

Supervised Feature Compression Based on Counterfactual Analysis

Piccialli, Veronica; Romero Morales, Dolores ; Salvatore, Cecilia

Document Version
Final published version

Published in:
European Journal of Operational Research

DOI:
[10.1016/j.ejor.2023.11.019](https://doi.org/10.1016/j.ejor.2023.11.019)

Publication date:
2024

License
CC BY

Citation for published version (APA):
Piccialli, V., Romero Morales, D., & Salvatore, C. (2024). Supervised Feature Compression Based on Counterfactual Analysis. *European Journal of Operational Research*, 317(2), 273-285.
<https://doi.org/10.1016/j.ejor.2023.11.019>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 23. Apr. 2025





Supervised feature compression based on counterfactual analysis

Veronica Piccialli^a, Dolores Romero Morales^b, Cecilia Salvatore^{c,*}

^a Department of Computer Control and Management Engineering Antonio Ruberti, Sapienza University of Rome, Rome, Italy

^b Department of Economics, Copenhagen Business School, Frederiksberg, Denmark

^c Department of Civil Engineering and Computer Science, University of Rome Tor Vergata, Rome, Italy

ARTICLE INFO

Keywords:

Machine learning
Supervised classification
Interpretability
Feature compression
Counterfactual analysis

ABSTRACT

Counterfactual Explanations are becoming a de-facto standard in post-hoc interpretable machine learning. For a given classifier and an instance classified in an undesired class, its counterfactual explanation corresponds to small perturbations of that instance that allows changing the classification outcome. This work aims to leverage Counterfactual Explanations to detect the important decision boundaries of a pre-trained black-box model. This information is used to build a supervised discretization of the features in the dataset with a tunable granularity. Using the discretized dataset, an optimal Decision Tree can be trained that resembles the black-box model, but that is more interpretable and compact. Numerical results on real-world datasets show the effectiveness of the approach in terms of accuracy and sparsity.

1. Introduction

Classification systems based on Machine Learning algorithms are often used to support decision-making in real-world applications such as healthcare (Babic, Gerke, Evgeniou, & Cohen, 2021), credit approval (Bastos & Matos, 2022; Dumitrescu, Hué, Hurlin, & Tokpavi, 2022; Kozodoi, Jacob, & Lessmann, 2022; Silva, Pereira, & Magalhães, 2022), or criminal justice (Ridgeway, 2013). These systems often act as black-boxes that lack of interpretability. Making Machine Learning systems trustworthy has become imperative, and interpretability, robustness, and fairness are often essential requirements for deployment (European Commission, 2020; Goodman & Flaxman, 2017; Rudin, Chen, Chen, Huang, Semenova, & Zhong, 2022).

This paper is devoted to enhancing the interpretability of black-box classifiers. Without loss of generality, we will focus on binary classification problems. We are given a black-box classification model, hereafter, the target model. The goal is to discretize the features by detecting their most critical values. Discretization techniques have often been proposed in the literature as a preprocessing step that allows the transformation of continuous data into categorical ones (Dash, Paramguru, & Dash, 2011; Dougherty, Kohavi, & Sahami, 1995; García, Luengo, Sáez, López, & Herrera, 2013; Ramírez-Gallego et al., 2016); the objective is to make the representation of the knowledge more concise. Note that a trivial way to discretize continuous features consists of creating a dummy variable for each splitting value along a feature; a splitting value is defined as the middle value between two consecutive training points on a feature. However, this results in a

very large number of dummy variables. In our approach we aim at a meaningful compressed feature representation, that gives us a more interpretable view of the data since the only relevant values are the critical ones defining the discretization. In addition, this procedure acts as a Feature Selection technique (Piramuthu, 2004), meaning that features on which no critical values are extracted are considered not important for detecting the input–output relationship, and can thus be filtered out. Discretizing features can also help in reducing noise that can cause overfitting to training data. Finally, we make it more affordable to train an optimal classification tree (Carrizosa, Molero-Río & Romero Morales, 2021) as the interpretable surrogate of the target model, and to certify its optimality, thanks to the reduction in the number of thresholds that need to be considered, namely, the critical ones.

In this paper, the feature discretization will be supervised by a Counterfactual Analysis on the target model. Counterfactual Analysis is a post-hoc local explainability technique (Karimi, Barthe, Schölkopf, & Valera, 2022; Martens & Provost, 2014; Molnar, Casalicchio, & Bischl, 2020; Wachter, Mittelstadt, & Russell, 2017) that has gained a lot of attraction especially in Supervised Classification. The starting point is an already trained classification model and an observation that has been classified by the model in the undesired class, e.g., as a *bad* payer in a credit approval application (Doumpos, Zopounidis, Gounopoulos, Platanakis, & Zhang, 2023; Fethi & Pasiouras, 2010). Counterfactual Analysis provides feedback on how to change the features of the observation in order to change the prediction given by the model to the

* Corresponding author.

E-mail address: cecilia.salvatore@uniroma2.it (C. Salvatore).

desired class, e.g., as a *good* payer in the credit approval application above. The Counterfactual Explanation depends on the cost function used to measure the changes, which is minimized, and the constraints imposed on the explanation, such as upper bounds on the changes to continuous variables or the correct modeling of changes to categorical features.

We propose to use Counterfactual Analysis to detect important decision boundaries of the target model. For this, we extract a set of univariate decision boundaries, i.e., axis-parallel hyperplanes, from Counterfactual Explanations. We use observations that have been correctly classified by the target model. To each of those, we find their counterfactual explanation, yielding a set of axis-parallel hyperplanes. By putting together all these hyperplanes from all the counterfactual explanations, we can identify for each feature a set of thresholds. This step allows us to derive a meaningful supervised discretization of the original dataset, where the splitting values on each feature are the thresholds, and hence are related to the decision boundary of the target model. With this, we can exploit the important decision boundaries to build a classification tree.

A major strength of our approach is that our supervised discretization is based on an optimization problem, the Counterfactual Explanation Problem. First, this allows us to control the importance of the features extracted, and hence the granularity of the discretization, through the counterfactual cost function. Furthermore, we can impose desirable properties of the boundary, by modeling constraints on the counterfactuals, like plausibility or fairness (Carrizosa, Ramírez-Ayerbe, & Romero Morales, 2023; Maragno, Röber, & Birbil, 2022).

In summary, given a black-box classification model, Counterfactual Explanations can help us to find a supervised discretization where the compression focuses on the important decision boundaries of the target model. With this new representation of the data we can build an interpretable surrogate model, namely an optimal univariate decision tree. Our approach has the following advantages:

- We produce a supervised discretization of the original dataset whose granularity can be tuned. The discretization is driven by an optimization problem, the Counterfactual Explanation Problem, that allows us to control the importance of the thresholds extracted and to impose properties on the boundary.
- The interpretable model that we build uses thresholds extracted from the black-box model. This guarantees the use of features and thresholds that are relevant for the problem, since they are part of the decision boundaries of a more complex classification model.
- Our methodology is suitable for any type of classifier for which the Counterfactual Explanation Problem can be solved efficiently, such as tree ensemble models (e.g., random forests, gradient boosting) or linear models (e.g., logistic regression, linear support vector machines).
- Our approach allows to strongly compress the original dataset, making it feasible to train an optimal classification tree and to certify its optimality (Lin, Zhong, Hu, Rudin, & Seltzer, 2020). The reduced dataset can be used to train other machine learning models.
- Once all the thresholds have been computed using a counterfactual analysis, the user can easily visualize the accuracy and granularity tradeoff by changing one parameter of the discretization. Tuning the granularity just results in selecting a subset of the thresholds computed by counterfactual analysis and is thus computationally cheap.
- Our discretization shows a good tradeoff between out-of-sample accuracy and standard metrics for evaluating discretization procedures (García et al., 2013), such as compression rate and inconsistency rate.

The remainder of the paper is structured as follows. In Section 2 we analyze the existing literature on Decision Trees, Counterfactual

Explanations and discretization procedures. In Section 3 we formalize our method. In Section 4 we describe the experimental setup and the obtained results. Finally, in Section 5 we draw some conclusions, and propose some lines for future research.

2. Literature review

The most popular and inherently interpretable models are univariate Decision Trees (Carrizosa, Molero-Río et al., 2021). Training an optimal Decision Tree is known to be an NP-complete problem (Laurent & Rivest, 1976). Many approaches for training Decision Trees thus rely on heuristics; the most used heuristic is based on a top-down greedy strategy in which the tree structure is recursively grown from the root to the terminal nodes (Breiman, Friedman, Stone, & Olshen, 1984; Quinlan, 1986, 2014). However, because of their greedy nature, these heuristics may result in poor generalization capabilities.

Recently, there has been an increasing interest in developing mathematical optimization formulations and numerical solution approaches to train Optimal Classification Trees (Carrizosa, Molero-Río et al., 2021). One of the first approaches for Optimal Classification Trees was proposed in Bertsimas and Dunn (2017), where a Mixed-Integer Linear Programming problem (MILP) for both univariate and multivariate classification trees is formulated. This formulation, however, has a weak linear relaxation and, thus, it is hard to solve to optimality on real-sized instances; the authors thus propose a local search algorithm to try to overcome this issue, returning a local solution that is not guaranteed to be the global minimum (Dunn, 2018). Recently, new approaches based on binary features have been proposed, and ad-hoc algorithms have been designed with the purpose of being scalable and fast (Aghaei, Gómez, & Vayanos, 2021; Günlük, Kalagnanam, Li, Menickelly, & Scheinberg, 2021; Lin et al., 2020; Verwer & Zhang, 2019). Handling continuous features thus requires to introduce a set of dummy binary variables. The dynamic programming approach proposed in Lin et al. (2020) is better suited for realistic datasets with a few continuous features, thanks to a new representation of the dynamic programming search space.

Still, the number of continuous features that can be handled is limited as each continuous feature is associated with a large number of dummy variables, thus enlarging the search space and slowing down the optimization process. In order to extend the applicability of the procedure, the authors propose some strategies to exploit knowledge extracted from black-box models to make the Decision Tree's optimization faster (McTavish et al., 2022). The most effective strategy is the "Guessing Thresholds", which allows the user to limit the number of dummy variables modeling continuous features. The "Guessing Thresholds" uses Tree Ensembles as a black-box model, and it considers as thresholds the ones used in the various splits of each tree in the ensemble; the thresholds are sorted by their importance (e.g. their Gini index) so that the least important can be removed; the Tree Ensemble is then re-fitted with the remaining features, and the procedure is repeated until there is a too large drop in the training performance. The idea of leveraging black-box models, and in particular Tree Ensembles, to build Decision Trees is also proposed in Vidal and Schiffer (2020), in which the authors design a dynamic programming approach where a single Decision Tree is built to reproduce exactly the decision function of a Tree Ensemble; however, in general, the resulting Decision Tree can be large and thus can lose some interpretability.

Our research collocates in between the approach in Vidal and Schiffer (2020) and the one in McTavish et al. (2022). We use optimization as in Vidal and Schiffer (2020), but we do not reproduce exactly the original black-box model. Instead, we exploit the boundary of the black-box model to compress the original dataset in a meaningful way by means of optimization. Then, we use the compressed dataset to efficiently train an optimal decision tree as in McTavish et al. (2022). We force interpretability by limiting the depth of the optimal decision tree. Our discretization is driven by a procedure based on optimization for

computing Counterfactuals Explanations, that are local explainability techniques, that can be used to explain a single decision of a black-box model. Indeed, Counterfactual Explanations allow providing feedback to users on how to change their features in order to change the outcome of the decision (Guidotti, 2022; Karimi et al., 2022). Formally, the Counterfactual Explanation of a datapoint x^0 , namely x^{CE} , is defined as the perturbation of minimal cost (w.r.t. some cost function) that allows changing the classification outcome. An optimization problem for computing the Counterfactual Explanations x^{CE} of a given point x^0 , namely the Counterfactual Explanation Problem was proposed in Wachter et al. (2017):

$$\arg \min_{x^{CE}} C(x^0, x^{CE}) \quad \text{s.t.} \quad f(x^{CE}) = y^{CE}, \tag{1}$$

where C is a cost function, f is the classification function and $y^{CE} \neq f(x^0)$ is the required label for the Counterfactual Explanation. The problem is then reformulated as an unconstrained problem with a differentiable objective function, composed of two terms: the first term of the objective requires the classification function to be as close as possible to the required label y^{CE} , while the second term requires minimizing the distance between the Counterfactual Explanation and the initial point.

Later in the literature, Verma, Boonsanong, Hoang, Hines, Dickerson, and Shah (2022) identifies some additional constraints Counterfactual Explanations should satisfy. Proximity requires that a valid Counterfactual Explanation must be a *small* change with respect to the initial point. Actionability implies that the Counterfactual Explanation can modify some features (e.g., income), while others must be *immutable* (e.g., sex, race). Sparsity requires a Counterfactual Explanation to be *sparse*, i.e., as few features as possible should change. This makes the Counterfactual Explanation more effective because simpler explanations can be better understood by users. Data Manifold Closeness suggests that a Counterfactual Explanation should be *realistic*, meaning that it should be close to training data. Finally, Causality requires that the Counterfactual Explanation adheres to observed *correlations* between features. Examples of constraints modeling domain knowledge and actionability can be found in Maragno et al. (2022), Parmentier and Vidal (2021). Robustness is also an important requirement, as it imposes that Counterfactual Explanations remain valid after retraining the model or slightly perturbing the input features (Fernández, Aledo, Gamez, & Puerta, 2022; Forel, Parmentier, & Vidal, 2022; Maragno et al., 2023).

For some classifiers such as Linear Support Vector Machines (SVMs) or Tree Ensembles, it is possible to derive an explicit expression of the classification function f , allowing to directly write Problem (1) as a Linear Programming problem or at most a convex quadratic problem. Integer variables can be added to represent the l_0 -norm in the objective function. For SVMs with non-linear kernels or for Neural Networks it is not possible to do so. To overcome this issue, Maragno et al. (2022) suggest the idea that the Counterfactual Explanation problem is a special case of Optimization with Constraint Learning, in which some of the constraints are learnt through a predictive model.

Recent literature states that Counterfactual Explanations can provide useful insights into the classification model, allowing them to be used not only for post-hoc explainability but also for debugging and detecting bias in models (Sokol & Flach, 2019). For example, if it turns out that without imposing the actionability constraints the Counterfactual Explanation changes a sensitive feature (e.g., gender) by saying that a woman would receive the loan if she were a man, then the classification model may be biased. This observation highlights that Counterfactual Explanations can be used to detect biases in Machine Learning models, opening the possibility of designing new fairness metrics that rely on Counterfactual Explanations (Goethals, Martens, & Calders, 2023; Kusner, Loftus, Russell, & Silva, 2017).

Some recent literature (Aïvodji, Bolot, & Gambs, 2020; Kuppa & Le-Khac, 2021; Mothilal, Sharma, & Tan, 2020; Zhao, Zhang, Xiao, &




Fig. 1. Closeness of a Counterfactual Explanation to the Decision Boundaries of the Target model produced by a Random Forest projected on two features (Income and # years of credit).

Lim, 2021) focuses on the use of Counterfactual Explanations in an adversarial setting for detecting the decision boundaries of machine learning models.

In this work, we propose to use Counterfactual Explanations to enhance the interpretability of the model itself. We show that using Counterfactual Explanations to detect the important decision boundaries of a black-box model allows us to design a supervised discretization technique that helps to build an optimal Decision Tree. Using a discretization technique as a preprocessing step can lead to many advantages: (1) some Machine Learning algorithms prefer categorical variables, e.g., the Naive Bayes classifier (Flores, Gámez, Martínez, & Puerta, 2011; Yang & Webb, 2009); (2) discretized data are easier to understand and to explain; (3) discretization can decrease the granularity of data, potentially decreasing the noise in the dataset (García et al., 2013). Nevertheless, any discretization process generally leads to a loss of information, making it crucial the minimization of the loss of information. In García et al. (2013), a taxonomy for categorizing discretizing methods is introduced. The effectiveness of a discretization procedure can then be evaluated according to different aspects: (a) the discretization should be able to compress the information as much as possible, by detecting as few intervals as possible; (b) the inconsistency rate produced by the discretization, i.e., the unavoidable error due to multiple points associated with the same discretization but with different labels; (c) the classification rate obtained by an algorithm trained on discretized data compared with one provided by the same algorithm on the initial representation of the dataset.

3. Method

We assume we have at hand a binary-classification training set:

$$D_{tr} = \{(x_i, y_i) : x_i \in \mathfrak{X}^m, y_i \in \{0, 1\} \quad \forall i = 1 \dots n_{tr}\},$$

taken out from a sample of n individuals. This procedure can, however, be easily extended to multi-class problems. For simplicity, we assume w.l.o.g. all features to be scaled between 0 and 1. Categorical features are assumed to be either preprocessed via the one-hot encoding procedure or treated as ordinal and then discretized.

We train a black-box model \mathcal{T} using D_{tr} , that we use as the target model for our procedure. For our purpose, \mathcal{T} can implement any classification algorithm as long as we can efficiently compute the Counterfactual Explanation associated with each instance $x \in D_{tr}$ with respect to model \mathcal{T} . Our objective is to train a small, compact and interpretable decision tree that acts as a surrogate for \mathcal{T} . At each branch node s , a univariate Decision Tree takes a univariate decision: if the input x on a given feature j is less or equal than a threshold τ_s the point is directed to the left child of s , otherwise to the right child.

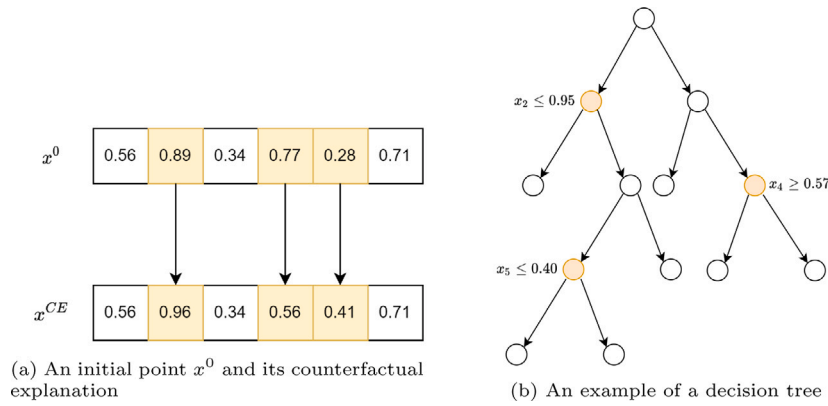


Fig. 2. Assume that x^{CE} is computed by solving Problem (1) with x^0 as input for a given Tree Ensemble acting as a target model. A perturbation ($\pm\epsilon$) of the values of the features where x^0 and x^{CE} differ can be used as splitting values for the nodes of a univariate Decision Tree.

By following the path of each point x from the root of the tree to the last level, a set of points is assigned to each leaf l of the tree; the classification outcome at each leaf l depends on the most frequent label of the points assigned to l . Training a Decision Tree results in choosing which feature to consider at each node s and the value of the threshold τ_s .

As shown in Fig. 1, the intuition behind our procedure is that a Counterfactual Explanation is very close ($\pm\epsilon$) to some decision boundaries of the target model. So, if we generate a large set of Counterfactual Explanations, we should be able to identify the most critical decision boundaries of the model and use them to guide the training procedure of the surrogate Decision Tree. Each Counterfactual Explanation x^{CE} perturbs just a subset of the features of the corresponding initial point x^0 (Fig. 2(a)); the values these features assume can be used to mimic some decision boundaries of the Target model, i.e., these values can be used as splitting values in the nodes of the surrogate Decision Tree (Fig. 2(b)).

Following Carrizosa, Ramírez Ayerbe, and Romero Morales (2024) for computing Counterfactual Explanations, we solve a problem derived from (1) under the assumption that the Target model is a Tree Ensemble. This problem takes the following parameters: $y^{CE} \in \{0, 1\}$, with $y^{CE} \neq y^0$, is the required outcome for the Counterfactual Explanation; T is the set of trees in the Tree Ensemble; \mathcal{L}_t and \mathcal{N}_t denote respectively the leaves and the internal nodes of each tree t in the tree ensemble; $v_{t,s}$ and $c_{t,s}$ denote respectively which feature and which threshold is used to split at node s in tree t ; $\mathcal{A}_L(t, l)$ and $\mathcal{A}_R(t, l)$ denote respectively the ancestors s of leaf l in tree t whose left/right path leads from s to l . We denote by $p_k(x)$ the probability for point x of being in class k ; the specific expression of $p_k(x)$ depends on the tree ensemble method used (e.g. Random Forest (Breiman, 2001) or Gradient Boosting (Chen et al., 2015)). The variables define the point x^{CE} and a binary variable $z_{t,l}$ that denotes the assignment of x^{CE} to one of the leaves for each of the trees in the tree ensemble. The formulation is the following:

$$\min_{x^{CE}} C(x^0, x^{CE}) := \lambda_0 \|x^0 - x^{CE}\|_0 + \lambda_1 \|x^0 - x^{CE}\|_1 + \lambda_2 \|x^0 - x^{CE}\|_2^2 \quad (2)$$

$$x_{v_{t,s}}^{CE} - M_0(1 - z_{t,l}) + \epsilon_{v_{t,s}} \leq c_{t,s} \quad \forall t \in T, l \in \mathcal{L}_t, s \in \mathcal{N}_t : s \in \mathcal{A}_L(t, l) \quad (3)$$

$$x_{v_{t,s}}^{CE} + M_1(1 - z_{t,l}) - \epsilon_{v_{t,s}} \geq c_{t,s} \quad \forall t \in T, l \in \mathcal{L}_t, s \in \mathcal{N}_t : s \in \mathcal{A}_R(t, l) \quad (4)$$

$$\sum_{l \in \mathcal{L}_t} z_{t,l} = 1 \quad \forall t \in T \quad (5)$$

$$p_{y^{CE}}(x^{CE}) \geq p_{y^0}(x^{CE}) + \epsilon \quad (6)$$

$$x^{CE} \in \mathcal{X}^0. \quad (7)$$

In the objective function, we consider the weighted combination (with non-negative coefficients) of l_0 -, l_1 - and l_2 -norm; the l_0 -norm is used

for feature sparsity while the l_1 - and l_2 - norms are used to measure proximity. The l_0 - and l_1 -norms are modeled in a standard way by introducing respectively binary variables and linear constraints. By increasing the value of parameter λ_0 , we encourage changing as few features as possible. Both M_0 and M_1 can be set to 1 since the features are assumed to be scaled between 0 and 1. Constraint (7) imposes that the Counterfactual Explanation should belong to a set \mathcal{X}^0 that represents a plausibility set for the initial point x^0 ; actionability, data manifold closeness, causality and other requirements can thus be expressed via additional constraints. Eqs. (3)–(6) impose that the label assigned to the Counterfactual Explanation should be the required label y^{CE} ; in this set of constraints, binary variables z model the assignment of x^{CE} to one of the leaves, for each tree in the tree ensemble. We can notice that constraints (3) and (4) depend on a threshold ϵ_j for each feature $j \in [1 \dots m]$. In our paper, we set the value of ϵ_j as the smallest difference between the two consecutive values that feature j assumes on the datapoints of D_{tr} . For the value of ϵ that appears in constraint (6), we use a fixed value, while in Forel et al. (2022) it is analyzed how to set it for requiring Counterfactual Explanation robustness in Tree Ensembles. The expression of constraint (6) depends on the specific tree ensemble method. We report the expressions for Random Forest:

$$\frac{1}{|T|} \sum_{t \in T} \sum_{l \in \mathcal{L}_t} w_{t,l,y^{CE}} z_{t,l} \geq \frac{1}{|T|} \sum_{t \in T} \sum_{l \in \mathcal{L}_t} w_{t,l,y^0} z_{t,l} + \epsilon, \quad (8)$$

where $w_{t,l,k}$ denotes the classification weight of leaf l in tree t for class k , and Gradient Boosting:

$$\begin{cases} \omega_0 + lr \sum_{t \in T} \sum_{l \in \mathcal{L}_t} w_{t,l} z_{t,l} \geq \epsilon & \text{if } y^{CE} = 1 \\ \omega_0 + lr \sum_{t \in T} \sum_{l \in \mathcal{L}_t} w_{t,l} z_{t,l} \leq -\epsilon & \text{if } y^{CE} = 0, \end{cases} \quad (9)$$

where ω_0 is the initial prediction of Gradient Boosting computed as the fraction of points in the majority class in the training set, lr is the learning rate, and $w_{t,l}$ is the value predicted by leaf l of tree t .

For each counterfactual couple (x^0, x^{CE}) , composed by a point x^0 and its counterfactual explanation x^{CE} , we restrict our attention to the features that change significantly, i.e., $|x_j^0 - x_j^{CE}| > \epsilon_j$. Then, we can compute a possible splitting threshold for feature j :

$$t_j = x_j^{CE} + \epsilon_j * \text{sign}(x_j^0 - x_j^{CE}) \quad \text{if } |x_j^0 - x_j^{CE}| > \epsilon_j. \quad (10)$$

Generating a set of counterfactual couples thus results in computing for each feature j a set of thresholds:

$$\tau_j = \{t_j = x_j^{CE} + \epsilon_j * \text{sign}(x_j^0 - x_j^{CE}), \quad \forall (x^0, x^{CE}) : |x_j^0 - x_j^{CE}| > \epsilon_j\}. \quad (11)$$

The set of thresholds across all features is denoted by $\tau = \cup_{j \in [1 \dots m]} \tau_j$. Our objective is to identify the most important decision boundaries of \mathcal{T} . The importance of a feature j can be measured by looking at the number of counterfactual couples for which a threshold of feature j is extracted, that is the size of τ_j . Note that increasing parameter λ_0 in

the cost function (2) results in principle in selecting the most important features. Let us denote by π_{t_j} the multiplicity of $t_j \in \tau_j \forall j \in [1 \dots m]$. We denote by $\tau_j^Q = \{t \in \tau_j : \pi_t \geq F_Q\}$ and by $\tau^Q = \cup_{j \in [1 \dots m]} \tau_j^Q$, where Q is a quantile value and F_Q is the Q th quantile of the frequency distribution $\pi = \{\pi_{t_j}\}_{\forall t_j \in \tau_j, \forall j \in [1 \dots m]}$.

After fixing a quantile value Q , we want to use the thresholds in τ^Q as splitting values in the nodes of the surrogate Decision Tree. This can be translated into a Feature Discretization procedure of the data in D_{tr} . Parameter Q allows thus to control the number of thresholds considered in τ^Q across all features.

Each feature j can be replaced with a set of binary variables, one for each threshold $t \in \tau_j^Q$. The binary variable b_{it} associated with each threshold $t \in \tau_j^Q$ represents whether feature j for data point x_i lies before or after t :

$$b_{it} = \begin{cases} 0 & \text{if } x_{ij} \leq t \\ 1 & \text{otherwise} \end{cases} \quad \forall t \in \tau_j^Q. \quad (12)$$

With this procedure we transform all the numerical features into binary variables. If we have no threshold on a feature, we remove it from the discretized dataset. Note that the higher the Q , the more compressed is the discretized dataset.

3.1. Algorithm

In this section, we describe the details of our procedure, namely FCCA (Feature Compression based on Counterfactual Analysis). As Target system \mathcal{T} we can consider any black-box model for which it is possible to solve the Counterfactual Explanation Problem (1).

After training the Target system, the second step is to extract from D_{tr} a large set of points \mathcal{M} for computing their Counterfactual Explanation. Our numerical experiments seem to suggest that the time for solving Problem (2)–(7) depends on the “closeness” of the initial point to the decision boundary of \mathcal{T} , where the closeness relates to how many features need to be changed in order to change the classification outcome. In fact, if the initial point is far from the decision boundary of \mathcal{T} , a large perturbation may be needed to cross the decision boundary i.e., change the classification outcome. A proxy for the time needed for solving Problem (2)–(7) is thus the classification probability of the tree ensemble. Indeed, a point that is classified with a high probability in Tree Ensemble model, will be classified with the same label by a large subset of the trees of the ensemble. This means that changing the classification outcome is likely to require changing many features. Thus it is likely to be computationally more expensive to compute the Counterfactual Explanation of a point classified with a higher classification probability. On the other hand, it is possible that considering points too close to the decision boundary (i.e., classified with a very low probability) can introduce some noise and make the procedure less robust. We can thus define \mathcal{M} to be composed by the points in the training set D_{tr} which are correctly classified and where the classification probability is bounded between two values $0.5 \leq p_0 \leq 1$ and $p_0 \leq p_1 \leq 1$:

$$\mathcal{M} = \{(x_i, y_i) \in D_{tr} \mid f_{\mathcal{T}}(x_i) = y_i \quad p_0 \leq \Pi_{\mathcal{T}}(x_i) \leq p_1\}, \quad (13)$$

where $f_{\mathcal{T}}(x_i)$ and $\Pi_{\mathcal{T}}(x_i)$ return respectively the classification label and the classification probability for x_i . We compute the set C of Counterfactual Explanations for all points in \mathcal{M} .

We can use Eq. (10) for extracting, from all couples in $(\mathcal{M}[i], C[i])_{i \in [1 \dots |\mathcal{M}|]}$, the set of thresholds τ . We set a value of Q between 0 and 1 and use this value to restrict τ to include only the more frequent thresholds τ^Q . The output of the procedure is thus the discretization of both data in the training and test set, namely $D'_{tr}(\tau^Q)$ and $D'_{ts}(\tau^Q)$. The discretized data can later be used to train and test the performance of a Decision Tree, which can be trained both by using the heuristic CART algorithm or by using an optimal approach as the one proposed in Lin et al. (2020). Note that training a Decision Tree on the discretized dataset implies that the decision tree uses exactly some of the thresholds found by the counterfactual computation.

The described procedure is summarized in Algorithm 1.

3.2. Discretization effectiveness

As a side product, our procedure returns different discretizations by changing Q . In order to evaluate the effectiveness of these discretizations in terms of compression ability, we introduce two metrics:

Compression rate When we apply the discretization, some points collapse to the same discretization. The compression rate is defined as $\eta = 1 - r$, where r is the ratio between the number of points in D_{tr} with different discretizations and the total number of points in D_{tr} .

Inconsistency rate As a downside of the compression, when multiple points collapse to the same discretization, it may happen that not all of them have the same label. For each feature j , we denote by ξ_j the number of thresholds in τ_j^Q ; the number of possible values that each discretized point assumes on feature j is thus $\xi_j + 1$. The number of possible discretized points \mathcal{N}_{τ^Q} thus depends on how many thresholds we have for each feature:

$$\mathcal{N}_{\tau^Q} = \prod_{j=1}^m (\xi_j + 1). \quad (14)$$

For each possible discretization $l \in [1 \dots \mathcal{N}_{\tau^Q}]$, we denote by $\Omega_l \subseteq D_{tr}$ the set of points that fall into that discretization. We denote by $\Omega_l^0 = \{x_i \in \Omega_l : y_i = 0\}$ and by $\Omega_l^1 = \{x_i \in \Omega_l : y_i = 1\}$. The number of inconsistencies δ_l in Ω_l is thus equal to the number of points in Ω_l with minority label: $\delta_l = \min\{|\Omega_l^0|, |\Omega_l^1|\}$. The inconsistency rate produced by the discretization procedure is thus expressed as:

$$\delta = \frac{1}{|D_{tr}|} \sum_{l=1}^{\mathcal{N}_{\tau^Q}} \delta_{\Omega_l}. \quad (15)$$

The inconsistency rate represents an irreducible error; thus, $1 - \delta$ is an upper bound to the accuracy that any classifier built on the discretized dataset is able to achieve. Note that in principle it is always possible to train on the discretized dataset a univariate decision tree, without limiting the depth, achieving an accuracy on the training set equal to $1 - \delta$.

Both the compression rate and the inconsistency rate depend on the value of Q . For high values of Q we consider a low granularity discretization that results in a high compression rate, related to high interpretability; but on the other hand, the discretization could produce a large number of inconsistencies, that represents a lower bound on the error rate of any classification method on the discretized dataset. The objective is thus to choose Q in order to keep a good trade-off between the compression rate and the inconsistency rate. Summarizing, in our methodology we have two parameters allowing to control the quality of the discretized dataset: the parameter λ_0 in the cost function (2) influences the number of features involved in the discretization (feature sparsity), whereas the parameter Q controls the number of thresholds involved across all the features (threshold sparsity).

4. Experimental setup

We tested the procedure summarized by Algorithm 1 on several binary classification datasets with continuous features, whose characteristics are summarized in Table 1. As Target black-box model, we can use any black-box model for which it is possible to efficiently solve Problem (1). Possible choices are Tree Ensembles (e.g. Random Forest, Gradient Boosting) or Support Vector Machines with a linear kernel. Since the objective is to learn the decision boundaries of \mathcal{T} , it is important to choose a model with good performance on the dataset considered. In our analysis, we use Gradient Boosting with 100 estimators of depth 1 and learning rate equal to 0.1: the performance

Algorithm 1 Pseudocode for the FCCA procedure

```

1: Input data:  $p_0 \geq 0.5, p_0 \leq p_1 \leq 1, \lambda_0, \lambda_1, \lambda_2 \geq 0, 0 \leq Q \leq 1$ 
2: Default:  $p_0 := 0.5, p_1 := 1, \lambda_0 := 0.1, \lambda_1 := 1, \lambda_2 := 0, Q := 0$ 
3:  $D_{tr} \leftarrow \{(x_i, y_i) : x_i \in \mathcal{X}^m, y_i \in \{0, 1\} \quad \forall i = 1 \dots n_{tr}\}$ 
4:  $D_{ts} \leftarrow \{(x_{\ell}, y_{\ell}) : x_{\ell} \in \mathcal{X}^m, y_{\ell} \in \{0, 1\} \quad \forall \ell = 1 \dots n_{ts}\}$ 
5:  $\mathcal{T} \leftarrow$  black-box model trained on  $D_{tr}$ 
6:  $f_{\mathcal{T}} \leftarrow$  predictions function of  $\mathcal{T}$ 
7:  $\Pi_{\mathcal{T}} \leftarrow$  classification probability function of  $\mathcal{T}$ 

Phase 1 – Computing  $\mathcal{M}$  and  $C$ 

8:  $\mathcal{M} = \{(x_i, y_i) \in D_{tr} : f_{\mathcal{T}}(x_i) = y_i \ \& \ p_0 \leq \Pi_{\mathcal{T}}(x_i) \leq p_1\}$ 
9:  $C \leftarrow$  Counterfactual Explanations of  $\mathcal{M}$  with parameters  $\lambda_0, \lambda_1, \lambda_2$ 

Phase 2 – Computing the thresholds  $\tau$ 

10:  $\tau_j = \{\} \quad \forall j = 1 \dots m$ 
11: for  $i = 1 \dots |\mathcal{M}|$  do
12:    $(x^0, y^0) = \mathcal{M}[i]$ 
13:    $(x^{CE}, y^{CE}) = C[i]$ 
14:   for  $j = 1 \dots m$  do
15:     if  $|x_j^0 - x_j^{CE}| > \epsilon_j$  then
16:        $t_j = x_j^{CE} + \epsilon_j * \text{sign}(x_j^0 - x_j^{CE})$ 
17:       add  $t_j$  to  $\tau_j$ 
18:     end if
19:   end for
20: end for
21:  $\pi \leftarrow$  frequency distribution of the thresholds in  $\cup_{j=1 \dots m} \tau_j$ 

Phase 3 – Discretizing the dataset

22:  $F_Q = \text{quantile}(\pi, Q)$ 
23:  $\tau_j^Q = \{t \in \tau_j : \pi_t \geq F_Q\} \quad \forall j = 1 \dots m$ 
24:  $D'_{tr}(\tau^Q) \leftarrow \emptyset, \quad D'_{ts}(\tau^Q) \leftarrow \emptyset$ 
25: for  $i = 1 \dots n_{tr}$  do
26:    $b_i \leftarrow \{0 \text{ if } x_{ij} \leq t \text{ else } 1 \quad \forall t \in \tau_j \quad \forall j = 1 \dots m\}$ 
27:    $D'_{tr}(\tau^Q) \leftarrow D'_{tr}(\tau^Q) \cup \{(b_i, y_i)\}$ 
28: end for
29: for  $\ell = 1 \dots n_{ts}$  do
30:    $b_{\ell} \leftarrow \{0 \text{ if } x_{\ell j} \leq t \text{ else } 1 \quad \forall t \in \tau_j \quad \forall j = 1 \dots m\}$ 
31:    $D'_{ts}(\tau^Q) \leftarrow D'_{ts}(\tau^Q) \cup \{(b_{\ell}, y_{\ell})\}$ 
32: end for
33: return  $D'_{tr}(\tau^Q), D'_{ts}(\tau^Q)$ 

```

of this algorithm on the datasets computed in k -fold crossvalidation is shown in Table 2. There are several advantages of using this small depth. First, solving Problem (2)–(7) is extremely fast (order of 10^{-2} seconds). Second, recall that the ultimate goal of our work is to build an interpretable univariate decision tree of small depth, so that the number of thresholds in the final model is extremely low. Therefore, it may be enough to select among a small number of thresholds, as is the case with depth 1. Nevertheless, in the Supplementary Material, which is available on github at <https://github.com/ceciliasalvatore/sFCCA/blob/master/Supplementary%20Material.pdf>, we have performed further experiments where we include larger values of the depth of the Gradient Boosting, showing that our method remains competitive both in time and accuracy.

The experiments have been run on an Intel i7-1165G7 2.80 GHz CPU with 16 GB of available RAM, running Windows 11. The procedure was implemented in Python by using scikit-learn v1.2.1 for training our models; the optimization Problem (2)–(7) for computing Counterfactual Explanations was solved with Gurobi 10.0.1 (Gurobi, 2021). The value of ϵ used in constraint (6) is set to 10^{-4} . The code of the experiments is available at <https://github.com/ceciliasalvatore/sFCCA.git>. The experimental settings used in our experiments are described in Table 2.

Table 1

Summary of the datasets used in the experimental phase. We report the dataset name; the total number of data points n ; the number of features m ; the number of data points used as training set; the value of k used for the k -fold crossvalidation; the number of data points used as an external test set.

Name	n	m	# training points	k	# external test points
<i>boston</i>	506	14	506	5	0
<i>arrhythmia</i>	453	191	453	5	0
<i>ionosphere</i>	351	32	351	5	0
<i>magic</i>	19 020	10	5000	5	14 020
<i>particle</i>	86 209	50	5000	5	81 209
<i>vehicle</i>	98 928	100	5000	5	93 928

Table 2

Experimental setting in the different datasets. The Target model \mathcal{T} used is the Gradient Boosting algorithm with 100 estimators and depth 1. We also report the computational time needed to run the FCCA procedure on a single fold of the dataset.

Name	p_0	p_1	λ_0	λ_1	λ_2	Accuracy \mathcal{T} (%)	Time (s)
<i>boston</i>	0.5	1	0.1	1	0	82.81	7.73
<i>arrhythmia</i>	0.5	1	0.1	1	0	79.44	21.54
<i>ionosphere</i>	0.5	1	0.1	1	0	91.46	8.15
<i>magic</i>	0.5	0.7	0.1	1	0	83.13	23.81
<i>particle</i>	0.5	0.7	0.1	1	0	88.94	16.80
<i>vehicle</i>	0.5	0.7	0.1	1	0	85.43	25.23

As shown in Table 1, we considered both small datasets with less than 5000 data points (*boston*, *arrhythmia* and *ionosphere*) and big datasets with more than 5000 data points (*magic*, *particle*, *vehicle*). For each dataset we compute the performance in a 5-fold crossvalidation. For big datasets we restrict the dataset to a random subset of 5000 data points; the remaining ones are used as an external test set to additionally validate the performance of the system. The number of 5000 data points provides a good compromise between training time and the generalization capability, as shown by the performance on the large external test set. Please note that all methods are trained and tested over the same sets of data.

The procedure requires setting a group of parameters:

- p_0 : lower bound on the classification probability for points in the set \mathcal{M} . We recommend using $p_0 = 0.5$. A higher value can be set if the dataset is very noisy or unbalanced.
- p_1 : upper bound on the classification probability for points in the set \mathcal{M} . We recommend using $p_1 = 1.0$. A lower value can be used to reduce the number of points of which computing the Counterfactual Explanation if (a) the cardinality of \mathcal{M} is large (more than 1000 points) or (b) the Target model \mathcal{T} is of higher complexity (e.g. it is an ensemble of trees with high depth) so that the time for computing each Counterfactual Explanation can be high.
- $\lambda_0, \lambda_1, \lambda_2$: hyperparameters for the Counterfactual Explanation Problem (2)–(7). We recommend setting $\lambda_2 = 0$ to significantly reduce the computational time for the Counterfactual Explanation problem. λ_0 and λ_1 must be set to trade-off between sparsity and proximity; we can set $\lambda_1 = 1$ without loss of generality so that we only need to tune λ_0 . By setting $\lambda_0 = 0.1$ and $\lambda_1 = 1$, we mean that changing one additional feature has the same weight as changing the absolute value of one feature of 0.1, which is the 10% of the scale range (recall that we assume features to be scaled between 0 and 1). We performed also a set of experiments for further values of $\lambda_0 \in \{0.001, 10, 1000\}$, and report the results in the Supplementary Material. When λ_0 increases, the set of thresholds extracted is more sparse, and some features do not present thresholds anymore. However, the final accuracy of the optimal decision tree does not change since the most important thresholds are selected in any case. However, a too large value of λ_0 could lead to a set of thresholds too sparse that can underfit the classification outcome, especially if the final tree is deeper.

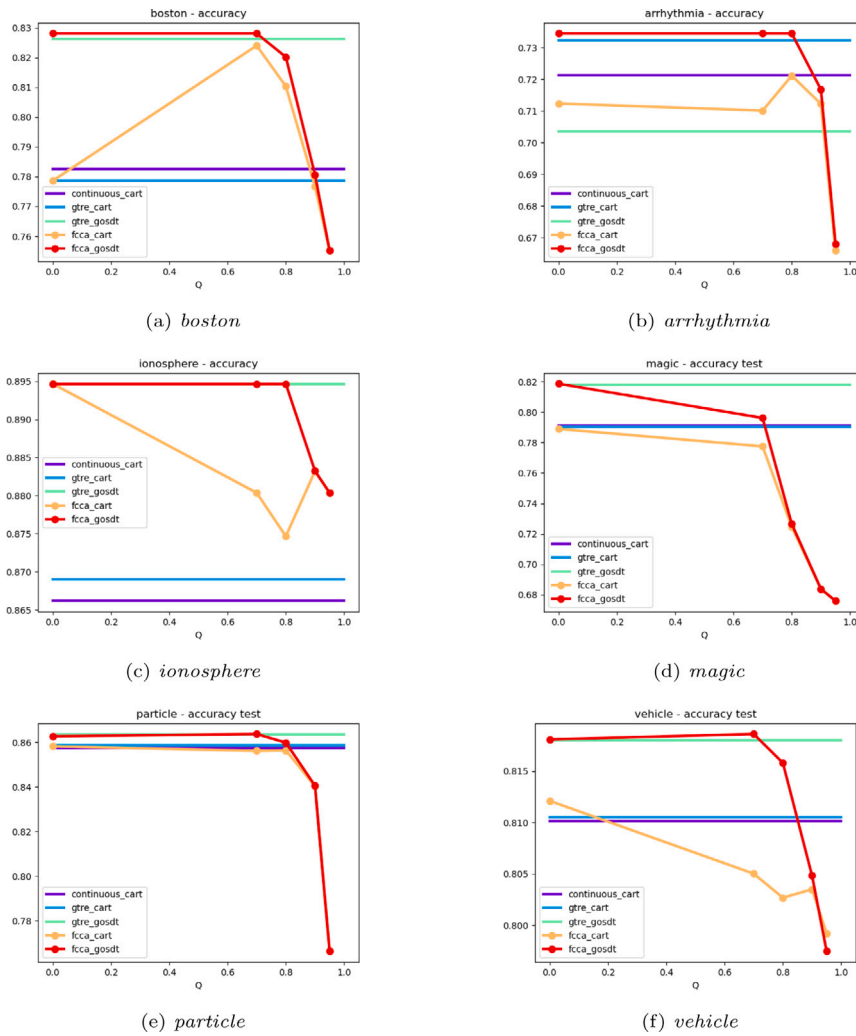


Fig. 3. Accuracy results on the benchmark datasets. We compare the performance of CART and GOSDT trained on the initial dataset with continuous features, the dataset discretized with the GTRE procedure, and the dataset discretized with the FCCA procedure. It is not possible to apply GOSDT directly to the dataset with continuous features. As reported in Table 1, for datasets with few observations (*boston*, *arrhythmia* and *ionosphere*) the accuracy is computed in a k -fold crossvalidation, while for datasets with many observations (*magic*, *particle* and *vehicle*) the accuracy is computed as the average result of the k classifiers trained in k -fold crossvalidation on the external test set.

- Q : the value of Q represents the granularity for the thresholds selected. The standard value is $Q = 0$, meaning that we consider all the thresholds extracted from the Counterfactual Analysis. Increasing the value of Q makes the discretized dataset more sparse, but, at the same time, it can degrade the representation of the input–output relationship. The trade-off between compression rate and inconsistency rate must be carefully considered when Q is increased.

4.1. Performance evaluation

In order to evaluate the effectiveness of the FCCA procedure, we look at two points of view:

- The performance of a classification model trained on the discretized dataset. In particular, our objective is to train a small Decision Tree; we can train Decision Trees either by using a heuristic algorithm such as CART or by using an optimal approach. In the second case, the discretization procedure allows us to use the GOSDT approach proposed by Lin et al. (2020), which requires binary features as input. For both CART and GOSDT, we limit to 3 the maximum depth allowed. For GOSDT, we also set the regularization parameter to $\frac{10}{n_{tr}}$; in our experiments, in fact, we noticed that using this value allows us to produce Decision Trees

that are shallow but can properly represent the input–output relationship.

- The compression and inconsistency rates of the discretized dataset.

Results obtained by the FCCA procedure are compared with the ones obtained by the initial dataset with continuous features and by the dataset discretized following the “Guessing Thresholds” procedure proposed in McTavish et al. (2022), that we will refer to as GTRE (Guessing Thresholds via Reference Ensemble). The GTRE procedure is also based on the idea of leveraging a Tree Ensemble model to extract relevant thresholds that discretize the input dataset in a meaningful way; the GTRE procedure uses as Tree Ensemble the Gradient Boosting algorithm. In all the experiments, when we run the two methods we set the same parameters for the Gradient Boosting to have a fair comparison (100 trees of depth 1 and learning rate 0.1).

4.2. Results

In this section, we analyze the results obtained in our experiments in terms of accuracy, sparsity, compression rate and inconsistency rate.

4.2.1. Accuracy

In Fig. 3 we present the results in terms of accuracy on the benchmark datasets obtained by training both CART and GOSDT on the initial

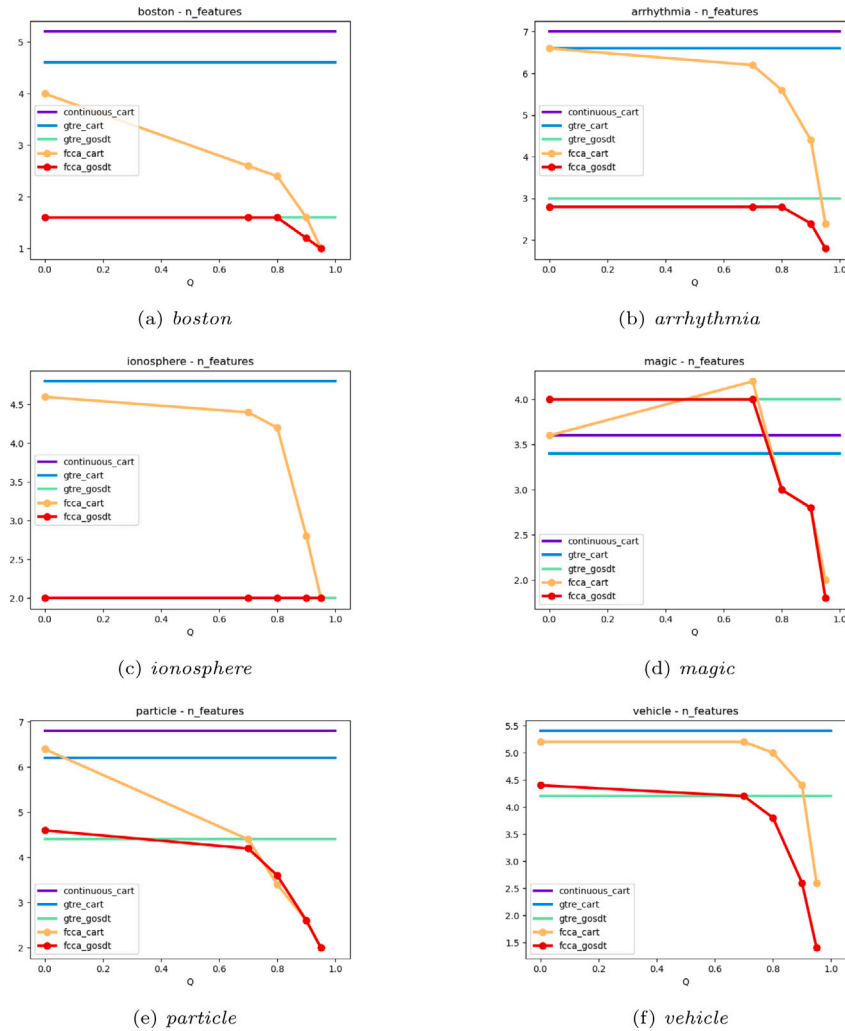


Fig. 4. Number of features used on the benchmark datasets. We compare the performance of CART and GOSDT trained on the initial dataset with continuous features, the dataset discretized with the GTRE procedure, and the dataset discretized with the FCCA procedure.

dataset with continuous features, the dataset discretized with the GTRE procedure, and the dataset discretized with the FCCA procedure at different levels of quantile Q . Please note that it is not possible to apply the GOSDT approach directly to the dataset with continuous features: in this case, it would, in fact, be necessary to introduce a threshold in the middle of any two consecutive data points for each feature; the resulting dataset is too big, and it is thus not affordable to compute GOSDT. In all datasets, we can notice that using a discretization technique is beneficial because it allows us to compute the optimal tree with the GOSDT algorithm; using this algorithm results in a more robust, accurate and sparse (Fig. 4) decision tree with respect to the heuristic CART. Combining GOSDT with the two discretization techniques GTRE and FCCA for $Q = 0$ achieves similar performance on all datasets; increasing the value of Q to 0.7 preserves a high accuracy level on almost all datasets, except for *magic*.

4.2.2. Compression and inconsistency rates

In Figs. 5 and 6, we plot respectively the compression rate η and the inconsistency rate δ of the initial dataset with continuous features, the dataset discretized with the GTRE procedure, and the dataset discretized with the FCCA procedure at different levels of the quantile Q . On all datasets, we can notice that the initial dataset with continuous features has zero compression and zero inconsistency, while both GTRE and FCCA have a significant compression rate that, as a downside,

leads to a non-zero inconsistency. For $Q = 0$, the compression and inconsistency rates of FCCA are very similar to the ones obtained by GTRE. In the FCCA procedure, compression and inconsistency rates are proportional to the quantile value Q . In fact, when Q tends to 1, also the compression rate η tends to 1. A high compression rate is positive because implies that we are able to summarize the data using information at a low granularity. Therefore it is easier to build a small and interpretable decision tree for learning the input–output relationship of this representation of the dataset. As a downside, however, also the inconsistency rate increases when Q tends to 1. The trade-off between compression rate and inconsistency rate can help us in deciding whether it is possible to increase the level of Q in order to obtain a more sparse dataset without affecting accuracy. In fact, in all datasets except for *magic* we can observe that increasing Q to 0.7 leads to an acceptable level of inconsistency when compared to the accuracy level that is generally reached on that dataset. On the contrary, on *magic* the inconsistency rate is quite high already for $Q = 0$, and it arrives at almost 20% for $Q = 0.7$: for this reason, for $Q = 0.7$ we record a decrease in the accuracy (Fig. 3(d)).

4.2.3. Sparsity

As already stated, increasing the value of Q in the FCCA procedure leads to selecting the thresholds that, according to the Counterfactual Analysis, are more relevant in describing the input–output relationship.

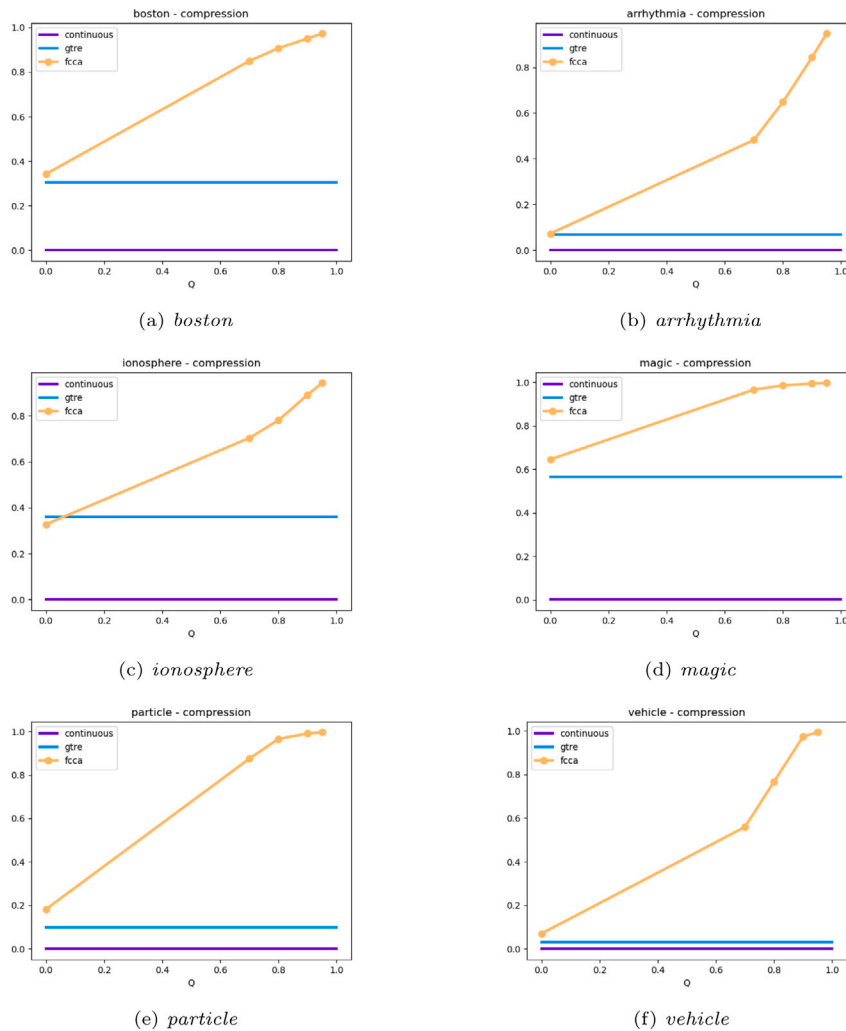


Fig. 5. Compression rate on the benchmark datasets. We compare the performance of CART and GOSDT trained on the initial dataset with continuous features, the dataset discretized with the GTRE procedure, and the dataset discretized with the FCCA procedure.

The discretized dataset derived is thus smaller because it is described through a smaller set of binary features. If Q is too high, we can lose the ability to represent the input–output relationship, ending up building low-performance surrogate classification models. In our previous analysis, we observed that $Q = 0.7$ is a good value in almost all of the benchmark datasets considered. In Figs. 7–12, we analyze the thresholds extracted by the FCCA procedure at different levels of Q and compare them to the thresholds extracted by the GTRE procedure. In these figures, we represent the thresholds extracted by each discretization procedure through a heatmap. For the FCCA procedure, the heatmap also represents the quantile. Please note that thresholds that appear at a certain quantile also appear at lower levels of quantiles. We can notice that the heatmaps obtained through the GTRE procedure and the ones obtained through the FCCA with $Q = 0$ are very similar, meaning that the two approaches extract very similar sets of thresholds. When Q grows, instead, the set of thresholds extracted by the FCCA procedure is much smaller.

4.3. Other target algorithms

In this section, we briefly analyze what happens by changing the algorithm implemented by the Target model \mathcal{T} . As alternative algorithms to Gradient Boosting, we consider Random Forest and Linear Support Vector Machines. For Random Forest, it is possible to compute Counterfactual Explanations analogously to what is done for Gradient

Boosting by solving Problem (2)–(7); for Support Vector Machines, instead, Eq. (1) leads to the following formulation:

$$\min_{x^{CE}} C(x^0, x^{CE}) := \lambda_0 \|x^0 - x^{CE}\|_0 + \lambda_1 \|x^0 - x^{CE}\|_1 + \lambda_2 \|x^0 - x^{CE}\|_2^2 \tag{16}$$

$$y^{CE}(w^T x^{CE} + b) \geq 1 \tag{17}$$

$$x^{CE} \in \mathcal{X}^0. \tag{18}$$

For the sake of brevity, we only show this comparison on one benchmark dataset, *boston*. The hyperparameters for both Random Forest and Linear SVMs are computed in crossvalidation: we thus set the maximum depth of the Random Forest to 4 and the value of C for SVM to 1. The other parameters used in this experiment are analogous to the one used in Section 4.2.

Fig. 13 presents the comparison of the FCCA method with different Target algorithms in terms of accuracy, sparsity, compression and inconsistency, while Fig. 14 presents the heatmap of the thresholds extracted by the FCCA method starting from the Random Forest and Linear SVM algorithms. We can notice that when $Q = 0$, the three Target models all lead to similar results in terms of accuracy; however, comparing Figs. 7 and 14, we can notice that the set of thresholds extracted is quite different. In particular, while the behavior for Gradient Boosting and Random Forest is quite similar, for the SVM we can notice that the thresholds extracted are distributed over almost the

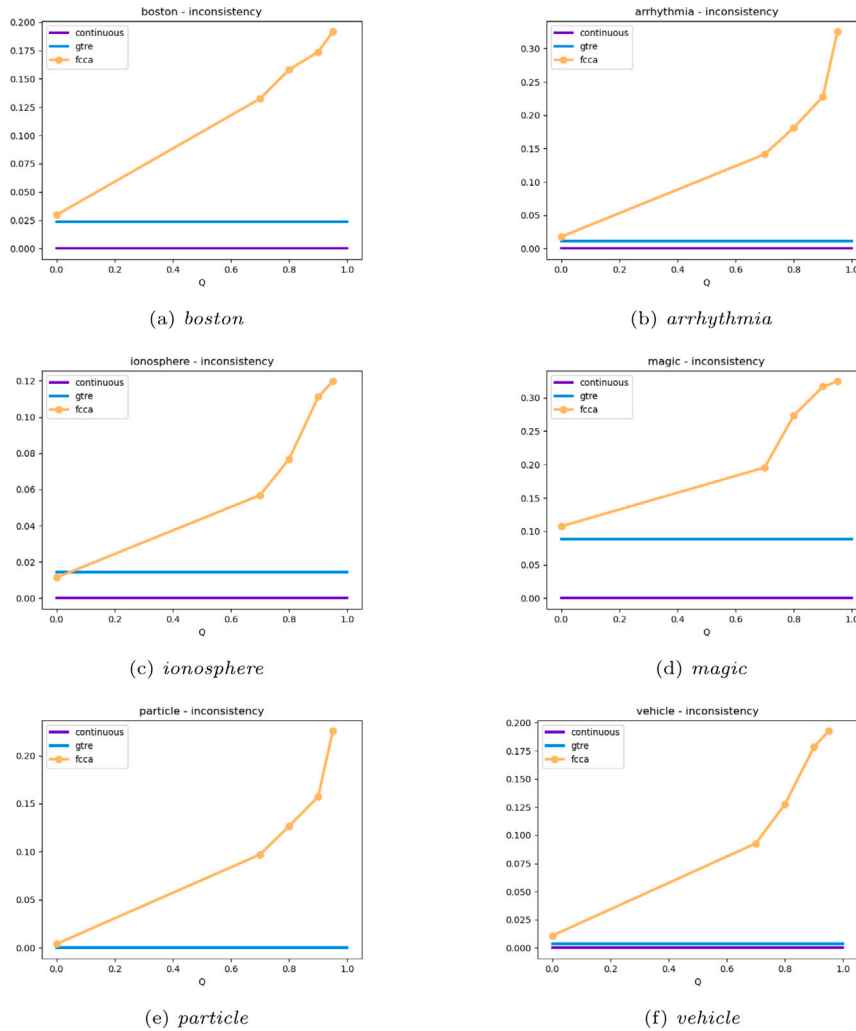


Fig. 6. Inconsistency rate on the benchmark datasets. We compare the performance of CART and GOSDT trained on the initial dataset with continuous features, the dataset discretized with the GTRE procedure, and the dataset discretized with the FCCA procedure.




Fig. 7. Heatmap of the thresholds extracted by the GTRE procedure and the FCCA procedure on *boston*. For the GTRE procedure, the heatmap only represents whether, in a given interval, there are some thresholds (Y) or not (N). For the FCCA procedure, and in case there are thresholds in the interval, the heatmap also represents the quantile. Please note that any threshold selected for a given quantile Q will also be selected for lower values of the quantile.

whole interval of possible values; this is due to the continuous nature of SVMs. For SVMs, Counterfactual Analysis thus results in a feature selection technique.

5. Conclusions

In this paper, we use Counterfactual Analysis to derive a supervised discretization of a dataset driven by a black-box model’s classification

function. Counterfactual Explanations are computed by solving an optimization problem that identifies the most relevant features and the corresponding cutting points for the black-box model. The Counterfactual Explanations provide a set of univariate decision boundaries that allows discretizing the original dataset into a set of binary variables. Having a compact dataset of binary variables makes it affordable to train an interpretable optimal decision tree by, e.g., the recent approach proposed in Lin et al. (2020). Our procedure allows us to discretize the dataset with a tunable granularity, controlled by means of the

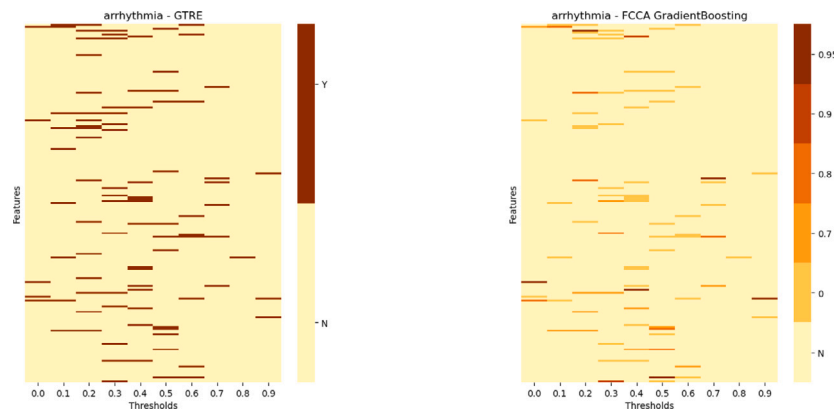


Fig. 8. Heatmap of the thresholds extracted by the GTRE procedure and the FCCA procedure on *arrhythmia*.

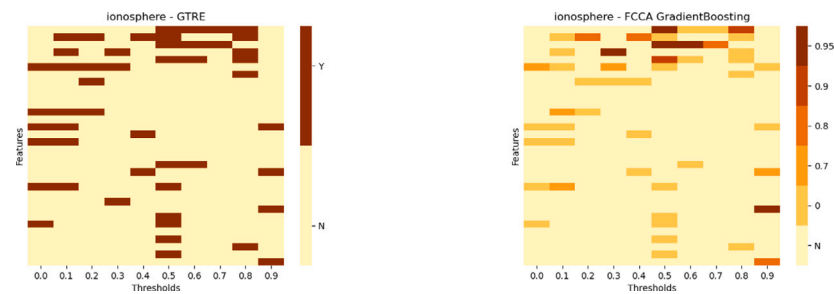


Fig. 9. Heatmap of the thresholds extracted by the GTRE procedure and the FCCA procedure on *ionosphere*.




Fig. 10. Heatmap of the thresholds extracted by the GTRE procedure and the FCCA procedure on *magic*.




Fig. 11. Heatmap of the thresholds extracted by the GTRE procedure and the FCCA procedure on *particle*.

parameter Q with $0 \leq Q \leq 1$. A high granularity, corresponding to low values of Q , results in a high level of detail in the dataset, where we consider a higher number of features, each represented by a large number of binary variables. A lower granularity, corresponding to high

values of Q , results in a sparse dataset, where we only select the most relevant features and model each of them with few binary variables.

Therefore, tuning the value of Q is needed to trade-off between the performance of a classification model built on this dataset (somehow




Fig. 12. Heatmap of the thresholds extracted by the GTRE procedure and the FCCA procedure on *vehicle*.




Fig. 13. Comparison on the use of different Target algorithms in the FCCA procedure in terms of accuracy, sparsity, compression and inconsistency on *boston*.




Fig. 14. Heatmap of the thresholds extracted by using the Random Forest and Linear SVM algorithms in the FCCA procedure on *boston*.

inversely proportional to Q) and its sparsity (that is a measure of interpretability and directly proportional to Q).

In the numerical section, we demonstrate the viability of our method on several datasets different in size, both in terms of the number of data points and the number of features. We compare our approach with the one proposed in [McTavish et al. \(2022\)](#), where they discretize the dataset by extracting thresholds retraining smaller and smaller version of the original black-box model until the accuracy drops. Our method has the advantage to detect the thresholds by solving an optimization problem, allowing us to tune both the sparsity in terms of features, controlling the cost function of the optimization problem, and the

sparsity in terms of the number of thresholds, controlling the parameter Q .

As future lines of research, we plan to use the optimization problem defined in [Carrizosa et al. \(2023\)](#) for computing counterfactuals of groups of individuals. We can add fairness and plausibility constraints to this problem to extract “fair” decision boundaries, and hence indirectly impose fairness of the optimal decision tree built by means of those boundaries.

Finally, we plan to formulate the problem of finding an optimal supervised discretization by a combinatorial optimization problem and study efficient algorithms for solving it.

Acknowledgments

This research has been financed in part by research project EC H2020 MSCA RISE NeEDS (Grant agreement ID: 822214). This support is gratefully acknowledged. Veronica Piccialli acknowledges financial support from PNRR MUR project PE000013-FAIR.

References

- Aghaei, S., Gómez, A., & Vayanos, P. (2021). Strong optimal classification trees. arXiv preprint arXiv:2103.15965.
- Aivodji, U., Bolot, A., & Gams, S. (2020). Model extraction from counterfactual explanations. arXiv preprint arXiv:2009.01884.
- Babic, B., Gerke, S., Evgeniou, T., & Cohen, I. G. (2021). Beware explanations from AI in health care. *Science*, 373(6552), 284–286.
- Bastos, J. A., & Matos, S. M. (2022). Explainable models of credit losses. *European Journal of Operational Research*, 301(1), 386–394.
- Bertsimas, D., & Dunn, J. (2017). Optimal classification trees. *Machine Learning*, 106, 1039–1082.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- Breiman, L., Friedman, J., Stone, C. J., & Olshen, R. (1984). *Classification and regression trees*. Chapman and Hall/CRC.
- Carrizosa, E., Molero-Río, C., & Romero Morales, D. (2021). Mathematical optimization in classification and regression trees. *TOP*, 29(1), 5–33.
- Carrizosa, E., Ramírez-Ayerbe, J., & Romero Morales, D. (2023). Mathematical optimization modelling for group counterfactual explanations. *Technical report*, IMUS, Sevilla, Spain, https://www.researchgate.net/publication/368958766_Mathematical_Optimization_Modelling_for_Group_Counterfactual_Explanations.
- Carrizosa, E., Ramírez Ayerbe, J., & Romero Morales, D. (2024). Generating collective counterfactual explanations in score-based classification via mathematical optimization. *Expert Systems With Applications*, 238, 121954.
- Chen, T., He, T., Benesty, M., Khotilovich, V., Tang, Y., Cho, H., et al. (2015). XGBoost: extreme gradient boosting. *R package version 0.4-2*, 1(4), 1–4.
- Dash, R., Paramguru, R. L., & Dash, R. (2011). Comparative analysis of supervised and unsupervised discretization techniques. *International Journal of Advances in Science and Technology*, 2(3), 29–37.
- Dougherty, J., Kohavi, R., & Sahami, M. (1995). Supervised and unsupervised discretization of continuous features. In A. Prieditis, & S. Russell (Eds.), *Machine learning proceedings 1995* (pp. 194–202). San Francisco (CA): Morgan Kaufmann, URL <https://www.sciencedirect.com/science/article/pii/B9781558603776500323>.
- Doumpos, M., Zopounidis, C., Gounopoulos, D., Platanakis, E., & Zhang, W. (2023). Operational research and artificial intelligence methods in banking. *European Journal of Operational Research*, 306(1), 1–16, URL <https://www.sciencedirect.com/science/article/pii/S037722172200337X>.
- Dumitrescu, E., Hué, S., Hurlin, C., & Tokpavi, S. (2022). Machine learning for credit scoring: Improving logistic regression with non-linear decision-tree effects. *European Journal of Operational Research*, 297(3), 1178–1192, URL <https://www.sciencedirect.com/science/article/pii/S0377221721005695>.
- Dunn, J. W. (2018). *Optimal trees for prediction and prescription* (Ph.D. thesis), Massachusetts Institute of Technology.
- European Commission (2020). White paper on artificial intelligence: a European approach to excellence and trust. https://ec.europa.eu/info/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en.
- Fernández, G., Aledo, J. A., Gamez, J. A., & Puerta, J. M. (2022). Factual and counterfactual explanations in fuzzy classification trees. *IEEE Transactions on Fuzzy Systems*, 30(12), 5484–5495.
- Fethi, M. D., & Pasiouras, F. (2010). Assessing bank efficiency and performance with operational research and artificial intelligence techniques: A survey. *European Journal of Operational Research*, 204(2), 189–198, URL <https://www.sciencedirect.com/science/article/pii/S0377221709005438>.
- Flores, M. J., Gámez, J. A., Martínez, A. M., & Puerta, J. M. (2011). Handling numeric attributes when comparing Bayesian network classifiers: does the discretization method matter? *Applied Intelligence*, 34(3), 372–385.
- Forel, A., Parmentier, A., & Vidal, T. (2022). Robust counterfactual explanations for random forests. arXiv preprint arXiv:2205.14116.
- García, S., Luengo, J., Sáez, J. A., López, V., & Herrera, F. (2013). A survey of discretization techniques: Taxonomy and empirical analysis in supervised learning. *IEEE Transactions on Knowledge and Data Engineering*, 25(4), 734–750.
- Goethals, S., Martens, D., & Calders, T. (2023). PreCoF: counterfactual explanations for fairness. *Machine Learning*, 1–32.
- Goodman, B., & Flaxman, S. (2017). European union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine*, 38(3), 50–57.
- Guidotti, R. (2022). Counterfactual explanations and how to find them: literature review and benchmarking. *Data Mining and Knowledge Discovery*, 1–55.
- Günlük, O., Kalaganam, J., Li, M., Menickelly, M., & Scheinberg, K. (2021). Optimal decision trees for categorical data via integer programming. *Journal of Global Optimization*, 81(1), 233–260.
- Gurobi (2021). Gurobi optimizer reference manual. <http://www.gurobi.com>.
- Karimi, A.-H., Barthe, G., Schölkopf, B., & Valera, I. (2022). A survey of algorithmic recourse: contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5), 1–29.
- Kozodoi, N., Jacob, J., & Lessmann, S. (2022). Fairness in credit scoring: Assessment, implementation and profit implications. *European Journal of Operational Research*, 297(3), 1083–1094, URL <https://www.sciencedirect.com/science/article/pii/S0377221721005385>.
- Kuppa, A., & Le-Khac, N.-A. (2021). Adversarial XAI methods in cybersecurity. *IEEE Transactions on Information Forensics and Security*, 16, 4924–4938.
- Kusner, M. J., Loftus, J., Russell, C., & Silva, R. (2017). Counterfactual fairness. *Advances in Neural Information Processing Systems*, 30, 4066–4076.
- Laurent, H., & Rivest, R. L. (1976). Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5(1), 15–17.
- Lin, J., Zhong, C., Hu, D., Rudin, C., & Seltzer, M. (2020). Generalized and scalable optimal sparse decision trees. In H. D. III, & A. Singh (Eds.), *Proceedings of machine learning research: Vol. 119, Proceedings of the 37th international conference on machine learning* (pp. 6150–6160). PMLR, URL <https://proceedings.mlr.press/v119/lin20g.html>.
- Maragno, D., Kurtz, J., Röber, T., Goedhart, R., Birbil, Ş., & den Hertog, D. (2023). Finding regions of counterfactual explanations via robust optimization. arXiv preprint arXiv:2301.11113.
- Maragno, D., Röber, T. E., & Birbil, I. (2022). Counterfactual explanations using optimization with constraint learning. arXiv preprint arXiv:2209.10997.
- Martens, D., & Provost, F. (2014). Explaining data-driven document classifications. *MIS Quarterly*, 38(1), 73–99.
- McTavish, H., Zhong, C., Achermann, R., Karimalis, I., Chen, J., Rudin, C., et al. (2022). Fast sparse decision tree optimization via reference ensembles. Vol. 36, In *Proceedings of the AAAI conference on artificial intelligence* (9), (pp. 9604–9613). URL <https://ojs.aaai.org/index.php/AAAI/article/view/21194>.
- Molnar, C., Casalicchio, G., & Bischl, B. (2020). Interpretable machine learning - a brief history, state-of-the-art and challenges. In *Joint European conference on machine learning and knowledge discovery in databases* (pp. 417–431). Springer.
- Mothilal, R. K., Sharma, A., & Tan, C. (2020). Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 607–617). Association for Computing Machinery.
- Parmentier, A., & Vidal, T. (2021). Optimal counterfactual explanations in tree ensembles. In M. Meila, & T. Zhang (Eds.), Vol. 139, *Proceedings of the 38th international conference on machine learning* (pp. 8422–8431). PMLR, URL <https://proceedings.mlr.press/v139/parmentier21a.html>.
- Piramuthu, S. (2004). Evaluating feature selection methods for learning in data mining applications. *European Journal of Operational Research*, 156(2), 483–494, URL <https://www.sciencedirect.com/science/article/pii/S0377221702009116>.
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1(1), 81–106.
- Quinlan, J. R. (2014). *C4.5: Programs for machine learning*. Elsevier.
- Ramírez-Gallego, S., García, S., Mouriño-Talín, H., Martínez-Rego, D., Bolón-Canedo, V., Alonso-Betanzos, A., et al. (2016). Data discretization: taxonomy and big data challenge. *WIREs Data Mining and Knowledge Discovery*, 6(1), 5–21.
- Ridgeway, G. (2013). The pitfalls of prediction. *National Institute of Justice Journal*, 271, 34–40.
- Rudin, C., Chen, C., Chen, Z., Huang, H., Semenova, L., & Zhong, C. (2022). Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys*, 16, 1–85.
- Silva, D. M., Pereira, G. H., & Magalhães, T. M. (2022). A class of categorization methods for credit scoring models. *European Journal of Operational Research*, 296(1), 323–331.
- Sokol, K., & Flach, P. A. (2019). Counterfactual explanations of machine learning predictions: opportunities and challenges for AI safety. *SafeAI@ AAAI*.
- Verma, S., Boonsanong, V., Hoang, M., Hines, K., Dickerson, J., & Shah, C. (2022). Counterfactual explanations and algorithmic recourses for machine learning: A review. arXiv preprint arXiv:2010.10596.
- Verwer, S., & Zhang, Y. (2019). Learning optimal classification trees using a binary linear program formulation. Vol. 33, In *Proceedings of the AAAI conference on artificial intelligence* (01), (pp. 1625–1632).
- Vidal, T., & Schiffer, M. (2020). Born-again tree ensembles. In H. D. III, & A. Singh (Eds.), Vol. 119, *Proceedings of the 37th international conference on machine learning* (pp. 9743–9753). PMLR, URL <https://proceedings.mlr.press/v119/vidal20a.html>.
- Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887.
- Yang, Y., & Webb, G. I. (2009). Discretization for naive-Bayes learning: managing discretization bias and variance. *Machine Learning*, 74(1), 39–74.
- Zhao, X., Zhang, W., Xiao, X., & Lim, B. (2021). Exploiting explanations for model inversion attacks. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 682–692).