
REPID: Regional Effect Plots with implicit Interaction Detection

Julia Herbinger
LMU Munich

Bernd Bischl
LMU Munich

Giuseppe Casalicchio
LMU Munich

Abstract

Machine learning models can automatically learn complex relationships, such as non-linear and interaction effects. Interpretable machine learning methods such as partial dependence plots visualize marginal feature effects but may lead to misleading interpretations when feature interactions are present. Hence, employing additional methods that can detect and measure the strength of interactions is paramount to better understand the inner workings of machine learning models. We demonstrate several drawbacks of existing global interaction detection approaches, characterize them theoretically, and evaluate them empirically. Furthermore, we introduce *regional effect plots with implicit interaction detection*, a novel framework to detect interactions between a feature of interest and other features. The framework also quantifies the strength of interactions and provides interpretable and distinct regions in which feature effects can be interpreted more reliably, as they are less confounded by interactions. We prove the theoretical eligibility of our method and show its applicability on various simulation and real-world examples.

1 INTRODUCTION

Many machine learning (ML) models are considered black-boxes, as they do not provide insights into how the model’s prediction function is composed and which features or interactions¹ are actually used by

¹Interactions describe to what extent a feature’s effect on the model prediction is influenced by other features.

the model. This lack of transparency has been partially addressed by recent developments in the field of interpretable ML. In general, the literature distinguishes between local and global interpretation methods (Molnar et al., 2020). Global interpretation methods aim at explaining the overall behavior of an ML model. Examples include the partial dependence (PD) plot (Friedman, 2001), which visualizes the effect of a feature on the model’s prediction, and the permutation feature importance, which quantifies the relevance of features (Fisher et al., 2019). However, many of these global interpretation methods are confounded by feature interactions, meaning that they can produce misleading explanations when feature interactions are present because they often aggregate over individual effects of local interpretation methods and thereby obfuscate heterogeneous effects induced by feature interactions (Molnar et al., 2021b). This so-called *aggregation bias* (Mehrabi et al., 2021) is responsible for producing global explanations that are usually not representative or not valid for many individuals. Instead of explaining the ML model on a global level, local interpretation methods – such as individual conditional expectation (ICE) curves (Goldstein et al., 2015), LIME (Ribeiro et al., 2016), or Shapley values (Strumbelj and Kononenko, 2014) – can be used to understand how a feature influences an individual prediction. However, many local interpretation methods do not provide a global understanding of the ML model due to their local view (i.e., their explanations only refer to individual observations). Thus, it is often recommended to consider both local and global interpretation methods. For example, in the case of PD plots, looking additionally at ICE curves (Goldstein et al., 2015) can help to reveal interactions when the ICE curves are heterogeneous (see Figure 1). Yet, ICE curves are not able to quantify the strength of the underlying feature interactions, nor can they tell exactly which features interact with each other. On the other hand, other methods that quantify the interaction strength between features are available. However, they do not provide any visual component of how these interactions influence the effect of a feature of interest (Friedman et al., 2008; Greenwell et al., 2018). The

work in this paper is motivated by subgroup analysis (Su et al., 2009) as a trade-off between local and global explanations. We aim to uncover a possible *aggregation bias* in the PD plot by finding interpretable subgroups in the data with differing influences of a feature on the predictions. Hence, for well-performing ML models, this might also reveal a possible bias in the data (e.g., when the influence of a feature on the prediction strongly differs for certain subgroups, although it should not) and thus might be helpful to uncover possible negative societal impacts.

Contributions: We introduce *regional effect plots with implicit interaction detection* (REPID), a model-agnostic interpretation method that produces regional effect plots (REPs) in which feature effects are less confounded by interactions. Regions are obtained by a decision tree and thus represent interpretable and distinct subgroups in the feature space. We also propose a new measure to detect and quantify interactions with a feature of interest, which can be used to rank interactions according to their strength. To receive a broader and more competitive comparability, we derive another global interaction index based on SHAP interaction values (Lundberg et al., 2018). We mathematically prove the theoretical meaningfulness of our method and demonstrate its advantages compared to not only the well-known H-statistic (Friedman et al., 2008), but also to Greenwell’s interaction index (Greenwell et al., 2018) and our derived global SHAP interaction index. Finally, we demonstrate the usefulness of our method on real-world data.

Open Science: The implementation of the proposed method and the fully reproducible code for all experiments are provided in a public repository².

2 BACKGROUND AND RELATED WORK

Notation: Consider a p -dimensional feature space $\mathcal{X} \in \mathbb{R}^p$ and a target space \mathcal{Y} (e.g., $\mathcal{Y} = \mathbb{R}$ for regression). The corresponding random variables are $X = (X_1, \dots, X_p)$ for the features and Y for the target. ML algorithms learn a prediction model \hat{f} using training data $\mathcal{D} = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^n$ sampled i.i.d. from the unknown joint distribution $\mathbb{P}_{X,Y}$. In our notation, the i -th observation is denoted by $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_p^{(i)})^T$, and $\mathbf{x}_j = (x_j^{(1)}, \dots, x_j^{(n)})^T$ denotes the realizations of the j -th feature X_j .

PD Plot (Friedman, 2001): The marginal relationship of features on model predictions can be visual-

ized by PD plots. Consider a set of feature indices $S \subseteq \{1, \dots, p\}$ and its complement $C = S^c$. Each observation $\mathbf{x}^{(i)}$ can be partitioned into $\mathbf{x}_S^{(i)}$ and $\mathbf{x}_C^{(i)}$ containing only features indexed by S and C , respectively. X_S and X_C refer to the corresponding random variables. The PD function of features indexed by S marginalizes over features in C and is defined as $f_S^{PD}(\mathbf{x}_S) = E_{X_C}[\hat{f}(\mathbf{x}_S, X_C)]$. The PD function is estimated by Monte-Carlo integration:

$$\hat{f}_S^{PD}(\mathbf{x}_S) = \frac{1}{n} \sum_{i=1}^n \hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)}). \quad (1)$$

Here, $\hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)})$ can be read as the prediction of the i -th observation where features in S were replaced by \mathbf{x}_S . Plotting the pairs $\{(\mathbf{x}_S^{(k)}, \hat{f}_S(\mathbf{x}_S^{(k)}))\}_{k=1}^m$ using grid points³ denoted by $\mathbf{x}_S^{(1)}, \dots, \mathbf{x}_S^{(m)}$ yields a PD curve. The mean-centered PD function can be estimated by

$$\hat{f}_S^{PD,c}(\mathbf{x}_S) = \hat{f}_S^{PD}(\mathbf{x}_S) - \frac{1}{m} \sum_{k=1}^m \hat{f}_S^{PD}(\mathbf{x}_S^{(k)}).$$

If $|S| = 2$, we get a 2-dimensional PD plot showing the joint marginal effect of the 2 features included in S .

ICE Plot (Goldstein et al., 2015): The averaging in Eq. (1) can obfuscate complex relationships resulting from feature interactions. ICE plots address this problem by directly visualizing individual curves for each observation, i.e., $\{(\mathbf{x}_S^{(k)}, \hat{f}_S(\mathbf{x}_S^{(k)}, \mathbf{x}_C^{(i)}))\}_{k=1}^m$ for all $i \in \{1, \dots, n\}$. ICE curves will usually have different shapes if interactions with other features in C are present. To facilitate the visual identification of heterogeneous ICE curves and, consequently, the presence of interactions, the authors propose the derivative-ICE (d-ICE) plot. Assuming that there are no interactions between features \mathbf{x}_S and \mathbf{x}_C , the prediction function can be written as $\hat{f}(\mathbf{x}) = \hat{f}(\mathbf{x}_S, \mathbf{x}_C) = g(\mathbf{x}_S) + h(\mathbf{x}_C)$. Hence, the partial derivatives of all ICE curves $\frac{\delta \hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)})}{\delta \mathbf{x}_S} = g'(\mathbf{x}_S)$ do not depend on $\mathbf{x}_C^{(i)}$, which means that d-ICE curves have the same shape if there are no interactions. The d-ICE plot visualizes the partial derivatives of ICE curves along with their standard deviation to highlight regions in \mathbf{x}_S where the d-ICE curves are heterogeneous (see Figure 1).

Visual INTERaction Effects (VINE) (Britton, 2019): The principle of VINE is to cluster similar slopes of ICE curves to obtain clusters where the curves are less affected by interactions based on a three-step approach: (1) for a feature of interest, find clusters where the ICE curves of that feature have similar slopes using, e.g., agglomerative clustering, (2) for each found cluster, create a binary label containing the information of whether an observation belongs to the considered cluster or any other cluster and apply a tree

²<https://github.com/JuliaHerbinger/repid>

³Common choices are randomly selected feature values, quantiles, or equidistant values (Molnar et al., 2021b).

stump, (3) identify the split feature and its split point and merge clusters that use the same feature and a similar split point. Although VINE is based on a similar strategy as our approach, its three-step approach has several disadvantages (see Section 3.1.1). Approaches to group ICE curves to reduce feature dependencies instead of feature interactions is introduced in Molnar et al. (2021a) and Grömping (2020).

H-Statistic (Friedman et al., 2008): The H-Statistic is based on the assumption that if two features do not interact, the 2-dimensional mean-centered PD function of two features \mathbf{x}_j and \mathbf{x}_l is additively separable and can be decomposed into the sum of their mean-centered 1-dimensional PDs, i.e.,

$$f_S^{PD,c}(\mathbf{x}_S) = f_j^{PD,c}(\mathbf{x}_j) + f_l^{PD,c}(\mathbf{x}_l), \text{ with } S = \{j, l\}.$$

The stronger an interaction effect, the more the sum of $f_j^{PD,c}(\mathbf{x}_j)$ and $f_l^{PD,c}(\mathbf{x}_l)$ will deviate from $f_S^{PD,c}(\mathbf{x}_S)$. Hence, the H-statistic computes the interaction strength between two features \mathbf{x}_j and \mathbf{x}_l by quantifying the degree of this deviation using

$$\hat{H}_S^2 = \frac{\sum_{i=1}^n (\hat{f}_S^{PD,c}(\mathbf{x}_S^{(i)}) - \sum_{k \in S} \hat{f}_k^{PD,c}(\mathbf{x}_k^{(i)}))^2}{\sum_{i=1}^n (\hat{f}_S^{PD,c}(\mathbf{x}_S^{(i)}))^2}. \quad (2)$$

Greenwell's interaction index (Greenwell et al., 2018): The interaction strength between two features \mathbf{x}_j and \mathbf{x}_l is quantified based on the variability of the PD function of \mathbf{x}_j conditioned on a fixed value of \mathbf{x}_l (see Appendix A.3.1).

However, the H-Statistic and the Greenwell's interaction index only quantify interaction effects and do not visualize how interactions influence the marginal effect of a feature. Moreover, both methods are sensitive to varying main effects (see Section 3.2.1 and 4.1).

Functional ANOVA (fANOVA) (Hooker, 2004): The fANOVA decomposes the prediction function as follows:

$$\hat{f}(\mathbf{x}) = g_0 + \sum_{k=1}^p \sum_{W \subseteq \{1, \dots, p\}, |W|=k} g_W(\mathbf{x}_W) \quad (3)$$

where $E_X[g_W(\mathbf{x}_W)] = 0$ for all feature index sets W (zero-means property). While $g_W(X_W)$ with $|W| = 1$ refers to main (or *first-order*) effects, $g_W(X_W)$ with $|W| > 1$ refers to interactions (or *higher-order*) effects. Based on the decomposition in Eq. (3), the authors detect interactions of any order by applying an efficient search algorithm and visualize them in an interaction network graph. However, the network only shows the presence of feature interactions and does not quantify the interaction strength or illustrate how they influence the prediction. A discussion on the assumptions and application of the fANOVA decomposition in the context of this paper is provided in Appendix A.1.

SHAP interaction values (Lundberg et al., 2018): The method is based on Shapley values (Shapley, 1953) and Shapley interaction indices (Fujimoto et al., 2006) from game theory. In the ML context, SHAP interaction values of two features quantify the pure interaction effect after accounting for the individual feature effects. The SHAP interaction values separate the interaction effect from the main effects of two features indexed by j and l (for $j \neq l$) for an observation \mathbf{x} :

$$\Phi_{j,l}(\mathbf{x}) = \sum_{S \subseteq \{1, \dots, p\} \setminus \{j, l\}} \frac{|S|!(p-|S|-2)!}{2^{(p-1)}!} \nabla_{j,l}(\mathbf{x}_S),$$

where $\nabla_{j,l}(\mathbf{x}_S) = f_{S \cup \{j, l\}}^{PD}(\mathbf{x}_{S \cup \{j, l\}}) - f_{S \cup \{j\}}^{PD}(\mathbf{x}_{S \cup \{j\}}) - f_{S \cup \{l\}}^{PD}(\mathbf{x}_{S \cup \{l\}}) + f_S^{PD}(\mathbf{x}_S)$. The SHAP interaction values have only been introduced on an observational level, where the final plot over all observations shows the influence of the interaction effect on the prediction.

3 THE REPID METHOD

REPID visualizes regional marginal effects of a certain feature of interest \mathbf{x}_S with $|S| = 1$ depending on its interactions with other features and quantifies the underlying interaction strength. The following simulation example demonstrates the benefits of our method compared to existing ones. We draw $n = 500$ samples for 6 independent random variables, which are distributed as follows: $X_1, X_2 \sim \mathcal{U}(-1, 1)$, $X_3, X_5 \sim \mathcal{B}(n, 0.5)$, $X_4 \sim \mathcal{B}(n, 0.7)$ and $X_6 \sim \mathcal{N}(1, 5)$. The true relationship is described by

$$f(\mathbf{x}) = 0.2\mathbf{x}_1 - 8\mathbf{x}_2 + 8\mathbf{x}_2 \mathbb{1}_{(\mathbf{x}_1 > 0)} + 16\mathbf{x}_2 \mathbb{1}_{(\mathbf{x}_3 = 0)} + \epsilon \quad (4)$$

with $\epsilon \sim \mathcal{N}(0, 1)$. We fit a random forest (RF) with 500 trees on the data. Due to the linear relationship, we can assume that the interaction strength between \mathbf{x}_2 and \mathbf{x}_3 is higher than the one between \mathbf{x}_2 and \mathbf{x}_1 .

3.1 Regional Effect Plots

3.1.1 Motivation

PD plots are often shown together with their underlying ICE curves (see Figure 1). The heterogeneous shapes of ICE curves imply the presence of feature interactions. Although ICE or d-ICE plots indicate interactions, they do not provide any information on which other features are responsible for these interactions and how the underlying interaction influences the marginal effect of \mathbf{x}_S (see Figure 1). Grouping homogeneous ICE curves will reduce the presence of individual interaction effects within a group. This leads to regional PD plots that actually reflect the pure marginal effect of \mathbf{x}_S within this group. VINE (Britton, 2019) implements this idea by clustering ICE curves with similar slopes (see Section 2). However, VINE is only




Figure 1: Left: ICE curves (black) and PD plot (blue) for x_2 . Right: Smoothed d-ICE curves (upper plot) and standard deviation of d-ICE curves (lower plot).

a visual tool and does not quantify or rank feature interactions. Furthermore, VINE is an unsupervised approach, and its solution depends on the number of clusters k that must be chosen (which is not trivial). Another drawback is that VINE “finds” feature interactions in an inconvenient second step by fitting a separate tree stump for each cluster (see Section 2). Due to the different tree stumps used in VINE, the derived decision rules are often not distinct and therefore difficult to interpret. In a third step, VINE introduces a post-hoc merging of clusters based on similar decision rules. In Figure 2, we show that this three-step approach does not always lead to meaningful groupings. While in the left plot, the ICE curves are divided meaningfully into 2 clusters based on the most interacting feature x_3 (according to Eq. (4)), the clusters in the right plot do not divide the ICE curves into visually meaningful groups with homogeneous ICE curves.




Figure 2: ICE and regional PD (dashed) plot of x_2 clustered by VINE for $k = 2$ (left) and $k = 5$ (right). The 5 clusters are reduced to 3 by post-hoc merging. Cluster numbers 0 and 3 still contain differing individual interaction effects, which are averaged and hence not represented well by the regional PD plot.

3.1.2 Methodology

Here, we derive a new tree-based approach to determine optimal REPs for any feature of interest x_S . REPs are regional PD plots that aggregate ICE curves

within automatically identified regions where feature effects are less confounded by interactions. Our aim is to recursively split the entire data referred to by index set $\mathcal{N} = \{1, \dots, n\}$ into interpretable regions to obtain more homogeneous ICE curves for x_S within the split regions denoted by \mathcal{N}_g (where $g \in \{1, \dots, G\}$ indexes a certain node of the tree and G is the number of all tree nodes). Hence, we want to split \mathcal{N} in such a way that ICE curves within the obtained regions have a similar shape, meaning that the distance of these ICE curves to the REP estimate (i.e., $\hat{f}_{S|\mathcal{N}_g}^{PD}(x_S) := \frac{1}{|\mathcal{N}_g|} \sum_{i \in \mathcal{N}_g} \hat{f}(x_S, x_C^{(i)})$) is small. To that end, we propose a tree-based partitioning in Algorithm 1, which refers only to a single binary split and is inspired by the CART algorithm (Breiman et al., 1984)⁴. The splitting is recursively repeated until the split criterion (denoted by $\mathcal{I}(\hat{t}, \hat{j})$ in Algorithm 1) does not improve anymore compared to the previous split or until a pre-specified stop criterion is met. The split criterion is based on a suitable risk function \mathcal{R} that operates on ICE curves (see also Eq. (6)).

Algorithm 1: Tree-based Partitioning

input: index set \mathcal{N} , risk \mathcal{R}_{L2} (see, e.g., Eq. (6))

output: child nodes $\mathcal{N}_l^{\hat{t}, \hat{j}}$ and $\mathcal{N}_r^{\hat{t}, \hat{j}}$

for each feature indexed by $j \in C$ **do**

for every split t on feature x_j **do**

$$\mathcal{N}_l^{t,j} = \{i \in \mathcal{N}\}_{x_j^{(i)} \leq t}; \mathcal{N}_r^{t,j} = \{i \in \mathcal{N}\}_{x_j^{(i)} > t}$$

$$\mathcal{I}(t, j) = \mathcal{R}_{L2}(\mathcal{N}_l^{t,j}) + \mathcal{R}_{L2}(\mathcal{N}_r^{t,j})$$

end for

end for

Choose $\hat{t}, \hat{j} \in \arg \min_{t,j} \mathcal{I}(t, j)$

We first estimate the mean-centered ICE curves by $\hat{f}^c(x_S, x_C^{(i)}) = \hat{f}(x_S, x_C^{(i)}) - \frac{1}{m} \sum_{k=1}^m \hat{f}(x_S^{(k)}, x_C^{(i)})$. Since we want to minimize the shape differences between ICE curves in the regions, we then define the risk function \mathcal{R}_{L2} in Eq. (6)⁵ such that the variance (L2 loss) of the mean-centered ICE curves is minimized. This can be estimated by calculating the L2 loss of the mean-centered ICE curves at each grid point (see Eq. (5)) and aggregating it over all grid points:

$$\mathcal{L}(\mathcal{N}_g, x_S) = \sum_{i \in \mathcal{N}_g} \left(\hat{f}^c(x_S, x_C^{(i)}) - \hat{f}_{S|\mathcal{N}_g}^{PD,c}(x_S) \right)^2 \quad (5)$$

$$\mathcal{R}_{L2}(\mathcal{N}_g) = \sum_{k=1}^m \mathcal{L}(\mathcal{N}_g, x_S^{(k)}) \quad (6)$$

⁴Algorithm 1 is defined for numerical features. For categorical features, we use an exhaustive search as seen in CART. The computational feasibility of this procedure depends on the number of categories.

⁵Multiplying with $\frac{1}{m}$ to obtain the average loss can be neglected for optimization.

Theorem 1 If Eq. (3) holds, then $\hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)})$ with $|S| = 1$ can be decomposed into the mean-centered⁶ main effect of \mathbf{x}_S (i.e. $g_S^{cS}(\mathbf{x}_S)$) and the mean-centered interaction effect of \mathbf{x}_S with \mathbf{x}_C for the i -th observation (i.e., $g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)})$):

$$\hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)}) = g_S^{cS}(\mathbf{x}_S) + \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)}).$$

Corollary 1.1 If Eq. (3) holds, then $f_S^{PD,c}(\mathbf{x}_S) = E_{X_C}[\hat{f}^c(\mathbf{x}_S, X_C)]$ with $|S| = 1$ can be decomposed into

$$g_S^{cS}(\mathbf{x}_S) + \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} E_{X_C} [g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, X_{C_k})].$$

The proof can be found in Appendix A.1.1.

Based on Theorem 1 and Corollary 1.1 – where we show that the mean-centered ICE curves and PD function can be decomposed in first-order and higher-order terms which depend on \mathbf{x}_S – we can prove in Theorem 2, that our risk function of Eq. (6) only depends on the interaction effects between \mathbf{x}_S and features in \mathbf{x}_C . Hence, by minimizing this risk function, we minimize the individual interaction effects between the feature of interest and all other features. Thus, we minimize the shape differences between ICE curves in each region. Theorem 3 states that the theoretical minimum of our split criterion leads to the optimal solution we aim to achieve, meaning that for each final region, all ICE curves are best represented by the REP.

Theorem 2 The distance minimized by the risk function \mathcal{R}_{L2} of Eq. (6) only depends on the mean-centered interaction effects between \mathbf{x}_S with $|S| = 1$ and all features interacting with \mathbf{x}_S , i.e., for the i -th observation, the distance results in

$$\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)}) - E_{X_C} [g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, X_{C_k})].$$

The proof can be found in Appendix A.1.2.

Theorem 3 If $\mathcal{I}(t, j) = 0$, i.e., the theoretical minimum of the split criterion is reached for a split, then the ICE curves within each of the child nodes \mathcal{N}_l and \mathcal{N}_r are identical to the respective REP (e.g., $\hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)}) = \hat{f}_{S|\mathcal{N}_l}^{PD,c}(\mathbf{x}_S) \quad \forall i \in \mathcal{N}_l$).

Proof 3 Since $\mathcal{R}_{L2}(\mathcal{N}_g) \geq 0 \quad \forall g \in \{1, \dots, G\}$, $\mathcal{I}(t, j) = 0$ implies $\hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)}) = \hat{f}_{S|\mathcal{N}_g}^{PD,c}(\mathbf{x}_S)$,

⁶ $g_W^{cS}(X_W) = g_W(X_W) - E_{X_S}[g_W(X_W)]$ is the mean-centered counterpart of $g_W(X_W)$ of Eq. (3) regarding X_S .

$\forall i \in \mathcal{N}_g, \forall g \in \{l, r\}$.

Applying our method to the simulation example introduced at the beginning of Section 3.1 leads to the REPs shown in Figure 3 after two splits. The first binary split divides the ICE curves of \mathbf{x}_2 using feature \mathbf{x}_3 , which interacts most with \mathbf{x}_2 (according to Eq.(4)). Each of the 2 resulting regions is then split again into 2 groups by feature \mathbf{x}_1 , which also interacts with \mathbf{x}_2 . Hence, after the second split, we receive interpretable and distinct regions with REPs that represent each sub-population well.




Figure 3: ICE curves for \mathbf{x}_2 grouped by REPID (black) and REPs (blue).

3.2 Quantifying Interaction Strength

3.2.1 Motivation

Besides understanding how other features influence the marginal effect of \mathbf{x}_S , users might be interested in how strong these interactions are and how to rank these features regarding their interaction strength with \mathbf{x}_S . The H-Statistic defined in Section 2 is a global measure that quantifies the strength of interaction between two features. However, its values are influenced by the main effects of the two regarded features (see Theorem 4). Hence, the two-way interaction with the highest H-Statistic value is not necessarily the strongest interaction, which we demonstrate in Section 4.1.

Theorem 4 The variance of the 2-dimensional mean-centered PD plot of features \mathbf{x}_j and \mathbf{x}_l ($\text{Var}(f_S^{PD,c}(\mathbf{x}_S))$ with $S = \{j, l\}$) depends on the mean-centered main effects (i.e., $g_j^{cS}(\mathbf{x}_j)$ and $g_l^{cS}(\mathbf{x}_l)$) of the two features of interest \mathbf{x}_j and \mathbf{x}_l . Since $\text{Var}(f_S^{PD,c}(\mathbf{x}_S))$ is the denominator of the H-Statistic, which is estimated as in Eq. (2), the H-Statistic itself also depends on the main effects of features in S . The proof can be found in Appendix A.1.3.

The global interaction index proposed by Greenwell et al. (2018) suffers from the same problem that we illustrate in Section 4.1 (see also Appendix A.3.1). A third method of quantifying the two-way interaction strength between features is based on SHAP interaction values (see Section 2). To the best of our knowledge, SHAP interaction values have been only defined on an observational level. Similar to the global feature importance used in Lundberg et al. (2018) to rank features according to their global impact in their SHAP summary plots, we suggest summarizing the individual SHAP interaction values for two features \mathbf{x}_j and \mathbf{x}_l into a global SHAP interaction index by

$$I_{j,l}^{\text{rel}} = \frac{I_{j,l}}{\sum_{l \in \{1, \dots, p\} \setminus \{j\}} I_{j,l}} \text{ where } I_{j,l} = \sum_{i=1}^n |\Phi_{j,l}(\mathbf{x}^{(i)})|.$$

Since the absolute values $I_{j,l}$ are difficult to interpret, we prefer a relative version $I_{j,l}^{\text{rel}}$, which we call the SHAP interaction index and can be interpreted as the proportion of all two-way interactions with \mathbf{x}_j to which the l -th feature contributes. By definition, SHAP interaction values only contain the interaction effect between \mathbf{x}_j and \mathbf{x}_l . Hence, in contrast to the H-Statistic, varying main effects do not change the ranking of our proposed global SHAP interaction index $I_{j,l}^{\text{rel}}$. However, both SHAP interaction indices and the H-Statistic are based on the joint distribution of the two regarded features, and hence, correlations between \mathbf{x}_S and features in \mathbf{x}_C might bias the interaction value calculated by these methods, as demonstrated in Section 4.1.

3.2.2 Methodology

Here, We derive an interaction index based on the split criterion minimized in Algorithm 1 and Eq. (6), and we prove its advantages compared to alternatives mentioned in Section 3.2.1. Since the risk function of our split criterion is based on the variance of mean-centered ICE curves – which measures the degree of existing feature interactions with \mathbf{x}_S – we can use the achieved risk reduction after a split to quantify the interaction strength. For better interpretability and comparability, we define the relative interaction importance for each parent node \mathcal{N}_P by

$$\text{intImp}(\mathcal{N}_P) = \frac{\mathcal{R}_{L_2}(\mathcal{N}_P) - (\mathcal{R}_{L_2}(\mathcal{N}_l) + \mathcal{R}_{L_2}(\mathcal{N}_r))}{\mathcal{R}_{L_2}(\mathcal{N})} \quad (7)$$

with $l, r \in \{1, \dots, G\}$ denoting the left and right child node of a parent node \mathcal{N}_P and \mathcal{N} representing the root node. Hence, $\text{intImp}(\mathcal{N}_P)$ measures the relative risk reduction after splitting \mathcal{N}_P compared to the risk within the root node $\mathcal{R}_{L_2}(\mathcal{N})$. Let $\mathcal{B}_P \subset \{1, \dots, G\}$ denote the index set of all parent nodes (i.e., all nodes that have child nodes), and let $\mathcal{B}_j \subseteq \mathcal{B}_P$ denote the subset of these parent nodes that used the regarded

feature \mathbf{x}_j for splitting. To obtain the relative interaction importance of feature \mathbf{x}_j , we sum up the relative interaction importance over the parent nodes in \mathcal{B}_j :

$$\text{intImp}_j = \sum_{P \in \mathcal{B}_j} \text{intImp}(\mathcal{N}_P). \quad (8)$$

This principle of summing up the relative risk reduction of individual splits regarding a certain feature in order to measure the interaction strength is related to how a decision tree measures the Gini or mean decrease impurity (MDI) feature importance (Breiman et al., 1984). We obtain a measure that reports how important each of these features is for reducing interactions and thus obtaining more representative REPs for \mathbf{x}_S . Our proposed interaction importance in Eq. (8) only depends on the interaction effects between \mathbf{x}_j and \mathbf{x}_S and not on their main effects (see Theorem 2), as opposed to the H-Statistic or the interaction index of Greenwell et al. (2018). Furthermore, we show by Theorem 5 that intImp – in contrast to the H-Statistic and the SHAP interaction index $I_{j,l}^{\text{rel}}$ – is not influenced by correlations between \mathbf{x}_S and \mathbf{x}_j .

Theorem 5 *Correlations between X_S and X_C do not influence the splitting procedure of REPID, since the loss function \mathcal{L} of Eq. (5) does not contain a covariance term between X_S and features in X_C . The proof can be found in Appendix A.1.4.*

To determine how well the resulting REPs in the terminal nodes represent the underlying ICE curves, we derive an R^2 measure, which is commonly used in statistics. The R^2 can be calculated by $R^2 = 1 - \frac{\text{SSE}(\text{complex model})}{\text{SSE}(\text{baseline model})}$ where the baseline model is, e.g., a constant mean prediction and the SSE is the sum of squared errors of the model. The measure (usually) only takes values between 0 and 1 when applied on training data. While a value of 1 indicates that the complex model fits the data perfectly, a value of 0 implies that the complex model does not outperform the baseline model. Similar to this concept, we use the global PD plot as our baseline model. Our complex model is the additive combination of the REPs in the terminal nodes of the final tree. Hence, each additive functional component (REP) is only valid for the specified region. The SSE of each model is measured by the variability of the underlying ICE curves. Let $\mathcal{B}_t = \mathcal{B}_P^c$ denote the subset of terminal nodes in a symmetric tree. We derive an interaction-related R^2 measure by aggregating the interaction importance over all parent nodes \mathcal{B}_P :

$$R_{\text{int}}^2 = \sum_{P \in \mathcal{B}_P} \text{intImp}(\mathcal{N}_P) = 1 - \frac{\sum_{t \in \mathcal{B}_t} \mathcal{R}_{L_2}(\mathcal{N}_t)}{\mathcal{R}_{L_2}(\mathcal{N})} \quad (9)$$

A detailed derivation can be found in Appendix A.2.

For our example, we obtain the relative interaction importance values for \mathbf{x}_2 , as stated in Table 1. Since

both child nodes after the first split use \mathbf{x}_1 as the splitting feature, the relative interaction importance values of the two nodes can be aggregated to obtain $intImp_1 = 0.14$. It follows that REPID detects (only) the feature interactions with \mathbf{x}_2 that have been specified in the underlying data-generating process and also ranks them in the correct order. The total variance after the second split is reduced by $R_{int}^2 = 97.5\%$ compared to the root node, suggesting that resulting REPs are now meaningful representatives for the average marginal effect, as shown in Figure 3.

Table 1: Relative interaction importance on a node level (left) and on a feature level (right). Gray shadings indicate how $intImp_j$ is calculated from $intImp(\mathcal{N}_P)$. The parameters d and P indicate the tree depth and the index of the parent node, respectively.

d	P	\mathbf{x}_j	$intImp(\mathcal{N}_P)$	\mathbf{x}_j	$intImp_j$
0	1	\mathbf{x}_3	0.835	\mathbf{x}_3	0.835
1	2	\mathbf{x}_1	0.074	\mathbf{x}_1	0.14
1	3	\mathbf{x}_1	0.066		

Stop Criteria A possible stop criterion for the tree is to limit the maximum depth of the tree or to define a minimum number of observations for each node. Furthermore, we can apply a stop criterion based on the interaction importance $intImp$. Let \mathcal{N}_g be the node we want to split and let \mathcal{N}_P be its parent node. Then, we only split deeper if $intImp(\mathcal{N}_g) \geq \gamma \cdot intImp(\mathcal{N}_P)$, with $\gamma \in [0, 1]$. In other words, we only split deeper if the improvement of the current split is at least as large as a pre-specified proportion of the improvement of the previous split. The suggested criteria can also be combined and the hyperparameters must be chosen by the user and usually depend on the underlying setting.

4 SIMULATION EXAMPLES

For many model-agnostic interpretation techniques – including interaction detection methods – ground truth information is usually not available on real-world data. Therefore, well-constructed simulation experiments with a known ground truth are often used for empirical evaluations and comparisons, while only one or few real-world datasets are used to demonstrate practical applicability (e.g., see Friedman et al. (2008), Fisher et al. (2019), Goldstein et al. (2015), Greenwell et al. (2018), or Aas et al. (2021)). Hence, we follow this commonly used approach to evaluate our method using various simulation settings.

4.1 Weaknesses of other Methods

In Section 3.2.1, we described disadvantages of several interaction measures from a theoretical perspective. In the following simulation example, we provide further empirical evidence. To be able to modify the degree of the feature dependencies later on, we use a Gaussian copula to simulate the data in all settings. In the initial setting, we draw 1000 samples of four approximately i.i.d. random variables, which are marginally $X_1, \dots, X_4 \sim \mathcal{U}(-1, 1)$, and assume the true underlying function of $f(\mathbf{x}) = r(\mathbf{x}) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, (\sigma(r(\mathbf{x})) \cdot 0.1)^2)$. We define the remainder by $r(\mathbf{x}) = \sum_{j=1}^4 \mathbf{x}_j + \mathbf{x}_1 \mathbf{x}_2 + \mathbf{x}_2 \mathbf{x}_3 + \mathbf{x}_1 \mathbf{x}_3 + \mathbf{x}_1 \mathbf{x}_2 \mathbf{x}_3$. To avoid undefined interaction effects, we fit a correctly specified linear model on the data. We repeat the experiment 30 times, and each time, we measure the interaction strength between \mathbf{x}_2 and the other three features using REPID as well as the three alternatives (the H-statistic, the Greenwell’s interaction index, and the SHAP interaction index). On three adjusted settings, we then illustrate that already small modifications of main effect sizes or feature dependencies may produce misleading results for some of the alternatives when used as a measure to rank interactions, while REPID provides correct and stable results. For the computations, we used an equidistant grid of size 20 for REPID and Greenwell’s interaction index. For better comparability, we used a sample size of 20 for the H-Statistic. We calculated the SHAP interaction index by aggregating the individual interaction indices for 100 randomly sampled observations, which are approximated by using 20 random permutations for all possible feature coalitions. For REPID, we combine the stop criteria described in Section 3.2.2 as follows: We use a maximum depth of 6, a minimum number of 10 observations per node, and an improvement factor of $\gamma = 0.15$.

(1) *Initial Setting:* The plot on the top left of Figure 4 shows that, for the initial setting, all methods on average correctly assign the same interaction importance to \mathbf{x}_1 as to \mathbf{x}_3 , while \mathbf{x}_4 does not interact with \mathbf{x}_2 .

(2) *Small main effects:* If we reduce the main effect of \mathbf{x}_1 to 0.1, we observe in the top right plot of Figure 4 that its interaction strength with \mathbf{x}_2 increases on average when the H-Statistic is used. This effect can be explained by Theorem 4. Hence, when main effects decrease, the proportion of the variance that explains the interaction between \mathbf{x}_1 and \mathbf{x}_2 increases compared to the proportion of the variance that explains the respective main effects. Also the method of Greenwell’s interaction index depends on the main effect sizes. However, since Greenwell’s interaction index includes the main effects in the nominator, the effect on the resulting interaction index is opposite to the one of the H-Statistic which includes the main effects

in the denominator. On the other hand, the SHAP interaction index as well as REPID are only based on interaction effects, and hence, varying main effects do not change the ranking. The plot on the bottom right of Figure 4 illustrates how problematic small main effects can be when the H-Statistic is applied. The H-Statistic leads to average interaction values close to 1 for \mathbf{x}_1 and \mathbf{x}_3 , although the actual interaction effect of \mathbf{x}_1 with \mathbf{x}_2 is twice as high as that of \mathbf{x}_3 with \mathbf{x}_2 .

(3) *Dependencies between the feature of interest and other features:* In the lower left plot of Figure 4, the correlation between \mathbf{x}_1 and \mathbf{x}_2 has been set to $\rho_{12} \approx 0.9$. Since we face a positive linear interaction effect between \mathbf{x}_1 and \mathbf{x}_2 , a positive linear correlation between these features leads to an increasing denominator of the H-Statistic. Hence, the respective H-Statistic value decreases compared to features that are independent of \mathbf{x}_2 (here, \mathbf{x}_3). The SHAP interaction index for \mathbf{x}_1 is higher than for \mathbf{x}_3 , since in this case, it can be shown that the interaction strength is an additive combination of the interaction effect and the covariance of the interacting features. Conversely, Greenwell’s interaction index is based on the variance of conditional marginal effects, and hence, the interaction index is not influenced by dependencies between the feature of interest and other features. The same holds for REPID, as proven with Theorem 5.

A summary of the simulation settings and key results is provided in Appendix B.1. Detailed theoretical derivations and explanations can be found in Appendix A.3.

4.2 Comparison on More Complex Settings

The aim in this simulation is to show that REPID detects existing interactions correctly in a more complex non-linear setting and to compare the results to the H-Statistic. Analogous to Hu et al. (2020), we draw 2000 samples of 10 independently and uniformly distributed random variables $X_1, \dots, X_{10} \sim \mathcal{U}(-1, 1)$ and assume the following true underlying function:

$$\begin{aligned}
 f(\mathbf{x}) = & 6\mathbf{x}_1 + \mathbf{x}_2^2 - \pi^{\mathbf{x}_3} + \exp^{-2\mathbf{x}_4^2} + (2 + |\mathbf{x}_5|)^{-1} \\
 & + \mathbf{x}_6 \log |\mathbf{x}_6| + 2\mathbf{x}_3 \mathbb{1}_{(\mathbf{x}_1 > 0)} \mathbb{1}_{(\mathbf{x}_2 > 0)} + 2\mathbf{x}_2 \mathbb{1}_{(\mathbf{x}_4 > 0)} \\
 & + 4(\mathbf{x}_2 \mathbb{1}_{(\mathbf{x}_2 > 0)})^{|\mathbf{x}_6|} + |\mathbf{x}_2 + \mathbf{x}_8| + \epsilon
 \end{aligned}$$

with $\epsilon \sim \mathcal{N}(0, 0.25)$. Hence, \mathbf{x}_2 interacts with five other features in a more complex and non-linear way. To avoid undefined interaction effects in a fitted model, we fit a correctly specified generalized additive model (GAM) and a tree-based extreme gradient boosting model (XGBOOST) with correctly specified interaction constraints⁷, a learning rate of 0.1, a maximum

⁷The “xgboost” library (Chen and Guestrin, 2016) enables definition of which features are allowed to interact with each other.

number of iterations of 1000, and a maximum tree depth of 6 on the simulated data. The performance of each model is measured by a separately simulated test set with the same distributional assumptions of size 100,000 and is reported in Figure 5. We repeat the experiment 30 times, and each time, we measure the interaction strength between \mathbf{x}_2 and the other nine features using REPID and the H-Statistic. For both methods, we again use a grid size of 20. For REPID, we apply the same stop criteria as in Section 4.1 but with a maximum tree depth of 7 due to a more complex setting. The results are illustrated in Figure 5. REPID correctly identifies only the true interactions for both models. In most of the repetitions, the H-Statistic does not find an interaction between \mathbf{x}_1 and \mathbf{x}_2 for the GAM. A possible reason for this behavior is the rather high main effect of \mathbf{x}_1 compared to the interaction effect (Theorem 4). More experiments of different models and settings – including varying values of λ to obtain shallower or deeper trees – can be found in Appendix B.2. The experiments show that shallow trees produce fewer regions and are therefore easier to interpret. However, they might only detect the most important interactions. Deeper trees are more likely to also identify less important interactions but are less interpretable.

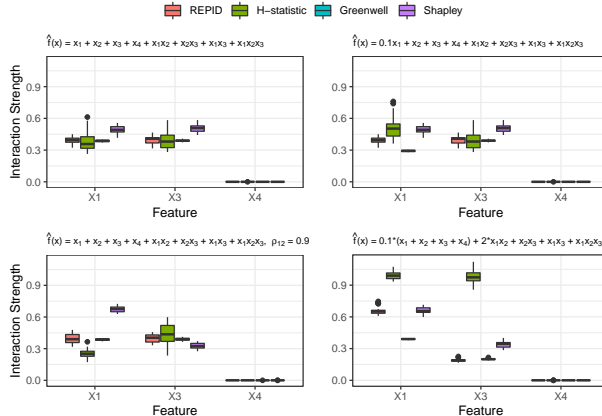


Figure 4: Comparison of REPID, the H-Statistic, Greenwell’s, and SHAP interaction indices for interactions between \mathbf{x}_2 and all other features for 30 repetitions. The upper left plot shows the initial setting (1). The upper and lower right plots adjust effect sizes (2), while the bottom left plot adjusts the correlation (3).

5 REAL-WORLD EXAMPLE

We now demonstrate the usefulness of REPID on the *titanic* data (Dawson, 1995). The labeled part of the dataset consists of 11 characteristics of 891 passengers of the ocean liner Titanic and a binary label if they survived. After some pre-processing steps that are de-

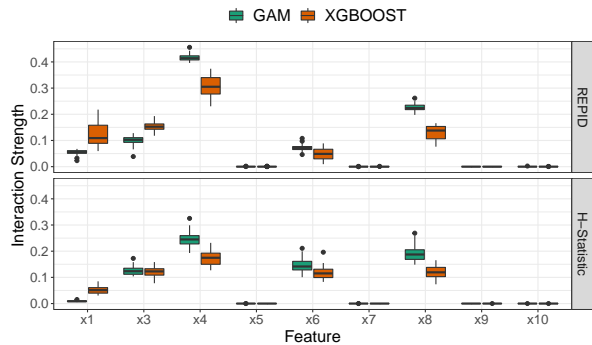


Figure 5: Comparison of the interaction strength between x_2 and all other features measured by REPID (top) and the H-Statistic (bottom) on 30 repetitions. The mean (standard deviation) of the models’ test performance (measured by the mean squared error) is: GAM: 0.36 (0.01), XGBOOST: 0.57 (0.11).

scribed in more detail in Appendix B.4, we train a RF with 500 trees on the dataset. Therefore, we obtain a balanced accuracy of 0.8 under 5-fold cross-validation. We are interested in how the age of the passengers affects the probability of survival. The left plot in Figure 6 shows that, from 0 to 20 years, the PD plot for passengers continuously decreases and then flattens above 20 years. The ICE curves indicate that age might influence the predicted survival probability for different passengers in different ways, and thus, interactions with other features might be present. The REPs after applying REPID by using a grid size of 20, a maximum depth of 3, a minimum number of 30 observations, and $\gamma = 0.2$ are shown in the right plot of Figure 6. The 3 most interacting features are Sex, Pclass (passenger class), and Fare. The green REPs show that the predicted survival probability of female passengers is on average higher compared to their male counterparts independent of their age. However, it is also visible that the probability strongly depends on the passenger’s class and the fare they paid. While female passengers who paid a high fare or who belong to an upper or middle class show an overall high survival probability independent of their age (even slightly increasing until 30), the survival probability of women with a low fare and Pclass drops with age. For men from middle and lower classes, the predicted survival probability drops dramatically from 0 until 20 to 30, meaning that for the sub-population of male passengers, the chances of survival are several factors higher for children than for adults.

More real-world examples for the *California housing* (Pace and Barry, 1997) and the *diabetes* (Smith et al., 1988) datasets are provided in Appendix B.4.




Figure 6: Global PD plot (blue) including ICE curves (left) and the REPs after applying REPID (right) for the feature of interest Age of the *titanic* dataset. The interaction importance $intImp_j$ between Age and the interacting features is 0.28 (Sex), 0.17 (Pclass), 0.13 (Fare), 0.06 (Embarked) and $R_{int}^2 = 0.64$.

6 DISCUSSION

We have introduced the interaction detection method REPID, which provides more representative PD plots on interpretable regions and enables quantification of feature interactions. We have proven its theoretical and empirical advantages and demonstrated how it out-performs alternatives presented in Section 3 and 4. Unlike the H-Statistic or SHAP interaction index, REPID is not influenced by correlations between the feature of interest x_S and other features x_C . However, like the other methods, it might be affected if features within x_C are correlated. Furthermore, the method might be limited if the feature of interest is, e.g., highly skewed, especially if an equidistant grid is used for computations. Possible solutions might be feature transformations or to use a sample or quantile-based grid. As our method is based on a tree-based partitioning algorithm that is known to be unstable (Breiman, 1996), the question arises whether the splitting procedure in Algorithm 1 is a potential limitation. However, with regards to the interaction quantification, we demonstrated in Section 4 that we obtain stable results when repeating the experiments multiple times. A more detailed analysis on the robustness of the method can be found in Appendix B.3.

Author Contributions (CRediT taxonomy)

Contributing authors: Julia Herbinger¹, Bernd Bischl², Giuseppe Casalicchio³. Conceptualization: 1,3; Methodology: 1,3; Project administration: 3; Formal analysis: 1,3; Writing - original draft preparation: 1; Writing - review and editing: 1,2,3; Investigation: 1,3; Visualization: 1; Validation: 1,3; Software: 1,3; Funding acquisition: 2,3; Supervision: 2,3.

Acknowledgements

This work has been partially supported by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A, the Bavarian Ministry of Economic Affairs, Regional Development and Energy as part of the program “Bayerischen Verbundförderprogramms (BayVFP) – Förderlinie Digitalisierung – Förderbereich Informations- und Kommunikationstechnik” under the grant DIK-2106-0007 // DIK0260/02. The authors of this work take full responsibility for its content.

References

- Aas, K., Jullum, M., and Løland, A. (2021). Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *Artificial Intelligence*, 298:103502.
- Apley, D. W. and Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4):1059–1086.
- Breiman, L. (1996). Bagging predictors. *Machine learning*, 24(2):123–140.
- Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984). *Classification and Regression Trees*. Wadsworth.
- Britton, M. (2019). Vine: Visualizing statistical interactions in black box models. *arXiv preprint arXiv:1904.00561*.
- Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, pages 785–794, New York, NY, USA. ACM.
- Dawson, R. J. M. (1995). The “unusual episode” data revisited. *Journal of Statistics Education*, 3(3).
- Fisher, A., Rudin, C., and Dominici, F. (2019). All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research*, 20(177):1–81.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232.
- Friedman, J. H., Popescu, B. E., et al. (2008). Predictive learning via rule ensembles. *Annals of Applied Statistics*, 2(3):916–954.
- Fujimoto, K., Kojadinovic, I., and Marichal, J.-L. (2006). Axiomatic characterizations of probabilistic and cardinal-probabilistic interaction indices. *Games and Economic Behavior*, 55(1):72–99.
- Goldstein, A., Kapelner, A., Bleich, J., and Pitkin, E. (2015). Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics*, 24(1):44–65.
- Greenwell, B. M., Boehmke, B. C., and McCarthy, A. J. (2018). A simple and effective model-based variable importance measure. *arXiv preprint arXiv:1805.04755*.
- Grömping, U. (2020). Model-agnostic effects plots for interpreting machine learning models. Report 1/2020, Reports in Mathematics, Physics and Chemistry. Department II, Beuth University of Applied Sciences Berlin.
- Hooker, G. (2004). Discovering additive structure in black box functions. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 575–580. ACM.
- Hooker, G. (2007). Generalized functional anova diagnostics for high-dimensional functions of dependent variables. *Journal of Computational and Graphical Statistics*, 16(3):709–732.
- Hu, L., Chen, J., Nair, V. N., and Sudjianto, A. (2020). Surrogate locally-interpretable models with supervised machine learning algorithms. *arXiv preprint arXiv:2007.14528*.
- Lundberg, S. M., Erion, G. G., and Lee, S.-I. (2018). Consistent individualized feature attribution for tree ensembles. *arXiv preprint arXiv:1802.03888*.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35.
- Molnar, C., Casalicchio, G., and Bischl, B. (2020). Interpretable machine learning—a brief history, state-of-the-art and challenges. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 417–431. Springer.
- Molnar, C., König, G., Bischl, B., and Casalicchio, G. (2021a). Model-agnostic feature importance and effects with dependent features – a conditional subgroup approach. *arXiv preprint arXiv:2006.04628*.
- Molnar, C., König, G., Herbringer, J., Freiesleben, T., Dandl, S., Scholbeck, C. A., Casalicchio, G., Grosse-Wentrup, M., and Bischl, B. (2021b). General pitfalls of model-agnostic interpretation methods for machine learning models. *arXiv preprint arXiv:2007.04131*.
- Pace, R. K. and Barry, R. (1997). Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297.

- Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). “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.
- Shapley, L. S. (1953). A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317.
- Smith, J. W., Everhart, J. E., Dickson, W., Knowler, W. C., and Johannes, R. S. (1988). Using the adap learning algorithm to forecast the onset of diabetes mellitus. In *Proceedings of the annual symposium on computer application in medical care*, page 261. American Medical Informatics Association.
- Strumbelj, E. and Kononenko, I. (2014). Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41(3):647–665.
- Su, X., Tsai, C.-L., Wang, H., Nickerson, D. M., and Li, B. (2009). Subgroup analysis via recursive partitioning. *Journal of Machine Learning Research*, 10(2).

Supplementary Material:

REPID: Regional Effect Plots with implicit Interaction Detection

A THEORETICAL EVIDENCE

A.1 Proofs

Here, we provide the proofs of the Theorems defined in Section 3. For each Theorem, we first provide a textual description in a proof sketch followed by the formal proof.

Note: For our proofs, we apply the concept of functional decomposition. One concept of functional decomposition has been introduced in Section 2. The so-called functional ANOVA (fANOVA) decomposition is a well-known approach to decompose a function in main and interaction effects. The fANOVA decomposition defined in Section 2 is based on Hooker (2004), and according to this definition, covariates must be independent to obtain a unique decomposition. However, we argue that this is not a relevant issue for our methods, since: (1) We do not try to estimate or calculate the decomposed mean-zero function terms g_W ; we only use the (valid) assumption that a function can be decomposed as in Eq. (3) to prove our theorems. Hence, we are not directly interested in a unique solution of the decomposition. (2) Still, it is possible to relax this assumption by using the generalized fANOVA (Hooker, 2007), which is a weighted version of the “normal” fANOVA to address the extrapolation problem when strong correlations are present. However, it is also possible to use another functional decomposition (e.g., as done in Apley and Zhu (2020)) for these proofs.

A.1.1 Proof of Theorem 1 and Corollary 1.1

Proof Sketch Since $E_{X_C} [\hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)})] = \hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)})$ and if Eq. (3) holds, the fANOVA decomposition can also be applied to the i -th ICE curve. Since $\mathbf{x}_C^{(i)}$ is constant in i , all fANOVA components that do not depend on \mathbf{x}_S can be summarized to an individual intercept shift of observation i and, thus, cancelled out by mean-centering an ICE curve. The remaining term is then defined by the mean-centered main and mean-centered individual interaction effect of \mathbf{x}_S for observation i . Taking the expected value w.r.t. X_C results in an analogous decomposition of the PD function and mean-centered PD function, respectively.

Proof 1 We first derive the fANOVA decomposition of the i -th ICE curve $\hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)})$ using Eq. (3) and use this decomposition to derive the mean-centered version $\hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)})$ for $|S| = 1$. Therefore, we first decompose the function into main and interaction effects depending on \mathbf{x}_S . *Note:* The term g_0 represents a constant intercept shift. This term is necessary to receive zero-mean functional components, i.e., e.g., $E_X[g_S(X_S)] = 0$.

$$\begin{aligned}
 \hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)}) &= E_{X_C|X_C} [\hat{f}(\mathbf{x}_S, X_C)|X_C = \mathbf{x}_C^{(i)}] \\
 &= \underbrace{g_0}_{\text{constant term}} + \underbrace{g_S(\mathbf{x}_S)}_{\text{main effect of } \mathbf{x}_S} + \underbrace{\sum_{j \in C} g_j(\mathbf{x}_j^{(i)})}_{\substack{\text{main effect of all other} \\ \text{features } \mathbf{x}_j \text{ for observation } i}} \\
 &+ \underbrace{\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)})}_{(k+1)\text{-order interaction between } \mathbf{x}_S \text{ and } \mathbf{x}_{C_k} \text{ for observation } i} + \underbrace{\sum_{k=2}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k}(\mathbf{x}_{C_k}^{(i)})}_{k\text{-order interaction between features within } C_k \text{ for observation } i}
 \end{aligned}$$

$$\begin{aligned}
 \hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)}) &= \hat{f}(\mathbf{x}_S, \mathbf{x}_C^{(i)}) - E_{X_S} [\hat{f}(X_S, \mathbf{x}_C^{(i)})] \\
 &= g_0 + g_S(\mathbf{x}_S) + \sum_{j \in C} g_j(\mathbf{x}_j^{(i)}) + \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)}) + \sum_{k=2}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k}(\mathbf{x}_{C_k}^{(i)}) \\
 &\quad - g_0 - \underbrace{E_{X_S} [g_S(X_S)]}_{=0} - \sum_{j \in C} g_j(\mathbf{x}_j^{(i)}) - E_{X_S} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(X_S, \mathbf{x}_{C_k}^{(i)}) \right] - \sum_{k=2}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k}(\mathbf{x}_{C_k}^{(i)}) \\
 &= \underbrace{g_S^{cS}(\mathbf{x}_S)}_{\substack{\text{mean-centered} \\ \text{main effect of } \mathbf{x}_S}} + \underbrace{\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)}) - E_{X_S} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(X_S, \mathbf{x}_{C_k}^{(i)}) \right]}_{\text{mean-centered interaction effect of } \mathbf{x}_S \text{ with } \mathbf{x}_C^{(i)} \text{ for observation } i} \\
 &= \underbrace{g_S^{cS}(\mathbf{x}_S)}_{\substack{\text{mean-centered} \\ \text{main effect of } \mathbf{x}_S}} + \underbrace{\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)})}_{\substack{\text{mean-centered interaction effect of} \\ \mathbf{x}_S \text{ with } \mathbf{x}_C^{(i)} \text{ for observation } i}}
 \end{aligned}$$

Proof 1.1 We first derive the fANOVA decomposition of the PD function $\hat{f}_S^{PD}(\mathbf{x}_S)$ using Eq. (3) and use this decomposition to derive its mean-centered version $f_S^{PD,c}(\mathbf{x}_S)$ for $|S| = 1$.

$$\begin{aligned}
 f_S^{PD}(\mathbf{x}_S) &= E_{X_C} [\hat{f}(\mathbf{x}_S, X_C)] \\
 &= E_{X_C} \left[g_0 + g_S(\mathbf{x}_S) + \sum_{j \in C} g_j(\mathbf{x}_j^{(i)}) + \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, X_{C_k}) + \sum_{k=2}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k}(X_{C_k}) \right] \\
 &= g_0 + g_S(\mathbf{x}_S) + E_{X_C} \left[\underbrace{\sum_{j \in C} g_j(\mathbf{x}_j^{(i)})}_{\substack{\text{expected main effect} \\ \text{of features in } \mathbf{x}_C (=0)}} \right] + E_{X_C} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, X_{C_k}) \right] + E_{X_C} \left[\underbrace{\sum_{k=2}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k}(X_{C_k})}_{\substack{\text{expected interaction effect} \\ \text{of features in } \mathbf{x}_C (=0)}} \right] \\
 &= g_0 + \underbrace{g_S(\mathbf{x}_S)}_{\text{main effect of } \mathbf{x}_S} + E_{X_C} \left[\underbrace{\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, X_{C_k})}_{\substack{\text{expected interaction effect} \\ \text{of } \mathbf{x}_S \text{ with } \mathbf{x}_C \text{ w.r.t. } \mathbf{x}_C}} \right]
 \end{aligned}$$

If the expected value of each decomposed term $g(\mathbf{x})$ exists and if the integral of the absolute value is finite, then Fubini's theorem can be applied, and the mean-centered PD function of \mathbf{x}_S for $|S| = 1$ can be derived by:

$$\begin{aligned}
 f_S^{PD,c}(\mathbf{x}_S) &= f_S^{PD}(\mathbf{x}_S) - E_{X_S} [f_S^{PD}(X_S)] \\
 &= E_{X_C} [\hat{f}(\mathbf{x}_S, X_C)] - E_{X_S} \left[g_0 + g_S(X_S) + E_{X_C} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(X_S, X_{C_k}) \right] \right]
 \end{aligned}$$

$$\begin{aligned}
 &= g_0 + g_S(\mathbf{x}_S) + E_{X_C} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(\mathbf{x}_S, X_{C_k}) \right] \\
 &- g_0 - \underbrace{E_{X_S} [g_S(X_S)]}_{=0} - E_{X_S} \left[E_{X_C} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}(X_S, X_{C_k}) \right] \right] \\
 &= \underbrace{g_S^{cS}(\mathbf{x}_S)}_{\text{mean-centered main effect of } \mathbf{x}_S} + E_{X_C} \left[\sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, X_{C_k}) \right] \\
 &\quad \underbrace{\hspace{10em}}_{\text{expected mean-centered interaction effect of } \mathbf{x}_S \text{ with } \mathbf{x}_C \text{ w.r.t. } \mathbf{x}_C}
 \end{aligned}$$

A.1.2 Proof of Theorem 2

Proof Sketch If the function $\hat{f}(\mathbf{x})$ can be decomposed as in Eq. (3), then Theorem 1 and Corollary 1.1 hold, and the main effect of \mathbf{x}_S is cancelled out when calculating $\mathcal{R}_{L2}(\mathcal{N}_g)$. The remaining term is given by the distance between the i -th centered interaction effect and the average centered interaction effect between \mathbf{x}_S and \mathbf{x}_C .

Proof 2 In the risk function of Eq. (6), the squared distance between the i -th mean-centered ICE curve $\hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)})$ and the respective PD function $f_S^{PD,c}(\mathbf{x}_S)$ is calculated. The distance can be reduced to the following term:

$$\begin{aligned}
 \hat{f}^c(\mathbf{x}_S, \mathbf{x}_C^{(i)}) - f_S^{PD,c}(\mathbf{x}_S) &= g_S^{cS}(\mathbf{x}_S) + \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)}) - g_S^{cS}(\mathbf{x}_S) - \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} E_{X_C} [g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, X_{C_k})] \\
 &= \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} (g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, \mathbf{x}_{C_k}^{(i)}) - E_{X_C} [g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, X_{C_k})])
 \end{aligned}$$

The first term is the mean-centered interaction effect of the i -th ICE curve, while the second term represents the mean-centered expected interaction effect over the joint distribution of \mathbf{x}_C (which is included in the mean-centered PD function, see also the decomposition of the mean-centered PD function $f_S^{PD,c}(\mathbf{x}_S)$ in the proof in Appendix A.1.1). The intuition behind our split criterion is that we search for the optimal split value of a feature in \mathbf{x}_C that reduces the aggregated variance over all curves the most if we split according to this optimal split value. Thus, we try to find regions in the feature space \mathbf{x}_C where the distance between the individual centered ICE curves in this region and the respective mean-centered PD plot is as small as possible. Hence, we want to minimize the deviation of the individual interaction effect of the ICE curves in a region from the average interaction effect in the considered region.

A.1.3 Proof of Theorem 4

Proof Sketch The two-way interaction index of the H-Statistic is calculated by dividing the variance of the difference between the centered 2-dimensional PD plot and the 1-dimensional PD plots of the two features of interest (nominator) by the variance of the centered 2-dimensional PD plot (denominator, see Eq. (2)). If Eq. (3) holds, we can apply Theorem 1 and Corollary 1.1, and it can be shown that the main effects of the two features of interest are cancelled out in the nominator, but are still present in the denominator (scaling factor) of the interaction index.

Proof 4 Let $S = \{j, l\}$ and $C = S^c$ its complement, then the 2-dimensional PD function $f_S^{PD}(\mathbf{x}_S)$ of \mathbf{x}_j and \mathbf{x}_l is given by

$$\begin{aligned}
 f_S^{PD}(\mathbf{x}_S) &= E_{X_C} [f(\mathbf{x}_S, X_C)] \\
 &= g_0 + g_j(\mathbf{x}_j) + g_l(\mathbf{x}_l) + g_{jl}(\mathbf{x}_j, \mathbf{x}_l) + E_{X_C} \left[\underbrace{\sum_{k=1}^{p-2} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k}(X_{C_k})}_{\text{expected interaction effect of features in } \mathbf{x}_C (=0)} \right] \\
 &+ E_{X_C} \left[\sum_{k=1}^{p-2} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{j\}}(\mathbf{x}_j, X_{C_k}) + g_{C_k \cup \{l\}}(\mathbf{x}_l, X_{C_k}) + g_{C_k \cup \{S\}}(\mathbf{x}_S, X_{C_k}) \right]
 \end{aligned}$$

If the expected value of each decomposed term $g(\mathbf{x})$ exists, and if the integral of the absolute value is finite, then Fubini's theorem can be applied, and the mean-centred 2-dimensional PD function $f_S^{PD,c}(\mathbf{x}_S)$ of features \mathbf{x}_j and \mathbf{x}_l can then be derived by

$$\begin{aligned}
 f_S^{PD,c}(\mathbf{x}_S) &= f_S^{PD}(\mathbf{x}_S) - E_{X_S} [f_S^{PD,c}(X_S)] \\
 &= g_0 + g_j(\mathbf{x}_j) + g_l(\mathbf{x}_l) + g_{jl}(\mathbf{x}_j, \mathbf{x}_l) \\
 &+ E_{X_C} \left[\sum_{k=1}^{p-2} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{j\}}(\mathbf{x}_j, X_{C_k}) + g_{C_k \cup \{l\}}(\mathbf{x}_l, X_{C_k}) + g_{C_k \cup \{S\}}(\mathbf{x}_S, X_{C_k}) \right] \\
 &- g_0 - \underbrace{E_{X_S} [g_j(X_j) + g_l(X_l) + g_{jl}(X_j, X_l)]}_{=0} \\
 &- E_{X_S} \left[\underbrace{E_{X_C} \left[\sum_{k=1}^{p-2} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{j\}}(X_j, X_{C_k}) + g_{C_k \cup \{l\}}(X_l, X_{C_k}) + g_{C_k \cup \{S\}}(X_S, X_{C_k}) \right]}_{\text{expected interaction effect between } \mathbf{x}_S \text{ and } \mathbf{x}_C (=0)} \right] \\
 &= \underbrace{g_j^{cS}(\mathbf{x}_j) + g_l^{cS}(\mathbf{x}_l)}_{\text{mean-centered main effects of } \mathbf{x}_S} + \underbrace{g_{jl}^{cS}(\mathbf{x}_j, \mathbf{x}_l)}_{\text{mean-centered interaction effect between } \mathbf{x}_j \text{ and } \mathbf{x}_l} \\
 &+ E_{X_C} \left[\sum_{k=1}^{p-2} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{j\}}^{cS}(\mathbf{x}_j, X_{C_k}) + g_{C_k \cup \{l\}}^{cS}(\mathbf{x}_l, X_{C_k}) + g_{C_k \cup \{S\}}^{cS}(\mathbf{x}_S, X_{C_k}) \right] \\
 &\underbrace{\hspace{10em}}_{\text{expected mean-centered interaction effects between features in } \mathbf{x}_S \text{ and features in } \mathbf{x}_C \text{ w.r.t. } \mathbf{x}_C}
 \end{aligned}$$

It follows that the H-Statistic still depends on the mean-centered main effects $g_j^{cS}(\mathbf{x}_j)$ and $g_l^{cS}(\mathbf{x}_l)$ of \mathbf{x}_j and \mathbf{x}_l in the denominator.

To calculate the nominator of the H-Statistic, we must subtract the 1-dimensional mean-centered PD functions of x_j and x_l as follows:

$$f_S^{PD,c}(\mathbf{x}_S) - f_j^{PD,c}(\mathbf{x}_j) - f_l^{PD,c}(\mathbf{x}_l) = g_j^{cS}(\mathbf{x}_j) + g_l^{cS}(\mathbf{x}_l) + g_{jl}^{cS}(\mathbf{x}_j, \mathbf{x}_l)$$

$$\begin{aligned}
 & + E_{X_C} \left[\sum_{k=1}^{p-2} \sum_{\substack{C_k \subseteq C, \\ |C_k|=k}} g_{C_k \cup \{j\}}^{c_S}(\mathbf{x}_j, X_{C_k}) + g_{C_k \cup \{l\}}^{c_S}(\mathbf{x}_l, X_{C_k}) + g_{C_k \cup \{S\}}^{c_S}(\mathbf{x}_S, X_{C_k}) \right] \\
 & - g_j^{c_S}(\mathbf{x}_j) - \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C \cup \{l\}, \\ |C_k|=k}} E_{X_{C \cup \{l\}}} \left[g_{C_k \cup \{j\}}^{c_S}(\mathbf{x}_j, X_{C_k}) \right] \\
 & - g_l^{c_S}(\mathbf{x}_l) - \sum_{k=1}^{p-1} \sum_{\substack{C_k \subseteq C \cup \{j\}, \\ |C_k|=k}} E_{X_{C \cup \{j\}}} \left[g_{C_k \cup \{l\}}^{c_S}(\mathbf{x}_l, X_{C_k}) \right]
 \end{aligned}$$

Thus, in the nominator of the H-Statistic, the variance of the calculated term is determined. This term only depends on interactions with features \mathbf{x}_j and \mathbf{x}_l , while the main effects $g_j^{c_S}(\mathbf{x}_j)$ and $g_l^{c_S}(\mathbf{x}_l)$ that are present in the denominator are cancelled out.

A.1.4 Proof of Theorem 5

Proof Sketch The loss function in Eq. (5), which is used for the splitting in Algorithm 1, is calculated grid-wise. This means that we calculate the variation measured by the estimated variance (L2 loss) for each grid point $x_S^{(k)}$ with $k \in \{1, \dots, m\}$. Hence, \mathbf{x}_S is not treated as a random variable but as a constant. It follows that when calculating the variance over all ICE curves on a specific grid point $x_S^{(k)}$, no covariance terms between X_S and features in X_C are considered.

Proof 5 $\mathcal{L}(\mathcal{N}_g, x_S)$ of Eq. (5) is estimated by taking the variance over all mean-centered ICE curves within a region \mathcal{N}_g for a fixed grid point of \mathbf{x}_S . Hence, for each grid point $k \in \{1, \dots, m\}$, we calculate:

$$\mathcal{L}(x_S^{(k)}, \mathcal{N}_g) = \text{Var}_{X|\mathcal{N}_g}[\hat{f}^c(X)|X_S = x_S^{(k)}] = \text{Var}_{X|\mathcal{N}_g}[\hat{f}^c(x_S^{(k)}, X_C)].$$

Since $x_S^{(k)}$ is constant, it follows $\text{Var}_{X|\mathcal{N}_g}[\hat{f}^c(x_S^{(k)}, X_C)] = \text{Var}_{X_C|\mathcal{N}_g}[\hat{f}^c(x_S^{(k)}, X_C)]$, and hence, the calculated variance only depends on features in C while there are no covariance terms between X_S and features in X_C included.

A.2 Derivation of R Squared Measure

Let $d = 0, \dots, D$ be the depth of the tree, where $d = 0$ is the depth of the root node and $d = D$ of the leaf nodes of a symmetric tree, and k defines the index of the node at each depth from left to right (starting from 0). With a slight abuse of notation, we denote \mathcal{R}_k^d as the risk of the k -th node at depth d . For example, \mathcal{R}_0^0 is the risk of the root node ($\mathcal{R}(\mathcal{N})$). Let $\mathcal{B}_t = \mathcal{B}_P^0$ denote the subset of terminal nodes in a symmetric tree. We can derive an interaction-related R^2 measure by aggregating the interaction importance over all parent nodes \mathcal{B}_P :

$$\begin{aligned}
 R_{int}^2 &= \sum_{P \in \mathcal{B}_P} \text{intImp}(\mathcal{N}_P) \\
 &= \frac{1}{\mathcal{R}_0^0} \cdot \sum_{d=0}^{D-1} \sum_{k=0}^d (\mathcal{R}_k^d - \mathcal{R}_{2k}^{d+1} - \mathcal{R}_{2k+1}^{d+1}) \\
 &= \frac{1}{\mathcal{R}_0^0} \cdot (\mathcal{R}_0^0 - \sum_{k=0}^{D-1} (\mathcal{R}_{2k}^D + \mathcal{R}_{2k+1}^D)) \\
 &= 1 - \frac{\sum_{k=0}^{D-1} (\mathcal{R}_{2k}^D + \mathcal{R}_{2k+1}^D)}{\mathcal{R}_0^0} \\
 &= 1 - \frac{\sum_{t \in \mathcal{B}_t} \mathcal{R}(\mathcal{N}_t)}{\mathcal{R}(\mathcal{N})}
 \end{aligned}$$

Explanation: According to Eq. (7), $\text{intImp}(\mathcal{N}_P)$ is defined by $\text{intImp}(\mathcal{N}_P) = \frac{\mathcal{R}(\mathcal{N}_P) - (\mathcal{R}(\mathcal{N}_l) + \mathcal{R}(\mathcal{N}_r))}{\mathcal{R}(\mathcal{N})}$ which is, e.g., for the first split (using the new notation defined in this section) the same as $\text{intImp}(\mathcal{N}) = \frac{\mathcal{R}_0^0 - (\mathcal{R}_0^1 + \mathcal{R}_1^1)}{\mathcal{R}_0^0}$ and for the split of the first left and right child nodes (which we denote here by \mathcal{N}_l and \mathcal{N}_r , respectively), we obtain $\text{intImp}(\mathcal{N}_l) = \frac{\mathcal{R}_0^1 - (\mathcal{R}_0^2 + \mathcal{R}_1^2)}{\mathcal{R}_0^0}$ and $\text{intImp}(\mathcal{N}_r) = \frac{\mathcal{R}_1^1 - (\mathcal{R}_2^2 + \mathcal{R}_3^2)}{\mathcal{R}_0^0}$. It follows that, after the second split ($D = 2$), R_{int}^2 can be calculated by

$$\begin{aligned}
 R_{int}^2 &= \text{intImp}(\mathcal{N}) + \text{intImp}(\mathcal{N}_l) + \text{intImp}(\mathcal{N}_r) \\
 &= \frac{1}{\mathcal{R}_0^0} (\mathcal{R}_0^0 - (\mathcal{R}_0^1 + \mathcal{R}_1^1) + \mathcal{R}_0^1 - (\mathcal{R}_0^2 + \mathcal{R}_1^2) + \mathcal{R}_1^1 - (\mathcal{R}_2^2 + \mathcal{R}_3^2)) \\
 &= \frac{1}{\mathcal{R}_0^0} \cdot \sum_{d=0}^1 \sum_{k=0}^d (\mathcal{R}_k^d - \mathcal{R}_{2k}^{d+1} - \mathcal{R}_{2k+1}^{d+1}) \\
 &= \frac{1}{\mathcal{R}_0^0} (\mathcal{R}_0^0 - (\mathcal{R}_0^2 + \mathcal{R}_1^2)) - (\mathcal{R}_2^2 + \mathcal{R}_3^2) \\
 &= \frac{1}{\mathcal{R}_0^0} \cdot (\mathcal{R}_0^0 - \sum_{k=0}^1 (\mathcal{R}_{2k}^{D=2} + \mathcal{R}_{2k+1}^{D=2})) \\
 &= 1 - \frac{\sum_{k=0}^{D-1} (\mathcal{R}_{2k}^D + \mathcal{R}_{2k+1}^D)}{\mathcal{R}_0^0} \\
 &= 1 - \frac{\sum_{t \in \mathcal{B}_t} \mathcal{R}(\mathcal{N}_t)}{\mathcal{R}(\mathcal{N})}
 \end{aligned}$$

From the second to the fourth line of the equation, we can see that the parent nodes (besides the root node) are cancelled out when aggregating the interaction importance over all nodes. It follows that only the deviation between the root node risk and the sum over all terminal node risks remains in the nominator. The denominator is always the root node (baseline) risk.

A.3 Explanations for Weaknesses of other Methods

A.3.1 Small Main Effects

For REPID, we proved with Theorem 2 that the split criterion only depends on interaction effects with the feature of interest \mathbf{x}_S and is independent of main effects. On the other hand, according to Theorem 4, the H-Statistic depends on main effects in the denominator of the H-Statistic. Since the main effect of feature \mathbf{x}_1 is reduced from 1 to 0.1 in the adjusted example of Section 4.1, the denominator of H-Statistic decreases, and hence, the overall H-Statistic value increases for feature \mathbf{x}_1 .

Since we provided proofs for REPID and for the H-Statistic, we will not go into more detail here, but instead derive explanations for the SHAP and Greenwell's interaction indices with regards to varying main effects.

SHAP interaction index By definition, SHAP interaction values only contain the interaction effect between the two features of interest and do not contain their main effects. Since we only sum up the absolute interaction values and divide them by the total amount of two-way interaction values between the feature of interest and all other features, there are also no main effects included in the global SHAP interaction index. Hence, varying main effects does not change the interaction strength / ranking calculated by the SHAP interaction index.

Example: Due to the complexity of an increasing number of feature permutations, we show this relationship on the following simple model: $\hat{f}(\mathbf{x}) = \hat{\beta}_1 \mathbf{x}_1 + \hat{\beta}_2 \mathbf{x}_2 + \hat{\beta}_{12} \mathbf{x}_1 \mathbf{x}_2$ with $E(X_1) = E(X_2) = 0$.

In this case, we can straightforwardly calculate the individual components of the SHAP interaction value with $S = \emptyset$:

$$f_{S \cup \{1,2\}}^{PD}(\mathbf{x}_{S \cup \{1,2\}}) = \hat{\beta}_1 \mathbf{x}_1 + \hat{\beta}_2 \mathbf{x}_2 + \hat{\beta}_{12} \mathbf{x}_1 \mathbf{x}_2$$

Since $E(X_1) = E(X_2) = 0$, it follows:

$$f_{S \cup \{1\}}^{PD}(\mathbf{x}_{S \cup \{1\}}) = \hat{\beta}_1 \mathbf{x}_1 \text{ and } f_{S \cup \{2\}}^{PD}(\mathbf{x}_{S \cup \{2\}}) = \hat{\beta}_2 \mathbf{x}_2 \text{ and } f_S^{PD}(\mathbf{x}_S) = E_X [\hat{f}(X)] = \hat{\beta}_{12} E_X [X_1 X_2]$$

and hence, the SHAP interaction value between \mathbf{x}_1 and \mathbf{x}_2 is given by

$$\begin{aligned}\Phi_{1,2}(\mathbf{x}) &= \frac{1}{2}(f_{S \cup \{1,2\}}^{PD} - f_{S \cup \{1\}}^{PD}(\mathbf{x}_{S \cup \{1\}}) - f_{S \cup \{2\}}^{PD}(\mathbf{x}_{S \cup \{2\}}) + f_S^{PD}(\mathbf{x}_S)) \\ &= \frac{1}{2}(\hat{\beta}_1 x_1 + \hat{\beta}_2 x_2 + \hat{\beta}_{12} x_1 x_2 - \hat{\beta}_1 x_1 - \hat{\beta}_2 x_2 + \hat{\beta}_{12} E_X [X_1 X_2]) \\ &= \frac{1}{2}(\hat{\beta}_{12} x_1 x_2 + \hat{\beta}_{12} E_X [X_1 X_2])\end{aligned}$$

Greenwell's interaction index Greenwell et al. (2018) defines feature importance $i(\mathbf{x}_j)$ as the standard deviation over the PD function of a feature \mathbf{x}_j with m_j unique values as follows:

$$i(\mathbf{x}_j)^2 = \frac{1}{m_j - 1} \sum_{k=1}^{m_j} \left(\hat{f}_j^{PD}(x_j^{(k)}) - \frac{1}{m_j} \sum_{k=1}^{m_j} \hat{f}_j^{PD}(x_j^{(k)}) \right)^2$$

To calculate the interaction between \mathbf{x}_j and \mathbf{x}_l , they define the conditional importance $i(\mathbf{x}_j | \mathbf{x}_l = \mathbf{x}_l^{(t)})$ of a feature \mathbf{x}_j given the t -th unique value of \mathbf{x}_l as follows:

$$i(\mathbf{x}_j | \mathbf{x}_l = x_l^{(t)})^2 = \frac{1}{m_j - 1} \sum_{k=1}^{m_j} \left(\hat{f}_j^{PD}(x_j^{(k)} | \mathbf{x}_l = x_l^{(t)}) - \frac{1}{m_j} \sum_{k=1}^{m_j} \hat{f}_j^{PD}(x_j^{(k)} | \mathbf{x}_l = x_l^{(t)}) \right)^2$$

With m_j and m_l being the number of unique values of \mathbf{x}_j and \mathbf{x}_l , respectively, the interaction measure $i(\mathbf{x}_j, \mathbf{x}_l)$ between these two features is then defined by:

$$\begin{aligned}i(\mathbf{x}_j, \mathbf{x}_l) &= \frac{1}{2} \sqrt{\frac{1}{m_l - 1} \sum_{t=1}^{m_l} \left[i(\mathbf{x}_j | \mathbf{x}_l = x_l^{(t)}) - \frac{1}{m_l} \sum_{t=1}^{m_l} i(\mathbf{x}_j | \mathbf{x}_l = x_l^{(t)}) \right]^2} \\ &\quad + \frac{1}{2} \sqrt{\frac{1}{m_j - 1} \sum_{k=1}^{m_j} \left[i(\mathbf{x}_l | \mathbf{x}_j = x_j^{(k)}) - \frac{1}{m_j} \sum_{k=1}^{m_j} i(\mathbf{x}_l | \mathbf{x}_j = x_j^{(k)}) \right]^2}\end{aligned}$$

Instead of conditioning on all features in C as done for ICE curves, Greenwell et al. (2018) conditions only on the second feature of interest (e.g., \mathbf{x}_l) to calculate the variation of PD curves for the first feature of interest (e.g., \mathbf{x}_j). Hence, they first take the variation of each conditioned curve and then calculate the variation over all these curves. Since they calculate the squared distance of each conditioned PD curve to its mean, the distance still contains the main effects of the two features of interest (see Theorem 1).

A.3.2 Dependencies between the Feature of Interest and other Features

For REPID, we proved with Theorem 5 that the loss function of Eq. (5) (which is used for splitting) is not affected by dependencies between the feature of interest \mathbf{x}_S and features in \mathbf{x}_C .

Hence, we will now derive explanations for the H-Statistic, the SHAP, and the Greenwell's interaction indices with regards to dependencies between the feature of interest and other features.

The H-Statistic The H-Statistic (which is estimated as in Eq. (2)) divides the variance of the difference between the mean-centered 2-dimensional PD plot and the two mean-centered 1-dimensional PD plots by the variance of the mean-centered 2-dimensional PD plot. Both the nominator and the denominator depend on the joint distribution of the two features of interest and, hence, also on the dependency between the two features.

Example Considering our simulation example in Section 4.1 with $E(X_1) = E(X_2) = E(X_3) = E(X_4) = 0$, the mean-centered 2-dimensional PD function between \mathbf{x}_1 and \mathbf{x}_2 with $S = \{1, 2\}$ is given by:

$$\hat{f}_S^{PD,c}(\mathbf{x}_1, \mathbf{x}_2) = \hat{\beta}_1 \mathbf{x}_1 + \hat{\beta}_2 \mathbf{x}_2 + \hat{\beta}_3 E(X_3) + \hat{\beta}_{12} \mathbf{x}_1 \mathbf{x}_2 + \hat{\beta}_{23} E(X_3) \mathbf{x}_2 + \hat{\beta}_{13} \mathbf{x}_1 E(X_3) + \hat{\beta}_{123} \mathbf{x}_1 E(X_3) \mathbf{x}_2$$

$$\begin{aligned}
& - \hat{\beta}_1 E(X_1) - \hat{\beta}_2 E(X_2) - \hat{\beta}_3 E(X_3) - \hat{\beta}_{12} E_{X_S} [X_1 X_2] - \hat{\beta}_{23} E(X_3) E(X_2) - \hat{\beta}_{13} E(X_1) E(X_3) \\
& - \hat{\beta}_{123} E_{X_S} [X_1 X_2] E(X_3) \\
& = \hat{\beta}_1 \mathbf{x}_1 + \hat{\beta}_2 \mathbf{x}_2 + \hat{\beta}_{12} (\mathbf{x}_1 \mathbf{x}_2 - E_{X_S} [X_1 X_2])
\end{aligned}$$

Calculating the denominator by taking the variance

$$\begin{aligned}
Var(\hat{f}_S^{PD,c}(\mathbf{x}_1, \mathbf{x}_2)) &= E \left[(\hat{\beta}_1 X_1 + \hat{\beta}_2 X_2 + \hat{\beta}_{12} (X_1 X_2 - E_{X_S} [X_1 X_2]))^2 \right] \\
& - E \left[\hat{\beta}_1 X_1 + \hat{\beta}_2 X_2 + \hat{\beta}_{12} (X_1 X_2 - E_{X_S} [X_1 X_2]) \right]^2 \\
&= E \left[\hat{\beta}_1^2 X_1^2 + 2\hat{\beta}_1 \hat{\beta}_2 X_1 X_2 + \hat{\beta}_2^2 X_2^2 + 2\hat{\beta}_1 \hat{\beta}_{12} X_1^2 X_2 + 2\hat{\beta}_2 \hat{\beta}_{12} X_1 X_2^2 \right] \\
&+ E \left[-2\hat{\beta}_1 \hat{\beta}_{12} X_1 E_{X_S} [X_1 X_2] - 2\hat{\beta}_2 \hat{\beta}_{12} X_2 E_{X_S} [X_1 X_2] + \hat{\beta}_{12}^2 X_1^2 X_2^2 \right] \\
&+ E \left[-2\hat{\beta}_{12}^2 X_1 X_2 E_{X_S} [X_1 X_2] + \hat{\beta}_{12}^2 E_{X_S} [X_1 X_2]^2 \right] \\
&= \hat{\beta}_1^2 Var(X_1) + \hat{\beta}_2^2 Var(X_2) + \hat{\beta}_{12}^2 Var(X_1 X_2) \\
&+ 2\hat{\beta}_1 \hat{\beta}_2 Cov(X_1, X_2) + 2\hat{\beta}_1 \hat{\beta}_{12} Cov(X_1^2, X_2) + 2\hat{\beta}_2 \hat{\beta}_{12} Cov(X_1, X_2^2)
\end{aligned}$$

in the nominator, we subtract the mean-centered 1-dimensional PD functions (i.e., $\hat{f}_1^{PD,c}(\mathbf{x}_1) = \hat{\beta}_1 \mathbf{x}_1$ and $\hat{f}_2^{PD,c}(\mathbf{x}_2) = \hat{\beta}_2 \mathbf{x}_2$) and take the variance, which results in

$$\begin{aligned}
& E_X \left[\hat{\beta}_{12}^2 (X_1 X_2 - E_{X_S} [X_1 X_2])^2 \right] - E_X \left[\hat{\beta}_{12} (X_1 X_2 - E_{X_S} [X_1 X_2]) \right]^2 \\
&= E_X \left[\hat{\beta}_{12}^2 X_1^2 X_2^2 - 2\hat{\beta}_{12}^2 X_1 X_2 E_{X_S} [X_1 X_2] + \hat{\beta}_{12}^2 E_{X_S} [X_1 X_2]^2 \right] \\
&= \hat{\beta}_{12}^2 Var(X_1 X_2) \\
&= \hat{\beta}_{12}^2 (Var(X_1)Var(X_2) - Cov(X_1, X_2)^2 + Cov(X_1^2, X_2^2))
\end{aligned}$$

It follows that by increasing the correlation between \mathbf{x}_1 and \mathbf{x}_2 to $\rho_{12} = 0.9$, the denominator of the H-Statistic increases compared to the nominator for the given example, and hence, the H-Statistic value between \mathbf{x}_1 and \mathbf{x}_2 decreases compared to the H-Statistic value between \mathbf{x}_2 and \mathbf{x}_3 .

Some general rules that were applied here:

- 1 Rearrangement of variance formula for functions: $Var(g(X)) = E[g(X)^2] - (E[g(X)])^2$
- 2 Expected value of a product of two random variables: $E[X_1 X_2] = E[X_1]E[X_2] + Cov(X_1, X_2)$ which reduces for $E(X_1) = E(X_2) = 0$ to $E[X_1 X_2] = Cov(X_1, X_2)$
- 3 Variance of a product of two random variables: $V(XY) = E[X^2 Y^2] - (E[XY])^2 = Cov(X^2, Y^2) + (V(X) + (E[X])^2)(V(Y) + (E[Y])^2) - (Cov(X, Y) + E[X]E[Y])^2$ which reduces for $E[X] = E[Y] = 0$ to $V(XY) = Cov(X^2, Y^2) + V(X)V(Y) - Cov(X, Y)^2$

SHAP interaction index SHAP interaction values – and with that, also the (global) SHAP Interaction index – depend on the correlation between the two features of interest, since we consider the joint distribution of the features as we do for the H-Statistic.

Example In Appendix A.3.1, we derived the SHAP interaction value for a simple linear model of two features with a positive linear interaction between these features, which resulted in

$$\Phi_{1,2}(\mathbf{x}) = \frac{1}{2}(\hat{\beta}_{12} \mathbf{x}_1 \mathbf{x}_2 + \hat{\beta}_{12} E_X [X_1 X_2])$$

Hence, if \mathbf{x}_1 and \mathbf{x}_2 are positively correlated as in our example in Section 4.1, then $E_X [X_1 X_2] > 0$, while this term is 0 if the two features are independent. This is why \mathbf{x}_1 shows a higher interaction value than \mathbf{x}_3 in the referred simulation study.

Greenwell’s interaction index Similarly to our approach, the Greenwell’s interaction index conditions on one of the two features of interest. They calculate the variance w.r.t. the other feature of interest, and vice versa. Hence, the dependency between the two regarded features does not influence the resulting interaction index.

B EMPIRICAL EVIDENCE

In this section, we provide more empirical evidence for the usefulness of REPID. We will further analyze the nonlinear simulation setting described in Section 4.2 and will also look at a linear example where interactions can clearly be ranked. Furthermore, we analyze the influence of the improvement parameter γ used as stop criterion and provide some evidence for the robustness of our method in Section B.3. In Section B.4, we clarify the pre-processing steps of the real-world example that was analyzed in Section 5.

Infrastructure All experiments only require CPUs (and no GPUs) and were computed on a Linux cluster (see Table 2).

Table 2: Description of the infrastructure used for the experiments in this paper.

Computing Infrastructure	
Type	Linux CPU Cluster
Architecture	28-way Haswell-EP nodes
Cores per Node	1
Memory limit (per core)	2.2 GB

B.1 Overview on Weaknesses of other Methods

In Table 3, we provide a brief overview of the simulation setting, including a sensitivity analysis that we performed in Section 4.1. The table shows that only REPID provides on average correct ranks for all settings, while the other state-of-the-art methods provide for at least one of the settings a wrong ranking (on average).

Table 3: Summary table of settings and key results of the simulation study in Section 4.1. The column “Setting” refers to the setting number in Section 4.1. The second column refers to the adjustments made in the setting compared to the initial setting. The other four columns show if the average ranks (r) of the feature interactions with the feature of interest (\mathbf{x}_2) are correct (meaning that the ranks are the same as the ranks of the underlying data-generating process and fitted linear model) or if they are wrong (different from the ranks in the data-generating process and fitted linear model).

Setting	Adjustment	REPID	H-Statistic	Greenwell	Shapley
(2)	$\beta_1 = 0.1$ (initial: 1)	correct $r(x_1) = r(x_3)$	wrong $r(x_1) > r(x_3)$	wrong $r(x_1) < r(x_3)$	correct
(2)	$\beta_1 = \beta_2 = \beta_3 = \beta_4 = 0.1$ and $\beta_{12} = 2$ (initial: 1)	correct $r(x_1) > r(x_3)$	wrong $r(x_1) = r(x_3)$	correct	correct
(3)	$\rho_{12} = 0.9$ (initial: 0)	correct $r(x_1) = r(x_3)$	wrong $r(x_1) < r(x_3)$	correct	wrong $r(x_1) > r(x_3)$

B.2 Further experiments

Nonlinear example In Section 4.2, we compared REPID and the H-Statistic for the interactions between the most interacting feature \mathbf{x}_2 and the other nine features of the simulation setting described in the referred section. In addition to the correctly specified GAM and XGBOOST model from Section 4.2, we now also compare the results to two other ML models: an RF with 500 trees – the mean and standard deviation of the models’ test performance (measured by the mean squared error) is 1.01 and 0.16 – and a support vector machine (SVM) using epsilon support vector regression with a Gauss kernel, $C = 1$ and $\epsilon = 0.1$ – the mean and standard deviation of the models’ test performance (measured by the mean squared error) is 0.76 and 0.07. The left plot in Figure 7 shows the same illustration as in Figure 5 for the interactions between the non-influential feature \mathbf{x}_{10} and all

other features. For the correctly specified GAM and XGBOOST model, both methods do – as expected – on average not find any interactions. While REPID on average also recognizes that there are no interactions present between \mathbf{x}_{10} and all other features for the SVM and RF models, the H-Statistic finds some higher interactions, especially for the SVM. A possible explanation is that \mathbf{x}_{10} does also not influence the target by a main effect in the underlying function, and hence, possible small found interaction effects might lead to high H-Statistic values. The outliers for some features when REPID is applied are possibly because the total variation of mean-centered ICE curves for non-influential features are rather small, and hence, relative loss reduction values might be high, although the absolute values are small. A potential solution to prevent these outliers is to extend the stop criterion by, e.g., a minimum absolute loss reduction constraint.

In the left plot in Figure 8, we analyzed the influence of the improvement parameter γ on the interaction strength. The difference between the threshold $\gamma = 0.15$, which we chose in Section 4.2, and $\gamma = 0.1$ is rather small, while it becomes more difficult to detect the smaller interactions with $\gamma = 0.2$. The smaller we choose γ to be, the deeper we split, and the less interaction variance remains in the final terminal nodes. Therefore, the obtained interaction strengths are more precise, and hence, our results seem to be more robust for different repetitions⁸. However, the deeper we split, the more final regions we obtain, which makes it more difficult to visually analyze the influence of the interactions on the marginal effect of the feature of interest. Hence, how to set the improvement parameter γ depends on the question the user would like to answer.

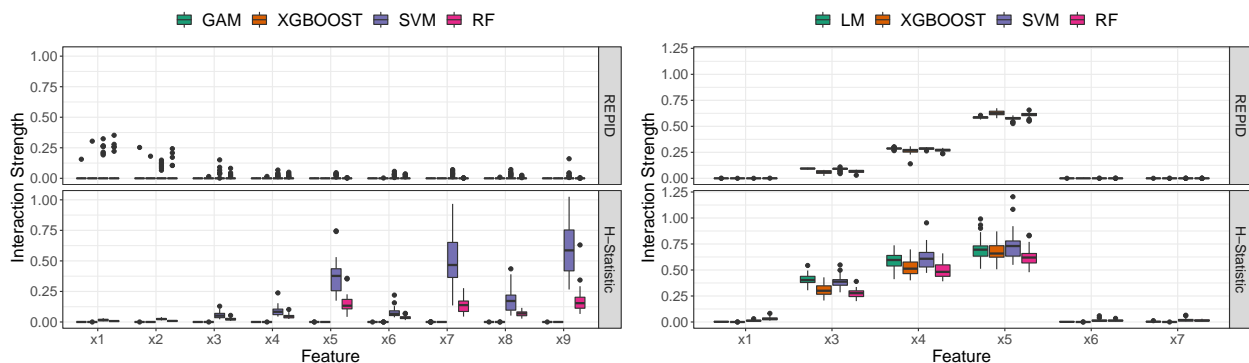


Figure 7: Left (right): The figure compares the interaction strength between \mathbf{x}_{10} (\mathbf{x}_2) and all other features measured by REPID and the H-Statistic for 4 different models on 30 repetitions of the described nonlinear (linear) simulation setting.

Linear example We now look at a further simulation example with only linear interaction effects between numeric features, which makes it possible to clearly rank the interactions between the feature of interest and all other features. Therefore, we draw 2000 samples of seven independent random variables, which are distributed as follows: $X_1, \dots, X_5 \sim \mathcal{U}(-1, 1)$, $X_6 \sim \mathcal{N}(0, 4)$ and $X_7 \sim \mathcal{N}(2, 9)$. The true underlying relationship is defined by $f(\mathbf{x}) = r(\mathbf{x}) + \epsilon$, where the remainder $r(\mathbf{x})$ is given by

$$r(\mathbf{x}) = \mathbf{x}_1 + 4\mathbf{x}_2 + 3\mathbf{x}_2\mathbf{x}_3 + 5\mathbf{x}_2\mathbf{x}_4 + 7\mathbf{x}_2\mathbf{x}_5$$

and $\epsilon \sim \mathcal{N}(0, (\sigma(r(\mathbf{x})) \cdot 0.1)^2)$. Hence, \mathbf{x}_5 interacts most with \mathbf{x}_2 , followed by \mathbf{x}_4 and then \mathbf{x}_3 . We fitted a linear model (LM) and an XGBOOST model with interaction constraints as well as an SVM and RF using the same configurations as for the nonlinear example on the simulated data. We repeated the experiment 30 times to quantify the interaction strength between \mathbf{x}_2 and all other features using REPID and the H-Statistic.⁹ We use the same specifications for the models' and interaction detection methods' hyperparameters as used in Section 4.2. The right plot in Figure 7 illustrates that both methods on average find the correct ranking of the feature interactions. However, REPID shows almost no variation over all repetitions and hence leads to more stable and clearer ranking results than the H-Statistic. In the right plot of Figure 8, the impact of the improvement parameter γ is shown. However, for this example, we barely see a difference between the different choices of γ , which might be due to the simplicity of the setting and hence that no deep trees are necessary to receive stable results for the interaction strength.

⁸The more robust results are shown by smaller interquartile ranges of boxplots in Figure 8.

⁹The mean (standard deviation) of the models' test performance (measured by the mean squared error) is for the LM: 0.15 (0.002), XGBOOST: 0.6 (0.22), SVM: 0.31 (0.069) and RF: 1.43 (0.34).

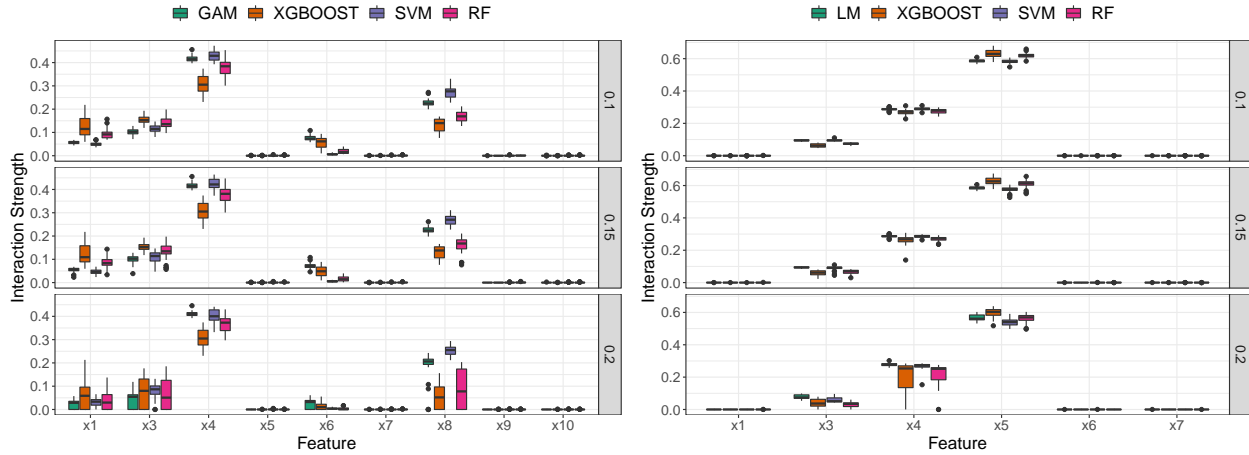


Figure 8: Left (right): The figure compares the interaction strength between \mathbf{x}_2 and all other features measured by REPID for 3 different improvement parameter values: $\gamma = 0.1$ (top), $\gamma = 0.15$ (middle), and $\gamma = 0.2$ (bottom) for 4 different models on 30 repetitions of the described nonlinear (linear) simulation setting.

B.3 Robustness analysis

An oft-stated limitation of the usage of decision trees is that they do not provide robust results. In Section 4 and B.2, we already showed that REPID provides robust results with regards to quantifying the interaction strength for different simulation settings. To investigate the robustness itself of the fitted trees, we extract and analyze the splits of the first three levels (depths) of the tree for the nonlinear example of Section 4.2, which is the most complex analyzed example of all examples in this paper. The frequencies of the features used at each split for the 30 repetitions is shown in Table 4 for each of the fitted models. For all repetitions and for all models, \mathbf{x}_4 was always chosen as the first splitting feature, with an average split value very close to 0, which shows only small variations (sd values). Furthermore, all models chose most often \mathbf{x}_8 for all nodes in the second level and \mathbf{x}_3 for all nodes in the third level of the tree. For the GAM that was correctly specified according to the true underlying function, the splits for the first three levels of the fitted decision tree barely differ. On the other hand, the SVM and the RF show higher variations. However, these models might have learned different interaction effects for different repetitions, and hence, it might be reasonable to receive different splits and REPs. The XGBOOST model also varies more than the GAM, which might be due to the fact that the GAM has a better and less variable model performance compared to the XGBOOST model, and hence, effect sizes might also vary less (see Figure 5). However, for all models, the feature chosen most often in each node is the same. It follows that REPID seems to provide robust results with regards to the interaction strength and the upper levels of the fitted tree if the same interactions have been learned by the ML models we want to explain.

B.4 Real-World Examples

Titanic dataset In Section 5, we applied REPID on the *titanic* dataset (Dawson, 1995). The labeled part of the dataset consists of 11 features and the binary survival target variable of 891 passengers. The features of the raw dataset include: *PassengerId*, *Name*, *Pclass*, *Sex*, *Age*, *SibSp*, *Parch*, *Ticket*, *Fare*, *Cabin*, *Embarked*, a detailed definition of each feature can be found at <https://www.kaggle.com/c/titanic/data>. To fit the RF model and analyze the predictions, we first pre-processed the data according to the following kaggle notebook <https://www.kaggle.com/nitinar1/titanic-solution-using-random-forest-tool-r>. The pre-processing steps can be summarized as follows:

- 1 We extract a title from the feature *Name* and categorize them into 5 categories (Master, Miss, Mr, Mrs and Rare Title).
- 2 We create a family size feature *FsizeD* from the features *Sibsp* as the number of siblings and *Parch* as the number of parents and children, and we categorize it into singleton, small and large family size.
- 3 We impute missing values of feature *Embarked* based on the fare price they paid.

- 4 We impute missing values of feature *Fare* by its median value of the respective *Pclass* and *Embarked* categories.
- 5 We impute the feature *Age* using a random forest imputation via multivariate imputation by chained equations.
- 6 We exclude the features *PassengerId*, *Name*, *Ticket*, *Cabin* from the dataset, which leaves us with nine features: *Pclass*, *Sex*, *Age*, *SibSp*, *Parch*, *Fare*, *Embarked*, *Title*, *FsizeD*.

California housing dataset As a second example, we applied REPID on the *California housing* dataset (Pace and Barry, 1997). The dataset contains information from the 1990 U.S. Census in California. Each of the 20640 observations provides information of a block group (small geographical unit), with an average population of around 1425 on the median house value (target), eight numeric features, and one categorical feature describing the ocean proximity. The features of the dataset include: *Longitude*, *Latitude*, *Housing median age*, *Total rooms*, *Total Bedrooms*, *Population*, *Households*, *Median Income* and *Ocean proximity*. A detailed definition of each feature can be found at <https://www.kaggle.com/camnugent/california-housing-prices>. Only the feature *Total bedroom* contains 207 missing values, which we imputed by the median value of *Total bedroom* of all other observations. Before applying the neural network on the data, we log transformed the target variable with a base of 10 and log transformed the features *Total rooms*, *Total Bedrooms*, *Population*, *Households*, *Median Income* using the natural logarithm. After pre-processing the data, we fit a neural net with one hidden layer of size 20, a weight decay of 0.1, and a maximum number of iterations of 1000. Thus, we obtain a mean absolute error (R-squared) of 0.08 (0.78) under 5-fold cross-validation. The left plot in Figure 9 shows that the median house value on average decreases the farther west a house is. The effect of individual observations seems to vary. However, visualizing ICE curves for such a high number of observations is not very insightful. In the right plot, we illustrate the resulting REPs after applying REPID with the same configurations as used for the titanic example in Section 5 but with $\gamma = 0.25$. The REPs show that the marginal effect of Longitude on the predicted median house value highly depends on how far north a house is (Latitude: the higher the value the farther north) and how close the house is to the ocean (Ocean proximity). For example, median values of houses that are farther north decrease with Longitude (light orange), while median values of houses farther south and not in the inland increase with Longitude (red).




Figure 9: The figure shows the global PD plot (blue), including ICE curves (left) and the REPs after applying REPID (right) for the feature of interest Longitude of the *California housing* dataset. The interaction importance $intImp_j$ between Longitude and the interacting features is 0.49 (Latitude), 0.18 (Ocean proximity), and $R_{int}^2 = 0.67$.

Diabetes dataset As a third real-world example, we apply REPID on the *Diabetes dataset*, which analyzes diabetes in Pima Indian women and is available in the *MASS* package in R. The dataset consists of seven numeric features and the binary target variable *type*, which indicates if a woman is diabetic. The features for the 332 women contained in the dataset include: *Npreg* (number of pregnancies), *Glu* (plasma glucose concentration), *Bp* (diastolic blood pressure in mm Hg), *Skin* (triceps skin fold thickness in mm), *Bmi* (body mass index, *ped* (diabetes pedigree function), *Age*. We trained an SVM using epsilon support vector regression with a Gauss

kernel, $C = 1$ and $\epsilon = 0.1$. Subsequently, we obtained a balanced accuracy of 0.72 using a 5-fold cross-validation. We are interested in how the feature *Skin* influences the predicted probability for diabetes. When looking at the global PDP in Figure 10, one would assume that the skin fold thickness does not effect the predicted probability for diabetes, however, the ICE curves in the left plot indicate heterogeneous effects and, hence, interactions. We apply REPID with the same configurations as used in the *titanic* example in Section 5 and obtain the REPs shown in the right plot of Figure 10. While the risk of diabetes is in general higher for women with a glucose concentration higher than 133 than for women with a lower glucose concentration, the REPs also show that the risk for women with high glucose concentration values first increases with skin fold thickness and then decreases (green and light green curves), while the risk of diabetes for women with lower glucose concentration values and a maximum of five pregnancies first slightly decreases until a thickness of approximately 20 mm and then increases with skin fold thickness (orange and red curve).

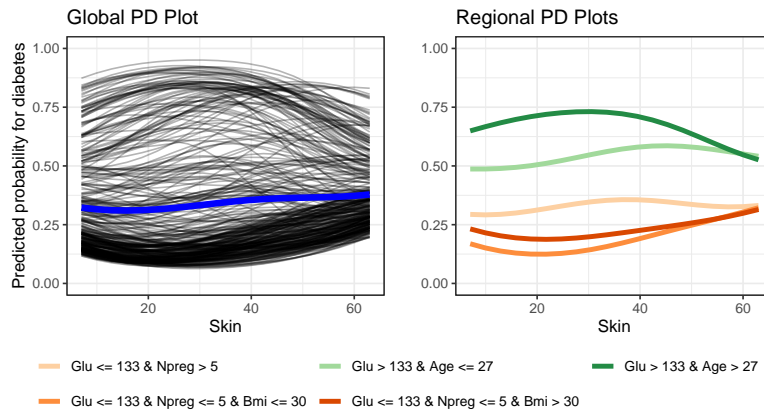


Figure 10: The figure shows the global PD plot (blue), including ICE curves (left), and the REPs after applying REPID (right) for the feature of interest *Skin* of the *diabetes* dataset. The interaction importance $intImp_j$ between *Skin* and the interacting features is 0.29 (Glu), 0.09 (Age), 0.08 (Npreg), 0.03 (Bmi) and $R_{int}^2 = 0.49$.

Table 4: Summary of the split information of the first three levels (depths) of the trees fitted by applying REPID to the simulation example stated in Section 4.2 for the 30 repetitions of the 4 models (GAM, XGBOOST, SVM, RF). The column “Depth” indicates the tree depth while the column “Node ID” indicates the respective node of this depth from left to right. The columns “Feature” and “Share” provide information of how often which feature was chosen for splitting in the respective node. The last two columns contain the mean and standard deviation of the respective split value. The coloring indicates the feature that was chosen most often for each node, where the different colors belong to the different tree depths.

Model	Depth	Node ID	Feature	Share	Split value mean	Split value sd
GAM	1	1	x4	1.00	0.01	0.02
	2	1	x8	1.00	0.02	0.10
	2	2	x8	1.00	-0.00	0.10
	3	1	x3	1.00	-0.04	0.14
	3	2	x8	0.03	0.31	
	3	2	x3	0.97	0.03	0.16
	3	3	x8	0.03	-0.37	
	3	3	x3	0.97	-0.05	0.16
	3	4	x8	0.07	0.38	0.07
3	4	x3	0.93	-0.01	0.14	
XGBOOST	1	1	x4	1.00	0.00	0.01
	2	1	x8	0.63	-0.06	0.11
	2	1	x3	0.37	-0.02	0.25
	2	2	x8	0.77	-0.05	0.11
	2	2	x3	0.23	0.01	0.11
	3	1	x8	0.17	-0.06	0.16
	3	1	x3	0.63	0.03	0.14
	3	1	x1	0.20	-0.01	0.10
	3	2	x8	0.07	-0.19	0.01
	3	2	x3	0.63	0.03	0.21
	3	2	x1	0.30	0.01	0.06
	3	3	x8	0.10	0.13	0.22
	3	3	x3	0.77	-0.04	0.19
	3	3	x1	0.13	-0.02	0.06
	3	4	x8	0.03	0.00	
3	4	x3	0.77	0.08	0.20	
3	4	x1	0.20	0.06	0.07	
SVM	1	1	x4	1.00	-0.03	0.07
	2	1	x8	1.00	-0.02	0.10
	2	2	x8	1.00	-0.11	0.10
	3	1	x4	0.23	-0.45	0.06
	3	1	x8	0.03	-0.55	
	3	1	x3	0.73	-0.02	0.15
	3	2	x4	0.37	-0.50	0.09
	3	2	x3	0.63	-0.15	0.13
	3	3	x4	0.20	0.35	0.07
	3	3	x8	0.07	-0.61	0.00
	3	3	x3	0.73	-0.18	0.16
	3	4	x4	0.07	0.24	0.05
	3	4	x3	0.93	-0.21	0.15
RF	1	1	x4	1.00	0.00	0.02
	2	1	x8	0.70	-0.12	0.09
	2	1	x3	0.30	0.21	0.18
	2	2	x8	1.00	-0.11	0.15
	3	1	x8	0.23	-0.08	0.13
	3	1	x3	0.70	0.17	0.18
	3	1	x1	0.07	0.04	0.17
	3	2	x8	0.30	-0.18	0.17
	3	2	x3	0.70	0.20	0.18
	3	3	x8	0.03	-0.48	
	3	3	x3	0.97	0.10	0.18
	3	4	x3	0.97	0.08	0.23