



On the Trade-Off Between Consistency and Coverage in Multi-label Rule Learning Heuristics

Michael Rapp^(✉), Eneldo Loza Mencía, and Johannes Fürnkranz

Knowledge Engineering Group, TU Darmstadt, Darmstadt, Germany
{mrapp, eneldo, juffi}@ke.tu-darmstadt.de

Abstract. Recently, several authors have advocated the use of rule learning algorithms to model multi-label data, as rules are interpretable and can be comprehended, analyzed, or qualitatively evaluated by domain experts. Many rule learning algorithms employ a heuristic-guided search for rules that model regularities contained in the training data and it is commonly accepted that the choice of the heuristic has a significant impact on the predictive performance of the learner. Whereas the properties of rule learning heuristics have been studied in the realm of single-label classification, there is no such work taking into account the particularities of multi-label classification. This is surprising, as the quality of multi-label predictions is usually assessed in terms of a variety of different, potentially competing, performance measures that cannot all be optimized by a single learner at the same time. In this work, we show empirically that it is crucial to trade off the consistency and coverage of rules differently, depending on which multi-label measure should be optimized by a model. Based on these findings, we emphasize the need for configurable learners that can flexibly use different heuristics. As our experiments reveal, the choice of the heuristic is not straight-forward, because a search for rules that optimize a measure locally does usually not result in a model that maximizes that measure globally.

Keywords: Multi-label classification · Rule learning · Heuristics

1 Introduction

As many real-world classification problems require to assign more than one label to an instance, multi-label classification (MLC) has become a well-established topic in the machine learning community. There are various applications of MLC such as text categorization [16], the annotation of images [4, 18] and music [24, 26], as well as use cases in bioinformatics [8] and medicine [20].

Rule learning algorithms are a well-researched approach to solve classification problems [13]. In comparison to complex statistical methods, like for example support vector machines or artificial neural networks, their main advantage is the interpretability of the resulting models. Rule-based models can easily be

understood by humans and form a structured hypothesis space that can be analyzed and modified by domain experts. Ideally, rule-based approaches are able to yield insight into the application domain by revealing patterns and regularities hidden in the data and allow to reason why individual predictions have been made by a system. This is especially relevant in safety-critical domains, such as medicine, power systems, or financial markets, where malfunctions and unexpected behavior may entail the risk of health damage or financial harm.

Motivation and Goals. To assess the quality of multi-label predictions in terms of a single score, several commonly used performance measures exist. Even though some of them originate from measures used in binary or multi-class classification, different ways to aggregate and average the predictions for individual labels and instances—most prominently *micro-* and *macro-averaging*—exist in MLC. Some measures like *subset accuracy* are even unique to the multi-label setting. No studies that investigate the effects of using different rule learning heuristics in MLC and discuss how they affect different multi-label performance measures have been published so far.

In accordance with previous publications in single-label classification, we argue that all common rule learning heuristics basically trade off between two aspects, *consistency* and *coverage* [12]. Our long-term goal is to better understand how these two aspects should be weighed to assess the quality of candidate rules during training if one is interested in a model that optimizes a certain multi-label performance measure. As a first step towards this goal, we present a method for flexibly creating rule-based models that are built with respect to certain heuristics. Using this method, we empirically analyze how different heuristics affect the models in terms of predictive performance and model characteristics. We demonstrate how models that aim to optimize a given multi-label performance measure can deliberately be trained by choosing a suitable heuristic. By comparing our results to a state-of-the-art rule learner, we emphasize the need for configurable approaches that can flexibly be tailored to different multi-label measures. Due to space limitations, we restrict ourselves to micro-averaged measures, as well as to Hamming and subset accuracy.

Structure of This Work. We start in Sect. 2 by giving a formal definition of multi-label classification tasks as well as an overview of inductive rule learning and the rule evaluation measures that are relevant to this work. Based on these foundations, in Sect. 3, we discuss our approach for flexibly creating rule-based classifiers that are built with respect to said measures. In Sect. 4, we present the results of the empirical study we have conducted, before we provide an overview of related work in Sect. 5. Finally, we conclude in Sect. 6 by recapitulating our results and giving an outlook on planned future work.

2 Preliminaries

MLC is a supervised learning problem in which the task is to associate an instance with one or several labels λ_i out of a finite label space $\mathbb{L} = \{\lambda_1, \dots, \lambda_n\}$, with $n = |\mathbb{L}|$ being the total number of predefined labels. An individual instance \mathbf{x}_j is represented in attribute-value form, i.e., it consists of a vector $\mathbf{x}_j = (v_1, \dots, v_l) \in \mathbb{D} = A_1 \times \dots \times A_l$, where A_i is a numeric or nominal attribute. Additionally, each instance \mathbf{x}_j is associated with a binary label vector $\mathbf{y}_j = (y_1, \dots, y_n) = \{0, 1\}^n$, where y_i indicates the presence (1) or absence (0) of label λ_i . Consequently, the training data set of a MLC problem can be defined as a set of tuples $T = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_m, \mathbf{y}_m)\}$, with $m = |T|$ being the number of available training instances. The classifier function $g(\cdot)$, that is deduced from a given training data set, maps an instance \mathbf{x} to a predicted label vector $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_n) = \{0, 1\}^n$.

2.1 Classification Rules

In this work, we are concerned with the induction of conjunctive, propositional rules $r : H \leftarrow B$. The body B of such a rule consists of one or several conditions that compare an attribute-value v_i of an instance to a constant by using a relational operator such as $=$ (in case of nominal attributes), or $<$ and \geq (in case of numerical attributes). On the one hand, the body of a conjunctive rule can be viewed as a predicate $B : \mathbf{x} \rightarrow \{\text{true}, \text{false}\}$ that states whether an instance \mathbf{x} satisfies all of the given conditions, i.e., whether the instance is *covered* by the rule or not. On the other hand, the head H of a (single-label head) rule consists of a single label assignment ($\hat{y}_i = 0$ or $\hat{y}_i = 1$) that specifies whether the label λ_i should be predicted as present (1) or absent (0).

2.2 Binary Relevance Method

In the present work, we use the *binary relevance* transformation method (cf. [4]), which reduces MLC to binary classification by treating each label $\lambda_i \in \mathbb{L}$ of a MLC problem independently. For each label λ_i , we aim at learning rules that predict the minority class $t_i \in \{0, 1\}$, i.e., rules that contain the label assignment $\hat{y}_i = t_i$ in their head. We define $t_i = 1$, if the corresponding label λ_i is associated with less than 50% of the training instances, or $t_i = 0$ otherwise.

A rule-based classifier—also referred to as a *theory*—combines several rules into a single model. In this work, we use (unordered) rule sets containing all rules that have been induced for the individual labels. Such a rule set can be considered as a disjunction of conjunctive rules (DNF). At prediction time, all rules that cover a given instance are taken into account to determine the predicted label vector $\hat{\mathbf{y}}$. An individual element $\hat{y}_i \in \hat{\mathbf{y}}$, that corresponds to the label λ_i , is set to the minority class t_i if at least one of the covering rules contains the label assignment $\hat{y}_i = t_i$ in its head. Otherwise, the element is set to the majority class $1 - t_i$. As all rules that have been induced for a label λ_i have the same head, no conflicts may arise in the process.

2.3 Bipartition Evaluation Functions

To assess the quality of individual rules, usually bipartition evaluation functions $\delta : \mathbb{N}^{2 \times 2} \rightarrow \mathbb{R}$ are used [25]. Such functions—also called *heuristics*—map a two-dimensional confusion matrix to a heuristic value $h \in [0, 1]$. A confusion matrix consists of the number of *true positive* (*TP*), *false positive* (*FP*), *true negative* (*TN*), and *false negative* (*FN*) labels that are predicted by a rule. We calculate the example-wise aggregated confusion matrix C_r for a rule $r : \hat{y}_i \leftarrow B$ as

$$C_r := \begin{pmatrix} TP & FP \\ FN & TN \end{pmatrix} = C_i^1 \oplus \dots \oplus C_i^j \oplus \dots \oplus C_i^m \quad (1)$$

where \oplus denotes the cell-wise addition of atomic confusion matrices C_i^j that correspond to label λ_i and instance \mathbf{x}_j .

Further, let y_i^j and \hat{y}_i^j denote the absence (0) or presence (1) of label λ_i for an instance \mathbf{y}_j according to the ground truth and a rule’s prediction, respectively. Based on these variables, we calculate the elements of C_i^j as

$$\begin{aligned} TP_i^j &= \llbracket y_i^j = t_i \wedge \hat{y}_i^j = t_i \rrbracket & FP_i^j &= \llbracket y_i^j \neq t_i \wedge \hat{y}_i^j = t_i \rrbracket \\ FN_i^j &= \llbracket y_i^j = t_i \wedge \hat{y}_i^j \neq t_i \rrbracket & TN_i^j &= \llbracket y_i^j \neq t_i \wedge \hat{y}_i^j \neq t_i \rrbracket \end{aligned} \quad (2)$$

where $\llbracket x \rrbracket = 1$, if x is true, 0 otherwise.

2.4 Rule Learning Heuristics

A good rule learning heuristic should (among other aspects) take both, the *consistency* and *coverage* of a rule, into account [13, 15]. On the one hand, rules should be consistent, i.e., their prediction should be correct for as many of the covered instances as possible. On the other hand, rules with great coverage, i.e., rules that cover a large number of instances, tend to be more reliable, even though they may be less consistent.

The *precision* metric exclusively focuses on the consistency of a rule. It calculates as the fraction of correct predictions among all covered instances:

$$\delta_{prec}(C) := \frac{TP}{TP + FP} \quad (3)$$

In contrast, *recall* focuses on the coverage of a rule. It measures the fraction of covered instances among all—covered and uncovered—instances for which the label assignment in the rule’s head is correct:

$$\delta_{rec}(C) := \frac{TP}{TP + FN} \quad (4)$$

The *F-measure* calculates as the (weighted) harmonic mean of precision and recall. It allows to trade off the consistency and coverage of a rule depending on the user-configurable parameter β :

$$\delta_F(C) := \frac{\beta^2 + 1}{\frac{\beta^2}{\delta_{rec}(C)} + \frac{1}{\delta_{prec}(C)}}, \text{ with } \beta \in [0, +\infty] \quad (5)$$

As an alternative to the F-measure, we use different parameterizations of the *m-estimate* in this work. It is defined as

$$\delta_m(C) := \frac{TP + m \cdot \frac{P}{P+N}}{TP + FP + m}, \text{ with } m \geq 0 \quad (6)$$

where $P = TP + FN$ and $N = FP + TN$. Depending on the parameter m , this measure trades off precision and *weighted relative accuracy* (WRA). If $m = 0$, it is equivalent to precision and therefore focuses on consistency. As m approaches $+\infty$, it converges to WRA and puts more emphasis on coverage, respectively [13].

3 Induction of Rule-Based Theories

For our experimental study, we implemented a method that allows to generate a large number of rules for a given training data set in a short amount of time (cf. Sect. 3.1).¹ The rules should ideally be unbiased, i.e., they should not be biased in favor of a certain heuristic, and they should be diverse, i.e., general rules should be included as well as specific rules. Given that these requirements are met, we consider the generated rules to be representative samples for the space of all possible rules, which is way too large to be explored exhaustively. We use the generated candidate rules as a starting point for building different theories. They consist of a subset of rules that are selected with respect to a specific heuristic (cf. Sect. 3.2) and filtered according to a threshold (cf. Sect. 3.3). Whereas the first step yields a theory with great coverage, the threshold selection aims at improving its consistency.

3.1 Generation of Candidate Rules

As noted in Sect. 2.2, we consider each label $\lambda_i \in \mathbb{L}$ of a MLC problem independently. For each of the labels we train multiple random forests [5], using varying configuration parameters, and extract rules from their decision trees.² As illustrated in Algorithm 1, we repeat the process until a predefined number of rules γ has been generated.

Each random forest consists of a predefined number of decision trees (we specify $I = 10$). To ensure that we are able to generate diverse rules later on, we vary the configuration parameter *depth* $\in [0, 8]$ that specifies the maximum depth of trees (unrestricted, if *depth* = 0) (cf. Algorithm 1, `trainForest`). For building individual trees, we only take a subset of the available training instances and attributes into account, which guarantees a diverse set of trees. Bagging is used for sampling the training instances, i.e., if m instances are available in total, $m \cdot P$ instances ($P = 100\%$, by default) are drawn randomly with replacement. Additionally, each time a new node is added to a decision tree, only a random selection of K out of l attributes ($K = \log_2(l - 1) + 1$, by default) is considered.

¹ Source code available at <https://github.com/mrapp-ke/RuleGeneration>.

² We use the random forest implementation provided by Weka 3.9.3, which is available at <https://www.cs.waikato.ac.nz/ml/weka>.

Algorithm 1: Iterative generation of rules from random forests

```

input : min. number of rules to be generated  $\gamma$ 
output: rule set  $R$ 
 $R = \emptyset$ 
while  $|R| < \gamma$  do
  foreach  $\lambda_i \in \mathbb{L}$  and  $depth \in [0, 8]$  do
     $rf = \text{trainForest}(\lambda_i, depth)$ 
     $R = R \cup \text{extractRules}(rf)$ 
return  $R$ 

```

To extract rules from a random forest (cf. Algorithm 1, `extractRules`), we traverse all paths from the root node to a leaf in each of its decision trees. We only consider paths that lead to a leaf where the minority class t_i is predicted. As a consequence, all rules that are generated with respect to a certain label λ_i have the same head $\hat{y}_i = t_i$. The body of a rule consists of a conjunction of all conditions encountered on the path from the root to the corresponding leaf.

3.2 Candidate Subset Selection

Like many traditional rule learning algorithms, we use a *separate-and-conquer* (SeCo) strategy for selecting candidate rules, i.e., new rules are added to the theory until all training instances are covered (or until it describes the training data sufficiently according to some stopping criterion). Whenever a new rule is added to the theory the training instances it covers are removed (“separate” step), and the next rule is chosen according to its performance on the remaining instances (“conquer” step).

To create different theories, we select subsets of the rules that have been generated earlier (cf. Sect. 3.1). We therefore apply the SeCo strategy for each label independently, i.e., for each label λ_i we take all rules with head $\hat{y}_i = t_i$ into account. Among these candidates we successively select the best rule according to a heuristic δ (cf. Sect. 2.4) until all *positive* training instances $P_i = \{(\mathbf{x}, \mathbf{y}) \in T \mid y_i = t_i\}$, with respect to label λ_i , are covered. To measure the quality of a candidate \mathbf{r} according to δ , we only take yet uncovered instances into account for computing the confusion matrix $C_{\mathbf{r}}$. If two candidates evaluate to the same heuristic value, we prefer the one that (a) covers more true positives, or (b) contains fewer conditions in its body. Whenever a new rule is added, the overall coverage of the theory increases, as more positive training instances are covered. However, the rule may also cover some of the *negative* instances $N_i = T \setminus P_i$. As the rule’s prediction is incorrect in such cases, the consistency of the theory may decrease.

3.3 Threshold Selection

As described in Sect. 3.2, we use a SeCo strategy to select more rules until all positive training instances are covered for each label. In this way, the coverage

of the resulting theory is maximized at the expense of consistency, because each rule contributes to the overall coverage, but might introduce wrong predictions for some instances. To trade off between these aspects, we allow to (optionally) specify a threshold ϕ that aims at diminishing the effects of inconsistent rules. It is compared to a heuristic value that is calculated for each rule according to the heuristic δ . For calculating the heuristic value, the rule’s predictions on the entire training data set are taken into account. This is different from the candidate selection discussed in Sect. 3.2, where instances that are already covered by previously selected rules are not considered. Because the candidate selection aims at selecting non-redundant rules, that cover the positive training instances as uniformly as possible, it considers rules in the context of their predecessors. In contrast, the threshold ϕ is applied at prediction time when no order is imposed on the rules, i.e., all rules whose heuristic value exceeds the threshold equally contribute to the prediction.

4 Evaluation

In this section, we present an empirical study that emphasises the need to use varying heuristics for candidate selection and filtering to learn theories that are tailored to specific multi-label measures. We further compare our method to different baselines to demonstrate the benefits of being able to flexibly adjust a learner to different measures, rather than employing a general-purpose learner.

4.1 Experimental Setup

We applied our method to eight different data sets taken from the Mulan project.³ We set the minimum number of rules to be generated to 300.000 (cf. Algorithm 1, parameter γ). For candidate selection according to Sect. 3.2, we used the m-estimate (cf. Eq. 6) with $m = 0, 2^1, 2^2, \dots, 2^{19}$. For each of these variants, we applied varying thresholds ϕ according to Sect. 3.3. The thresholds have been chosen such that they are satisfied by at least 100%, 95%, \dots , 5% of the selected rules. All results have been obtained using 10-fold cross validation.

In addition to the m-estimate, we also used the F-measure (cf. Eq. 5) with varying β -parameters. As the conclusions drawn from these experiments are very similar to those for the m-estimate, we focus on the latter at this point.

Among the performance measures that we report are micro-averaged precision and recall. Given a global confusion matrix $C := C_1^1 \oplus \dots \oplus C_i^j \oplus \dots \oplus C_n^m$ that consists of the TP , FP , TN , and FN aggregated over all test instances \mathbf{x}_j and labels λ_i , these two measures are calculated as defined in Eqs. 3 and 4. Moreover, we report the micro-averaged F1 score (cf. Eq. 5 with $\beta = 1$) as well as Hamming and subset accuracy. Hamming accuracy calculates as

$$\delta_{Hamming}(C) := \frac{TP + TN}{TP + FP + TN + FN} \quad (7)$$

³ Data sets and detailed statistics available at <http://mulan.sourceforge.net/datasets-mlc.html>.

whereas subset accuracy differs from the other measures, because it is computed instance-wise. Given true label vectors $Y = (\mathbf{y}_1, \dots, \mathbf{y}_m)$ and predicted label vectors $\hat{Y} = (\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_m)$, it measures the fraction of perfectly labeled instances:

$$\delta_{acc}(Y, \hat{Y}) := \frac{1}{m} \sum_j \llbracket \mathbf{y}_j = \hat{\mathbf{y}}_j \rrbracket \quad (8)$$

4.2 Analysis of Different Parameter Settings

For a broad analysis, we trained $20^2 = 400$ theories per data set using the same candidate rules, but selecting and filtering them differently by using varying combinations of the parameters m and ϕ as discussed in Sect. 4.1. We visualize the performance and characteristics of the resulting models as two-dimensional matrices of scores (cf. e.g. Fig. 1). One dimension corresponds to the used m -parameter, the other refers to the threshold ϕ , respectively.

Some of the used data sets (CAL500, FLAGS, and YEAST) contain very frequent labels for which the minority class $t_i = 0$. This is rather atypical in MLC and causes the unintuitive effect that the removal of individual rules results in a theory with greater recall and/or lower precision. To be able to compare different parameter settings across multiple data sets, we worked around this effect by altering affected data sets., i.e., inverting all labels for which $t_i = 0$.

Predictive Performance. In Figs. 1 and 2 the average ranks of the tested configurations according to different performance measures are depicted. The rank of each of the 400 parameter settings was determined for each data set separately and then averaged over all data sets. The depicted standard deviations show that the optimal parameter settings for a respective measure may vary depending on the data set. However, for each measure there is an area in the parameter space where a good setting can be found with high certainty.

As it can clearly be seen, precision and recall are competing measures. The first is maximized by choosing small values for m and filtering extensively, the latter benefits from large values for m and no filtering. Interestingly, setting $m = 0$, i.e., selecting candidates according to the precision metric, does not result in models with the highest overall precision. This is in accordance with Fig. 3, where the models with the highest F1 score do not result from using the F1-measure for candidate selection. Instead, optimizing the F1 score requires to choose small values for m to trade off between consistency and coverage. The same applies to Hamming and subset accuracy, albeit both of these measure demand to put even more weight on consistency and filtering more extensively compared to F1.

Model Characteristics. Besides the predictive performance, we are also interested in the characteristics of the theories. Figure 4 shows how the number of rules in a theory as well as the average number of conditions are affected by varying parameter settings. The number of rules independently declines when




Fig. 1. Ranks and standard deviation of average ranks over all data sets according to Hamming and subset accuracy using different parameters m (horizontal axis) and ϕ (vertical axis). Best parameters for different data sets specified by red + signs. (Color figure online)

using greater values for the parameter m and/or smaller values for ϕ , resulting in less complex theories that can be comprehended by humans more easily. The average number of conditions is mostly affected by the parameter m .

Figure 5 provides an example of how different parameters affect the model characteristics. It shows the rules for predicting the same label as induced by two fundamentally different approaches. The first approach ($m = 16, \phi = 0.3$) reaches high scores according to the F1-measure, Hamming accuracy, and subset accuracy, whereas the second one ($m = 262144, \phi = 1.0$) results in high recall.

4.3 Baseline Comparison

Although the goal of this work is not to develop a method that generally outperforms existing rule learners, we want to ensure that we achieve competitive results. For this reason, we compared our method to JRip, Weka’s re-




Fig. 2. Ranks and standard deviation of average ranks over all data sets according to micro-averaged precision, recall, and F1-measure. Best parameters for different data sets specified by red + signs. (Color figure online)




Fig. 3. Ranks and standard deviation of average ranks over all data sets according to micro-averaged F1-measure, when using the F-measure with varying β -parameters (horizontal axis) instead of the m-estimate for candidate selection. Best parameters for different data sets specified by red + signs. (Color figure online)




Fig. 4. Ranks and standard deviation of average ranks over all data sets regarding the number of rules and conditions. A smaller rank means more rules or conditions.

$m = 16, \phi = 0.3$	Mi. Precision = 74.07%, Mi. Recall = 78.26%
$Cough \leftarrow \overline{\text{“cough”}} \wedge \overline{\text{“aldrich”}} \wedge \overline{\text{“opacity”}} \wedge \overline{\text{“tachypnea”}} \wedge \overline{\text{“streaky”}} \wedge \overline{\text{“side”}} \wedge \overline{\text{“distal”}} \wedge \overline{\text{“diaphragm”}}$	
$Cough \leftarrow \overline{\text{“cough”}} \wedge \overline{\text{“x-rays”}} \wedge \overline{\text{“vomiting”}} \wedge \overline{\text{“proximity”}} \wedge \overline{\text{“hematuria”}} \wedge \overline{\text{“focal”}}$	
$Cough \leftarrow \overline{\text{“cough”}} \wedge \overline{\text{“group”}} \wedge \overline{\text{“edema”}} \wedge \overline{\text{“fever”}}$	
$Cough \leftarrow \overline{\text{“cough”}} \wedge \overline{\text{“lobe”}} \wedge \overline{\text{“breathing”}}$	
$Cough \leftarrow \overline{\text{“coughing”}}$	
$m = 262144, \phi = 1.0$	Mi. Precision = 65.61%, Mi. Recall = 89.57%
$Cough \leftarrow \overline{\text{“cough”}} \wedge \overline{\text{“ureteral”}} \wedge \overline{\text{“stones”}} \wedge \overline{\text{“contrast”}}$	
$Cough \leftarrow \overline{\text{“coughing”}}$	
$Cough \leftarrow \overline{\text{“code”}}$	
$Cough \leftarrow \overline{\text{“substance”}}$	

Fig. 5. Exemplary rule sets predicting the label *786.2:Cough* of the data set MEDICAL, which contains textual radiology reports that were categorized into diseases.

Table 1. Predictive performance of Ripper using IREP and post-processing (R_3), without using post-processing (R_2), and using neither IREP nor post-processing (R_1) compared to approaches trying to optimize micro-averaged F1 (M_F), Hamming accuracy (M_H), and subset accuracy (M_S).

	F1				Hamming acc.				Subset acc.			
	R_1	R_2	R_3	M_F	R_1	R_2	R_3	M_H	R_1	R_2	R_3	M_S
BIRDS	43.65	41.12	46.01	45.33	94.39	94.48	95.17	95.10	44.20	45.57	51.48	48.85
CAL500	33.63	33.18	33.76	40.10	82.14	83.66	85.39	86.02	0.00	0.00	0.00	0.00
EMOTIONS	56.96	58.68	60.97	65.20	75.12	75.38	77.21	77.65	18.04	20.40	23.60	22.42
ENRON	50.57	53.05	55.33	51.07	94.35	94.70	94.93	94.54	6.17	7.99	9.16	7.81
FLAGS	71.81	72.96	74.85	72.83	73.02	74.08	75.20	73.39	15.47	17.05	21.00	9.82
GENBASE	98.83	98.68	98.68	99.14	99.89	99.88	99.88	99.92	97.28	96.83	96.83	97.89
MEDICAL	81.40	83.67	84.81	81.67	99.01	99.10	99.15	98.98	66.74	69.91	72.16	66.43
SCENE	63.97	63.25	64.55	67.44	87.87	87.25	88.03	88.93	46.61	44.54	46.24	49.73
YEAST	58.65	60.41	61.19	64.25	78.50	78.29	78.77	79.24	8.73	7.86	9.18	11.75
Avg. rank	3.44	3.00	1.67	1.78	3.44	2.89	1.67	1.89	2.89	2.67	1.56	2.11

implementation of Ripper [7], using the binary relevance method. By default, Ripper uses *incremental reduced error pruning* (IREP) and post-processes the induced rule set. Although our approach could make use of such optimizations, this is out of the scope of this work. For a fair comparison, we also report the results of JRip without using IREP ($P = false$) and/or with post-processing turned off ($O = 0$).

Note that we do not consider the random forests from which we generate rules (cf. Sect. 3.1) to be relevant baselines. This is, because random forests use voting for making a prediction, which is fundamentally different than rule learners that model a DNF. Also, we train random forests consisting of a very large number

of trees with varying depths to generate diverse rules. In our experience, these random forests perform badly compared to commonly used configurations.

We tested three different configurations of our approach. The parameters m and ϕ used by these approaches have been determined on a validation set by using nested 5-fold cross validation on the training data. For the approach M_F , the parameters have been chosen such that the F1-measure is maximized. M_H and M_S were tuned with respect to Hamming and subset accuracy, respectively.

According to Table 1, our method is able to achieve reasonable predictive performances. With respect to the measure they try to optimize, our approaches generally rank before JRip with optimizations turned off (R_1), which is the competitor that is conceptually closest to our method. Although IREP definitely has a positive effect on the predictive performance, our approaches also tend to outperform JRip with IREP enabled, but without using post-processing (R_2). Despite the absence of advanced pruning and post-processing techniques, our approaches are even able to surpass the fully fledged variant of JRip (R_1) on some data sets. We consider these results as a clear indication that it is indispensable to be able to flexibly adapt the heuristic used by a rule learner if one aims at deliberately optimizing a specific multi-label performance measure.

5 Related Work

Several rule-based approaches to multi-label classification have been proposed in the literature. On the one hand, there are methods based on descriptive rule learning, such as association rule discovery [17, 18, 22, 23], genetic algorithms [1, 6], or evolutionary classification systems [2, 3]. On the other hand, there are algorithms that adopt the separate-and-conquer strategy used by many traditional rule learners for binary or multi-class classification, e.g. by Ripper [7], and transfer it to MLC [19, 21]. Whereas in descriptive rule learning one does usually not aim at discovering rules that minimize a certain (multi-label) loss, the latter approaches employ a heuristic-guided search for rules that optimize a given rule learning heuristic and hence could benefit from the results of this work.

Similar to our experiments, empirical studies aimed at discovering optimal rule learning heuristics have been published in the realm of single-label classification [14, 15]. Moreover, to investigate the properties of bipartition evaluation functions, ROC space isometrics have been proven to be a helpful tool [9, 10]. They have successfully been used to study the effects of using different heuristics in separate-and-conquer algorithms [12], or for ranking and filtering rules [11].

6 Conclusions

In this work, we presented a first empirically study that thoroughly investigates the effects of using different rule learning heuristics for candidate selection and filtering in the context of multi-label classification. As commonly used multi-label measures, such as micro-averaged F1, Hamming accuracy, or subset accuracy,

require to put more weight on the consistency of rules rather than on their coverage, models that perform well with respect to these measures are usually small and tend to contain specific rules. This is beneficial in terms of interpretability as less complex models are assumed to be easier to understand by humans.

As our main contribution, we emphasise the need to flexibly trade off the consistency and coverage of rules, e.g., by using parameterized heuristics like the m-estimate, depending on the multi-label measure that should be optimized by the model. Our study revealed that the choice of the heuristic is not straightforward, because selecting rules that minimize a certain loss functions locally does not necessarily result in that loss being optimized globally. E.g., selecting rules according to the F1-measure does not result in the overall F1 score to be maximized. For optimal results, the trade-off between consistency and coverage should be fine-tuned depending on the data set at hand. However, our results indicate that, even across different domains, the optimal settings for maximizing a measure can often be found in the same region of the parameter space.

In this work, we restricted our study to DNFs, i.e., models that consist of non-conflicting rules predicting the same outcome for a label. This restriction simplifies the implementation and comprehensibility of the learner, as no conflicts may arise at prediction time. However, we expect that including both, rules that model the presence as well as the absence of labels, could be beneficial in terms of robustness and could have similar, positive effects on the consistency of the models as the threshold selection used in this work. Furthermore, we leave the empirical analysis of macro-averaged performance measures for future work.

Acknowledgments. This research was supported by the German Research Foundation (DFG) (grant number FU 580/11).

References

1. Allamanis, M., Tzima, F.A., Mitkas, P.A.: Effective rule-based multi-label classification with learning classifier systems. In: Tomassini, M., Antonioni, A., Daolio, F., Buesser, P. (eds.) ICANNGA 2013. LNCS, vol. 7824, pp. 466–476. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-37213-1_48
2. Arunadevi, J., Rajamani, V.: An evolutionary multi label classification using associative rule mining for spatial preferences. In: IJCA Special Issue on Artificial Intelligence Techniques–Novel Approaches and Practical Applications (2011)
3. Ávila-Jiménez, J.L., Gibaja, E., Ventura, S.: Evolving multi-label classification rules with gene expression programming: a preliminary study. In: Corchado, E., Graña Romay, M., Manhaes Savio, A. (eds.) HAIS 2010. LNCS (LNAI), vol. 6077, pp. 9–16. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13803-4_2
4. Boutell, M.R., Luo, J., Shen, X., Brown, C.M.: Learning multi-label scene classification. *Pattern Recogn.* **37**(9), 1757–1771 (2004)
5. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001)

6. Cano, A., Zafra, A., Gibaja, E.L., Ventura, S.: A grammar-guided genetic programming algorithm for multi-label classification. In: Krawiec, K., Moraglio, A., Hu, T., Etaner-Uyar, A.Ş., Hu, B. (eds.) EuroGP 2013. LNCS, vol. 7831, pp. 217–228. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-37207-0_19
7. Cohen, W.W.: Fast effective rule induction. In: Proceedings of the 12th International Conference on International Conference on Machine Learning (1995)
8. Diplaris, S., Tsoumakas, G., Mitkas, P.A., Vlahavas, I.: Protein classification with multiple algorithms. In: Bozanis, P., Houstis, E.N. (eds.) PCI 2005. LNCS, vol. 3746, pp. 448–456. Springer, Heidelberg (2005). https://doi.org/10.1007/11573036_42
9. Flach, P.A.: The geometry of ROC space: understanding machine learning metrics through ROC isometrics. In: Proceedings of the 20th International Conference on Machine Learning (2003)
10. Fürnkranz, J., Flach, P.A.: An analysis of rule evaluation metrics. In: Proceedings of the 20th International Conference on Machine Learning (2003)
11. Fürnkranz, J., Flach, P.: An analysis of stopping and filtering criteria for rule learning. In: Boulicaut, J.-F., Esposito, F., Giannotti, F., Pedreschi, D. (eds.) ECML 2004. LNCS (LNAI), vol. 3201, pp. 123–133. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30115-8_14
12. Fürnkranz, J., Flach, P.A.: ROC 'n' rule learning-towards a better understanding of covering algorithms. *Mach. Learn.* **58**(1), 39–77 (2005)
13. Fürnkranz, J., Gamberger, D., Lavrač, N.: Foundations of Rule Learning. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-3-540-75197-7>
14. Janssen, F., Fürnkranz, J.: An empirical investigation of the trade-off between consistency and coverage in rule learning heuristics. In: Jean-Fran, J.-F., Berthold, M.R., Horváth, T. (eds.) DS 2008. LNCS (LNAI), vol. 5255, pp. 40–51. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88411-8_7
15. Janssen, F., Fürnkranz, J.: On the quest for optimal rule learning heuristics. *Mach. Learn.* **78**(3), 343–379 (2010)
16. Klimt, B., Yang, Y.: The enron corpus: a new dataset for email classification research. In: Boulicaut, J.-F., Esposito, F., Giannotti, F., Pedreschi, D. (eds.) ECML 2004. LNCS (LNAI), vol. 3201, pp. 217–226. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30115-8_22
17. Lakkaraju, H., Bach, S.H., Leskovec, J.: Interpretable decision sets: a joint framework for description and prediction. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2016)
18. Li, B., Li, H., Wu, M., Li, P.: Multi-label classification based on association rules with application to scene classification. In: The 9th International Conference for Young Computer Scientists (2008)
19. Mencía, E.L., Janssen, F.: Learning rules for multi-label classification: a stacking and a separate-and-conquer approach. *Mach. Learn.* **105**(1), 77–216 (2016)
20. Pestian, J.P., et al.: A shared task involving multi-label classification of clinical free text. In: Proceedings of the Workshop on BioNLP 2007: Biological, Translational, and Clinical Language Processing (2007)
21. Rapp, M., Loza Mencía, E., Fürnkranz, J.: Exploiting anti-monotonicity of multi-label evaluation measures for inducing multi-label rules. In: Phung, D., Tseng, V.S., Webb, G.I., Ho, B., Ganji, M., Rashidi, L. (eds.) PAKDD 2018. LNCS (LNAI), vol. 10937, pp. 29–42. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93034-3_3

22. Thabtah, F.A., Cowling, P., Peng, Y.: MMAC: a new multi-class, multi-label associative classification approach. In: 4th IEEE International Conference on Data Mining (2004)
23. Thabtah, F.A., Cowling, P., Peng, Y.: Multiple labels associative classification. *Knowl. Inf. Syst.* **9**(1), 109–129 (2006)
24. Trohidis, K., Tsoumakas, G., Kalliris, G., Vlahavas, I.P.: Multi-label classification of music into emotions. In: International Society for Music Information Retrieval (2008)
25. Tsoumakas, G., Katakis, I., Vlahavas, I.: Mining multi-label data. In: Maimon, O., Rokach, L. (eds.) *Data Mining and Knowledge Discovery Handbook*. Springer, Boston (2009). https://doi.org/10.1007/978-0-387-09823-4_34
26. Turnbull, D., Barrington, L., Torres, D., Lanckriet, G.: Semantic annotation and retrieval of music and sound effects. *IEEE Trans. Audio Speech Lang. Process.* **16**(2), 467–476 (2008)