in the ensemble differ in terms the structure of their models. For instance, a random seed can define how the random trees are constructed in RF, RT and REPTree, how the linear models are defined in SGD, and how the network connection weights are firstly defined in MLP. Nevertheless, all classifiers are constructed using the same data, differently from Bagging and RSS. Thus, at the end, the final prediction is based on the average of the class probabilities generated by the base classifiers. **Parameters:**

- Classifier (c)[-W]: The classifier to be used for each member of the ensemble. It is restricted in one of the five (5) algorithms aforementioned, i.e., RF, RT, REPTree, SGD and MLP.
- Number of iterations (ni)[-I]: It defines the number of iterations to be performed, i.e., the number of classifiers in the ensemble. It may take values in the range: $\{ni \in \mathbb{Z} \mid 2 \leq ni \leq 64\}$.
Default value: 10.

There are no dependencies/constraints between the parameters of RC.

6.5 Ada Boost M1

The method for boosting a nominal class classifier using the Adaboost M1 (AdaM1) approach [23]. **Parameters:**

- Classifier (c)[-W]: The classifier to be used. It can be one classification algorithm from Section 5, except for the algorithms LMT, OneR, K*, SGD and VP, as these the classifiers produced by these algorithms do not handle weighted instances. It is important to mention that Auto-WEKA allows any classifier at the base level of AdaM1, including those which can not handle weights in the instances. In this case, Auto-WEKA ignores that algorithm (with its configuration) and proceeds with the search.
- Weight threshold (wt)[-P]: It defines the weight threshold for weighted pruning, i.e., it only selects instances with weights that contribute to the specified quantile of the weight distribution. It may take values in the range: $\{wt \in \mathbb{Z} \mid 50 \leq wt \leq 100\}$.
Default value: 100.
- Number of iterations (ni)[-I]: It defines the number of iterations to be performed, i.e., the number of classifiers in the ensemble. It may take the values in the range: $\{ni \in \mathbb{Z} \mid 2 \leq ni \leq 128\}$.
Default value: 10.

- Use resampling (ur)[-Q]: It decides whether AdaM1 will use resampling instead of reweighting. Thus, it is possible to generate an unweighted dataset from the weighted data by resampling. In this case, instances are chosen with probability proportional to their weight. As a result, instances with high weight are replicated frequently, and the ones with low weight may never be selected. Once the new dataset becomes as large as the original one, it is fed into the learning approach instead of the weighted data [64]. It can take

Boolean values (true or false).

Default value: false.

There are no dependencies/constraints between the parameters of AdaM1.

7 SEARCH SPACE – PREPROCESSING ALGORITHMS FROM WEKA

In this section, the *search space* of (single label) preprocessing classification algorithms from WEKA [27] is studied. This is also done in order to extend and improve the *search space* of multi-label methods. Instead of using just a single-label classification (SLC) algorithm at the SLC base level, a wrapper containing preprocessing methods is firstly used and, just after that, SLC is performed.

7.1 Attribute Selection Classifier

The method that reduces the dimensionality of training and test data by performing attribute selection (using the training set only) before the data is set as input to a classifier [64], constructing an attribute selection classifier (ASC). **Parameters:**

- Classifier (*c*)[-W]: The classifier to be used. It can be any classification algorithm from Section 5.
- Search method (*sm*)[-S]: The search method for selecting the attribute subset to be used as input by the classifier. It may take two values:
  - (1) Best First: It searches the space of attribute subsets by greedy hill-climbing augmented with a backtracking facility.
  - (2) Greedy Stepwise : It performs a greedy forward search through the space of attribute subsets.

Both methods use the evaluator “CfsSubsetEval”, which evaluates the worth of a subset of attributes by considering the individual predictive ability of each attribute along with the degree of redundancy between them. Hence, ASC is conceptually equivalent to using the CFS (Correlation-based Feature Selection) attribute selection method followed by the use of the chosen classifier with the attributes selected by CFS.

Default value: Best First.

There are no dependencies/constraints between the parameters of ASC.

8 A FORMAL DESCRIPTION OF THE MLC SEARCH SPACE

In this section, we describe the context-free grammar [58] used to formally specify the *search spaces* specified in this supplementary material.

Formally, a grammar G is represented by a four-tuple $\langle N, T, P, S \rangle$, where N represents a set of non-terminals, T a set of terminals, P a set of production rules and S (a member of N) the start symbol. In this thesis, we will use the Backus Naur Form (BNF) to represent grammars. This means that each production rule has, for instance, the following form $\langle \text{Start} \rangle ::= [\langle A \rangle] \langle B \rangle \mid \langle C \rangle \ (d \mid e)$. Symbols wrapped in “ $\langle \rangle$ ” represent non-terminals, whereas terminals (such as d and e) are not bounded by “ $\langle \rangle$ ”. The special symbols “ $|$ ”, “ $[]$ ” and “ $()$ ” represent, respectively, a choice, an optional element and a set of grouped elements that should be used

together. Additionally, the symbol “ $\#$ ” represents a comment in the grammar, i.e., it is ignored by the grammar’s parser. The choice of one among all elements connected by “ $|$ ” is made using a uniform probability distribution (i.e., all elements are equally likely to occur in an individual).

The proposed grammar has 125 production rules, in a total of 124 non-terminals and 213 terminals. Figures 1-7 present the produced grammar (in the Backus Naur Form) that encompasses the knowledge about multi-label classification in MEKA. This grammar models the *search space* Large, as this *search space* includes all learning algorithms and hyper-parameters from the other *search spaces* (i.e., Small and Medium). To consider the grammar for the other two *search spaces*, it is just needed to exclude the respective production rules, non-terminals and terminals – which represent the learning algorithms and their respective hyper-parameters – that do not make part of them. For sake of simplicity, we will not do that here.

In Figure 1, the first production rule ($\langle \text{Start} \rangle$) is used to describe the multi-label classification (MLC) *search space*. In this grammar rule, $\langle \text{MLC-PT} \rangle$ denotes problem transformation, $\langle \text{MLC-AA} \rangle$ denotes algorithm adaptation, and $\langle \text{META-MLC-LEVEL} \rangle$ denotes the multi-label meta-algorithms. We designed the grammar in such a way that all the 26 MLC algorithms have the same probability of being chosen (i.e., each MLC algorithm has $\approx 3.846\%$ of chance of being chosen). Furthermore, the MLC algorithms must use a prediction threshold ($\langle \text{pred\_tshd} \rangle$), which defines the threshold to perform the classification using the model’s confidence outputs [4].

The grammar rule defining the problem transformation methods, i.e., $\langle \text{MLC-PT} \rangle$, has two components in the right-hand side, namely the actual problem transformation algorithm $\langle \text{ALGS-PT} \rangle$ (defined in Figure 5) and the single-label classification algorithm (SLC, which is represented by the rule $\langle \text{ALG-SLC} \rangle$ in the grammar) to perform the single-label classification task(s). This happens because the problem transformation method transforms the multi-label task into one or more single-label tasks. We start discussing the rule defining $\langle \text{ALG-SLC} \rangle$.

We divided the SLC algorithms in six (6) types for the grammar following the WEKA software: Trees, Rules, Lazy, Functions, Bayes and Others. The last type was created just to simplify the grammar (i.e., it is not an inherent WEKA’s category). Figure 1 shows the grammar rules for Tree algorithms. Figure 2 shows the grammar rules for Rules and Lazy algorithms. Figure 3 shows the grammar rules for the other three types of SLC algorithms. Figure 1 also defines the Attribute Selection Classifier (ASC), a wrapper which can be used together with the SCL algorithms. In this case, a preprocessing method is used before the classification step is performed.

It is also important to mention that some methods at the single-label level, such as Decision Stump and ZeroR, do not have user-defined hyper-parameters. Others, such as the Bayesian Network Classifier algorithms, do not have user-defined hyper-parameters in the Auto-WEKA software, even though they have user-defined parameters in WEKA. That is, the developers of Auto-WEKA have chosen to use a fixed predefined number of parameter settings for some algorithms. As we are following Auto-WEKA to define the parameters at this level (because its robustness to select SLC algorithms), the absence of user-defined parameters in some methods was maintained.

```

<Start> ::= (<MLC-PT> | <MLC-AA> | <META-MLC-LEVEL> ) <pred_tshd>

<pred_tshd> ::= PCut1 | PCutL | RANDOM-REAL(>0.0, <1.0)    #pred_tshd='prediction threshold'
                                                         #PCut1='P-Cut method',PCutL='P-Cut method by Label'

<MLC-PT> ::= <ALGS-PT> <ALGS-SLC>

<ALGS-SLC> ::= <ALG-TYPE> | <META1> <ALG-WEIGHTED-TYPE> | <META2> <ALG-RANDOM-TYPE> | <META3> <ALG-TYPE>

<ALG-TYPE> ::= [ASC <sm>] (<TREES> | <RULES> | <LAZY> | <FUNCTIONS> | <BAYES> | <OTHERS>)
                                                         #ASC='Attribute Selection Classifier'

<sm> ::= GreedyStepwise | BestFirst                        #sm='search method'

<TREES> ::= <C4.5> | DecisionStump | ( ( (RandomForest <nt> | <RandomTree>) <nf> ) | <REPTree> ) <md>

<C4.5> ::= <C4.5-Basics> ( (<cf> [sr]) | u )                #sr='subtree raising', u='unpruned'
<C4.5-Basics> ::= <mno> [ct] [bs] [umc] [ul]                #ct='collapse tree', bs='binary splits'
                                                         #umc='use MDL correction', ul='use Laplace'
                                                         #cf='confidence factor'
                                                         #mno='minimum number of objects'

<cf> ::= RANDOM-REAL(0.0, 1.0)
<mno> ::= RANDOM-INT(1, 64)

<nt> ::= RANDOM-INT(2, 256)                                #nt='number of trees'
<nf> ::= RANDOM-INT(2, 32)                                  #nf='number of features'
<md> ::= RANDOM-INT(2, 20)                                  #md='maximum depth'

<RandomTree> ::= <mw> <nfbgt>
<mw> ::= RANDOM-INT(1,64)                                   #mw='minimum weight for instances in a leaf'
<nfbgt> ::= 2 | 3 | 4 | 5                                   #nfbgt='number of folds for back-fitting and
                                                         # for growing the tree'

<REPTree> ::= <mw> [up]                                     #up='use pruning'
                                                         #mw is not included in the same rule for Random Tree
                                                         #and for REPTree because of the grammar's constraints

```

Figure 1: Defined Grammar – Part 1: General and SLC Trees Algorithms.

```

<RULES> ::= <DT> | <JRip> | OneR <mbs> | <PART> | ZeroR

<DT> ::= <em> [uibk] <sm> <crv>                            #uibk='use IBk'
                                                         #sm='search method -- defined earlier'
                                                         #em='evaluation measure'
                                                         #crv='number of folds for cross-validation'

<em> ::= acc | rmse | mae | auc
<crv> ::= 1 | 2 | 3 | 4

<JRip> ::= <mtw> [cer] [up] <o>                             #cer='check error rate', up='use pruning'
<mtw> ::= RANDOM-REAL(1.0, 5.0)                            #mtw='minimum total weight for instances
                                                         # covered by a rule'
                                                         #o='number of optimization runs'

<o> ::= RANDOM-INT(1,5)

<mbs> ::= RANDOM-INT(1,32)                                   #mbs='minimum bucket size'

<PART> ::= <PART-BASICS> (rep <nr> | ebp)                   #rep='use reduced-error pruning'
                                                         #nr='number of folds for reduced-error pruning'
                                                         #ebp='use error-based pruning'
                                                         #mno='minimum number of objects'

<PART-BASICS> ::= <mno> [bs]
<nr> ::= RANDOM-INT(2,5)

<LAZY> ::= <KNN> | <K*>

<KNN> ::= <k_nn> [loo] [<dw>]                               #loo='leave-one-out to set the k value given the range'
<k_nn> ::= RANDOM-INT(1,64)                                  #k_nn='number of neighbors'
<dw> ::= F | I                                               #dw='distance weighting'

<K*> ::= <gb> [eab] <mm>                                     #eab='entropic auto-blending'
<gb> ::= RANDOM-INT(1,100)                                   #gb='global blending'
<mm> ::= a | d | m | n                                       #mm='missing mode to deal with missing values'

```

Figure 2: Defined Grammar – Part 2: SLC Rules and Lazy Algorithms .

At the single-label level, we also have meta-algorithms, divided in three (3) types: <META1>, <META2> and <META3>. These three

categories of meta-algorithms are firstly called in Figure 1. As shown

```

<FUNCTIONS> ::= <VotedPerceptron> | <MultiLayerPerc> |
               (<StocGradDescent> | LogisticRegression) <r> | <SeqMinOptimization>

<VotedPerceptron> ::= <i> <mk> <e>
<i> ::= RANDOM-INT(1,10)                #i='number of iterations'
<mk> ::= RANDOM-INT(5000, 50000)         #mk='maximum number of alterations to the perceptrons'
<e> ::= RANDOM-REAL(0.2, 5.0)           #e='The exponent for the polynomial kernel'

<MultiLayerPerc> ::= <lr> <m> <nhn> [n2b] [r] [d]
#n2b='nominal to binary filter',
#r='use reset approach',
#d='decay in the learning rate'
<lr> ::= RANDOM-REAL(0.1, 1.0)          #lr='learning rate'
<m> ::= RANDOM-REAL(0.0, 1.0)          #m='momentum'
<nhn> ::= a | i | o | t                #nhl='rules to define the number of hidden nodes'

<StocGradDescent> ::= <lf> <lr_sgd> [nn] [nrmv]
#nn='do not normalize',
#nrmv='do not replace missing values'
<lf> ::= 0 | 1 | 2                     #lf='loss function'
<lr_sgd> ::= RANDOM-REAL(0.00001, 1.0) #lr_sgd='learning rate for SGD'
<r> ::= RANDOM-REAL(0.000000000001,10.0) #r='ridge value in the log-likelihood'

<SeqMinOptimization> ::= <c> <ft> [bcm] <kernel>
#bcm='build calibration models'
<c> ::= RANDOM-REAL(0.5,1.5)           #c='the cost, i.e., complexity parameter'
<ft> ::= 0 | 1 | 2                     #ft='filter type'
<kernel> ::= ( NormPolyKernel |
               PolyKernel
               ) <exp> [ulo] |
               Puk <om> <sig> | RBF <g>
#ulo='use lower order'
<exp> ::= RANDOM-REAL(0.2, 5.0)         #exp='the exponent'
<om> ::= RANDOM-REAL(0.1, 1.0)         #om='the omega value'
<sig> ::= RANDOM-REAL(0.1, 10.0)       #sig='the sigma value'
<g> ::= RANDOM-REAL(0.001, 1.0)       #g='the gamma value'

<BAYES> ::= NaiveBayes [<NB-Parameters>] | <BayesianNetworkClassifiers> | NaiveBayesMultinomial
<NB-Parameters> ::= uke | usd
#uke='use kernel estimator'
#usd='use supervised distribution'

<BayesianNetworkClassifiers> ::= TAN | K2 | HillClimber | LAGDHillClimber | SimulatedAnnealing | TabuSearch

<OTHERS> ::= (SimpleLogistic [ucv] |
               <LogisticModelTrees>
               ) [uaic] [<wtb>]
#ucv='use cross-validation'
#uaic='use AIC measure as stopping criteria'

<LogisticModelTrees> ::= <mno> [cn] [sor] [fr] [eop]
#cn='convert nominal to binary'
#sor='split on residuals'
#fr='fast regression', eop='error on probabilities'

<wtb> ::= RANDOM-REAL(0.0, 1.0)       #wtb='weight trim beta'

```

Figure 3: Defined Grammar – Part 3: SLC Functions, Bayes and Others Algorithms.

in Figure 4, <META1> may take the two (2) meta-algorithms Ada Boost M1 (AdaM1) and Locally Weighted Learning (LWL), which need a base classifier at the SLC level that handles weighted instances. This is the reason the rule <ALG-WEIGHTED-TYPE> is defined. On the other hand, <META2> may take just one algorithm, i.e., Random Committee. The reason for that is because this SLC meta-algorithm can only be used with randomizable base classifiers. The rule <ALG-RANDOM-TYPE> expresses these randomizable classifiers. The least restricted meta-algorithms are Random Subspace and Bagging, which are specified by <META3>, being able to use any SLC base classifier (from <ALG-TYPE>).

It is important to emphasize that all 28 SLC methods (traditional, meta and preprocessing) have the same chance of being chosen by a *search method* that follows this grammar. Therefore, each SLC algorithm has a probability of $\approx 3.571\%$ of being selected.

The second component of problem transformation methods is the actual problem transformation algorithm to deal with the single-label classification. In other words, this step defines the choice of the MLC algorithm to handle the results created by the single-label classification models. For this component, we divided its respective algorithms into three categories, i.e., three production rules in the grammar (see Figure 5): <ALGS-PT1>, <ALGS-PT2> and <ALGS-PT3>. The main reason for the creation of these (sub-)categories is related to the constraints of the multi-label meta-algorithms in the MEKA software. Although all MLC algorithms can be used in a standalone fashion, they can also be combined with multi-label meta-algorithms. In MEKA, some MLC algorithms work very well at the multi-label base level of meta-algorithms, whereas others do not. Thus, we had to create rules in the grammar to overcome the limitations in the used software. The next paragraphs will refer to the Figures 5 and 7 to explain these links

```

<META1> ::= <LWL> | <AdaM1>

<LWL> ::= <k_lwl> [<wk>]                                #LWL='Locally Weighted Learning'
<k_lwl> ::= -1 | 10 | 30 | 60 | 90 | 120                #k_lwl='number of neighbors in LWL'
<wk> ::= 0 | 1 | 2 | 3 | 4                             #wk='weighting kernel'

<AdaM1> ::= <wt> [<ur>] <ni_ada_and_bagging>            #ur='use resampling'
<wt> ::= RANDOM-INT(50, 100) | 100                     #wt='weight threshold'
<ni_ada_and_bagging> ::= RANDOM-INT(2, 128)              #ni_ada_and_bagging='number of iterations for
                                                         # AdaM1 and Bagging'

<ALG-WEIGHTED-TYPE> ::= <TREES> | <RULES-PARTIAL> | <KNN> | <BAYES> | <FUNCTIONS-PARTIAL>
<RULES-PARTIAL> ::= <DT> | <JRip> | <PART> | ZeroR
<FUNCTIONS-PARTIAL> ::= <MultiLayerPerc> | <SeqMinOptimization> | <SimpleLogistic> <uaic> <wtb_activate>

<META2> ::= RandomCommittee <ni_random_methods>
<ni_random_methods> ::= RANDOM-INT(2, 64)               #ni_random_methods='number of iterations for
                                                         # random methods'
<ALG-RANDOM-TYPE> ::= ( ( (RandomForest <nt> | <RandomTree>) <nf> ) | <REPTree> ) <md> |
                     <StocGradDescent> <r> | <MultiLayerPerc>

<META3> ::= <Bagging> | <RandomSubspace>

<Bagging> ::= (<bsp> | 100 coob) <ni_ada_and_bagging>   #coob='calculate out-of-bag'
                                                         #when coob is true, bag percent size must be 100
<bsp> ::= RANDOM-INT(10, 100)                          #bsp='bag size percent'
<RandomSubspace> ::= <sss> <ni_random_methods>
<sss> ::= RANDOM-REAL(0.1, 1.0)                         #sss='subspace size'

```

Figure 4: Defined Grammar – Part 4: SLC Meta-Algorithms.

and constraints between problem transformation algorithms and multi-label meta-algorithms.

We referred to the first production rule to define problem transformation methods as <ALGS-PT1> in Figure 5. This rule encompasses the traditional algorithms BR, CC and LC (which is also known as LP). Besides, it includes the quick versions of BR and CC (i.e. BRq and CCq), all the complex classifier chains and trellis algorithms (which are defined by the rule <ComplexCC_Trellis>), Four-class pairWise (FW), Ranking and Threshold (RT), and all the label powerset based algorithms (which are defined by the rule <LP_based>). The production rule <ALGS-PT1> is presented in the following rules in Figure 7: <META-MLC1> (via <ALGS-PT>), <META-MLC2> and <META-MLC3>. This means that this category of PT methods describes the majority of the MLC algorithms in MEKA (84.21% of the cases, i.e., 16 of the 19 MLC algorithms) and, in addition, all these algorithms can be combined with all meta-algorithms in the MEKA software. Thus, <ALGS-PT1> can be considered the least restrictive of the PT method rules in the grammar.

<ALGS-PT2>, in Figure 5, is the production rule to describe solely the Bayesian Classifier Chain (BCC) algorithm, one of the most constrained algorithms in the MEKA software. The BCC algorithm can only be executed in a standalone fashion or combined with the algorithms described by the production rules <META-MLC1> (via <ALGS-PT>) and <META-MLC3>. This means that BCC can be used with four (4) of the seven (7) meta-algorithms (in Figure 7): Subset Mapper (SM), Random Subspace Multi-Label (RSML), Expectation Maximization (EM) and Classification Maximization (CM). In other cases of trying to use BCC, this will result in errors in MEKA's output and, therefore, this was not allowed in the grammar.

Similarly to <ALGS-PT2>, we have <ALGS-PT3>, a problem transformation rule that represents the Population of Monte-Carlo Classifier Chains (PMCC) algorithm. This algorithm can only be used by itself and at the multi-label base level of five (5) of the seven (7) meta-algorithms: Subset Mapper (SM), Random Subspace Multi-Label (RSML), Bagging of Multi-Label methods (BaggingML), Bagging of Multi-Label methods with Duplicates (BaggingMLDup) and Ensemble of Multi-Label methods (EnsembleML). These five multi-label meta-algorithms are defined by the production rules <META-MLC1> and <META-MLC2>. Therefore, the creation of <ALGS-PT2> is justified by the fact that PMCC algorithm can only be combined with these meta-algorithms, i.e., a constraint that did not appear in the other rules of the grammar.

Besides the problem transformation methods, we also have a multi-label version of the back propagation algorithm for training neural networks, called ML-BPNN. This algorithm can be seen in the Figure 6 and is the only one (for now) representing the algorithm adaptation (AA) methods, defined by the production rule <MLC-AA>. ML-BPNN can also be associated to meta-algorithms. As we can see in Figure 7, this MLC algorithm can be linked to the meta-algorithms defined by the production rules <META-MLC1>, <META-MLC2> and <META-MLC3>.

Finally, Figure 7 covers all the multi-label meta-algorithms, which are defined by the production rule <META-MLC-LEVEL>. As we explained previously, we created the production rules <META-MLC1>, <META-MLC2> and <META-MLC3> in order to expand these five rules into <META-MLC-LEVEL> to control the limitations, constraints and dependencies of the MEKA software between meta-algorithms and multi-label algorithms (problem transformation and algorithm adaptation methods).

```

<ALGS-PT> ::= <ALGS-PT1> | <ALGS-PT2> | <ALGS-PT3>

<ALGS-PT1> ::= BR | CC | LC | (BRq | CCq) <dsr> |
               <ComplexCC_Trellis> | FW | RT | <LP_based>
               #BR='Binary Relevance', CC='Classifier Chain'
               #LC='Label Combination'
               #BRq and CCq = 'quick versions for BR and CC'
               #FW='Four-class pairWise', RT='Ranking-Threshold'

<ALGS-PT2> ::= BCC <dp_complete>
               #BCC='Bayesian Classifier Chain'
<ALGS-PT3> ::= PMCC <B> <ts> <ii> <chi_PMCC> <ps> <pof>
               #PMCC='Population of Monte-Carlo Classifier Chains'

<dsr> ::= RANDOM-REAL(0.2, 0.8)
               #dsr='down-sample ratio'

<ComplexCC_Trellis> ::= PCC | (MCC <chi_MCC> | <CT>) <ii> <pof> |
               (CDN | <CDT>) <i_cdn_cdt> <ci>
               #PCC='Probabilistic Classifier Chains'
               #MCC='Monte-Carlo Classifier Chains'
               #CT='Classification Trellis'
               #CDN='Conditional Dependency Networks'
               #CDT='Conditional Dependency Trellis'

<chi_MCC> ::= <chi_CT> | 0
               #chi_MCC='number of chain iterations for MCC'
<ii> ::= RANDOM-INT(2, 100)
               #ii='number of inference iterations'
<pof> ::= Accuracy | Jaccard index | Hamming score | Exact match | Jaccard distance | Rank loss |
               Hamming loss | Zero One loss | Harmonic score | Log Loss lim:L | Micro Recall | One error |
               Log Loss lim:D | Micro Precision | Macro Precision | Macro Recall | F1 micro averaged |
               Avg precision | F1 macro averaged by example | F1 macro averaged by label | AUPRC macro averaged |
               AUROC macro averaged | Levenshtein distance
               #pof='Payoff function'

<CT> ::= <chi_CT> <w> <dp>
               #dp='dependency type'
<dp> ::= C | I | Ib | IbF | H | Hbf | X | F | None
               #chi_CT='number of chain iterations for CT'
<chi_CT> ::= RANDOM-INT(2, 1500)
               #w='width of the trellis'
<w> ::= 0 | 1 | -1 <d>
               #d='neighborhood density'
<d> ::= RANDOM-INT(1, SQRT(L) + 1)
               #Where L is the number of labels
               #parameters defined earlier

<CDT> ::= <w> <dp>
               #parameters defined earlier

<i_cdn_cdt> ::= RANDOM-INT(101, 1000)
               #i_cdn_cdt='total number of iterations'
<ci> ::= RANDOM-INT(1, 100)
               #ci='collection iterations'

<LP_based> ::= (PS | PSt | <RAKEL-based>) <sv> <pv>
               #PS='Pruned Sets'
               #PSt='Pruned Sets with Threshold'
               #sv='subsampling value'
               #pv='pruning value'

<RAKEL-based> ::= (RAKEL <sre> | RAKELd) <les>
               #RAKEL='Random k-labEL Pruned Sets'
               #RAKELd='Random k-labEL Disjoint Pruned Sets'
               #sre='number of subsets to run in an ensemble'
               #les='number of labels in each label subset'
               #Where L is the number of labels

<dp_complete> ::= <dp> | LEAD
               #dp='complete dependency type for BCC'

<B> ::= RANDOM-REAL(0.01, 0.99)
               #B='Beta factor for deacresing the temperature'
<ts> ::= 0 | 1
               #ts='Temperature switch'
<ps> ::= RANDOM-INT(1, 50)
               #ps='population size'
<chi_PMCC> ::= RANDOM-INT(51, 1500)
               #chi_PMCC='number of chain iterations for PMCC'

```

Figure 5: Defined Grammar – Part 5: MLC Problem Transformation Methods.

```

<MLC-AA> ::= <ML-BPNN>

<ML-BPNN> ::= <ne> <nhu_bpnn> <lr_bpnn> <m_bpnn>
               #ML-BPNN='Multi-Label Back Propagation
               #Neural Network'
               #ne='number of epochs'
               #nhu_bpnn='number of hidden units, that
               #is a parameter that depends on the
               #number of attributes of the dataset'

<ne> ::= RANDOM-INT(10, 1000)
<nhu_bpnn> ::= RANDOM-REAL(0.2, 1.0) * n_attributes

<lr_bpnn> ::= RANDOM-REAL(0.001, 0.1)
               #lr_bpnn='learning rate for BPNN/DBPNN'
<m_bpnn> ::= RANDOM-REAL(0.2, 0.8)
               #m_bpnn='momentum for BPNN and DBPNN'

```

Figure 6: Defined Grammar – Part 6: MLC Algorithm Adaptation Methods.

```

<META-MLC-LEVEL> ::= <META-MLC1> | <META-MLC2> | <META-MLC3>
                                     #META-MLC 1-3=`meta MLC algorithms
                                     # with different constraints'

<META-MLC1> ::= (SM | <RSML>) (<ALGS-PT> <ALGS-SLC> | <ML-BPNN>)
                                     #SM=`Subset Mapper -- MLC method as parameter'

<RSML> ::= <bsp> <i_metamlc> <ap>
                                     #RSML=`Random Subspace Multi-Label'
<bsp> ::= RANDOM-INT(10, 100)
                                     #bsp=`bag size percent'
<i_metamlc> ::= RANDOM-INT(10, 50)
                                     #i_metamlc=`number of iterations for
                                     #meta MLC methods'
<ap> ::= RANDOM-INT(10, 100)
                                     #ap=`attribute percent'

<META-MLC2> ::= <alg-meta-mlc2> (<ALGS-PT1> | <ALGS-PT3>) <ALGS-SLC> | <ML-BPNN>

<alg-meta-mlc2> ::= ((BaggingML | BaggingMLDup <bsp> ) | EnsembleML <bsp_ensembleML>) <i_metamlc>
                                     #BaggingML=`Bagging of Multi-Label methods'
                                     #BaggingMLDup=`BaggingML with duplicates'
                                     #EnsembleML=`Ensemble of Multi-Label methods'
                                     #bsp=`bag size percent -- defined earlier'

<bsp_ensembleML> ::= RANDOM-INT(52, 72)
                                     #bsp_ensembleML=`specific bsp for EnsembleML'

<META-MLC3> ::= ( (EM | CM ) <i_metamlc> ) (<ALGS-PT1> | <ALGS-PT2>) <ALGS-SLC> | <ML-BPNN>
                                     #EM=`Expectation Maximization'
                                     #CM=`Classification Maximization'

```

Figure 7: Defined Grammar – Part 7: MLC Meta-Algorithms.

REFERENCES

- [1] Michael Abramovici, Manuel Neubach, Madjid Fathi, and Alexander Holland. 2008. Competing Fusion for Bayesian Applications. In *Proceedings of Information Processing and Management of Uncertainty in Knowledge-Based Systems*. 378–385.
- [2] Charu C. Aggarwal and ChengXiang Zhai. 2012. A Survey of Text Classification Algorithms. In *Mining Text Data*. Springer US, Boston, MA, USA, 163–222.
- [3] David W. Aha, Dennis Kibler, and Marc K. Albert. 1991. Instance-Based Learning Algorithms. *Machine Learning* 6, 1 (1991), 37–66.
- [4] Reem Al-Otaibi, Peter Flach, and Meelis Kull. 2014. Multi-Label Classification: A Comparative Study on Threshold Selection Methods. In *Proceedings of the International Workshop on Learning over Multiple Contexts, collocated with the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (LMCE/ECML-PKDD'14)*. 1–8.
- [5] Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. 1997. Locally Weighted Learning. *Artificial Intelligence Review* 11, 1 (1997), 11–73.
- [6] Remco R. Bouckaert. 1995. *Bayesian Belief Networks: From Construction to Inference*. Ph.D. Dissertation. University of Utrecht, Utrecht, Netherlands.
- [7] Remco R. Bouckaert. 2007. *Bayesian Network Classifiers in WEKA*. Technical Report. Department of Computer Science, University of Waikato, Hamilton, New Zealand.
- [8] Jeffrey P. Bradford, Clayton Kunz, Ron Kohavi, Clifford Brunk, and Carla E. Brodley. 1998. Pruning Decision Trees with Misclassification Costs. In *Proceedings of the European Conference on Machine Learning (ECML'98)*. Springer-Verlag, Berlin/Heidelberg, Germany, 131–136.
- [9] Leo Breiman. 1996. Bagging Predictors. *Machine Learning* 24, 2 (1996), 123–140.
- [10] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (2001), 5–32.
- [11] Saskia L. Cessie and Johannes C. van Houwelingen. 1992. Ridge Estimators in Logistic Regression. *Applied Statistics* 41, 1 (1992), 191–201.
- [12] John G. Cleary and Leonard E. Trigg. 1995. K*: An Instance-based Learner Using an Entropic Distance Measure. In *Proceedings of the International Conference on Machine Learning (ICML'95)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 108–114.
- [13] William W. Cohen. 1995. Fast Effective Rule Induction. In *Proceedings of the International Conference on Machine Learning (ICML'95)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 115–123.
- [14] Gregory F. Cooper and Edward Herskovits. 1992. A Bayesian Method for the Induction of Probabilistic Networks from Data. *Machine Learning* 9, 4 (1992), 309–347.
- [15] Krzysztof Dembczyński, Weiwei Cheng, and Eyke Hüllermeier. 2010. Bayes Optimal Multilabel Classification via Probabilistic Classifier Chains. In *Proceedings of the International Conference on Machine Learning (ICML'10)*. Omnipress, Madison, WI, USA, 279–287.
- [16] Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society, Series B* 39, 1 (1977), 1–38.
- [17] James Dougherty, Ron Kohavi, and Mehran Sahami. 1995. Supervised and Unsupervised Discretization of Continuous Features. In *Machine Learning Proceedings 1995*. Elsevier, Amsterdam, Netherlands, 194–202.
- [18] Rong-En Fan and Chih-Jen Lin. 2007. *A Study on Threshold Selection for Multi-Label Classification*. Technical Report. Department of Computer Science, National Taiwan University, Taipei City, Taiwan. 23 pages.
- [19] Usama M. Fayyad and Keki B. Irani. 1993. Multi-interval Discretization of Continuous-Valued Attributes for Classification Learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'93)*, Vol. 2. Morgan Kaufmann Publishers, San Francisco, CA, USA, 1022–1027.
- [20] Eibe Frank and Remco R. Bouckaert. 2006. Naive Bayes for Text Classification with Unbalanced Classes. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD'06)*. Springer-Verlag, Berlin/Heidelberg, Germany, 503–510.
- [21] Eibe Frank, Mark Hall, and Bernhard Pfahringer. 2003. Locally Weighted Naive Bayes. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI'03)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 249–256.
- [22] Eibe Frank and Ian H. Witten. 1998. Generating Accurate Rule Sets Without Global Optimization. In *Proceedings of the International Conference on Machine Learning (ICML'98)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 144–151.
- [23] Yoav Freund and Robert E. Schapire. 1996. Experiments with a New Boosting Algorithm. In *Proceedings of the International Conference on Machine Learning (ICML'96)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 148–156.
- [24] Yoav Freund and Robert E. Schapire. 1999. Large Margin Classification Using the Perceptron Algorithm. *Machine Learning* 37, 3 (1999), 277–296.
- [25] Nir Friedman, Dan Geiger, and Moises Goldszmidt. 1997. Bayesian Network Classifiers. *Machine Learning* 29, 2–3 (1997), 131–163.
- [26] Yuhong Guo and Suicheng Gu. 2011. Multi-label Classification Using Conditional Dependency Networks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'11)*, Vol. 2. AAAI Press, Palo Alto, CA, USA, 1300–1305.
- [27] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA Data Mining Software: An Update. *ACM SIGKDD Explorations Newsletter* 11, 1 (2009), 10–18.
- [28] Trevor Hastie and Robert Tibshirani. 1998. Classification by Pairwise Coupling. In *Proceedings of the Conference on Advances in Neural Information Processing Systems (NIPS'97)*. MIT Press, Cambridge, MA, USA, 507–513.
- [29] David Heckerman, Dan Geiger, and David M. Chickering. 1994. Learning Bayesian Networks: The Combination of Knowledge and Statistical Data. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*. AAAI Press, Palo Alto, CA, USA, 85–96.
- [30] Alireza S. Hesar, Hamid Tabatabaee, and Mehrdad Jalali. 2012. Structure Learning of Bayesian Networks Using Heuristic Methods. In *Proceedings of International Conference on Information and Knowledge Management (ICKM'12)*. Engineering & Technology Digital Library.

- [31] Tin Kam Ho. 1998. The Random Subspace Method for Constructing Decision Forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 8 (1998), 832–844.
- [32] Robert C. Holte. 1993. Very Simple Classification Rules Perform Well on Most Commonly Used Datasets. *Machine learning* 11, 1 (1993), 63–90.
- [33] Thorsten Joachims. 2002. *Learning to Classify Text Using Support Vector Machines: Methods, Theory and Algorithms*. Kluwer Academic Publishers, Hingham, MA, USA.
- [34] George H. John and Pat Langley. 1995. Estimating Continuous Distributions in Bayesian Classifiers. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI'95)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 338–345.
- [35] S. S. Keerthi, S. K. Shevade, C. Bhattacharyya, and K. R. K. Murthy. 2001. Improvements to Platt's SMO Algorithm for SVM Classifier Design. *Neural Computing* 13, 3 (2001), 637–649.
- [36] Scott Kirkpatrick, Charles D. Gelatt, and Mario P. Vecchi. 1983. Optimization by Simulated Annealing. *Science* 220, 4598 (1983), 671–680.
- [37] Ron Kohavi. 1995. The Power of Decision Tables. In *Proceedings of the European Conference on Machine Learning (ECML'95)*. Springer-Verlag, Berlin/Heidelberg, Germany, 174–189.
- [38] Lars Kotthoff, Chris Thornton, Holger H. Hoos, Frank Hutter, and Kevin Leyton-Brown. 2017. Auto-WEKA 2.0: Automatic Model Selection and Hyperparameter Optimization in WEKA. *Journal of Machine Learning Research* 18, 1 (2017), 826–830.
- [39] Lars Kotthoff, Chris Thornton, and Frank Hutter. 2017. *User Guide for Auto-WEKA Version 2.6*. Technical Report. Computer Science Department at University British Columbia, Vancouver, BC, Canada. 15 pages.
- [40] Niels Landwehr, Mark Hall, and Eibe Frank. 2005. Logistic Model Trees. *Machine Learning* 59, 1 (2005), 161–205.
- [41] David D. Lewis. 1998. Naïve (Bayes) at Forty: The Independence Assumption in Information Retrieval. In *Proceedings of the European Conference on Machine Learning (ECML'1998)*. Springer-Verlag, Berlin/Heidelberg, German, 4–15.
- [42] Gjorgji Madjarov, Dragi Kocev, Dejan Gjorgjevikj, and Sašo Deroski. 2012. An Extensive Experimental Comparison of Methods for Multi-Label Learning. *Pattern Recognition* 45, 9 (2012), 3084–3104.
- [43] Andrew McCallum and Kamal Nigam. 1998. A Comparison of Event Models for Naïve Bayes Text Classification. In *Proceedings of the Workshop on Learning for Text Categorization*. AAAI Press, Palo Alto, CA, USA, 41–48.
- [44] John C. Platt. 1999. Advances in Kernel Methods. In *Advances in Kernel Methods*, Bernhard Schölkopf, Christopher J. C. Burges, and Alexander J. Smola (Eds.). MIT Press, Cambridge, MA, USA, Chapter Fast Training of Support Vector Machines Using Sequential Minimal Optimization, 185–208.
- [45] John Ross Quinlan. 1993. *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [46] Jesse Read. 2008. A Pruned Problem Transformation Method for Multi-Label Classification. In *Proceedings of the New Zealand Computer Science Research Student Conference (NZCSRS'08)*. University of Canterbury, Christchurch, Canterbury, New Zealand, 143–150.
- [47] Jesse Read. 2010. *Scalable Multi-label Classification*. Ph.D. Dissertation. The University of Waikato, Hamilton, New Zealand.
- [48] Jesse Read, Luca Martino, and David Luengo. 2013. Efficient Monte Carlo Optimization for Multi-label Classifier Chains. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP'13)*. IEEE, New York, NY, USA, 3457–3461.
- [49] Jesse Read, Luca Martino, and David Luengo. 2014. Efficient Monte Carlo Methods for Multi-dimensional Learning with Classifier Chains. *Pattern Recognition* 47, 3 (2014), 1535–1546.
- [50] Jesse Read, Luca Martino, Pablo M. Olmos, and David Luengo. 2015. Scalable Multi-Output Label Prediction: From Classifier Chains to Classifier Trellises. *Pattern Recognition* 48, 6 (2015), 2096–2109.
- [51] Jesse Read and Fernando Perez-Cruz. 2014. Deep Learning for Multi-Label Classification. *ArXiv arXiv:1502.05988* (2014), 1–8.
- [52] Jesse Read, Bernhard Pfahringer, and Geoff Holmes. 2008. Multi-label Classification Using Ensembles of Pruned Sets. In *Proceedings of the International Conference on Data Mining*. IEEE, New York, NY, USA, 995–1000.
- [53] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. 2011. Classifier Chains for Multi-label Classification. *Machine Learning* 85, 3 (2011), 333–359.
- [54] 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.
- [55] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. 1986. Learning Representations by Back-Propagating Errors. *Nature* 323, 6088 (1986), 533–536.
- [56] Steven L. Salzberg. 1994. C4.5: Programs for Machine Learning by J. Ross Quinlan. Morgan Kaufmann Publishers, Inc., 1993. *Machine Learning* 16, 3 (1994), 235–240.
- [57] Robert E. Schapire and Yoram Singer. 1999. Improved Boosting Algorithms Using Confidence-Rated Predictions. *Machine Learning* 37, 3 (1999), 297–336.
- [58] Michael Sipser. 2012. *Introduction to the Theory of Computation* (3rd ed.). Cengage Learning, Boston, MA, USA.
- [59] Marc Sumner, Eibe Frank, and Mark Hall. 2005. Speeding Up Logistic Model Tree Induction. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD'05)*. Springer-Verlag, Berlin/Heidelberg, Germany, 675–683.
- [60] Chris Thornton, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. 2013. Auto-WEKA: Combined Selection and Hyperparameter Optimization of Classification Algorithms. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'13)*. ACM, New York, NY, USA, 847–855.
- [61] Grigoris Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. 2010. *Mining Multi-label Data*. Springer, Boston, MA, USA, 667–685.
- [62] 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.
- [63] Bülent Üstün, Willem J. Melssen, and Lutgarde M. C. Buydens. 2006. Facilitating the Application of Support Vector Regression by Using a Universal Pearson VII Function Vased Kernel. *Chemometrics and Intelligent Laboratory Systems* 81, 1 (2006), 29–40.
- [64] Ian H. Witten, Eibe Frank, Mark A. Hall, and Christopher J. Pal. 2016. *Data Mining: Practical Machine Learning Tools and Techniques* (4th ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [65] Julio H. Zaragoza, Luis E. Sucar, Eduardo F. Morales, Concha Bielza, and Pedro Larrañaga. 2011. Bayesian Chain Classifiers for Multidimensional Classification. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'11)*. AAAI Press, Palo Alto, CA, USA, 2192–2197.
- [66] Min-Ling Zhang and Zhi-Hua Zhou. 2006. Multilabel Neural Networks with Applications to Functional Genomics and Text Categorization. *IEEE Transactions on Knowledge and Data Engineering* 18, 10 (2006), 1338–1351.