

Interpretable classifiers for tabular data via feature selection and discretization

Reijo Jaakkola, Tomi Janhunen, Antti Kuusisto, Masood Feyzbakhsh Rankooh and Miikka Vilander^{*,†}

Tampere University, Finland

Abstract

We introduce a method for computing immediately human interpretable yet accurate classifiers from tabular data. The classifiers obtained are short Boolean formulas, computed via first discretizing the original data and then using feature selection coupled with a very fast algorithm for producing the best possible Boolean classifier for the setting. We demonstrate the approach via 12 experiments, obtaining results with accuracies comparable to ones obtained via random forests, XGBoost, and existing results for the same datasets in the literature. In most cases, the accuracy of our method is in fact similar to that of the reference methods, even though the main objective of our study is the immediate interpretability of our classifiers. We also prove a new result on the probability that the classifier we obtain from real-life data corresponds to the ideally best classifier with respect to the background distribution the data comes from.

Keywords

Interpretable AI, Boolean logic, Overfitting

1. Introduction

Explainability and human interpretability are becoming an increasingly important part of research on machine learning. In addition to the immediate benefits of explanations and interpretability in scientific contexts, the capacity to provide explanations behind automated decisions has already been widely addressed also on the level of legislation. For example, the European General Data Protection Regulation [1] and California Consumer Privacy Act [2] both refer to the right of individuals to get explanations of automated decisions concerning them.

This article investigates interpretability in the framework of tabular data. Data in tabular form is important for numerous scientific and real-life contexts, being often even regarded as the *most important* data form: see, e.g., [3, 4]. While explainable AI (or XAI) methods custom-made for pictures and text cannot be readily used as such in the setting of tabular data [5], numerous successful XAI methods for the tabular setting exist; see the survey [6] for an overview.

In the current paper, we focus on producing *highly interpretable* classifiers for tabular data. By “interpretable” we here mean that the inner workings of the classifier are directly readable

DAO-XAI 2024: 4th International Workshop on Data meets Applied Ontologies in Explainable AI, October 19, 2024, Santiago de Compostela, Spain

^{*}Corresponding author.

[†]The authors are listed in alphabetical order.

ORCID: 0000-0003-4714-4637 (R. Jaakkola); 0000-0002-2029-7708 (T. Janhunen); 0000-0003-1356-8749 (A. Kuusisto); 0000-0001-5660-3052 (M. F. Rankooh); 0000-0002-7301-939X (M. Vilander)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

from the classifier itself, which contrasts with explaining the operation of an external, black-box classifier [7].¹ For most of our experiments, the classifiers we obtain are *globally* interpretable, meaning that the classifier itself immediately reveals how it works on every possible input. Global interpretability stems from the classifiers being extremely short logical formulas and thereby directly intelligible. Now, globally interpretable classifiers obtained by our method are automatically also *locally* interpretable, meaning that they provide an easy way of explicating why any particular input was classified in a particular way. For a reasonably small part of our experiments, the classifiers obtained are only locally but not quite globally interpretable.

The classifiers we produce are given in the form of short Boolean DNF-formulas which can naturally be conceived as simply Boolean concepts. One of the key issues that makes our approach possible is the surprising power of successful feature selection. Indeed, in our formulas, we use very small numbers of attributes, making them short and thus interpretable. Already in [8], it was observed that by using a single attribute, one could get, for many of the 16 datasets studied in that paper, an accuracy not drastically different from the then state-of-the-art decision trees. We apply this “simplicity first” approach in our method, together with a very rough discretization of numerical attributes to Boolean form. While we focus on interpretability, we obtain also relatively *accurate* classifiers. Indeed, our classifiers perform surprisingly well in comparison to widely recognized methods despite the fact that generally, there exists an obvious trade-off between accuracy and interpretability.

1.1. Overview of our method and contributions

Our method works as follows. We have a tabular dataset S with numerical and categorical attributes X_1, \dots, X_k and a binary target attribute q . Our goal is to produce a classifier for q based on X_1, \dots, X_k . Our method is based on the following steps; see Section 3 for more details.

1. We first discretize the data to Boolean form using a very rough method of chopping numerical predicates at the median.
2. Then, for increasing numbers ℓ , we use feature selection to choose ℓ suitable attributes. After that, we compute the *best possible* Boolean formula for predicting q . By “best possible” we mean the Boolean formula has the least percentage of misclassified points among all Boolean formulas using the ℓ chosen attributes. In other words, the formula has the least empirical error with respect to the 0-1 loss function. The formula is given in DNF-form.
3. We use nested cross-validation to test the accuracy of the formulas and find the most accurate formula using up to 10 attributes (the parameter 10 can be adjusted, but the choice 10 turned out sufficient for our experiments: see Section 3 for a discussion). We look through the sequence of formulas obtained above for different numbers ℓ of features. Among the formulas with accuracy within one percentage point of the most accurate formula in the sequence, we select the one with the least number of features. We use the

¹However, our method could of course also be used to explicate existing black-box models in a model agnostic way via first generating data with the black-box model to be explained and then using the method to produce a closely corresponding interpretable classifier.

attributes in this formula to train the final formula via the entire training data. Finally, we slightly simplify the final obtained formula—keeping it in DNF-form—via using a standard method from the literature.

To demonstrate the robustness of our method, the discretization step is performed very roughly. Also the feature selection procedures are not critical to our approach; we use three readily available ones and choose the best final formula.

Concerning step 2, for computing the best possible DNF-formula, we present a very efficient algorithm running in time $\mathcal{O}(|W||\tau||T|)$, where $|W|$ is the number of rows in the data, $|\tau|$ the number of *selected* features and $|T| \leq \min(|W|, 2^{|\tau|})$ the number of different *row types* realized in the data with the selected features. The algorithm is fixed-parameter linear when $|\tau|$ is bounded by a constant—which it is in our scenario. If $|\tau|$ was not constant, the algorithm would have running time $\mathcal{O}(|W|^2)$. Now, all our tests ran fast and smoothly on a laptop even with large datasets containing up to 423 680 rows. For six of the datasets, the runs took less than 30 seconds per split. Of the remaining datasets, five took less than 12 minutes per split. Even the two largest datasets had a runtime of less than an hour. The hyperparameter optimizations of random forests and XGBoost took systematically longer than running our method. For example, for the largest dataset, the hyperparameter optimization for XGBoost took roughly two and half hours per split.

As already indicated, we test our method on 12 tabular datasets: seven binary classification tasks from the UCI machine learning repository; a high-dimensional binary classification task with biological data; and four benchmark sets from [9]. See Section 4 for a detailed list. In each test, we compare the accuracy of our classifier to a result obtained by state-of-the-art classifiers. Now, taking into account that tabular data is still a challenge to deep learning methods [9, 4], we use both XGBoost [10] and Sklearn’s implementation of random forests [11] as reference methods. These models are widely recognized for their performance on tabular data [3, 9]. For the UCI datasets we additionally compare to the method of [12] based on the length of formulas. This method is computationally too inefficient to use on the other datasets we consider. In addition to the reference methods, in relation to 10 of the experiments, we also report results found in the literature for the used datasets.

For 11 datasets we use standard ten-fold cross-validation and report the average accuracy as well as the standard deviation over the ten splits. For the high-dimensional dataset Colon with less than 100 data points we instead use leave-one-out cross-validation.

All formulas encountered in our experiments are small enough to be easily *locally* interpretable. In the local interpretation process, the interpreter has a row of classified data which is then compared to the conjunctions of the DNF-classifier our method produces. A positively classified row will match with precisely one conjunction of the DNF-classifier, and a negatively interpreted one will clash with at least one attribute of each conjunction.

In addition to local interpretability, most of the formulas produced in the experiments are in fact so simple that they can be readily *globally* interpreted, meaning that their behaviour on any input is clear simply by the form of the formula. Globally interpreting a short DNF-formula involves looking at each of the few conjunctions separately and interpreting the meaning of their particular combination of attributes. The behaviour of the entire formula is then given by the fact that it accepts the cases given by the conjunctions and rejects everything else.

As a concrete example of an interpretable classifier, we get the formula

$$\neg p_1 \vee (\neg p_2 \wedge \neg p_3)$$

from an experiment on a Breast Cancer dataset concerning the benignity of breast tumors. The attributes p_1 , p_2 and p_3 relate to measures of uniformity of cell size, bare nuclei and bland chromatin. The accuracy of this formula on the corresponding test data is 94.1 percent. The average accuracy of our method over ten splits of the Breast Cancer data is 95.9 percent, while XGBoost obtains 97.1 percent and random forests 97.4 percent. We stress that our formula is, indeed, highly interpretable, while the classifiers obtained by XGBoost and random forests are of very large and of a black box nature.

As another example, from a Colon dataset we get the formula

$$p_1 \vee p_2$$

from the majority of our tests. The attributes p_1 and p_2 have been selected from a total of 5997 Booleanized attributes related to genes. The accuracy of our method on this dataset is 80.6 percent, while random forests get 83.9 percent and XGBoost 72.6 percent.

In relation to accuracy, the experiments are summarized in Figure 1. In Table 1 we report the average number of attributes (features) used by the final formulas and the dimensions of the datasets.

In general, our method has similar accuracy to the reference methods. For the UCI datasets reported on the left in Figure 1, our method obtains accuracies on par with and sometimes better than the state-of-the-art reference methods. For example, on the Hepatitis dataset we obtain an average accuracy of 80.7 percent compared to the 78.2 percent of random forests, 79 percent of XGBoost and 79.4 percent of the formula-size method of [12]. The four benchmark datasets from [9] reported on the right in Figure 1 were the most difficult for our method, but even for these we obtained some surprisingly accurate and reasonably interpretable formulas. For example we obtain a formula of the form

$$\begin{aligned} & (p_1 \wedge p_2 \wedge p_3 \wedge \neg p_4) \vee (p_2 \wedge \neg p_1 \wedge \neg p_5 \wedge \neg p_4) \\ & \vee (p_4 \wedge p_3 \wedge \neg p_1 \wedge \neg p_5 \wedge \neg p_2) \end{aligned}$$

from an experiment on the RoadSafety dataset. This formula is quite short and has a test accuracy of 73.2 percent compared to the average 83.2 percent achieved by black box classifiers.

Concerning our experiments, we stress once more that the main advantage of our method is interpretability. As Table 1 indicates, most of our classifiers use very small numbers of features, thus being interpretable. Some classifiers obtained in the tests are longer, but still within the bounds of local (while not necessarily global) interpretability.

Finally, a key insight behind our method is the idea of computing the ideal classifiers (i.e., the above mentioned best possible Boolean classifiers) and the fact that this can be done fast using the $\mathcal{O}(|W||\tau||T|)$ algorithm (where $|T| \leq \min(|W|, 2^{|\tau|})$) when τ is small. Since small τ is often sufficient—which is another key insight in the method—the approach indeed works quite accurately and fast. The ideal classifiers being central to the approach, we call the method *the ideal classifier method*. A possible alternative for this is would be the *ideal DNF-method*.

In addition to experiments, we also prove a novel theoretical sample bound result that can be used for estimating whether a Boolean classifier obtained from data is in fact an ideal Boolean classifier with respect to the background distribution the data comes from. See Section 3.1 for the related theorem and the Appendix for the proof. Our result is in flavour similar to the various results in statistics that estimate the sizes of samples needed for obtaining a given confidence interval; see, e.g., [13] for further details. Results of this form can indeed be useful for estimating if the classifiers we obtain from datasets are in fact best possible classifiers (for the given features) with respect to the underlying probability distribution the data originates from. In Section 3.1 we illustrate how our result can be used in practice in relation to the datasets of the current paper.

1.2. Further related work

Concerning further related work, while the literature on explainability is rather extensive, only a relatively small part of it is primarily based on logic. See [14] for a survey on logic-based explainability, called *formal explainable AI* (or FXAI) there. We mention here the two prominent works dealing with minimality notions for Boolean classifiers and pioneering much of the recent work in FXAI, [15] and [16].

Like the articles [15] and [16], in fact most of logic-based explainability differs significantly from the current paper. Firstly, most papers in the field concern local rather than global explanations, and also inherent interpretability (as opposed to explainability of existing classifiers) is rarely the main focus. However, there is one method that is close enough to ours to require an explicit and direct analysis in the current paper. That method—let us here call it the *formula-size method* (FSM)—is investigated in [12]. Just like the current work, FSM uses a validation-based approach to avoid overfitting and find Boolean formulas with a small error over real-life datasets. A major difference between our approach and FSM is that in our method, we investigate increasing *numbers of features* as opposed to increasing *length bounds* on formulas as in FSM. The algorithm of FSM searches through the space of all possible formulas of increasing lengths, making it impossible to use with larger datasets. This contrasts with the fixed-parameter linear $\mathcal{O}(|W|2^{|\tau|})$ algorithm we use with $|\tau|$ being a constant.² Our algorithm outputs classifiers with the minimum error in relation to the set of input features used, whereas FSM optimizes with respect to formula size. Together with other experiments, we also describe in Section 4.2 tests we ran to compare FSM with our method.

Concerning yet further related work on interpretable AI, the articles [17, 18] investigate the use of sparse rule lists and scoring systems which are optimized with respect to their error and size. These models are sparse in the sense that they try to use a small number of features, which makes them interpretable. The empirical results reported in these papers also demonstrate the surprising effectiveness of these interpretable models on real-world tabular data. Using the methods proposed in the articles [17, 18] requires—as in the case of FSM in [12]—solving very resource-consuming combinatorial optimization problems, since in addition to their error, classifiers are also optimized with respect to their size. For yet further related work on interpretable AI, see [7].

²In our experiments we constrain $|\tau|$ to be at most 10.

2. Preliminaries

A **vocabulary** is a finite set of symbols p_i referred to as **proposition symbols**. For a vocabulary $\sigma = \{p_1, \dots, p_k\}$ the syntax of propositional logic $\text{PL}[\sigma]$ over σ is given by the grammar $\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$, where $p \in \sigma$. We also define the exclusive or $\varphi \oplus \psi := (\varphi \vee \psi) \wedge \neg(\varphi \wedge \psi)$ as an abbreviation. A formula $\varphi \in \text{PL}[\sigma]$ is in **disjunctive normal form** (DNF) if $\varphi = \bigvee_{i=1}^m \psi_i$, where each ψ_i is a conjunction of **literals** (i.e., formulas p or $\neg p$ where $p \in \sigma$).

A σ -model is a structure $M = (W, V)$, where W is a finite non-empty set called the **domain** of M and $V : \sigma \rightarrow \mathcal{P}(W)$ is a **valuation** which assigns to each $p \in \sigma$ the set of points $V(p) \subseteq W$ where p is true. Such a valuation extends in the usual way into a valuation $V : \text{PL}[\sigma] \rightarrow \mathcal{P}(W)$, with \wedge, \vee and \neg corresponding to the intersection, union and complementation (with respect to W) operations. A formula $\varphi \in \text{PL}[\sigma]$ is true in the point $w \in W$ of a σ -model $M = (W, V)$, denoted $M, w \models \varphi$, if $w \in V(\varphi)$. Note that thereby each formula φ corresponds to a subset $V(\varphi)$ of W . For $\varphi, \psi \in \text{PL}[\tau]$, we define $\varphi \models \psi$ if for all models M and all $w \in W$, $M, w \models \varphi$ implies $M, w \models \psi$.

A σ -**type** t is a conjunction such that for each $p \in \sigma$, precisely one of the literals p and $\neg p$ is a conjunct of t and t has no other conjuncts. We assume some standard bracketing and ordering of literals so that there are exactly $2^{|\sigma|}$ σ -types. We denote the set of σ -types by T_σ . Note that each point $w \in W$ of a σ -model $M = (W, V)$ satisfies exactly one σ -type. Thus, σ -types induce a partition of the domain W . On the other hand, from the truth table of a formula $\varphi \in \text{PL}[\sigma]$ one can obtain an equivalent formula ψ that is a disjunction of types and thus in DNF.

For a vocabulary τ , let $\mu : T_{\tau \cup \{q\}} \rightarrow [0, 1]$ be a probability distribution. Here q is the separate target attribute of the classification task. For $t \in T_\tau$, we define $\mu(t) = \mu(t \wedge q) + \mu(t \wedge \neg q)$. The distribution μ corresponds to the real-world phenomenon that gives rise to practical data. Thus, we generally assume that μ is unknown and define the **true error** (or **risk**) of a formula $\varphi \in \text{PL}[\tau]$ with respect to μ as

$$\text{err}_\mu(\varphi) := \Pr_{t \sim \mu}[t \models \varphi \oplus q] = \sum_{\substack{t \in T_{\tau \cup \{q\}} \\ t \models \varphi \oplus q}} \mu(t).$$

This is the probability that φ disagrees with q . Our goal is to obtain formulas with a small true error. This is made difficult by the fact that μ is unknown. We can, however, estimate the true error via available data.

Let $M = (W, V)$ be a $\tau \cup \{q\}$ -model. For us, M corresponds to the available tabular data. Given a propositional formula $\varphi \in \text{PL}[\tau]$, we define the **empirical error** (or **empirical risk**) of φ with respect to M as

$$\text{err}_M(\varphi) := \frac{|V(\varphi \oplus q)|}{|W|}.$$

The empirical error $\text{err}_M(\varphi)$ is easily computable as the proportion of points where φ disagrees with q . If M is fairly sampled from μ , then by the law of large numbers $\text{err}_M(\varphi) \rightarrow \text{err}_\mu(\varphi)$ almost surely when $|W| \rightarrow \infty$.

Given a distribution μ , the formula

$$\varphi_{\text{id}}^\mu := \bigvee \left\{ t \in T_\tau \mid \frac{\mu(t \wedge q)}{\mu(t)} \geq 1/2 \right\},$$

which we call the **ideal classifier**, has the smallest true error with respect to μ among the formulas in $\text{PL}[\tau]$. This is a syntactic, logic-based representation of what is known as the **Bayes classifier** in the literature [19], not to be confused with naive Bayesian classifiers. A Bayes classifier always gives the best possible prediction, and clearly so does an ideal classifier. Now, as μ is unknown, the ideal classifier is again a theoretical goal for us to approximate.

Given a $\tau \cup \{q\}$ -model $M = (W, V)$, the formula

$$\varphi_{\text{id}}^M := \bigvee \left\{ t \in T_\tau \mid \frac{|V(t \wedge q)|}{|V(t)|} \geq 1/2 \right\},$$

which we call the **empirical ideal classifier**, has the smallest empirical error with respect to M among the formulas in $\text{PL}[\tau]$. This formula is easily computable and an essential tool of our study.

3. Feature selection and overfitting

Our general goal is to use, for a suitable set τ of attributes, the empirical ideal classifier to approximate the ideal classifier. In this section, we specify our methodology and show bounds on sufficient sample size to guarantee that the two classifiers are identical with high probability.

Let τ be a small set of promising attributes chosen from the initially possibly large set of all attributes. We describe a quadratic time algorithm to obtain the empirical ideal classifier. The pseudocode **Algorithm 1** below describes a formal implementation of this algorithm. Basically, we scan the points $w \in W$ once. For every τ -type t that is realized in $M = (W, V)$, we initiate and maintain two counters, n_t and c_t . The first counter n_t counts how many times t is realized in M , while c_t counts how many times $t \wedge q$ is realized in M . The number c_t/n_t is then the probability $|V(t \wedge q)|/|V(t)|$. The empirical ideal classifier φ_{id}^M can be constructed by taking a disjunction over all the types t which are realized in M and for which $c_t/n_t \geq 1/2$.

Algorithm 1 Compute the ideal classifier φ_{id}^M

Input: a $(\tau \cup \{q\})$ -model $M = (W, V)$

```

1:  $T_M \leftarrow \emptyset$  {All the  $\tau$ -types realized in  $M$  will be stored in the set  $T_M$ }
2: for  $w \in W$  do
3:    $t \leftarrow$  the  $\tau$ -type of  $w$ 
4:   if  $t \notin T_M$  then
5:      $T_M \leftarrow T_M \cup \{t\}$ 
6:      $n_t, c_t \leftarrow 0, 0$ 
7:      $n_t \leftarrow n_t + 1$ 
8:     if  $w \in V(q)$  then
9:        $c_t \leftarrow c_t + 1$ 
10:  $\varphi_{\text{id}}^M \leftarrow \perp$ 
11: for  $t \in T_M$  do
12:   if  $c_t/n_t \geq 1/2$  then
13:      $\varphi_{\text{id}}^M \leftarrow \varphi_{\text{id}}^M \vee t$ 
14: return  $\varphi_{\text{id}}^M$ 

```

It is clear that this algorithm runs in polynomial time with respect to the size of M , the size being $\mathcal{O}(|W||\tau|)$. A more precise analysis shows that the running time of this algorithm is $\mathcal{O}(|W||T_M||\tau|)$, where $|T_M|$ counts the number of τ -types that are realized in M . Since $|T_M| \leq |W|$, this gives a worst case time complexity of $\mathcal{O}(|W|^2|\tau|)$. If τ has a fixed size bound, as in our experiments below, then this reduces to a linear time algorithm. Moreover, we note that clearly the size $|\varphi_{id}^M|$ is $\mathcal{O}(|W||\tau|)$.

A full step-by-step description of our method follows:

1. We begin with a tabular dataset with features X_1, \dots, X_m, q , where q is a Boolean target attribute. We denote this full training data by W_0 .
2. We randomly separate 30 percent of W_0 as validation data W_{val} . The remaining part we call the training data W_{train} . We then discretize (or Booleanize) both of these datasets, ending up with tabular datasets with strictly Boolean attributes p_1, \dots, p_k, q . To this end, we simply chop the numerical attributes at the median value of the training data W_{train} , above median meaning “yes” and at most median corresponding to “no”. For categorical attributes we use one-hot encoding.
3. We iterate steps 4 and 5 for increasing numbers ℓ of features from 1 to 10.
4. We run three feature selection procedures on the training data, each selecting ℓ of the attributes p_1, \dots, p_k to be used for classification.
5. Suppose one procedure selected the set $F = \{p_{i_1}, \dots, p_{i_\ell}\}$ of features. We use Algorithm 1 to compute the empirical ideal classifier for q on the data W_{train} using the attributes in F . Note that this is the *best possible* classifier for q over F , given in DNF-form. We do this step for all three sets of features given by the procedures. We select the formula with the highest validation accuracy r on the data W_{val} and record the tuple (ℓ, F, r) for future steps.
6. Once step 5 has halted, we look back at the tuples (ℓ, F, r) recorded. We select the feature set F_ℓ corresponding to the smallest number ℓ with validation accuracy r_ℓ within one percentage point of the best accuracy in the full sequence

$$(1, F_1, r_1), \dots, (10, F_{10}, r_{10}).$$

We run the Booleanization again using the median of the full training data W_0 and compute the best possible classifier one last time using the selected set F_ℓ of features and the data W_0 .

7. Finally, we simplify the selected DNF-formula via `simplify_logic` from SymPy. This gives a potentially simpler logically equivalent DNF-formula.

To demonstrate the robustness of our method, the discretization steps are performed very roughly, using simply the medians. Also the specific feature selection procedures are not in any

way custom-made for our method; we use three readily available ones (see Sect. 4 for details) and use the one with the best accuracy.

Regarding step 5, choosing the formula with the best validation accuracy is done to avoid overfitting to the training data. With the increase of the number of features, the validation accuracy generally first improves and at some point starts to decline. This decline is a sign of overfitting. As overfitting often occurs already for small numbers of selected attributes, we may regard overfitting as *useful* for the method, leading to *short* formulas. However, in some cases, the accuracy stagnates and the overfitting point seems hard to find. In these cases, we nevertheless stop at ten features, the maximum considered in the method. For some datasets it could be necessary to go further before the accuracy stagnates, but for our experiments, ten features was (more than) enough.

While the last formula obtained in step 5 can be quite long, step 6 helps to obtain shorter formulas. By choosing an earlier formula in the obtained sequence with almost the same accuracy, we can reduce the number of features used, drastically improving the interpretability of our formulas without sacrificing much accuracy.

The formula obtained from step 6 is often very short and in most cases even globally interpretable. Nevertheless, as a final step, a standard formula simplifying tool `simplify_logic` from SymPy is used, and this gives a potentially even shorter DNF-formula. We note that even without using `simplify_logic`, all formulas encountered in our experiments are readily *locally* interpretable. Recall from the Introduction that in the local interpretation process, a positively classified input will match with precisely one conjunction of the DNF-classifier, and a negatively interpreted one will clash with at least one attribute of each conjunction.

3.1. Bounds on sample size

As our method consists of using the empirical ideal classifier as an approximation of the ideal classifier, we would like to have some guarantees on when the two classifiers are the same. We next present a theorem which tells us how large samples we need in order to be confident that the empirical ideal classifier is also the true ideal classifier.

The lower bound provided by our theorem depends in a crucial way on how “difficult” the underlying distribution μ is. More formally, we say that a probability distribution $\mu : T_{\tau \cup \{q\}} \rightarrow [0, 1]$ ε -**separates** $t \in T_\tau$, if we have that $|\mu(t \wedge q) - \mu(t \wedge \neg q)| \geq \varepsilon$. The larger the ε is, the easier it is to detect via sampling which of the types $t \wedge q$ and $t \wedge \neg q$ has the higher probability of occurring.

The formal statement of the theorem is now as follows. See Appendix 6.1 for the proof.

Theorem 3.1. *Fix a vocabulary τ , a proposition symbol $q \notin \tau$ and a probability distribution $\mu : T_{\tau \cup \{q\}} \rightarrow [0, 1]$. Let $\varepsilon, \delta > 0$ and*

$$n \geq \frac{2 \ln(2^{|\tau|+1}/\delta)}{\varepsilon^2}.$$

Then with probability at least $1 - \delta$, the empirical ideal classifier with respect to a sample M of size n agrees with the ideal classifier with respect to μ on every $t \in T_\tau$ which is ε -separated by μ . In particular, if μ ε -separates every $t \in T_\tau$, then the empirical ideal classifier is the ideal classifier with probability at least $1 - \delta$.

Corollary 3.2. Fix a vocabulary τ , a proposition symbol $q \notin \tau$ and a probability distribution $\mu : T_{\tau \cup \{q\}} \rightarrow [0, 1]$. Let $\varepsilon, \delta > 0$ and

$$n \geq \frac{2^{2|\tau|+1} \ln(2^{|\tau|+1}/\delta)}{\varepsilon^2}.$$

Then with probability at least $1 - \delta$, we have that

$$\text{err}_\mu(\varphi_{id}^M) < \text{err}_\mu(\varphi_{id}^\mu) + \varepsilon.$$

Proof. Fix $\varepsilon, \delta > 0$. Given $\eta > 0$ we know that if φ_{id}^M agrees with φ_{id}^μ on every $t \in T_\tau$ which is η -separated by μ , then

$$\begin{aligned} \text{err}_\mu(\varphi_{id}^\mu) - \text{err}_\mu(\varphi_{id}^M) &= \sum_{\substack{t \in T_\tau \\ \mu \text{ does not } \eta\text{-separate } t}} |\mu(t \wedge q) - \mu(t \wedge \neg q)| \\ &< |T_\tau| \eta. \end{aligned}$$

Setting $\eta := \varepsilon/|T_\tau|$ and applying Theorem 3.1 gives us the desired result. \square

To illustrate the use of this latter bound, suppose that $|\tau| = 3$, $\delta = 0.01$ and $\varepsilon = 0.05$. Corollary 3.2 shows that we need a sample M of size at least 18887 to know that with probability at least 0.99 the theoretical error of φ_{id}^M is less than that of φ_{id}^μ plus 0.05. The datasets that we consider in Section 4 contain three datasets (Covertypes, Electricity, RoadSafety) that have at least 18887 data points. Thus we can be fairly confident that for these datasets any empirical ideal classifier that uses at most three features obtains an accuracy which is similar to that of the true ideal classifier on those features.

4. Experiments

4.1. Experimental setup

We compared our method empirically to random forests and XGBoost. These two comparison methods are very commonly used and state-of-the-art for tabular data. We tested all three methods on 12 tabular datasets. The datasets can be categorized as follows.

1. 7 binary classification tasks from the UCI machine learning repository. Five of these were selected arbitrarily, while two further ones (BreastCancer and GermanCredit) were randomly chosen among the ones used in [12]. One of the 7 datasets (StudentDropout) was originally a ternary classification task; we converted it into a binary one.
2. 4 tabular-data benchmarks out of 7 binary classification benchmarks that were presented in the paper [9]. These datasets were also originally from the UCI machine learning repository.
3. a high-dimensional binary classification task containing biological data from an open-source feature selection repository presented in [20].

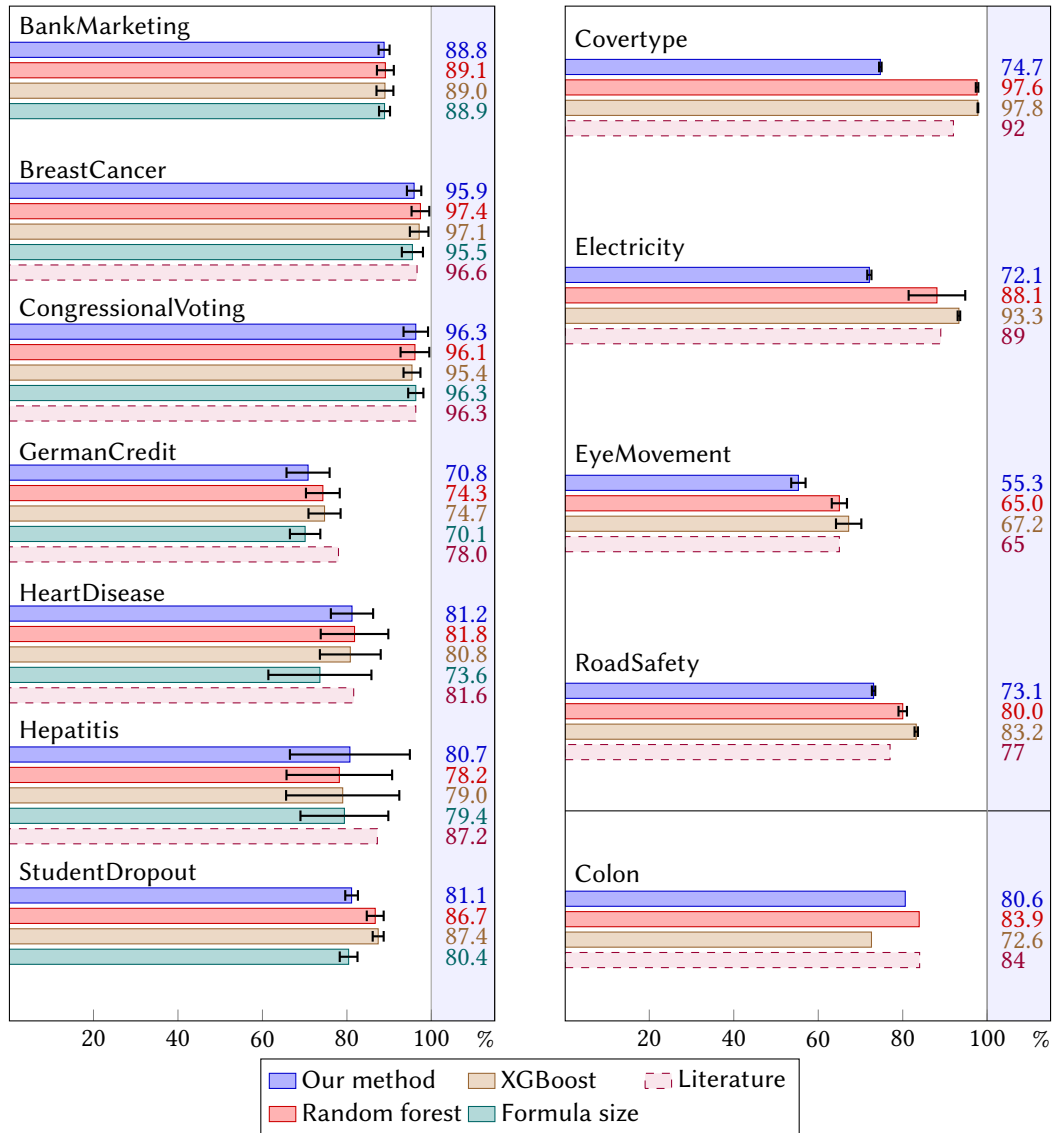


Figure 1: The average test accuracies obtained with our method, random forests and XGBoost for all datasets. For all but the Colon dataset, also standard deviations are reported. For the UCI datasets we also include a comparison to the formula size method. When available, we have also included accuracies reported in literature, though these are not directly comparable due to different technical particularities in the experiments.

All rows with missing values were removed from these datasets.³ Most of the datasets contained both categorical and numerical features. While our method includes a rough median-based Booleanization of the data as discussed in Section 3, random forests and XGBoost used the original non-Booleanized data.

³In the case of the Hepatitis dataset we removed two columns that had several missing values.

For all but one of the datasets we use ten-fold cross-validation for all methods. That is, we split the data into ten equal parts and for each such 10 percent part we input the remaining 90 percent into the method being tested, obtain a classifier (using the chosen 90 percent of the data) and record its accuracy on the 10 percent part. We report the averages and standard deviations of the accuracies over the ten 90/10 splits.

For the high-dimensional dataset Colon with less than 100 data points, we use leave-one-out cross-validation. That is, for each data point, we set the point aside and input the remaining data into the method. We test whether the obtained classifier classifies the omitted point correctly or not. We report the average accuracy over all data points. This is a well-known standard practice for small datasets.

As mentioned in Section 3, our method utilizes readily available feature selection methods. We used Scikit-learn’s `SelectKBest`, which returns k features that have the highest scores based on a given univariate statistical test. `SelectKBest` supports three methods for calculating the scores for the features: F-test, mutual information, and χ^2 -test. When generating formulas, we tested all three of these methods and selected the best in terms of validation accuracy. We emphasize that feature selection was performed after Booleanizing the data.

If the empirical ideal classifier is allowed to use too many features, it will most likely overfit on its training data. As already discussed, we used nested cross-validation to determine how many features the empirical ideal classifier can use without overfitting. For nested cross-validation we used a 70/30-split. In all of the experiments we allowed the empirical ideal classifier to use at most ten attributes. This is a hyperparameter one could optimize for each dataset separately. For simplicity we used the same limit for all datasets. In all of our experiments, the validation accuracy either declined or stagnated well before ten features, indicating that a larger number is not needed. On the other hand, for interpretability, ten features is perhaps even a bit excessive.

For random forests we used Scikit-learn’s implementation. For both random forests and XGBoost, nested cross-validation was used for tuning the hyperparameters. As for our method, we used a 70/30-split. For hyperparameter optimization, we used Optuna [21] with the number of trials being 100. The hyperparameter spaces used for random forests and XGBoost were the same as in [9]. For the readers convenience, we also report the used hyperparameter spaces in Appendix 6.2.

We also ran the method of [12] for some datasets. As in the Introduction, we refer to their approach as the *formula-size method*. This method is computationally resource consuming, so we only ran it for the seven UCI repository datasets. The high-dimensional biological data and benchmark data would be too difficult for the method.

For the formula-size method we used the openly available implementation from <https://github.com/asptools/benchmarks>. The runs were conducted on a Linux cluster featuring Intel Xeon 2.40 GHz CPUs with 8 CPU cores per run, employing a timeout of 72 hours per instance and a memory limit of 64 GB.

4.2. Results

Figure 1 contains the average test accuracies and standard deviations that were obtained by using our new method, random forests and XGBoost using ten-fold cross-validation. The left side of Figure 1 also contains results obtained with the formula-size method of [12]. We also

report accuracies that we found in the literature for these datasets. For the left side of Figure 1, we include the accuracies that were reported for these datasets in the UCI machine learning repository. For the right side of Figure 1, for the case of Colon, we used the result reported in [22], while for remaining datasets we reported the accuracies given in the full-version of [9].

We emphasize that the accuracies found from the literature are not directly comparable to the accuracies that we obtained with our method. For example, in some cases a single 70/30-split was used instead of our ten-fold cross-validation. The conventions concerning the handling of missing values could also differ from ours. Furthermore, the accuracies in [9] are given to the nearest percentage point.

From Figure 1 we see that our method produces classifiers that mostly have accuracies similar to the ones obtained by the alternative methods. In particular, on medical data such as BreastCancer and Hepatitis, our method compares well to the reference methods. Perhaps the biggest gaps are found in the cases of CoverType and Electricity, two of the benchmark datasets. However, we emphasize that in [9], it was mentioned that all of the benchmark datasets were specifically selected to be “not too easy”. Hence, one should not perhaps expect that “simple” classifiers with focus on interpretability will perform too well with these datasets.

We list the average runtime of our method on a single split of the data in the Appendix 6.3. For six of the datasets, the average runtime was less than 30 seconds. Of the remaining datasets, four had runtimes of less than 12 minutes. Even for the two largest datasets, Covertype and RoadSafety, the runtime is still less than an hour and, notably, systematically less than that of the hyperparameter optimizations of random forests and XGBoost. The experiments on our method, random forests and XGBoost were run on a laptop.

We turn our attention to our main goal of interpretability. In Table 1 we report the average of the *smallest number of features* sufficient to reach the reported accuracy. Table 1 also gives the number of rows and attributes in the Booleanized datasets as well as the original datasets. We see from Table 1 that for eight out of 12 datasets, the best accuracy was reached already with on average less than four features, leading to very short formulas. Such short DNF-formulas are globally interpretable as follows. To interpret a short DNF-formula, one looks at each of the few conjunctions in the formula individually to determine their meaning. This is easy as each conjunction has very few attributes. The meaning of the entire formula is then given by the fact that it accepts only the cases given by the conjunctions and rejects everything else.

Next we discuss the results obtained using the formula-size method of [12]. By the size of a propositional formula they mean the number of proposition symbols, conjunctions, disjunctions and negations in the formula. The results are given in Figure 1.

We see that for most datasets on the left side in Figure 1, the accuracies obtained by the formula-size method are similar to our method. A meaningful difference can be seen in the HeartDisease dataset, where both the accuracy of 73.6 percent and standard deviation of 12.2 percent are significantly behind other methods. The formulas obtained are short, with formula size mostly less than 10.

The downside of the formula-size method is computational resources. As described in the previous subsection, we used a Linux cluster featuring Intel Xeon 2.40 GHz CPUs with 8 CPU cores per run, employing a timeout of 72 hours per instance and a memory limit of 64 GB. For all but one dataset we ran the formula-size method on, the computation took more than 24 hours per split of the data with this high efficiency setup. Even for the easiest one,

Table 1

The average number of features used by the final formulas for each dataset. We also report the total numbers of features in the Booleanized and original datasets for comparison as well as the numbers of data points.

Data set	Selected features	Boolean features	Original features	Data points
BankMarketing	1.3	51	51	4521
BreastCancer	3.8	9	9	683
CongressionalVoting	1.0	48	16	435
GermanCredit	3.9	61	20	1000
HeartDisease	3.9	25	13	303
Hepatitis	2.7	74	17	155
StudentDropout	5.8	112	36	4424
Coverttype	2.5	54	54	423680
Electricity	7.5	14	8	38474
EyeMovement	5.5	26	23	7608
RoadSafety	5.4	324	32	111762
Colon	2.1	5997	2000	63

BreastCancer, the computation time was in the order of 12 hours per split. Furthermore, for the benchmark datasets and the high-dimensional biological dataset Colon, the formula-size method is unfeasible already for very small formula sizes and thus we do not report any results for those datasets.

In contrast to the formula-size method, for all but the two largest datasets, the runs of our method took less than 12 minutes per split, and in six cases only seconds. More detailed runtimes are reported in Appendix 6.3. We conclude that our method obtains similar or better results than the formula-size method with a fraction of the computational resources on datasets where both methods are feasible. Our method is also usable on larger datasets, where the formula-size method is not.

4.3. Examples of obtained formulas

Here we present selected examples of formulas that were generated by our method. Firstly, we refer the reader back to the example formulas concerning the BreastCancer and Colon datasets presented in the Introduction. Both of these formulas were very short and had high accuracies.

Another example formula is from the HeartDisease dataset, where the task is to determine, whether a patient has a heart disease. We obtain the formula

$$(p_1 \wedge p_2) \vee (p_1 \wedge \neg p_3) \vee (p_2 \wedge \neg p_3),$$

where the attributes p_1 , p_2 and p_3 relate to a color test of blood vessels, fixed thallium defects and chest pain, respectively. This formula is obtained from two different splits of the data and the respective accuracies are 76.7 percent and 86.7 percent. The average accuracy of our method

on this dataset is 81.2 percent. Random forests average at 81.8 percent, XGBoost at 80.8 percent and the formula size method at 73.6 percent.

The Hepatitis dataset concerns the mortality of patients with hepatitis. The formula

$$p_1 \vee \neg p_2 \vee \neg p_3$$

is obtained from five of the ten splits of the data. Here p_1 , p_2 and p_3 relate to albumin, bilirubin and a histology test respectively. The accuracies of this formula on the five different test datasets range from 61.5 percent to 92.3 percent. The average accuracy of our method on the Hepatitis dataset is 80.7 while random forests get 78.2 percent and XGBoost 79.0 percent. The high variance seems to be due to the data only consisting of 155 data points. This is supported by the fact that all of the tested methods obtain a high standard deviation of accuracy on this dataset.

In the Appendix 6.4, there are further examples of formulas obtained as outputs of our method. For each dataset, we list the formula obtained for the first 90/10 split of our ten-fold cross-validation. While some of these formulas are quite long, they should still be rather fast to use for local explanations explicating why a particular input was classified in a particular way (cf. the Introduction for details on local explanations). For these experiments, we used the same maximum number 10 of features and the same limit of 1 percentage point to choose an accurate enough formula. These are hyperparameters that one could optimize to obtain a better balance of interpretability and accuracy on any specific dataset. For some datasets, raising the number of maximum features could lead to more accurate, but longer, formulas. Raising the percentage threshold on the other hand would lead to shorter, more interpretable formulas at the cost of some accuracy.

5. Conclusions

We have introduced a new method to compute immediately interpretable Boolean classifiers for tabular data. While the main point is interpretability, even the accuracy of our formulas is similar to the ones obtained via widely recognized current methods. We have also established a theoretical result for estimating if the obtained formulas actually correspond to ideally accurate ones in relation to the background distribution. In the future, we will especially consider more custom-made procedures of discretization in the context of our method, as this time discretization was carried out in a very rough way to demonstrate the effectiveness and potential of our approach. We expect this to significantly improve our method over a notable class of tabular datasets.

Our approach need not limit to Boolean formulas only, as we can naturally extend our work to general relational data. We can use, e.g., description logics and compute concepts C_1, \dots, C_k and then perform our procedure using C_1, \dots, C_k , finding short Boolean combinations of concepts. This of course differs from the approach of computing empirical ideal classifiers in the original description logic, but can nevertheless be fruitful and interesting. We leave this for future work.

Acknowledgments

T. Janhunnen, A. Kuusisto, M. F. Rankooh and M. Vilander were supported by the Academy of Finland consortium project Explaining AI via Logic (XAILOG), grant numbers 345633 (Janhunnen) and 345612 (Kuusisto). A. Kuusisto and M. Vilander were also supported by the Academy of Finland project Theory of computational logics, grant numbers 352419, 352420, 353027, 324435, 328987.

References

- [1] EU-GDPR, Regulation (EU) 2016/679 of the European parliament and of the council, 2016. URL: <https://www.privacy-regulation.eu/en/index.htm>.
- [2] California-OAG, CCPA regulations: Final regulation text, 2021. URL: <https://oag.ca.gov/privacy/ccpa>.
- [3] R. Shwartz-Ziv, A. Armon, Tabular data: Deep learning is not all you need, *Inf. Fusion* 81 (2022) 84–90.
- [4] V. Borisov, T. Leemann, K. Seßler, J. Haug, M. Pawelczyk, G. Kasneci, Deep neural networks and tabular data: A survey, *CoRR* abs/2110.01889 (2021). URL: <https://arxiv.org/abs/2110.01889>.
- [5] M. Pawelczyk, K. Broelemann, G. Kasneci, Learning model-agnostic counterfactual explanations for tabular data, in: Y. Huang, I. King, T. Liu, M. van Steen (Eds.), *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020, ACM / IW3C2, 2020*, pp. 3126–3132.
- [6] M. Sahakyan, Z. Aung, T. Rahwan, Explainable artificial intelligence for tabular data: A survey, *IEEE Access* 9 (2021) 135392–135422.
- [7] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nature Machine Intelligence* 1 (2019) 206–215.
- [8] R. C. Holte, Very simple classification rules perform well on most commonly used datasets, *Mach. Learn.* 11 (1993) 63–91.
- [9] L. Grinsztajn, E. Oyallon, G. Varoquaux, Why do tree-based models still outperform deep learning on typical tabular data?, in: *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track, 2022*. URL: https://openreview.net/forum?id=Fp7__phQszn.
- [10] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, R. Rastogi (Eds.), *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016, ACM, 2016*, pp. 785–794.
- [11] L. Breiman, Random forests, *Machine Learning* 45 (2001) 5–32. doi:10.1023/A:1010933404324.
- [12] R. Jaakkola, T. Janhunnen, A. Kuusisto, M. F. Rankooh, M. Vilander, Short Boolean formulas as explanations in practice, in: S. A. Gaggl, M. V. Martinez, M. Ortiz (Eds.), *Logics in Artificial Intelligence - 18th European Conference, JELIA 2023, Dresden, Germany,*

- September 20-22, 2023, Proceedings, volume 14281 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 90–105.
- [13] F. M. Dekking, C. Kraaikamp, H. P. Lopuhaä, L. E. Meester, *A Modern Introduction to Probability and Statistics: Understanding Why and How*, Springer Texts in Statistics, Springer, 2005.
- [14] J. Marques-Silva, Logic-based explainability in machine learning, in: L. E. Bertossi, G. Xiao (Eds.), *Reasoning Web. Causality, Explanations and Declarative Knowledge - 18th International Summer School 2022*, Berlin, Germany, September 27-30, 2022, Tutorial Lectures, volume 13759 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 24–104.
- [15] A. Shih, A. Choi, A. Darwiche, A symbolic approach to explaining Bayesian network classifiers, in: J. Lang (Ed.), *IJCAI*, 2018, pp. 5103–5111.
- [16] A. Ignatiev, N. Narodytska, J. Marques-Silva, Abduction-based explanations for machine learning models, in: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*, Honolulu, Hawaii, USA, January 27 - February 1, 2019, AAAI Press, 2019, pp. 1511–1519.
- [17] E. Angelino, N. Larus-Stone, D. Alabi, M. Seltzer, C. Rudin, Learning certifiably optimal rule lists for categorical data, *Journal of Machine Learning Research* 18 (2018) 1–78.
- [18] C. Rudin, B. Ustun, Optimized scoring systems: Toward trust in machine learning for healthcare and criminal justice, *Interfaces* 48 (2018) 449–466. doi:10.1287/inte.2018.0957.
- [19] L. Devroye, L. Györfi, G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, volume 31 of *Stochastic Modelling and Applied Probability*, Springer, 1996.
- [20] J. Li, K. Cheng, S. Wang, F. Morstatter, R. P. Trevino, J. Tang, H. Liu, Feature selection: A data perspective, *ACM Computing Surveys (CSUR)* 50 (2018) 94.
- [21] T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: A next-generation hyperparameter optimization framework, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, p. 2623–2631. doi:10.1145/3292500.3330701.
- [22] J. Yang, O. Lindenbaum, Y. Kluger, Locally sparse neural networks for tabular biomedical data, in: *International Conference on Machine Learning*, 2021.
- [23] W. Hoeffding, Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association* 58 (1963) 13–30. doi:10.1080/01621459.1963.10500830.

6. Appendix.

6.1. Proof of Theorem 3.1

Let τ be a propositional vocabulary and $q \notin \tau$ the proposition symbol that we need to explain. Let $\tau^+ := \tau \cup \{q\}$ and fix a probability distribution $\mu : T_{\tau^+} \rightarrow [0, 1]$. In what follows, for notational simplicity we will assume that $\mu(t \wedge q) > \mu(t \wedge \neg q)$, for every τ -type t .

Let M be a sample of size n , i.e., a τ^+ -model of size n . For each τ^+ -type t we use $\llbracket t \rrbracket_M$ to denote the number of times t is realized in M . The main idea of the proof is to show that if

$$n \geq \frac{2 \ln(2^{|\tau^+|+1}/\delta)}{\varepsilon^2},$$

then

$$\Pr \left[\text{for every } \tau\text{-type } t \text{ for which } \mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon \text{ we have that } \llbracket t \wedge q \rrbracket_M > \llbracket t \wedge \neg q \rrbracket_M \right] \geq 1 - \delta.$$

Indeed, the above clearly implies that also with probability at least $1 - \delta$ we have that the empirical ideal classifier with respect to M agrees with the ideal classifier with respect to all τ -types t with $\mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon$.

Consider first a fixed τ -type t such that $\mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon$. Now, observe that if

$$n(\mu(t \wedge q) - \theta) < \llbracket t \wedge q \rrbracket_M \text{ and } \llbracket t \wedge \neg q \rrbracket_M < n(\mu(t \wedge \neg q) + \theta),$$

then

$$\llbracket t \wedge q \rrbracket_M > \llbracket t \wedge \neg q \rrbracket_M$$

holds provided that

$$\mu(t \wedge q) - \mu(t \wedge \neg q) \geq 2\theta.$$

Setting $\theta := \varepsilon/2$, we have by assumption that

$$\begin{aligned} & \Pr[\llbracket t \wedge q \rrbracket_M \geq \llbracket t \wedge \neg q \rrbracket_M] \\ & \geq \Pr[n(\mu(t \wedge q) - \theta) < \llbracket t \wedge q \rrbracket_M \text{ and } \llbracket t \wedge \neg q \rrbracket_M < n(\mu(t \wedge \neg q) + \theta)] \\ & \geq 1 - (\Pr[n(\mu(t \wedge q) - \theta) \geq \llbracket t \wedge q \rrbracket_M] + \Pr[\llbracket t \wedge \neg q \rrbracket_M \geq n(\mu(t \wedge \neg q) + \theta)]) \end{aligned}$$

For the second inequality we used the union bound. To get a lower bound of $1 - \delta$ on the latter probability, it suffices to show that

$$\Pr[n(\mu(t \wedge q) - \theta) \geq \llbracket t \wedge q \rrbracket_M] \leq \frac{\delta}{2}$$

and

$$\Pr[\llbracket t \wedge \neg q \rrbracket_M \geq n(\mu(t \wedge \neg q) + \theta)] \leq \frac{\delta}{2}$$

Hoeffding's inequalities [23] imply that these bounds hold provided that

$$n \geq \frac{\ln(2/\delta)}{2\theta^2} = \frac{2 \ln(2/\delta)}{\varepsilon^2}.$$

Thus, if we take sample M of at least size

$$\frac{2 \ln(2/\delta)}{\varepsilon^2},$$

then with probability at least $1 - \delta$ we have that $\llbracket t \wedge q \rrbracket_M > \llbracket t \wedge \neg q \rrbracket_M$.

So far we focused on a single τ -type t . Using union bound again we see that

$$\Pr \left[\text{for every } \tau\text{-type } t \text{ for which } \mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon \text{ we have that } \llbracket t \wedge q \rrbracket_M > \llbracket t \wedge \neg q \rrbracket_M \right] \\ \geq 1 - \sum \Pr[\llbracket t \wedge q \rrbracket_M \leq \llbracket t \wedge \neg q \rrbracket_M],$$

where the sum is performed with respect to τ -types t for which $\mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon$. Setting again $\theta := \varepsilon/2$ it follows from our previous calculations that for every $\tau \in T_\tau$ for which $\mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon$ we have that

$$\Pr[\llbracket t \wedge q \rrbracket_M \leq \llbracket t \wedge \neg q \rrbracket_M] \leq \delta/2^{|\tau|},$$

provided that

$$n \geq \frac{\ln(2^{|\tau|+1}/\delta)}{2\theta^2} = \frac{2 \ln(2^{|\tau|+1}/\delta)}{\varepsilon^2}.$$

In particular, we will then have that

$$\Pr \left[\text{for every } \tau\text{-type } t \text{ for which } \mu(t \wedge q) - \mu(t \wedge \neg q) \geq \varepsilon \text{ we have that } \llbracket t \wedge q \rrbracket_M > \llbracket t \wedge \neg q \rrbracket_M \right] \geq 1 - \delta,$$

which is what we wanted to show.

6.2. Hyperparameter spaces for Random Forest and XGBoost

Hyperparameters not listed here were kept at their default values.

6.2.1. Random Forest

- max_depth: None, 2, 3, 4
- n_estimators: Integer sampled from [9.5, 3000.5] using log domain
- criterion: gini, entropy
- max_features: sqrt, log2, None, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9
- min_samples_split: 2, 3
- min_samples_leaf: Integer sampled from [1.5, 50.5]
- bootstrap: True, False
- min_impurity_decrease: 0.0, 0.01, 0.02, 0.05

6.2.2. XGBoost

- `max_depth`: Integer sampled from [1,11]
- `n_estimators`: Integer sampled from [100,5900] with step-size being 200.
- `min_child_weight`: Float sampled from [1.0, 100.0] using log domain
- `subsample`: Float sampled from [0.5, 1.0]
- `learning_rate`: Float sampled from [1e-5, 0.7] using log domain
- `colsample_bylevel`: Float sampled from [0.5, 1.0]
- `gamma`: Float sampled from [1e-8, 7.0] using log domain
- `reg_lambda`: Float sampled from [1.0, 4.0] using log domain
- `reg_alpha`: Float sampled from [1e-8, 100.0] using log domain

6.3. Runtimes

We report the average runtime of our method over the different splits of the ten-fold (or in the case of Colon leave-one-out) cross-validation. The experiments were run on a standard laptop.

Data set	Runtime of our method
BankMarketing	1 minute 30 seconds
BreastCancer	13 seconds
CongressionalVoting	10 seconds
GermanCredit	26 seconds
HeartDisease	11 seconds
Hepatitis	2 seconds
StudentDropout	1 minute 26 seconds
Coverttype	45 minutes 10 seconds
Electricity	11 minutes 54 seconds
EyeMovement	3 minutes 0 seconds
RoadSafety	28 minutes 37 seconds
Colon	19 seconds

6.4. Example formula for each dataset

Here we report for each dataset the formula given by the first of the ten splits in our ten-fold cross-validation.

6.4.1. BankMarketing

poutcome_success

6.4.2. BreastCancer

$$\begin{aligned} & (\neg \text{bare_nuclei_above_median} \wedge \neg \text{bland_chromatin_above_median}) \\ & \vee (\neg \text{bland_chromatin_above_median} \wedge \neg \text{single_epithelial_cell_size_above_median}) \\ & \vee (\neg \text{bland_chromatin_above_median} \wedge \neg \text{uniformity_of_cell_size_above_median}) \\ & \vee (\text{bare_nuclei_above_median} \wedge \neg \text{single_epithelial_cell_size_above_median} \wedge \neg \text{uniformity_of_cell_size_above_median}) \\ & \vee (\text{single_epithelial_cell_size_above_median} \wedge \neg \text{bare_nuclei_above_median} \wedge \neg \text{uniformity_of_cell_size_above_median}) \\ & \vee (\text{uniformity_of_cell_size_above_median} \wedge \neg \text{bare_nuclei_above_median} \wedge \neg \text{single_epithelial_cell_size_above_median}) \end{aligned}$$

6.4.3. CongressionalVoting

physician_fee_freeze_y

6.4.4. GermanCredit

$$\begin{aligned} & (A1_A14 \wedge \neg A1_A11 \wedge \neg A3_A30) \\ & \vee (A1_A14 \wedge \neg A1_A11 \wedge \neg A3_A31) \\ & \vee (A6_A65 \wedge \neg A1_A11 \wedge \neg A3_A30) \\ & \vee (\neg A1_A14 \wedge \neg A3_A30 \wedge \neg A3_A31) \\ & \vee (A1_A11 \wedge A6_A65 \wedge \neg A1_A14 \wedge \neg A3_A31) \end{aligned}$$

6.4.5. HeartDisease

$$(\text{caa}_0 \wedge \text{thall}_2) \vee (\text{caa}_0 \wedge \neg \text{old_peak_above_median}) \vee (\text{thall}_2 \wedge \neg \text{old_peak_above_median})$$

6.4.6. Hepatitis

albumin_above_median \vee \neg *bilirubin_above_median* \vee \neg *histology*

6.4.7. StudentDropout

$$\begin{aligned} & (\neg \text{Curricular_units_1st_sem_grade_above_median} \wedge \neg \text{Tuition_fees_up_to_date}) \\ & \vee (\neg \text{Curricular_units_2nd_sem_approved_above_median} \wedge \neg \text{Tuition_fees_up_to_date}) \\ & \vee (\neg \text{Curricular_units_1st_sem_approved_above_median} \wedge \neg \text{Curricular_units_1st_sem_grade_above_median} \\ & \quad \wedge \neg \text{Curricular_units_2nd_sem_approved_above_median} \wedge \neg \text{Curricular_units_2nd_sem_grade_above_median}) \end{aligned}$$

6.4.8. Covertypes

$$(\text{Elevation_above_median} \wedge \neg \text{Soil_Type12}) \vee (\text{Soil_Type22} \wedge \neg \text{Soil_Type12})$$

6.4.9. Electricity

$(nswprice_above_median \wedge transfer_above_median \wedge vicprice_above_median)$
 $\vee (nswprice_above_median \wedge vicprice_above_median \wedge \neg nswdemand_above_median)$
 $\vee (day_0 \wedge nswdemand_above_median \wedge period_above_median \wedge transfer_above_median \wedge vicdemand_above_median)$
 $\vee (nswdemand_above_median \wedge nswprice_above_median \wedge \neg day_0 \wedge \neg period_above_median)$
 $\vee (nswdemand_above_median \wedge nswprice_above_median \wedge \neg vicdemand_above_median \wedge \neg vicprice_above_median)$
 $\vee (nswprice_above_median \wedge period_above_median \wedge \neg day_0 \wedge \neg vicdemand_above_median)$
 $\vee (nswprice_above_median \wedge vicdemand_above_median \wedge \neg day_0 \wedge \neg transfer_above_median)$
 $\vee (day_0 \wedge nswdemand_above_median \wedge period_above_median \wedge vicprice_above_median \wedge \neg transfer_above_median)$
 $\vee (nswdemand_above_median \wedge period_above_median \wedge transfer_above_median \wedge vicprice_above_median \wedge \neg day_0)$
 $\vee (day_0 \wedge period_above_median \wedge transfer_above_median \wedge \neg nswprice_above_median \wedge \neg vicprice_above_median)$
 $\vee (nswdemand_above_median \wedge vicdemand_above_median \wedge vicprice_above_median \wedge \neg day_0 \wedge \neg period_above_median)$
 $\vee (day_0 \wedge nswprice_above_median \wedge \neg period_above_median \wedge \neg vicdemand_above_median \wedge \neg vicprice_above_median)$
 $\vee (day_0 \wedge nswprice_above_median \wedge \neg transfer_above_median \wedge \neg vicdemand_above_median \wedge \neg vicprice_above_median)$
 $\vee (day_0 \wedge vicprice_above_median \wedge \neg nswdemand_above_median \wedge \neg period_above_median \wedge \neg vicdemand_above_median)$
 $\vee (nswdemand_above_median \wedge transfer_above_median \wedge \neg period_above_median \wedge \neg vicdemand_above_median \wedge \neg vicprice_above_median)$
 $\vee (vicdemand_above_median \wedge \neg day_0 \wedge \neg nswdemand_above_median \wedge \neg period_above_median \wedge \neg transfer_above_median)$
 $\vee (day_0 \wedge vicdemand_above_median \wedge \neg nswdemand_above_median \wedge \neg nswprice_above_median \wedge \neg transfer_above_median \wedge \neg vicprice_above_median)$

6.4.10. EyeMovement

$(nextWordRegress_0 \wedge \neg nextWordRegress_1 \wedge \neg wordNo_above_median)$
 $\vee (nextWordRegress_0 \wedge prevFixPos_above_median \wedge \neg nextWordRegress_1 \wedge \neg regressDur_above_median)$
 $\vee (nextWordRegress_1 \wedge \neg nextWordRegress_0 \wedge \neg prevFixPos_above_median \wedge \neg regressDur_above_median$
 $\wedge \neg wordNo_above_median)$

6.4.11. RoadSafety

$(Casualty_Type_9 \wedge Sex_of_Casualty_1 \wedge Vehicle_Type_9 \wedge \neg Sex_of_Casualty_0)$
 $\vee (Propulsion_Code_1 \wedge Sex_of_Casualty_1 \wedge Vehicle_Type_9 \wedge \neg Sex_of_Casualty_0)$
 $\vee (Sex_of_Casualty_1 \wedge Vehicle_Type_9 \wedge \neg Engine_Capacity_CC_above_median \wedge \neg Sex_of_Casualty_0)$
 $\vee (Propulsion_Code_1 \wedge Sex_of_Casualty_0 \wedge Vehicle_Type_9 \wedge \neg Casualty_Type_9 \wedge \neg Sex_of_Casualty_1)$
 $\vee (Propulsion_Code_1 \wedge Sex_of_Casualty_1 \wedge \neg Casualty_Type_9 \wedge \neg Engine_Capacity_CC_above_median \wedge \neg Sex_of_Casualty_0)$
 $\vee (Sex_of_Casualty_0 \wedge Vehicle_Type_9 \wedge \neg Casualty_Type_9 \wedge \neg Engine_Capacity_CC_above_median \wedge \neg Sex_of_Casualty_1)$

6.4.12. Colon

For the high-dimensional dataset Colon we used leave-one-out cross-validation; we report the formula obtained from leaving the first datapoint out.

v765_two