BETTER SAMPLING IN EXPLANATION METHODS CAN PREVENT DIESELGATE-LIKE DECEPTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Machine learning models are used in many sensitive areas where, besides predictive accuracy, their comprehensibility is also important. Interpretability of prediction models is necessary to determine their biases and causes of errors and is a necessary prerequisite for users' confidence. For complex state-of-the-art black-box models post-hoc model-independent explanation techniques are an established solution. Popular and effective techniques, such as IME, LIME, and SHAP, use perturbation of instance features to explain individual predictions. Recently, Slack et al. (2020) put their robustness into question by showing that their outcomes can be manipulated due to poor perturbation sampling employed. This weakness would allow dieselgate type cheating of owners of sensitive models who could deceive inspection and hide potentially unethical or illegal biases existing in their predictive models. This could undermine public trust in machine learning models and give rise to legal restrictions on their use.

We show that better sampling in these explanation methods prevents malicious manipulations. The proposed sampling uses data generators that learn the training set distribution and generate new perturbation instances much more similar to the training set. We show that the improved sampling increases the robustness of the LIME and SHAP, while previously untested method IME is already the most robust of all.

1 INTRODUCTION

Machine learning models are used in many areas where besides predictive performance, their comprehensibility is also important, e.g., in healthcare, legal domain, banking, insurance, consultancy, etc. Users in those areas often do not trust a machine learning model if they do not understand why it made a given decision. Some models, such as decision trees, linear regression, and naïve Bayes, are intrinsically easier to understand due to the simple representation used. However, complex models, which are mostly used in practice due to better accuracy, are incomprehensible and behave like black boxes, e.g., neural networks, support vector machines, random forests, and boosting. For these models, the area of explainable artificial intelligence (XAI) has developed post-hoc explanation methods that are model-independent and determine the importance of each feature for the predicted outcome. Frequently used methods of this type are IME (Štrumbelj & Kononenko, 2013), LIME (Ribeiro et al., 2016), and SHAP (Lundberg & Lee, 2017).

To determine the features' importance, these methods use perturbation sampling. Slack et al. (2020) recently noticed that the data distribution obtained in this way is significantly different from the original distribution of the training data as we illustrate in Figure 1a. They showed that this can be a serious weakness of these methods. Possibility to manipulate the post-hoc explanation methods is a critical problem for the ML community, as the reliability and robustness of explanation methods are essential for their use and public acceptance. These methods are used in the interpretation of otherwise black-box models, they help in debugging of models, and reveal models' biases, thereby establishing trust in their behaviour. Non-robust explanation methods that can be manipulated can lead to catastrophic consequences, as explanations do not detect racist, sexist, or otherwise biased model if the owner of the model wants to hide these biases. This would enable dieselgate-like cheating where owners of sensitive prediction models could hide the socially, morally, or legally unacceptable biases present in their models. As the schema of the attack on explanation methods

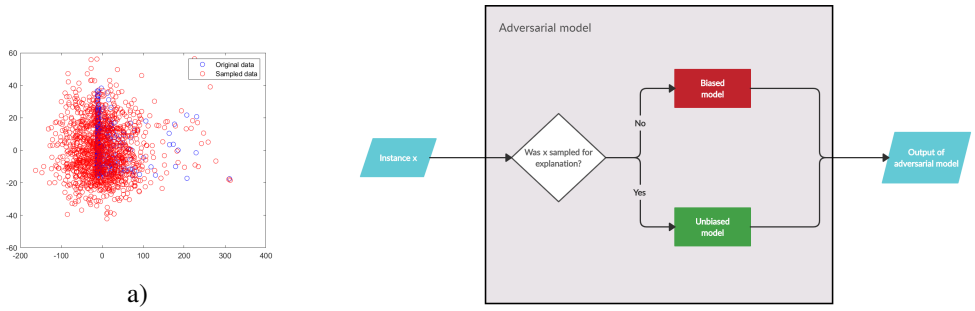

a) b)

Figure 1: a) PCA based visualization of a part of the COMPAS dataset. The blue points show the original instances and the red points represent instances generated with the perturbation sampling used in the LIME method. The distributions are notably different. b) The idea of the attack on explanation methods based on the difference of distributions. The attacker's adversarial model contains both the biased and unbiased model. The decision function that is part of the cheating model decides if the instance is outside the original distribution (i.e. used only for explanation), or it is an actual instance. If the case of an actual instance, the result of the adversarial model is equal to the result of the biased model, otherwise it is equal to the result of the unbiased model.

on Figure 1b shows, owners of prediction models could detect when their models are examined and return unbiased predictions in this case and biased predictions in normal use. This could have serious consequences in areas where the reliability and fairness of predictive models are essential, e.g., in healthcare or banking. Such weaknesses can undermine users' trust in machine learning models in general and slow down technological progress.

In this work, we propose to change the main perturbation-based explanation methods and make them more resistant to manipulation attempts. In our solution, the problematic perturbation-based sampling is replaced with more advanced sampling, which uses modern data generators that better capture the distribution of the training dataset. We test three generators, the RBF network based generator (Robnik-Šikonja, 2016), random forest-based generator, available in R library semiArtificial (Robnik-Šikonja, 2019), as well as the generator using variational autoencoders (Miok et al., 2019). We show that the modified gLIME and gSHAP methods are much more robust than their original versions. For the IME method, which previously was not analyzed, we show that it is already quite robust. We release the modified explanation methods under the open-source license[1]

The remainder of the paper is organized as follows. In Section 2, we present the necessary background and related work on explanation methods, attacks on them, and data generators. In Section 3, we propose a defence against the described weaknesses of explanation methods and in Section 4, we empirically evaluate the proposed solution. In Section 5, we draw conclusions and present ideas for further work.

2 Background and related work

In this section, we first briefly describe the background on post-hoc explanation methods and attacks on them, followed by data generators and related works on robustness of explanation methods.

2.1 Post-hoc explanation methods

The current state-of-the-art perturbation-based explanation methods, IME, LIME, and SHAP, explain predictions for individual instances. To form an explanation of a given instance, they measure the difference in prediction between the original instance and its neighbouring instances, obtained with perturbation sampling. Using the generated instances, the LIME method builds a local interpretable model, e.g., a linear model. The methods SHAP and IME determine the impact of the features as Shapley values from the coalitional game theory (Shapley, 1988). In this way, they assure that the produced explanations obey the four Shapley fairness axioms (Štrumbelj & Kononenko, 2013). Due to the exponential time complexity of Shapley value calculation, both methods try to approximate them. The three methods are explained in detail in the above references, and a formal overview is presented in Appendix A, while below, we present a brief description. In our exposition

[1] https://anonymous.4open.science/r/5d550c62-5c5c-4ee3-81ef-ab96fe0838ca/

of the explanation methods, we denote with f the predictive model and with x the instance we are explaining.

Explanations of instances with the LIME method is obtained with an interpretable model g. The model g has to be both locally accurate (so that it can obtain correct features contributions) and simple (so that it is interpretable). Therefore in LIME g is a linear regression model trained on the instances sampled in the neighbourhood of the explained instance x, which is first normalized.The perturbed explanation instances are generated by adding Gaussian noise to each feature of x independently. Using linear models to determine the contributions of features, LIME is fast but may produce poor explanations for instances close to complex decision boundaries.

SHAP estimates Shapley values using the approach of LIME (i.e. local linear models computed with the weighted least squares algorithm) but using a different neighbourhood and loss function for local accuracy which assure that Shapley values are obtained. The sampling of explanation instances in SHAP relies on the distribution set D, obtained from the training set, by computing the centroids of k-means clusters.

IME computes Shapley values using pure sampling approach (no proxy with an interpretable model). Its sampling replaces part of the explained instance with existing values of selected features and then averages differences in the prediction of explained instance and samples. This makes the method slower but also potentially more robust as the method does not assume the shape of the dependency.

2.2 Attacks on explanations

The idea of the attack on explanations presented in Slack et al. (2020) is that the attacker creates an adversarial (biased) model whose behaviour it wants to hide from explanation methods, e.g., a racist model that does not grant credits to blacks. As Figure 1b illustrates, a part of the attacker's adversarial model is an unbiased model, which, e.g., does not take disputed attributes such as race into account. The adversarial model tries to manipulate the explanation method by behaving like an unbiased model on perturbed instances used for the explanation – in this case, it uses the output of the unbiased model. On instances from the original distributions that are a part of a normal model use, the adversarial model uses the output of the biased model. To distinguish between the original and perturbation-sampled data, the adversarial model contains a decision model that selects the right type of model (biased or unbiased) based on the predicted type of instances.

The reason that the decision function can distinguish between the original data and sampled instances are the shortcomings of explanation methods that use the perturbation sampling (see Figure 1a). For that graph, we created new instances based on the COMPAS dataset (see Section 4) with the Gaussian noise as in the LIME explanation method. We mapped the original and sampled instances to 2D space using the PCA transformation. We can see that the sampled data is distributed differently than the original data. As the difference is revealed with PCA, other machine learning algorithms can also effectively distinguish between sampled and original data.

The attack can be formally defined in the following way. The attacker has a biased model b and a sample S from the distribution X_{dist}. The attacker has no additional information about the distribution X_{dist}. The attacker's goal is to produce an adversarial model e that behaves like b on the data from X_{dist} but does not reveal b's bias to the explanation methods. We assume that the attacker also has an unbiased model ψ that hides the f's bias. The attacker creates a decision model d, that should output 1 if the input instance x is from X_{dist} and 0 otherwise. The model d is trained on S and generated perturbation samples. The attacker creates the adversarial model e with its output defined by the following equation:

$$e(x) = \begin{cases} b(x), & d(x) = 1 \\ \psi(x), & d(x) = 0 \end{cases} \tag{1}$$

2.3 Data generators

We used three different data generators based on different algorithms, modelling the distribution of the training set: variational autoencoder with Monte Carlo dropout (Miok et al., 2019), RBF network (Robnik-Šikonja, 2016), and random forest ensemble (Robnik-Šikonja, 2019). In the remainder of the paper, we refer to the listed generators consecutively as MCD-VAE, rbfDataGen and TreeEnsemble.

Autoencoder (AE) consists of two neural networks called encoder and decoder. It aims to compress the input instances by passing them through the encoder and then reconstructing them to the original values with the decoder. Once the AE is trained, it can be used to generate new instances. Variational autoencoder (Doersch, 2016) is a special type of autoencoder, where the vectors z in the latent dimension (output of the encoder and input of the decoder) are normally distributed. Encoder is therefore approximating the posterior distribution $p(z|x)$, where we assume $p(z|x) \sim \mathcal{N}(\mu_x, \Sigma_x)$. The generator proposed by Miok et al. (2019) uses the Monte Carlo dropout (Gal & Ghahramani, 2016) on the trained decoder. The idea of this generator is to propagate the instance x through the encoder to obtain its latent encoding z. This can be propagated many times through the decoder, obtaining every time a different result due to the Monte Carlo dropout, but preserving similarity to the original instance x.

The RBF network (Moody & Darken, 1989) uses Gaussian kernels as hidden layer units in a neural network. Once the parameters of the network are learned, the rbfDataGen generator (Robnik-Šikonja, 2016) can sample from the normal distributions, defined with obtained Gaussian kernels, to generate new instances.

The TreeEnsemble generator (Robnik-Šikonja, 2019) builds a set of random trees (forest) that describe the data. When generating new instances, the generator traverses from the root to the leaves of a randomly chosen tree, setting values of features in the decision nodes on the way. When reaching a leaf, it assumes that it has captured the dependencies between features. The remaining features can, therefore, be generated independently according to the observed empirical distribution in this leaf. For each generated instance, all attributes can be generated in one leaf, or another tree can be randomly selected where unassigned feature values are filled in. By selecting different trees, different features are set in the interior nodes and leaves.

2.4 Related work on robustness of explanations

The adversarial attacks on perturbation based explanation methods were proposed by Slack et al. (2020), who show that LIME and SHAP are vulnerable due to the perturbation based sampling used. We propose the solution to the exposed weakness in SHAP and IME based on better sampling using data generators adapted to the training set distribution.

In general, the robustness of explanation methods has been so far poorly researched. There are claims that post-hoc explanation methods shall not be blindly trusted, as they can mislead users (deliberately or not) and disguise gender and racial discrimination (Lipton, 2016). Selbst & Barocas (2018) and Kroll et al. (2017) showed that even if a model is completely transparent, it is hard to detect and prevent bias due to the existence of correlated variables.

Specifically, for deep neural networks and images, there exist adversarial attacks on saliency map based interpretation of predictions, which can hide the model's bias (Dombrowski et al., 2019; Heo et al., 2019; Ghorbani et al., 2019). Dimanov et al. (2020) showed that a bias of a neural network could be hidden from post-hoc explanation methods by training a modified classifier, that has similar performance to the original one but the importance of the chosen feature is significantly lower.

The kNN-based explanation method, proposed by Chakraborty et al. (2020), tries to overcomes the inadequate perturbation based sampling used in explanation methods by finding similar instances to the explained one in the training set instead of generating new samples. This solution is inadequate for realistic problems as the problem space is not dense enough to get reliable explanations. Our defence of current post-hoc methods is based on superior sampling, which has not yet been tried. We are also not aware of any other defences against the adversarial attacks on post-hoc explanations.

3 Robustness through better sampling

We propose the defence against the adversarial attacks on explanation methods that replaces the problematic perturbation sampling with a better one, thereby making the explanation methods more robust. We want to generate the explanation data in such a way that the attacker cannot determine whether an instance is sampled or obtained from the original data. With an improved sampling, the adversarial model shown in Figure 1b shall not be able to determine whether the instance x was generated by the explanation method, or it is the original instance the model has to label. With

a better data generator, the adversarial model cannot adjust its output properly and the deception, described in Section 2.2, becomes ineffective.

The reason for the described weakness of LIME and SHAP is inadequate sampling used in these methods. Recall that LIME samples new instances by adding Gaussian noise to the normalized feature values. SHAP samples new instances from clusters obtained in advance with the k-means algorithm from the training data.

Instead of using the Gaussian noise with LIME, we generate explanation samples for each instance with one of the three better data generators, MCD-VAE, rbfDataGen, or TreeEnsemble (see Section 2.3). We call the improved explanation methods gLIME and gSHAP (g stands for generator-based). Using better generators in the explanation methods, the decision function in the adversarial model will less likely determine which predicted instances are original and which are generated for the explanation purposes.

Concerning gLIME, we generate data in the vicinity of the given instance using MCD-VAE, as the LIME method builds a local model. Using the TreeEnsemble and rbfDataGen generators, we do not generate data in the neighbourhood of the given instance but leave it to the proximity measure of the LIME method to give higher weights to instances closer to the explained one.

In SHAP, the perturbation sampling replaces the values of hidden features in the explained instance with the values from the distribution set D. The problem of this approach is that it ignores the dependencies between features. For example, in a simple dataset with two features, house size and house price, let us assume that we hide the house price, but not the house size. These two features are not independent, because the price of a house increases with its size. Suppose we are explaining an instance that represents a large house. Disregarding the dependency, using the sampled set D, SHAP creates several instances with a low price, as such instances appeared in the training set. In this way, the sampled set contains many instances with a small price assigned to a large house, from which the attacker can determine that these instances were created in perturbation sampling and serve only for the explanation.

In the proposed gSHAP, using the MCD-VAE and TreeEnsemble generators, the distribution set D is generated in the vicinity of the explained instance. In the sampled instances, some feature values of the explained instance are replaced with the generated values, but the well-informed generators take into account dependencies between features detected in the original distribution. This will make the distribution of the generated instances very similar to the original distribution. In our example, the proposed approach generates new instances around the instance that represents a large house, and most of these houses will be large. As the trained generators capture the dependencies from the original dataset, these instances will also have higher prices. This will make it difficult for the attacker to recognize the generated instances, used in explanations. The advantage of generating the distribution set close to the observed instance is demonstrated in Appendix B.

The rbfDataGen generator does not support the generation of instances around a given instance. We, therefore, generate the sampled set based on the whole training set and not for each instance separately (we have this option also for TreeEnsemble). This is worse than generating the distribution set around the explained instance but still better than generating it using the k-means sampling in SHAP. There are at least three advantages. First, the generated distribution set D can be larger. The size of the k-means distribution set cannot be arbitrarily large, because it is limited with the number of clusters in the training set. Second, the centroids of the clusters obtained by the k-means algorithm do not necessarily provide a good summary of the training set and may not be valid instances from training set distribution. They also do not capture well the density of instances (e.g., most of the data from the training set could be grouped in one cluster). Third, using the proposed generators, SHAP becomes easier to use compared to the k-means generator, where users have to determine the number of clusters, while the data generators we tested can be used in the default mode without parameters,

4 EVALUATION

To evaluate the proposed improvements in the explanation methods, we first present the used datasets in Section 4.1, followed by the experiments. We test the robustness of gLIME, gSHAP, and gIME against the adversarial models. To be more realistic, we equipped the adversarial models with the

Data set	# inst.	# features	# categorical	sensitive	unrelated 1	unrelated 2
COMPAS	6172	7	4	*race*	*random1*	*random2*
German	1000	25	15	*gender*	*pctOfIncome*	/
CC	1994	100	0	*racePctWhite*	*random1*	*random2*

Table 1: Basic information and sensitive features in the the used data sets. The target variable is not included in the number of features and is binary for all data sets. The *pctOfIncome* full name is *loanRateAsPercentOfIncome*.

same improved data generators we used in the explanation methods. It is reasonably to assume that attackers could also use better data generators when preparing their decision function as this would make their attacks much stronger. As the evaluation metric for the success of deception, we use the proportion of instances where the adversarial model deceived the explanation methods so that they did not detect sensitive attributes as being important in the prediction models.

4.1 Sensitive datasets prone to deception

Following (Slack et al., 2020), we conducted our experiments on three data sets from domains where a biased classifier could pose a critical problem, such as granting a credit, prediction of crime recidivism, and prediction of violent crime rate. The basic information on the data sets are presented in Table 1. The statistics were collected after removing the missing values from the data sets, and before we encoded categorical features as one-hot-encoded vectors.

COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) is a risk assessment used by the courts in some US states to determine the crime recurrence risk of a defendant. The dataset (Angwin et al., 2016)) includes criminal history, time in prison, demographic data (age, gender, race), and COMPAS risk assessment of the individual. The dataset contains data of 6,172 defendants from Broward Couty, Florida. The sensitive attribute in this dataset (the one which the adversarial model will be biased on) is the race. African Americans, whom biased model associates with a high risk of recidivism, represent 51.4% of instances from the data set. The target variable in this set is the COMPAS score, which is divided into two classes: a high and low risk. The majority class is the high risk, which represents 81.4% of the instances.

The German Credit dataset (German for the rest of the paper) from the UCI repository (Dua & Graff, 2019) includes financial (bank account information, loan history, loan application information, etc.) and demographic data (gender, age, marital status, etc.) for 1,000 loan applicants. A sensitive attribute in this data set is gender. Men, whom the biased model associates with a low-risk, represent 69% of instances. The target variable is the loan risk assessment, which is divided into two classes: a good and bad customer. The majority class is a good customer, which represents 70% of instances.

Communities and Crime (CC) data set (Redmond & Baveja, 2002) contains data about the rate of violent crime in US communities in 1994. Each instance represents one community. The features are numerical and represent the percentage of the population in the community with a certain property or the average of population in the community. Features include socio-economic (e.g., education, house size, etc.) and demographic (race, age) data. The sensitive attribute is the percentage of the white race. The biased model links instances where the percentage of whites is above average to a low rate of violent crime. The target variable is the rate of violent crime, which is divided into two classes: high and low. Both classes are equally represented in the data set.

4.2 Robustness of explanation methods

To evaluate the robustness of explanation methods with added generators (i.e. gLIME, gSHAP, and gIME), we split the data into training and evaluation set in the ratio $90\% : 10\%$. We used the same training set for training of adversarial models and explanation methods. We encoded categorical features as one-hot-encoded vectors.

We simulated the adversarial attack for every combination of generators used in explanation methods and adversarial models (except for method IME, where we did not use rbfDataGen which is unable to generate instances in the neighbourhood of a given instance). When testing SHAP, we used two variants of the TreeEnsemble generator for explanation: generating new instances around the explained instance and generating new instances according to the whole training set distribution. In the LIME testing, we used only the whole training set variant of TreeEnsemble inside the

explanation. In the IME testing, we only used the variant of Tree Ensemble that fills in the hidden values (called TEnsFillIn variant). For training of the adversarial models, we used the whole training set based variant of TreeEnsemble, except for IME, where we used the TEnsFillIn variant. These choices reflect the capabilities of different explanation method and attempt to make both defence and attack realistic (as strong as possible).

In all cases, the biased model b (see Section 2.2) was a simple function depending only on the value of the sensitive feature. The unbiased model ψ depended only on the values of unrelated features. The sensitive and unrelated features are shown in the right-hand side of Table 1. Features *random1* and *random2* were uniformly randomly generated from the $\{0, 1\}$ set. On COMPAS and CC, we simulated two attacks, one with a single unrelated feature (the result of ψ depends only on the value of the unrelated feature 1), and another one with two unrelated features.

For every instance in the evaluation set, we recorded the most important feature according to the used explanation method (i.e. the feature with the largest absolute value of the contribution as determined by the explanation method). The results are shown as a heatmap in Figure 2. The green colour means that the explanation method was deceived in less than 30% of cases, and the red means that it was deceived in more than 70% of cases. We consider deception successful if the sensitive feature was not recognized as the most important by the explanation method.

The gLIME method is more robust with the addition of rbfDataGen and TreeEnsemble than LIME and less robust with the addition of MCD-VAE. This suggests that parameters for MCD-VAE were not well-chosen, and this pattern repeats for other explanation methods. Both TreeEnsemble and rbfDataGen make gLIME considerably more robust on COMPAS and German datasets, but not on the CC dataset. We believe the reason for that is that features in CC are strongly interdependent, and many represent the same attribute as a fraction of value, e.g., we have the share of the white population, the share of the Asian population, and so on. This interdependence dictates that all fractions have to add up to 1. The data generators are unlikely to capture so strong conditional dependencies, but the decision function in the adversarial model is likely to approximate it.

The gSHAP method is most robust when using the TreeEnsemble generator, but it shows less robust behaviour than the gLIME method. The reason for that could be in the feature value replacement strategy used by the SHAP and gSHAP methods, which change only some of the feature's values with the help of the distribution set. This can lead to out of distribution instances if the instances in the distribution set are not close enough to the explained instance. With gLIME, we generate complete instances which are more likely to be in the distribution of the training set.

IME is quite robust even with the perturbation sampling, as we expected. This suggests that IME is the most robust of all three tested explanation methods and shall be considered the method of choice in sensitive situations. The gIME results when using the TreeEnsemble generator (TEnsFillIN variant) are comparable to the results of the original IME variant. This suggests that sampling from a smaller data set, which represents the neighbourhood of the explained instance does not decrease the robustness of the method.

5 Conclusions

We presented the defence against adversarial attacks on explanation methods. The attacks exploit the shortcomings of perturbation sampling in post-hoc explanation methods. This sampling used in these methods produces instances too different from the original distribution of the training set. This allows unethical owners of biased prediction models to detect which instances are meant for explanation and label them in an unbiased way. We replaced the perturbation sampling with data generators that better capture the distribution of a given data set. This prevents the detection of instances used in explanation and disarms attackers. We have shown that the modified gLIME and gSHAP explanation methods, which use better data generators, are more robust than the original variants, while IME is already quite robust. Our preliminary results in Appendix C show that using the TreeEnsemble generator, the gIME methods converges faster and requires from 30-50% fewer samples.

The effectiveness of the proposed defence depends on the choice of the data generator and its parameters. While the TreeEnsemble generator turned out the most effective in our evaluation, in practice several variants might need to to be tested to get a robust explanation method. Inspecting authorities

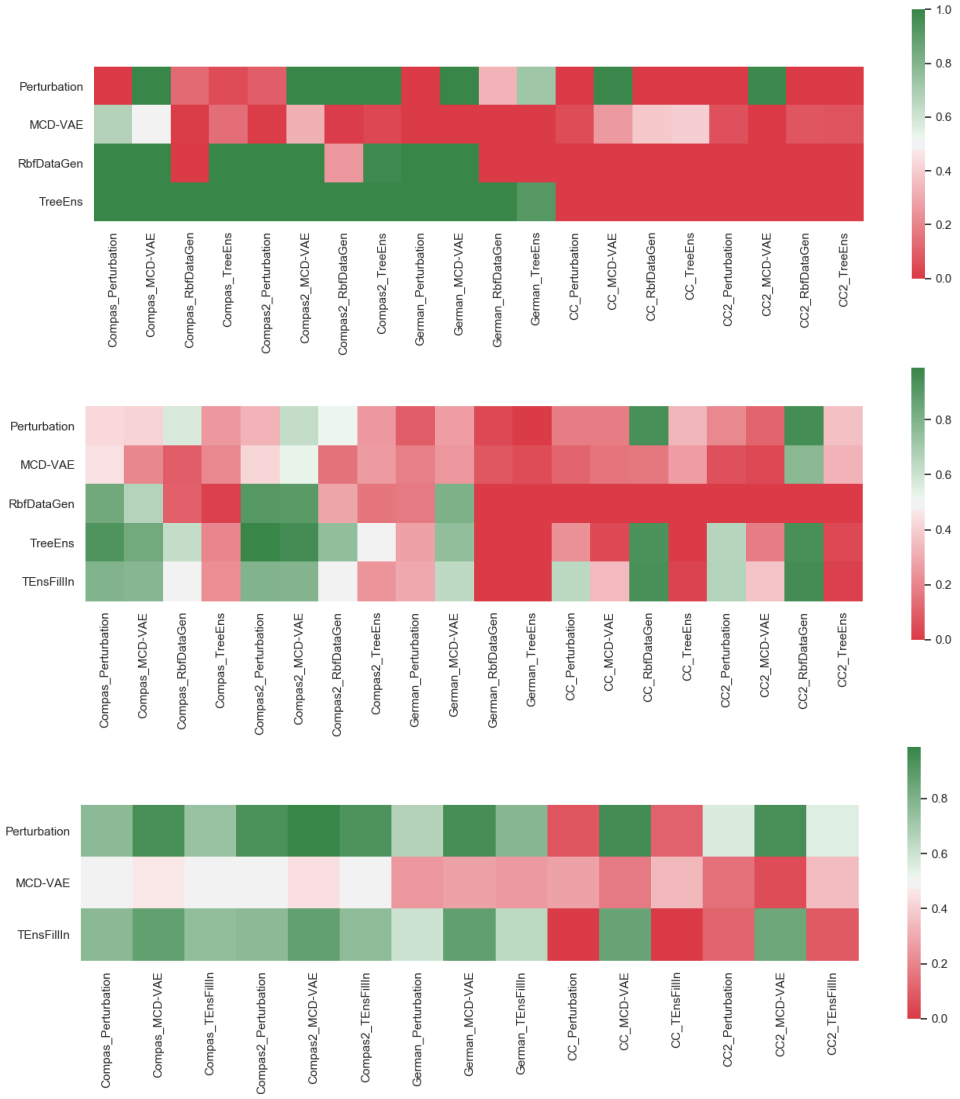

Figure 2: The robustness results for gLIME (top), gSHAP (middle), and gIME (bottom). The graphs show the proportion of evaluation set instances, where the sensitive feature was recognized as the most important by the used explanation method. Rows represent the generators used for explanations. The column labels consist of the name of the data set on which the experiment was performed and the name of the generator used for training of the adversarial model. Compas2 and CC2 denote an attack with two independent features. Perturbation represents the original sampling used in LIME, SHAP, and IME, TEnsFillIn represents the TreeEnsemble variant where new instances are generated around the given one, and TreeEns represents the generation from the whole distribution.

shall be aware of the need for good data generators and make the access to training data of sensitive prediction models a legal requirement. Luckily, even a few non-deceived instances would be enough to raise an alarm about the unethical model.

This work opens a range of possibilities for further research. The proposed defence and attacks shall be tested on other data sets with different number and types of features, with missing data, and on different types of problems such as text, images, and graphs. The work on useful generators shall be extended to find time-efficient generators with easy to set parameters and ability to generate new instances in the vicinity of a given one. Generative adversarial networks (Goodfellow et al., 2014) may be a promising line of research. The TreeEnsemble generator, currently written in pure R, shall be rewritten in a more efficient programming language. The MCD-VAE generator shall be studied to allow automatic selection of reasonable parameters for a given dataset. In SHAP sampling, we could first fix the values of the features we want to keep, and the values of the others would be generated using the TreeEnsemble generator.

REFERENCES

Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. *ProPublica, May 23*, 2016.

Joymallya Chakraborty, Kewen Peng, and Tim Menzies. Making fair ML software using trustworthy explanation. *arXiv preprint arXiv:2007.02893*, 2020.

Botty Dimanov, Umang Bhatt, Mateja Jamnik, and Adrian Weller. You shouldn't trust me: Learning models which conceal unfairness from multiple explanation methods. In *SafeAI@ AAAI*, pp. 63–73, 2020.

Carl Doersch. Tutorial on variational autoencoders. *ArXiv*, abs/1606.05908, 2016.

Ann-Kathrin Dombrowski, Maximillian Alber, Christopher Anders, Marcel Ackermann, Klaus-Robert Müller, and Pan Kessel. Explanations can be manipulated and geometry is to blame. In *Advances in Neural Information Processing Systems*, pp. 13589–13600, 2019.

Dheeru Dua and Casey Graff. UCI machine learning repository. `http://archive.ics.uci.edu/ml`, 2019. [Accessed: 9. 8. 2020].

Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pp. 1050–1059, 2016.

Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:3681–3688, 2019.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.

Juyeon Heo, Sunghwan Joo, and Taesup Moon. Fooling neural network interpretations via adversarial model manipulation. In *Advances in Neural Information Processing Systems*, pp. 2925–2936, 2019.

Joshua A Kroll, Joanna Huey, Solon Barocas, Edward W Felten, Joel R Reidenberg, David G Robinson, and Harlan Yu. Accountable algorithms. *University of Pennsylvania Law Review*, 165(3): 633–705, 2017.

Zachary Chase Lipton. The mythos of model interpretability. *CoRR*, abs/1606.03490, 2016.

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*, pp. 4765–4774. Curran Associates, Inc., 2017.

Kristian Miok, Deng Nguyen-Doan, Daniela Zaharie, and Marko Robnik-Šikonja. Generating data using Monte Carlo dropout. In *2019 IEEE 15th International Conference on Intelligent Computer Communication and Processing (ICCP)*, pp. 509–515, 2019.

J. Moody and C. J. Darken. Fast learning in networks of locally-tuned processing units. *Neural Computation*, 1:281–294, 1989.

Michael Redmond and Alok Baveja. A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research*, 141: 660–678, 2002.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144, 2016.

Marko Robnik-Šikonja. *semiArtificial: Generator of Semi-Artificial Data*, 2019. URL `https://cran.r-project.org/package=semiArtificial`. R package version 2.3.1.

Marko Robnik-Šikonja. Data generators for learning systems based on RBF networks. *IEEE Transactions on Neural Networks and Learning Systems*, 27(5):926–938, 2016.

Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53 – 65, 1987.

Andrew D Selbst and Solon Barocas. The intuitive appeal of explainable machines. *Fordham Law Review*, 87:1085, 2018.

Lloyd S. Shapley. A value for n-person games. In Alvin E. Roth (ed.), *The Shapley Value: Essays in Honor of Lloyd S. Shapley*, pp. 31–40. Cambridge University Press, 1988.

Dylan Slack, Sophie Hilgard, Emily Jia, Sameer Singh, and Himabindu Lakkaraju. Fooling LIME and SHAP: Adversarial attacks on post-hoc explanation methods. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2020.

Erik Štrumbelj and Igor Kononenko. Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41:647–665, 2013.

A DETAILS ON POST-HOC EXPLANATION METHODS

For the sake of completeness, we present further details on the explanation methods LIME (Ribeiro et al., 2016), SHAP (Lundberg & Lee, 2017), and IME (Štrumbelj & Kononenko, 2013). Their complete description can be found in the above-stated references. In our exposition of the explanation methods, we denote with f the predictive model, with x the instance we are explaining, and with n the number of features describing x.

A.1 LIME

Explanations of instances with the LIME method is obtained with an interpretable model g. The model g has to be both locally accurate (so that it can obtain correct features contributions) and simple (so that it is interpretable).

The explanation of the instance x for the predictive model f, obtained with the LIME method, is defined with the following equation:

$$\xi(x) = \arg\min_{g \in G}(\mathcal{L}(f, g, \pi_x) + \Omega(g)), \tag{2}$$

where $\Omega(g)$ denotes the measure of complexity for interpretable model, $\pi_x(z)$ denotes the proximity measure between x and generated instances z, $\mathcal{L}(f, g, \pi_x)$ denotes the measure of local fidelity of interpretable model g to the prediction model f, and G denotes the set of interpretable models. We use the linear version of the LIME method, where G represents the set of linear models. With x' we denote the normalized presentation of instance x, i.e. the numerical attributes have their mean set to 0 and variance to 1, and the categorical attributes contain value 1 if they have the same value as the exoplained instance and 0 otherwise. The proximity function π_x is defined on the normalized instances (hence we use the notation $\pi_{x'}$), and uses the exponential kernel: $\pi_{x'}(z') = e^{\frac{d(x', z')}{\sigma^2}}$, where $d(x', z')$ denotes a distance measure between x' and z'. The local fidelity measure $\mathcal{L}$ from Equation (2) is defined as:

$$\mathcal{L}(f, g, \pi_{x'}) = \sum_{z' \in \mathcal{Z}} \pi_{x'}(z')(f(z) - g(z'))^2, \tag{3}$$

where $\mathcal{Z}$ denotes the set of samples. In LIME, each generated sample is obtained by adding Gaussian noise to each feature of x' independently.

Using linear models as the set of interpretable models, LIME is relatively fast but may produce poor explanations for instances close to complex decision boundaries.

A.2 SHAP

We refer to SHAP as the method called Kernel SHAP by (Lundberg & Lee, 2017). SHAP essentially estimates Shapley values using the approach of LIME, which means that the calculation of explanations is fast due to the use of local linear models computed with the weighted least squares algorithm. The explanation of the instance x with the SHAP method are feature contributions ϕ_i, $i = 1, 2, ..., n$ that are coefficients of the linear model $g(x') = \phi_0 + \sum_{i=1}^{n} \phi_i \cdot x'_i$, obtained with LIME, where $x'_i \in \{0, 1\}$ for $i \in \{1, 2, ..., n\}$.

As LIME does not compute Shapley values, this property is assured with the proper selection of functions $\Omega(g), \pi_{x'}(z')$, and $\mathcal{L}(f, g, \pi_{x'})$. Let us first define the function $h_x(z')$ that maps from the $\{0, 1\}^n$ to the original feature space. The function h_x is defined implicitly with the equation $f(h_x(z')) = E[f(z)|z_i = x_i \; \forall i \in \{j; z'_j = 1\}]$. This is the expected value of the model f when we do not know the values of features at indices where z' equals 0 (these values are hidden). Functions $\Omega(g), \pi_{x'}(z')$ and $\mathcal{L}(f, g, \pi_{x'})$ that enforce the computation of Shapley values are:

$$\Omega(g) = 0,$$

$$\pi_{x'}(z') = \frac{n}{\binom{n}{|z'|} \cdot |z'| \cdot (m' - |z'|)},$$

$$\mathcal{L}(f, g, \pi_{x'}) = \sum_{z' \in \mathcal{Z}} (f(h_x(z')) - g(z'))^2 \cdot \pi_{x'}(z'),$$

where $|z'|$ denotes the number of nonzero features of z' and $\mathcal{Z} \subseteq 2^{\{0,1\}^n}$.

The main purpose of the sampling in this method is to determine the value $f(h_x(z'))$ because in general predictive models cannot work with hidden values. To determine $f(h_x(z'))$, SHAP uses the distribution set D which we obtain from the training set. For D, SHAP takes the centroids of clusters, obtained by k-means clustering algorithm on the training set. The number of clusters is set by the user. Value of $f(h_x(z'))$ is determined by the following sampling:

$$f(h_x(z')) = E[f(z)|z_i = x_i \; \forall i \in \{j; z'_j = 1\}] = \frac{1}{|D|} \sum_{d \in D} f(x_{[x_i = d_i, z'_i = 0]}), \tag{4}$$

where $x_{[x_i = d_i, z'_i = 0]}$ denotes instance x with features that are 0 in z' being set to the feature values from d.

A.3 IME

The explanation of x with the method IME are feature contributions ϕ_i, $i = 1, 2, ..., n$. Štrumbelj & Kononenko (2013) shoved that Shapley values are the only solution that takes into account contributions of interactions for every subset of features in a fair way. The i-th feature contribution can be calculated with the following expresiion:

$$\phi_i(x) = \frac{1}{n} \sum_{\pi \in S_n} \sum_{w \in \mathcal{X}} p(w) \cdot (f(w_{[w_i = x_i, i \in Pre^i(\pi) \cup \{i\}]}) - f(w_{[w_i = x_i, i \in Pre^i(\pi)]})), \tag{5}$$

where S_n denotes a group of permutations of n elements, $\mathcal{X}$ denotes the training set instances, $Pre^i(\pi)$ represents the set of indices that precedes i in the permutation π, i.e. $Pre^i(\pi) = \{j; \pi(j) < \pi(i)\})$. Let $p(w)$ denote the probability of the instance w in $\mathcal{X}$ and let $w_{[formula]}$ denote the instance w with some of its features values changed according to the *formula*. To calculate $\phi_i(x)$, we have to go through $|\mathcal{X}| \cdot n!$ iterations, which can be slow. Therefore, the method IME uses the following sampling. The sampling population of i-th feature is $V_{\pi,w} = (f(w_{[w_i = x_i, i \in Pre^i(\pi) \cup \{i\}]}) - f(w_{[w_i = x_i, i \in Pre^i(\pi)]}))$ for every combination of the permutation π and instance w. IME draws m_i samples $V_1, ..., V_{m_i}$ at random with repetition. The estimate of $\phi_i(x)$ is defined with the equation:

$$\hat{\phi}_i = \frac{1}{m_i} \sum_{i=1}^{m_i} V_i. \tag{6}$$

Contrary to SHAP, IME does not use approximation with linear models which compute the contributions of all features at once but has to compute the Shapley values by averaging over large enough sample for each feature separately. This makes the method slower but also potentially more robust as the method does not assume the shape of the dependency in the space of normalized features.

B Demonstration of better sampling in SHAP

The graphs in Figure 3 show the PCA based 2D space of the evaluation part of the COMPAS dataset (see Section 4 for the dataset description). The left-hand side shows the SHAP-generated sampled instances using the k-means algorithm (14 clusters determined by the silhouette score (Rousseeuw (1987)). The sample produced with the MCD-VAE generator in gSHAP is shown on the right-hand side. This sample is much more similar to the original distribution compared to the SHAP sampling.

C Improved IME convergence rate with the TreeEnsemble generator

Preliminary, we tested how the use of better generators affects the convergence rate of the IME explanation method. The preliminary results in Table 2 show a significant reduction in the number of needed samples and a slight increase in the error compared to the original perturbation sampling. Note that the error measure deployed is biased in favor of the perturbation sampling, which was used to determine the gold standard.

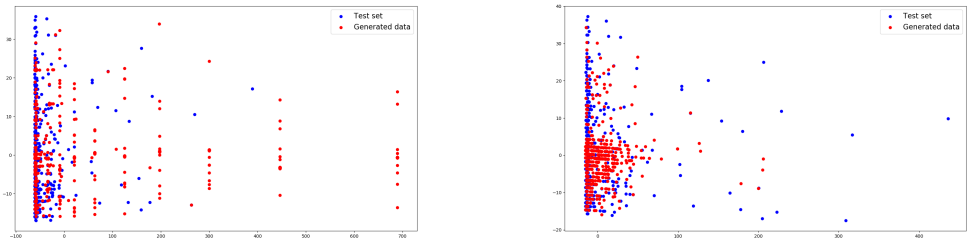

Figure 3: Visual comparison of original and sampled distributions for the COMPASS dataset. The SHAP k-means based generator (left) produces instances less similar to the original data, compared to the MCD-VAE generator (right).

COMPAS dataset

	Error		# samples			
Classifier	Perturb.	TEnsFillIn	Perturb.	TEnsFillIn	Reduction %	CA %
Naive Bayes	0.0076	0.0217	18571	11278	39	83
Linear SVM	0.0033	0.0080	8244	4423	46	84
Random forest	0.0049	0.0221	45372	26960	41	80
Neural network	0.0057	0.0130	16157	8841	45	84

German dataset

	Error		# samples			
Classifier	Perturb.	TEnsFillIn	Perturb.	TEnsFillIn	Reduction %	CA %
Naive Bayes	0.0076	0.0201	56052	39257	30	77
Linear SVM	0.0005	0.0013	3877	2157	44	69
Random forest	0.0046	0.0141	92478	66639	28	74
Neural network	0	0	0	26	/	69

CC dataset

	Error		# samples			
Classifier	Perturb.	TEnsFillIn	Perturb.	TEnsFillIn	Reduction %	CA %
Naive Bayes	0.0028	0.0046	32910	20117	39	70
Linear SVM	0.0009	0.0048	73324	39098	47	62
Random forest	0.0032	0.0045	109958	58852	46	79
Neural network	0.0012	0.0061	144183	70020	51	72

Table 2: Comparison of original perturbation sampling and the TreeEnsemble generator with data fill-in inside the IME method. The results show average scores for the evaluation set. The column Perturb. presents the perturbation based sampling and TEnsFillIn presents the sampling using TreeEnsemble generator with missing parts of instances filled in. The Reduction column shows the reduction in the number of samples using the TEnsFillIn method compared to perturbations. The CA stands for the classification accuracy of the explained classifier on the evaluation set.