

Extracting optimal explanations for ensemble trees via automated reasoning

Author

Zhang, G, Hóu, Z, Huang, Y, Shi, J, Bride, H, Dong, JS, Gao, Y

Published

2022

Journal Title

Applied Intelligence

Version

Accepted Manuscript (AM)

DOI

[10.1007/s10489-022-04180-1](https://doi.org/10.1007/s10489-022-04180-1)

Rights statement

© 2022 Springer US. This is an electronic version of an article published in Applied Intelligence, 2022. Applied Intelligence is available online at: <http://link.springer.com/> with the open URL of your article.

Downloaded from

<http://hdl.handle.net/10072/419413>

Griffith Research Online

<https://research-repository.griffith.edu.au>

Extracting Optimal Explanations for Ensemble Trees via Logical Reasoning

Gelin Zhang¹, Zhé Hóu², Yanhong Huang¹, Jianqi Shi¹, Hadrien Bride², Jin Song Dong^{2,3}, and Yongsheng Gao²

¹ National Trusted Embedded Software Engineering Technology Research Center,
East China Normal University, China

² Griffith University, Australia

³ National University of Singapore, Singapore

Abstract. Ensemble trees are a popular machine learning model which often yields high prediction performance when analysing structured data. Although individual small decision trees are deemed explainable by nature, an ensemble of large trees is often difficult to understand. In this work, we propose an approach called optimised explanation (OptExplain) that faithfully extracts global explanations of ensemble trees using a combination of logical reasoning, sampling and optimisation. Building on top of this, we propose a method called the profile of equivalent classes (ProClass), which uses MAX-SAT to simplify the explanation even further¹. Our experimental study on several datasets shows that our approach can provide high-quality explanations to large ensemble trees models, and it betters recent top-performers.

Keywords: Explainable Artificial Intelligence (XAI) · Random Forest · Classification · Rule Extraction.

1 Introduction

Background. Ensemble trees are a family of machine learning techniques that combine individual decision trees to form a better prediction model. Examples include random forest [14,2], which combines strong learners (e.g., large trees) to reduce variance and avoid overfitting. Boosting [7,9], on the other hand, combines weak learners (e.g., small trees) to reduce bias. Ensemble trees are very successful in today’s data analytics competitions and applications; they are especially suited to analyse *structured data* such as databases and spreadsheets, where they sometimes outperform deep learning [20].

Although decision trees are often deemed an explainable, or even a “white-box” model, such an impression usually refers to a single, short decision tree. In the context of ensemble trees such as the models generated by random forest or boosting, there can be a large number of trees and each tree can be gigantic. For example, to achieve a 0.76+ area-under-the-curve (AUC) for the 1 million

¹ The code is available at <https://github.com/GreenZhang/OptExplain>.

flight dataset [20], Silas [3,15] trains a model of 100 trees and each tree has more than 32,000 branches. Such a model certainly does not manifest itself in an explainable manner to the general user. The main goal of this work is to extract faithful explanations for such large-scale models.

There are several existing methods for analysing and interpreting machine learning models. For example, the LIME tool [22] and the SHAP values [16] are both promising techniques for solving this problem. We will discuss more details of related work in Section 5. However, most existing work is done from a statistics perspective. Such methods use a prediction model as a black-box and attempt to find statistical (e.g., linear) approximations of the model. By contrast, our philosophy is that we should analyse the internal working of the model and obtain an understanding of how it works logically. Further, many existing techniques are focused on *local explanations*, that is, how the model predicts for a particular data instance. This work is primarily about *global explanations*, which explains how the model behaves generally. Part of the reasons why we choose ensemble trees is that decision trees are no strangers to logicians. For example, binary decision diagrams, which have a similar form, are widely used in theorem proving [10] and model checking [25]. The tree structure is well-understood in the logic and verification community, and there are many possibilities to apply logical reasoning to analysing ensemble trees.

Our approach. In our previous work [15], we have used sampling and maximum satisfiable subset to extract the decision logic of the model. However, it is non-trivial to manually adjust the sampling parameters, which may lead to vastly different explanations. Default parameters often lead to very simple explanations that diverge from the original model. In this paper, we propose an integrated and automated framework for providing global explanations. Moreover, our goal is not just to give *an* explanation as is done in the literature, but to give *the optimal* explanation in terms of *simplicity* and *faithfulness*.

An outline of our approach follows: we extract logical formulae from a set of trees where each branch forms a “decision rule”. We then reduce the size of the model by filtering out low-quality nodes (i.e., sub-formulae) and branches. We also devise a customised formulae simplification algorithm to obtain logically equivalent smaller models. In case there are still too many decision rules, we group the rules into “equivalent classes” to further abstract the model. The parameters in the above process are optimised using particle swarm towards a sweet spot of simplicity and faithfulness. As an extra step, we can simplify the explanation using MAX-SAT to obtain even more abstract representations of each equivalent class, which we call the “profiles of classes”. Such profiles can provide straightforward and even visual explanations of the model.

The utilities of this work are manifold. First, our approach can provide human-understandable explanations that are very close to the original ensemble trees model in predictive behaviour. Second, such an explanation can also be used as an approximation of the original model. For example, verification of machine learning models is another popular topic, but a general sound and complete verification algorithm for ensemble trees have proven impractical [26]. As a step

back, we can look at the software testing scenario: since the explanation mimics the behaviour of the original model, if it violates a property, then it is likely that the original model would fail the verification, too. In such cases, we can use the explanation to constrain the search space when finding counterexamples. Finally, this work can serve as a stepping stone towards explaining deep learning. There are existing methods for converting neural networks to decision trees [13] exactly for explanation purposes. However, these methods only induce a single decision tree, whose predictive performance is incomparable to the deep learning model. One may convert neural networks to a set of decision trees instead, then this work can be directly applied to obtain explanations.

Contributions. The main contributions of this paper are as follows:

1. We formalise ensemble trees into logical formulae and develop simplification and abstraction algorithms that are specialised for machine learning.
2. We propose an automated explanation extraction method called *OptExplain*, which combines logical reasoning, sampling, and bio-inspired optimisation.
3. We also develop a method called *ProClass* that computes the abstractions of each (equivalent) class using MAX-SAT.
4. Through case studies and experiment, we show that our method is practical and useful on different datasets. It also outperforms similar tools.

The remainder of this paper is organised as follows: Section 2 describes the preliminary concepts, Section 3 details the proposed approach, Section 4 gives case studies and experiment, Section 5 discusses related work, and Section 6 concludes the paper.

2 Preliminaries

In this section, we redefine decision trees and their ensembles from a logical language point of view.

2.1 Decision Trees With a Logical Foundation

In supervised learning, a structured dataset for classification is defined as set of *instances* of the form (\vec{x}, y) where $\vec{x} = [x_1, \dots, x_n]$, $n \in \mathbb{N}$, is an input vector called *features* and y is an outcome value often called the *label*. We denote by X the feature space and Y the outcome space.

A *decision tree* is composed of internal nodes (diamonds in Fig. 1) and terminal nodes called leaves (ovals in Fig. 1). Each internal node is associated with a logical formula over




Fig. 1: An example decision tree.

a feature. Each leaf node contains a set of instances, which yield a *vote distribution* of the form (n_1, \dots, n_m) where m is the number of classes and n_i ($1 \leq i \leq m$) is the number of instances of the corresponding class. For example, in Fig. 1, the left-most leaf node $(0, 6)$ indicates that there are 0 `class0` instances and 6 `class1` instances. Without loss of generality, we focus on binary trees, in which internal nodes have two successors respectively called the left and right child nodes. By convention, the instances that satisfy the logical formula of an internal node go to the right child node, and those that do not satisfy go to the left child node. For example, in Fig. 1, let I be the set of training instances associated with the root node, $I_1 \subset I$ be the subset that satisfies the formula F_1 , then I_1 will be the set of instances associated with the right child node (with formula F_2), and $I_2 = I \setminus I_1$ will be the set of instances associated with the left child (leaf) node.

Given a decision tree, any input vector (or instance) is associated with a single leaf. A decision tree is, therefore, a compact representation of a function of the form $t : X \rightarrow \mathbb{N}^m$, where m is the number of classes. The output of a decision tree is a distribution of *votes* for each class. To obtain an outcome in Y , we take the class with the most votes.

2.2 Ensemble of Decision Trees

We adopt the definitions of Cui et al. [4]. Let an ensemble be a set of decision trees of size T . It gives the weighted sum of the trees as follows:

$$E(x) = \sum_{i=1}^T w_i \cdot t_i(x) \quad (1)$$

where E is the function for the ensemble, w_i and t_i are respectively the weight and function for each tree. The summation aggregates the weighted votes from each tree and obtains the final votes for each class. Thus, the ensemble is also a function of the signature $E : X \rightarrow \mathbb{N}^m$ and requires a voting mechanism to obtain the outcome. We give some famous examples of ensemble trees below.

Bagging. Each decision tree is trained using a subset of the dataset that is sampled uniformly with replacement. The remaining instances form the out-of-bag (OOB) set. When selecting the best formula at each decision node in a tree, only a subset of the features are considered. This is commonly found in algorithms such as Random Forest [2]. Bagging grows large trees with low bias and the ensemble reduces variance.

Boosting. Boosting trains weak learners, i.e., small trees, iteratively as follows:

$$E_{i+1}(x) = E_i(x) + \alpha_i \cdot t_i(x) \quad (2)$$

where t_i is the weak learner trained at iteration i and α_i is its weight. The final ensemble is thus a special case of $E(x)$ above where w_i is α_i . The ensemble reduces bias.

AdaBoost [7] focuses on training instances that are misclassified in the previous iteration by optimising α_i and t_i in the formula below:

$$\underset{\alpha_i, t_i}{\text{minimise}} \sum_{j=0}^N L(y^{(j)}, E_i(x^{(j)}) + \alpha_i \cdot t_i(x^{(j)})) \tag{3}$$

where L is a loss function measuring the difference between the actual outcome $y^{(j)}$ of instance j and $E_{i+1}(x^{(j)})$. AdaBoost often uses exponential loss $L(a, b) = e^{-a \cdot b}$ in which case the weak learners are trained by weighted instances.

Gradient Boosting [8] is a generalisation of the above minimisation where each tree is trained using

$$t_i(x^{(j)}) \approx -\frac{\partial L(y^{(j)}, E_i(x^{(j)}))}{\partial E_i(x^{(j)})} \tag{4}$$

which is equivalent to training a regression tree using original data points but with new outcome values defined by the negative gradient. In this case, α_i is the learning rate.

In this work, we evaluate our approach using random forest, but our approach can also be adapted to handle boosting models.

3 The Proposed Method: OptExplain and ProClass

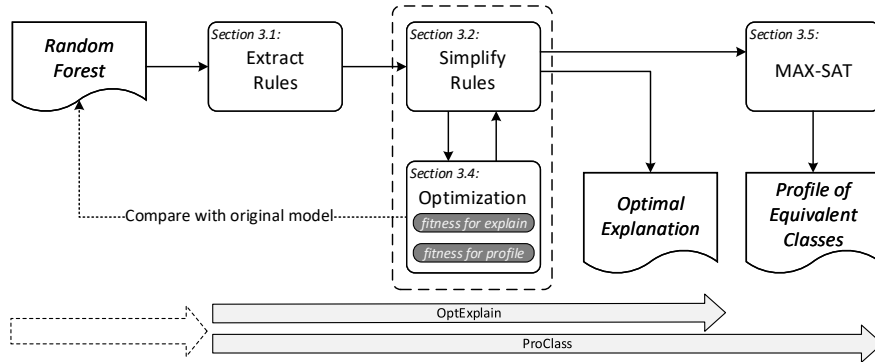


Fig. 2: An overview of the proposed method.

This section details our ensemble trees explanation extraction approach. We give an overview of our approach in Figure 2. First, we get a set of decision rules by traversing the random forest. This set of decision rules is made up of branch formulae for all the trees in the forest. The size of this set is usually large, and we need to simplify the set in order to obtain a smaller set that is as close as possible to the prediction ability of the original model. We keep the nodes and branches with good quality to get the final explanation set. The

preciseness of the simplification process and the scale of the explanation are related to the input parameters, and we get the optimized parameters through the bionic optimization algorithm PSO (Particle Swarm Optimization). There is two fitness function in optimization process, one dedicate to generate optimal explanations while the other one is to generate explanations used by profile.

3.1 From Decision Trees to Decision Rules

Given a decision tree defined in Section 2, it is straightforward to obtain the formula that is associated with each internal node. From there, we can obtain the “decision logic” of a branch via the *branch formula* of the following form:

$$(\bigwedge F) \rightarrow D \quad (5)$$

where $\bigwedge F$ is the conjunction of all the (possibly negated) internal node formula on the branch, and D is the vote distribution at the leaf node. That is, if the branch goes through a node F to the right branch, then we include F in the conjunction, otherwise we include $\neg F$. We refer to a formula of the above form as a *decision rule*.

It is worth noting that, by the construction of decision trees, the branches are *exclusive* to each other. That is, a data instance can never satisfy multiple branch formulae from the same tree at the same time. As a result, a decision tree t can be converted into a set of mutually exclusive decision rules, denoted as R_t . R_t can be used in classification tasks by finding the decision rule that satisfies an instance and outputting the class of the largest number of votes.

The above method can be extended to handle an ensemble of trees produced by random forest [2] or boosting [7,8]. In such cases, we need to consider a set E of trees, and we need to multiply the vote distribution D of each tree by its weight in the ensemble. The result, which we refer to as R_E , is the union of the set of decision rules from each tree.

Unlike the set of decision rules for a single tree, that for an ensemble of trees may contain decision rules that are *not* exclusive to each other. In fact, for an ensemble E of n trees and for any data instance, there should be exactly n decision rules in the set R_E that are satisfied by the instance — one from each tree. To use R_E in classification tasks, one can find the subset of decision rules that are satisfied by an instance x and aggregate the weighted vote counts for each class from those rules. The class with the highest weighted count is the output, which we refer to as $R_E(x)$. The following result is straightforward:

Proposition 1. *For any ensemble trees model E and any data instance x , suppose R_E is the set of decision rules of E derived by the method above, then $E(x) = R_E(x)$.*

In the sequel, we shall denote the original ensemble trees model as E and the converted set of decision rules as R_E .

3.2 Simplification of Decision Rules

As discussed in Section 1, some ensemble trees models used in real-life applications are enormous and complex. Consequently, the converted set of decision rules for such a model consists of a huge number of rules and each rule may be a very long formula. To reduce the complexity of the explanation, we consider simplifying the set of decision rules in two dimensions: the length of the rules and the number of the rules.

Continuing from the output R_E of Section 3.1, each formula in R_E is a branch formula, which we can simplify using a node filter as step one.

Parameter 1 (Node Filter θ) We measure the “quality” of a node (NQ) by information gain (IG):

$$NQ(n) = IG(n) \quad (6)$$

where n is the target node. We scan the nodes of each branch in each tree, and remove the nodes with NQ below θ , which is a real positive number.

The second step is to simplify each branch formula by merging the nodes.

Lemma 1. For any branch formula $(\bigwedge F_i) \rightarrow D$, where each F_i is the formula associated with a node on the branch, the left-hand side can be simplified to conjunction normal form (CNF) with at most n conjuncts, where n is the number of features of the dataset.

Proof. By the construction of decision trees, for any two conjuncts $v \geq l$ and $v \geq l'$ over a numeric feature v that appears in $\bigwedge F_i$, we can simplify them into one conjunct $v \geq l''$ where $l'' = \max(l, l')$.

For any two conjuncts $v \geq l$ and $\neg(v \geq l')$ over v , if $l < l'$, then they can be simplified to $l \leq v < l'$. Otherwise, we have $l \geq l'$. In this case, the two conjuncts do not have an intersection, and there would be no instances at the leaf node and the training algorithm should not let this case happen.

For any two conjuncts $\neg(v \geq l)$ and $\neg(v \geq l')$ over v , they can be simplified to $v < l''$ where $l'' = \min(l, l')$.

For any two conjuncts $v' \in C$ and $v' \in C'$ over a nominal feature v' , they can be simplified into one conjunct $v' \in C''$ where $C'' = C \cap C'$. If any of the two conjuncts is negated, e.g., $\neg(v' \in C)$, we can simply take the complement set $C''' = C_v \setminus C$, where C_v is the full set of permitted discrete values of the feature v , and convert the negated conjunct into $v' \in C'''$. The remainder of the proof is analogous.

Thus, the left-hand side of the branch formula can be simplified to one conjunct per feature. \square

Note that the simplification of Lemma 1 preserves logical equivalence of the set of decision rules while the other steps of this section do not. One may ask why do we use the node filter when Lemma 1 can merge the nodes. The reason is that using all the nodes in a branch may result in very narrow and focused explanations, so we use θ as a parameter to adjust the scope.

The above two steps aim to shorten the decision rules. The third step is to reduce the number of rules by filtering out those of low quality.

Parameter 2 (Rule Filter ϕ) We measure the “quality” of a decision rule (RQ) by the following formula:

$$RQ(r) = \frac{(\log_2(m) - H(l_r))}{\log_2(m)} \times Acc \quad (7)$$

where r is the target rule, m is the number of classes, $H(l_r)$ is the entropy of the leaf of the rule, and Acc is the accuracy of the corresponding tree on the OOB dataset. We remove a rule if its RQ is less than ϕ , which is a real number between 0 and 1.

The fourth step merges decision rules into groups of the same class signature.

Parameter 3 (Leaf Merger ψ) Given a vote distribution (n_1, \dots, n_m) , where m is the number of classes, we convert the distribution into ratios (ξ_1, \dots, ξ_m) where each $\xi_i, 1 \leq i \leq m$, is the ratio of class i in the leaf node. The class signature of this leaf node is defined as the tuple $(\lceil \xi_i/\psi \rceil, \dots, \lceil \xi_m/\psi \rceil)$, where ψ is a real number between 0 and 1.

Using the above definition, we divide the set of decision rules into a set of sets $\{G_1, \dots, G_j\}$. Each $G_i, 1 \leq i \leq j$, contains the set of rules of the same class signature. Intuitively, a larger ψ yields fewer distinct equivalent classes/groups and vice versa. We use this parameter to control the number of groups in the final explanation.

For a large-scale random forest, the filtered rules are still a large-scale formula set. In the last step, we control the number of decision rules we get.

Parameter 4 (Size Filter k) In each G_i , we take the number of instances in the leaf node of each rule as the **weight of the rule**, and we select k rules in a weight-first manner.

Associating each rule with the weight is crucial because now the vote distribution has been converted into ratios, and we have lost the information of the number of votes in the distribution. The weight retains this information. For example, we would prefer a ratio of $(0.7, 0.3)$ with 100 votes to overwhelm a ratio of $(0.1, 0.9)$ with 10 votes.

To summarise, we denote the composition of the above steps as *Simp*, and write the simplification procedure as follows:

$$Simp(R_E, \theta, \phi, \psi, k) = R'_E \quad (8)$$

where R'_E is a set of sets of decision rules where each rule has a weight. We use R'_E as our explanation to the original model.




Fig. 3: An example of the prediction procedure using the simplified rules R'_E .

3.3 Prediction

To use R'_E in a classification task, let x be a data instance, we first find all the decision rules in R'_E that are satisfied by x , then multiply the class signature of each rule by the corresponding weight, and finally add up to get a tuple. The class with the largest value is the output. This procedure is denoted as $R'_E(x) = c$ where c is a class. We give an example in Fig. 3 where we put the satisfied rules in red boxes.

3.4 Optimal Explanations

Now we consider a step further: how to evaluate explanations and find the optimal one? Intuitively, a good explanation should meet the following two criteria:

- The classification behaviour of the explanation R'_E should be similar to the original model E .
- The explanation should be concise and small.


We use *fidelity* to measure the first criterion. Fidelity is defined as the degree of similarity between the predictions of R'_E and E on unseen data [21]. First, we take the test set without labels as D . Then we use the classification results on D from the original model E as the “ground truth”. Lastly, we evaluate the classification accuracy of the explanation R'_E on D as fidelity. The fidelity component is denoted as $N(R'_E, E, D)$.

The second criterion is *scale* which is measured by the total number of conjuncts in the rules of R'_E and is denoted as $L(R'_E)$.

The *score* of an explanation R'_E is defined as

$$S_{opt}(R'_E) = N(R'_E, E, D) \times \epsilon + \frac{1 - \epsilon}{1 + e^{5 \times (\frac{L(R'_E)}{m \times n} - 1)}} \quad (9)$$

where ϵ is a real number between 0 and 1, n is the number of features, and m is the number of classes. The intuition is that the score grows linearly with fidelity, but it drops significantly when the explanation is too large. The second



(a) The effect of θ and ϕ on *scale* coverage. (The original scale is 65247.)

(b) The effect of k on *fidelity*.

Fig. 4: The effect of parameters on the two criterion.

component is a sigmoid-shaped function. Also, a large ϵ puts more importance on fidelity, and a small ϵ puts more importance on scale.

The parameters in Section 3.2 introduce a great variety of potential explanations. Particularly, the two parameters θ and ϕ control the strictness of the node filter and the rule filter respectively. We experiment on the diabetes dataset [6] to explore the effect of parameters on the two criteria. Fig. 4(a) shows the effect of the θ and ϕ on scale. As θ and ϕ rise, we get significantly smaller explanations. Even when θ and ϕ are both 0, the scale has dropped by about 20% compare to the original model due to Lemma 1. We also explore the effect of k on fidelity. We set θ to 0.65 and ϕ to 0.7. Fig. 4(b) shows the average of 50 tests. The results show that a relatively small value of k can get a fidelity close to all the rules.

To obtain an optimal explanation, we use the *linearly decreasing inertia weigh particle swarm optimization* algorithm (LDIW-PSO) [24] to optimise the parameters mentioned above with the S_{opt} as the fitness, which is the objective function to be optimised. Then we apply the optimal parameters to *Simp* and obtain the optimal explanation R_{opt} . We refer to the above procedure as *OptExplain*.

3.5 Profile of Equivalent Classes

An explanation with high fidelity usually has large scale, while a concise explanation can allow users to quickly understand the predictive behavior of the model at the sacrifice of fidelity. Sometimes the latter is preferred to draw a high-level conclusion of the classification behavior. We propose a new method called the *Profile of Equivalent Classes (ProClass)*, which is a more concise description of classes based on the extracted decision rules.

The *ProClass* process is based on the R'_E obtained by Equation 8. The R'_E may contain several groups, each has no more than k rules. Then we merge the

decision rules in each group by solving a weighted maximum satisfiable (MAX-SAT) problem to get a new explanation R''_E . Different from Section 3.4, *ProClass* requires more rules as input, and the number of groups is equal to the number of classes. The number of groups in R''_E is denoted as $M(R''_E)$. For the above needs, we defined another *score* for the optimization in *ProClass*:

$$S_{pro}(R''_E) = (m - M(R''_E) + 1) \times L(R''_E) \quad (10)$$

where m is the number of classes. Using S_{pro} as fitness, we obtain the optimized result is denoted as R'_{opt} . An ideal R'_{opt} has m groups $\{G_1, \dots, G_m\}$. For each group G_i ($1 \leq i \leq m$), we associate the weight of each rule to each of its conjuncts, and send all the weighed conjuncts in the group to a SAT solver such as Z3 [18]. The solver will return a subset of satisfied conjuncts that maximise the total weight. After solving the MAX-SAT problem for each G_i , we obtain a subset of conjuncts. Then we simplify the conjuncts into one rule r'_i using Lemma 1. Performing the above steps on all groups, we get the profile of equivalent classes denoted as R_{pro} , which has the form

$$R_{pro} = \{r'_1 \rightarrow S_1, \dots, r'_m \rightarrow S_m\}.$$

where r'_i is the logic for predicting the class S_i .

4 Case Studies and Experiment

In this section, we demonstrate our method through case studies. We used scikit-learn to train random forest models. We implemented our method in Python and evaluated it on multiple datasets: adult, credit, diabetes, german, mnist, spambase, all of which are available on OpenML [19]. There are two important parameters in LDIW-PSO algorithm: particle size and iteration period. In our experiments, both particle size and iteration period are set to 20 by default. Experiments were conducted on a machine with an Intel Core i9-7960X CPU and 32GB RAM.

4.1 Case Study 1: Diabetes Prediction

We first evaluate *OptExplain* on diabetes dataset [6], which has 8 features, 2 classes, and 768 samples. The eight features are the number of times pregnant (preg), plasma glucose concentration (plas), diastolic blood pressure (pres), 2-hour serum insulin (insu), triceps skinfold thickness (skin), body mass index (mass), diabetes pedigree function (pedi) and age. We randomly select 100 samples as the testing set, and the rest as the training set. Then we train a random forest with 100 trees and unlimited depth.

We set ϵ to 0.9, and *OptExplain* will produce an explanation R_{opt} . The result is shown in Table 1 and Table 2. Recall that R_E is equivalent to the original model E (Proposition 1). Both R_E and R_{opt} have 80% accuracy on the test set. R_E has 11106 rules with 102584 conjuncts, while R_{opt} has 6 rules with

(ϕ, θ, ψ, k)	R_E		R_{opt}		
	scale	accuracy	scale	accuracy	fidelity
(0.55, 0.45, 0.83, 3.0)	102584	80%	6	80%	92%

Table 1: Optimized parameters and predictive performance of the explanation.

Groups	Class Signature (negative,positive)	Rules	Weight
Group1	(2.0 , 0.0)	$pedi \leq 0.7$	30
		$plas \leq 130.0$	23
		$plas \leq 157.5$	21
Group2	(0.0 , 2.0)	$mass > 28.7$	30
		$age > 27.5$	22
		$plas > 122.5$	20

Table 2: The optimal explanation R_{opt} .

6 conjuncts. The fidelity of R_{opt} is 92%, which means it is very similar to the original model. By observing the decision rules, users can analyze what role each feature plays in the prediction process. In this explanation, if an instance has $plas > 157.5$, then the chance of prediction positive diabetes is high.

negative	positive
$2.0 < preg \leq 2.5$	$7.5 < preg \leq 8$
$90 < plas \leq 91.5$	$173.5 < plas \leq 175$
$74 < pres \leq 74.5$	$70 < pres \leq 71$
$28.5 < skin \leq 29$	$23.5 < skin \leq 24$
$61.5 < insu \leq 63$	$126.5 < insu \leq 127.5$
$25.9 < mass \leq 26$	$32.9 < mass \leq 33$
$pedi = 0.2$	$pedi = 0.5$
$23.5 < age \leq 24$	$33 < age \leq 33.5$

Table 3: The profile of equivalent classes of a random forest model for diabetes.

Table 3 gives the profile of equivalent classes derived from *ProClass*. The profile is divided into two groups corresponding to equivalent class labels: negative diabetes and positive diabetes. Medical practitioners can quickly obtain the salient feature values corresponding to each class according to the profile.

4.2 Case Study 2: Digit Recognition (MNIST)

In order to visually demonstrate the profile, we use the MNIST dataset [19] to illustrate *ProClass*. The MNIST dataset contains 70,000 images of size 28×28

pixels. In order to reduce the complexity, we randomly picked 10,000 samples from the dataset as the training set, and train a random forest with 100 trees and 15 max-depth. *ProClass* produces a profile R_{pro} which gives the range of some pixels. Then we take the median of the range as the pixel value. The remaining pixels that do not appear in the profile are set to light blue.




Fig. 5: As visualisation of R_{pro} on MNIST.

From Fig. 5 we can see the profile’s visualization of each digit. For peripheral pixels that are not significant in the prediction, the profile will not give a description. As we can see, the profiles indeed yield human-understandable visualisation of each class.

4.3 Experiment and Comparison

We compared the proposed method to a baseline method — Hara and Hayashi’s approach named defragTrees [11]. In their work, defragTrees has been compared with BATrees [1], inTrees [5] and Node Harvest [17]. Their result suggests that defragTrees generated smaller set of rules with higher fidelity than the other methods. We choose the following datasets: adult, credit, diabetes, german, spambase [6]. We split the datasets into two subsets at random: a 70% training set, and a 30% testing set. Then we train the ensemble trees with 100 trees and 10 max-depth for *OptExplain* and defragTrees. For *OptExplain*, we set two experimental groups and set ϵ as 0.5 and 0.9 respectively.

We compare the two methods with respect to the two criteria: scale and fidelity. We conducted the experiment over ten random data realizations for each dataset. Table 4 shows the average values of the ten tests. The number in bold is the best scale/fidelity value for each dataset. Table 4 shows that the scale of explanations generated by *OptExplain* with 0.5 ϵ is the smallest except on the german dataset, and the fidelity is generally better than defragTrees. The small-scale explanation on german generated by defragTrees has much lower fidelity.

Data	defragTrees		<i>OptExplain</i>			
			$\epsilon=0.5$		$\epsilon=0.9$	
	scale	fidelity%	scale	fidelity%	scale	fidelity%
adult	65.5	81.9	37.2	87.0	43.2	88.4
credit	8.7	85	2.8	94.2	9.8	94.6
diabetes	14.2	74.8	4	85.8	8.3	88.3
german	5.4	69.8	19	87.3	50.2	88.7
spambase	82	91.7	22	91.6	40.6	92.9

Table 4: *OptExplain* and defragTrees

OptExplain with 0.9 ϵ can generate the highest fidelity explanation with similar scale than defragTrees. In both settings of ϵ , *OptExplain* generates superior explanations in most cases.




Fig. 6: Mean elapsed time comparison. (The y-axis is logarithmic scale.)

Finally, we visually compare the computation time in Fig. 6, and show the mean time of 10 tests. We run *OptExplain* on models of 100 trees with depth of 10, 20, and 30 respectively. The method defragTrees on 100 trees and depth 10 is also used as a reference. The results show that *OptExplain* has better computational performance than defragTrees on adult, diabetes and spambase. *OptExplain* can also deal with very large models, and for the above three datasets, *OptExplain* can generate explanations for trees of depth 30 faster than defragTrees can for depth 10. For a concrete example, the mean computation time of *OptExplain* on the adult dataset is 1,949 s, while the time of defragTrees is 5,797 s (both depth 10).

5 Related Work

There are a variety of approaches for tackling machine learning interpretability. Some recent and popular ones are focused on local explanations. LIME [22] is

such a tool that finds linear approximations of the prediction model and gives importance weights for certain predicates used in the prediction. Anchors [23] generates “if-then” style explanations for predictions. Such explanations have similar forms as our decision rules and are considered intuitive and easy to understand by the user. Shapley (SHAP) Values [16] are often used to extract importance scores and impacts on features. Like LIME, SHAP also provides user-friendly graphical presentations (e.g., bar charts) for explaining predictions. It should be noted that SHAP can also be used to obtain global feature importance. The above three methods are *model-agnostic*, which means that in the process of providing explanations and making machine learning more “white-box”, they take prediction models as a “black-box” and attempt to find patterns of *features* when the model makes predictions. An advantage is that they can be applied to different machine learning techniques, including ensemble trees and neural networks. CHIRPS [12] is another technique for local explanations. In contrast to the above techniques, CHIRPS looks into decision trees and uses frequent pattern mining on decision nodes to obtain decision rules as the explanation.

Global explanations are more closely related to this work. Recent examples include Hara and Hayashi’s approach [11] that uses Bayesian model selection to extract decision rules. Deng’s inTrees [5] extracts, selects and prunes rules from a set of decision trees and uses frequent pattern analysis to summarise rules into a smaller prediction model.

Most relevant techniques come from a statistics perspective. Their simplification process often consists of selection, pruning and frequency analysis. By contrast, our work comes from a logician’s point of view and, on top of the usual selection and pruning operations, uses automated reasoning to simplify logical formulae and find abstractions of equivalent classes. We see many related methods as complementary ones rather than competitors because they output in different forms. For example, one can combine SHAP values and our work to form a more comprehensive explanation.

6 Conclusion and Future Work

This paper presents a streamlined procedure for extracting optimal *logical* explanations from ensemble trees models. As an additional feature, our method can also output the “profile” of each class so that the user can *see* how the model predicts in different cases. We give two case studies to illustrate how our method works and show that our method outperforms state-of-the-art through experimental results.

As future work, we plan to use this work to perform efficient verification tasks. In particular, we will use the optimal explanation as an approximation of the original model and try to “debug” the model against user-specified properties. The rules that violate the property can serve as constraints to narrow down the search space when finding counterexamples. Another important future direction is to convert deep neural networks to ensemble trees and extend this work to explain deep learning.

References

1. Breiman, L., Shang, N.: Born again trees (1996)
2. Breiman, L.: Random forests. *Machine Learning* **45**(1), 5–32 (Oct 2001)
3. Bride, H., Dong, J., Dong, J.S., Hóu, Z.: Towards dependable and explainable machine learning using automated reasoning. In: *Formal Methods and Software Engineering - 20th International Conference on Formal Engineering Methods, ICFEM 2018, Gold Coast, QLD, Australia, November 12-16, 2018, Proceedings*. pp. 412–416 (2018)
4. Cui, Z., Chen, W., He, Y., Chen, Y.: Optimal action extraction for random forests and boosted trees. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 179–188. KDD '15, ACM, New York, NY, USA (2015). <https://doi.org/10.1145/2783258.2783281>, <http://doi.acm.org/10.1145/2783258.2783281>
5. Deng, H.: Interpreting tree ensembles with intrees. *International Journal of Data Science and Analytics* **7**(4), 277–287 (2019)
6. Dua, D., Graff, C.: UCI machine learning repository (2017), <http://archive.ics.uci.edu/ml>
7. Freund, Y., E Schapire, R.: A short introduction to boosting. *Journal of Japanese Society for Artificial Intelligence* **14**, 771–780 (10 1999)
8. Friedman, J.H.: Greedy function approximation: A gradient boosting machine. *Annals of Statistics* **29**, 1189–1232 (2000)
9. Friedman, J.H.: Stochastic gradient boosting. *Computational statistics & data analysis* **38**(4), 367–378 (2002)
10. Goré, R., Olesen, K., Thomson, J.: Implementing tableau calculi using bdds: Bddtab system description. In: *Automated Reasoning - 7th International Joint Conference, IJCAR 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 19-22, 2014. Proceedings*. pp. 337–343 (2014)
11. Hara, S., Hayashi, K.: Making tree ensembles interpretable: A bayesian model selection approach. In: *International conference on artificial intelligence and statistics*. pp. 77–85. PMLR (2018)
12. Hatwell, J., Gaber, M.M., Azad, R.M.A.: Chirps: Explaining random forest classification. *Artificial Intelligence Review* **53**, 5747–5788 (2020)
13. Hinton, G., Frosst, N.: Distilling a neural network into a soft decision tree (2017), <https://arxiv.org/pdf/1711.09784.pdf>
14. Ho, T.K.: Random decision forests. In: *Proceedings of 3rd international conference on document analysis and recognition*. vol. 1, pp. 278–282. IEEE (1995)
15. Ltd, D.I.P.: Silas. <https://depintel.com/silas/> (2018)
16. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 4765–4774. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
17. Meinshausen, N.: Node harvest (2009)
18. de Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems*. pp. 337–340. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
19. OpenML: openml.org. <https://www.openml.org> (Accessed 2019)
20. Pafka, S.: A minimal benchmark for scalability, speed and accuracy of commonly used open source implementations of the top machine learning algorithms for binary classification. <https://github.com/szilard/benchm-ml> (2018)

21. Papenmeier, A., Englebienne, G., Seifert, C.: How model accuracy and explanation fidelity influence user trust. arXiv preprint arXiv:1907.12652 (2019)
22. Ribeiro, M.T., Singh, S., Guestrin, C.: "why should I trust you?": Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016. pp. 1135–1144 (2016)
23. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
24. Shi, Y., Eberhart, R.C.: Empirical study of particle swarm optimization. In: Proceedings of the 1999 Congress on Evolutionary Computation-CEC99 (Cat. No. 99TH8406). vol. 3, pp. 1945–1950 Vol. 3 (1999). <https://doi.org/10.1109/CEC.1999.785511>
25. Sun, J., Liu, Y., Dong, J.S., Pang, J.: PAT: Towards flexible verification under fairness. Lecture Notes in Computer Science, vol. 5643, pp. 709–714. Springer (2009)
26. Törnblom, J., Nadjm-Tehrani, S.: Formal Verification of Random Forests in Safety-Critical Applications: 6th International Workshop, FTSCS 2018, Gold Coast, Australia, November 16, 2018, Revised Selected Papers, pp. 55–71 (01 2019)