



Scalable Rule Lists Learning with Sampling

Leonardo Pellegrina

Dept. of Information Engineering, University of Padova
Padova, Italy

leonardo.pellegrina@unipd.it

Fabio Vandin

Dept. of Information Engineering, University of Padova
Padova, Italy

fabio.vandin@unipd.it

Abstract

Learning interpretable models has become a major focus of machine learning research, given the increasing prominence of machine learning in socially important decision-making. Among interpretable models, rule lists are among the best-known and easily interpretable ones. However, finding optimal rule lists is computationally challenging, and current approaches are impractical for large datasets.

We present a novel and scalable approach to learn nearly optimal rule lists from large datasets. Our algorithm uses sampling to efficiently obtain an approximation of the optimal rule list with rigorous guarantees on the quality of the approximation. In particular, our algorithm guarantees to find a rule list with accuracy very close to the optimal rule list when a rule list with high accuracy exists. Our algorithm builds on the VC-dimension of rule lists, for which we prove novel upper and lower bounds. Our experimental evaluation on large datasets shows that our algorithm identifies nearly optimal rule lists with a speed-up up to two orders of magnitude over state-of-the-art exact approaches. Moreover, our algorithm is as fast as, and sometimes faster than, recent heuristic approaches, while reporting higher quality rule lists. In addition, the rules reported by our algorithm are more similar to the rules in the optimal rule list than the rules from heuristic approaches.

CCS Concepts

• **Mathematics of computing** → **Probabilistic algorithms**; • **Information systems** → **Data mining**; • **Theory of computation** → **Sketching and sampling**; **Sample complexity and generalization bounds**.

Keywords

Rule list learning, Interpretable Machine Learning, VC-dimension, Sample complexity bounds

ACM Reference Format:

Leonardo Pellegrina and Fabio Vandin. 2024. Scalable Rule Lists Learning with Sampling. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3671989>



This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '24, August 25–29, 2024, Barcelona, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0490-1/24/08
<https://doi.org/10.1145/3637528.3671989>

1 Introduction

Interpretability is one of the characteristics of machine learning models that has become a major topic of research due to the ever increasing impact of machine learning models in socially important decision-making [22]. The goal is to build models that are easily understood by humans, while achieving high predictive power, in contrast to black-box models (e.g., large deep-learning models) that are highly predictive but are not transparent nor interpretable.

Rule lists [21], and more in general rule-based models such as decision trees, are among the best-known and easily interpretable models. A rule list is a sequence of rules, and the prediction from such a model is obtained by applying the first rule in the list whose condition is satisfied for the given input.

Rule-based models often display predictive power comparable to black-box models in several critical applications [22]. While advanced techniques have been recently developed to speed-up the discovery of optimal rule lists [1, 14, 23, 32], it remains a computationally challenging task, in particular for large datasets where differentiable models (such as neural networks) can be easily learned with gradient descent and its variants.

Fast heuristic methods for rule lists and other rule-based models have been developed [7, 30], but they do not provide guarantees in terms of the accuracy of the rule list they report. In several applications identifying an optimal, or close to optimal rule, is crucial, since the goal is to find models that are not only interpretable, but also have high predictive power and can therefore be used as interpretable alternatives to black-box models.

A natural solution to speed up the training on a large dataset, and to obtain an approximation of the best rule list, is to reduce the size of the training set, e.g., by only considering a small random sample of the dataset. On the other hand, it is clear that there is an intrinsic trade-off between the accuracy of this approximation and the size of the random sample, i.e., the computational cost incurred by the learning algorithm. This accuracy depends directly on the fluctuations of the losses estimated on the sample w.r.t. the losses measured on the entire dataset: a sample too small can be analyzed quickly but may provide inaccurate estimates. Moreover, it is possible that the best rule list, learned from a random sample, may be qualitatively different, thus not representative, of the optimal solution computed on the whole dataset.

Contributions. To address these challenges, we precisely quantify the maximum deviation between the accuracies of rule lists that are estimated on the random sample w.r.t. the respective exact counterparts. As a consequence, we obtain sufficient sample sizes that guarantee that an accurate rule list can be found from a small random sample. We then propose SAMRULE (*Sampling for Rules list Learning*), a scalable approach to approximate optimal rule lists from large datasets with rigorous guarantees on the quality of the approximation.

- SAMRULE is the first scalable approach to approximate optimal rule lists from large datasets while providing rigorous guarantees on the accuracy of the reported rule list. In particular, SAMRULE guarantees to find a rule list with accuracy very close to the optimal one when a rule list with high accuracy exists. As a result, SAMRULE makes it possible to assess the existence of high accuracy rule lists for large datasets.
- SAMRULE uses sampling to efficiently obtain an approximation of the optimal rule list with guarantees. We derive novel bounds relating the accuracy of a rule lists in the sample with its accuracy in the whole dataset *and* with the accuracy of the optimal rule list. Our bounds build on the VC-dimension of rule lists, for which we prove novel upper and lower bounds, and that allow us to define sample sizes leading to the desired quality guarantees.
- Our experimental evaluation shows that SAMRULE enables the identification of nearly optimal rule lists from large datasets, with a speed-up of up to two orders of magnitude over state-of-the-art exact approaches. Moreover, SAMRULE is as fast as, or faster than, recent heuristic approaches, while reporting higher quality rule lists, for which it can also certify their quality with respect to the optimal rule list. The rules reported by SAMRULE are also more similar to the rules in the optimal rule list than the rules from heuristic approaches.

2 Related Works

We now discuss the works most related to ours. We focus on methods to learn rule lists, in particular *sparse* rule lists. Sparsity is a crucial requirement for the interpretability of models. For an overview on the topic of interpretability in machine learning, we point the interested reader to the survey from Rudin et al. [22].

Finding optimal sparse rule lists is NP-hard, and several methods have been proposed to learn sparse rule lists. Such methods fall under two categories: exact methods, that guarantee to find the optimal rule list (i.e., with highest accuracy), and heuristic methods, that are faster but provide no guarantees on the accuracy of the reported rule list w.r.t. the optimal.

In the category of exact methods, Angelino et al. [1] propose CORELS, a branch-and-bound algorithm that achieves several orders of magnitude speedup in time and a significant reduction of memory consumption, by leveraging algorithmic bounds, efficient data structures, and computational reuse. Rudin and Ertekin [23] develop a mathematical programming approach to building rule lists. Okajima and Sadamasa [14] define a continuous relaxed version of a rule list, and propose an algorithm that optimizes rule lists based on such continuous relaxed versions. Yu et al. [32] propose a SAT-based approach to find optimal rule sets and lists. While highly optimized, exact methods such as the aforementioned ones only apply to problems of moderate size, and do not scale to large datasets.

In the category of heuristic methods, Rivest [21] proposed greedy splitting techniques, also used in subsequent heuristic algorithms for learning decision trees [5, 19]. RIPPER [7] builds rule sets in a greedy fashion, and a similar greedy strategy has been used for finding sets of robust rules in terms of minimum description length [9].

Other works considered different problems and classification models, such as probabilistic rules for multiclass classification [18], and, more recently, rule lists with preferred variables [15], sets of locally optimal rules [11], and multi-label rule sets [6]. Yang et al. [30] uses, instead, a Bayesian approach with Monte-Carlo search to explore the rule list solution space. All such approaches focus on improving efficiency with respect to exact approaches, but do not provide guarantees on the quality of the reported rule lists with respect to the optimal one. In several challenging instances, the gap with the optimal solution may be substantial, resulting in a predictive model with suboptimal performance.

A different line of research, related to the identification of almost-optimal models, is provided by the exploration of the *Rashomon set*, that is the set of almost-optimal models for a machine learning problem [4, 10, 26]. Works in this area have focused on the exploration or full enumeration of the Rashomon set, including the case of rule-based models [29], for example to study variable importance in well-performing models [8]. While some of our theoretical results may be useful for such applications as well, our focus is not the exploration or enumeration of *all* almost optimal models, but rather on efficiently finding one nearly optimal model while providing guarantees on its quality.

Our algorithm builds on novel upper bounds to the VC-dimension [28] of rule lists, a fundamental combinatorial measure of the complexity of classification models. We combine these bounds with sharp concentration inequalities to bound the maximum deviation between the prediction accuracy of rule lists from a random sample w.r.t. the entire dataset. Our VC-dimension upper bounds generalize previous results that can be derived from rule-based classification models [24], as they provide a more granular dependence on the rule list parameters compared to known results (see Section 4.1 for a detailed discussion). We establish the tightness of our analysis by proving almost matching lower bounds, that significantly improves over the best known results from Yildiz [31]. The VC-dimension has been used to develop generalization bounds and sampling algorithms for other data mining tasks, including association rules mining [20] and sequential pattern mining [25]. To the best of our knowledge, our work is the first to use VC-dimension in designing efficient sampling approaches for learning rule lists.

3 Preliminaries

We define a *dataset* \mathcal{D} as a collection $\mathcal{D} = \{(s_1, t_1), \dots, (s_n, t_n)\}$ of n training instances, where each training instance is a pair $(s, t) \in \mathcal{D}$ composed by a set s of d features, and a categorical label t . We denote with x_i the i -th feature, for $1 \leq i \leq d$. For simplicity, in this work we will focus on datasets with binary features, $x_i \in \{0, 1\}$ for $1 \leq i \leq d$, and binary labels $t \in \{0, 1\}$, but most of our results extend to the general case. We use \mathcal{D}_s to denote the collection of the d features of each training instance, ignoring the labels, that is $\mathcal{D}_s = \{s_1, s_2, \dots, s_n\}$.

We define a rule list $R = [r_1, r_2, \dots, r_k, r_\top]$ of length $|R| = k$ as a sequence of $k \geq 0$ rules $r_i, 1 \leq i \leq k$, plus a *default rule* r_\top . Each rule $r \in R$ is defined as a pair $r = (c, p)$ where c is a condition on a subset of the d features of the dataset, and $p \in \{0, 1\}$ is a prediction of the label. Similarly to previous rule-based models [1, 23, 30], we only consider conditions composed by conjunctions of monotone

variables of the type $c = \bigwedge_i "x_i = 1"$, i.e., we do not consider the negation of binary features (e.g., " $x_i = 0$ "). Note that negations can be included easily by considering additional binary features. The default rule $r_\top = (c_\top, p)$ is composed by the condition c_\top that is always true. Given an instance $(s, t) \in \mathcal{D}$, the predicted label $P(R, s)$ by the rule list R on s is computed as the predicted label p of the first rule r whose condition c is satisfied by s ; Figure 1 shows an example of a dataset with $n = 5$ training instances over $d = 4$ features, and a rule list R of length $k = 3$ with conjunctions with at most $z = 2$ terms.

\mathcal{D}	x_1	x_2	x_3	x_4	t	
s_1	0	1	0	0	1	
s_2	1	1	0	0	1	if $x_2 = 1 \rightarrow 1$
s_3	0	0	1	1	1	else if $x_1 = 1 \wedge x_3 = 1 \rightarrow 0$
s_4	0	0	0	0	0	else if $x_4 = 1 \rightarrow 1$
s_5	1	0	1	1	0	else $\rightarrow 0$

(a)

(b)

Figure 1: (a): example of a dataset \mathcal{D} with $n = 5$ instances and $d = 4$ features. (b): example of a rule list R with length $k = 3$ with conjunctions with at most $z = 2$ terms. The rule list R perfectly classifies the instances of \mathcal{D} .

The *regularized binary loss* $\ell(R, \mathcal{D})$ of a rule list R is defined, for some constant $\alpha \geq 0$, as

$$\ell(R, \mathcal{D}) = \frac{1}{n} \sum_{(s,t) \in \mathcal{D}} \mathbb{1}[P(R, s) \neq t] + \alpha|R|.$$

The use of the *regularized* loss is common practice in rule list learning [1], allowing to prefer shorter (i.e., simpler) rule lists over longer ones when their accuracy is similar, as quantified by the parameter α .

Define \mathcal{R}_k^z as the set of rule lists with length $\leq k$, where each rule in R is given by the conjunction of at most z terms. The rule list learning problem is to identify a rule R^* such that

$$R^* = \arg \min_{R \in \mathcal{R}_k^z} \{\ell(R, \mathcal{D})\}.$$

The values of k and z are usually fairly small, since large values would compromise the interpretability of the rule list [22]. Finding R^* is NP-hard, and to develop efficient algorithms one has to resort to approximations.

Our goal is to compute, for given accuracy parameters $\varepsilon, \theta \in (0, 1]$, a rule list \tilde{R} that provides an (ε, θ) -approximation of the optimal rule list R^* , defined as follows.

Definition 3.1. A rule list \tilde{R} provides an (ε, θ) -approximation of the optimal rule list R^* if it holds

$$\ell(\tilde{R}, \mathcal{D}) \leq \ell(R^*, \mathcal{D}) + \varepsilon \max\{\ell(R^*, \mathcal{D}), \theta\}.$$

Our definition of (ε, θ) -approximation is motivated by the fact that the optimal loss $\ell(R^*, \mathcal{D})$ is unknown a priori, and often extremely expensive to compute or even estimate from large datasets. Therefore, we need to design approximation guarantees that are sharp in all situations, i.e., for the all possible values of the optimal

loss. The (ε, θ) -approximation of Definition 3.1 allows to interpolate between an additive approximation, with small absolute error $\varepsilon\theta$, when the optimal loss $\ell(R^*, \mathcal{D})$ is small (i.e., $\leq \theta$), and a relative approximation (with parameter ε) when the optimal loss $\ell(R^*, \mathcal{D})$ is large (i.e., $> \theta$). This flexible design avoids statistical bottlenecks that are incurred when the loss of the best model is large, a well known issue in statistical learning theory [3]. Taking this effect into account allows SAMRULE to be extremely efficient in terms of sample sizes.

As we will formalize in Section 4, the main goal of SAMRULE is to compute an (ε, θ) -approximation of the optimal rule list from a randomly drawn samples with high probability.

4 SAMRULE Algorithm

As introduced in previous sections, the main idea of our approach is to compute an approximation of the best rule list performing the analysis of a small random sample, instead of processing a large dataset, which may be extremely expensive and unfeasible in practice. In this section we formally define our algorithm.

Define a sample \mathcal{S} as a collection of m instances of the dataset \mathcal{D} taken uniformly and independently at random. The regularized loss of a rule list R computed on the sample \mathcal{S} is defined as

$$\ell(R, \mathcal{S}) = \frac{1}{m} \sum_{(s,t) \in \mathcal{S}} \mathbb{1}[P(R, s) \neq t] + \alpha|R|.$$

Intuitively, $\ell(R, \mathcal{S})$ provides an *estimate* of $\ell(R, \mathcal{D})$. More precisely, the loss of any rule R computed on \mathcal{S} is an *unbiased* estimator of the loss computed on \mathcal{D} ; in fact, it holds

$$\mathbb{E}_{\mathcal{S}}[\ell(R, \mathcal{S})] = \ell(R, \mathcal{D}).$$

We remark, however, that the convergence of $\ell(R, \mathcal{S})$ towards $\ell(R, \mathcal{D})$ in *expectation* is not sufficient to derive rigorous approximations for the best rule list; instead, it is necessary to take into account the variance of the deviations, that indirectly depends the *complexity* of the class of rule list models. Therefore, our approach is based on properly bounding the *deviations* $|\ell(R, \mathcal{S}) - \ell(R, \mathcal{D})|$ with *high probability*, in order to guarantee that the best rule identified from the sample \mathcal{S} is a high-quality approximation of the optimal solution from \mathcal{D} . To obtain a scalable method, we wish to identify the smallest possible sample size that yields the above-mentioned guarantees. We address this challenging problem combining advanced VC-dimension bounds and sharp concentration inequalities.

We now describe our algorithm SAMRULE. The input to SAMRULE is: a dataset \mathcal{D} of m training instances over d features; the values k and z , representing the maximum number of rules in a rule list and the maximum number of conditions for each rule; constants $\varepsilon, \theta \in (0, 1]$, that are the values defining the quality of the (ε, θ) -approximation (see Section 3); and a constant $\delta \in (0, 1]$, that defines the required *confidence*, that is the probability that the output by SAMRULE is a (ε, θ) -approximation. (Note that the set \mathcal{R}_k^z of rules lists that can be produced in output by SAMRULE is defined by \mathcal{D} , k , and z .) Given the input, SAMRULE computes the quantity $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$, defined as the minimum value of $m \geq 1$ such that it holds

$$\sqrt{\frac{3\theta \ln(\frac{2}{\delta})}{m}} + \sqrt{\frac{2\left(\theta + \sqrt{\frac{3\theta \ln(\frac{2}{\delta})}{m}}\right)\left(\omega + \ln(\frac{2}{\delta})\right)}{m}} + \frac{2\left(\omega + \ln(\frac{2}{\delta})\right)}{m} \leq \varepsilon\theta,$$

where $\omega = kz \ln(2ed/z) + 2$. Note that $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ can be easily computed (e.g., with a binary search over m). We also show analytical bounds to $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ in our analysis (Section 4.2).

SAMRULE then draws a sample \mathcal{S} of $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ instances taken uniformly at random from \mathcal{D} , and outputs $\hat{R} = \arg \min_{R \in \mathcal{R}_k^z} \ell(R, \mathcal{S})$ by finding the optimal rule list on the sample \mathcal{S} . Note that SAMRULE can leverage any exact algorithm, such as CORELS [1], to find \hat{R} from \mathcal{S} .

In the following sections, we prove that SAMRULE provides an (ϵ, θ) -approximation of the optimal rule list $R^* = \arg \min_{R \in \mathcal{R}_k^z} \ell(R, \mathcal{S})$ with probability $\geq 1 - \delta$. Our proof leverage novel bounds on the VC-dimension of rule lists, that we present in Section 4.1, while the approximation guarantees from SAMRULE are proved in Section 4.2.

4.1 Complexity of Rule Lists

In this section we provide analytical results on the complexity of rule lists, by studying their VC-dimension [28]. Intuitively, the VC-dimension measures the capacity of a class of prediction models of predicting arbitrary labels assignments with perfect accuracy. As discussed in previous sections, these bounds will be instrumental to derive accurate approximation bounds and prove that the reported rule by SAMRULE is an (ϵ, θ) -approximation.

We start by providing the main concepts related to VC-dimension (see [12, 13, 27] for a more in depth introduction).

Given a dataset $\mathcal{D} = \{(s_1, t_1), \dots, (s_n, t_n)\}$, remember that $\mathcal{D}_s = \{s_1, \dots, s_n\}$. Given the set \mathcal{R}_k^z of rule lists with length $\leq k$, where each condition contains conjunctions with at most z terms, for each rule list $R \in \mathcal{R}_k^z$ we define the *projection* of R on the dataset \mathcal{D} as $X(R, \mathcal{D}) = \{s \in \mathcal{D}_s : P(R, s) = 1\}$. We define the *range space* associated to the dataset \mathcal{D} and to \mathcal{R}_k^z as $(\mathcal{D}, \tilde{\mathcal{R}}_k^z)$, where $\tilde{\mathcal{R}}_k^z = \{X(R, \mathcal{D}) : R \in \mathcal{R}_k^z\}$ is the *range set*. Note that in general $|\tilde{\mathcal{R}}_k^z|$ can be smaller than $|\mathcal{R}_k^z|$, since two rule lists R_1 and R_2 may provide the same predictions for all instances in a dataset \mathcal{D} , which implies $X(R_1, \mathcal{D}) = X(R_2, \mathcal{D})$. The *projection* of the range set $\tilde{\mathcal{R}}_k^z$ on a subset $\mathcal{D}' \subseteq \mathcal{D}$ of the dataset \mathcal{D} is defined as $\mathcal{P}(\tilde{\mathcal{R}}_k^z, \mathcal{D}') = \{X(R, \mathcal{D}') : R \in \mathcal{R}_k^z\}$.

If $\mathcal{P}(\tilde{\mathcal{R}}_k^z, \mathcal{D}') = 2^{|\mathcal{D}'|}$, i.e., the projection of $\tilde{\mathcal{R}}_k^z$ on \mathcal{D}' produces all possible dichotomies of the elements of \mathcal{D}' , then \mathcal{D}' is *shattered* by $\tilde{\mathcal{R}}_k^z$. The VC-dimension of the range space $(\mathcal{D}, \tilde{\mathcal{R}}_k^z)$ is the cardinality of the largest subset of \mathcal{D} that is shattered by $\tilde{\mathcal{R}}_k^z$. In what follows we will use $VC(\mathcal{R}_k^z)$ to denote the VC-dimension of the range space $(\mathcal{D}, \tilde{\mathcal{R}}_k^z)$ associated to the set of rule lists \mathcal{R}_k^z .

4.1.1 Upper bounds to the VC-dimension. We now prove refined upper bounds to the VC-dimension of rule lists. We start by providing upper bounds to the *growth function* (the maximum cardinality of distinct projections of a range set, see below) of rule lists \mathcal{R}_k^1 , for which $z = 1$ (i.e., each rule consists of a condition on exactly 1 feature). From such bound we derive an upper bound to the VC-dimension of the range space associated to \mathcal{R}_k^1 , that we will generalize to derive an upper bound to the VC-dimension of \mathcal{R}_k^z for general values of z . Due to space constraints, some of the proofs are in Appendix.

Define the *growth function* $\Lambda(\mathcal{R}_k^z, m)$ of the rule lists \mathcal{R}_k^z as the maximum number of *distinct* projections of rule lists in \mathcal{R}_k^z over any dataset \mathcal{D} with m instances, that is:

$$\Lambda(\mathcal{R}_k^z, m) = \max_{\mathcal{D}' \subseteq \mathcal{D}: |\mathcal{D}'|=m} |\mathcal{P}(\tilde{\mathcal{R}}_k^z, \mathcal{D}')|.$$

THEOREM 4.1. *It holds*

$$\Lambda(\mathcal{R}_k^1, m) \leq 2 + \sum_{j=1}^k \left[2^j \prod_{i=0}^{j-1} (d-i) \right].$$

PROOF. For $0 \leq j \leq k$, we define the set \mathcal{G}_j of rule lists with length exactly j as $\mathcal{G}_j = \{R \in \mathcal{R}_k^1 : |R| = j\}$, and define $\tilde{\mathcal{G}}_j = \{X(R, \mathcal{D}) : R \in \mathcal{G}_j\}$ the range set for \mathcal{G}_j , with $\tilde{\mathcal{R}}_k^1 = \cup_j \tilde{\mathcal{G}}_j$. From an union bound and Jensen's inequality, it holds

$$\begin{aligned} \Lambda(\mathcal{R}_k^1, m) &= \max_{\mathcal{D}' \subseteq \mathcal{D}: |\mathcal{D}'|=m} |\mathcal{P}(\cup_j \tilde{\mathcal{G}}_j, \mathcal{D}')| \\ &\leq \max_{\mathcal{D}' \subseteq \mathcal{D}: |\mathcal{D}'|=m} \sum_{j=0}^k |\mathcal{P}(\tilde{\mathcal{G}}_j, \mathcal{D}')| \leq \sum_{j=0}^k \Lambda(\mathcal{G}_j, m). \end{aligned}$$

We proceed to bound each $\Lambda(\mathcal{G}_j, m)$ separately, while excluding rule lists $R \in \mathcal{G}_j$ that have the same projection, on any \mathcal{D}' , of other rules $R' \in \mathcal{G}_{j'}$, for some $j' \neq j$. First, consider \mathcal{G}_0 ; it holds $\Lambda(\mathcal{G}_0, m) = 2$ since \mathcal{G}_0 contains $R_1 = [(c_\top, 1)]$ and $R_2 = [(c_\top, 0)]$. Then, consider \mathcal{G}_j for any $1 \leq j \leq k$, and let any rule list $R \in \mathcal{G}_j$, with $R = [(c_1, p_1), \dots, (c_j, p_j), (c_\top, p)]$. There are two possibilities: $p_j = p$ or $p_j \neq p$. In the first case, denote the rule $R' = [(c_1, p_1), \dots, (c_{j-1}, p_{j-1}), (c_\top, p)] \in \mathcal{G}_{j-1}$. It is easy to observe that the projections $X(R, \mathcal{D}')$ and $X(R', \mathcal{D}')$ are equal, for any \mathcal{D}' . Therefore, any rule list $R \in \mathcal{G}_j$ with $p_j = p$ can be ignored from the upper bound to $\Lambda(\mathcal{G}_j, m)$ as it is already covered by at least one element of the set \mathcal{G}_{j-1} . Then, consider a rule list $R = [(c_1, p_1), \dots, (c_j, p_j), (c_\top, p)] \in \mathcal{G}_j$ such that there exist two indices $1 \leq y < i \leq j$ with $c_y = c_i$. We observe that the rule $R' = [(c_1, p_1), \dots, (c_{i-1}, p_{i-1}), (c_{i+1}, p_{i+1}), \dots, (c_j, p_j), (c_\top, p)] \in \mathcal{G}_{j-1}$ is equivalent to R , since $X(R, \mathcal{D}') = X(R', \mathcal{D}')$ for any \mathcal{D}' . From these observations, the number of rules lists $R \in \mathcal{G}_j$ with distinct projections on a dataset with m elements cannot exceed $2^j \prod_{i=0}^{j-1} (d-i)$, as we consider all possible ordered choices of j distinct elements from the set of d features $\{x_1, \dots, x_d\}$, and all 2^j possible assignments of the predicted values $\{p_1, \dots, p_j\} \in \{0, 1\}^j$, while p is set to the unique value $\neq p_j$. The sum of these bounds yields the statement. \square

A consequence of Theorem 4.1 is the following upper bound to the VC-dimension $VC(\mathcal{R}_k^1)$ of \mathcal{R}_k^1 .

COROLLARY 4.2. *The VC-dimension $VC(\mathcal{R}_k^1)$ of \mathcal{R}_k^1 is*

$$VC(\mathcal{R}_k^1) \leq \lfloor k \log_2(2d) + 2 \rfloor.$$

PROOF. We first note that the VC-dimension of a range set Q is at most m if it holds $\log_2(\Lambda(Q, m)) < m$. Using the upper bound to

$\Lambda(\mathcal{R}_k^1, m)$ from Theorem 4.1, we have the following:

$$\begin{aligned} \log_2(\Lambda(\mathcal{R}_k^1, m)) &\leq \log_2 \left(2 + \sum_{j=1}^k \left[2^j \prod_{i=0}^{j-1} (d-i) \right] \right) \\ &\leq \log_2 \left(2 + \sum_{j=1}^k (2d)^j \right) \leq \log_2 \left(1 + \frac{(2d)^{k+1}}{2d-1} \right) \\ &\leq \log_2 \left(\frac{(2d)^{k+1}}{2d-1} \right) + \log_2(e) \frac{2d-1}{(2d)^{k+1}}, \end{aligned}$$

where in the last inequality we have used the fact that $\log_2(1+x) \leq \log_2(x) + \log_2(e)/x$, $\forall x \geq 0$. Then, we have

$$\begin{aligned} &\log_2 \left(\frac{(2d)^{k+1}}{2d-1} \right) + \log_2(e) \frac{2d-1}{(2d)^{k+1}} \\ &\leq (k+1) \log_2(2d) - \log_2(2d-1) + \log_2(e)(2d)^{-k} \\ &= k \log_2(2d) + \log_2 \left(\frac{2d}{2d-1} \right) + \log_2(e)(2d)^{-k} \\ &\leq k \log_2(2d) + 1 + \log_2(e)(2d)^{-k} < k \log_2(2d) + 2, \end{aligned}$$

proving the statement. \square

Using the above upper bound to $VC(\mathcal{R}_k^1)$, we prove the following upper bound to the VC-dimension $VC(\mathcal{R}_k^z)$ of \mathcal{R}_k^z .

COROLLARY 4.3. *The VC-dimension $VC(\mathcal{R}_k^z)$ of \mathcal{R}_k^z is*

$$VC(\mathcal{R}_k^z) \leq \left\lceil kz \log_2 \left(\frac{2ed}{z} \right) + 2 \right\rceil.$$

The bounds we derived in Corollary 4.2 and Corollary 4.3 scale linearly with k and z , and logarithmically with d . This implies that, since we are interested in *sparse* models with small values of k and z , the resulting complexity will not be large. This guarantees that accurate models will generalize well, and that the random sample \mathcal{S} should be representative of the dataset \mathcal{D} . Moreover, the weak dependence of the bounds on d allows SAMRULE to be extremely effective even on datasets with large sets of features.

Rudin et al. [24] study generalization bounds for classes of binary classification models built from sequences of association rules; this general model includes rule lists. Moreover, they show that the VC-dimension of the set of rule lists created using pre-mined rules is exactly the size of the set of pre-mined rules, which is $O(d^z)$ (see Theorem 3 in [24]). However, we remark that this result only holds for unconstrained rule lists (e.g., without any bound k on their length), and cannot be adapted to our setting. Our $O(kz \log(d/z))$ upper bound on the VC-dimension is more accurate, since it offers a more granular dependence on the parameters defining the rule list search space. More precisely, it explicitly depends on the maximum number k of rules in the list, the number of features d , and the maximum number z of conditions in each rule, while the bound $O(d^z)$ from [24] is not sensible to such constraints (as it assumes $k = d^z$). Since in most cases $k \ll d^z$, our upper bounds are orders of magnitude smaller.

4.1.2 Lower bounds to the VC-dimension. In this section we prove almost matching *lower bounds* to the VC-dimension of rule lists, confirming the tightness of our analysis. Our first result involves a lower bound to $VC(\mathcal{R}_k^1)$.

THEOREM 4.4. *The VC-dimension $VC(\mathcal{R}_k^1)$ of \mathcal{R}_k^1 is*

$$VC(\mathcal{R}_k^1) \geq \left\lfloor k \log_2 \left(\frac{d+k}{k} \right) \right\rfloor.$$

We may observe that the upper bound $O(k \log(d))$ given by Corollary 4.2 is almost tight, since Theorem 4.4 implies that $VC(\mathcal{R}_k^1) \in \Omega(k \log(\frac{d+k}{k}))$. Reducing this gap is an interesting open problem. We remark that the lower bound from Theorem 4.4 improves the best known result, which provides a (weaker) lower bound of rate $\Omega(k + \log(d-k))$ (see Theorem 3 of [31]).

We now prove a lower bound for the general case $z \geq 1$.

THEOREM 4.5. *The VC-dimension $VC(\mathcal{R}_k^z)$ of \mathcal{R}_k^z is*

$$VC(\mathcal{R}_k^z) \geq \left\lfloor kz \log_2 \left(\frac{d}{z \sqrt[k]{k}} \right) \right\rfloor.$$

From Theorem 4.5 and Corollary 4.3 we conclude that $VC(\mathcal{R}_k^z) \in O(kz \log(d/z))$ and $VC(\mathcal{R}_k^z) \in \Omega(kz \log(d/z \sqrt[k]{k}))$, which are almost matching. A further refinement of either the upper or lower bound is an interesting direction for future research.

4.2 Approximation guarantees

In this section we prove sharp approximation bounds for the estimated losses $\ell(R, \mathcal{S})$ of rules R on a random sample \mathcal{S} w.r.t. the losses $\ell(R, \mathcal{D})$ on the dataset \mathcal{D} , and derive sufficient sample sizes to compute (ε, θ) -approximations for the optimal rule list R^* . Due to space constraints, we provide most of the proofs in the Appendix. We define ω as $\omega = kz \ln(2ed/z) + 2$, a complexity parameter of our bounds derived from the results of Section 4.1 (Corollary 4.3).

THEOREM 4.6. *Let \mathcal{S} be an i.i.d. random sample of size $m \geq 1$ of the dataset \mathcal{D} . Then, for $\delta \in (0, 1)$, the following inequalities hold simultaneously with probability $\geq 1 - \delta$:*

$$\begin{aligned} \ell(R, \mathcal{D}) &\leq \ell(R, \mathcal{S}) + \sqrt{\frac{2\ell(R, \mathcal{S}) \left(\omega + \ln\left(\frac{2}{\delta}\right) \right)}{m}} + \frac{2 \left(\omega + \ln\left(\frac{2}{\delta}\right) \right)}{m}, \forall R \in \mathcal{R}_k^z, \\ \ell(R^*, \mathcal{S}) &\leq \ell(R^*, \mathcal{D}) + \sqrt{\frac{3\ell(R^*, \mathcal{D}) \ln\left(\frac{2}{\delta}\right)}{m}}. \end{aligned}$$

From the concentration bounds of Theorem 4.6 we conclude that, with high probability, all estimates $\ell(R, \mathcal{S})$ from the sample \mathcal{S} provide accurate upper bounds to the losses $\ell(R, \mathcal{D})$ in the dataset, while the loss $\ell(R^*, \mathcal{S})$ of the optimal rule list strongly concentrates to its expectation $\ell(R^*, \mathcal{D})$. We now combine these inequalities to obtain an easy-to-check condition; this condition yields a sufficient number of samples to obtain a high-quality approximation of the optimal rule list R^* from a random sample. Define $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ as the minimum value of $m \geq 1$ such that it holds

$$\sqrt{\frac{3\theta \ln\left(\frac{2}{\delta}\right)}{m}} + \sqrt{\frac{2 \left(\theta + \sqrt{\frac{3\theta \ln\left(\frac{2}{\delta}\right)}{m}} \right) \left(\omega + \ln\left(\frac{2}{\delta}\right) \right)}{m}} + \frac{2 \left(\omega + \ln\left(\frac{2}{\delta}\right) \right)}{m} \leq \varepsilon \theta. \quad (1)$$

We prove that a sample \mathcal{S} of size at least $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ provides an (ε, θ) -approximation of R^* with high probability.

THEOREM 4.7. *Let \mathcal{S} be an i.i.d. random sample of size $\geq \hat{m}(\mathcal{R}_k^z, \mathcal{D})$ of the dataset \mathcal{D} . The output $\tilde{R} = \arg \min_{R \in \mathcal{R}_k^z} \ell(R, \mathcal{S})$ of SAMRULE provides an (ε, θ) -approximation of R^* with probability $\geq 1 - \delta$.*

PROOF. Our proof is based on obtaining an upper bound to $\ell(\tilde{R}, \mathcal{D})$, where \tilde{R} is the rule list reported by SAMRULE, that is a function of the optimal loss $\ell(R^*, \mathcal{D})$, and then showing that it is sufficiently small to guarantee an (ε, θ) -approximation. First, we observe that $\ell(R^*, \mathcal{D}) \leq \ell(\tilde{R}, \mathcal{D})$, as R^* is one of the rule lists with minimum loss. Therefore, from the first set of inequalities of Theorem 4.6 setting $R = \tilde{R}$, we have the upper bound to $\ell(R^*, \mathcal{D})$

$$\ell(R^*, \mathcal{D}) \leq \ell(\tilde{R}, \mathcal{S}) + \sqrt{\frac{2\ell(\tilde{R}, \mathcal{S})\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}} + \frac{2\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}. \quad (2)$$

We now prove that, since \tilde{R} is chosen as the rule list with minimum loss on the sample, its estimated loss $\ell(\tilde{R}, \mathcal{S})$ should be concentrated towards $\ell(R^*, \mathcal{D})$. We have $\ell(\tilde{R}, \mathcal{S}) \leq \ell(R^*, \mathcal{S})$, as \tilde{R} is one of the rule lists with minimum loss in \mathcal{S} . Using the last inequality of Theorem 4.6 it holds

$$\ell(\tilde{R}, \mathcal{S}) \leq \ell(R^*, \mathcal{S}) \leq \ell(R^*, \mathcal{D}) + \sqrt{\frac{3\ell(R^*, \mathcal{D}) \ln\left(\frac{2}{\delta}\right)}{m}}. \quad (3)$$

Replacing all occurrences of $\ell(\tilde{R}, \mathcal{S})$ in (2) with its upper bound given by the r.h.s. of (3), we obtain

$$\ell(\tilde{R}, \mathcal{D}) \leq \ell(R^*, \mathcal{D}) + u(\ell(R^*, \mathcal{D}), m), \quad (4)$$

where the function $u(v, m)$ is

$$u(v, m) = \sqrt{\frac{3v \ln\left(\frac{2}{\delta}\right)}{m}} + \sqrt{\frac{2\left(v + \sqrt{\frac{3v \ln\left(\frac{2}{\delta}\right)}{m}}\right)\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}} + \frac{2\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}.$$

We now seek to identify the minimum m such that $\ell(\tilde{R}, \mathcal{D})$ is guaranteed to be $\leq \ell(R^*, \mathcal{D}) + \varepsilon \max\{\theta, \ell(R^*, \mathcal{D})\}$. Note that, from the derivations above, it is sufficient to verify that $u(\ell(R^*, \mathcal{D}), m) \leq \varepsilon \max\{\theta, \ell(R^*, \mathcal{D})\}$, for any $\ell(R^*, \mathcal{D}) \in [0, 1]$. First, consider the case $\ell(R^*, \mathcal{D}) \leq \theta$. We have that

$$u(\ell(R^*, \mathcal{D}), m) \leq u(\theta, m) \leq \varepsilon \theta$$

is true when $m \geq \hat{m}(\mathcal{R}_k^z, \mathcal{D})$ by definition of $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ given in the statement, since $u(\theta, m) \leq \varepsilon \theta$ is guaranteed by (1). We now consider the case $\ell(R^*, \mathcal{D}) \geq \theta$. We need to verify that

$$u(v, m) \leq \varepsilon v, \forall \theta \leq v \leq 1. \quad (5)$$

From the definition of $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$, we know that $u(\theta, m) \leq \varepsilon \theta$ holds for all $m \geq \hat{m}(\mathcal{R}_k^z, \mathcal{D})$, therefore (5) is verified for $v = \theta$. By taking the derivative w.r.t. v of both sides of (5), it is simple to check that the derivative of the r.h.s. is constant, while the derivative of the l.h.s. is monotonically decreasing with v . Therefore, (5) also holds for all $\theta < v \leq 1$, obtaining the statement. \square

We note that, while $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ can be easily computed (e.g., with a binary search over m), from (1) it may not be simple to interpret

its dependence on the several parameters. To do so, we prove the following upper bound to $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$.

THEOREM 4.8. *It holds $\hat{m}(\mathcal{R}_k^z, \mathcal{D}) \in \mathcal{O}\left(\frac{\omega + \ln\left(\frac{1}{\delta}\right)}{\varepsilon^2 \theta}\right)$.*

PROOF. First, we prove that $\hat{m}(\mathcal{R}_k^z, \mathcal{D}) \geq 3 \ln(2/\delta)/\theta$. This follows easily from $\sqrt{\frac{3\theta \ln(2/\delta)}{m}} \leq \varepsilon \theta$, i.e., only considering the left-most term of (1). Then, it holds

$$\begin{aligned} & \sqrt{\frac{3\theta \ln\left(\frac{2}{\delta}\right)}{m}} + \sqrt{\frac{2\left(\theta + \sqrt{\frac{3\theta \ln\left(\frac{2}{\delta}\right)}{m}}\right)\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}} + \frac{2\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m} \\ & \leq \sqrt{\frac{3\theta \ln\left(\frac{2}{\delta}\right)}{m}} + \sqrt{\frac{4\theta\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}} + \frac{2\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m} \\ & \leq \sqrt{\frac{14\theta\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}} + \frac{2\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}, \end{aligned}$$

where the second-last inequality holds for all $m \geq 3 \ln(2/\delta)/\theta$. By solving the quadratic inequality

$$\sqrt{14m\theta\left(\omega + \ln(2/\delta)\right)} + 2\left(\omega + \ln(2/\delta)\right) \leq m\varepsilon\theta,$$

we observe that (1) holds for all m with

$$m \geq \frac{\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)\sqrt{14}\sqrt{8\varepsilon + 14}}{2\varepsilon^2\theta} + \frac{2\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{\varepsilon\theta} + \frac{7\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{\varepsilon^2\theta}. \quad (6)$$

The r.h.s. of (6) is in $\mathcal{O}\left(\frac{\omega + \ln\left(\frac{1}{\delta}\right)}{\varepsilon^2\theta}\right)$; moreover, it upper bounds $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$, as (6) provides a value of m that make (1) true (not necessarily the minimum), obtaining the statement. \square

Interestingly, the sample size $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ is completely independent of the size n of the dataset \mathcal{D} : it only depends on the parameters $\omega = kz \ln(2ed/z) + 2$ of the rule lists search space, and the desired approximation accuracy ε , θ and confidence δ . This characteristic allows SAMRULE to be applied to massive datasets, of arbitrarily large size n .

5 Experiments

This section presents the results of our experiments. The main goal of our experimental evaluation is to test the scalability of SAMRULE in analyzing large datasets compared to exact approaches (Section 5.1). To do so, we measure the number of samples used by SAMRULE to obtain an accurate approximation of the best rule, its running time, and the accuracy of the reported rule list compared to the optimal solution. Then, we also compare SAMRULE with state-of-the-art heuristics for rule list training (Section 5.2), as such methods, while not offering theoretical guarantees, may still provide accurate rule lists in practice. Finally, we quantify the performance of SAMRULE under several settings of the rule list parameters z and k (Section 5.3).

Datasets. We tested SAMRULE on 8 benchmark datasets from UCI¹. We binarized the datasets containing continuous features by considering 4 thresholds at equally spaced quantiles: for each continuous feature f and each threshold t , we created two binary features “ $f \geq t$ ” and “ $f < t$ ”. The statistics of the resulting binary

¹<https://archive.ics.uci.edu/>

datasets are described in Table 1. Since some of these datasets are quite small, we replicate them r times, i.e., each training sample of \mathcal{D} is copied r times (see the column r of Table 1). Note that this preprocessing allows to obtain larger datasets, all with approximately 10^6 training instances, while preserving the rule list distribution and search space structure, as the losses and covered instances of all rule lists are the same on the original and replicated datasets.

Compared methods. Our main goal is to compare SAMRULE with CORELS [1], the state-of-the-art exact method to identify the rule list with minimum loss from a dataset \mathcal{D} . SAMRULE, as described in Section 4, generates the random sample \mathcal{S} and then runs an exact algorithm to search for the rule list with minimum loss on \mathcal{S} (using the same settings and parameters). Note that by using CORELS as exact algorithm within SAMRULE, all differences between SAMRULE and CORELS are due to the use of sampling, i.e., there are no other confounding factors in the comparison between the two methods. We also compare SAMRULE with SBRL [30] and RIPPER [7], two state-of-the-art heuristic approaches for rule list learning. SBRL uses a scalable Monte-Carlo approach to approximately search for accurate rule lists, while RIPPER leverages a greedy selection of conditions. Our comparison is motivated by the fact that, even if these methods do not provide guarantees in terms of solution quality, they are often the methods of choice for practitioners since they obtain good solutions and are designed to scale to large datasets.

Rule Lists Parameters. To compare SAMRULE with exact and heuristic methods (Sections 5.1 and 5.2), for each dataset we set the parameter k as shown in Table 1 and $z = 1$. In Section 5.3 we evaluate the impact of different choices of k and z . Regarding SAMRULE, we compute (ε, θ) -approximations for several combinations of the parameters ε and θ . We consider $\varepsilon \in \{1, 0.5, 0.25\}$, and vary θ in the interval $[0.005, 0.05]$. For all experiments we fix $\delta = 0.05$, as we did not observe significant differences for other values (given the exponential dependence of the bounds w.r.t. δ).

Experimental setup. We implemented SAMRULE in Python. The code and the scripts to reproduce all experiments are available online². For CORELS, we have used the implementation available online³. We made minor modifications to the original implementation to limit its exploration to rule lists of length at most k (i.e., pruning all rules of length $> k$ instead of performing an unbounded search). We made similar minor changes to the implementation of SBLR⁴. We evaluated RIPPER with a recent efficient implementation⁵. All the code was compiled and executed on a machine equipped with 2.30 GHz Intel Xeon CPU, 1 TB of RAM, on Ubuntu 20.04. We repeated all experiments 10 times, and report averages \pm stds over the 10 repetitions.

5.1 Comparison to exact method

In this section we describe the experimental comparison between SAMRULE and CORELS, the state-of-the-art method to identify the optimal rule list R^* with minimum loss. Our main goal is to evaluate the scalability of SAMRULE in terms of number of samples and running time required to obtain an accurate approximation of the best rule. Furthermore, we evaluate, both quantitatively and

Table 1: Statistics of the datasets considered in our experiments. n is the number of transactions, d is the number of binary features, r is the replication factor, k is the maximum rule list length.

\mathcal{D}	n	d	r	k
a9a	32561	124	100	4
adult	32561	175	100	4
bank	41188	152	100	4
higgs	11000000	531	1	3
ijcnn1	91701	35	100	8
mushroom	8124	118	200	5
phishing	11050	69	100	8
susy	5000000	179	1	5

qualitatively, the accuracy of the rule list found by SAMRULE w.r.t. the optimal solution returned by CORELS. We ran both methods on all datasets, setting $z = 1$ and k as in Table 1. For SAMRULE, we vary the parameters θ and ε as described at the beginning of Section 5.

Figure 2 shows the results for these experiments. In Figure 2.(a) and (b) we compare the running time of SAMRULE with the time needed by CORELS on the adult and bank datasets. The results for other datasets are very similar, and shown in Figure 6 (in the Appendix). From these results, we can immediately conclude that SAMRULE requires a *small fraction* of the time needed by CORELS, with an improvement of up to 2 orders of magnitude. The reason for this significant speedup is that SAMRULE searches for the rule with minimum loss on a small sample \mathcal{S} , which is all cases orders of magnitude smaller than the size of original dataset \mathcal{D} . We show the number of samples $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ used by SAMRULE for all datasets and for all parameters in Figure 5 (in the Appendix).

Then, we evaluated the quality of the solutions returned by SAMRULE with the optimal rule list computed by CORELS. Figures 2.(c) and (d) show the average deviations $|\ell(\tilde{R}, \mathcal{D}) - \ell(R^*, \mathcal{D})|$ of the loss of the rule list \tilde{R} found by SAMRULE w.r.t. the optimal rule list R^* found by CORELS on the dataset \mathcal{D} . To verify the validity of SAMRULE’s theoretical guarantees, the plots also show the loss of the optimal solution $\ell(R^*, \mathcal{D})$ found by CORELS (purple horizontal line), and upper error bars (+std) for the average deviations (see Figures 6 and 7 in the Appendix for the plots for all datasets with both bars). From these results, we observe that the rule lists reported by SAMRULE are extremely accurate in terms of prediction accuracy, since the deviations $|\ell(\tilde{R}, \mathcal{D}) - \ell(R^*, \mathcal{D})|$ are *orders of magnitude* smaller than $\ell(R^*, \mathcal{D})$, and smaller than guaranteed by our theoretical analysis. This confirms that SAMRULE outputs extremely accurate rule lists, even when trained on random samples that are orders of magnitude smaller than the entire dataset. Furthermore, it is likely that the guarantees of the (ε, θ) -approximations hold for samples smaller than what guaranteed by our analysis; this leaves significant opportunities for further improvements of our algorithm. Then, we quantify the deviations between the loss $\ell(\tilde{R}, \mathcal{S})$ of \tilde{R} estimated on the sample w.r.t. to the loss $\ell(\tilde{R}, \mathcal{D})$ on the dataset, i.e., the approximation error incurred by SAMRULE due to analyzing \mathcal{S} instead of the entire dataset \mathcal{D} . In Figure 8 we show the average loss approximation error $|\ell(\tilde{R}, \mathcal{D}) - \ell(\tilde{R}, \mathcal{S})|$, which is

²<https://github.com/VandinLab/SamRule>

³<https://github.com/corels/corels>

⁴<https://github.com/Hongyuy/sbrlmod/>

⁵<https://github.com/imoscovitz/wittgenstein>

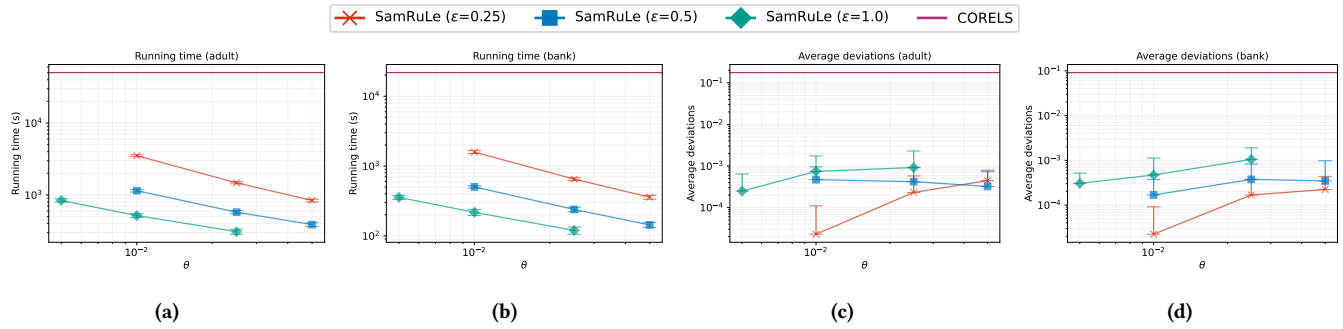


Figure 2: Performance and accuracy comparison between SAMRuLE and CORELS on adult and bank datasets, for different values of ϵ and θ . (a)-(b): running times of CORELS and SAMRuLE. (c)-(d): average deviations $|\ell(\tilde{R}, \mathcal{D}) - \ell(R^*, \mathcal{D})|$ of the loss of the rule list \tilde{R} found by SAMRuLE with the optimal rule list R^* found by CORELS (purple horizontal line drawn at $y = \ell(R^*, \mathcal{D})$). The deviation plots only show upper errors bars at +std to improve readability. See Figures 6 and 7 in the Appendix for the plots for all datasets and with \pm std bars.




Figure 3: (a): optimal rule R^* computed by CORELS on the adult dataset ($\ell(R^*, \mathcal{D}) = 0.176$) to predict high income (the label 1 denotes “ $\geq 50K$ ”). (b): set of rule lists computed by SAMRuLE over 10 runs. SAMRuLE identified the optimal rule and slight variations \tilde{R}_1 and \tilde{R}_2 that differ in the second rule of the list: they predict a lower outcome using the age (\tilde{R}_1) and the per-week work hours features (\tilde{R}_2) with respective loss $\ell(\tilde{R}_1, \mathcal{D}) = 0.1763$ and $\ell(\tilde{R}_2, \mathcal{D}) = 0.1775$.

also extremely small (i.e., 1 to 2 orders of magnitude smaller than $\ell(\tilde{R}, \mathcal{D})$), confirming that the conclusions that can be drawn from the sample using SAMRuLE, e.g., from the estimated loss $\ell(\tilde{R}, \mathcal{S})$, are very close to the corresponding exact ones. Finally, we compared the logical conditions in the rule lists returned by SAMRuLE with the ones in the best rule list computed by CORELS. Our goal is to verify that the insights gained from the approximated prediction models from SAMRuLE were similar to the optimal ones, i.e., that SAMRuLE allows a qualitative interpretation of the reported rule list that was stable over the different experimental runs and similar to what obtainable from the exact analysis. Figure 3 reports the optimal rule list computed by CORELS on the dataset adult (a), and the set of rule lists computed by SAMRuLE (b) over all runs. Interestingly, we observe that the approximations from SAMRuLE either match the optimal solution, or are very similar to it. In fact, all rules reported by SAMRuLE either share the same conditions found in the optimal rule list, or replace one of the feature with alternative reasonable insights (e.g., predicting a lower income for young individuals and limited weekly working hours). In general, we found the solutions reported by SAMRuLE to be extremely stable and similar to the respective optimal solutions also for all other dataset.

From these observations we conclude that SAMRuLE computes extremely accurate rule lists using a fraction of the resources needed by exact approaches, therefore scaling effectively to large datasets.

5.2 Comparison to heuristic methods

In this set of experiments, we compare SAMRuLE with two state-of-the-art heuristic methods SBRL and RIPPER. For SAMRuLE we fix $\theta = 0.025$ and $\epsilon = 0.5$. We ran SBRL for 10^4 iterations using default parameters (as suggested by [1, 30]). Also for these experiments, for each dataset we fix $z = 1$ and k to the values in Table 1.

We show the results of these experiments in Figure 4.

Figure 4.(a) shows the running times for the three methods. We observe that in all but one case, RIPPER is the slowest method. On two datasets (ijcnn1 and higgs) we stopped it as it could not complete after more than 12 hours, while requiring a very large memory footprint (e.g., more than 400 GB of memory for higgs). Regarding SBRL and SAMRuLE, both methods are fast, e.g., requiring always less than 18 minutes. More precisely, SAMRuLE is faster on 3 datasets, up to 3 orders of magnitude for the susy dataset. In other 2 datasets, the running times of the two methods are comparable, while for phishing and adult SBRL is faster by a factor at most 4. This experiment confirms that SAMRuLE is very practical, requiring a lower or comparable amount of resources of state-of-the-art scalable heuristics.

Figure 4.(b) compares all methods in terms of accuracy. The plots show the values of the loss $\ell(\tilde{R}, \mathcal{D})$ for the rule list \tilde{R} reported by all runs of the methods. We may observe that RIPPER is the worst approach, as it always reports rule lists with the highest loss

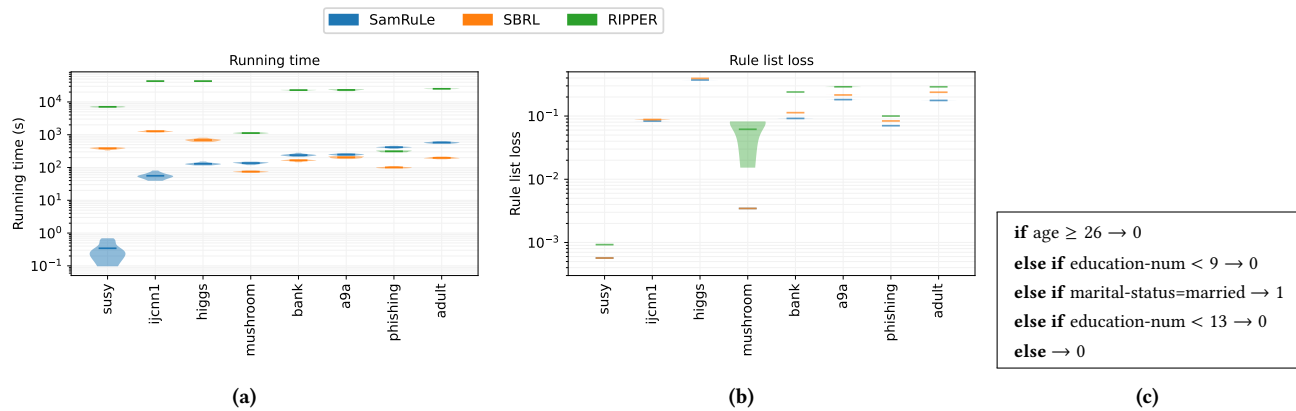


Figure 4: Comparison in terms of running time (a) and accuracy (b) between SAMRULE, SBRL, and RIPPER. (c): rule R computed by SBRL on adult with loss $\ell(R, \mathcal{D}) = 0.238$ over all 10 runs.

(results for ijcn1 and higgs are not shown as RIPPER could not complete in reasonable time, as discussed before), while SAMRULE always provides a rule list with the smallest loss. Regarding SBRL, we observe that it reports a rule list with the same, or almost the same, loss of SAMRULE on 4 datasets, while it provides suboptimal solutions for other cases, with losses up to 30% higher than SAMRULE. This suggests that, while SBRL scales to large datasets, it often provides solutions that are sensibly less accurate than the optimal one. Instead, as discussed previously, SAMRULE outputs a rule list with guaranteed gap with the optimal solution, and always very close to it in practice. We remark that providing a suboptimal solution may also impact the interpretation for the predictions, e.g., missing relevant factors for the model. In fact, we report the rule list computed by SBRL on adult (see Figure 4.(c)) for all the 10 runs: such rule list does not involve the capital-gain feature, a key condition for the optimal solution to predict high income (Figure 3.(a)), thus obtaining a higher loss (0.238). In contrast, this feature was always found in the rule lists reported by SAMRULE (Figure 3.(b)), which have loss always very close to the optimal (0.176).

Overall, compared to RIPPER and SBRL, SAMRULE outputs an equally accurate solution in less time, or a sensibly better rule list using comparable resources. We conclude that SAMRULE provides an excellent combination of scalability and high accuracy for rule list learning from large datasets, achieving a better trade-off than state-of-the-art heuristic methods with no theoretical guarantees.

5.3 Impact of rule list parameters

In this final set of experiments, we evaluate the impact to the performance of SAMRULE of the rule list search space parameters z and k . We focus on the datasets mushroom and phishing, as the results for other datasets were similar. We test all values of $1 \leq z \leq 3$ and $1 \leq k \leq 5$, fixing $\theta = 0.025$ and $\varepsilon = 0.5$, measuring the number of samples required by SAMRULE, its running time, and the accuracy of the rule lists provided in output.

We show these results in Figure 9. When increasing k and z , the number of samples $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ considered by SAMRULE grows following the expect trend $\hat{m}(\mathcal{R}_k^z, \mathcal{D}) \in O((\omega + \log(1/\delta))/(\varepsilon^2\theta))$ proved in our analysis (Theorem 4.8), resulting in sample sizes that

are always a small fraction of the size of the dataset. Consequently, the running time of SAMRULE increases roughly linearly with the sample size, remaining practical for all settings. Regarding the accuracy of the rule lists, we observed the parameter k to have the largest impact on the loss, that remains fairly stable w.r.t. z .

These results demonstrate that SAMRULE is applicable to complex analysis involving larger values of k and z while scaling to large datasets, that are in most cases out of reach of exact approaches.

6 Conclusions

We introduced SAMRULE, a novel and scalable algorithm to find nearly optimal rule lists. SAMRULE uses sampling to scale to large datasets, and provides rigorous guarantees on the quality of the rule lists it reports. Our approach builds on the VC-dimension of rule lists, for which we proved novel upper and lower bounds. Our experimental evaluation shows that SAMRULE enables learning highly accurate rule lists on large datasets, is up to two orders of magnitude faster than state-of-the-art exact approaches, and is as fast as, and sometimes faster than, recent heuristic approaches, while reporting higher quality rule lists.

Our work opens several interesting directions for future research, including the use of sampling to scale approaches for learning other rule-based models while providing rigorous guarantees on the quality of the learned model. Moreover, the efficient computation of advanced *data-dependent* complexity measures, such as Rademacher averages [2, 16, 17], may be useful to obtain even sharper approximation guarantees for our problem.

Acknowledgments

This work was supported by the “National Center for HPC, Big Data, and Quantum Computing”, project CN00000013, and by the PRIN Project n. 2022TS4Y3N - EXPAND: scalable algorithms for EXPLORATORY ANALYSES OF heterogeneous and dynamic Networked Data, funded by the Italian Ministry of University and Research (MUR), and by the project BRAINTEASER (Bringing Artificial Intelligence home for a better care of amyotrophic lateral sclerosis and multiple sclerosis), funded by European Union’s Horizon 2020 (grant agreement No. GA101017598).

References

- [1] Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. 2018. Learning certifiably optimal rule lists for categorical data. *Journal of Machine Learning Research* 18, 234 (2018), 1–78.
- [2] Peter L. Bartlett and Shahar Mendelson. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3, Nov (2002), 463–482.
- [3] Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. 2005. Theory of classification: A survey of some recent advances. *ESAIM: probability and statistics* 9 (2005), 323–375.
- [4] Leo Breiman. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science* 16, 3 (2001), 199–231.
- [5] L. Breiman, J. Friedman, R. A. Olshen, and C. J. Stone. 1984. *Classification and Regression Trees*. Routledge.
- [6] Martino Ciaperoni, Han Xiao, and Aristides Gionis. 2023. Concise and interpretable multi-label rule sets. *Knowledge and Information Systems* 65, 12 (2023), 5657–5694.
- [7] William W. Cohen. 1995. Fast effective rule induction. In *Machine learning proceedings 1995*. Elsevier, 115–123.
- [8] Jiayun Dong and Cynthia Rudin. 2020. Exploring the cloud of variable importance for the set of all good models. *Nature Machine Intelligence* 2, 12 (2020), 810–824.
- [9] Jonas Fischer and Jilles Vreeken. 2019. Sets of robust rules, and how to find them. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 38–54.
- [10] Aaron Fisher, Cynthia Rudin, and Francesca Dominici. 2019. All Models are Wrong, but Many are Useful: Learning a Variable’s Importance by Studying an Entire Class of Prediction Models Simultaneously. *J. Mach. Learn. Res.* 20, 177 (2019), 1–81.
- [11] Van Quoc Phung Huynh, Johannes Fürnkranz, and Florian Beck. 2023. Efficient learning of large sets of locally optimal classification rules. *Machine Learning* 112, 2 (2023), 571–610.
- [12] Michael Mitzenmacher and Eli Upfal. 2017. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press.
- [13] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. *Foundations of machine learning*. MIT press.
- [14] Yuzuru Okajima and Kunihiro Sadamasa. 2019. Decision list optimization based on continuous relaxation. In *Proceedings of the 2019 SIAM International Conference on Data Mining*. SIAM, 315–323.
- [15] Ioanna Papagianni and Matthijs van Leeuwen. 2023. Discovering Rule Lists with Preferred Variables. In *International Symposium on Intelligent Data Analysis*. Springer, 340–352.
- [16] Leonardo Pellegrina. 2023. Efficient Centrality Maximization with Rademacher Averages. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*. Association for Computing Machinery, New York, NY, USA, 1872–1884. <https://doi.org/10.1145/3580305.3599325>
- [17] Leonardo Pellegrina, Cyrus Cousins, Fabio Vandin, and Matteo Riondato. 2022. MCRapper: Monte-Carlo Rademacher averages for poset families and approximate pattern mining. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 6 (2022), 1–29.
- [18] Hugo M Proença and Matthijs van Leeuwen. 2020. Interpretable multiclass classification by MDL-based rule lists. *Information Sciences* 512 (2020), 1372–1393.
- [19] J. Ross Quinlan. 1993. *C4.5: programs for machine learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. <http://portal.acm.org/citation.cfm?id=152181>
- [20] Matteo Riondato and Eli Upfal. 2014. Efficient Discovery of Association Rules and Frequent Itemsets through Sampling with Tight Performance Guarantees. *ACM Trans. Knowl. Disc. from Data* 8, 4 (2014), 20. <https://doi.org/10.1145/2629586>
- [21] Ronald L Rivest. 1987. Learning decision lists. *Machine learning* 2 (1987), 229–246.
- [22] Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. 2022. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys* 16, none (2022), 1 – 85.
- [23] Cynthia Rudin and Şeyda Ertekin. 2018. Learning customized and optimized lists of rules with mathematical programming. *Mathematical Programming Computation* 10 (2018), 659–702.
- [24] Cynthia Rudin, Benjamin Letham, and David B Madigan. 2013. Learning theory analysis for association rules and sequential event prediction. (2013).
- [25] Diego Santoro, Andrea Tonon, and Fabio Vandin. 2020. Mining Sequential Patterns with VC-Dimension and Rademacher Complexity. *Algorithms* 13, 5 (2020), 123.
- [26] Lesia Semenova, Cynthia Rudin, and Ronald Parr. 2022. On the existence of simpler machine learning models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1827–1858.
- [27] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- [28] V. N. Vapnik and A. Ya. Chervonenkis. 1971. On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. *Theory of Probability & Its Applications* 16, 2 (1971), 264. <https://doi.org/10.1137/1116025>
- [29] Rui Xin, Chudi Zhong, Zhi Chen, Takuya Takagi, Margo Seltzer, and Cynthia Rudin. 2022. Exploring the whole rashomon set of sparse decision trees. *Advances in Neural Information Processing Systems* 35 (2022), 14071–14084.
- [30] Hongyu Yang, Cynthia Rudin, and Margo Seltzer. 2017. Scalable Bayesian rule lists. In *International conference on machine learning*. PMLR, 3921–3930.
- [31] Olcay Taner Yildiz. 2014. VC-Dimension of Rule Sets. In *2014 22nd International Conference on Pattern Recognition*. IEEE, 3576–3581.
- [32] Jinqiang Yu, Alexey Ignatiev, Peter J Stuckey, and Pierre Le Bodic. 2021. Learning optimal decision sets and lists with sat. *Journal of Artificial Intelligence Research* 72 (2021), 1251–1279.

A Appendix

In this Appendix we provide proofs and additional experimental results that could not fit in the main paper due to space constraints. Some figures, and some of the proofs for the results of Section 4.1 are deferred to the online extended version, that is available at <https://arxiv.org/abs/2406.12803>.

A.1 Proofs of Section 4.1

PROOF OF COROLLARY 4.3. Given a dataset \mathcal{D} with d features, we create a new dataset \mathcal{D}^z built as follows. Let $C = \{“x_1 = 1”, “x_2 = 1”, \dots, “x_d = 1”\}$ be the set of all d possible conditions on the d binary features of \mathcal{D} ; for any non-empty subset $A \subseteq C$ with $|A| \leq z$, we add to \mathcal{D}^z the binary feature x'_A that is equal to 1 for all training instances of \mathcal{D} such that $\bigwedge_{c \in A} c$ is true. Equivalently, the feature values of x'_A are obtained from the logical AND of the evaluations of the conditions in A . It follows that the total number of features of \mathcal{D}^z is $d' = \sum_{i=1}^z \binom{d}{i}$.

We now observe that the set $\mathcal{P}(\tilde{\mathcal{R}}_k^z, \mathcal{D})$ of projections of rule lists in \mathcal{R}_k^z on \mathcal{D} is contained in the set $\mathcal{P}(\tilde{\mathcal{R}}_k^1, \mathcal{D}^z)$ of projections of rules in \mathcal{R}_k^1 on \mathcal{D}^z , since we can replace each rule with a conjunctions with t terms, with $1 \leq t \leq z$, in any rule list $\in \mathcal{R}_k^z$ on \mathcal{D} by a rule with a single condition on one of the features of \mathcal{D}^z , obtaining an equivalent rule list from \mathcal{R}_k^1 . Therefore, from Corollary 4.2 applied to \mathcal{R}_k^1 over a dataset with d' features, and the fact $d' = \sum_{i=1}^z \binom{d}{i} \leq \left(\frac{ed}{z}\right)^z$, we obtain

$$\begin{aligned} VC(\mathcal{R}_k^z) &\leq \lfloor k \log_2(2d') + 2 \rfloor = \left\lfloor k \log_2 \left(2 \sum_{i=1}^z \binom{d}{i} \right) + 2 \right\rfloor \\ &\leq \left\lfloor k \log_2 \left(2 \left(\frac{ed}{z} \right)^z \right) + 2 \right\rfloor \leq \left\lfloor kz \log_2 \left(\frac{2ed}{z} \right) + 2 \right\rfloor, \end{aligned}$$

and the statement follows. \square

A.2 Proofs of Section 4.2

To prove our results we use the following Chernoff bounds (see Theorem 4.4 and 4.5 of [12]).

THEOREM A.1. *Let X_1, \dots, X_m be independent Poisson trials such that $\Pr(X_i = 1) = p_i$. Let $X = \sum_{i=1}^m X_i$ and $\mu = \mathbb{E}[X]$. Then the following Chernoff bounds hold for any $0 < \gamma < 1$:*

$$\Pr(X \geq (1 + \gamma)\mu) \leq \exp(-\mu\gamma^2/3),$$

$$\Pr(X \leq (1 - \gamma)\mu) \leq \exp(-\mu\gamma^2/2).$$

PROOF OF THEOREM 4.6. We prove the first set of inequalities. Let R be an arbitrary rule list from \mathcal{R}_k^z , and define the functions

$$g(R, \mathcal{S}) = \ell(R, \mathcal{S}) - \alpha|R|, \quad (7)$$

$$g(R, \mathcal{D}) = \ell(R, \mathcal{D}) - \alpha|R|. \quad (8)$$

Note that $g(R, \mathcal{D})$ and $g(R, \mathcal{S})$ are the non-regularized variants of the loss functions $\ell(R, \mathcal{D})$ and $\ell(R, \mathcal{S})$. It easy to show that $\mathbb{E}_{\mathcal{S}}[g(R, \mathcal{S})] = g(R, \mathcal{D}) \leq \ell(R, \mathcal{D})$, and that $g(R, \mathcal{S})$ is an average of m binary random variables with expectation $g(R, \mathcal{D})$. From an application of the Chernoff bound (Theorem A.1) to the random variable $Z = mg(R, \mathcal{S})$ with $\mathbb{E}_{\mathcal{S}}[Z] = mg(R, \mathcal{D})$, we have that, for any $0 < \gamma < 1$,

$$\Pr\left(Z \leq (1 - \gamma)\mathbb{E}[Z]\right) \leq \exp\left(-\frac{\mathbb{E}[Z]\gamma^2}{2}\right). \quad (9)$$

Fixing $0 < \delta' < 1$, imposing the r.h.s. of (9) to be $\leq \delta'$, and solving for γ , gives

$$\Pr\left(g(R, \mathcal{S}) + \sqrt{\frac{2g(R, \mathcal{D}) \ln(1/\delta')}{m}} \leq g(R, \mathcal{D})\right) \leq \delta'.$$

Define the events

$$E_R = "g(R, \mathcal{S}) + \sqrt{\frac{2g(R, \mathcal{D}) \ln(1/\delta')}{m}} \leq g(R, \mathcal{D}), \forall R \in \mathcal{R}_k^z,$$

and the event $E = "\exists R : E_R \text{ is true}"$. Note that $\Pr(E_R) \leq \delta'$. We want to prove that $\Pr(E) \leq \delta/2$. First, we observe that

$$\Pr(E) = \Pr\left(\bigcup_{R \in \mathcal{R}_k^z} E_R\right).$$

Denote with R_1 and R_2 two rule lists $\in \mathcal{R}_k^z$ such that the projections $X(R_i, \mathcal{D})$ of R_i on \mathcal{D} are equal: $X(R_1, \mathcal{D}) = X(R_2, \mathcal{D})$. Consequently, it holds $g(R_1, \mathcal{D}) = g(R_2, \mathcal{D})$. Moreover, for all possible samples \mathcal{S} , it holds $X(R_1, \mathcal{S}) = X(R_2, \mathcal{S})$ and $g(R_1, \mathcal{S}) = g(R_2, \mathcal{S})$. The observations above imply that the events E_{R_1} and E_{R_2} are identical.

We now define the set $C(\mathcal{R}_k^z, \mathcal{D}) \subseteq \mathcal{R}_k^z$ as a cover of the rangeset $\tilde{\mathcal{R}}_k^z$ as follows: fix an arbitrary, deterministic total order over all elements of \mathcal{R}_k^z ; then, for every distinct projection $H \in \tilde{\mathcal{R}}_k^z$, let $\{R : X(R, \mathcal{D}) = H, R \in \mathcal{R}_k^z\}$ be the set of rule lists with projection equal to H . From this set, we pick the minimum element in the total order over \mathcal{R}_k^z and we include it in $C(\mathcal{R}_k^z, \mathcal{D})$. Therefore, $C(\mathcal{R}_k^z, \mathcal{D})$ contains a unique rule list for each distinct projection of the rangeset $\tilde{\mathcal{R}}_k^z$ over the dataset \mathcal{D} . This implies that, for any $R \in \mathcal{R}_k^z$, there is an unique rule list $R' \in C(\mathcal{R}_k^z, \mathcal{D})$ such that $X(R, \mathcal{D}) = X(R', \mathcal{D})$

and that the events E_R and $E_{R'}$ are identical. It follows

$$\Pr\left(\bigcup_{R \in \mathcal{R}_k^z} E_R\right) = \Pr\left(\bigcup_{R \in C(\mathcal{R}_k^z, \mathcal{D})} E_R\right),$$

as we can replace the union over all $R \in \mathcal{R}_k^z$ with the union over the cover $C(\mathcal{R}_k^z, \mathcal{D})$ containing rule lists with unique projections on \mathcal{D} . From the definitions of $\Lambda(\mathcal{R}_k^z, n)$ and $C(\mathcal{R}_k^z, \mathcal{D})$, it holds

$$|C(\mathcal{R}_k^z, \mathcal{D})| = |\tilde{\mathcal{R}}_k^z(\mathcal{D})| \leq \Lambda(\mathcal{R}_k^z, n).$$

From an union bound, we have

$$\Pr(E) = \Pr\left(\bigcup_{R \in C(\mathcal{R}_k^z, \mathcal{D})} E_R\right) \leq \sum_{R \in C(\mathcal{R}_k^z, \mathcal{D})} \Pr(E_R) \leq \Lambda(\mathcal{R}_k^z, n)\delta'.$$

Choosing $\delta' = \delta/(2\Lambda(\mathcal{R}_k^z, n))$, we obtain that

$$g(R, \mathcal{D}) \leq g(R, \mathcal{S}) + \sqrt{\frac{2g(R, \mathcal{D})\left(\omega + \ln\left(\frac{2}{\delta}\right)\right)}{m}} \quad (10)$$

holds for all $R \in \mathcal{R}_k^z$ with probability $\geq 1 - \delta/2$, since $\ln(\delta') \leq \omega + \ln(\frac{2}{\delta})$ (following analogous derivations for the proof of Corollary 4.2). The first set of inequalities is obtained from (10) after adding $\alpha|R|$ on both sides, from the fact that $g(R, \mathcal{D}) \leq \ell(R, \mathcal{D})$, after solving the quadratic inequality $y \leq u + \sqrt{uy}$ w.r.t. y , and after straightforward computations.

We now prove the last inequality and the statement. Let R^* be an arbitrary rule list with $\ell(R^*, \mathcal{D}) = \min_{R \in \mathcal{R}_k^z} \ell(R, \mathcal{D})$. Note that R^* is fixed and independent of the choice of \mathcal{S} . We apply the Chernoff bound to the random variable $Z = mg(R^*, \mathcal{S})$ with $\mathbb{E}_{\mathcal{S}}[Z] = mg(R^*, \mathcal{D})$, obtaining for any $0 < \gamma < 1$

$$\Pr\left(Z \geq (1 + \gamma)\mathbb{E}[Z]\right) \leq \exp\left(-\frac{\mathbb{E}[Z]\gamma^2}{3}\right). \quad (11)$$

Setting the r.h.s. of (11) $\leq \delta/2$, and solving for γ using $\alpha \geq 0$, gives

$$\Pr\left(g(R^*, \mathcal{S}) \geq g(R^*, \mathcal{D}) + \sqrt{\frac{3g(R^*, \mathcal{D}) \ln(2/\delta)}{m}}\right) \leq \delta/2,$$

and, equivalently,

$$g(R^*, \mathcal{S}) \leq g(R^*, \mathcal{D}) + \sqrt{\frac{3g(R^*, \mathcal{D}) \ln(2/\delta)}{m}}$$

with probability $\geq 1 - \delta/2$. We obtain the last inequality of the statement after adding $\alpha|R^*|$ to both sides, and from the fact $g(R^*, \mathcal{D}) \leq \ell(R^*, \mathcal{D})$. From an union bound, all inequalities of the theorem are simultaneously valid with probability $\geq 1 - \delta$, obtaining the statement. \square

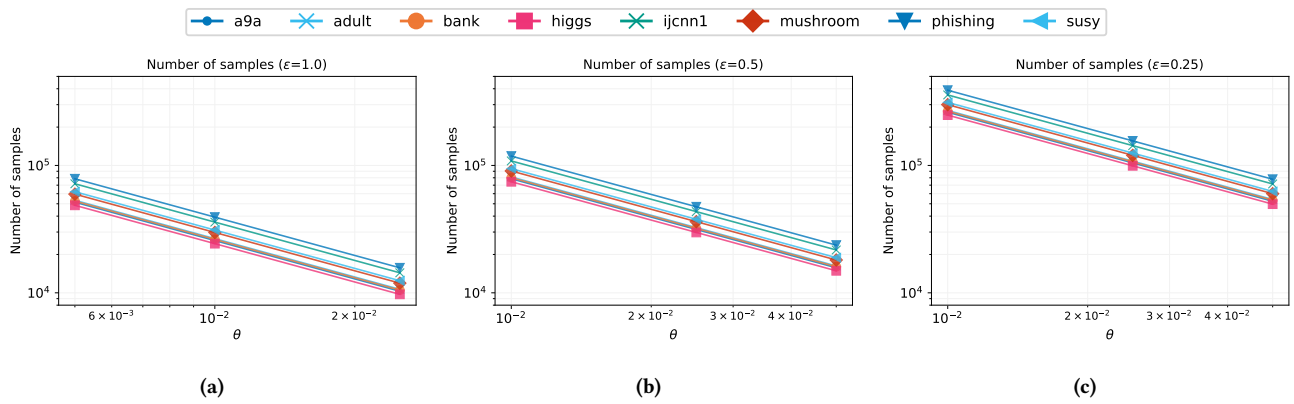


Figure 5: Number of samples $\hat{m}(\mathcal{R}_k^z, \mathcal{D})$ used by SAMRuLE varying ϵ and θ for all datasets. k is set as in Table 1 and $z = 1$.

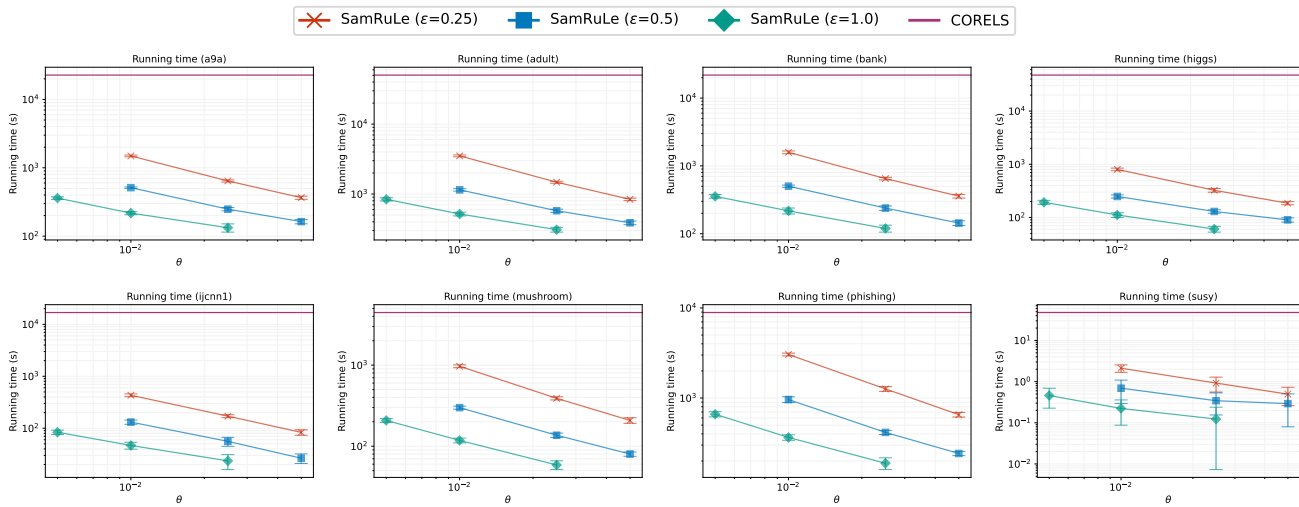


Figure 6: Running times of SAMRuLE and CORELS for different values of ϵ and θ on all datasets.

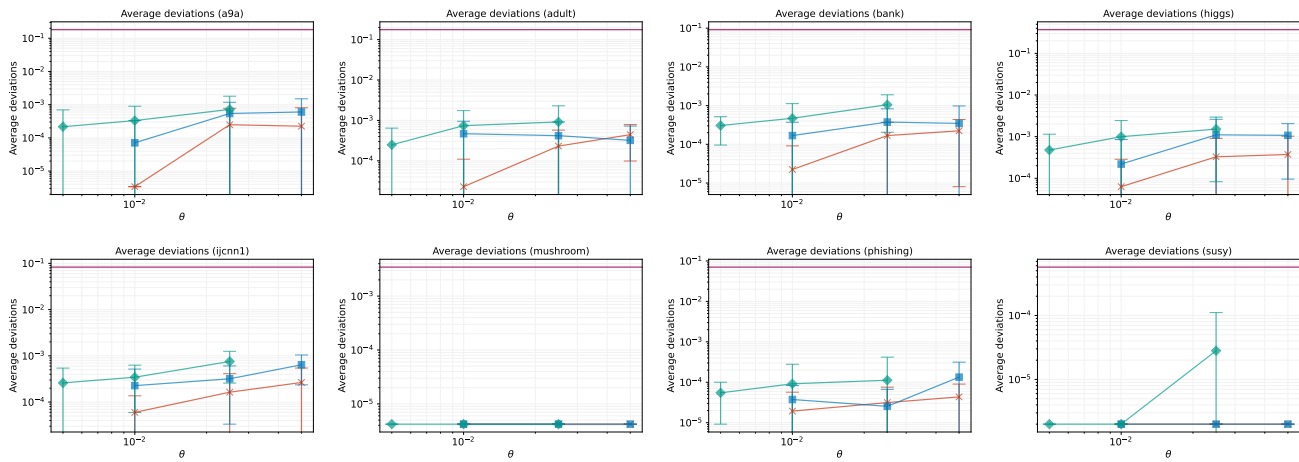


Figure 7: Average deviations $|\ell(\tilde{R}, \mathcal{D}) - \ell(R^*, \mathcal{D})|$ over all runs for different values of ϵ and θ and for all datasets. The purple horizontal line is drawn at $y = \ell(R^*, \mathcal{D})$ (the loss of the optimal rule list R^* found by CORELS). (same legend of Figure 6.)