

IMLI: An Incremental Framework for MaxSAT-Based Learning of Interpretable Classification Rules

Bishwamitra Ghosh
National University of Singapore
Singapore

Kuldeep S. Meel
National University of Singapore
Singapore

ABSTRACT

The wide adoption of machine learning in the critical domains such as medical diagnosis, law, education had propelled the need for interpretable techniques due to the need for end users to understand the reasoning behind decisions due to learning systems. The computational intractability of interpretable learning led practitioners to design heuristic techniques, which fail to provide sound handles to tradeoff accuracy and interpretability.

Motivated by the success of MaxSAT solvers over the past decade, recently MaxSAT-based approach, called MLIC, was proposed that seeks to reduce the problem of learning interpretable rules expressed in Conjunctive Normal Form (CNF) to a MaxSAT query. While MLIC was shown to achieve accuracy similar to that of other state of the art *black-box* classifiers while generating small interpretable CNF formulas, the runtime performance of MLIC is significantly lagging and renders approach unusable in practice. In this context, authors raised the question: *Is it possible to achieve the best of both worlds, i.e., a sound framework for interpretable learning that can take advantage of MaxSAT solvers while scaling to real-world instances?*

In this paper, we take a step towards answering the above question in affirmation. We propose IMLI: an incremental approach to MaxSAT based framework that achieves scalable runtime performance via partition-based training methodology. Extensive experiments on benchmarks arising from UCI repository demonstrate that IMLI achieves up to three orders of magnitude runtime improvement without loss of accuracy and interpretability.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning; Supervised learning; Rule learning.**

KEYWORDS

Interpretable Model, Classification Rules, MaxSAT-based Formulation

ACM Reference Format:

Bishwamitra Ghosh and Kuldeep S. Meel. 2019. IMLI: An Incremental Framework for MaxSAT-Based Learning of Interpretable Classification Rules.

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

AIES '19, January 27–28, 2019, Honolulu, HI, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6324-2/19/01...\$15.00

<https://doi.org/10.1145/3306618.3314283>

In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*, January 27–28, 2019, Honolulu, HI, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3306618.3314283>

1 INTRODUCTION

The recent advances in the machine learning techniques have led autonomous decision making systems be adopted in wide range of domains to perform data-driven decision making. As such the domains range from movie recommendations, ad predictions to legal, medical, and judicial. The diversity of domains mandate different criteria for the machine learning techniques. For domains such as movie recommendations and ad predictions, accuracy is usually the primary objective but for safety critical domains [17] such as medical and legal, interpretability, privacy, and fairness [1] are of paramount importance.

It has been long observed that the interpretable techniques are typically trusted and adopted by decision makers as interpretability provides them understanding of reasoning behind a tool's decision making [21]. At this point, it is important to acknowledge that formalizing interpretability is a major challenge [11] and we do not claim to have final word on this. In this context, it is worth noting that for several domains such as medical domain, which was the motivation for our investigation, decision rules with small number of rules tend to be most interpretable [12].

Since the problem of rule learning is known to be in NP-hard, the earliest efforts focused on heuristic approaches that sought to combine heuristically chosen optimization functions with greedy algorithmic techniques. Recently, there has been surge of effort to achieve balance between accuracy and rule size via principled objective functions and usage of combinatorial optimization techniques such as linear programming (LP) relaxations, sub-modular optimization, or Bayesian methods [2, 4, 14, 15, 27]¹. Motivated by the success of MaxSAT solving over the past decade, Malioutov and Meel proposed a MaxSAT-based approach, called MLIC [13], that provides a precise control of accuracy vs. interpretability. The said approach was shown to provide interpretable Boolean formulas without significant loss of accuracy compared to the state of the art classifiers. MLIC, however, has poor scalability in terms of training time and times out for most instances beyond hundreds of samples. In this context, we ask: *Can we design a MaxSAT-based framework to efficiently construct interpretable rules without loss of accuracy and scaling to large real-world instances?*

The primary contribution of this paper is an affirmative answer to the above question. We first investigate the reason for poor scalability of MLIC and attribute it to large size (i.e., number of clauses) of MaxSAT queries constructed by MLIC. In particular, for training

¹An extensive survey of related work is presented in Appendix A.

data of n samples over m boolean features, MLIC constructs a formula of size $O(n \cdot m \cdot k)$ to construct a k -clause Boolean formula. We empirically observe that the performance of MaxSAT solvers has worse than quadratic degradation in runtime with increase in the size of query. This leads us to propose a novel incremental framework, called IMLI, for learning interpretable rules using MaxSAT. In contrast to MLIC, IMLI makes p queries to MaxSAT solvers with each query of the size $O(\frac{n}{p} \cdot m \cdot k)$. IMLI relies on first partitioning the data into p partitions and then incrementally learning rules on the p partitions in a linear order such that rule learned for the i -th partition not only uses the current partition but regularizes itself with respect to the rules learned from the first $i - 1$ partitions. We conduct a comprehensive experimental study over the large set of benchmarks and show that IMLI significantly improves upon the runtime performance of MLIC by achieving speedup of up to three orders of magnitude. Furthermore, the rules learned by IMLI are significantly small and easy to interpret compared to that of the state of the art classifiers such as RIPPER and MLIC.

Similar to Malioutov and Meel [13], we hope that IMLI will excite researchers in machine learning and CP/SAT (Constraint Programming/Satisfiability) communities to consider this topic further: in designing new MaxSAT-based formulations and in turn designing the MaxSAT solvers tuned for interpretable machine learning.

2 PRELIMINARIES

We use capital boldface letters such as \mathbf{X} to denote matrices while lower boldface letters \mathbf{y} are reserved for vectors/sets. For a matrix \mathbf{X} , \mathbf{X}_i represents the i -th row of \mathbf{X} while for a vector/set \mathbf{y} , y_i represents the i -th element of \mathbf{y} .

Let F be a Boolean formula and $\mathbf{b} = \{b_1, b_2, \dots, b_m\}$ be the set of variables appearing in F . A literal is a variable (b_i) or its complement ($\neg b_i$). A *satisfying assignment* or a *witness* of F is an assignment of variables in \mathbf{b} that makes F evaluate to *true*. If σ is an assignment of variables and $b_i \in \mathbf{b}$, we use $\sigma(b_i)$ to denote the value assigned to b_i in σ . F is in Conjunctive Normal Form (CNF) if $F := C_1 \wedge C_2 \cdots \wedge C_k$, where each clause C_i is represented as disjunction of literals. We use $|C_i|$ to denote the number of literals in C_i . For two vectors \mathbf{u} and \mathbf{v} over propositional variables or constants ($0, 1, \text{true}, \text{false}$ etc.), we define $\mathbf{u} \vee \mathbf{v} = \bigvee_i (u_i \vee v_i)$, where u_i and v_i denote a variable/constant at the i -th index of \mathbf{u} and \mathbf{v} respectively. In this context, note that the operation \wedge between a variable and a constant follows the standard interpretation, i.e., $0 \wedge b = 0$ and $1 \wedge b = b$.

We consider a standard binary classification, where we are given a collection of training samples $\{\mathbf{X}_i, y_i\}$ where each vector $\mathbf{X}_i \in \mathcal{X}$ contains the valuation of the features $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$ for sample i , and $y_i \in \{0, 1\}$ is the binary label for sample i . A classifier \mathcal{R} is a mapping that takes in a feature vector \mathbf{x} and return a class y , i.e., $y = \mathcal{R}(\mathbf{x})$. The goal is not only to design \mathcal{R} to approximate our training set, but also to generalize to unseen samples arising from the same distribution. We define two rules \mathcal{R}_1 and \mathcal{R}_2 to be equivalent if $\forall i, \mathcal{R}_1(\mathbf{X}_i) = \mathcal{R}_2(\mathbf{X}_i)$. In this work, we restrict \mathbf{x} and y to be Boolean (we discuss in Sect. 4.2 that such a restriction can be achieved without loss of generality) and focus on classifiers that can be expressed compactly in CNF. We use $\text{clause}(\mathcal{R}, i)$ to denote the i -th clause of \mathcal{R} . Furthermore, we use $|\mathcal{R}|$ to denote the rule-size of

classifier \mathcal{R} that is the sum of the count of literals in all the clauses, i.e., $|\mathcal{R}| = \sum_i |\text{clause}(\mathcal{R}, i)|$.

In this work, we focus on the weighted variant of CNF wherein a weight function is defined over clauses. For a clause C_i and weight function $W(\cdot)$, we use $W(C_i)$ to denote the weight of clause C_i . We say that a clause C_i is hard if $W(C_i) = \infty$, otherwise C_i is called a soft clause. To avoid notational clutter, we overload $W(\cdot)$ to denote the weight of an assignment or clause, depending on the context. We define weight of an assignment σ as the sum of weight of clauses that σ does not satisfy. Formally, $W(\sigma) = \sum_{i|\sigma \not\models C_i} W(C_i)$.

Given F and weight function $W(\cdot)$, the problem of MaxSAT is to find an assignment σ^* that has the minimum weight, i.e., $\sigma^* = \text{MaxSAT}(F, W)$ if $\forall \sigma \neq \sigma^*, W(\sigma^*) \leq W(\sigma)$. Our formulation will have positive clause weights, hence MaxSAT corresponds to satisfying as many clauses as possible, and picking the strongest clauses among the unsatisfied ones. Borrowing terminology of community focused on developing MaxSAT solvers, we are solving a partial weighted MaxSAT instance wherein we mark all the clauses with ∞ weight as hard and clauses with other positive value less than ∞ weight as soft and ask for a solution that optimizes the partial weighted MaxSAT formula. The knowledge of inner working of MaxSAT solvers and encoding of our representation into weighted MaxSAT is not required for this paper.

3 PROBLEM FORMULATION

Given a training set $\{\mathbf{X}, \mathbf{y}\}$, our goal is to find an interpretable rule that is as accurate as possible. As noted earlier, there are several notions of interpretability. We follow the notion employed in Malioutov and Meel [13], which focuses on the construction of rules involving few clauses each with few literals².

In particular, suppose \mathcal{R} classifies all samples correctly, i.e., $\forall i, y_i = \mathcal{R}(\mathbf{X}_i)$. Among all the rules that classify all samples correctly, we choose \mathcal{R} which is the sparsest (most interpretable) one.

$$\min_{\mathcal{R}} |\mathcal{R}| \text{ such that } \forall i, y_i = \mathcal{R}(\mathbf{X}_i)$$

A classifier rule, however, can not classify all samples correctly. Hence we choose a classifier that makes less prediction error. $\mathcal{E}_{\mathcal{R}}$ is the set of samples which are misclassified by \mathcal{R} , i.e., $\mathcal{E}_{\mathcal{R}} = \{\mathbf{X}_i | y_i \neq \mathcal{R}(\mathbf{X}_i)\}$. Hence we aim to find \mathcal{R} as follows.

$$\min_{\mathcal{R}} |\mathcal{R}| + \lambda |\mathcal{E}_{\mathcal{R}}| \text{ such that } \forall \mathbf{X}_i \notin \mathcal{E}_{\mathcal{R}}, y_i = \mathcal{R}(\mathbf{X}_i)$$

λ is the data fidelity parameter balancing the trade-off between classifier complexity and prediction accuracy. Higher value of λ guarantees less prediction error while sacrificing the sparsity of \mathcal{R} by adding more literals in \mathcal{R} , and vice versa. Therefore λ is an inverse of regularization.

4 IMLI: MAXSAT-BASED INCREMENTAL LEARNING FRAMEWORK OF INTERPRETABLE RULES

In this section, we present the primary contribution of this paper, IMLI, which is a MaxSAT-based incremental learning framework for interpretable classification rules. The core technical idea behind

²An advantage of Malioutov and Meel's formulation is a formal notion of interpretability, which is amenable to formal analysis. We do not wish to claim that Malioutov and Meel's notion is the only formal definition of interpretability.

IMLI is to divide the training data into a fixed number p of partitions and employ MaxSAT based learning framework for each partition such that the MaxSAT query constructed for partition i is based on the training data for partition i and the rule learned until partition $i-1$. To this end, we use the notation $(\mathbf{X}^i, \mathbf{y}^i)$ to refer to the training data for the i -th partition. We assume that $\forall i, |\mathbf{X}^i| = |\mathbf{X}^{i-1}|$.

The rest of the section is organized as follows: we first describe the construction of MaxSAT query for the i -th partition in Sect. 4.1 to learn CNF rules, and then discuss the discretization techniques for real-world datasets in Sect. 4.2. The incrementality of IMLI gives rise to the challenge of having redundant literals in the learned rules; we address such redundancy in Sect. 4.3 and finally we discuss, in Sect. 4.4, how our framework for learning CNF rules can be easily extended to learn DNF rules as well. We provide an illustration of rule learning of IMLI in Appendix B.1.

4.1 Construction of MaxSAT Query

We now discuss the construction of a MaxSAT query, denoted by Q_i , for the i -th partition ($i \in [1, p]$). To construct the MaxSAT query for the i -th partition, we assume an access to the rule learned from the $(i-1)$ -th partition (where \mathcal{R}_0 is an empty formula).

The construction of Q_i takes in four parameters: (i) k , the desired number of clauses in CNF rule, (ii) λ , the data fidelity parameter, (iii) a matrix $\mathbf{X}^i \in \{0, 1\}^{n \times m}$ describing the binary value of m features for each of n samples with \mathbf{X}_q^i being a binary valued vector for the q -th sample corresponding to feature vector $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$, (iv) a label vector $\mathbf{y}^i \in \{0, 1\}^n$ containing a class label y_q^i for the sample \mathbf{X}_q^i . Consequently, IMLI constructs a MaxSAT query for the i -th partition and invokes an off-the-shelf MaxSAT solver to compute the underlying rule \mathcal{R}_i .

IMLI considers two types of propositional variables: (i) feature variables and (ii) noise (classification error) variables. For the i -th partition, IMLI formulates a classifier rule \mathcal{R}_i based on following intuition. Recall, a k -clause CNF rule $\mathcal{R}_i = \bigwedge_{l=1}^k C_l$ is represented as the conjunction of k clauses where clause C_l is the disjunction of feature variables. A sample \mathbf{X}_q^i satisfies C_l if \mathbf{X}_q^i has at least one similar feature whose representative variable is present in C_l . If \mathbf{X}_q^i satisfies $\forall l, C_l$, then $\mathcal{R}_i(\mathbf{X}_q^i) = 1$ otherwise $\mathcal{R}_i(\mathbf{X}_q^i) = 0$. Since feature x_j can be present or not present in each of k clauses, IMLI considers k boolean variables, each denoted by b_j^l ($l \in [1, k]$) for feature x_j to denote its participation in the l -th clause. A sample \mathbf{X}_q^i , however, can be misclassified by \mathcal{R}_i i.e., $\mathcal{R}_i(\mathbf{X}_q^i) \oplus y_q^i = 1$. IMLI introduces a noise variable η_q corresponding to sample \mathbf{X}_q^i so that the assignment of η_q can be interpreted whether \mathbf{X}_q^i is misclassified by \mathcal{R}_i or not. Hence the key idea of IMLI for learning the i -th partition is to define a MaxSAT query over $k \times m + n$ propositional variables, denoted by $\{b_1^1, b_2^1, \dots, b_m^1, \dots, b_m^k, \eta_1, \dots, \eta_n\}$. The MaxSAT query of IMLI consists of the following three sets of constraints:

- (1) Since our objective is to find sparser rules, the default objective of IMLI would be to add a constraint to falsify as many b_j^l as possible. As noted earlier, rule \mathcal{R}_{i-1} from the $(i-1)$ -th partition plays an important role in the construction of MaxSAT constraints of the i -th partition. Therefore, if $x_j \in \text{clause}(\mathcal{R}_{i-1}, l)$, IMLI would deviate from its default

behavior by adding a constraint to keep the corresponding literal *true* in the optimal assignment. The weight corresponding to this clause is 1. We formalize our discussion as follows:

$$V_j^l := \begin{cases} b_j^l & \text{if } x_j \in \text{clause}(\mathcal{R}_{i-1}, l) \\ -b_j^l & \text{otherwise} \end{cases}; \quad W(V_j^l) = 1$$

- (2) We use noise variables to handle mis-classifications and therefore, IMLI tries to falsify as many noise variables as possible. Since data fidelity parameter λ is proportionate to accuracy, IMLI puts λ weight to each following soft clause.

$$N_q := (\neg \eta_q); \quad W(N_q) = \lambda$$

- (3) Let $\mathbf{B}_l = \{b_j^l \mid j \in [1, m]\}$. Here we provide the third set of constraints of IMLI.

$$D_q := (\neg \eta_q \rightarrow (y_q \leftrightarrow \bigwedge_{l=1}^k (\mathbf{X}_q \vee \mathbf{B}_l))); \quad W(D_q) = \infty$$

Every hard clause D_q can be interpreted as follows. If η_q is assigned to *false* ($\neg \eta_q = \text{true}$) then $y_q^i = \mathcal{R}_i(\mathbf{X}_q^i) = \bigwedge_{l=1}^k (\mathbf{X}_q \vee \mathbf{B}_l)$. The operator “ \vee ” is defined in Sect. 2.

Finally, the set of constraints Q_i for the i -th partition constructed by IMLI is defined as follows:

$$Q_i := \bigwedge_{j=1, l=1}^{j=m, l=k} V_j^l \wedge \bigwedge_{q=1}^n N_q \wedge \bigwedge_{q=1}^n D_q$$

Next, we extract \mathcal{R}_i from the solution of Q_i as follows.

CONSTRUCTION 1. Let $\sigma^* = \text{MaxSAT}(Q_i, W)$, then $x_j \in \text{clause}(\mathcal{R}_i, l)$ iff $\sigma^*(b_j^l) = 1$.

In the rest of the manuscript, we will use \mathcal{R} to denote \mathcal{R}_p .

4.2 Beyond Binary Features

We have considered that the feature value of a training sample is binary. Real-world datasets, however, contain categorical, real-valued or numerical features. We use the standard discretization technique to convert categorical and continuous (real or integer value) features to boolean features. We use *one hot encoding* to convert categorical features to binary features by introducing a boolean vector with the cardinality equal to the number of distinct categories of individual categorical features.

EXAMPLE 4.1. Consider a categorical feature with three categories: “red”, “green”, “yellow”. One hot encoding would convert this feature to three binary variables, which take values 100, 010, and 001 for the three categories.

Furthermore, we can discretize the continuous-valued features into binary features by comparing the feature value to a collection of thresholds within range and introducing a boolean feature vector with cardinality proportional to the number of considered thresholds [13]. Specifically, for a continuous feature x_c we consider a number of thresholds $\{\tau_1, \dots, \tau_t\}$ where $\tau_i < \tau_{i+1}$ and define two separate Boolean features $I[x_c \geq \tau_i]$ and $I[x_c < \tau_i]$ for each τ_i . We present the following definitions based on the discretization of continuous features.

DEFINITION 2. $tval(b) : b \rightarrow \tau$ is a function over boolean variables corresponding to discretized binary features (from a continuous feature) and outputs the compared threshold value.

DEFINITION 3. $op(b) : b \rightarrow \{\geq, <\}$ is a function over boolean variables corresponding to discretized binary features (from a continuous feature) and outputs the comparison operator between continuous feature value and $tval(b)$.

DEFINITION 4. $siblings(b_i, b_j) : (b_i, b_j) \rightarrow \{true, false\}$ is a function over pair of boolean variables b_i, b_j and outputs true if the boolean features corresponding to b_i, b_j are constructed by discretizing the same continuous feature and $op(b_i) = op(b_j)$.

EXAMPLE 4.2. Consider a continuous feature x_c with range $(0, 100)$ and three thresholds $\{25, 50, 75\}$ associated with this feature. IMLI introduces 6 new boolean features $\{x_1 : I[x_c \geq 25], x_2 : I[x_c < 50], x_3 : I[x_c \geq 75], x_4 : I[x_c < 25], x_5 : I[x_c < 50], x_6 : I[x_c < 75]\}$. Following this discretization technique, the binary feature vector of a sample with feature value $x_c = 37.5$ is 100011, because among the 6 introduced boolean features $x_1 : I[37.5 \geq 25] = 1, x_5 : I[37.5 < 50] = 1$, and $x_6 : I[37.5 < 75] = 1$.

EXAMPLE 4.3. In Example 4.2, b_i is a boolean variable corresponding to feature x_i . Now $tval(b_1) = 25, op(b_1) = "\geq", siblings(b_1, b_2) = true$, and $siblings(b_1, b_4) = false$.

4.3 Redundancy Removal

Given the incremental procedure of learning \mathcal{R} where the constraints for the i -th partition are influenced from the rule learned until the $(i - 1)$ -th partition, one key challenge is to address potential redundancy in the learned rules. In particular, we observe that redundancy manifests itself in binary features corresponding to continuous-valued features as the $(i - 1)$ -th partition might suggest inclusion of feature $I[x_c < \tau_u]$ while the i -th partition also suggests inclusion of feature $I[x_c < \tau_v]$ where $\tau_u \neq \tau_v$. To this end, we present Algorithm 1 to remove redundant literals.

Algorithm 1 Remove Redundancy

```

1: procedure REMOVEREDUNDANTLITERALS( $\mathcal{R}$ )
2:   for each clause  $C_l$  of  $\mathcal{R}$  do
3:     for each pair  $\langle b_i^l, b_j^l \rangle$  where  $\sigma(b_i^l) = \sigma(b_j^l) = 1$ ,
        $siblings(b_i^l, b_j^l) = true$ , and  $tval(b_i^l) < tval(b_j^l)$  do
4:       if  $op(b_i^l) = op(b_j^l) = "\geq"$  then
5:          $\mathcal{R}' = \mathcal{R}[\sigma(b_j^l) \mapsto 0]$        $\triangleright b_j^l$  is redundant
6:       else
7:          $\mathcal{R}' = \mathcal{R}[\sigma(b_i^l) \mapsto 0]$ 
8:   return  $\mathcal{R}'$ 

```

LEMMA 5. $|\mathcal{R}'| \leq |\mathcal{R}|$ and \mathcal{R}' is equivalent to \mathcal{R} .

PROOF. See Appendix C. □

4.4 Learning DNF Rules

Primarily we focus on learning rule \mathcal{R} which is in CNF form. We can also apply incremental technique for learning DNF rules. Suppose,

we want to learn a rule $y = S(x)$ where $S(x)$ is expressible in DNF. We show that $y = S(x) \leftrightarrow \neg(y = \neg S(x))$. Here $\neg S(x)$ is in CNF. Therefore, to learn DNF rule $S(x)$, we simply call IMLI with $\neg y$ as input for all p batches, learn CNF rule, and finally negate the learned rule. Hence Algorithm 1 can be directly applied.

EXAMPLE 4.4. “(is Male \vee Age < 50) \wedge (Education = Graduate \vee Income ≥ 1500)”—rule is learned for negated class label. The resultant DNF rule is “(is not Male \wedge Age ≥ 50) \vee (Education \neq Graduate \wedge Income < 1500)”

5 EXPERIMENT

We have implemented a prototype implementation in Python to evaluate the performance of IMLI³. The experiment has been conducted on high performance computer cluster, where each node consists of E5-2690 v3 CPU with 24 cores, 96GB of RAM, and in total 130,000 CPU hours. We have conducted an extensive set of experiments on publicly available benchmarks (detailed description in Appendix D.1) from UCI repository [10] to answer the following questions.

- (1) How do the training time and accuracy of IMLI compare to that of state of the art classifiers including both interpretable and non-interpretable ones?
- (2) How do accuracy, rule size, and training time of IMLI vary with data fidelity parameter λ and the number of partitions p ?
- (3) How interpretable are the rules generated by IMLI?

In summary, the experimental results demonstrate that IMLI can scale to large datasets involving tens of thousands of samples with hundreds of binary features. In contrast to MLIC, IMLI achieves up to three orders of magnitude improvement in training time without loss of accuracy and interpretability. IMLI generates rules which are not only interpretable but also accurate compared to other classifiers, which often produce non-interpretable models for the sake of accuracy.

5.1 Experiment Methodology:

To measure the performance gain over MLIC, we measure the accuracy and training time of IMLI vis-a-vis MLIC. We also perform comparisons with another state of the art classifier RIPPER and other (mostly) non-interpretable classifiers such as random forest (RF), support vector classifier (SVC), Nearest Neighbors classifier (NN), l_1 -penalized Logistic Regression (LR).

The number of parameter values is comparable (10) for each technique. For RF and RIPPER, we use control based on the cutoff of the number of examples in the leaf node. For SVC, NN, and LR we discretize the regularization parameter on a logarithmic grid. For both IMLI and MLIC, we have two choices of $\lambda \in \{5, 10\}$, three choices of $k \in \{1, 2, 3\}$, and two choices of the type of rule as {CNF, DNF}. For IMLI we vary the number of partitions p for each dataset such that each partition has at least eight samples and at most 512 samples. For all classifiers, we set the training time cutoff to be 1000 seconds.

We perform an assessment of test accuracy on a holdout set and mean validation accuracy on a 10-fold cross-validation set (holdout

³ <https://github.com/meelgroup/mlic>

set 10%, validation set 9%, training set 81%). We compute test accuracy and mean validation accuracy across the ten folds for each choice of the parameters for each technique, and report test accuracy, mean validation accuracy, and mean training time for a choice of the parameters which incurs the best test accuracy. To remove the bias of a particular holdout set we perform ten repetitions with different holdout sets and present the mean statistics.

For MLIC and IMLI, we experimented with different MaxSAT solvers and finally chose MaxHS [9] for MLIC since MaxSAT queries generated by MLIC timeout for all the solvers and MaxHS is the only solver to return the best found answer so far. In contrast, queries constructed by IMLI are easier and the best runtime performance is obtained by using Open-WBO solver [16].

5.2 Results

5.2.1 Comparison Among Different Classifiers: Table 1 presents the comparison of IMLI vis-a-vis typical interpretable and non-interpretable classifiers. The first three columns list the name, size (number of samples), and the number of binary features (discretized) for each dataset. The next seven columns present test accuracy, validation accuracy, and training time of the classifiers.

In Table 1 we observe that MLIC and RIPPER have slightly higher accuracy than IMLI. Specifically considering all datasets MLIC (resp. RIPPER) has on average 1.12% (resp. 0.12%) higher test accuracy and 3.09% (resp. 2.29%) higher validation accuracy than that of IMLI. In contrast, IMLI takes up to three order of magnitude less training time compared to MLIC and upto one order of magnitude less time compared to RIPPER. Interestingly, IMLI is competitive to black-box classifiers, e.g. SVC and NN for large datasets. In this context, we think IMLI achieves a sweet spot in achieving significant runtime improvement in training without losing accuracy.

At this point, one may wonder as to whether minor loss in accuracy also leads to loss of interpretability. To this end, we illustrate a detailed comparison among the generated rules of IMLI, RIPPER, and MLIC in Table 2. We observe that rule size of IMLI is significantly smaller than that of RIPPER and MLIC. In particular, note that IMLI can generate rules with size less than eight for all the datasets (exception in Adult dataset where IMLI still has the most sparse rule), thereby demonstrating the sparsity of generated rules. In contrast, MLIC and RIPPER generate rules of significantly larger size than IMLI. As indicated earlier, sparsity is only one of several possible approaches to quantify interpretability. Therefore, we also decided to observe the generated rules and interestingly, the generated rules seem very intuitive. We have listed the generated rules in Appendix E

5.2.2 Varying Data Fidelity λ : In Figure 1 we present the result for varying λ . Our experiment result finds a similar observation in all the datasets, and here we present result for Parkinsons dataset.

Recall that size of a rule is the total number of literals appearing in \mathcal{R} . As we increase the value of λ , rule size (Figure 1a) and the time taken to solve the MaxSAT query (Figure 1b) decreases. When $\lambda = 1$, all the soft clauses have equal weight. However, when λ is higher, soft clause N_q is put a higher weight than V_j^l , which turns out in finding the solution of the query requiring less time because of the priority among soft clauses. Therefore, the generated rule becomes sparser. We find a similar trend for DNF rules too. In




Figure 1: Effect of data fidelity parameter λ on rule size and training time. The number within parenthesis denotes the number of clause k for the respective rule.

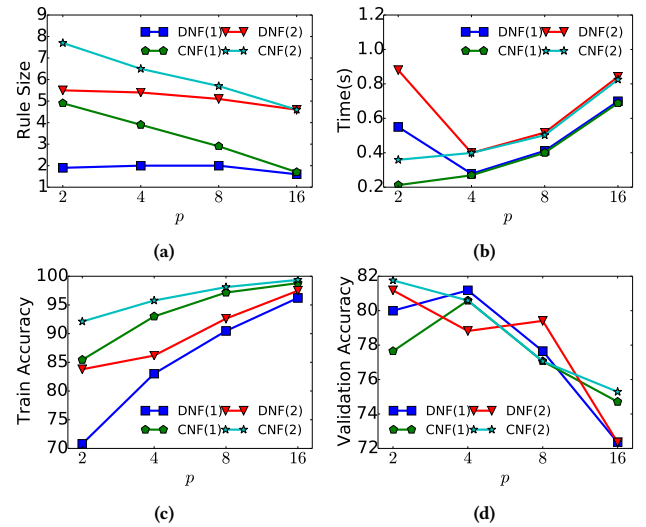


Figure 2: Effect of the number of partitions p on rule size, training time, train accuracy, and validation accuracy ($k \in \{1, 2\}$, $\lambda = 10$). The number within parenthesis denotes the number of clause k for the respective rule.

empirical study we find that as we increase λ , training accuracy increases gradually but validation accuracy and test accuracy do not follow a monotonic behavior in the partition-based learning.

5.2.3 Varying The Number of Training Partitions p : Figure 2 presents the effect on rule size, training time, train accuracy, and validation accuracy as we vary p . For all the datasets we find a similar observation and here present the result for Parkinsons dataset for ease of exposition.

In Figure 2a we observe that the size of the rule decreases as p increases. This observation can be attributed to the decrease in the number of training samples per partition with the increase in the number of partitions and consequently, smaller rules suffice. In Figure 2b IMLI empirically shows that the training time at first decreases significantly and then increases slowly with the increase in p . This observation can be attributed to the combined effect of the number of queries and the size of queries. Initially, we achieve a significant reduction in the size of query while the number of queries eventually dominate the runtime.

Dataset	Size	Features	LR	NN	RF	SVC	RIPPER	MLIC	IMLI
Parkinsons	195	392	97.5 [69.76] (0.22)	90 [82.68] (0.14)	97.5 [80.98] (1.7)	87.5 [83.5] (0.18)	85.00 [76.74] (2.92)	97.5 [82.35] (114.75)	95 [79.41] (0.37)
Ionosphere	351	540	93.06 [90.64] (0.32)	86.11 [87.18] (0.26)	95.83 [92.77] (1.72)	95.83 [90.65] (0.26)	93.06 [85.73] (3.3)	93.06 [91.94] (917.13)	91.67 [85.48] (0.5)
WDBC	569	540	98.28 [96.77] (0.33)	95.69 [97.27] (0.45)	96.55 [96.68] (1.8)	96.55 [97.16] (0.27)	92.24 [93.54] (3.53)	93.97 [95.1] Timeout	89.66 [91.18] (0.78)
Blood	748	64	80 [75.92] (0.2)	76.67 [76.14] (0.2)	76 [76.22] (1.68)	76 [76.22] (0.18)	76 [76.22] (2.23)	75.33 [77.61] (5.96)	76 [76.12] (0.24)
PIMA	768	134	75.32 [74.75] (0.3)	77.92 [73.23] (0.32)	76.62 [75.54] (1.99)	75.32 [76.63] (0.37)	75.32 [74.36] (2.58)	75.97 [71.74] Timeout	73.38 [68.12] (0.74)
Tom's HW	28179	844	96.98 [97.12] (2.24)	94.11 [93.91] (910.36)	97.11 [97.35] (27.11)	96.83 [97.1] (354.15)	96.75 [97.12] (37.81)	96.61 [96.55] Timeout	96.86 [96.23] (23.67)
Adult	32561	262	84.58 [84.99] (5.8)	83.46 [83.62] (640.81)	84.31 [84.68] (36.64)	84.39 [84.69] (918.26)	83.72 [83.49] (37.66)	79.72 [79.53] Timeout	80.84 [77.43] (25.07)
Credit-default	30000	334	80.81 [82.16] (6.87)	79.61 [80.2] (872.97)	80.87 [82.13] (37.72)	80.69 [82.1] (847.93)	80.97 [82.07] (20.37)	80.72 [82.14] Timeout	79.41 [75.81] (32.58)
Twitter	49999	1050	95.67 [96.32] (3.99)	Timeout	95.16 [96.46] (67.83)	Timeout	95.56 [96.16] (98.21)	94.78 [95.69] Timeout	94.69 [95.08] (59.67)

Table 1: Comparisons of classification accuracy with 10-fold cross validation for different classifiers. For every cell in the last seven columns the top value represents the test accuracy (%) on unseen data, the middle value surrounded by square bracket represents average validation accuracy (%) of 10-fold, and the bottom value surrounded by parenthesis represents the average training time in seconds.

Dataset	RIPPER	MLIC	IMLI
Parkinsons	2.6	2	8
Ionosphere	9.6	13	5
WDBC	7.6	14.5	2
Blood	1	3	3.5
Adult	107.55	44.5	28
PIMA	8.25	16	3.5
Tom's HW	30.33	2	2.5
Twitter	21.6	20.5	6
Credit	14.25	6	3

Table 2: Size of the rule of interpretable classifiers.

In Figure 2c we observe that IMLI tends to make less training error as p goes higher because IMLI learns on fewer samples with fixed λ value. Moreover, we observe that CNF rules have higher train accuracy than DNF rules, and 2-clause rules have higher train accuracy than 1-clause rules for both CNF and DNF rules.

In Figure 2d we notice a decrease in validation accuracy as we allow more partitions because learning on fewer samples results in a rule that has less predictive power on validation set. For small p , IMLI, however, ensures higher validation accuracy if the rule has more clauses. Moreover, effect of the number of partitions on test accuracy computed on unseen data does not follow any specific pattern.

In summary, we observe that the number of partitions gives a sound handle to the end user to tradeoff the training time, validation accuracy, and interpretability of the rules.

6 CONCLUSION

In this paper, we present IMLI: an incremental framework for MaxSAT-based learning of interpretable classification rules. Extensive experiments on UCI datasets demonstrate that IMLI achieves up to three orders of magnitude improvement in training time with only a minor loss of accuracy. We think IMLI highlights the promise of MaxSAT-based approach and opens up several interesting directions of future research at the intersection of AI and SAT/SMT community. In particular, it would be an interesting direction of future research if the MaxSAT solvers can be designed to take advantage of incrementality of IMLI.

ACKNOWLEDGMENTS

This work was supported in part by NUS ODPRT Grant, R-252-000-685-133. The computational work for this article was performed on resources of the National Supercomputing Centre, Singapore <https://www.nsc.sg>.

REFERENCES

- [1] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2017. Fairness and Machine Learning. *NIPS Tutorial* (2017).
- [2] D. Bertsimas, A. Chang, and C. Rudin. 2012. An Integer Optimization Approach to Associative Classification. In *Proc. of NIPS*.
- [3] Christian Bessiere, Emmanuel Hebrard, and Barry O’Sullivan. 2009. Minimizing decision tree size as combinatorial optimisation. In *Proc. of CP*.
- [4] E. Boros, P.L. Hammer, T. Ibaraki, A. Kogan, E. Mayoraz, and I. Muchnik. 2000. An implementation of logical analysis of data. *Proc. of TKDE* (2000).
- [5] Gert Cauwenberghs and Tomaso Poggio. 2001. Incremental and decremental support vector machine learning. In *Proc. of NIPS*.
- [6] P. Clark and T. Niblett. 1989. The CN2 Induction Algorithm. *Mach. Learn.* (Mar. 1989).
- [7] William W Cohen. 1995. Fast effective rule induction. In *Machine Learning Proceedings 1995*. Elsevier.
- [8] W. W. Cohen and Y. Singer. 1999. A Simple, Fast, and Effective Rule Learner. In *Proc. of AAAI*. Orlando, FL.
- [9] Jessica Davies and Fahiem Bacchus. 2011. Solving MAXSAT by solving a sequence of simpler SAT instances. In *Proc. of CP*.
- [10] Dua Dheeru and Efi Karra Taniskidou. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [11] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. (2017).
- [12] Benjamin Letham, Cynthia Rudin, Tyler H McCormick, David Madigan, et al. 2015. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics* (2015).
- [13] Dmitry Malioutov and Kuldeep S Meel. 2018. MLIC: A MaxSAT-Based Framework for Learning Interpretable Classification Rules. In *Proc. of CP*.
- [14] D. M. Malioutov and K. R. Varshney. 2013. Exact Rule Learning via Boolean Compressed Sensing. In *Proc. of ICML*.
- [15] M. Marchand and J. Shawe-Taylor. 2002. The set covering machine. *Journal of Machine Learning Research* Dec (2002).
- [16] Ruben Martins, Vasco Manquinho, and Inês Lynce. 2014. Open-WBO: A modular MaxSAT solver. In *Proc. of SAT*.
- [17] Clemens Otte. 2013. Safe and interpretable machine learning: a methodological review. In *Computational intelligence in intelligent data analysis*.
- [18] J Ross Quinlan. 1993. C4. 5: Programming for machine learning. *Morgan Kaufmann* (1993).
- [19] J Ross Quinlan. 2014. *C4. 5: programs for machine learning*. Elsevier.
- [20] Liva Ralaivola and Florence d’Alché Buc. 2001. Incremental support vector machine learning: A local approach. In *Proc. of ICANN*.
- [21] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Why should i trust you?: Explaining the predictions of any classifier. In *Proc. of KDD*.
- [22] Ronald L Rivest. 1987. Learning decision lists. *Machine learning* (1987).
- [23] Stefan Ruppig. 2001. Incremental learning with support vector machines. In *Proc. of ICDM*.
- [24] Guolong Su, Dennis Wei, Kush R Varshney, and Dmitry M Malioutov. 2015. Interpretable two-level Boolean rule learning for classification. (2015).
- [25] Nadeem Ahmed Syed, Syed Huan, Liu Kah, and Kay Sung. 1999. Incremental learning with support vector machines. (1999).
- [26] Fulton Wang and Cynthia Rudin. 2015. Falling rule lists. In *Proc. of AISTATS*.
- [27] T. Wang, C. Rudin, F. Doshi-Velez, Y. Liu, E. Klampfl, and P. MacNeille. 2015. Or’s of And’s for Interpretable Classification, with Application to Context-Aware Recommender Systems. (2015).

A RELATED WORK

The study of designing interpretable machine learning classifiers can find its root in the development of popular models such as decision trees [3, 19], decision lists [22], classification rules [7] etc. Apart from designing models with the purpose of generating interpretable rules, various studies have been conducted in order to improve the efficiency and scalability of the model. Specifically, decision rule approaches such as C4.5 rules [18], CN2[6], RIPPER [7], SLIPPER [8] rely on heuristics, branch pruning, ad-hoc local criteria e.g., maximizing information gain, coverage, etc. because these models consider an intractable combinatorial optimization function.

In recent work, Malioutov et al. has proposed rule based classification system by borrowing ideas from Boolean compressed sensing [14]. Two-level Boolean rules [24] has been proposed to

trade classification accuracy and interpretability, where hamming loss is used to characterize accuracy and sparsity is used to characterize interpretability. Wang et al. [26] has proposed a Bayesian framework for learning falling rule lists which is an ordered list of if-then rules. Chen et al. designs an optimization approach to learning falling rule lists and “softly” falling rule lists, along with Monte-Carlo search algorithms that use bounds on the optimal solution to prune the search space.

Incremental learning techniques are one possible solution to the scalability problem, where data is processed in parts, and the result combined so as to use less memory [25]. Incremental framework has been studied in SVM [23] to improve the existing approach. Specifically, an on-line recursive algorithm for SVM has been studied to facilitate learning one vector at a time [5] and a local incremental approach has been proposed [20] to learn a SVM based on Radial Basis Function Kernel.

B EXAMPLES

B.1 Illustration of Incremental Learning

We illustrate an interpretable rule generated by IMLI with step by step formulation over partitions on iris dataset⁴. Iris dataset has four attributes: sepal length, sepal width, petal length, and petal width. All feature values are scaled in centimeter. Iris dataset has three classes: Iris Setosa, Iris Versicolour, and Iris Virginica. We consider the binary problem of classifying Iris Versicolour from the other two species: Setosa and Verginica. Here we consider that \mathcal{R} is a single clause DNF rule and learned over four partitions (e.g. $\mathcal{R}_1, \dots, \mathcal{R}_4$). The final rule \mathcal{R} is equivalent to \mathcal{R}_4 .

$$\mathcal{R}_1 := \text{petal length} \leq 5.32 \wedge \text{petal length} > 1.7 \wedge \text{petal width} \leq 1.8$$

$$\mathcal{R}_2 := \text{sepal width} \leq 3.1 \wedge \text{petal length} \leq 5.32 \wedge \text{petal length} > 1.7 \wedge \text{petal width} \leq 1.5$$

$$\mathcal{R}_3 := \text{sepal width} \leq 3.1 \wedge \text{petal length} \leq 5.0 \wedge \text{petal length} > 1.7 \wedge \text{petal width} \leq 1.5$$

$$\mathcal{R} := \text{sepal width} \leq 3.1 \wedge \text{petal length} \leq 5.0 \wedge \text{petal length} > 1.7 \wedge \text{petal width} \leq 1.8$$

\mathcal{R} can be interpreted as: a sample which satisfies all of the four constraints is predicted as Iris Versicolour. Here the rule size $|\mathcal{R}| = 4$. Specifically, \mathcal{R}_1 is learned on the first partition of training data. \mathcal{R}_2 learns two literals (2^{nd} and 3^{rd}) which also appear in \mathcal{R}_1 , introduces two new literals (1^{st} and 4^{th}) while learning on the second training partition, and falsifies a previously learned literal from \mathcal{R}_1 (i.e., (petal width ≤ 1.8)).

Since the dataset contains continuous valued features, IMLI removes redundant literals at each step of learning by applying Algorithm 1. For example, we show the generated rule \mathcal{R}'_3 for the 3^{rd} partition if we do not apply Algorithm 1.

$$\mathcal{R}'_3 := \text{sepal width} \leq 3.1 \wedge \text{petal length} \leq 5.0 \wedge \text{petal length} \leq 5.32 \wedge \text{petal length} > 1.7 \wedge \text{petal width} \leq 1.5 \wedge \text{petal width} \leq 1.8$$

⁴<https://archive.ics.uci.edu/ml/datasets/iris>

The underline marked literals are redundant, hence removed in \mathcal{R}_3 .

C PROOF

PROOF. (Lemma 5) The case $|\mathcal{R}'| = |\mathcal{R}|$ is trivial because no literal is removed from \mathcal{R}' , and \mathcal{R}' is equivalent to \mathcal{R} .

When $|\mathcal{R}'| < |\mathcal{R}|$, $\exists \langle b_i^l, b_j^l \rangle$ where $l \in [1, k]$, $siblings(b_i^l, b_j^l) = true$, $tval(b_i^l) < tval(b_j^l)$, and x_c is the considered continuous feature.

Suppose $op(b_i^l) = op(b_j^l) = \geq$. Consider two sets of real number: $S_i = \{x_i : x_i \geq tval(b_i^l)\}$ and $S_j = \{x_j : x_j \geq tval(b_j^l)\}$. As \mathcal{R} is in CNF, \mathcal{R} checks $I[x_c \in S_i \cup S_j]$ to classify $\forall q, X_q$. Here $S_i \cup S_j = S_i$, so b_j^l can be pruned.

We can prove similarly when $op(b_i^l) = op(b_j^l) = <$.

Therefore, $|\mathcal{R}'| \leq |\mathcal{R}|$ and \mathcal{R}' is equivalent to \mathcal{R} . \square

D EXPERIMENT

D.1 Dataset Description

We use nine publicly available datasets of various size from UCI repository for conducting experiments for IMLI. The datasets contain both real and categorical valued features. The datasets are buzz events from two different social networks: Twitter and Tom's HW (Tom), Adult data (Adult), Parkinson's Disease detection dataset (Parkinsons), Ionosphere (Ion), Pima Indians Diabetes (PIMA), Blood service centers (Blood), breast cancer Wisconsin diagnostic (WDBC), and Credit-default approval dataset (Credit-default).

E INTERPRETABLE RULES

In this section we are presenting the rules generated by IMLI for the datasets we use in experiment.

E.1 Rule for Credit Default Dataset:

A client will default if :=
(education=others OR
repayment status September: payment delay > 1 month OR
repayment status August: payment delay > 2 months OR
repayment status June: payment delay > 2 months)

E.2 Rule for Adult Dataset

A person's income is greater than 50k if :=
(workclass is not Federal-gov AND workclass is not State-gov AND
education is not 11th AND education is not 5th-6th AND education
is not 7th-8th AND education-num > 10.0 AND marital-status is
not Divorced AND marital-status is not Married-AF-spouse AND
marital-status is not Married-spouse-absent AND marital-status is
not Never-married AND marital-status is not Separated AND occu-
pation is not Handlers-cleaners AND occupation is not Machine-
op-inspct AND occupation is not Priv-house-serv AND occupation
is not Protective-serv AND relationship is not Own-child AND re-
lationship is not Unmarried AND native-country is not Cambodia
AND native-country is not Columbia AND native-country is not
Dominican-Republic AND native-country is not Guatemala AND
native-country is not Hungary AND native-country is not Jamaica

AND native-country is not Laos AND native-country is not Mex-
ico AND native-country is not Outlying-US,Guam-USVI-etc AND
native-country is not Poland AND native-country is not Vietnam)

E.3 Rule for WDBC Dataset

Tumor is diagnosed as malignant if :=
(standard area of tumor > 38.43 OR
largest perimeter of tumor > 115.9 OR
largest number of concave points of tumor > 0.1508)

E.4 Rule for Blood Transfusion Service Center Dataset

He/she will donate blood if :=
(Months since last donation \leq 4 AND
total number of donations > 3 AND
total donated blood \leq 750.0 c.c. AND
months since first donation \leq 45)

E.5 Rule for Pima Indians Diabetes Database

Tested positive for diabetes if :=
(Plasma glucose concentration > 125 AND
Triceps skin fold thickness \leq 35 mm AND
Diabetes pedigree function > 0.259 AND
Age > 25 years)

E.6 Rule for Parkinson's Disease Dataset

A person has Parkinson's disease if :=
(minimum vocal fundamental frequency \leq 87.57 Hz OR minimum
vocal fundamental frequency > 121.38 Hz OR Shimmer:APQ3 \leq
0.01 OR MDVP:APQ > 0.02 OR D2 \leq 1.93 OR NHR > 0.01 OR HNR
> 26.5 OR spread2 > 0.3)
AND
(Maximum vocal fundamental frequency \leq 200.41 Hz OR HNR \leq
18.8 OR spread2 > 0.18 OR D2 > 2.92)

E.7 Rule for Ionosphere Dataset

A radar is "Good" if:=
($x_1 = 1$ AND $x_2 > 0$ AND $x_4 > 0$ AND $x_5 > -0.23$)
Here "x" represents the set of columns of the dataset.

E.8 Rule for Tom's Hardware Dataset

A topic is popular if :=
(Number of displays at time 2 > 1936 OR Number of displays at
time 7 > 1250.6)

E.9 Rule for Twitter Dataset

A topic is popular if :=
(Number of Created Discussions at time 1 > 78 OR
Attention Level measured with number of authors at time 6 >
0.000365 OR Attention Level measured with number of contri-
butions at time 0 > 0.00014 OR Attention Level measured with
number of contributions at time 1 > 0.000136 OR Number of
Authors at time 0 > 147 OR Average Discussions Length at time
3 > 205.4 OR Average Discussions Length at time 5 > 654.0)