Increasing Performance And Sample Efficiency With Model-agnostic Interactive Feature Attributions

Joran Michiels *[1], Maarten De Vos[1,2], and Johan Suykens[1]

[1]ESAT-STADIUS, KU Leuven
[2]Department of Development and Regeneration, KU Leuven
{joran.michiels,maarten.devos,johan.suykens}@kuleuven.be

June 28, 2023

Abstract

Model-agnostic feature attributions can provide local insights in complex ML models. If the explanation is correct, a domain expert can validate and trust the model's decision. However, if it contradicts the expert's knowledge, related work only corrects irrelevant features to improve the model. To allow for unlimited interaction, in this paper we provide model-agnostic implementations for two popular explanation methods (Occlusion and Shapley values) to enforce entirely different attributions in the complex model. For a particular set of samples, we use the corrected feature attributions to generate extra local data, which is used to retrain the model to have the right explanation for the samples. Through simulated and real data experiments on a variety of models we show how our proposed approach can significantly improve the model's performance only by augmenting its training dataset based on corrected explanations. Adding our interactive explanations to active learning settings increases the sample efficiency significantly and outperforms existing explanatory interactive strategies. Additionally we explore how a domain expert can provide feature attributions which are sufficiently correct to improve the model.

1 Introduction

With machine learning models becoming more complex and their application more widespread, pressure to understand the model's decisions is increasing. This has resulted in the proposal of many approaches to increase insights into

*Corresponding author

the model. To not compromise the performance of complex models and allow for straight-forward model comparison, model-agnostic post-hoc explanations are preferred. Their generality and ease of use has made them a very popular choice to validate the decision of the model and/or increase trust. But what if the explanation is actually wrong according to the expert? This is a perfect opportunity to improve the model by changing its explanation, especially when data is scarce or during active learning when labelling is difficult. In this work we show how model-agnostic post-hoc explanations can be made fully interactive, in particular how corrections to these explanations can be enforced in any model by modifying the training data. This extends the existing state-of-the-art [1] which only allows limited correction of irrelevant features. Besides, this work creates a direct connection between feature attributions and model improvement, providing a use case to test human interpretation of feature importance in further work.

The specific contributions of this paper are:

- a model optimization setup using interactive explanations, incorporating standard supervised learning and the previously proposed active learning setup [1, 2];
- implementations of two novel interactive feature attribution methods, respectively for Occlusion [3] and the much used SHAP explanations [4], that significantly improve model performance compared to related work [1] and standard active learning (increasing sample efficiency by more than 50% in all experiments);
- theoretical validation of our methods and comparison with related work [1];
- a novel experiment exploring how domain knowledge can be encoded in feature attributions which can be successfully used to improve model performance using our approach.

This paper will first discuss the necessary background and related work. Next we introduce our interactive feature attributions and provide some theoretical validation. Finally we discuss our experiments and conclude with an outlook on future work.

2 Background

Interactive explanations have mostly appeared in the context of active learning [1]. The main goal of *active learning* is reducing the amount of labelling a domain expert has to perform. This is achieved by only asking the expert to label the samples which are difficult to the model, e.g. in the context of classification samples who are close to the decision boundary are probably more interesting to the model. We denote the set of interesting samples as the *query*. The procedure is repeated iteratively to achieve maximal sample efficiency. For a

complete overview of active learning we refer to [5]. [1] proposed the addition of explanations in the active learning loop to increase trust (of the expert) in the model and further increase sample efficiency. An explanation accompanies each interesting sample. Therefore it is *local* in contrast to global, which attempts to explain the whole model.

In this paper the focus will be on a specific type of local explanations, *feature attributions*. They explain the model output $f(\boldsymbol{x})$ of a single sample $\boldsymbol{x}$ (note the bold font) by assigning an attribution R_i to each feature value $\boldsymbol{x}_i$. The attributions can be interpreted as signed feature importances or local feature effects on the model output. They can be either model-specific or model-agnostic. Model-specific methods may require certain model characteristics (differentiability for Integrated Gradients [6]) or can be efficient implementations of a model-agnostic concept (SHAP values for trees [7]). Model-agnostic implementations treat the model as a black-box and generate explanations through sampling it.

A simple model-agnostic attribution is Occlusion [3]:

$$R_i = f(\boldsymbol{x}) - f([\boldsymbol{x}_{\bar{i}}, b_i]), \tag{1}$$

with b a background input and $\bar{i}$ all features except i. This background input is an important design choice in many model-agnostic implementations and can be seen as the baseline to which we compare the importance of the feature value. Equation (1) can be interpreted as the change in model output if feature value $\boldsymbol{x}_i$ becomes known. One could imagine other attributions that are equivalent:

$$R_{i,S} = f([\boldsymbol{x}_{S_i}, b_{\bar{S}_i}]) - f([\boldsymbol{x}_S, b_{\bar{S}}]), \tag{2}$$

with S a subset of the features, S_i the same set including i and $\bar{S}$ all features except the ones in S. Considering an ordering p in which each feature gets known, $R_{i,S}$ is the attribution if feature i becomes known after the features in S where the unknown values are replaced by corresponding values from the background b. If it is unclear which ordering p (or which corresponding set of features S^p) is the best one, the average over all $p \in P$ can be used:

$$R_i = \frac{1}{|P|} \sum_p f\left(\left[\boldsymbol{x}_{S_i^p}, b_{\bar{S}_i^p}\right]\right) - f\left([\boldsymbol{x}_{S^p}, b_{\bar{S}^p}]\right). \tag{3}$$

These are known as SHAP values (Shapley additive explanations) [4]. They are the only feature attributions satisfying a desirable set of properties including *local accuracy*:

$$f(\boldsymbol{x}) = f(b) + \sum_{i=1}^{M} R_i, \tag{4}$$

meaning that the feature attributions (M is total number of features) sum to the difference in the model output between the explained sample $\boldsymbol{x}$ and the background sample b.

Both Occlusion and SHAP can also be averaged over many background inputs. In fact, the original SHAP value definition computes the expectation

of R_i over all b [4]. Our paper actually uses so-called Baseline Shapley values [8]. As will become clear later, assuming a single background input allows for a more straight-forward interactive implementation.

Lastly, note that for linear models $f(x) = \sum_i w_i x_i + w_0$ Occlusion and SHAP values give the same attribution namely:

$$R_i = w_i(\boldsymbol{x}_i - b_i). \tag{5}$$

3 Related work

In [1, 2] interactive local post-hoc explanations were explored for the first time. [1] introduces the framework of *explanatory interactive learning* (XIL) combining active learning with explanations. Our proposed methods do not distinguish between regular supervised learning and active learning, although the active learning seems a natural setting to exploit our methods.

[1] proposes a model-agnostic setup named CAIPI. If the explanation of a sample $\boldsymbol{x}$ wrongly identified a subset of features C as relevant (as per expert feedback), the data is augmented with so-called counterexamples of the form $([\boldsymbol{x}_{\bar{c}}, b_c], \boldsymbol{y})$ for every $c \in C$ with $\boldsymbol{y}$ the corrected target and with b either a random sample, fixed alternative sample or another sample from the training dataset. They show how their counterexamples can easily correct models which are fooled by features (confounders) who are wrongly correlated with the output, also known as the 'right for the wrong reason' case or the *Clever Hans* scenario [2]. Through classification experiments on data with known irrelevant features they show how their approach significantly improves the explanations of the models w.r.t. to standard interactive learning. Our approach allows for correcting all feature attribution in stead of only correcting irrelevant features. We will compare their method with ours in a series of experiments. While [1] provides results using LIME explanations [9], their approach can work with any type of explanation method. For maximum comparability, we implement their method with the same explanations as our approach.

Other similar work on specific models (thus not comparable to our model-agnostic approach) includes [10] where input gradients are constrained to improve differentiable models effectively and efficiently. [11, 12] explore how domain knowledge can be used to constrain a support vector machine in a supervised learning setting while [13] examines how experts can improve a Naïve Bayes model. For more related work, we refer to [14].

[2] further investigates XIL in the Clever Hans scenario and successfully uses the method of [10] on a plant phenotyping task. In [15] several methods are combined in a framework and benchmarked on specific tasks to correct the shortcut behaviour from Clever Hans models.

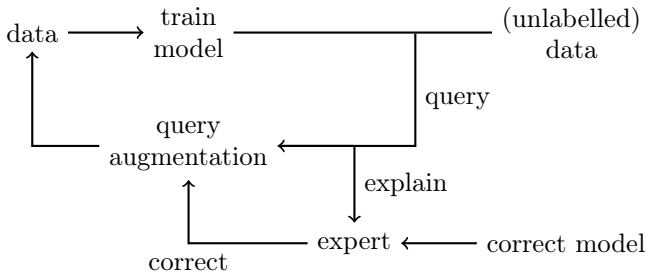

Figure 1: Model optimization setup with interactive explanations.

4 Improving the model through interactive feature attributions

We propose the model optimization setup in Figure 1 where a model can be iteratively improved through interactive explanations of selected samples from the complete dataset (including unlabelled). The model is first trained on the available labelled training data, whereafter an interesting query is selected and explained to the expert. For each sample of the query, the expert is allowed to correct the predicted label (note that if the sample was part of the labelled dataset, its correct label is already in the training data) and fully change its explanation.

It is assumed that a data generating process exists and we denote the model mimicking this data generating process perfectly on the available features as a *correct model* f^*. Note that this model might not achieve perfect accuracy on the test set, even without the presence of noise, because not all necessary features are available to make the model deterministic. On the other hand, multiple correct models might exist if more features are available to the model than necessary. In practice a correct model is unknown (otherwise machine learning would not be necessary) but we assume the domain expert has some knowledge of it which can be expressed through the chosen local explanation method.

To keep our method widely applicable and model-agnostic, the correct explanations are enforced on the model by retraining it on augmented data.

4.1 Query augmentation

For each sample $\boldsymbol{x}$ in the query, new samples are distilled from its (corrected) target $\boldsymbol{y}$ and its correct feature attributions. Knowledge of background input b is also assumed. Unless mentioned otherwise, the augmentation is identical for regression ($f(\cdot), \boldsymbol{y} \in \mathbb{R}$) and classification ($f(\cdot), \boldsymbol{y} \in \{0, 1\}$).

Correct explanations R_i^* for a sample $\boldsymbol{x}$ can be enforced in model f by retraining f so that its explanations R_i are equal to R_i^*. In the case of Occlusion,

by Equation (1) we know that:

$$\begin{aligned} R_i &= R_i^* \\ f(\boldsymbol{x}) - f([\boldsymbol{x}_{\bar{i}}, b_i]) &= R_i^*. \end{aligned} \tag{6}$$

Combining this constraint with $f(\boldsymbol{x}) = \boldsymbol{y}$ (either the target was already known or has been corrected by the expert) we get:

$$f([\boldsymbol{x}_{\bar{i}}, b_i]) = \boldsymbol{y} - R_i^* \tag{7}$$

Thus the sample $([\boldsymbol{x}_{\bar{i}}, b_i], \boldsymbol{y} - R_i^*)$ should be added to the training dataset. Note the similarity with the state-of-the-art [1], if $R_i^* = 0$ our sample augmentations are identical to theirs.

Repeating the procedure above for SHAP values (this time for all i), we get the following constraints:

$$\begin{cases} \forall i : \frac{1}{|P|} \sum_p f\left(\left[\boldsymbol{x}_{S_i^p}, b_{\bar{S}_i^p}\right]\right) - f\left([\boldsymbol{x}_{S^p}, b_{\bar{S}^p}]\right) = R_i^* \\ f(\boldsymbol{x}) = \boldsymbol{y}. \end{cases} \tag{8}$$

This set of equations is underdetermined so extra constraints have to be imposed. A sensible choice would be enforcing the same attribution for every subset. In the Appendix we show how this leads to training samples of the form $([\boldsymbol{x}_{\bar{S}_i}, b_{S_i}], \boldsymbol{y} - \sum_{j \in S_i} R_j^*)$ or $([\boldsymbol{x}_{\bar{S}_i}, b_{S_i}], H(\boldsymbol{y} - \sum_{j \in S_i} R_j^* - 0.5))$, with H the Heaviside function, in the case of classification.

Note that in classification settings it is customary to compute explanations on the decision function or class probability to get continuous and more fine-grained feature importances [16]. For example, for logistic regression it is easy to see why this is preferred: if we compute the explanation on the log odds, we are essentially explaining a linear model which has very intuitive (see Equation (5)) attributions. However in the context of model-agnostic interactive explanations it makes little sense, since most classifiers require samples with discrete targets and domain experts tend to provide a corrected class and not a class probability.

4.2 Theoretical validation

In the case of linear models, we can examine how the described query augmentation allows to correct the model. Consider a correct model f^* and incorrect model f,

$$f^*(x) = \sum_i w_i^* x_i + w_0^* \quad f(x) = \sum_i w_i x_i + w_0 \tag{9}$$

We aim to correct f by correcting a number of local attributions with background input b. For query $\boldsymbol{x}$, correct target $\boldsymbol{y}$ and feature attribution R_k^* the sample $([\boldsymbol{x}_{\bar{k}}, b_k], \boldsymbol{y} - R_k^*))$ is added to the training data. Including the original (corrected) sample, model f should adhere to the following constraints:

$$\begin{cases} \boldsymbol{y} = \sum_i w_i \boldsymbol{x}_i + w_0 \\ \boldsymbol{y} - R_k^* = \sum_{i \neq k} w_i \boldsymbol{x}_i + w_k b_k + w_0. \end{cases} \tag{10}$$

Subtracting the constraints, we find that $R_k^* = w_k(\boldsymbol{x}_k - b_k)$, so if we aim to correct $w' \to w_k^*$, $R_k^* \equiv w_k^*(\boldsymbol{x}_k - b_k)$. This is the SHAP and Occlusion attribution for linear models. This means that a single augmentation can correct the corresponding model parameter. One correct explanation can completely correct a linear model. As explained above, the approach from [1] correspond to setting $R_k^* = 0$ and is thus only correct if $w_k^* = 0$, otherwise it incorrectly pushes $w_k \to 0$.

For non-linear models, the approach with Occlusion comes down to adding an extra correct sample to the training data, and is thus expected to be beneficial. In case of SHAP, since the extra samples were produced after imposing constraints, they do not necessarily correspond to correct samples from the correct model. In the Clever Hans scenario [2] a confounded feature c has an attribution $R_c \neq 0$. Since $\boldsymbol{x}_c$ will not be meaningful, an expert should set $R_c^* = 0$, making our method equivalent to [1].

4.3 Implementation

To eliminate any bias towards the newest queries, the model is retrained on all data (including the augmented data) every iteration. To maintain a balance between adding new samples and adding corrections, for each correction, we propose to add the new sample to the dataset (thus oversampling it). In the case of Occlusion we also allow for incomplete explanations, if the expert has partial domain knowledge. Complete implementations are in Algorithm 1 and 2. Note that if we set $R_i = 0$ for all i who are irrelevant, Algorithm 1 can and will be used as an implementation for CAIPI [1].

5 Experiments

The goals of these experiments are to confirm our theoretical validation; to show the generality of our methods in different training tasks (linear and non-linear complex models, classification and regression); and examine the advantages over related work [1] on interactive explanations.

We compare our proposed methods which allow for full correction of the explanation (denoted by *Interactive Occlusion* or *Interactive SHAP*) against the state-of-the-art which only correct irrelevant features (denoted by *CAIPI*) and against regular active learning (only correcting the targets, denoted by *Baseline*). In the experiments, the interesting queries are selected by random. As a background value b either the average or the median (in case of discrete features) of the complete dataset is used. To fairly compare all strategies, the number of extra (augmented) samples per iteration is the same for all strategies. Our Occlusion implementation adds $2|R^*|$ samples per sample $\boldsymbol{x}$ per iteration (see Algorithm 1), thus in the results of standard active learning, $\boldsymbol{x}$ is oversampled $2|R^*|$ times. Similarly, in the results with SHAP values, $K = |R^*|$ in Algorithm 2. If R^* is not completely determined (some attributions are missing, this is the

Algorithm 1: Augmentation Occlusion

input : $X_{tr}, Y_{tr}, X_{query}, b$
for $\boldsymbol{x} \in X_{query}$ **do**
  if $\boldsymbol{x}$ *is unlabelled* **then**
    get $\boldsymbol{y}$ from expert;
  get R^* from expert;
  for $R_i^* \in R^*$ **do**
    if R_i^* *is available* **then**
      $\boldsymbol{x}_{aug} \leftarrow \boldsymbol{x}$;
      $\boldsymbol{x}_{aug,i} \leftarrow b_i$;
      $\boldsymbol{y}_{aug} \leftarrow \boldsymbol{y} - R_i^*$;
      add $\boldsymbol{x}_{aug}, \boldsymbol{y}_{aug}$ to (X_{tr}, Y_{tr});
      add $\boldsymbol{x}, \boldsymbol{y}$ to (X_{tr}, Y_{tr});
retrain model on (X_{tr}, Y_{tr});

Algorithm 2: Augmentation SHAP

input : $X_{tr}, Y_{tr}, X_{query}, b, K$
for $\boldsymbol{x} \in X_{query}$ **do**
  if $\boldsymbol{x}$ *is unlabelled* **then**
    get $\boldsymbol{y}$ from expert;
  get R^* from expert;
  repeat K **times**
    $p \leftarrow Permutation(\{1, 2, \dots M\})$;
    $\boldsymbol{x}_{aug} \leftarrow \boldsymbol{x}$;
    $\boldsymbol{x}_{aug,S_i^p} \leftarrow b_{S_i^p}$;
    $\boldsymbol{y}_{aug} \leftarrow \boldsymbol{y} - \sum_{j \in S_i^p} R_j^*$;
    append $\boldsymbol{x}_{aug}, \boldsymbol{y}_{aug}$ to (X_{tr}, Y_{tr});
    append $\boldsymbol{x}, \boldsymbol{y}$ to (X_{tr}, Y_{tr});
retrain model on (X_{tr}, Y_{tr});

case for CAIPI), we again oversample $\boldsymbol{x}$ to make sure the amount of augmented samples is the same.

As is customary in active learning, we simulate the domain expert by querying a completely labelled dataset. Correct explanation are also simulated in related work [1, 10]. Our first set of experiments continues this practice, testing our first hypotheses that **correct feature attributions can be used to improve models**. In a final experiment the authors' (basic) domain knowledge is encoded into approximately correct explanations, shedding light on our second hypotheses: **domain experts can supply feature attributions that are sufficiently correct**.

All experiments were run on a server with 24 CPU nodes within an hour. Code for our implementations and all experiments is available on request. We enjoyed using the Python packages `scikit-learn` [17] and `shap` [4] (solely KernelSHAP is used).

5.1 Correct feature attributions can be used to improve models

These experiments take the expert out of the equation using a known data generating model. First, to confirm our theoretical validation, a known linear model with five uniformly generated coefficients $w_i^* \in [-1, 1]$ is considered which

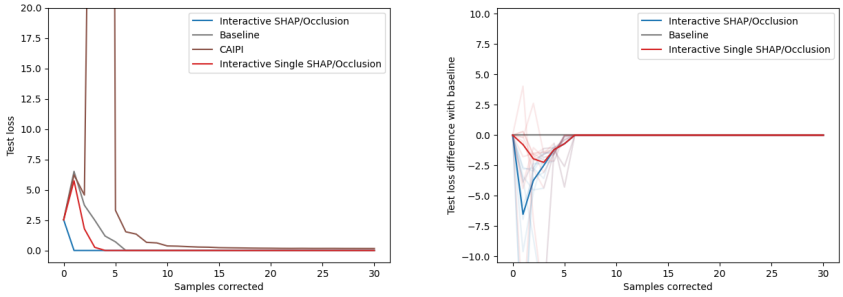

(a) (b)

Figure 2: Comparison of test loss evolution for different active learning strategies. The experiment was repeated for five random linear models and five randomly shuffled datasets per model. (a) shows the average test loss for our strategies (*Interactive SHAP/Occlusion* and *Interactive Single SHAP/Occlusion*), the current state-of-the-art [1] (*CAIPI*) and the standard approach without explanations (*Baseline*). To determine the significance Figure (b) shows the average of the per run difference in test loss of our approaches w.r.t. the baseline. The faded lines are the actual differences.

is used to generate noiseless targets for normal distributed Gaussian data:

$$\begin{aligned} x_i &\sim \mathcal{N}(0,1) \text{ for } i = 1 \dots 5 \\ y &= \sum_i w_i^* x_i + w_0^*. \end{aligned} \tag{11}$$

An untrained linear model with all coefficients zero is then iteratively corrected by considering one queried sample per iteration. The evolution of the test loss on 100 hold-out samples for different strategies is shown in Figure 2. As already mentioned in the background SHAP and Occlusion attributions are equivalent for linear models. Algorithm 1 is used to augment every sample. To provide a comparison with the state-of-the art [1], the least important feature u of the queried sample is considered irrelevant, i.e. only $R_u^* = 0$ in Algorithm 1. This can be directly compared to only setting R_u^* to its correct SHAP/Occlusion attribution (denoted by *Interactive Single SHAP/Occlusion*).

A similar experiment is repeated for a logistic regression model. Again, samples are generated according to Equation (11) but this time $y = H(\sum_i w_i^* x_i + w_0^*)$ with H the Heaviside function. In this case, initially the model is trained on two random samples of different classes (since the training dataset has to have two different classes) whereafter each iteration a new sample is added and corrected. In this setting CAIPI, which considers all features with Occlusion attribution $R_k = 0$ to be irrelevant, should be directly compared to our methods. The evolution of the test accuracy on 100 hold-out samples is shown in Figure 3.

To show the generality of our method, we next consider a more complex

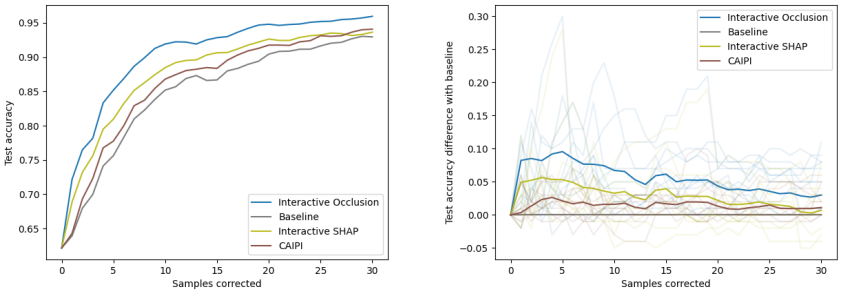

(a) (b)

Figure 3: Comparison of test accuracy evolution for different active learning strategies. The experiment was repeated for five random logistic regression models and five randomly shuffled datasets per model. (a) shows the average test accuracy for our strategies (*Interactive Occlusion* and *Interactive SHAP*), the state-of-the-art *CAIPI* and the standard approach without explanations (*Baseline*) . To determine the significance, Figure (b) shows the average of the per run difference in test accuracy of all approaches w.r.t. the baseline. The faded lines are the actual differences.

dataset, the Boston Housing dataset [18]. In this case, the real data generating model is unknown. To accurately provide correct attributions we train a boosted forest with 10 trees on the complete dataset and use the predicted house prices as the labels, and the correct explanations to iteratively improve an untrained boosted forest (with the same hyperparameters). The evolution of test loss on 253 hold out samples (half of the dataset) is shown in Figure 4. This time we consider a random query of 5 new samples each iteration. As in the experiment with linear models, we also include a *Interactive Single SHAP/Occlusion* as an extra comparison with the state-of-the-art.

The Appendix holds additional experiments with different models.

5.2 Domain experts can supply feature attributions that are sufficiently correct

To motivate the practical use of our proposed methods, we show how expert knowledge about the data can lead to an increased performance through interactive feature attributions. The task is to predict the survival of Titanic passengers based on class, sex, age and number of siblings (*SibSp*). To determine an upper bound, the dataset is again simulated as in the above experiment and explanations from the correct model (a boosted forest with 10 trees) are used to improve an untrained boosted forest (same hyperparameters as correct model, trained on two random samples of different classes). While the authors are no experts on the survival of passengers aboard the Titanic, they are familiar with

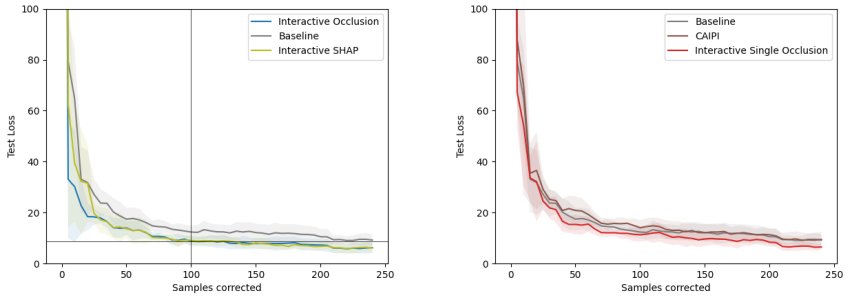

(a) (b)

Figure 4: Comparison of test loss evolution for different active learning strategies on the Boston Housing dataset. The experiment was repeated for five randomly shuffled datasets. (a) shows the average test loss for our strategies (*Interactive Occlusion* and *Interactive SHAP*) and the standard approach without explanations (*Baseline*). (b) shows the average test loss for the current state-of-the-art [1] *CAIPI* and *Interactive Single SHAP/Occlusion*.

the phrase 'women and children first'. This piece of knowledge can be encoded in Occlusion feature attributions: with regard to the known background sample $(class = 2, sex = male, age = 28, SibSp = 0)$, a procedure to set the correct attributions is shown in Figure 5a. For the state-of-the-art approach CAIPI [1] the irrelevant features get an attribution of zero. Figure 5b shows the results when correcting a random query of 10 samples each iteration.

6 Discussion

6.1 Correct feature attributions can be used to improve models

The results with linear models (Figure 2) show that our methods significantly outperform the state-of-the-art [1] and standard active learning on sample efficiency. Also, correcting one complete explanation is indeed sufficient to fully correct a linear model (*Interactive SHAP/Occlusion* goes to zero after one sample). CAIPI diverges outside the figure boundary, most likely because every iteration, theoretically the coefficient associated with the irrelevant feature w_u gets pushed to zero. This is in stark contrast to setting R_u to its correct SHAP/Occlusion attribution: *Interactive Single SHAP/Occlusion* still converges rather rapidly to zero loss.

With logistic regression we see very similar results, our methods outperform the state-of-the art significantly. Algorithmically speaking, the only difference between *Interactive Occlusion* and *CAIPI* is that our approach also enforces nonzero feature attributions (in the classification setting either -1 or 1). At

input : $\boldsymbol{x}, \boldsymbol{y}$
if $\boldsymbol{y} = 1$ **then**
  if $\boldsymbol{x}_{sex} = female$ **then**
    if *Occlusion* **then**
      $R^*_{sex} \leftarrow 1$;
    if *CAIPI* **then**
      $R^*_{/sex} \leftarrow 0$;
  else if $\boldsymbol{x}_{age} < 12$ **then**
    if *Occlusion* **then**
      $R^*_{age} \leftarrow 1$;
    if *CAIPI* **then**
      $R^*_{/age} \leftarrow 0$;

(a)

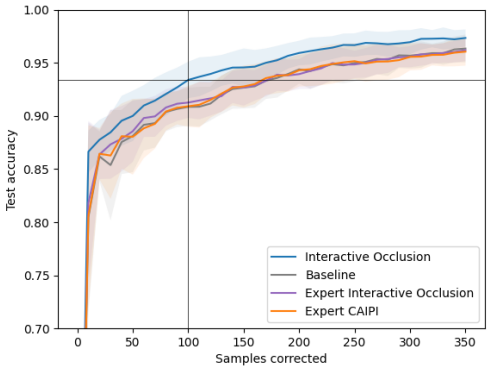

(b)

Figure 5: Figure (a) holds a procedure to encode Titanic domain knowledge in feature attributions. Figure (b) holds the comparison of test accuracy evolution for different active learning strategies on the Titanic dataset. The experiment was repeated for twenty randomly shuffled datasets. The figure shows the average test accuracy for entirely correct explanations (*Interactive Occlusion*), explanations generated by experts (*Expert Interactive Occlusion, Expert CAIPI*) and standard active learning.

sample level this results in extra samples containing feature values that change the class, which can play a big role in training classification models efficiently. This also motivates exploring interactive counterfactual explanations ([19] is a review). Making the same mistake as CAIPI, [10] also only enforces zero input gradients.

On the Boston Housing data, our methods also do better (increasing sample efficiency by more than 100%) than CAIPI, which does worse than the baseline. However, the method does not diverge such as was the case for linear models. This is probably because of the increased complexity of the boosted tree model, meaning extra local data has less effect on the global model. A linear model is fully defined by its explanation (see our theoretical validation). Moreover the approach with SHAP values, while still being significantly better than standard active learning, is worse than using Occlusion attributions (we see the same with logistic regression). While it is argued that the generality and intuitive properties (such as local accuracy) of the SHAP value increase its understandability [4], our proposed interactive implementation does not outperform the much more simple Occlusion attribution. This was partly expected because the added complexity of SHAP results in extra constraints to generate the local augmented data, which means the extra samples are not necessarily samples from the correct model. Therefore in the next experiment we will only use Occlusion attributions.

Note that all of the active learning results can be easily interpreted as regular supervised learning results if they are read in the vertical direction for different fixed number of corrected samples. Doing so shows that adding augmented data (according to our implementations) to particular fixed sets of samples always increased the performance of the model.

6.2 Domain experts can supply feature attributions that are sufficiently correct

The experiment on the Titanic dataset shows that general knowledge can indeed be encoded in Occlusion feature attributions leading to an increase in performance. Our approach outperforms the state-of-the-art [1]. It seems more intuitive to provide correct attributions based on the domain knowledge of the Titanic dataset, than to denote the irrelevant features. Enforcing that features are completely irrelevant is too drastic: e.g. although sex might be the most relevant feature, other features such as age can also be relevant, but to a lesser degree.

7 Limitations

Most of our experiments use exactly correct attributions (also in [1]). In reality, domain experts will make mistakes in correcting attributions, still our last experiment in Subsection 5.2 shows that even rudimentary domain knowledge is beneficial. While our approach significantly increases sample efficiency, it is unclear if is time efficient: correcting attributions might take more time than simply correcting extra samples. However, usually the expert has a reason to

correct a target, if this reason can be efficiently encoded in attributions, it should only add a little extra time.

8 Conclusion and future work

This paper introduced the first interactive feature attributions which allow for complete correction. We theoretically discuss our approach and provide connections with related work. Implementations for two popular attributions (Occlusion and SHAP) significantly improve the sample efficiency and performance of several models in regression and classification tasks compared to standard active/supervised learning and the state-of-the-art. Additionally we provide an experiment that showcases how actual domain knowledge can increase the performance of a trained model through our interactive attributions. In further work we plan to involve a real domain expert to further motivate our approach. Overall, our contribution is widely applicable and could significantly improve model performance when data is scarce and increase the sample efficiency of active learning tasks.

As a rare testable practical use case for feature attributions, this work also opens up the discussion on comparing different feature attributions and their human interpretation. Existing work mostly focuses on theory/intuition ([20] is a recent interesting effort), benchmarks ([21] provides an overview of different metrics in the image domain). In reality, explanations are to be used by experts and their evaluation should take this into account, e.g. [22] examines how different explanation can influence decision accuracy and time. Another evaluation could be through a task involving interactive explanations. Our results already hint that the current state-of-the-art in local explanations, SHAP values, might perform worse than basic Occlusion attributions. The added complexity of advanced explanation methods might not be preferred in interactive settings. Of course, different interactive implementations for SHAP values exist. Further work could consist on joint optimization of an explanation loss (e.g. distance to the correct explanation) and the regular classification or regression loss, similar to [10]. Additionally, it would be very interesting to further explore how domain experts can encode their knowledge into different types of explanation. In that case, the intuitive properties of SHAP values such as local accuracy might be useful.

Acknowledgements

This project has received funding from the Flemish Government (AI Research Program) and from the FWO ('Artificial Intelligence (AI) for data-driven personalised medicine', G0C9623N and 'Deep, personalized epileptic seizure detection', G0D8321N) and Leuven.AI Institute.

References

[1] Stefano Teso and Kristian Kersting. Explanatory interactive machine learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 239–245, 2019.

[2] Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger, Franziska Herbert, Xiaoting Shao, Hans-Georg Luigs, Anne-Katrin Mahlein, and Kristian Kersting. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nature Machine Intelligence*, 2(8):476–486, 2020.

[3] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I 13*, pages 818–833. Springer, 2014.

[4] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

[5] Burr Settles. Active learning literature survey. 2009.

[6] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR, 2017.

[7] Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. From local explanations to global understanding with explainable ai for trees. *Nature machine intelligence*, 2(1):56–67, 2020.

[8] Mukund Sundararajan and Amir Najmi. The many shapley values for model explanation. In *International conference on machine learning*, pages 9269–9278. PMLR, 2020.

[9] 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*, pages 1135–1144, 2016.

[10] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.

[11] Omar Zaidan, Jason Eisner, and Christine Piatko. Using "annotator rationales" to improve machine learning for text categorization. In *Human language technologies 2007: The conference of the North American chapter of the association for computational linguistics; proceedings of the main conference*, pages 260–267, 2007.

[12] Kevin Small, Byron C Wallace, Carla E Brodley, and Thomas A Trikalinos. The constrained weight space svm: learning with ranked features. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 865–872, 2011.

[13] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th international conference on intelligent user interfaces*, pages 126–137, 2015.

[14] Stefano Teso, Öznur Alkan, Wolfang Stammer, and Elizabeth Daly. Leveraging explanations in interactive machine learning: An overview. *arXiv preprint arXiv:2207.14526*, 2022.

[15] Felix Friedrich, Wolfgang Stammer, Patrick Schramowski, and Kristian Kersting. A typology for exploring the mitigation of shortcut behaviour. *Nature Machine Intelligence*, pages 1–12, 2023.

[16] Scott M Lundberg, Bala Nair, Monica S Vavilala, Mayumi Horibe, Michael J Eisses, Trevor Adams, David E Liston, Daniel King-Wai Low, Shu-Fang Newman, Jerry Kim, et al. Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature biomedical engineering*, 2(10):749–760, 2018.

[17] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[18] David Harrison Jr and Daniel L Rubinfeld. Hedonic housing prices and the demand for clean air. *Journal of environmental economics and management*, 5(1):81–102, 1978.

[19] Sahil Verma, Varich Boonsanong, Minh Hoang, Keegan E Hines, John P Dickerson, and Chirag Shah. Counterfactual explanations and algorithmic recourses for machine learning: a review. *arXiv preprint arXiv:2010.10596*, 2020.

[20] Ian C Covert, Scott Lundberg, and Su-In Lee. Explaining by removing: A unified framework for model explanation. *The Journal of Machine Learning Research*, 22(1):9477–9566, 2021.

[21] Richard Tomsett, Dan Harborne, Supriyo Chakraborty, Prudhvi Gurram, and Alun Preece. Sanity checks for saliency metrics. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 6021–6029, 2020.

[22] Sérgio Jesus, Catarina Belém, Vladimir Balayan, João Bento, Pedro Saleiro, Pedro Bizarro, and João Gama. How can i choose an explainer? an application-grounded evaluation of post-hoc explanations. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 805–815, 2021.

A Derivation SHAP implementation

As mentioned above the correct SHAP values R_i^* for sample $\boldsymbol{x}$ and corrected target $\boldsymbol{y}$ impose the following constraints on model f:

$$\begin{cases} \forall i \in F : \frac{1}{|P|} \sum_p f\left(\left[\boldsymbol{x}_{S_i^p}, b_{\bar{S}_i^p}\right]\right) - f\left([\boldsymbol{x}_{S^p}, b_{\bar{S}^p}]\right) = R_i^* \\ f(\boldsymbol{x}) = \boldsymbol{y}, \end{cases} \tag{12}$$

with $F = \{1 \dots M\}$ the set of all features. From this set of constraints we like to distil extra local data (i.e. extra function values) that can be added to the training data to enforce the correct SHAP values on the model. However, this set of equations is underdetermined: there are $|S|$ unknowns, $f\left([\boldsymbol{x}_S, b_{\bar{S}}]\right)\ \forall S \subseteq F$, but only $M+1$ constraints. Note that $|S| = 2^M > M+1$ for $M \geq 2$. A sensible choice would be enforcing the same attribution for every subset. This results in the following constraints:

$$\begin{cases} \forall i \in F, S \subset F : f\left(\left[\boldsymbol{x}_{S_i}, b_{\bar{S}_i}\right]\right) - f\left([\boldsymbol{x}_S, b_{\bar{S}}]\right) = R_i^* \\ f(\boldsymbol{x}) = \boldsymbol{y}. \end{cases} \tag{13}$$

This set of equations is solvable. To determine $f\left([\boldsymbol{x}_S, b_{\bar{S}}]\right)$ we can use the first constraint of Equation (13) to arrive at:

$$f\left([\boldsymbol{x}_S, b_{\bar{S}}]\right) = f(\boldsymbol{x}) - \sum_{j \in \bar{S}} R_j^*. \tag{14}$$

Enforcing the second constraint of Equation (13), we get:

$$f\left([\boldsymbol{x}_S, b_{\bar{S}}]\right) = \boldsymbol{y} - \sum_{j \in \bar{S}} R_j^*. \tag{15}$$

Thus enforcing correct SHAP values R_i^* results in samples of the form $([\boldsymbol{x}_S, b_{\bar{S}}], \boldsymbol{y} - \sum_{j \in \bar{S}} R_j^*)$. Since there are 2^M possible subsets, we will not use *all* training samples of this form. Instead, equivalent to approximating SHAP values [4], a number of feature orderings p will be sampled and the corresponding samples $([\boldsymbol{x}_{S_i^p}, b_{\bar{S}_i^p}], \boldsymbol{y} - \sum_{j \in \bar{S}_i^p} R_j^*)\ \forall i$ will be added to the training dataset.

B Additional experiments

To further motivate the generality of our methods, we conduct extra experiments on two other model types.

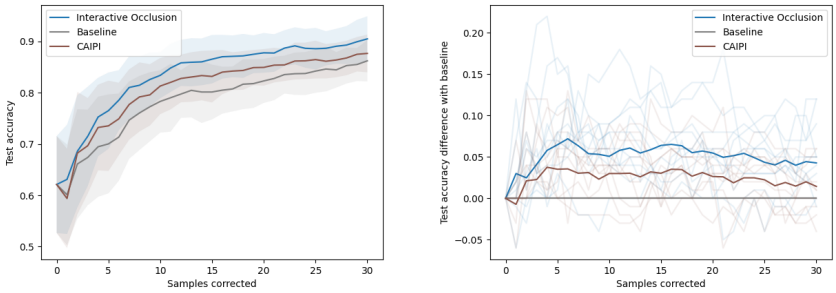

(a) (b)

Figure 6: Comparison of test accuracy evolution of a support vector machine for different active learning strategies. The experiment was repeated for five random data-generating logistic regression models and five randomly shuffled datasets per model. (a) shows the average test accuracy for our strategy (*Interactive Occlusion*), the state-of-the-art *CAIPI* and the standard approach without explanations (*Baseline*) . To determine the significance, Figure (b) shows the average of the per run difference in test accuracy of all approaches w.r.t. the baseline. The faded lines are the actual differences.

B.1 Experiment support vector machines

We use a support vector machine with RBF kernel on the data generated by random logistic regression models (the experiment shown in Figure 3) and achieve the results in Figure 6. Our approach outperforms the state-the-art [1] and the baseline.

B.2 Experiment multilayer perceptron

Lastly we train a multilayer perceptron with 9 hidden nodes (using stochastic gradient descent) iteratively on the Boston housing dataset [18] (similar experiment to Figure 4). In Figure 7 and Figure 8 the data generator is a multilayer perception with the same hyperparameters. In Figure 9, it is the same model as in the experiment of our main paper, namely a boosted forest with 10 trees. In Figure 7 and Figure 9 the multilayer perceptron is trained for 100 epochs. In Figure 8 it is trained for 200 epochs. We see that in all cases our approach outperforms the state-of-the-art and the baseline. Comparing Figure 7 and 8, it is seems that our approach can also prevent overfitting.

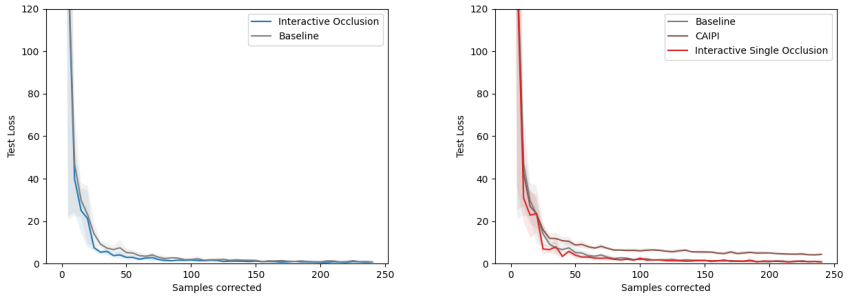

(a) (b)

Figure 7: Comparison of test loss evolution for different active learning strategies on the Boston Housing dataset. The experiment was repeated for five randomly shuffled datasets. The data-generating model and training model are both multilayer perceptrons. They are trained for 100 epochs. (a) shows the average test loss for our strategy (*Interactive Occlusion*) and the standard approach without explanations (*Baseline*). (b) shows the average test loss for the current state-of-the-art [1] *CAIPI* and *Interactive Single Occlusion*.

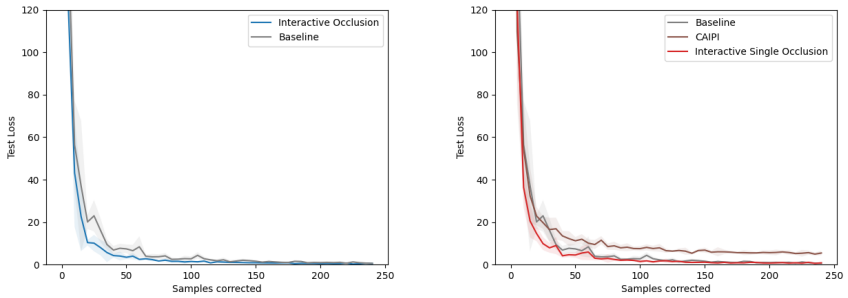

(a) (b)

Figure 8: Comparison of test loss evolution for different active learning strategies on the Boston Housing dataset. The experiment was repeated for five randomly shuffled datasets. The data-generating model and training model are both multilayer perceptrons. They are trained for 200 epochs. (a) shows the average test loss for our strategy (*Interactive Occlusion*) and the standard approach without explanations (*Baseline*). (b) shows the average test loss for the current state-of-the-art [1] *CAIPI* and *Interactive Single Occlusion*.

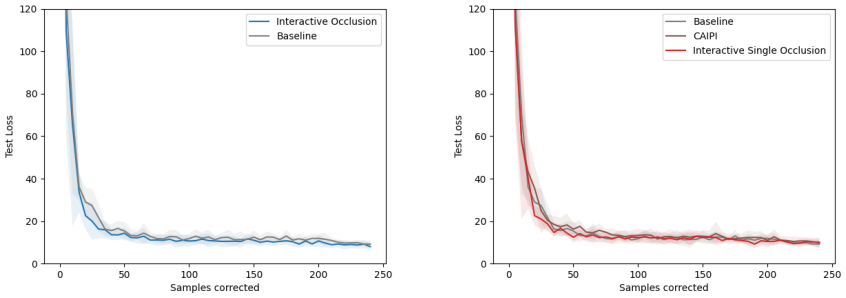

(a) (b)

Figure 9: Comparison of test loss evolution for different active learning strategies on the Boston Housing dataset. The experiment was repeated for five randomly shuffled datasets. The data-generating model is a boosted forest with 10 trees. The training model is a multilayer perceptron trained for 100 epochs. (a) shows the average test loss for our strategy (*Interactive Occlusion*) and the standard approach without explanations (*Baseline*). (b) shows the average test loss for the current state-of-the-art [1] *CAIPI* and *Interactive Single Occlusion*.