

# Factual and Counterfactual Explanations for Black Box Decision Making

**Riccardo Guidotti**

ISTI CNR

**Anna Monreale**

University of Pisa

**Fosca Giannotti**

ISTI CNR

**Dino Pedreschi**

University of Pisa

**Salvatore Ruggieri**

University of Pisa

**Franco Turini**

University of Pisa

**Abstract**—The rise of sophisticated machine learning models has brought accurate but obscure decision systems, which hide their logic, thus undermining transparency, trust, and the adoption of artificial intelligence (AI) in socially sensitive and safety-critical contexts. We introduce a local rule-based explanation method, providing faithful explanations of the decision made by a black box classifier on a specific instance. The proposed method first learns an interpretable, local classifier on a synthetic neighborhood of the instance under investigation, generated by a genetic algorithm. Then, it derives from the interpretable classifier an explanation consisting of a decision rule, explaining the factual reasons of the decision, and a set of counterfactuals, suggesting the changes in the instance features that would lead to a different outcome. Experimental results show that the proposed method outperforms existing approaches in terms of the quality of the explanations and of the accuracy in mimicking the black box.

■ **NEWSPAPERS ARE FULL** of commentaries about algorithms taking critical decisions that heavily

impact on our life and society, from loan concession in bank systems to pedestrian detection in self-driving cars. The worry is not only due to the increasing automation of decision making, but mostly to the fact that the algorithms are opaque and their logic unexplained. The main cause of this lack of transparency is that

*Digital Object Identifier 10.1109/MIS.2019.2957223*

*Date of publication 3 December 2019; date of current version 22 January 2020.*

decision-making algorithms are generated often from data through machine learning (ML). ML allows building predictive models that map user features into a decision, obtained by generalizing from a dataset of examples. The process of inferring a classification model from examples cannot be easily controlled because the size of training data and the complexity of the learned model are too big for humans. The inability to obtain an explanation for a decision is a profound drawback of learning from data, limiting social acceptance, and trust in its adoption in many sensitive contexts.

In this article, we address the problem of explaining the decision outcome taken by an algorithm providing “meaningful explanations of the logic involved” when automated decision making takes place, as prescribed by the “right to explanation” of the European General Data Protection Regulation. We perform our research under some specific assumptions. First, we assume that an explanation is interesting if it clarifies why a *specific decision* has been made, i.e., we aim for *local* explanations, not general, *global*, descriptions of how the overall system works.<sup>1,2</sup> Second, we assume that the explanations should be as close as possible to the language of reasoning, which is formal logic. Thus, we assume the user can understand elementary logic rules, but it should also be considered that from logic rules it is easy to construct narratives that are understood by users with diverse expertise. Finally, we assume that the black box decision system can be queried as many times as necessary, to probe its decision behavior to the scope of reconstructing its logic; this is certainly the case in a legal argumentation in court, or in an industrial setting. On the other hand, we make no assumptions on the algorithms used in the obscure classifier: we aim at an *agnostic* explanation method analyzing the input–output behavior of the black box, disregarding its internals.<sup>3</sup>

We propose LORE, a LOCAL rule-based explanation method for tabular data. Given a black box binary predictor  $b$  and a specific instance  $x$  labeled with outcome  $y$  by  $b$ , we build a simple, interpretable predictor by first generating a balanced set of neighbor instances of the given  $x$  through an ad-hoc genetic algorithm, and then

extracting from such a set labeled with  $b$  a decision tree classifier. A *local explanation* is then extracted from the obtained decision tree. The local explanation is a pair composed by 1) a—factual—*logic rule*, corresponding to the path in the tree that explains why  $x$  has been labeled as  $y$  by  $b$ , and 2) a set of *counterfactual rules*, explaining which changes in  $x$  would invert the class  $y$  assigned by  $b$ . For example, from the compas dataset, we may have the following explanation: the rule  $\{age \leq 39, race = African-American, recidivist = True\} \rightarrow HighRisk$  and the counterfactuals  $\{age > 40\}$  and  $\{race = White-American\}$ . Here, the factual explanation is that the high risk of recidivism is predicted for a black person younger than 40 with prior recidivism; the counterfactuals explain that a lower risk would be predicted if the person were either older than 40 or white. The usefulness of the explanation depends on the stakeholder: it may make sense to a judge that wants to understand and evaluate the suggestion by the decision support system and possibly discover that it is biased against black people.

The intuition behind our method, common to other *local* approaches, such as LIME<sup>3</sup> and ANCHOR,<sup>4</sup> is that the decision boundary for the black box can be arbitrarily complex over the whole data space, but in the neighborhood of a data point there is a high chance that the decision boundary is clear and simple, hence amenable to be captured by an interpretable model. These methods are named *local* because they focus on the behavior of the black box in the neighborhood of the specific instance  $x$ , without providing a single description of the logic of the black box for all possible instances. On the other hand, global methods such as by Grover *et al.*,<sup>2</sup> aims at retrieving explanations for the whole logic of the model. The novelty of our method is twofold. First, the high expressiveness of the proposed explanation surpasses state-of-the-art methods providing not only succinct evidence why an instance has been assigned a specific label, but also counterfactuals suggesting what should be different in the vicinity of the instance to reverse the predicted outcome. In other words, our inferred explanations are both factual and counterfactual, in line with

the cognitive psychology literature maintaining that counterfactuals help people to reason on explanations that identify cause-effect or reason-action relations between events.<sup>5,6</sup> Similarly, Grover *et al.*<sup>2</sup> produces “balanced” explanations supporting and opposing to a fact. Second, the local decision boundary in the neighborhood of the instance to explain is explored through a focused genetic algorithm, which produces high-quality training data to learn the local decision tree.

We propose extensive experiments to assess the goodness of our explanation method with respect to existing linear, rule-based, and counter-factual-based explanation approaches.

## LOCAL RULE-BASED EXPLANATION METHOD

Given the black box  $b$ , and an instance  $x$  in the feature space  $\mathcal{X}^{(m)}$ , we aim to solve the *black box outcome explanation problem* which consists in providing an explanation  $e$  for the decision  $b(x) = y$ . We assume that some knowledge is available about the feature space  $\mathcal{X}^{(m)}$ , i.e., the empirical distribution of the  $m$  features. Nothing is assumed about the process of constructing the black box  $b$ .

As a solution to the black box outcome explanation problem, we propose a LORE method. LORE learns an interpretable predictor  $c$  that reproduces and accurately mimics the local behavior of  $b$  in the neighborhood  $Z$  of  $x$ . The neighborhood  $Z$  is generated by LORE as part of the explanation process through a *genetic algorithm* in order to accurately explore the local decision boundary of  $b$ . The synthetic instances in  $Z$  are then labeled using  $b$ , and this set is used to train an interpretable local predictor  $c$ . LORE adopts a *decision tree* as interpretable predictor  $c$ . Finally, a *factual and counter-factual explanation*  $e$  for the decision  $b(x) = y$  is derived from the structure of  $c$ . The explanation consists of a decision rule  $r$  and a set of counterfactual rules  $\Phi$ . Details are discussed in the rest of this section.

### Factual and Counter-Factual Explanation

We define an explanation  $e$  as a pair of objects:  $e = \langle r, \Phi \rangle$ , where  $r = p \rightarrow y$  is a *factual* decision rule describing the reason for the decision value

$y = b(x)$ , whereas  $\Phi$  is a set of *counterfactual* rules, namely rules describing the minimal number of changes in the feature values of  $x$  that would change the decision of the predictor to  $y' \neq y$ . Given the instance  $x = \{(age=22), (job=clerk), (income=800), (car=no)\}$ , we consider the following explanation for a loan request:

$$e = \langle r = \{age \leq 25, job = clerk, income \leq 900\} \rightarrow deny \\ \Phi = \{(\{income > 900\} \rightarrow grant) \\ (\{job = employer\} \rightarrow grant)\} \rangle$$

In a factual decision rule  $r$  of the form  $p \rightarrow y$ , the decision  $y$  is the *consequence* of the rule, while the *premise*  $p$  is a boolean condition on feature values. We assume that  $p$  is a conjunction of split conditions of the form  $a_i \in [v_i^{(l)}, v_i^{(u)}]$ , where  $a_i$  is a feature and  $v_i^{(l)}, v_i^{(u)}$  are lower and upper bound values in the domain of  $a_i$  extended with  $\pm\infty$ . An instance  $x$  *satisfies*  $r$ , or  $r$  *covers*  $x$ , if the boolean condition  $p$  evaluates to true for  $x$ , i.e., if the split conditions  $sc(x)$  is true for every condition in  $p$ . The rule  $r$  in the example above is satisfied by  $x = \{(age=22), (job=clerk), (income=800), (car=no)\}$ . When the instance  $x$  for which we have to explain the decision satisfies  $p$ , the rule  $p \rightarrow y$  represents a *motivation for taking* the decision, i.e.,  $p$  explains the fact why  $b$  returned  $y$ .

Consider now a set  $\delta$  of split conditions. We denote the update of  $p$  by  $\delta$  as  $p[\delta] = \delta \cup \{(a \in [v_i^{(l)}, v_i^{(u)}]) \in p \mid \nexists w_i^{(l)}, w_i^{(u)}. (a \in [w_i^{(l)}, w_i^{(u)}]) \in \delta\}$ . Intuitively,  $p[\delta]$  is the logical condition  $p$  with ranges for attributes overwritten as stated in  $\delta$ , e.g.,  $\{age \leq 25, job = clerk\} \{age > 25\}$  is  $\{age > 25, job = clerk\}$ . A *counterfactual rule* for  $p$  is a rule of the form  $p[\delta] \rightarrow y'$ , for  $y' \neq y$ . We call  $\delta$  a *counterfactual*. A counterfactual  $\delta$  describes *what* features to change and *how* to change them to get an outcome different from  $y$ . Continuing the example, changing the income feature of  $x$  to any value  $> 900$  it will change the predicted outcome from *deny* to *grant*. An expected property of a consistent counterfactual rule  $p[\delta] \rightarrow y'$  is that it should be *minimal* with respect to  $x$ . Minimality is measured with respect to the number of split conditions  $sc$  in  $p[\delta]$  not satisfied by  $x$ . We define  $nf(p[\delta], x) = |\{sc \in p[\delta] \mid \neg sc(x)\}|$ , where  $nf(\cdot, \cdot)$  stands for number of falsified split conditions. For example,  $\{income > 900\} \rightarrow grant$  is a

parent 1	25	clerk	800	yes
parent 2	30	empl.	2k	no
			↓	
children 1	25	empl.	2k	yes
children 2	30	clerk	2k	no

**Figure 1.** Crossover.

minimal counterfactual with one condition falsified by  $x$ . In summary, a counterfactual  $\delta$  is a (minimal) *motivation for reversing* the decision outcome.

### Neighborhood Generation

The first step of LORE to extract an explanation  $e$  is the *neighborhood generation* aiming to identify a set of instances  $Z$ , with feature characteristics close to the ones of  $x$ , that is able to reproduce the local decision behavior of the black box  $b$ . Since the objective is to learn a predictor, the neighborhood should be flexible enough to include instances with both decision values, namely  $Z = Z_{=} \cup Z_{\neq}$  where instances  $z \in Z_{=}$  are such that  $b(z) = b(x)$ , and instances  $z \in Z_{\neq}$  are such that  $b(z) \neq b(x)$ . We extract balanced subsets  $Z_{=}$  and  $Z_{\neq}$ , and then put  $Z = Z_{=} \cup Z_{\neq}$ . This task differs from approaches to instance *selection* based on genetic algorithms.<sup>7</sup> In our case, we cannot assume the availability of the training set of  $b$ , or not even that  $b$  is a supervised ML predictor for which a training set exists. Our task is instead similar to instance *generation* in the field of active learning, including evolutionary approaches.<sup>8</sup> We adopt an approach based on a *genetic algorithm*, which generates  $z \in Z_{=} \cup Z_{\neq}$  by maximizing the following fitness functions:

$$fitness_{=}^x(z) = I_{b(x)=b(z)} + (1 - d(x, z)) - I_{x=z}$$

$$fitness_{\neq}^x(z) = I_{b(x)\neq b(z)} + (1 - d(x, z)) - I_{x=z}$$


where  $d: \mathcal{X}^{(m)} \rightarrow [0, 1]$  is a distance function,  $I_{true}=1$ , and  $I_{false}=0$ . The first fitness function looks for instances  $z$  similar to  $x$  (term  $1-d(x, z)$ ), but not equal to  $x$  (term  $I_{x=z}$ ) for which  $b$  produces the same outcome as  $x$  (term  $I_{b(x)=b(z)}$ ). Thus, the maximization of  $fitness_{=}^x$  occurs for instances different from  $x$  and whose prediction is equal to  $b(x)$ . The second fitness function leads to the generation of instances  $z$  similar to  $x$ , but not equal to it, for which  $b$  returns a different decision.

parent	25	clerk	800	yes
	↓		↓	
children	27	clerk	1k	yes

**Figure 2.** Mutation.

LORE generates  $Z$  by instantiating the evolutionary approach of Bäck *et al.*<sup>9</sup> Using the terminology of Derrac *et al.*,<sup>8</sup> it is an instance of generational genetic algorithms for evolutionary prototype generation. However, prototypes are a condensed subset of a training set that enable optimization in predictor learning. We aim instead to generate new instances that separate well the decision boundary of the black box  $b$ . The *neighborhood generation* function first initializes the population  $P_0$  with  $N$  copies of the instance  $x$  to explain. Then, it enters the evolution loop that begins with the *selection*  $A$  of the  $P_i$  population having the highest fitness score. After that, the crossover operator is applied to a proportion of  $A$  according to the  $pc$  probability. The resulting and the untouched individuals are placed in  $B$ . We use a *two-point crossover* which selects two parents and two crossover features at random, and then swap the crossover feature values of the parents (see Figure 1). Thereafter, a proportion of  $B$ , determined by  $pm$ , is mutated and placed in  $C$ . The unmutated individuals are also added to  $C$ . Mutation consists of replacing features values at random according to the empirical distribution of a feature (see Figure 2). In experiments, we derive such distribution from the test set of instances to explain. Individuals in  $C = P_{i+1}$  are evaluated according to the fitness function, and the evolution loop continues until  $G$  generations are completed. The best individuals, according to the fitness function, are returned. The *neighborhood generation* function is run twice, once using  $fitness_{=}^x$  to derive  $Z_{=}$ , and once using  $fitness_{\neq}^x$  to derive  $Z_{\neq}$ .

Figure 3 shows an example of neighborhood generation for a black box consisting of a random forest (RF) and a bidimensional feature space. The figure contrasts uniform random generation around a specific instance  $x$  (starred) to our genetic approach. The latter yields a neighborhood that is considerably denser than the former one in the boundary region of the predictor. The density of the generated instances will



**Figure 3.** Black box decision: purple versus green. Starred instance  $x$ . Uniformly random (1st) and genetic generation (2nd). Density of random (3rd) and genetic generation (4th). The color bar in the last two plots indicates the density level (best view in color).

be a key factor in extracting a good local interpretable predictors.

Local Classifier and Explanation Extraction

Given the neighborhood  $Z$  of  $x$ , the second step of LORE is to build an interpretable predictor  $c$  trained on the instances  $z \in Z$  labeled with the black box decision  $b(z)$  to locally mimic the behavior of  $b$  in  $Z$ . Since  $c$  must be interpretable and able to provide a factual and counter-factual explanation  $e$ , LORE considers decision tree classifiers as interpretable predictor: 1) decision rules can naturally be derived from a root-leaf path in a decision tree; and 2) counterfactuals can be extracted by symbolic reasoning over the tree.

Once the decision tree  $c$  has been trained on  $Z$  labeled with  $b$ , LORE derives the explanation  $e = \langle r, \Phi \rangle$  as follows. The decision rule  $r = p \rightarrow y$  is formed by including in  $p$  the split conditions on the path from the root to the leaf node that is satisfied by  $x$ , and setting  $y = c(x)$ . By construction,  $r$  is consistent with  $c$  and satisfied by  $x$ . Consider now the counterfactual rules in  $\Phi$ . LORE looks for all paths in the decision tree  $c$  leading to a decision  $y' \neq y$ . Fix one of such paths, and let  $q$  be the conjunction of split conditions in it. Again by construction,  $q \rightarrow y'$  is a counterfactual rule consistent with  $c$ . Notice that, since we are using a decision tree, the counterfactual  $\delta$  for which  $q = p[\delta]$  has not to be explicitly computed—this is a benefit of using decision trees. Among all such  $q$ 's, only those with the minimum number of split conditions  $sc$  not satisfied by  $x$  are kept in  $\Phi$ .


As an example, consider the decision tree in Figure 4, and the instance  $x = \{(age=22), (job=clerk), (income=800), (car=no)\}$  for which

the decision *deny* (e.g., of a loan) has to be explained. The path followed by  $x$  is the leftmost one in the tree. The *factual* decision rule extracted from the path is  $\{age \leq 25, job = clerk, income \leq 900\} \rightarrow deny$ . There are four paths leading to the opposite decision:  $q_1 = \{age \leq 25, job = clerk, income > 900\}$ ,  $q_2 = \{17 < age \leq 25, job = employer\}$ ,  $q_3 = \{age > 25, income \leq 1500, car = yes\}$ , and  $q_4 = \{age > 25, income > 1500\}$ . It turns out:  $nf(q_1, x) = 1$ ,  $nf(q_2, x) = 1$ ,  $nf(q_3, x) = 2$ ,  $nf(q_4, x) = 2$ , and  $\Phi = \{q_1 \rightarrow grant, q_2 \rightarrow grant\}$ .

A *counterfactual instance* can also be generated from  $x$  and from counterfactual rule  $q \rightarrow y'$ . Among all instances that satisfy  $q$ , we choose the one that minimally changes attributes from  $x$  according to  $q$ .

EXPERIMENTS

In this section, we evaluate LORE against state-of-the-art competitors. (The source code and the datasets for reproducing the experiments are publicly available at <https://github.com/riccotti/LORExperiments> were performed on Ubuntu 16.04 LTS, 32 GB RAM, 3.30 GHz Intel Core i7.). LORE has been developed in Python, using the *deap* (<https://github.com/DEAP/deap>) library for the genetic neighborhood generation, and the *YaDT* system (<http://pages.di.unipi.it/ruggieri/YaDT/>) for the decision tree induction.



**Figure 4.** Decision tree mimicking the local behavior of a black box.

**Table 1. Lore versus lime: *hit* scores.**

Dataset	german		compas		adult	
Black Box	LORE	LIME	LORE	LIME	LORE	LIME
RF	<b>.925 ± .2</b>	.880 ± .3	<b>.941 ± .2</b>	.826 ± .4	<b>.901 ± .3</b>	.824 ± .4
NN	.980 ± .1	<b>1.00 ± .0</b>	<b>.987 ± .1</b>	.902 ± .3	.918 ± .3	<b>.998 ± .1</b>
SVM	<b>1.00 ± .0</b>	.966 ± .1	<b>.997 ± .1</b>	.900 ± .3	.985 ± .1	<b>.987 ± .1</b>

## Experimental Setup

We ran experiments on three tabular real-world open source datasets: adult (<https://archive.ics.uci.edu/ml/datasets/adult>), compas (<https://github.com/propublica/compas-analysis>), and german ([https://archive.ics.uci.edu/ml/datasets/statlog+\(german+credit+data\)](https://archive.ics.uci.edu/ml/datasets/statlog+(german+credit+data))). Each dataset was randomly split into two parts: 70% was used to train the black box classifiers, 30%, denoted by  $X$ , was used as instances to be explained. We denote by  $\hat{Y}$  the decisions provided by the black box  $b$  and by  $Y$  the set of decisions provided by the interpretable predictor  $c$ . We trained and explained away the following black box classifiers: RF, support vector machine (SVM), and multilayer perceptron (NN) as implemented by the *scikit-learn* Python library. Default parameters were used for both the black boxes and the libraries of LORE (the parameters of the genetic procedure, namely probabilities of crossover and mutation, number of generations, and population size are set with the default values<sup>9</sup> of 0.7, 0.2, 10, and 1000, respectively). We consider the following performance indicators to evaluate the quality of the explanations:

- The  $fidelity(Y, \hat{Y}) \in [0, 1]$  compares the predictions of  $c$  and  $b$  on  $Z$  measuring how good is  $c$  at mimicking  $b$ .
- The  $l-fidelity(Y, \hat{Y}) \in [0, 1]$  compares the predictions of  $c$  and  $b$  on the local (hence “l-”) instances of  $Z$  covered by  $r$  measuring how good is  $r$  at mimicking  $b$ .
- The  $hit(y, \hat{y}) \in \{0, 1\}$  compares the predictions of  $c$  and  $b$  on the instance  $x$  under analysis. It returns 1 if  $c(x)$  is equal to  $b(x)$ , and 0 otherwise.


We adopt the F1-measure for the first two. Aggregated values are reported by averaging performance indicators over  $X$ .

## Rules Versus Linear Regression

We present a qualitative and quantitative comparison with the linear explanations of LIME (<https://github.com/marcotcr/lime>).<sup>3</sup> A first crucial difference is that in LIME, the number of features composing an explanation is an input parameter that must be specified by the user. LORE, instead, automatically provides the user with an explanation including only the features useful to justify the black box decision. This is a clear improvement over LIME. In experiments, unless otherwise stated, we vary the number of features of LIME explanations from two to ten, and we consider the performance with the highest hit score.

Table 1 reports the mean and standard deviation of *hit*, while Figure 5 details the box plots of *fidelity* (top) and *l-fidelity* (bottom). The results show that LORE definitely outperforms LIME under various viewpoints. Regarding the *hit* score, LORE is clearly better than LIME in six out of nine cases, is very close to it in two cases, and performs clearly worse in 1 case. LORE has better *fidelity* scores and is more robust than LIME, which, instead, exhibits very high variability in the neighborhood (i.e., *l-fidelity*). This result can be attributed to the genetic approach of LORE. Figure 6 reports a multidimensional scaling of the neighborhoods of an instance  $x$  generated by the two approaches. LORE computes a dense and compact neighborhood. The instances generated by LIME, instead, can be very distant from each other and with low density around  $x$ .

We claim that the explanations provided by LORE are more abstract and comprehensible than the ones of LIME. Consider the example in Figure 7. The top part reports a LORE local explanation for an instance  $x$  from german. The central part is a LIME explanation. Weights in  $w$  are associated with the categorical values in the instance  $x$  to explain, and



**Figure 5.** LORE versus LIME: box plots of *fidelity* and *l-fidelity*. Top numbers are mean values.


to continuous upper/lower bounds where the bounding values are taken from  $x$ . Each weight tells the user how much the decision would have changed for different values of a specific feature. In the example, the weight 0.11 has the following meaning<sup>3</sup>: “if the duration in months had been higher than the value it is for  $x$ , the prediction would have been, on average, 0.11 less  $\theta$  (or 0.11 more  $I$ )” (See <https://marcotcr.github.io/lime/tutorials/Tutorial%20-%20continuous%20and%20categorical%20features.html>) *lime-tutorial* for more details about LIME explanations. A not very easy logic to follow when compared to a single decision rule which characterizes the contextual conditions for the decision of the black box. Another advantage of LORE explanation consists in the set of counterfactual rules. LIME provides a rough indication of where to look for a different decision: different categorical values or lower/higher continuous values of some feature. LORE’s counterfactual rules provide high-level and minimal-change contexts for reversing the outcome prediction.

called *anchors* (<https://github.com/marcotcr/anchor>), as explanations. Anchors are computed by incrementally adding equality conditions in the premise, while an estimate of the anchor precision is above a minimum threshold (set to 95%). Such an estimation relies on neighborhood generation through pure-exploration multiarmed bandit. On a qualitative level of comparison, ANCHOR requires the *a priori* discretization of continuous features, while LORE benefits of the capabilities of decision trees to split continuous features (see Figure 7).

On a quantitative level of comparison, since Anchor produces a single rule model starting from the instance to explain, hit is 100% by construction. Moreover, l-fidelity boils down to rule precision, namely the fraction of instances in the neighborhood set that is correctly classified by the rule, i.e., that have the same black box prediction as the instance to explain. Figure 8 reports the average precision of the decision rules for both ANCHOR and LORE. By construction,

Rules Versus Anchors

ANCHOR<sup>4</sup> is a rule-based local explainer inspired to LIME, which provides decision rules,




**Figure 6.** Neighborhoods multidimensional scaling of an instance  $x$ . LORE (left) computes a dense and compact neighborhood. LIME (right) generates instances that can be very distant from each other and with low density around  $x$ .

```

- LORE
r = ({credit_amount > 836, other_debtors = none,
      credit_history = critical account} → decision = 0)
Φ = ({({credit_amount ≤ 836, other_debtors = none,
        credit_history = critical account} → decision = 1),
      ({credit_amount > 836, other_debtors = none,
        credit_history = all paid back} → decision = 1) }
- LIME
w = [({duration_month ≤ 12:0.11, status= male_single: 0.07,
      credit_history= critical account: 0.06}, decision = 0),
      ({account_check_status = 0 ≤ ... < 200 DM: 0.09,
        installment_as_income = 1: 0.07}, decision = 1)]
- ANCHOR
a = ({credit_history = critical account,
      duration_month ∈ [0, 18.00]} → decision = 0)
    
```

**Figure 7.** Examples of explanations from LORE, LIME, and ANCHOR.



**Figure 8.** LORE VERSUS ANCHOR: box plots of *precision* and *coverage*. Top numbers are mean values.

**Table 2.** Lore versus anchor: Jaccard measure of stability.

<i>Dataset</i>	german		compas		adult	
<i>Black box</i>	LORE	ANCHOR	LORE	ANCHOR	LORE	ANCHOR
RF	<b>.76 ± .15</b>	.61 ± .15	<b>.75 ± .12</b>	.73 ± .14	<b>.70 ± .15</b>	.69 ± .15
NN	<b>.69 ± .18</b>	.53 ± .21	<b>.83 ± .13</b>	.79 ± .16	<b>.81 ± .12</b>	.65 ± .16
SVM	<b>.82 ± .16</b>	.32 ± .16	<b>.71 ± .16</b>	.70 ± .20	<b>.87 ± .14</b>	.67 ± .13

rule precision in ANCHOR is very high, since an estimator of such precision is constrained to be at least 95%. We also evaluate the rule coverage, namely the fraction of instances to explain covered by the rule: large values of coverage means better rule generalization.<sup>4</sup> Figure 8 also shows the average coverage of the decision rules, with LORE showing a consistently better coverage than ANCHOR. We also measure the possible overspecialization of the decision rules by evaluating their stability with respect to randomness in the neighborhood generation. We measure stability using Jaccard coefficient of feature sets used in the 10 decision rules computed for the same instances in 10 runs of the system. Table 2 reports mean and standard deviation of Jaccard coefficient. LORE has a better stability than ANCHOR for all datasets and black boxes. In summary, ANCHOR shows better precision than LORE at the expenses of generality and stability of the produced explanations. Such two properties are, however, essential for a general acceptance of an explanation methodology.

#### Rule-Based Versus Stochastic Counterfactuals

We compare LORE with the stochastic optimization counterfactual (SOC) approach<sup>10</sup> returning an instance  $x'$  as close as possible to a given  $x$ , but

for which the black box outputs a different prediction. To make a fair comparison, we have implemented the SOC as an alternative fitness function of the genetic neighborhood generation. Table 3 shows the performances of the two methods on  $nf$  (number of falsified conditions in counterfactual rules),  $c$ -hit (rate of agreement of black box and counterfactual decision for counterfactual instance), and  $cl$ -fidelity (F1-score of agreement of black box and counterfactual decision). Results show that LORE returns shorter explanations, i.e., simpler explanations, and provides counterfactual rules with an higher fidelity than SOC.

#### FUTURE WORKS

Future research directions include multivalued classification, going beyond relation data towards image and text, going beyond decision

**Table 3.** LORE versus SOC: performance of counterfactual rules.

<i>dataset</i>	<i>method</i>	<i>nf</i>	<i>c</i> -hit	<i>cl</i> -fidelity
german	LORE	<b>1.52 ± 1.18</b>	<b>.7765 ± .38</b>	<b>.6355 ± .43</b>
	SOC	14.80 ± 1.59	.3118 ± .47	.2297 ± .36
compas	LORE	<b>1.84 ± 0.78</b>	<b>.8694 ± .37</b>	<b>.8611 ± .41</b>
	SOC	6.24 ± 1.45	.8036 ± .38	.7555 ± .34



trees, and considering alternative models such as rule sets and rule lists. Moreover, even though LORE should be able to deal with high dimensional datasets via the genetic process, we plan to develop further experiments to verify its robustness. We are fully aware that the quest towards meaningful explanations of black box systems is at an embryonic stage. How to turn our proposed factual and counterfactual rules into substantive narratives, which can empower human stakeholders with diverse expertise by boosting their causal and what-if reasoning, is a fascinating challenge of great practical relevance for the successful adoption of many AI innovations.

## ACKNOWLEDGMENTS

This work was supported by the European Commission through the H2020 project INFRAIA-1-2014-2015: Research Infrastructure G.A. 654024 *SoBigData*, G.A. 825619 *A4EU*, G.A. 761758 *Humane AI*, and the ERC-2018-ADG G.A. 834756 “XAI: Science and technology for the eXplanation of AI decision making.”

## REFERENCES

1. R. Guidotti *et al.*, “A survey of methods for explaining black box models,” *ACM Comput. Surv.*, vol. 51, no. 5, 2018, Art. no. 93.
2. S. Grover, C. Pulice, G. I. Simari, and V. Subrahmanian, “Beef: Balanced english explanations of forecasts,” *IEEE Trans. Comput. Social Syst.*, vol. 6, no. 2, pp. 350–364, Apr. 2019.
3. M. T. Ribeiro *et al.*, ““Why should I trust you?”: Explaining the predictions of any classifier,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 1135–1144.
4. M. T. Ribeiro *et al.*, “Anchors: High-precision model-agnostic explanations,” in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018, pp. 1527–1535.
5. R. M. J. Byrne, “Counterfactuals in explainable artificial intelligence (XAI): Evidence from human reasoning,” in *Proc. 28th Int. Joint Conf. Artif. Intell.*, 2019, pp. 6276–6282.
6. T. Miller, “Explanation in artificial intelligence: Insights from the social sciences,” *Artif. Intell.*, vol. 267, pp. 1–38, 2019.
7. J. A. Olvera-López *et al.*, “A review of instance selection methods,” *Artif. Intell. Rev.*, vol. 34, no. 2, pp. 133–143, 2010.
8. J. Derrac *et al.*, “A survey on evolutionary instance selection and generation,” in *Proc. Model., Anal., Appl. Metaheuristic Comput.*, 2012, pp. 233–266.
9. T. Bäck, D. B. Fogel, and Z. Michalewicz, *Evolutionary Computation 1: Basic Algorithms and Operators*, vol. 1. Boca Raton, FL, USA: CRC Press, 2000.
10. S. Wachter *et al.*, “Counterfactual explanations without opening the black box: Automated decisions and the GDPR,” *Harvard J. Law Technol.*, vol. 31, pp. 842–887, 2018.

**Riccardo Guidotti** is currently a Researcher with Knowledge Discovery and Data Mining Laboratory, ISTI-CNR, Pisa, Italy. He received the Ph.D. degree in computer science from the University of Pisa, Pisa, Italy, with a thesis on Personal Data Analytics. He received an IBM fellowship Award in 2014/2015, and spent an Internship at IBM Research Dublin. He has been awarded with the ISTI-CNR Young Researcher Award from 2016 to 2019, and with the NGDSA at DSAA 2018. Contact him at [riccardo.guidotti@isti.cnr.it](mailto:riccardo.guidotti@isti.cnr.it).

**Anna Monreale** is currently an Assistant Professor with the Computer Science Department, University of Pisa, Pisa, Italy, and a member of Knowledge Discovery and Data Mining Laboratory. She has been a Visiting Student with Stevens Institute of Technology in 2010. Her research interests include data analytics, privacy, and ethical issues in mining human data. She is a Privacy-by-Design Ambassador and a member of the EU Panel of Experts on the Open Science Cloud. Contact her at [anna.monreale@unipi.it](mailto:anna.monreale@unipi.it).

**Fosca Giannotti** is currently a Director of Research with ISTI-CNR, Pisa, Italy. She leads the Knowledge Discovery and Data Mining Laboratory, Pisa, Italy, a joint research initiative of ISTI-CNR and University of Pisa, founded in 1994. She is the coordinator of the European research infrastructure SoBigData (<http://www.sobigdata.eu>). She received the ERC Advanced Grant XAI—*Science and technology for the explanation of AI decision making*. Contact her at [fosca.giannotti@isti.cnr.it](mailto:fosca.giannotti@isti.cnr.it).

**Dino Pedreschi** is currently a Professor of Computer Science with the University of Pisa, Pisa, Italy, a co-founder of Knowledge Discovery and Data Mining Laboratory, and a Pioneering Scientist in data science. His research focus is on big data analytics and mining and their impact on society. He is the Director of the Data Science Ph.D. program with Scuola Normale Superiore, Pisa, Italy. Contact him at [dino.pedreschi@unipi.it](mailto:dino.pedreschi@unipi.it).

**Salvatore Ruggieri** is currently a Professor of Computer Science with the University of Pisa, Pisa, Italy. He is a member of Knowledge Discovery and Data Mining Laboratory, with research interests in data mining and AI, including: algorithmic fairness, explainable AI, systems for knowledge discovery process; sequential and parallel classification algorithms. He is a program Co-Chair of the ACM conference on Fairness, Accountability, and Transparency, 2020. Contact him at [salvatore.ruggieri@unipi.it](mailto:salvatore.ruggieri@unipi.it).

**Franco Turini** is currently a Professor with the Department of Computer Science, University of Pisa, Pisa, Italy, and a co-founder of Knowledge Discovery and Data Mining Laboratory. From 1978–1980, he was a Visiting Scientist with the Carnegie-Mellon University and the IBM Research Center S. Jose, afterward. From 1992–1993, he was a Visiting Professor with the University of Utah. His contributions include discrimination-aware data mining and ML Contact him at [franco.turini@unipi.it](mailto:franco.turini@unipi.it).



## IEEE TRANSACTIONS ON BIG DATA

### ▶ SUBSCRIBE AND SUBMIT

For more information on paper submission, featured articles, calls for papers, and subscription links visit: [www.computer.org/tbd](http://www.computer.org/tbd)

*TBD* is financially cosponsored by IEEE Computer Society, IEEE Communications Society, IEEE Computational Intelligence Society, IEEE Sensors Council, IEEE Consumer Electronics Society, IEEE Signal Processing Society, IEEE Systems, Man & Cybernetics Society, IEEE Systems Council, and IEEE Vehicular Technology Society

*TBD* is technically cosponsored by IEEE Control Systems Society, IEEE Photonics Society, IEEE Engineering in Medicine & Biology Society, IEEE Power & Energy Society, and IEEE Biometrics Council

